MegaRAID®
SAS Software

User Guide

80-00156-01
Rev. K
January 2011



80-00156-01K

Revision History

Version and Date	Description of Changes
80-00156-01 Rev. K, January 2011	Added and updated sections in the WebBIOS, MegaCLI, and MegaRAID Storage Manager chapters.
80-00156-01 Rev. J, September 2010	• Added section *Enabling the Snapshot Scheduler on page 209*. • Added *MegaRAID Software Licensing on page 301*. • Added *Enabling Drive Security using EKM on page 337*. • Added *Managing Software Licensing on page 54* and *SafeStore Security Options on page 171* details.
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 2.1	Updated operating system support for MegaCLI.
80-00156-01 Rev. B, June 2007 Version 2.0	Updated the WebBIOS Configuration Utility and the MegaRAID Storage Manager. Updated the MegaCLI commands. Added the RAID introduction chapter.
80-00156-01 Rev. A, August 2006 Version 1.1	Corrected the procedure for creating RAID 10 and RAID 50 drive groups in the WebBIOS Configuration Utility.
DB15-000339-00, December 2005 Version 1.0	Initial release of this document.

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.

This preliminary document describes a preproduction product and contains information that may change substantially for any final commercial release of the product. LSI Corporation makes no express or implied representation or warranty as to the accuracy, quality, or completeness of information contained in this document, and neither the release of this document nor any information included in it obligates LSI Corporation to make a commercial release of the product. 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

Website
www.lsi.com

Document Number: 80-00156-01 Rev. K
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 II Features 14
1.4 Solid State Drive Features 15
1.4.1 Solid State Drive Guard 15
1.5 Dimmer Switch Feature 15
1.6 UEFI 2.0 Support 16
1.7 Configuration Scenarios 16
1.7.1 Valid Drive Mix Configurations with HDDs and CacheCade software 17
1.8 Technical Support 18

Chapter 2: Introduction to RAID 19
2.1 RAID Description 19
2.2 RAID Benefits 19
2.3 RAID Functions 19
2.4 Components and Features 19
2.4.1 Drive Group 20
2.4.2 Virtual Drive 20
2.4.3 Fault Tolerance 20
2.4.4 Consistency Check 21
2.4.5 Copyback 21
2.4.6 Background Initialization 22
2.4.7 Patrol Read 23
2.4.8 Disk Striping 23
2.4.9 Disk Mirroring 23
2.4.10 Parity 24
2.4.11 Disk Spanning 25
2.4.12 Hot Spares 26
2.4.13 Disk Rebuilds 27
2.4.14 Rebuild Rate 28
2.4.15 Hot Swap 28
2.4.16 Drive States 28
2.4.17 Virtual Drive States 29
2.4.18 Beep Codes 29
2.4.19 Enclosure Management 29
2.5 RAID Levels 29
2.5.1 Summary of RAID Levels 30
2.5.2 Selecting a RAID Level 30
2.5.3 RAID 0 31
2.5.4 RAID 1 31
2.5.5 RAID 5 32
2.5.6 RAID 6 33
2.5.7 RAID 00 34
2.5.8 RAID 10 35
2.5.9 RAID 50 36
2.5.10 RAID 60 37
2.6 RAID Configuration Strategies 38

2.6.1 Maximizing Fault Tolerance 39
2.6.2 Maximizing Performance 40
2.6.3 Maximizing Storage Capacity 41
2.7 RAID Availability 42
2.7.1 RAID Availability Concept 42
2.8 Configuration Planning 43
2.9 Number of Drives 43
2.9.1 Drive Group Purpose 43

Chapter 3: SafeStore Disk Encryption 45
3.1 Overview 45
3.2 Purpose and Benefits 45
3.3 Terminology 46
3.4 Workflow 46
3.4.1 Enable Security 46
3.4.2 Change Security 47
3.4.3 Create Secure Virtual Drives 47
3.4.4 Import a Foreign Configuration 48
3.5 Instant Secure Erase 48

Chapter 4: WebBIOS Configuration Utility 51
4.1 Overview 51
4.2 Starting the WebBIOS CU 51
4.3 WebBIOS CU Main Screen Options 52
4.4 Managing Software Licensing 54
4.4.1 Managing MegaRAID Advanced Software Options 54
4.4.2 Reusing the Activation Key 57
4.4.3 Managing Advanced Software Summary 58
4.4.4 Activating an Unlimited Key Over a Trial Key 59
4.4.5 Activating a Trial Software 59
4.4.6 Activating an Unlimited Key 59
4.4.7 Securing MegaRAID Advanced Software 60
4.4.8 Confirm Rehosting Process 60
4.5 Creating a Storage Configuration 62
4.5.1 Selecting the Configuration with the Configuration Wizard 62
4.5.2 Converting JBOD to Unconfigured Good Drives 64
4.5.3 Using Automatic Configuration 65
4.5.4 Using Manual Configuration 66
4.6 Creating a CacheCade Configuration 104
4.7 Selecting SafeStore Encryption Services Security Options 109
4.7.1 Enabling the Security Key Identifier, Security Key, and Password 109
4.7.2 Enabling Drive Security using EKM 114
4.7.3 Changing the Security Key Identifier, Security Key, and Pass Phrase 115
4.7.4 Change Security from EKM to LKM 115
4.7.5 Changing Security from LKM to EKM 121
4.7.6 Import Foreign Drives in EKM /EKM Secured Locked Drives 122
4.7.7 Import Foreign Drives for LKM Secured Locked Drives 122
4.7.8 Import Foreign Drives in LKM /EKM Secured Locked Drives 123
4.7.9 Disabling the Drive Security Settings 123
4.8 Viewing and Changing Device Properties 125
4.8.1 Viewing Controller Properties 125

4.8.2 Viewing Virtual Drive Properties, Policies, and Operations 129
4.8.3 Viewing Drive Properties 131
4.8.4 Viewing and Changing Battery Backup Unit Information 133
4.8.5 Managing Link Speed 136
4.9 Expanding a Virtual Drive 138
4.10 Using MegaRAID Recovery 139
4.10.1 Recovery Scenarios 140
4.10.2 Enabling the Recovery Advanced Software 140
4.10.3 Creating Snapshots and Views 143
4.10.4 Creating Concurrent Snapshots 146
4.10.5 Selecting the Snapshot Settings 147
4.10.6 Viewing Snapshot Properties 149
4.10.7 Restoring a Virtual Drive by Rolling Back to a Snapshot 150
4.10.8 Cleaning up a Snapshot Repository 152
4.11 Viewing System Event Information 154
4.12 Managing Configurations 155
4.12.1 Running a Consistency Check 156
4.12.2 Deleting a Virtual Drive 156
4.12.3 Importing or Clearing a Foreign Configuration 156
4.12.4 Importing Foreign Configurations 160
4.12.5 Migrating the RAID Level of a Virtual Drive 160
4.12.6 New Drives Attached to a MegaRAID Controller 162

Chapter 5: MegaRAID Command Tool 165
5.1 Product Overview 165
5.2 Novell NetWare, SCO, Solaris, FreeBSD, and DOS Operating System Support 166
5.3 Command Line Abbreviations and Conventions 167
5.3.1 Abbreviations Used in the Command Line 167
5.3.2 Conventions 167
5.4 Pre-boot MegaCLI 168
5.5 CacheCade-Related Options 169
5.5.1 Create a Solid State Drive Cache Drive to Use as Secondary Cache 170
5.5.2 Delete a Solid State Drive Cache Drive 170
5.6 Software License Key 170
5.7 SafeStore Security Options 171
5.7.1 Use Instant Secure Erase on a Physical Drive 171
5.7.2 Secure Data on a Virtual Drive 171
5.7.3 Destroy the Security Key 172
5.7.4 Create a Security Key 172
5.7.5 Drive Security Key 172
5.7.6 Change the Security Key 173
5.7.7 Get the Security Key ID 173
5.7.8 Set the Security Key ID 174
5.7.9 Verify the Security Key 174
5.8 Controller Property-Related Options 174
5.8.1 Display Controller Properties 174
5.8.2 Display Number of Controllers Supported 174
5.8.3 Enable or Disable Automatic Rebuild 175
5.8.4 Flush Controller Cache 175
5.8.5 Set Controller Properties 175

5.8.6 Display Specified Controller Properties 177
5.8.7 Set Factory Defaults 177
5.8.8 Set SAS Address 177
5.8.9 Set Time and Date on Controller 177
5.8.10 Display Time and Date on Controller 178
5.8.11 Get Connector Mode 178
5.8.12 Set Connector Mode 178
5.9 Patrol Read-Related Controller Properties 178
5.9.1 Set Patrol Read Options 179
5.9.2 Set Patrol Read Delay Interval 179
5.10 BIOS-Related Properties 179
5.10.1 Set or Display Bootable Virtual Drive ID 179
5.10.2 Select BIOS Status Options 180
5.11 Battery Backup Unit-Related Properties 180
5.11.1 Display BBU Information 180
5.11.2 Display BBU Status Information 181
5.11.3 Display BBU Capacity 182
5.11.4 Display BBU Design Parameters 182
5.11.5 Display Current BBU Properties 183
5.11.6 Start BBU Learning Cycle 183
5.11.7 Place Battery in Low-Power Storage Mode 183
5.11.8 Set BBU Properties 184
5.12 Options for Displaying Logs Kept at the Firmware Level 184
5.12.1 Event Log Management 184
5.12.2 Set BBU Terminal Logging 185
5.13 Configuration-Related Options 185
5.13.1 Create a RAID Drive Group from All Unconfigured Good Drives 185
5.13.2 Add RAID 0, 1, 5, or 6 Configuration 187
5.13.3 Add RAID 10, 50, or 60 Configuration 188
5.13.4 Clear the Existing Configuration 188
5.13.5 Save the Configuration on the Controller 188
5.13.6 Restore the Configuration Data from File 189
5.13.7 Manage Foreign Configuration Information 189
5.13.8 Delete Specified Virtual Drive(s) 190
5.13.9 Display the Free Space 190
5.14 Virtual Drive-Related Options 190
5.14.1 Display Virtual Drive Information 190
5.14.2 Change the Virtual Drive Cache and Access Parameters 191
5.14.3 Display the Virtual Drive Cache and Access Parameters 191
5.14.4 Manage Virtual Drives Initialization 192
5.14.5 Manage a Consistency Check 192
5.14.6 Schedule a Consistency Check 193
5.14.7 Manage a Background Initialization 193
5.14.8 Perform a Virtual Drive Reconstruction 194
5.14.9 Display Information about Virtual Drives and Drives 194
5.14.10 Display the Number of Virtual Drives 194
5.14.11 Clear the LDBBM Table Entries 195
5.14.12 Display the List of Virtual Drives with Preserved Cache 195
5.14.13 Discard the Preserved Cache of a Virtual Drive(s) 195
5.14.14 Expand a Virtual Drive 196
5.15 Drive-Related Options 196
5.15.1 Display Drive Information 196
5.15.2 Set the Drive State to Online 197

5.15.3 Set the Drive State to Offline ... 197
5.15.4 Change the Drive State to Unconfigured Good ... 197
5.15.5 Change the Drive State ... 198
5.15.6 Manage a Drive Initialization ... 198
5.15.7 Rebuild a Drive ... 199
5.15.8 Locate the Drive(s) and Activate LED ... 199
5.15.9 Mark the Configured Drive as Missing ... 199
5.15.10 Display the Drives in Missing Status ... 200
5.15.11 Replace the Configured Drives and Start an Automatic Rebuild ... 200
5.15.12 Prepare the Unconfigured Drive for Removal ... 200
5.15.13 Display Total Number of Drives ... 201
5.15.14 Display List of Physical Devices ... 201
5.15.15 Download Firmware to the Physical Devices ... 201
5.15.16 Configure All Free Drives into a RAID 0, 1, 5, or 6 Configuration for a Specific Controller ... 202
5.15.17 Set the Mapping Mode of the Drives to the Selected Controller(s) ... 203
5.15.18 Perform the Copyback Operation on the Selected Drive ... 203
5.16 Enclosure-Related Options ... 203
5.16.1 Display Enclosure Information ... 204
5.16.2 Display Enclosure Status ... 204
5.17 Flashing the Firmware ... 204
5.17.1 Flash the Firmware with the ROM File ... 204
5.17.2 Flash the Firmware in Mode 0 with the ROM File ... 204
5.18 SAS Topology ... 205
5.19 Diagnostic-Related Options ... 205
5.19.1 Start Controller Diagnostics ... 205
5.19.2 Start Battery Test ... 205
5.20 Recovery (Snapshot)-Related Options ... 205
5.20.1 Enable the Snapshot Feature ... 206
5.20.2 Disable the Snapshot Feature ... 206
5.20.3 Take Snapshot of Volume ... 206
5.20.4 Set the Snapshot Properties ... 207
5.20.5 Delete a Snapshot ... 207
5.20.6 Create a View ... 208
5.20.7 Delete a View ... 208
5.20.8 Rollback to an Old Snapshot ... 208
5.20.9 Display Snapshot and View Information ... 209
5.20.10 Clean the Recoverable Free Space on the Drives in a Virtual Drive ... 209
5.20.11 Display the Information for a Specific View ... 209
5.20.12 Enabling the Snapshot Scheduler ... 209
5.21 FastPath-related Options ... 210
5.22 Miscellaneous Options ... 210
5.22.1 Display the MegaCLI Version ... 210
5.22.2 Display Help for MegaCLI ... 210
5.22.3 Summary Information ... 210

Chapter 6: MegaRAID Storage Manager Overview and Installation ... 211
6.1 Overview ... 211
6.1.1 Creating Storage Configurations ... 211
6.1.2 Monitoring Storage Devices ... 211
6.1.3 Maintaining Storage Configurations ... 211
6.2 Hardware and Software Requirements ... 211

6.3 Prerequisites to Running MSM Remote Administration ... 212
6.4 Installing MegaRAID Storage Manager ... 212
6.4.1 Prerequisite for MSM Installation ... 212
6.4.2 Installing MegaRAID Storage Manager Software on Microsoft Windows ... 213
6.4.3 Installing MegaRAID Storage Manager for SPARC ... 216
6.4.4 Installing MegaRAID Storage Manager Software for Linux ... 217
6.4.5 Prerequisites for Installing MSM on RHEL6.0 x64 Operating System ... 218
6.4.6 Linux Error Messages ... 218
6.4.7 MSM Customization ... 218
6.5 MegaRAID Storage Manager Support and Installation on VMWare ... 219
6.5.1 Installing MegaRAID Storage Manager for VMWare Classic ... 219
6.5.2 Uninstalling MegaRAID Storage Manager for VMWare ... 219
6.5.3 MegaRAID Storage Manager Support on the VMWare ESXi Operating System ... 220
6.5.4 Limitations ... 222
6.6 Installing and Configuring a CIM Provider ... 224
6.6.1 Installing a CIM SAS Storage Provider on Linux ... 224
6.6.2 Installing a CIM SAS Storage Provider on Windows ... 225
6.7 Installing and Configuring an SNMP Agent ... 226
6.7.1 Prerequisite for LSI SNMP Agent RPM Installation ... 226
6.7.2 Installing and Configuring an SNMP Agent on Linux ... 226
6.7.3 Installing and Configuring an SNMP Agent on Solaris ... 228
6.7.4 Installing an SNMP Agent on Windows ... 230
6.8 MegaRAID Storage Manager Support and Installation on Solaris 10 ... 231
6.8.1 Installing MegaRAID Storage Manager Software for Solaris 10 ... 231
6.8.2 Uninstalling MegaRAID Storage Manager Software for Solaris 10 ... 231
6.9 Prerequisites to Running MSM Remote Administration ... 232

Chapter 7: MegaRAID Storage Manager Windows and Menus ... 233
7.1 Starting MegaRAID Storage Manager Software ... 233
7.2 MegaRAID Storage Manager Main Menu ... 236
7.2.1 Dashboard/PhysicalView/Logical View ... 236
7.2.2 Properties/Graphical View Tabs ... 240
7.2.3 Event Log Panel ... 241
7.2.4 Menu Bar ... 241

Chapter 8: Configuration ... 243
8.1 Creating a New Storage Configuration ... 243
8.1.1 Selecting Virtual Drive Settings ... 243
8.1.2 Optimum Controller settings for CacheCade ... 245
8.1.3 Optimum Controller settings for FastPath ... 245
8.1.4 Creating a Virtual Drive Using Simple Configuration ... 245
8.1.5 Creating a Virtual Drive Using Advanced Configuration ... 249
8.2 Converting JBOD Drives to Unconfigured Good ... 255
8.2.1 Converting JBOD to Unconfigured Good from the MSM Window ... 256
8.3 Adding Hot Spare Drives ... 257
8.4 Changing Adjustable Task Rates ... 258
8.5 Changing Power Settings ... 259
8.5.1 Changing Power Settings Using the Dimmer Switch Enhancement ... 261
8.5.2 Appearance for Controller Defined ... 263
8.6 Changing Virtual Drive Properties ... 264
8.7 Changing a Virtual Drive Configuration ... 265

8.7.1 Accessing the Modify Drive Group Wizard . . . 266
8.7.2 Adding a Drive or Drives to a Configuration . . . 267
8.7.3 Removing a Drive from a Configuration . . . 269
8.7.4 Replacing a Drive . . . 270
8.7.5 Migrating the RAID Level of a Virtual Drive . . . 270
8.7.6 New Drives Attached to a MegaRAID Controller . . . 273
8.8 Deleting a Virtual Drive . . . 273

Chapter 9: Monitoring System Events and Storage Devices . . . 275
9.1 Monitoring System Events . . . 275
9.2 Configuring Alert Notifications . . . 276
9.2.1 Setting Alert Delivery Methods . . . 278
9.2.2 Changing Alert Delivery Methods for Individual Events . . . 278
9.2.3 Changing the Severity Level for Individual Events . . . 279
9.2.4 Multiple Events Displayed in a Single Pop-up Window . . . 280
9.2.5 Entering or Editing the Sender Email Address and SMTP Server . . . 280
9.2.6 Authenticating a Server . . . 281
9.2.7 Saving Backup Configurations . . . 281
9.2.8 Loading Backup Configurations . . . 282
9.2.9 Adding Email Addresses of Recipients of Alert Notifications . . . 282
9.2.10 Testing Email Addresses of Recipients of Alert Notifications . . . 283
9.2.11 Removing Email Addresses of Recipients of Alert Notifications . . . 283
9.3 Monitoring Controllers . . . 284
9.4 Monitoring Drives . . . 284
9.5 Running a Patrol Read . . . 286
9.5.1 Patrol Read Task Rates . . . 288
9.6 Monitoring Virtual Drives . . . 288
9.7 Monitoring Enclosures . . . 289
9.8 Monitoring Battery Backup Units . . . 290
9.8.1 Battery Learn Cycle . . . 290
9.9 Monitoring Rebuilds and Other Processes . . . 291

Chapter 10: Maintaining and Managing Storage Configurations . . . 293
10.1 Initializing a Virtual Drive . . . 293
10.1.1 Running a Group Initialization . . . 293
10.2 Running a Consistency Check . . . 294
10.2.1 Setting the Consistency Check Settings . . . 295
10.2.2 Scheduling a Consistency Check . . . 296
10.2.3 Running a Group Consistency Check . . . 297
10.3 Scanning for New Drives . . . 298
10.4 Rebuilding a Drive . . . 298
10.5 Making a Drive Offline or Missing . . . 299
10.6 Upgrading the Firmware . . . 299

Chapter 11: Using MegaRAID® Advanced Software . . . 301
11.1 MegaRAID Advanced Software . . . 301
11.2 Recovery Advanced Software . . . 301
11.2.1 MegaRAID Software Licensing . . . 301
11.2.2 Managing MegaRAID Advanced Software . . . 301

11.2.3 Activation Key . . . 303
11.2.4 Advanced MegaRAID Software Status Summary . . . 304
11.2.5 Application Scenarios and Messages . . . 305
11.2.6 Activating an Unlimited Key Over a Trial Key . . . 306
11.2.7 Configuring Key Vault (Re-hosting process) . . . 309
11.2.8 Rehosting Complete . . . 310
11.2.9 Deactivate Trial Software . . . 311
11.2.10 MegaRAID Recovery . . . 312
11.2.11 Recovery Scenarios . . . 312
11.2.12 Enabling the Recovery Advanced Software . . . 313
11.2.13 Snapshot Repository . . . 314
11.2.14 Selecting the Virtual Drive . . . 316
11.2.15 Scheduling Snapshots . . . 317
11.2.16 Editing Snapshots . . . 318
11.2.17 Snapshot Base Details . . . 319
11.2.18 Manage Snapshots . . . 320
11.2.19 Editing the Snapshot Schedule . . . 321
11.2.20 Advanced Settings . . . 322
11.2.21 Create View . . . 323
11.2.22 Viewing Snapshot Details . . . 324
11.2.23 No View Details for Snapshot . . . 325
11.2.24 No Snapshot Schedule . . . 326
11.2.25 Creating Views . . . 327
11.2.26 Restoring by Rolling Back to a Snapshot . . . 328
11.2.27 Restoring from a View . . . 329
11.2.28 Deleting a Snapshot . . . 329
11.3 Join Mirror . . . 329
11.4 CacheCade Advanced Software . . . 331
11.4.1 Using the CacheCade Advanced Software . . . 331
11.5 FastPath Advanced Software . . . 335
11.5.1 Setting FastPath Options . . . 335
11.6 LSI SafeStore Encryption Services . . . 336
11.6.1 Enabling Drive Security using EKM . . . 337
11.6.2 **Supporting EKM mode . . . 337**
11.6.3 Change Security Settings- LKM . . . 339
11.6.4 Change Security Settings - EKM . . . 342
11.6.5 Importing Foreign Drives . . . 343
11.6.6 Importing Foreign Drives to LKM . . . 344
11.6.7 Importing Foreign Drives to EKM . . . 344
11.6.8 Importing Foreign Drives to EKM . . . 345
11.6.9 Enabling Drive Security using LKM . . . 346
11.6.10 Changing the Drive Security Settings . . . 349
11.6.11 Disabling Drive Security . . . 350
11.6.12 Importing or Clearing a Foreign Configuration . . . 351
11.7 Managing Link Speed . . . 354

Appendix A: Events and Messages . . . 357
A.1 Error Levels . . . 357
A.2 Event Messages . . . 357

Appendix B: MegaCLI Error Messages . . . 371
B.1 Error Messages and Descriptions . . . 371

Appendix C: Glossary 375

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, WebBIOS configuration utility, and MegaRAID command line interface (CLI). This chapter documents the SAS technology, Serial ATA (SATA) technology, CacheCade software, Dimmer Switch™, EUFI 2.0, configuration scenarios, and drive types.

1.1 SAS Technology

The MegaRAID 6Gb/s SAS RAID controllers are high-performance intelligent PCI Express-to-SCSI/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 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 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 1.1*. In addition, the controller supports the SATA II protocol defined by the *Serial ATA specification, version 1.0a*. Supporting both the SAS and SATA II interfaces, the SAS controller is a versatile controller that provides the backbone of both server environments and high-end workstation environments.

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

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

1.2 Serial-attached SCSI Device Interface

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

- Supports SATA II data transfers of 3.0 Gb/s
- Supports STP data transfers of 3.0 Gb/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 CacheCade software drives attached to MegaRAID SAS controllers. These drives are expected to behave like SATA HDDs or SAS HDDs. The major advantages of CacheCade software drives include:

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

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

NOTE: MegaRAID implements support for only those SATA CacheCade software drives which support ATA-8 ACS compliance.

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

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

1.4.1 Solid State Drive Guard

CacheCade software is known for its reliability and performance. CacheCade software Guard™, a feature that is unique to MegaRAID, increases the reliability of CacheCade software by automatically copying data from a drive with potential to fail to a designated hot spare or newly inserted drive. Because CacheCade software are very reliable, non-redundant RAID 0 configurations are much more common than in the past. CacheCade software Guard offers added data protection for RAID 0 configurations.

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

1.5 Dimmer Switch Feature

Powering and cooling drives represents a major cost for data centers. The new MegaRAID Dimmer Switch reduces the power consumption of the devices connected to a MegaRAID controller. This helps to share resources more efficiently and lower costs.

With Dimmer Switch, any unconfigured drive connected to a MegaRAID controller is spun down after 30 minutes of inactivity, reducing its power usage. Spun down drives are spun up automatically when you create a configuration using those drives.

1.6 UEFI 2.0 Support

Significant challenges face operating system and platform developers to innovate using the legacy PC-AT BIOS boot environment. These include memory constraints, maintenance challenges, and increased complexities due to a lack of industry-wide standards.

To handle these challenges, the Unified Extensible Firmware Interface (UEFI) was developed to do the following:

- Define a clean interface between operating systems and the hardware platform at boot time.
- Support an architecture-independent mechanism for initializing add-in cards.

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 commands (SPT), 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

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

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

Figure 1 shows a direct-connect configuration. The Inter-IC (I^2C) 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.

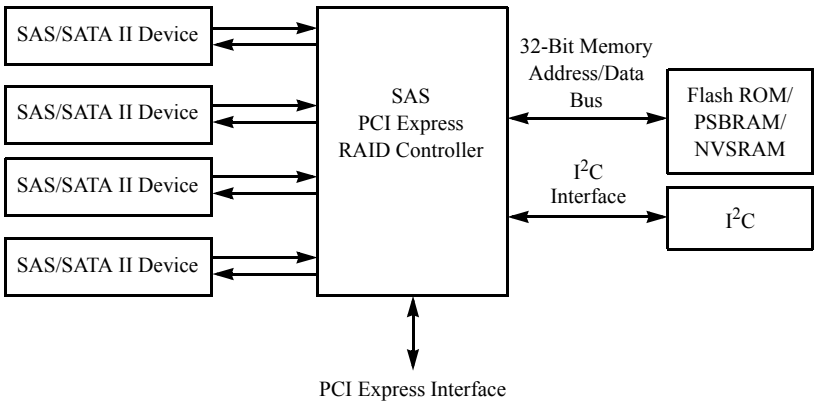


Figure 1: Example of an LSI SAS Direct-Connect Application

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.

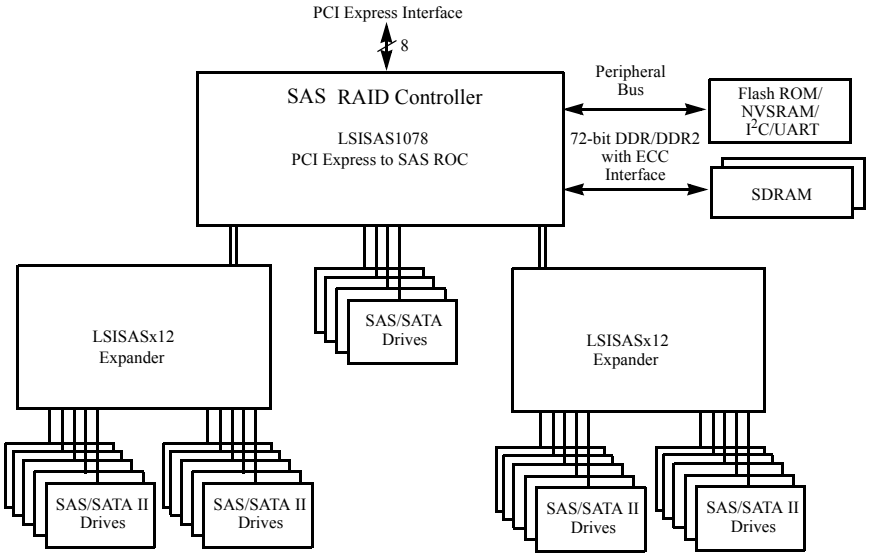


Figure 2: Example of an LSI SAS RAID Controller Configured with an LSISASx12 Expander

1.7.1 Valid Drive Mix Configurations with HDDs and CacheCade software

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

You can choose whether to allow a virtual drive to consist of both CacheCade software devices and HDDs. For a virtual drive that consists of CacheCade software only, you can choose whether to allow SAS CacheCade software drives and SATA CacheCade software drives in that virtual drive. For virtual drives that have both CacheCade software and HDDs, you can choose whether to mix SAS and SATA HDD drives with SAS and SATA CacheCade 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 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 software (SATA-only configuration)
3.	SAS HDD with a mix of SAS and SATA CacheCade software (a SATA HDD cannot be added)
4.	SATA HDD with a mix of SAS and SATA CacheCade software (a SAS HDD cannot be added)
5.	SAS CacheCade software with a mix of SAS and SATA HDD (a SATA CacheCade software cannot be added)
6.	SATA CacheCade software with a mix of SAS and SATA HDD (a SAS CacheCade software cannot be added)
7.	A mix of SAS and SATA HDD with a mix of SAS and SATA CacheCade software
8.	A CacheCade 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 26*.

1.8 Technical Support

For assistance with installing, configuring, or running your MegaRAID 6Gb/s SAS RAID controller, 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 email 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. Below are some common RAID functions:

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

2.4 Components and Features

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

2.4.1 Drive Group

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 whole 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, this means that it is also important for the system to be available during the repair of the problem.

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

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

2.4.3.1 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 issue events appropriately.

In a system with two ESMs, you can replace one of the ESMs without affecting the virtual drive availability. For example, the controller can run heavy I/Os, and when you replace one of the ESM modules, I/Os should not stop. The controller uses different paths to balance the load on the entire system.

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

2.4.4 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 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:
  - Where the hot spare is not configured (or unavailable) in the system.
  - There are two drives (both members of virtual drives), 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 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 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* or Section 9.5, *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 8KB to 1024KB. 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.

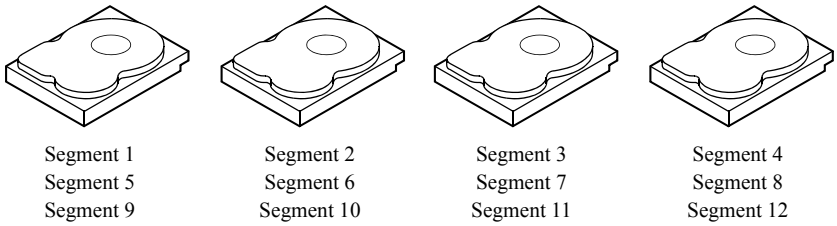


Figure 3: Example of Disk Striping (RAID 0)

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 is expensive because each drive in the system must be duplicated. Figure 4 shows an example of disk mirroring.

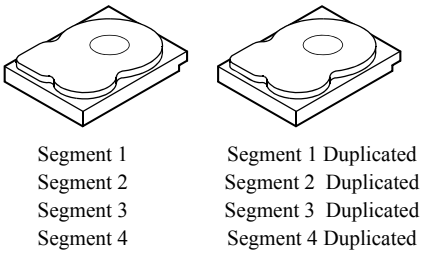


Figure 4: Example of Disk Mirroring (RAID 1)

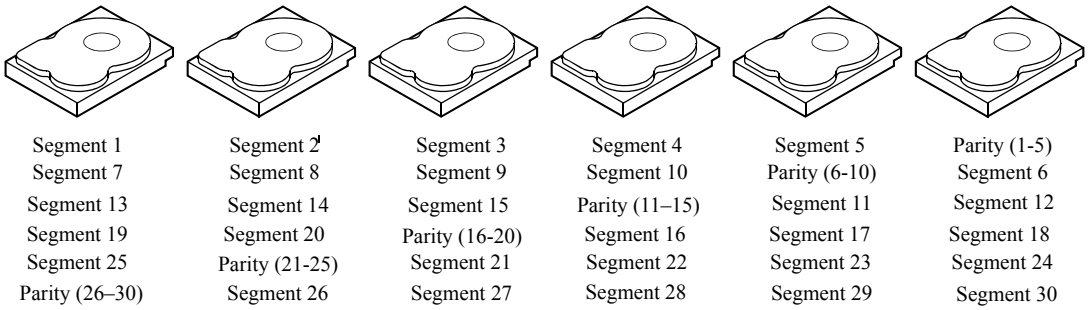
2.4.10 Parity

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.

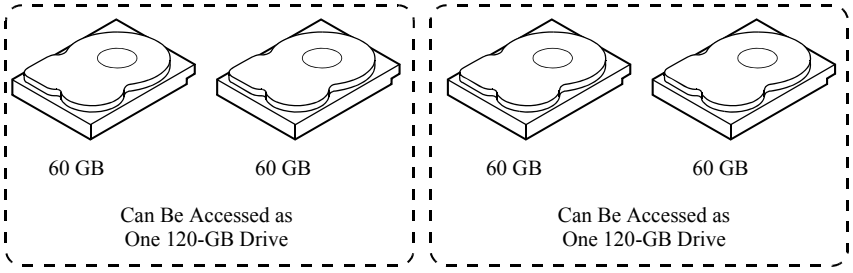


Figure 6: Example of Disk Spanning

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 eight. 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 eight spans. You must use an even number of drives in each RAID virtual drive in the span. The RAID 1 virtual drives must have the same stripe size.
50	Configure RAID 50 by spanning two contiguous RAID 5 virtual drives. The RAID 5 virtual drives must have the same stripe size.
60	Configure RAID 60 by spanning two contiguous RAID 6 virtual drives. The RAID 6 virtual drives must have the same stripe size.

NOTE: In a spanned VD (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.

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

The RAID management software allows you to specify drives as hot spares. When a hot spare is needed, the RAID controller assigns the hot spare that has a capacity closest to and at least as great as that of the failed drive to take the place of the failed drive. The failed drive is removed from the virtual drive and marked ready awaiting removal once the rebuild to a hot spare begins. You can make hot spares of the drives that are not in a RAID virtual drive.

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

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

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

There are two types of hot spares:

- Global hot spare
- Dedicated hot spare

2.4.12.1 Global Hot Spare

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

2.4.12.2 Dedicated Hot Spare

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

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

Observe the following parameters when using hot spares:

- Hot spares are used only in drive groups with redundancy: RAID levels 1, 5, 6, 10, 50, and 60.
- A hot spare connected to a specific RAID controller can be used to rebuild a drive that is connected to the same controller only.
- You must assign the hot spare to one or more drives through the controller BIOS or use drive group management software to place it in the hot spare pool.
- A hot spare must have free space equal to or greater than the drive it replaces. For example, to replace an 18-GB drive, the hot spare must be 18 GB or larger.

2.4.13 Disk Rebuilds

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

The RAID controller uses hot spares to rebuild failed drives automatically and transparently, at user-defined rebuild rates. If a hot spare is available, the rebuild can start automatically when a drive fails. If a hot spare is not available, the failed drive must be replaced with a new drive so that the data on the failed drive can be rebuilt.

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

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

NOTE: If a source drive fails during a rebuild to a hot spare, the rebuild fails, and the failed source drive is marked as offline. In addition, the rebuilding hot spare drive is changed back to a hot spare. After a rebuild fails because of a source drive failure, the dedicated hot spare is still dedicated and assigned to the correct drive group, and the global hot spare is still global.

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

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

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

2.4.16 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.
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. When a virtual drive with cached data goes offline, the cache for the virtual drive is discarded. Because the virtual drive is offline, the cache cannot be saved.

2.4.17 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 1 or more 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 1 drive	Degraded	1 second on and 1 second off
RAID 5 loses 2 or more drives	Offline	3 seconds on and 1 second off
RAID 6 loses 1 drive	Partially Degraded	1 second on and 1 second off
RAID 6 loses 2 drives	Degraded	1 second on and 1 second off
RAID 6 loses more than 2 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 and/or hardware. The disk subsystem can be part of the host computer or can reside in an external disk enclosure. Enclosure management helps you stay informed of events in the disk subsystem, such as a drive or power supply failure. Enclosure management increases the fault tolerance of the disk subsystem.

2.5 RAID Levels

The RAID controller supports RAID levels 0, 00, 1, 5, 6, 10, 50, and 60. The supported RAID levels are summarized in the following section.

In addition, it supports independent drives (configured as RAID 0 and RAID 00.) The following sections describe the RAID levels in detail.

2.5.1 Summary of RAID Levels

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

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

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

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

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 eight spans. You must use an even number of drives in each RAID virtual drive in the span. The RAID 1 virtual drives must have the same stripe size. RAID 10 provides high data throughput and complete data redundancy but uses a larger number of spans.

RAID 50, a combination of RAID 0 and RAID 5, uses distributed parity and disk striping. A RAID 50 drive group is a spanned drive group in which data is striped across multiple RAID 5 drive groups. RAID 50 works best with data that requires high reliability, high request rates, high data transfers, and medium-to-large capacity.

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

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

2.5.2 Selecting a RAID Level

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

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

2.5.3 RAID 0

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

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

By breaking up a large file into smaller segments, the RAID controller can use both SAS drives and SATA drives to read or write the file faster. RAID 0 involves no parity calculations to complicate the write operation. This makes RAID 0 ideal for applications that require high bandwidth but do not require fault tolerance. Table 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 lost if any drive fails.
Drives	1 to 32

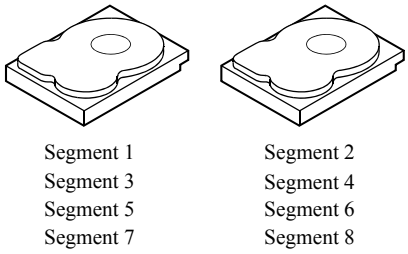


Figure 7: RAID 0 Drive Group Example with Two Drives

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 to 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.

Table 8: RAID 1 Overview

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 - 32 (must be an even number of drives)

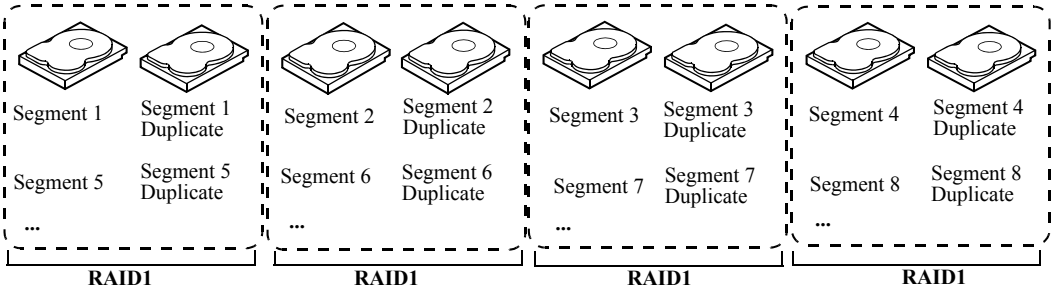


Figure 8: RAID 1 Drive Group

2.5.5 RAID 5

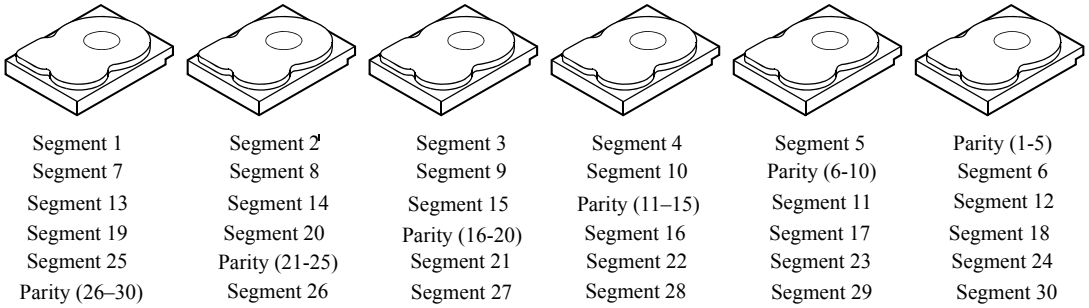
RAID 5 includes disk striping at the block level and parity. Parity is the data's property of being odd or even, and parity checking is used to detect errors in the data. In RAID 5, the parity information is written to all drives. RAID 5 is best suited for networks that perform a lot of small input/output (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 recreate all missing information. Use also for office automation and online customer service that requires fault tolerance. Use for any application that has high read request rates but low write request rates.
Strong Points	Provides data redundancy, high read rates, and good performance in most environments. Provides redundancy with lowest loss of capacity.
Weak Points	Not well-suited to tasks requiring lot of writes. Suffers more impact if no cache is used (clustering). Drive performance will be reduced if a drive is being rebuilt. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.
Drives	3 to 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 two drives in a virtual drive without losing data. Provides a high level of data protection through the use of a second parity block in each stripe. Use RAID 6 for data that requires a very high level of protection from loss.

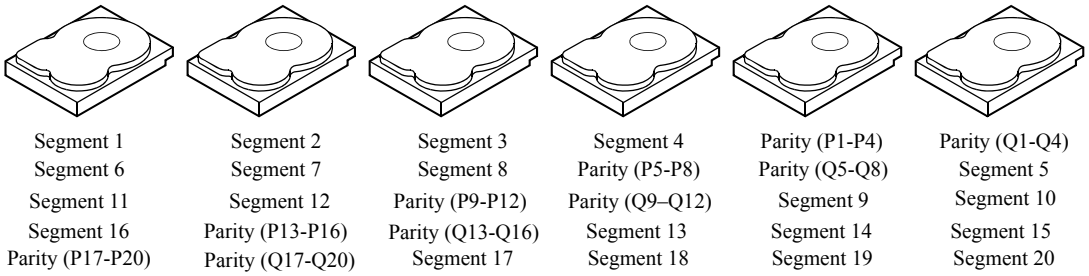
In the case of a failure of one drive or two drives in a virtual drive, the RAID controller uses the parity blocks to recreate all of the missing information. If two drives in a RAID 6 virtual drive fail, two drive rebuilds are required, one for each drive. These rebuilds do not occur at the same time.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

Table 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 to 32

Figure 2.5.7 shows a RAID 6 data layout. The second set of parity drives are denoted by *Q*. The *P* drives follow the RAID 5 parity scheme.



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 whole virtual drive (all drives associated with the virtual drive) will fail.

By breaking up a large file into smaller segments, the RAID controller can use both SAS drives and SATA drives to read or write the file faster. RAID 00 involves no parity calculations to complicate the write operation. This makes RAID 00 ideal for applications that require high bandwidth but do not require fault tolerance. Table 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 to 256

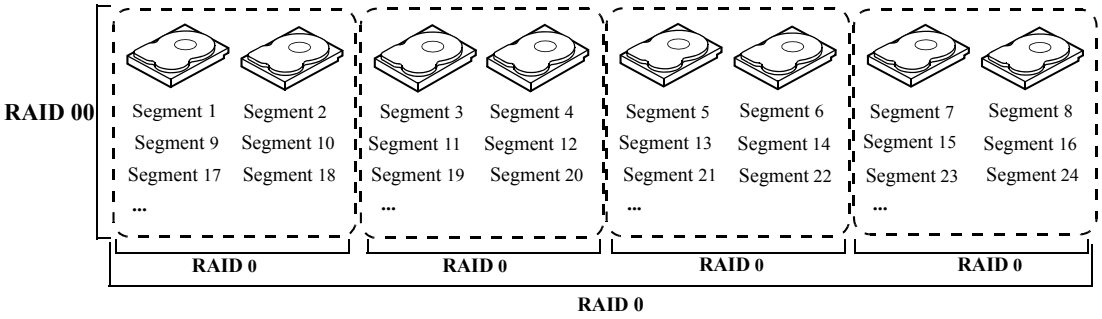


Figure 11: RAID 00 Drive Group Example with Two Drives

2.5.8 RAID 10

RAID 10 is a combination of RAID 0 and RAID 1, and consists of stripes across mirrored drives. RAID 10 breaks up data into smaller blocks and then mirrors the blocks of data to each RAID 1 drive group. The first RAID 1 drive in each drive group then duplicates its data to the second drive. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set. The RAID 1 virtual drives must have the same stripe size.

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

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

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

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

Table 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 - 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).

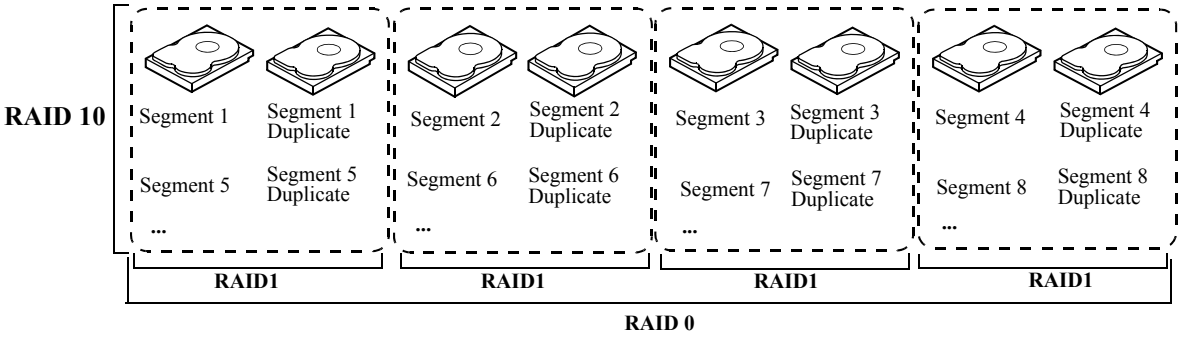


Figure 12: RAID 10 Level Virtual Drive

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 eight spans and tolerate up to eight drive failures, though less than total drive capacity is available. Though multiple drive failures can be tolerated, only one drive failure can be tolerated in each RAID 5 level drive group.

Table 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 to 8 times as many parity drives as RAID 5.
Drives	Eight spans of RAID 5 drive groups containing 3-32 drives each (limited by the maximum number of devices supported by the controller)

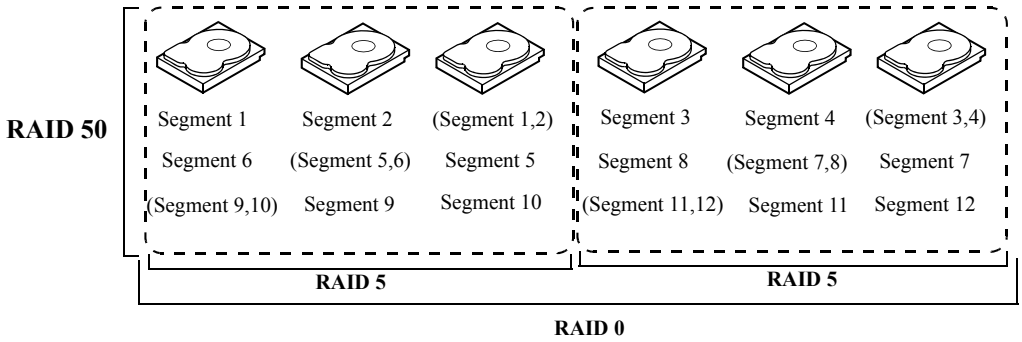


Figure 13: RAID 50 Level Virtual Drive

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.

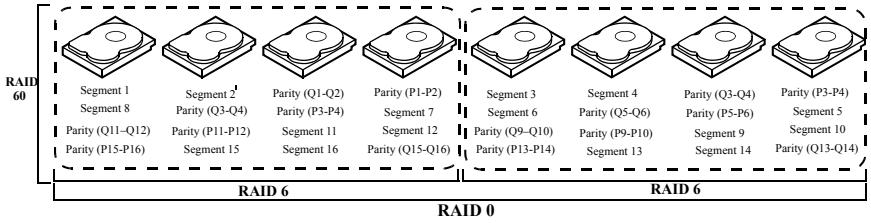
RAID 60 breaks up data into smaller blocks, and then stripes the blocks of data to each RAID 6 disk set. RAID 6 breaks up data into smaller blocks, calculates parity by performing an exclusive-or on the blocks and then writes the blocks of data and parity to each drive in the drive group. The size of 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	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. In the case of a failure of one drive or two drives in a RAID set in a virtual drive, the RAID controller uses the parity blocks to recreate all of the missing information. If two drives in a RAID 6 set in a RAID 60 virtual drive fail, two drive rebuilds are required, one for each drive. These rebuilds can occur at the same time. 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. 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.
Weak Points	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.
Drives	A minimum of 8

Figure 14 shows a RAID 6 data layout. The second set of parity drives are denoted by *Q*. The *P* drives follow the RAID 5 parity scheme.



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

Figure 14: RAID 60 Level Virtual Drive

2.6 RAID Configuration Strategies

The most important factors in RAID drive group configuration are:

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

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

2.6.1 Maximizing Fault Tolerance

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

A hot swap is the manual substitution of a replacement unit in a disk subsystem for a defective one, where the substitution can be performed while the subsystem is runninghot 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 bandwidth but do not require fault tolerance.
1	Provides complete data redundancy. If one drive fails, the contents of the other drive in the drive group can be used to run the system and reconstruct the failed drive. The primary advantage of disk mirroring is that it provides 100 percent data redundancy. Since the contents of the drive are completely written to a second drive, no data is lost if one of the drives fails. Both drives contain the same data at all times. RAID 1 is ideal for any application that requires fault tolerance and minimal capacity.
5	Combines distributed parity with disk striping. Parity provides redundancy for one drive failure without duplicating the contents of entire drives. If a drive fails, the RAID controller uses the parity data to reconstruct all missing information. In RAID 5, this method is applied to entire drives or stripes across all drives in a drive group. Using distributed partiy, 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 partiy, 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 drive integrity.
50	Provides data redundancy using distributed parity across spanned RAID 5 drive groups. RAID 50 includes both parity and disk striping across multiple drives. If a drive fails, the RAID controller uses the parity data to recreate all missing information. RAID 50 can sustain one drive failure per RAID 5 drive group and still maintain data integrity.
60	Provides data redundancy using distributed parity across spanned RAID 6 drive groups. RAID 60 can sustain two drive failures per RAID 6 drive group and still maintain data integrity. It provides the highest level of protection against drive failures of all of the RAID levels. RAID 60 includes both parity and disk striping across multiple drives. If a drive fails, the RAID controller uses the parity data to recreate all missing information.

2.6.2 Maximizing Performance

A RAID disk subsystem improves I/O performance. The RAID drive group appears to the host computer as a single storage unit or as multiple virtual units. I/O is faster because drives can be accessed simultaneously. Table 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. Since each drive contains both data and parity, numerous writes can take place concurrently. In addition, robust caching algorithms and hardware based exclusive-or assist make RAID 5 performance exceptional in many different environments. Parity generation can slow the write process, making write performance significantly lower for RAID 5 than for RAID 0 or RAID 1. 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 eight.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans and RAID performance degrades to that of a RAID 1 or RAID 5 drive group.
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 eight.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans and RAID performance degrades to that of a RAID 1 or RAID 5 drive group.
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 eight.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans and RAID performance degrades to that of a RAID 1 or RAID 6 drive group. 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 then 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.
1	With RAID 1 (mirroring), data written to one drive is simultaneously written to another drive, which doubles the required data storage capacity. This is expensive because each drive in the system must be duplicated.
5	RAID 5 provides redundancy for one drive failure without duplicating the contents of entire drives. RAID 5 breaks up data into smaller blocks, calculates parity by performing an exclusive-or on the blocks, then writes the blocks of data and parity to each drive in the drive group. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set.
6	RAID 6 provides redundancy for two drive failures without duplicating the contents of entire drives. However, it requires extra capacity because it uses two parity blocks per stripe. This makes RAID 60 more expensive to implement.
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 big drive. Spanning overcomes lack of disk space and simplifies storage management by combining existing resources or adding relatively inexpensive resources.
50	RAID 50 requires two to four times as many parity drives as RAID 5. This RAID level works best when used with data that requires medium to large capacity.
60	RAID 60 provides redundancy for two drive failures in each RAID set without duplicating the contents of entire drives. However, it requires extra capacity because a RAID 60 virtual drive has to generate two sets of parity data for each write operation. This makes RAID 60 more expensive to implement.

2.7 RAID Availability

2.7.1 RAID Availability Concept

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

2.7.1.1 Spare Drives

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

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

Hot spare drives are drives that power up along with the RAID drives and operate in a standby state. If a drive used in a RAID virtual drive fails, a hot spare automatically takes its place and the data on the failed drive is rebuilt on the hot spare. Hot spares can be used for RAID levels 1, 5, 6, 10, 50, and 60.

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

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

2.7.1.2 Rebuilding

If a drive fails in a drive group that is configured as a RAID 1, 5, 6, 10, 50, or 60 virtual drive, you can recover the lost data by rebuilding the drive. If you have configured hot spares, the RAID controller automatically tries to use them to rebuild failed drives. Manual rebuild is necessary if no hot spares with enough capacity to rebuild the failed drives are available.You must insert a drive with enough storage into the subsystem before rebuilding the failed drive.

2.8 Configuration Planning

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

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

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

2.9 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 Level(s)
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 it is housed in, the data on that drive is encrypted and useless to anyone who attempts to access without the 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 MegaRAID Storage Manager (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 key(s) 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, 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 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 (FW) uses a unique 512-bit key for each virtual disk. The VEK for the VDs 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 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 will be 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.

There are three procedures you can perform 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 MegaRAID Storage Manager.

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 Pass Word

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 their 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 MSM 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 MegaRAID Storage Manager.

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 exists 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.

There are three major reasons for using instant secure erase.

If there is a 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 the drive can be moved to another server or area without concern that old data could be found.

If there is a 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 there is a 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 needs to 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 CU, unlike the MegaRAID Storage Manager™ software, resides in the SAS controller BIOS and operates independently of the operating system.

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

- Create drive groups and virtual drives for storage configurations
- Create a CacheCade configuration
- Import a foreign configuration
- Migrate a storage configuration to a different RAID level
- Delete virtual drives
- Display controller, virtual drive, drive, and battery backup unit (BBU) properties, and change parameters
- Detect configuration mismatches
- Check configurations for data consistency
- Scan devices connected to the controller
- Initialize virtual drives

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

4.2 Starting the WebBIOS CU

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

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

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

   The Controller Selection screen appears.

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

   The main WebBIOS CU screen appears.

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 CU Main Screen Options

Figure 15 shows the screen that appears when you start the WebBIOS CU and select a controller.

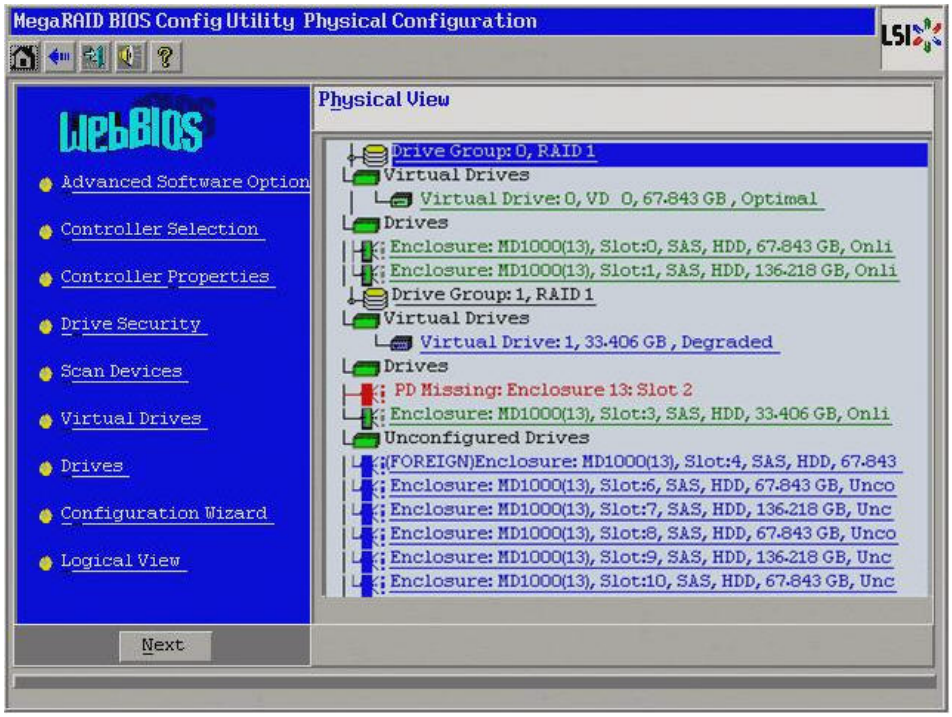


Figure 15: WebBIOS CU Main Screen

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

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

NOTE: The minimum screen resolution for WebBIOS is 640x480.

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

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

- Enclosure
- Slot
- Interface type (such as SAS or SATA)
- Drive type (HDD or CacheCade software)

- Drive size
- Drive status (such as **Online** or **Unconfigured Good**)

The toolbar at the top of the WebBIOS CU has the following buttons, as listed in Table 20.

Table 20: WebBIOS CU Toolbar Icons

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

Here is a description of the options listed on the left of the main WebBIOS CU screen (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 Section 4.4.1, *Managing MegaRAID Advanced Software Options*.
- **Controller Selection: (Alt+c)** Select this option to view the Controller Selection screen, where you can select a different SAS controller. You can then view information about the controller and the devices connected to it, or create a new configuration on the controller.
- **Controller Properties: (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 CU re-scan the physical and virtual drives for any changes in the drive status or the physical configuration. The WebBIOS CU displays the results of the scan in the physical and virtual drive descriptions.

- **Virtual Drives: (Alt+v)** Select this option to view the Virtual Drives screen, where you can change and view virtual drive properties, delete virtual drives, initialize drives, and perform other tasks. For more information, see Section 4.8.2, *Viewing Virtual Drive Properties, Policies, and Operations*.
- **Drives: (Alt+d)** Select this option to view the Drives screen, where you can view drive properties, create hot spares, and perform other tasks. For more information, see Section 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 Logical View; **Alt+h** for Physical View) Select this option to toggle between the Physical View screen and the Logical View screen.
- **Events: (Alt+e)** Select this option to view system events in the Event Information screen. For more information, see Section 4.11, *Viewing System Event Information*.
- **Exit: (Alt+x)** Select this option to exit the WebBIOS CU and continue with system boot.

4.4 Managing Software Licensing

This section explains the procedure of obtaining the software license key to enable the advanced options in WebBIOS. These advanced features are offered by the MegaRAID advanced software. 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 have to configure the **Advanced Software Options** menu present in the WebBIOS main screen 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 screen. The **Advanced Software Options** wizard appears as shown in the following figure.

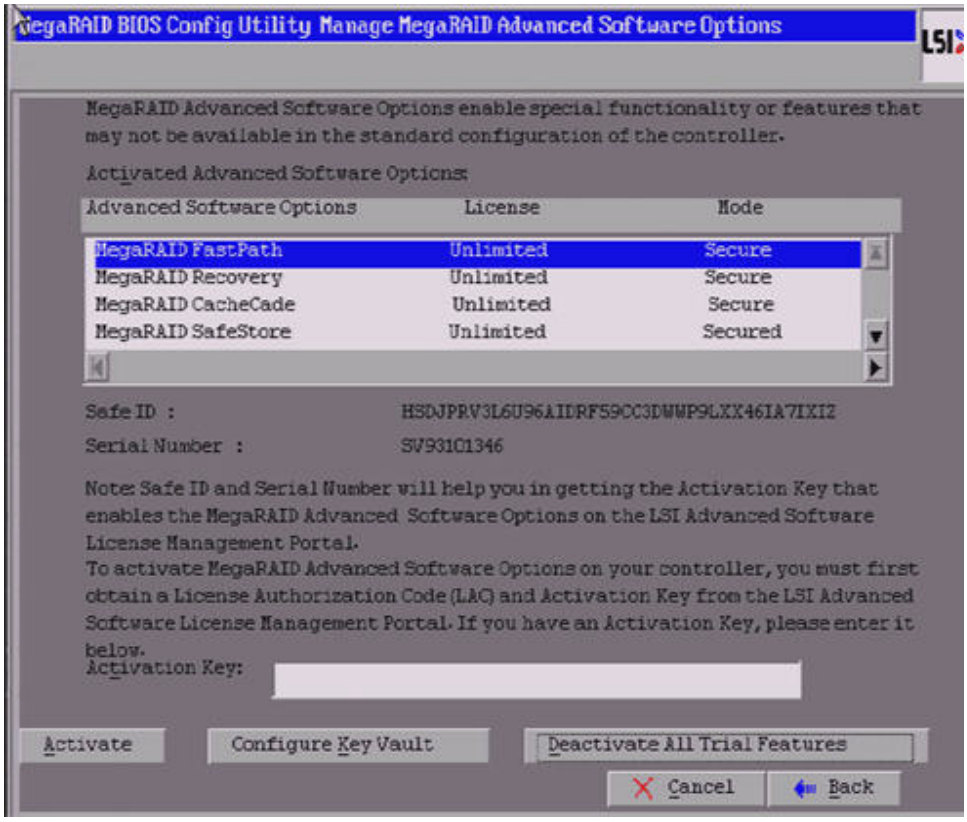


Figure 16: Manage MegaRAID Advanced Software Options Wizard

NOTE: When you click the **Advanced Software Options** menu in the main WebBIOS screen, if rehosting is not required the **Manage MegaRAID Advanced Software Options** screen as shown in Figure 16 will appear, otherwise if the user decides to go for the rehosting process, the Section 4.4.8, *Confirm Rehosting Process* screen will appear.

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

- The **Advanced Software** 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 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.

2. Click the **LSI Advanced Software License Management Portal** link to obtain the license authorization code and activation key.

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

3. Click **Activate**.

   The **Advanced Software Options Summary** wizard appears as shown in Figure 22.

4. Click **Configure Key Vault**.

   The **Confirm Rehosting process** wizard appears as shown in Figure 27.

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

- Scenario #1

  When features have been transferred from NVRAM to keyvault, and no rehosting is required, the **Configure Key Vault** button will be displayed.

- Scenario# 2

  When the rehosting process needs to be completed, the **Configure Key Vault** button will be displayed.

5. Click **Deactivate All Trial Software**.

   The **Deactivate All Trial Advanced Software Options** dialog box appears as shown in the following figure.

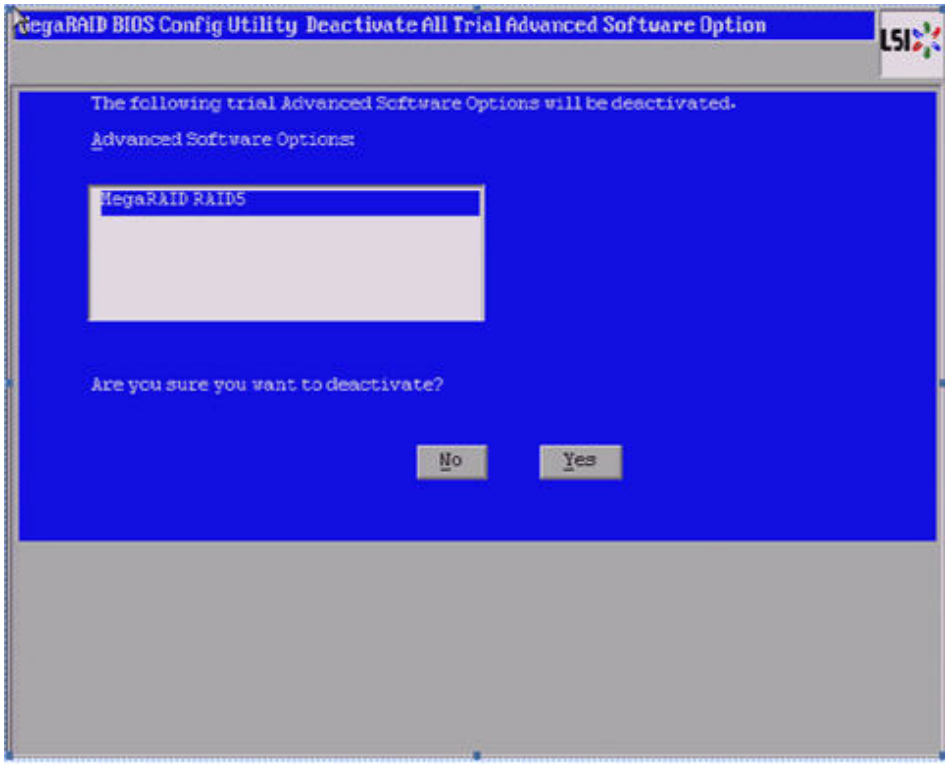


Figure 17: Deactivate All Trial Advanced Software Options

If you want to *deactivate* the software that is being used with a trial key, then click *Yes*; otherwise, click *No*.

When the activation key is improper in the Activation field in the **Advanced Software Options** wizard (Figure 16), 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

- Scenario # 2

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

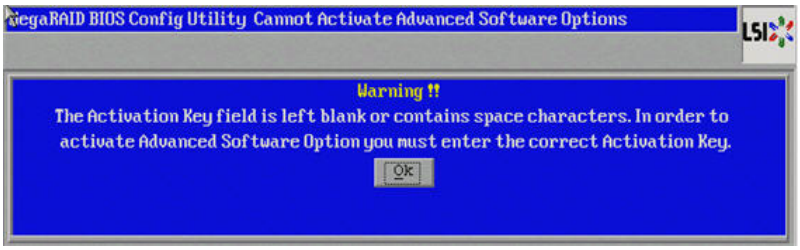


Figure 19: Activation key Blank or Space

- 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

4.4.2 Reusing the Activation Key

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

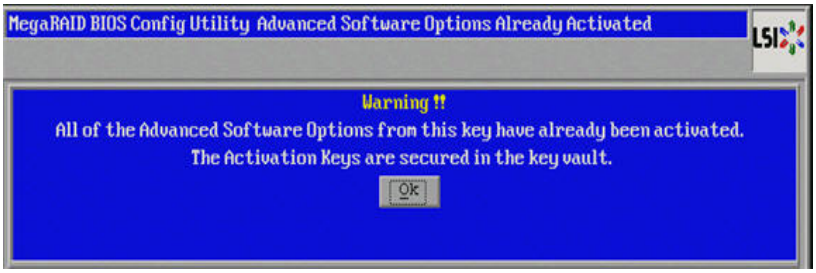


Figure 21: Reusing the Activation Key

4.4.3 Managing Advanced Software Summary

When you click *Activate* in **Advanced Software Options** wizard (Figure 16), the **Advanced Software Options Summary** wizard appears as shown in the following figure.

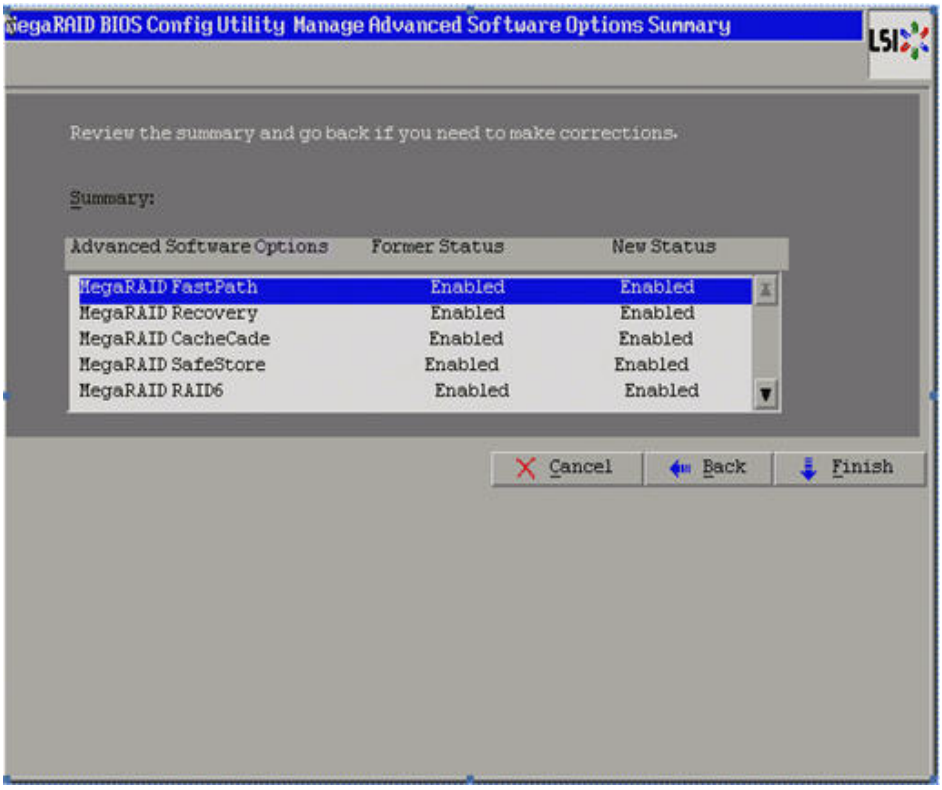


Figure 22: Advanced Software Options Summary Wizard

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 message, "*Review the summary and go back if you need to make corrections*" appears as shown in the following figure.

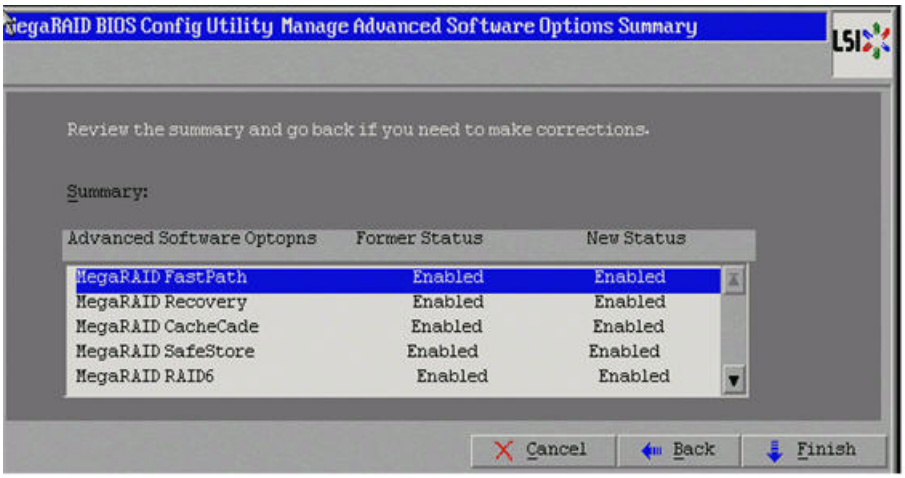


Figure 23: Activating an Unlimited Key over a Trial Key

4.4.5 Activating a Trial Software

When you activate a trial software, the message "*This trial software expires in 30 days*", appears as indicated (yellow text) in the following figure.

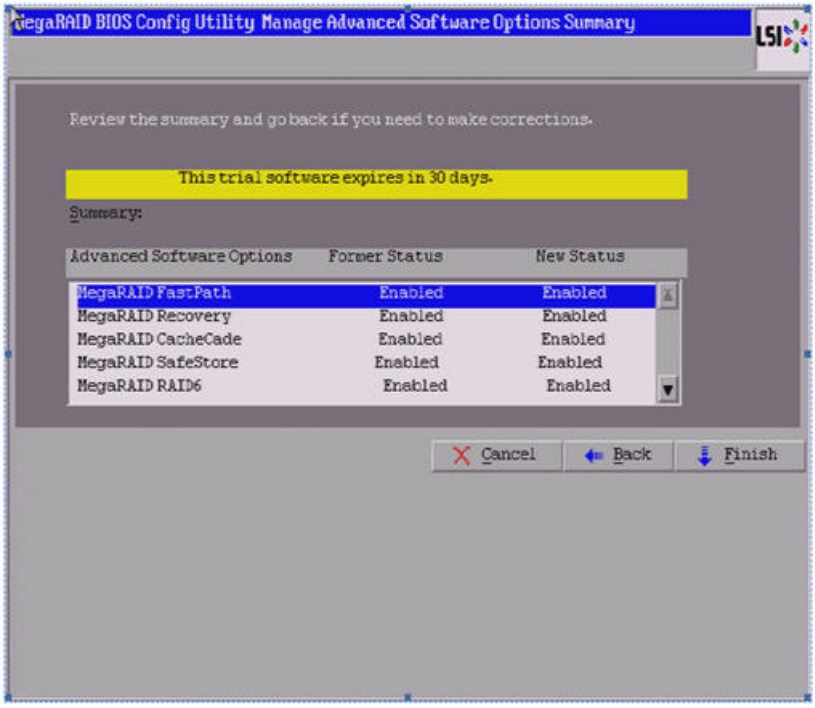


Figure 24: Activating a Trial Software

4.4.6 Activating an Unlimited Key

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

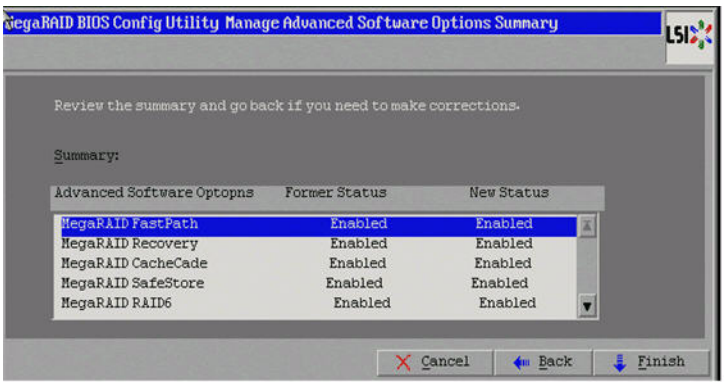


Figure 25: Activating an Unlimited Key

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 in Figure 16, the **Secure MegaRAID Advanced Software Options** dialog box appears as shown in the following figure.

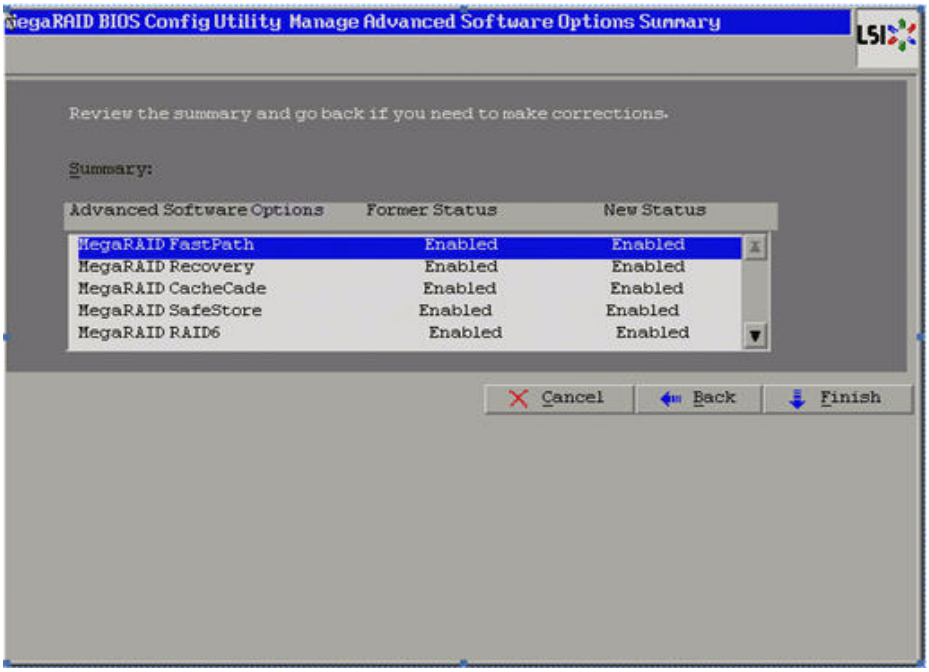


Figure 26: Securing MegaRAID Advanced Software

4.4.8 Confirm Rehosting Process

Confirming rehosting process involves the process of transferring or rehosting the advanced software features from one controller to another.

Perform the following steps to confirm the rehosting process.

1. Click the **Configure Key Vault** button in the **Advanced Software Options** wizard in Figure 16.

   The **Confirm Rehosting Process** wizard appears as shown in the following figure.

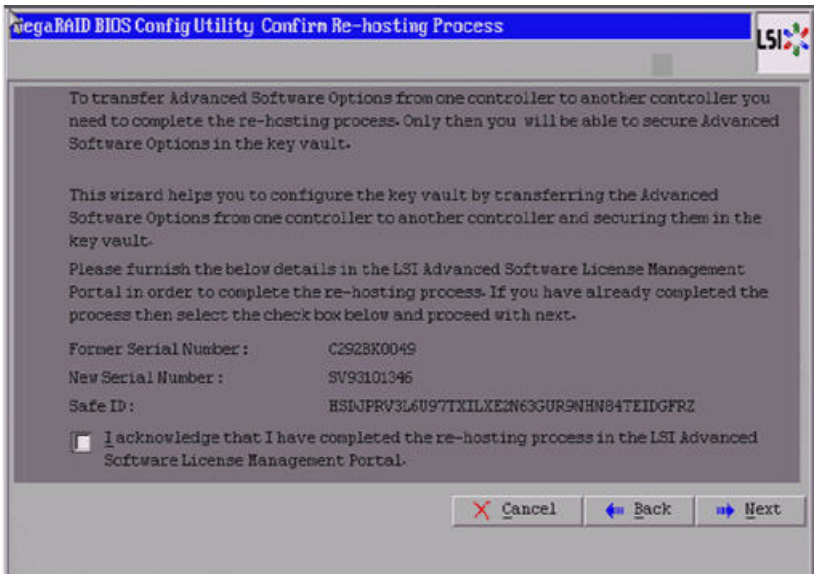


Figure 27: Confirm Rehosting Process

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

   When you click **Next**, either of the following scenarios might occur.

- Scenario # 1

  During the rehosting process, when no features need to be transferred from NVRAM to Key Vault, but a new Key Vault is connected from controller 1 to controller 2, and in the controller 2 NVRAM no features are to be transferred, the MegaRAID Advanced Software Options wizard appear.

- Scenario # 2

  When you need to transfer the features from controller 1 to controller 2, and in controller 2 NVRAM, if there are some features that need to be transferred to Key Vault, the Rehosting Process - Complete screen appears, as shown in Figure 28.

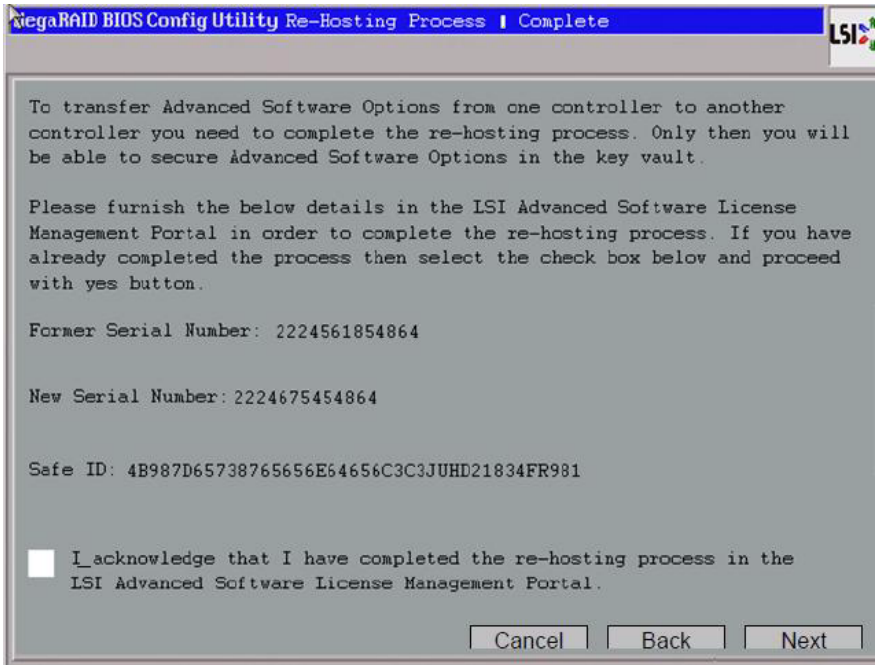


Figure 28: Rehosting Process Complete

4. Select the check box next to the message "*I acknowledge that I have completed the re-hosting process in the LSI Advanced Software License Management Portal*".

   The rehosting process is completed.

NOTE: If you click **Next** in the **Rehosting Process Complete** screen, if rehosting is not complete, the features will not copied into the keyvault, and the features will remain in the itself , but the user can still use the advanced features.

4.5 Creating a Storage Configuration

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

4.5.1 Selecting the Configuration with the Configuration Wizard

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

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

   The first Configuration Wizard screen appears, as shown in Figure 29.

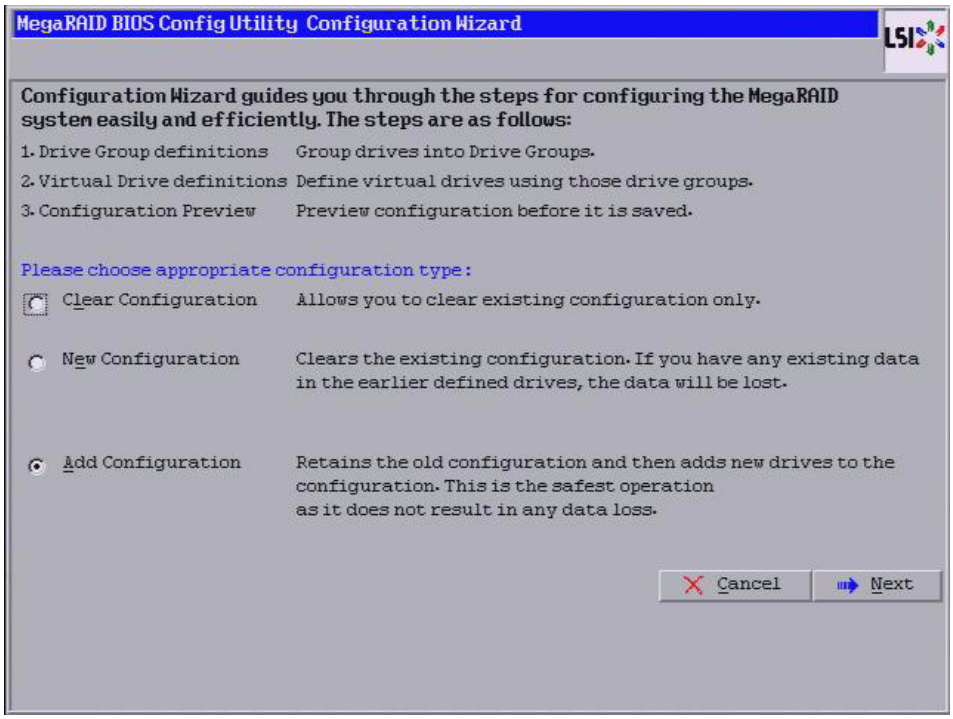


Figure 29: WebBIOS Configuration Wizard Screen

2. Select a configuration option.

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

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

3. Click **Next**.

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

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

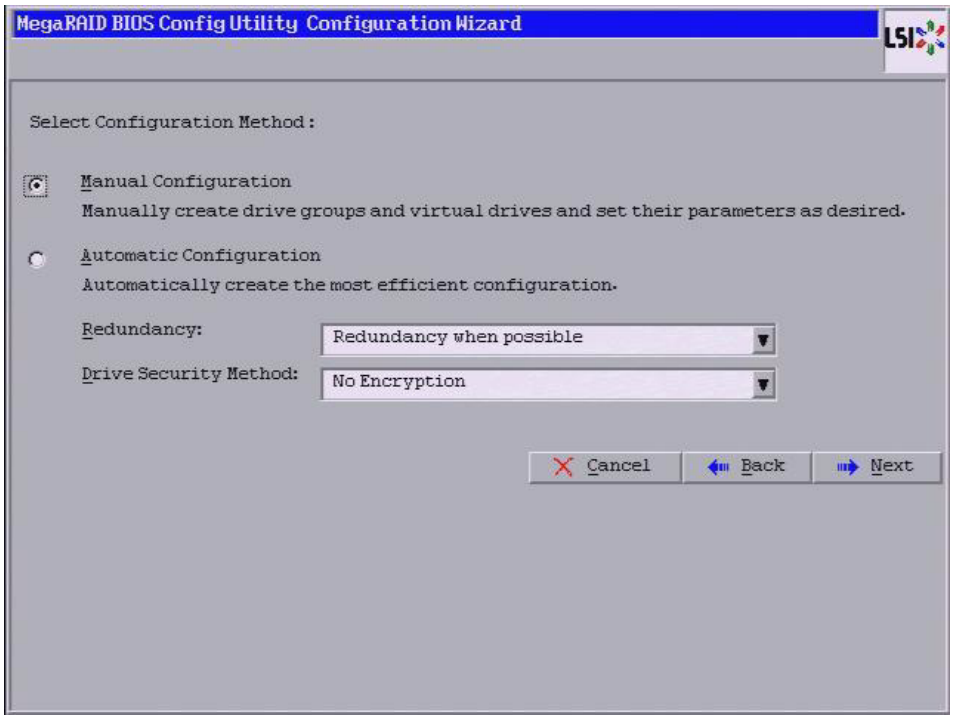


Figure 30: WebBIOS Configuration Method Screen

4. On this screen, select a configuration mode:
   - **Manual Configuration:** Allows you to control all attributes of the new storage configuration as you create drive groups and virtual drives, and set their parameters.
   - **Automatic Configuration:** Automatically creates an optimal RAID configuration.
   - If you select Automatic Configuration, you can choose whether to create a redundant RAID drive group or a non-redundant RAID 0 drive group. Select one of the following options in the Redundancy field:

     Redundancy when possible

     No redundancy
   - 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 field:

     No Encryption

     Drive Encryption
5. Click **Next** to continue.

If you select the Automatic Configuration option, continue with Section 4.5.3, *Using Automatic Configuration*. If you select Manual Configuration, continue with Section 4.5.4, *Using Manual Configuration*.

4.5.2 Converting JBOD to Unconfigured Good Drives

Perform the following steps to convert JBOD to Unconfigured Good drives.

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

2. Click **Next** to view the second Controller Properties screen.
3. Click **Next** to view the third Controller Properties screen.
4. Click **Manage JBOD** from the Third Controller Properties Screen (See Figure 87).

   The **JBOD Conversion** screen appears as shown in the following figure.

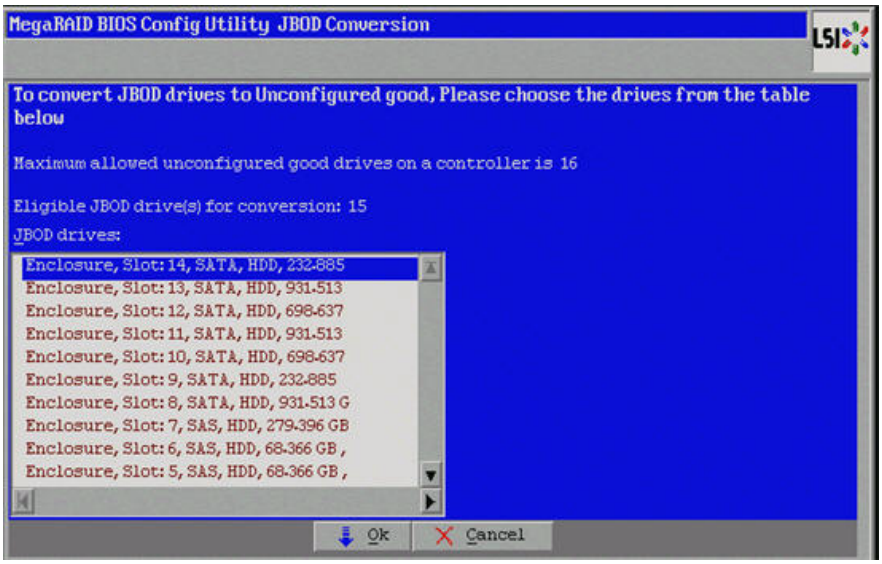


Figure 31: JBOD Conversion Screen

5. Select the **drives** which you want to be converted as Unconfigured Good.
6. Click **OK**.

   The selected drives will be converted to Unconfigured Good.

NOTE: For more information read the tips provided on the JBOD drives screen.

4.5.3 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 screen, and click **Accept** to accept it. (Or click **Back** to go back and change the configuration.)
   - **RAID 0:** If you select Automatic Configuration and No Redundancy, WebBIOS creates a RAID 0 configuration.
   - **RAID 1:** If you select **Automatic Configuration** and **Redundancy when possible**, and only two drives are available, WebBIOS creates a RAID 1 configuration.
   - **RAID 5**: If you select **Automatic Configuration** and **Redundancy when possible**, and three or more drives are available, WebBIOS creates a RAID 5 configuration.
   - **RAID 6**: If you select **Automatic Configuration** and **Redundancy when possible**, and the RAID 6 option is enabled, and three or more drives are available, WebBIOS creates a RAID 6 configuration.

2. Click **Yes** when you are prompted to save the configuration.
3. Click **Yes** when you are prompted to initialize the new virtual drive(s).

   WebBIOS CU begins a background initialization of the virtual drives.

   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.

4.5.4 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.

4.5.4.1 Using Manual Configuration: RAID 0

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

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

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

1. Hold <Ctrl> while selecting two or more ready drives in the Drives panel on the left until you have selected all desired drives for the drive group.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the 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 drive encryption.

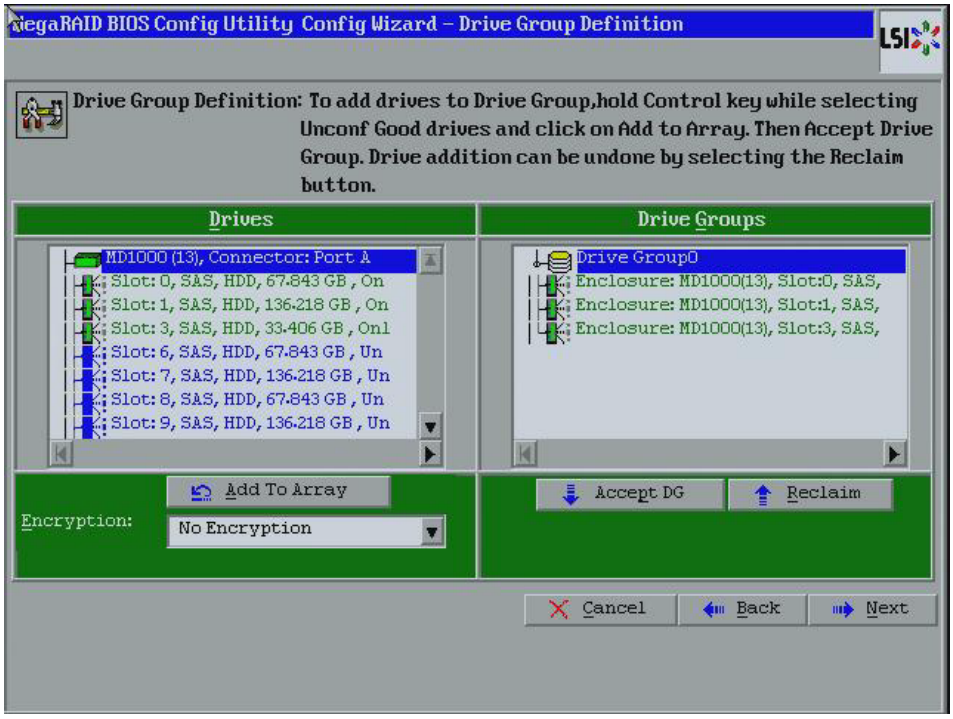


Figure 32: WebBIOS Drive Group Definition Screen

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

   The Virtual Drive Definition screen appears, as shown in Figure 33. This screen lists the possible RAID levels for the drive group.

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

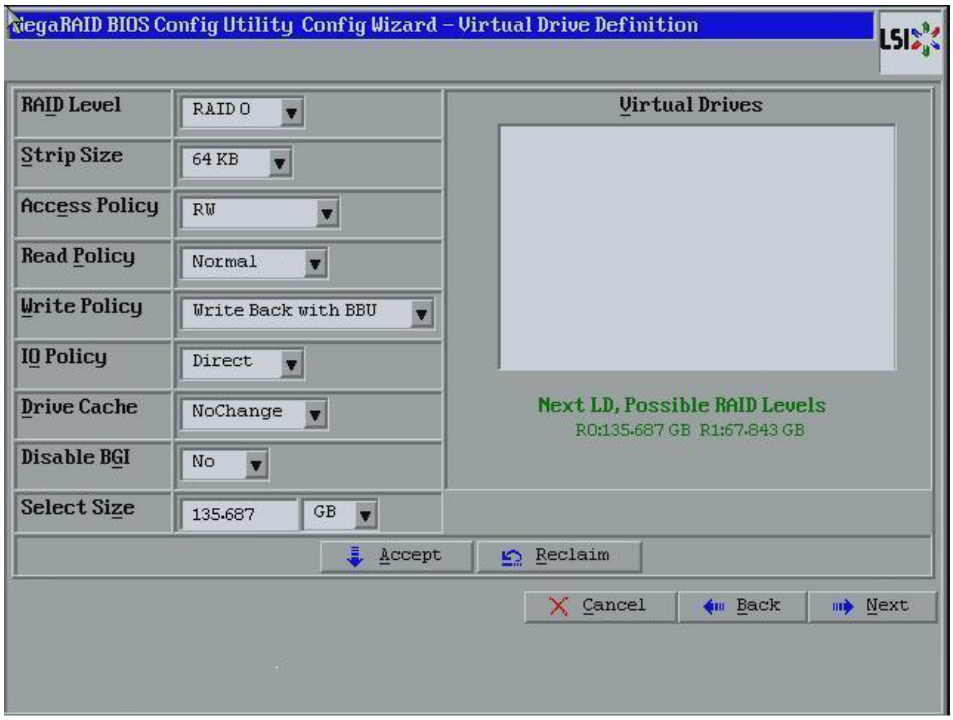


Figure 33: WebBIOS Virtual Drive Definition Screen

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

   Here are brief explanations of the virtual drive options:

   — **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 0.

   — **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, 16, 32, 64, 128, 256, 512, 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 is the default.

   Ahead: This option enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.

— **Write Policy:** Specify the write policy for this virtual drive:

WBack: In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.

WThru: In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.

Bad BBU: Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

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

— **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.

Direct: In Direct I/O mode, the reads are not automatically cached in controller cache and the data is transferred directly to the host except when 1) the data is buffered in controller memory in order to execute parity; or 2) the data is transferred from the controller cache because of read-ahead algorithms. This is the default.

Cached: In Cached I/O mode, all reads are buffered in the controller's cache memory or in the CacheCade virtual drive.

— **Drive Cache:** Specify the drive cache policy:

Enable: Enable the drive cache.

Disable: Disable the drive cache.

NoChange: Leave the current drive cache policy as is. This is the default.

— **Disable BGI:** Specify the background initialization status:

No: Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.

Yes: Select *Yes* if you do not want to allow background initializations for configurations on this controller.

— **Select Size:** Specify the size of the virtual drive in megabytes. Normally, this would be the full size for RAID 0 shown in the Configuration panel on the right. You can specify a smaller size if you want to create other virtual drives on the same drive group.

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

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

8. Click **Next** after you finish defining the virtual drives.

   The Configuration Preview screen appears, as shown in Figure 34.

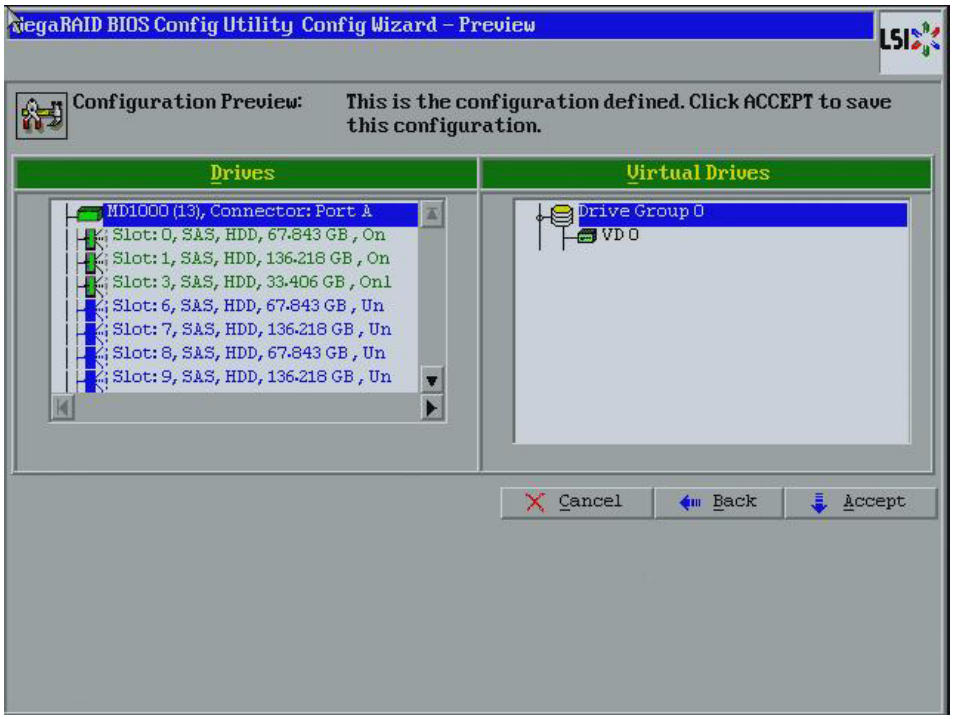


Figure 34: RAID 0 Configuration Preview

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

    The WebBIOS main menu appears.

4.5.4.2 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 screen appears. You use this screen to select drives to create drive groups.

1. Hold <Ctrl> while you select two ready drives in the Drives panel on the left. You must select an even number of drives.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the Drive Groups panel on the right, as shown in Figure 35.

   If you need to undo the changes, click **Reclaim**.
3. 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.

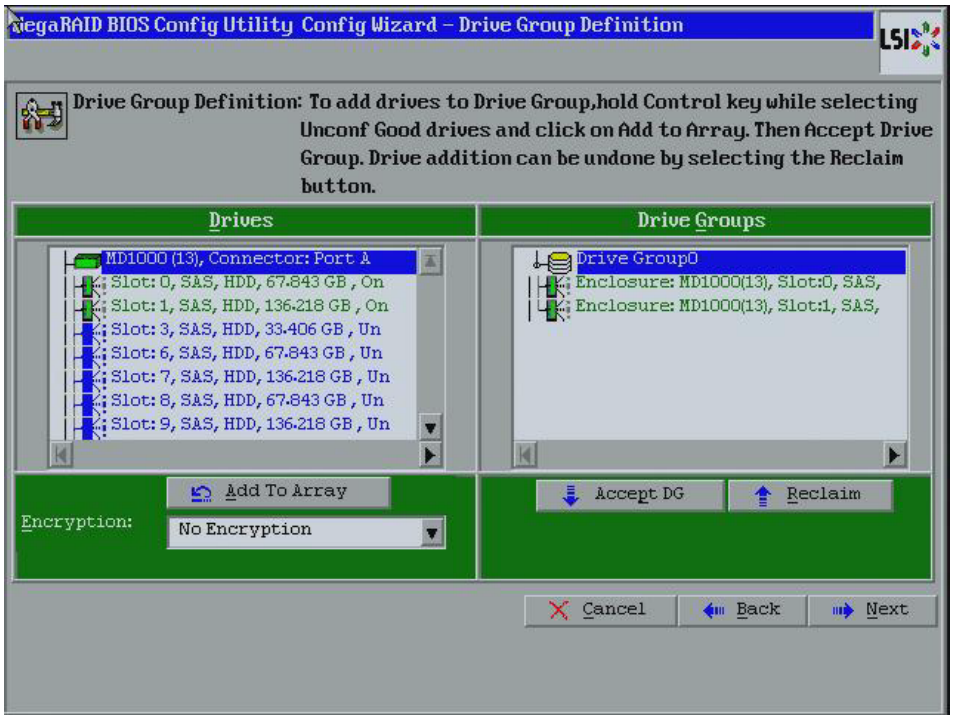


Figure 35: WebBIOS Drive Group Definition Screen

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

   The Virtual Drive Definition screen appears, as shown in Figure 36. You use this screen to select the RAID level, strip size, read policy, and other attributes for the new virtual drives.

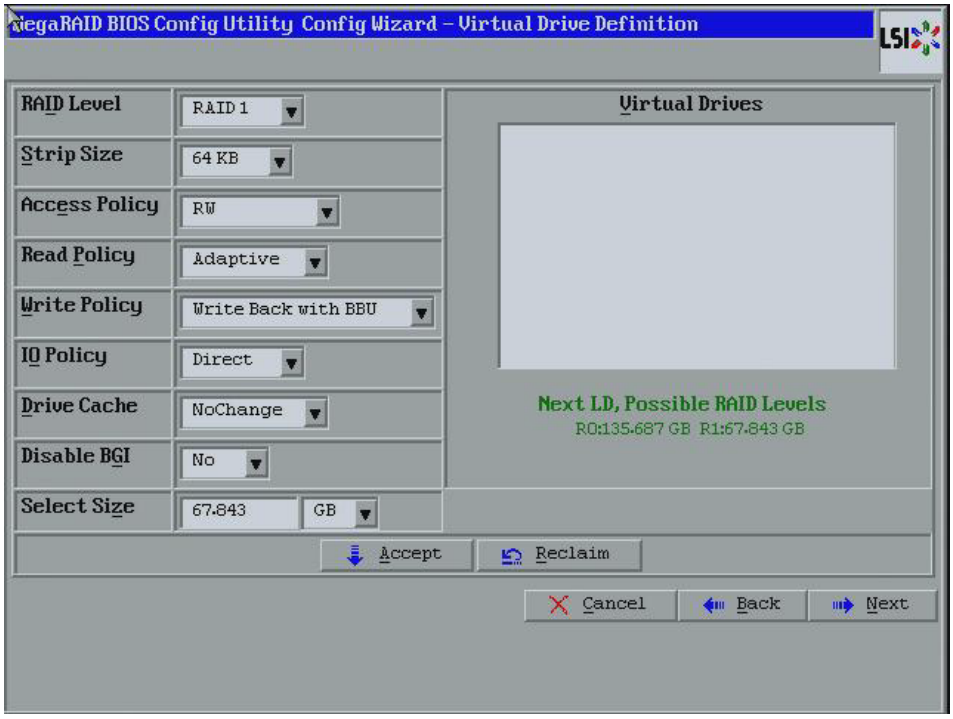


Figure 36: WebBIOS Virtual Drive Definition Screen

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

   Here are brief explanations of the virtual drive options:

   — **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 1.

   — **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, 16, 32, 64, 128, 256, 512, 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 is the default.

   Ahead: This option enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.

— **Write Policy:** Specify the write policy for this virtual drive:

WBack: In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction.
This setting is recommended in Standard mode.

WThru: In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction.
This is the default.

Bad BBU: Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

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

— **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.

Direct: In Direct I/O mode, the reads are not automatically cached in controller cache and the data is transferred directly to the host except when 1) the data is buffered in controller memory in order to execute parity; or 2) the data is transferred from the controller cache because of read-ahead algorithms. This is the default.

Cached: In Cached I/O mode, all reads are buffered in the controller's cache memory or in the CacheCade virtual drive.

— **Drive Policy:** Specify the drive cache policy:

Enable: Enable the drive cache.

Disable: Disable the drive cache.

NoChange: Leave the current drive cache policy as is. This drive policy is the default.

— **Disable BGI:** Specify the background initialization status:

No: Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.

Yes: Select *Yes* if you do not want to allow background initializations for configurations on this controller.

— **Select Size:** Specify the size of the virtual drive(s) in megabytes. Normally, 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.

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

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

8. Click **Next** after you finish defining the virtual drives.

   The Configuration Preview screen appears, as shown in Figure 37.

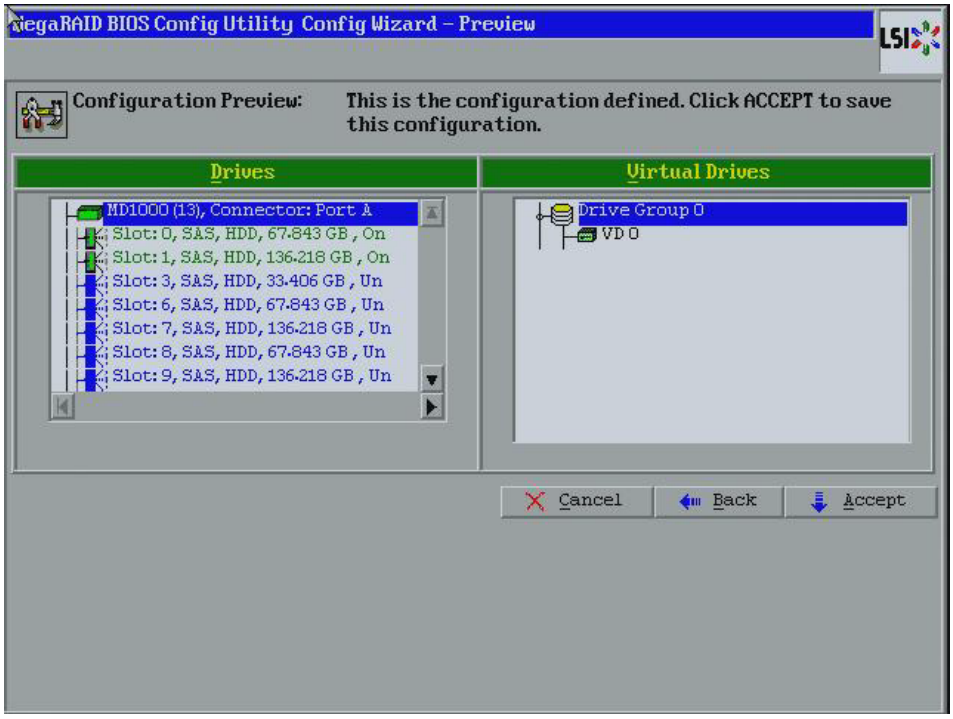


Figure 37: RAID 1 Configuration Preview

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

    The WebBIOS main menu appears.

4.5.4.3 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 recreate all missing information. You can use RAID 5 for office automation and online customer service that require fault tolerance.

In addition, RAID 5 is good for any application that has high read request rates but low write request rates.

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

1. Hold <Ctrl> while you select at least three ready drives in the Physical Drives panel on the left.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the Drive Groups panel on the right, as shown in Figure 38.

   If you need to undo the changes, click **Reclaim**.
3. Choose whether to use drive encryption.

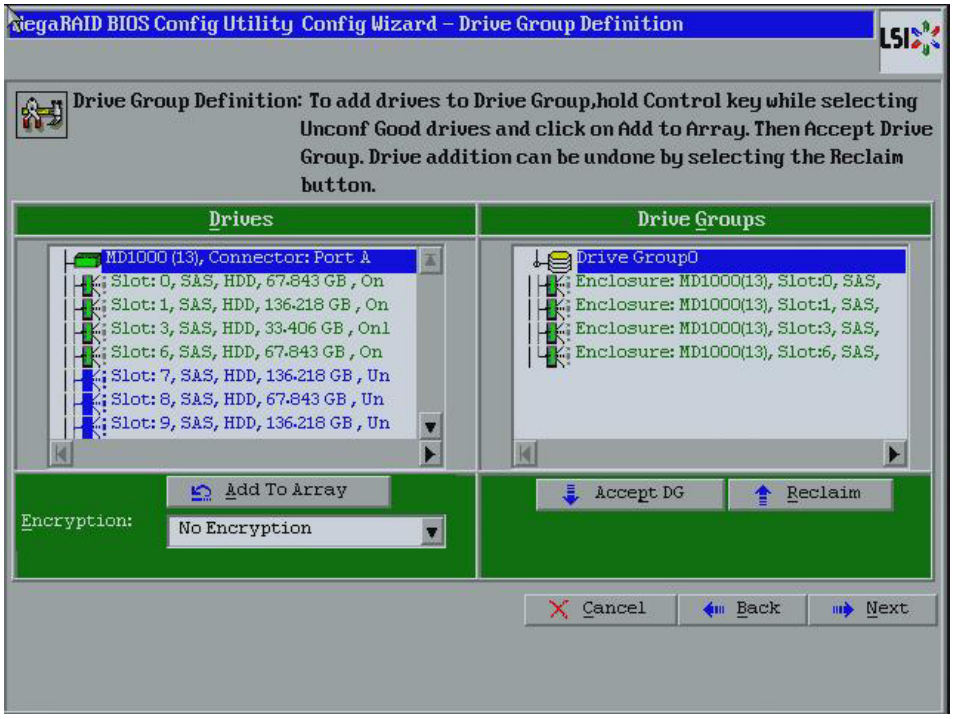


Figure 38: WebBIOS Drive Group Definition Screen

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

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

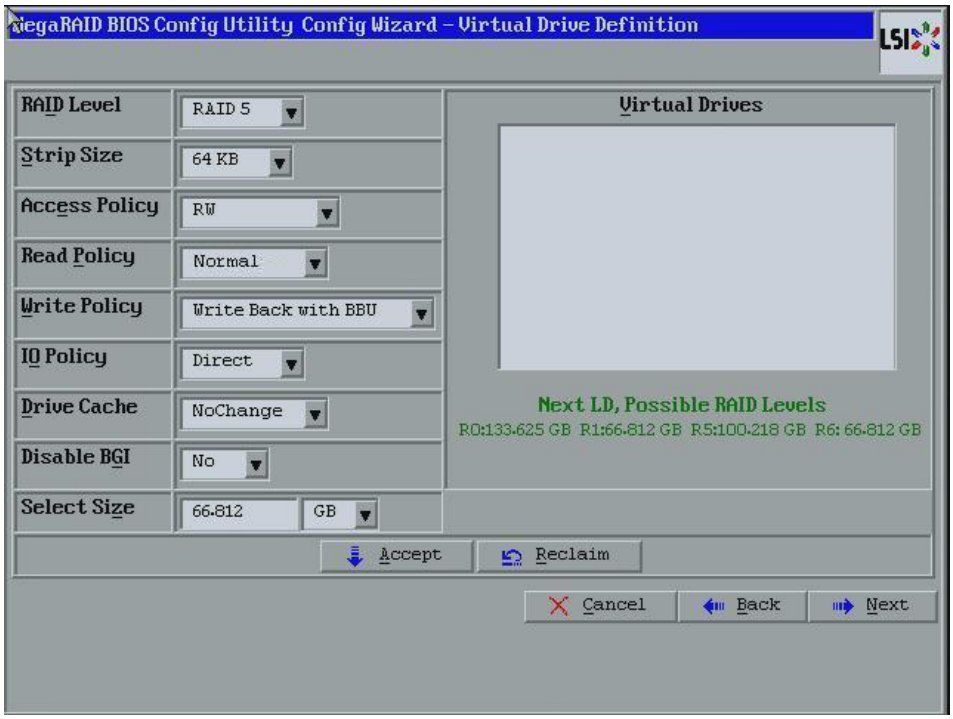


Figure 39: WebBIOS Virtual Drive Definition Screen

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

   Here are brief explanations of the virtual drive options:

   — **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 5.

   — **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, 16, 32, 64, 128, 256, 512, 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 disables the read ahead capability. This is the default.

Ahead: This enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.

— **Write Policy:** Specify the write policy for this virtual drive:

WBack: In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.

WThru: In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.

Bad BBU: Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

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

— **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.

Direct: In Direct I/O mode, the reads are not automatically cached in controller cache and the data is transferred directly to the host except when 1) the data is buffered in controller memory in order to execute parity; or 2) the data is transferred from the controller cache because of read-ahead algorithms. This is the default.

Cached: In Cached I/O mode, all reads are buffered in the controller's cache memory or in the CacheCade virtual drive.

— **Drive Policy:** Specify the drive cache policy:

Enable: Enable the drive cache.

Disable: Disable the drive cache.

NoChange: Leave the current drive cache policy as is.
This drive policy is the default.

— **Disable BGI:** Specify the background initialization status:

No: Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.

Yes: Select *Yes* if you do not want to allow background initializations for configurations on this controller.

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 megabytes. Normally, this 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.

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

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

8. Click **Next** after you finish defining the virtual drives.

   The Configuration Preview screen appears, as shown in Figure 40.

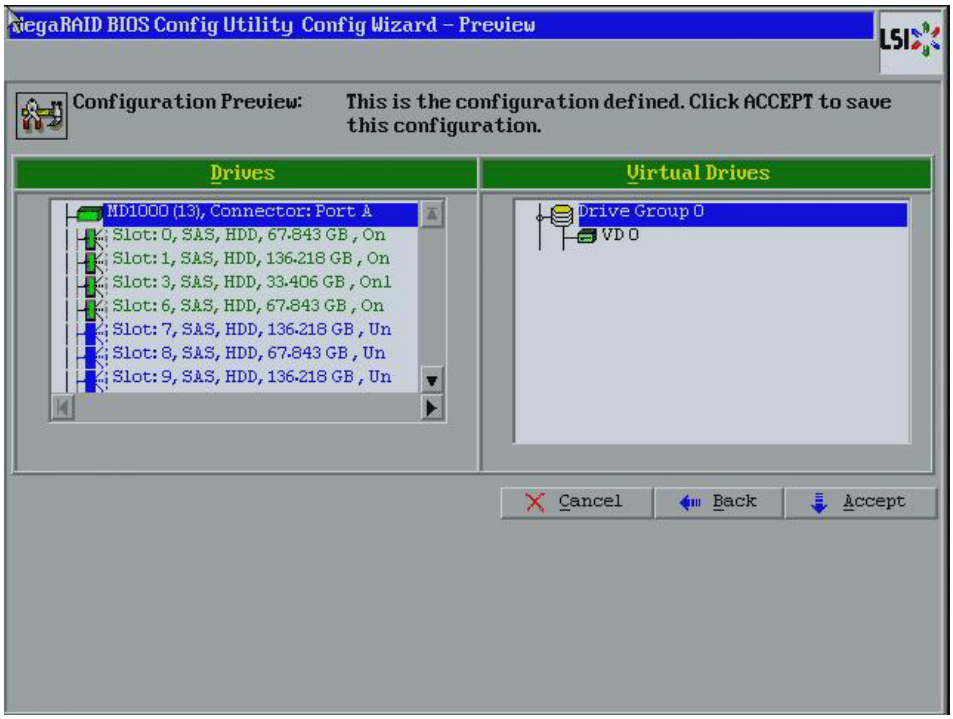


Figure 40: RAID 5 Configuration Preview

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

    The WebBIOS main menu appears.

4.5.4.4 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 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 (IMR) displays new drives as JBOD (Just a Bunch of Disks). For MegaRAID, unless the inserted drive contains valid DDF metadata, new drives display as JBOD for MegaRAID Entry level controllers, such as the SAS 9240-4i/8i. 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 drive Group Definition screen appears. You use this screen to select drives to create drive groups.

1. Hold <Ctrl> while selecting at least three ready drives in the Drives panel on the left.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the Drive Groups panel on the right, as shown in Figure 41.

   If you need to undo the changes, click **Reclaim**.
3. Choose whether to use drive encryption.

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

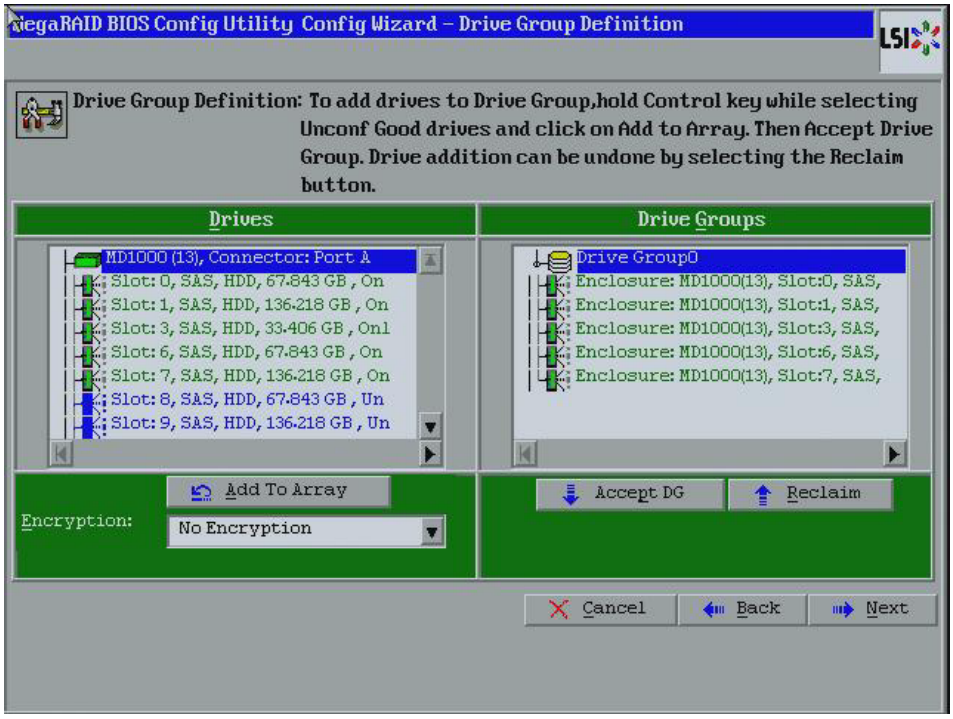


Figure 41: WebBIOS Drive Group Definition Screen

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

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

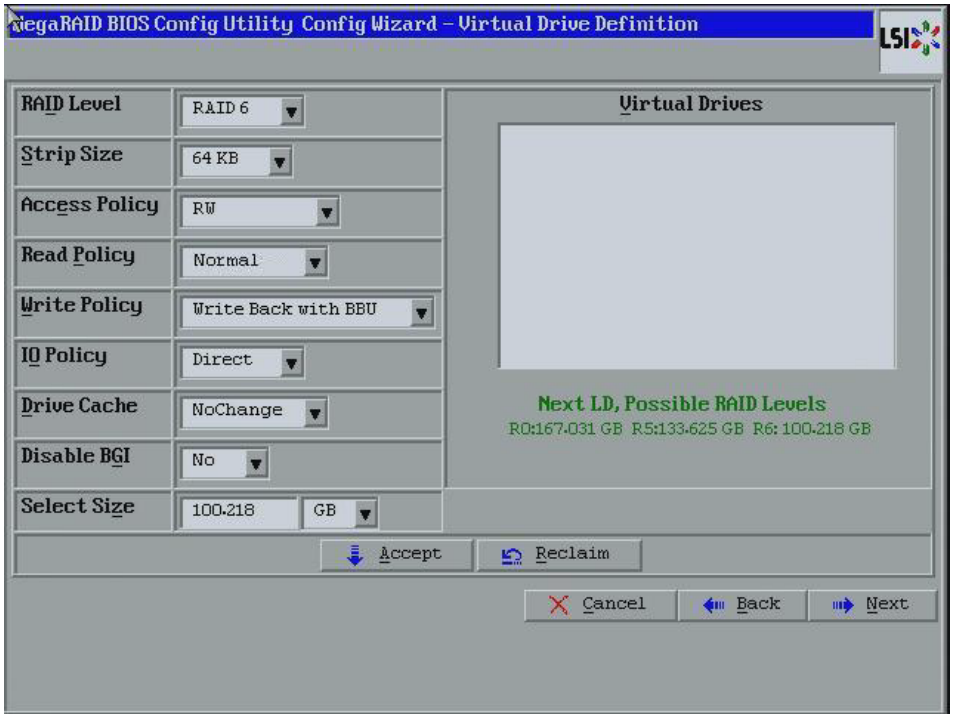


Figure 42: WebBIOS Virtual Drive Definition Screen

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

   Here are brief explanations of the virtual drive options:

   — **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 6.

   — **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, 16, 32, 64, 128, 256, 512, 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.

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

   — **Access Policy:** Select the type of data access that is allowed for this virtual drive:

   RW: Allow read/write access. This is the default.

   Read Only: Allow read-only access.

   Blocked: Do not allow access.

   — **Read Policy:** Specify the read policy for this virtual drive:

   Normal: This disables the read ahead capability. This is the default.

Ahead: This enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.

— **Write Policy:** Specify the write policy for this virtual drive:

WBack: In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction.
This setting is recommended in Standard mode.

WThru: In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.

Bad BBU: Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

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

— **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.

Direct: In Direct I/O mode, the reads are not automatically cached in controller cache and the data is transferred directly to the host except when 1) the data is buffered in controller memory in order to execute parity; or 2) the data is transferred from the controller cache because of read-ahead algorithms. This is the default.

Cached: In Cached I/O mode, all reads are buffered in the controller's cache memory or in the CacheCade virtual drive.

— **Drive Policy:** Specify the drive cache policy:

Enable: Enable the drive cache.

Disable: Disable the drive cache.

NoChange: Leave the current drive cache policy as is. This drive policy is the default.

— **Disable BGI:** Specify the background initialization status:

No: Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.

Yes: Select *Yes* if you do not want to allow background initializations for configurations on this controller.

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 megabytes. Normally, 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.

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

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

3. Click **Next** after you finish defining the virtual drives.

   The Configuration Preview screen appears, as shown in Figure 43.

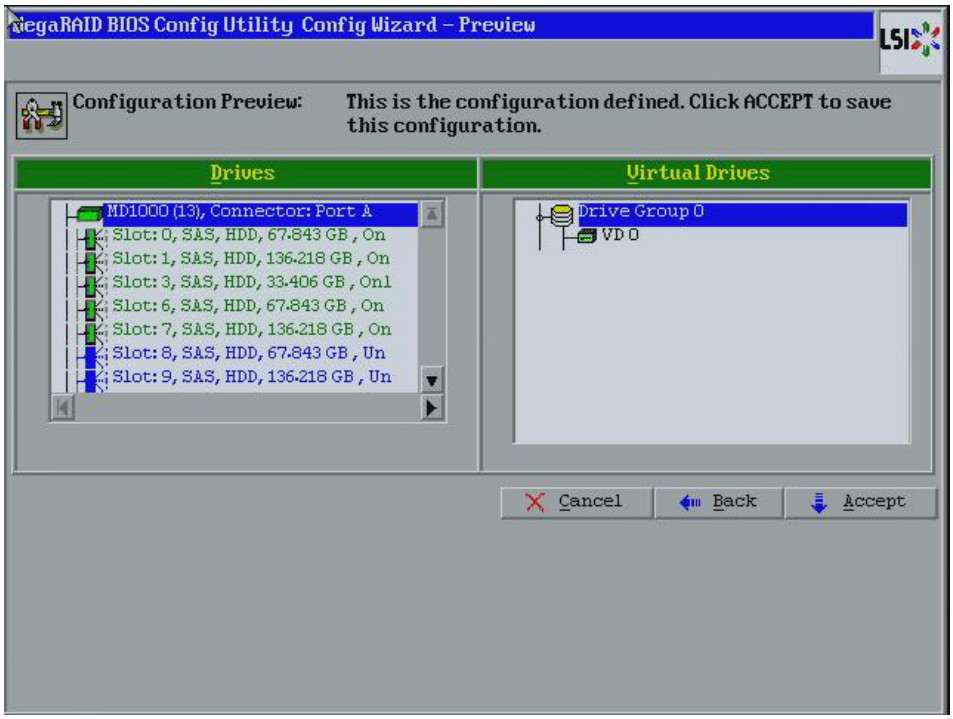


Figure 43: RAID 6 Configuration Preview

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

   The WebBIOS main menu appears.

4.5.4.5 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 Drive Group Definition screen appears.

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

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

   If you need to undo the changes, click **Reclaim**.
3. Click **Accept DG** to create a RAID 0 drive group.

   An icon for the next drive group appears in the right panel.
4. Hold <Ctrl> while you select more ready drives in the Drives panel to create a second RAID 0 drive group.
5. Click **Add To Array** to move the drives to a second drive group configuration in the 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.)

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

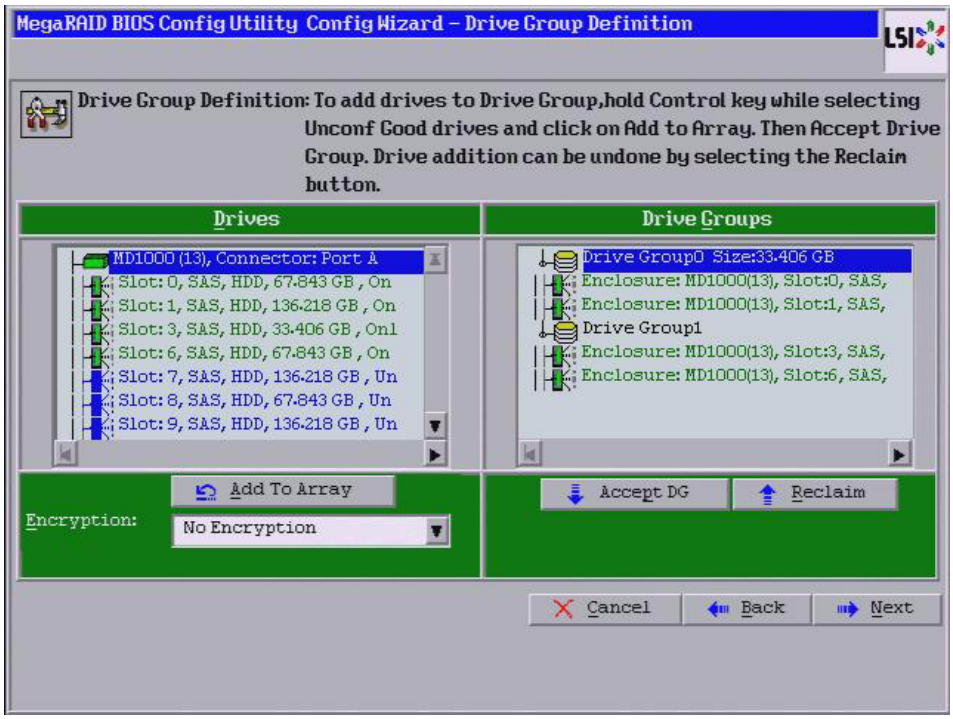


Figure 44: WebBIOS Drive Group Definition Screen

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

    The Span Definition screen appears, as shown in Figure 45. This screen shows the drive group holes that you can select to add to a span.

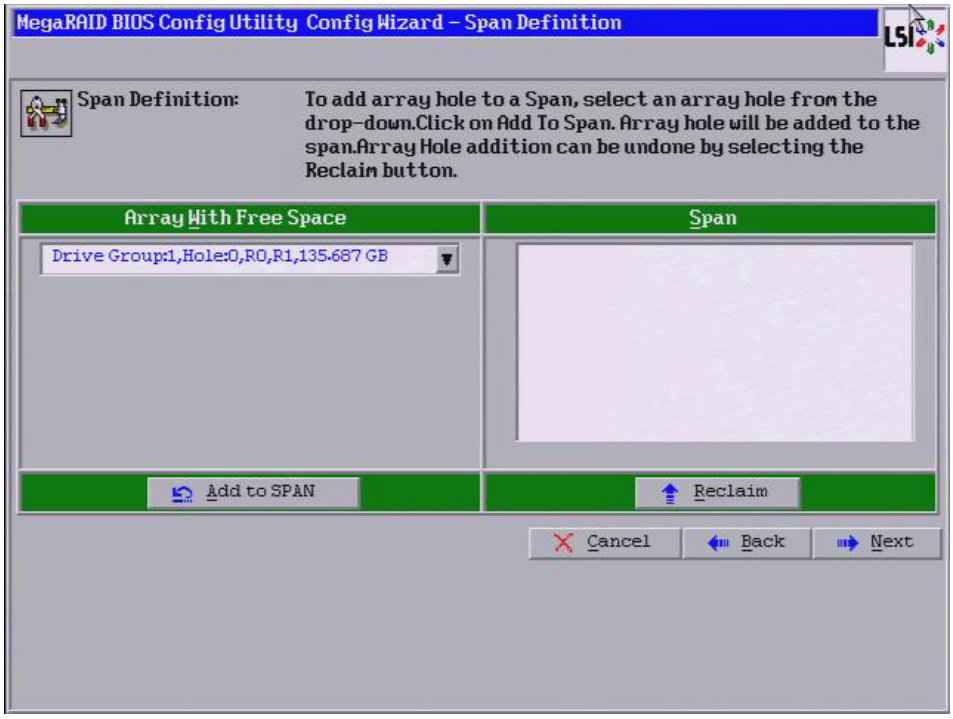


Figure 45: WebBIOS Span Definition Screen

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

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

12. Hold <Ctrl> while you select a second drive group, and then click **Add to SPAN**.
13. Repeat the previous two steps until you have selected all of the drive groups that you want.
14. Click **Next**.

    The Virtual Drive Definition screen appears, as shown in Figure 46. You use this screen to select the RAID level, strip size, read policy, and other attributes for the new virtual drives.

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

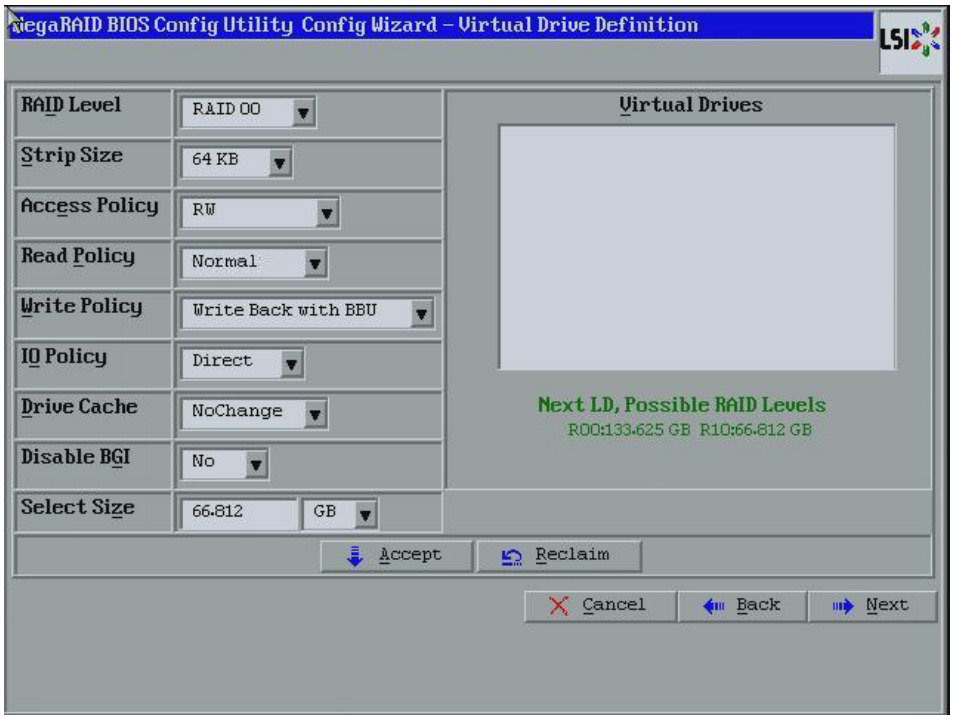


Figure 46: WebBIOS Virtual Drive Definition Screen

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

    Here are brief explanations of the virtual drive options:

    — **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 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, 16, 32, 64, 128, 256, 512, 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 type of access is the default.

    Blocked: Do not allow access.

    — **Read Policy:** Specify the read policy for this virtual drive:

    Normal: This option disables the read ahead capability. This is the default.

    Ahead: This option enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.

— **Write Policy:** Specify the write policy for this virtual drive:

WBack: In Writeback mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.

WThru: In Writethrough mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.

Bad BBU: Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

CAUTION: LSI allows Writeback mode to be used with or without a battery. To protect the entire system, LSI recommends that you use **either** a battery to protect the controller cache or an uninterruptible power supply (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, the reads are not automatically cached in controller cache and the data is transferred directly to the host except when 1) the data is buffered in controller memory in order to execute parity; or 2) the data is transferred from the controller cache because of read-ahead algorithms. This is the default.

Cached: In Cached I/O mode, all reads are buffered in the controller's cache memory or in the CacheCade virtual drive.

— **Drive Policy:** Specify the drive cache policy:

Enable: Enable the drive cache.

Disable: Disable the drive cache.

NoChange: Leave the current drive cache policy as is. This setting is the default.

— **Disable BGI:** Specify the background initialization status:

No: Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This 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 megabytes. Normally, 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.

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

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

18. After you finish defining the virtual drives, click **Next**.

    The Configuration Preview screen appears, as shown in Figure 47.

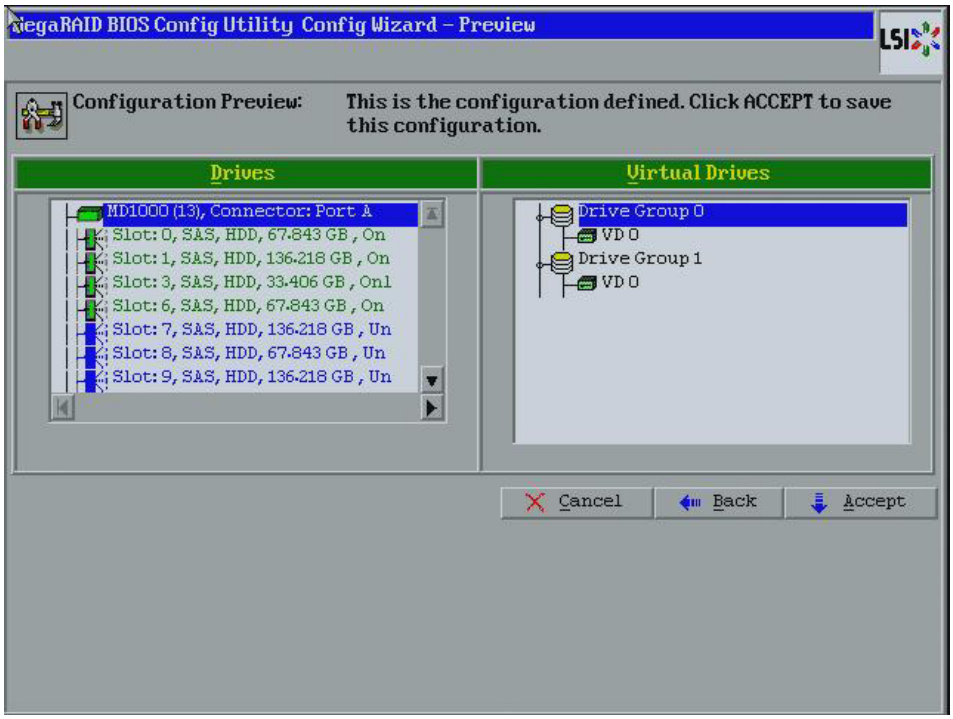


Figure 47: RAID 00 Configuration Preview

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

    The WebBIOS main menu appears.

4.5.4.6 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 screen appears.

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

1. Hold <Ctrl> while selecting two ready drives in the Drives panel on the left.

2. Click **Add To Array** to move the drives to a proposed two-drive drive group configuration in the Drive Groups panel on the right.

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

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

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

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

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

7. 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.

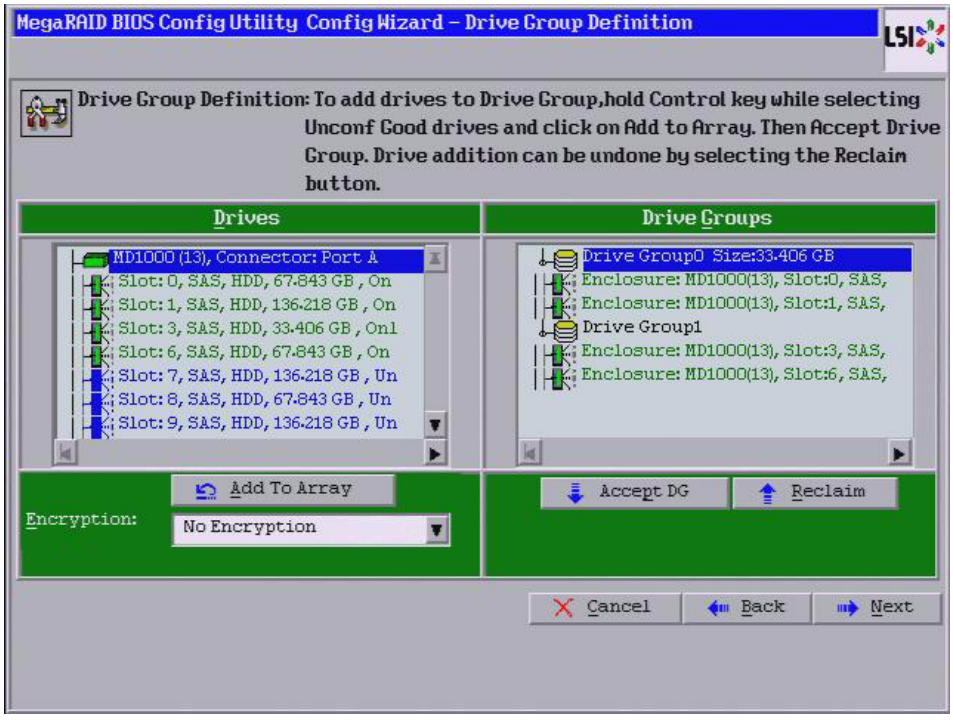


Figure 48: WebBIOS Drive Group Definition Screen

8. Repeat the previous three steps until you have selected all the drives you want for the drive groups.
9. After you finish selecting drives for the drive groups, select each drive group and click **Accept DG** for each.

10. Click **Next**.

    The Span Definition screen appears, as shown in Figure 49.
    This screen displays the drive group holes you can select to add to a span.

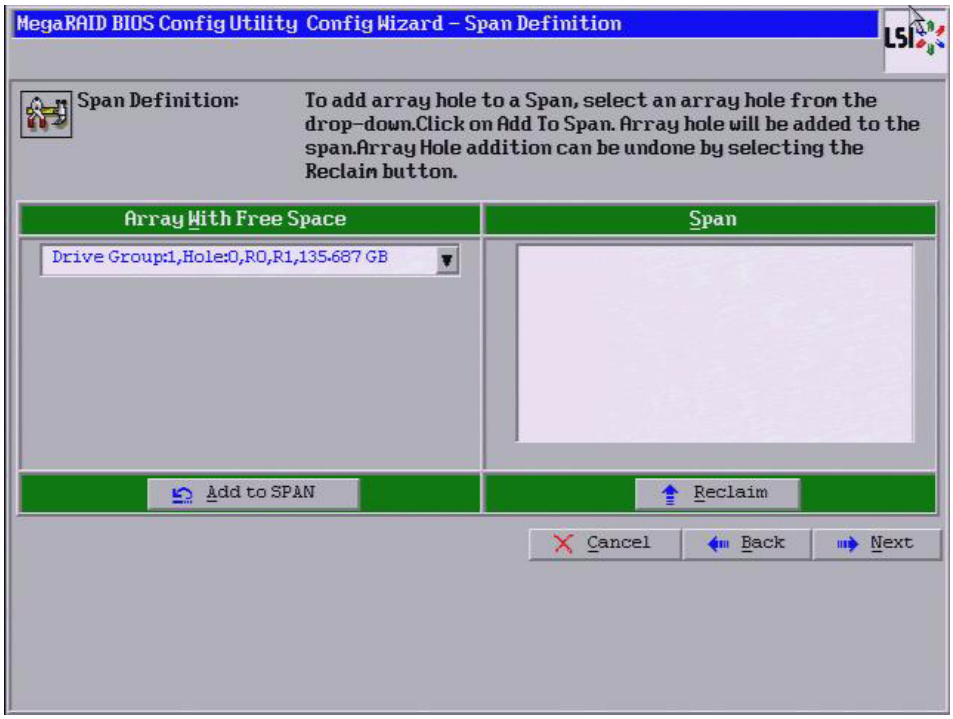


Figure 49: WebBIOS Span Definition Screen

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

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

12. Hold <Ctrl> while you select a second drive group with two drives, and click **Add to SPAN**.

    Both drive groups display in the right frame under Span.

13. If there are additional drive groups with two drives each, you can add them to the virtual drive.
14. Click **Next**.

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

15. Hold <Ctrl> while you select two drive groups with two drives in the Configuration panel on the right.

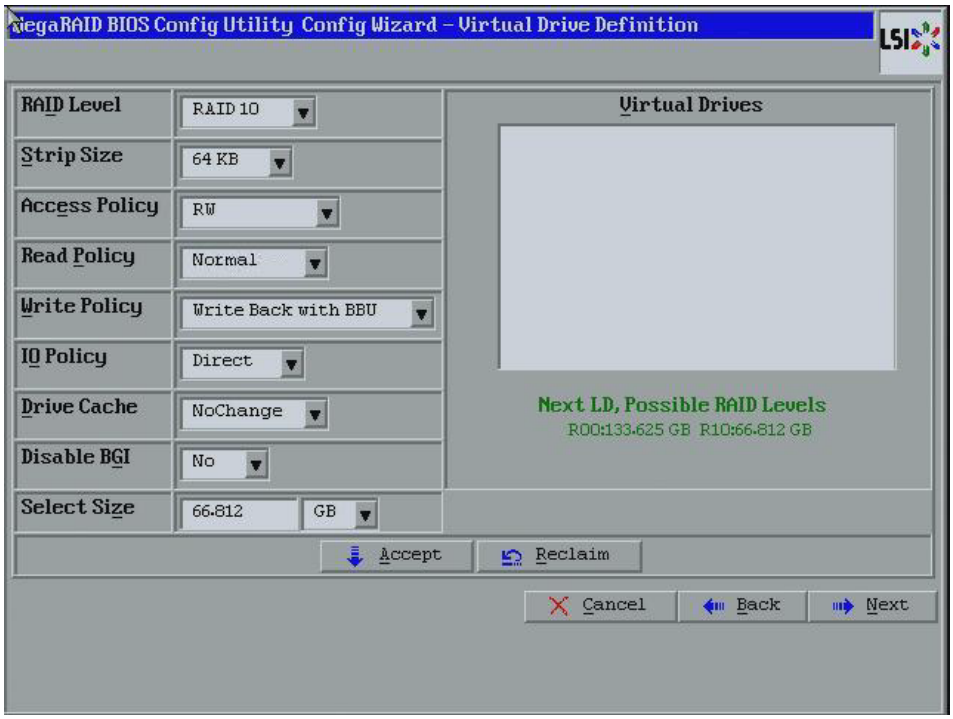


Figure 50: WebBIOS Virtual Drive Definition Screen

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 two.

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

    Here are brief explanations of the virtual drive options:

    — **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 10.

    — **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, 16, 32, 64, 128, 256, 512, 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 is the default.

    Blocked: Do not allow access.

— **Read Policy:** Specify the read policy for this virtual drive:

Normal: This disables the read ahead capability. This is the default.

Ahead: This enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.

— **Write Policy:** Specify the write policy for this virtual drive:

WBack: In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.

WThru: In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.

Bad BBU: Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

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

— **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.

Direct: In Direct I/O mode, the reads are not automatically cached in controller cache and the data is transferred directly to the host except when 1) the data is buffered in controller memory in order to execute parity; or 2) the data is transferred from the controller cache because of read-ahead algorithms. This is the default.

Cached: In Cached I/O mode, all reads are buffered in the controller's cache memory or in the CacheCade virtual drive.

— **Drive Policy:** Specify the drive cache policy:

Enable: Enable the drive cache.

Disable: Disable the drive cache.

NoChange: Leave the current drive cache policy as is. This drive policy is the default.

— **Disable BGI:** Specify the background initialization status:

No: Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.

Yes: Select *Yes* if you do not want to allow background initializations for configurations on this controller.

— **Select Size:** Specify the size of the virtual drive in megabytes. Normally, this would be the full size for RAID 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.

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

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

18. After you finish defining the virtual drives, click **Next** .

    The Configuration Preview screen appears, as shown in Figure 51.

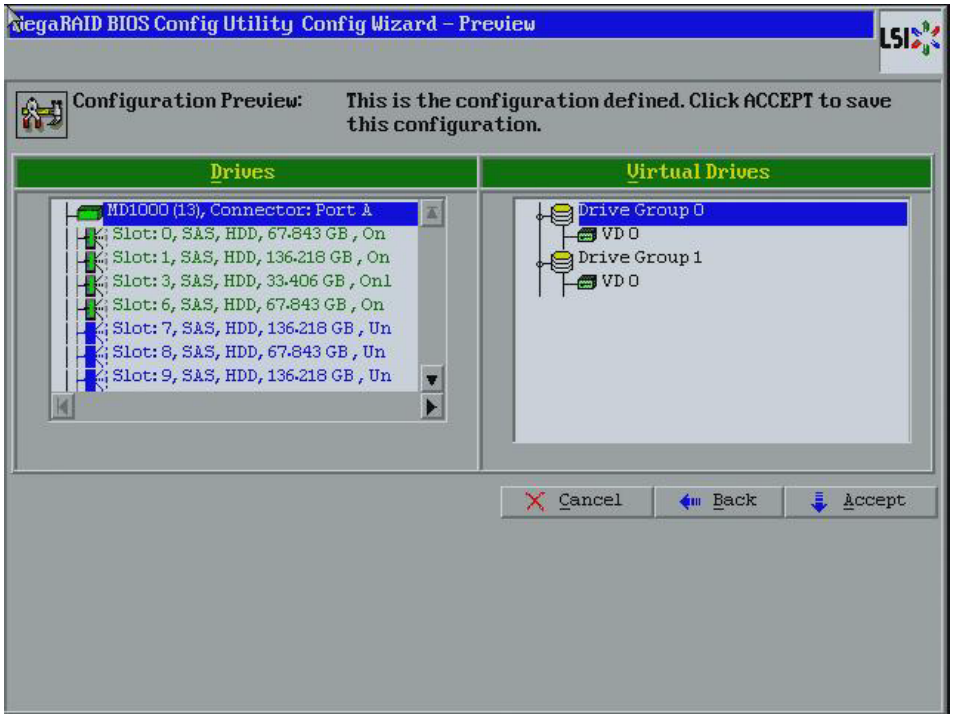


Figure 51: RAID 10 Configuration Preview

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

    The WebBIOS main menu appears.

4.5.4.7 Using Manual Configuration: RAID 50

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

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

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

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

   If you need to undo the changes, click **Reclaim**.
3. Click **Accept DG** to create a RAID 5 drive group.

   An icon for a second drive group displays in the right panel.
4. Click on the icon for the second drive group to select it.
5. Hold <Ctrl> while selecting at least three more ready drives in the Drives panel to create a second drive group.
6. Click **Add To Array** to move the drives to a proposed drive group configuration in the Drive Groups panel on the right, as shown in Figure 52.

   If you need to undo the changes, click **Reclaim**.
7. Choose whether to use drive encryption.

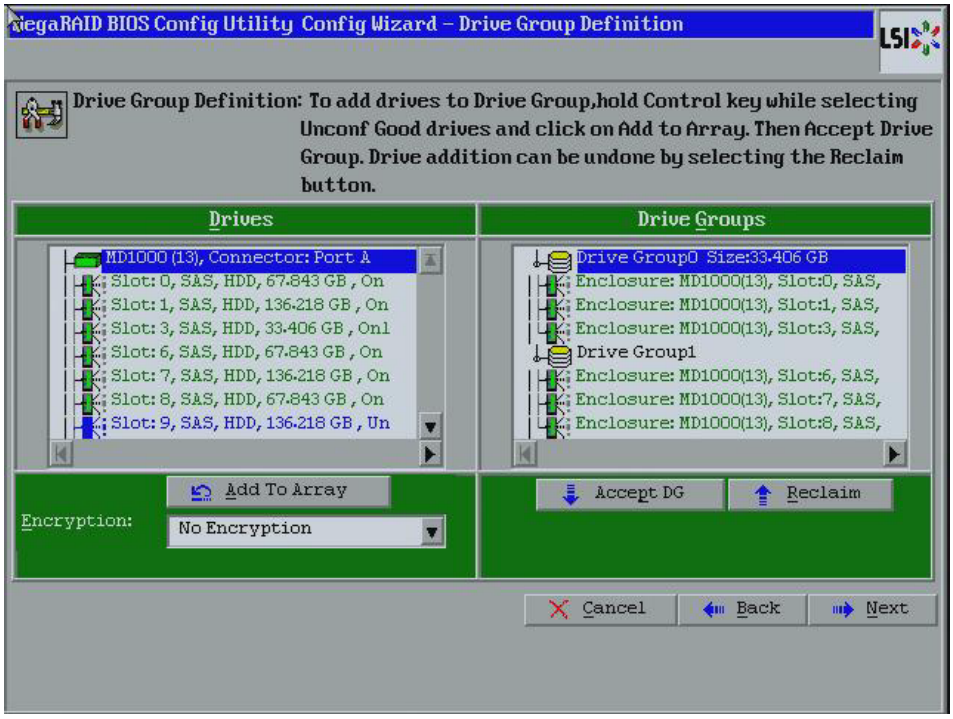


Figure 52: WebBIOS Drive Group Definition Screen

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

The Span Definition screen appears, as shown in Figure 53.
This screen displays the drive group holes you can select to add to a span.

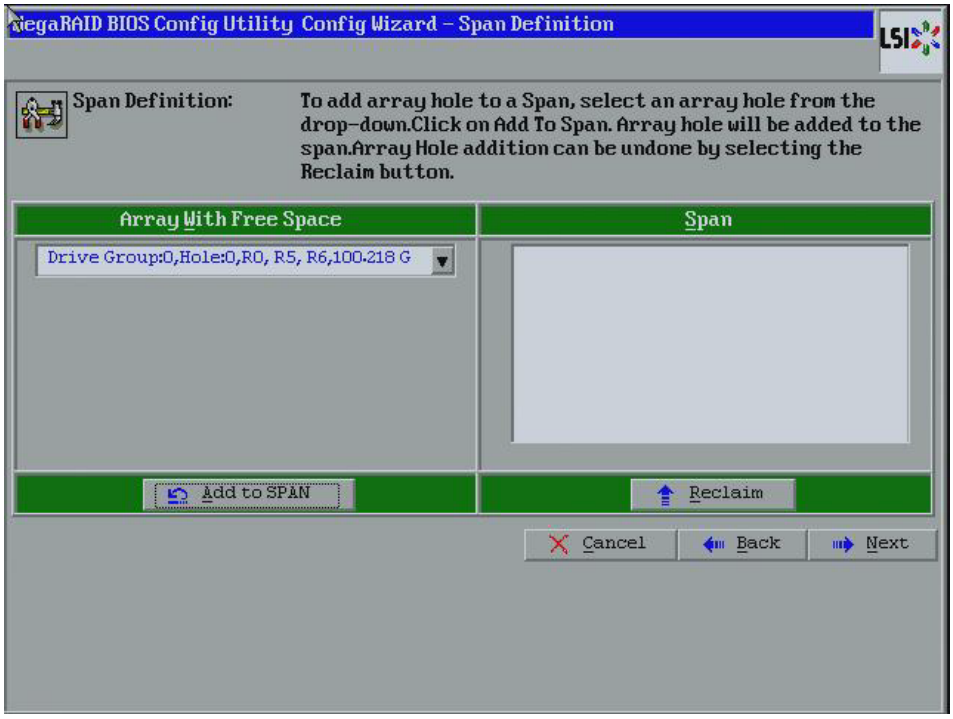


Figure 53: WebBIOS Span Definition Screen

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

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

11. Hold <Ctrl> while you select a second drive group of three or more drives, and click **Add to SPAN**.

    Both drive groups display in the right frame under Span.

12. Click **Next**.

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

13. Hold <Ctrl> while you select two drive groups with three or more drives each in the Configuration panel on the right.

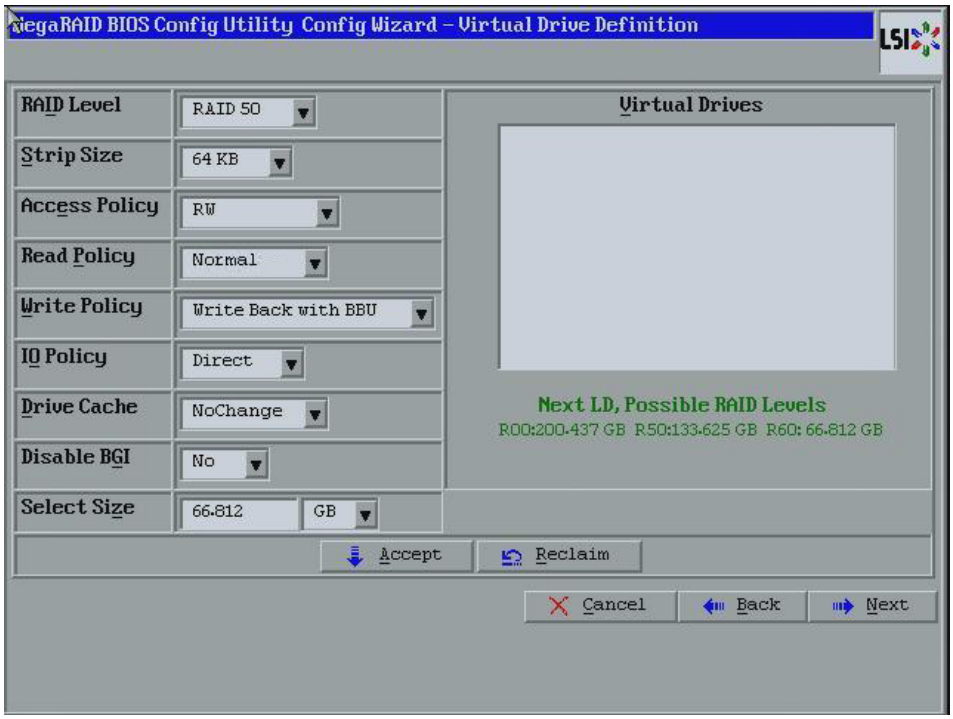


Figure 54: WebBIOS Virtual Drive Definition Screen

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

    Here are brief explanations of the virtual drive options:

    — **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 50.

    — **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, 16, 32, 64, 128, 256, 512, 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 is the default.

      Blocked: Do not allow access.

    — **Read Policy:** Specify the read policy for this virtual drive:

      Normal: This disables the read ahead capability. This is the default.

      Ahead: This enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.

— **Write Policy:** Specify the write policy for this virtual drive:

WBack: In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction.
This setting is recommended in Standard mode.

WThru: In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction.
This is the default.

Bad BBU: Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

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

— **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.

Direct: In Direct I/O mode, the reads are not automatically cached in controller cache and the data is transferred directly to the host except when 1) the data is buffered in controller memory in order to execute parity; or 2) the data is transferred from the controller cache because of read-ahead algorithms. This is the default.

Cached: In Cached I/O mode, all reads are buffered in the controller's cache memory or in the CacheCade virtual drive.

— **Drive Policy**: Specify the drive cache policy:

Enable: Enable the drive cache.

Disable: Disable the drive cache. This drive policy is the default.

NoChange: Leave the current drive cache policy as is.
This is the default.

— **Disable BGI:** Specify the background initialization status:

No: Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.

Yes: Select *Yes* if you do not want to allow background initializations for configurations on this controller.

— **Select Size:** Specify the size of the virtual drive in megabytes. Normally, this would be the full size for RAID 50 shown in the Configuration Panel on the right. You can specify a smaller size if you want to create other virtual drives on the same drive group.

15. Click **Accept** to accept the changes to the virtual drive definition or click **Reclaim** to undo the changess

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

    The Configuration Preview screen appears, as shown in Figure 55.

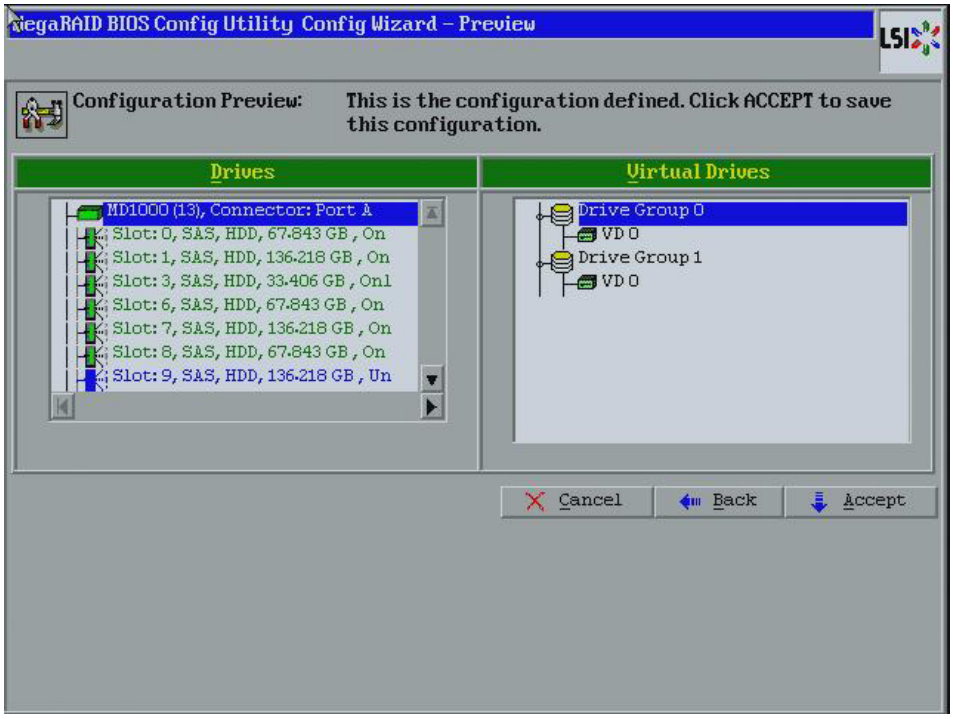


Figure 55: RAID 50 Configuration Preview

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

    The WebBIOS main menu appears.

4.5.4.8 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 screen appears. You use this screen to select drives to create drive groups.

1. Hold <Ctrl> while selecting at least three ready drives in the Drives panel on the left.

2. Click **Add To Array** to move the drives to a proposed drive group configuration in the Drive Groups panel on the right.

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

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

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

4. Click on the icon for the second drive group to select it.
5. Hold <Ctrl> while selecting at least three more ready drives in the Drives panel to create a second drive group.
6. Click **Add To Array** to move the drives to a proposed drive group configuration in the Drive Groups panel on the right, as shown in Figure 56.

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

7. Choose whether to use drive encryption.

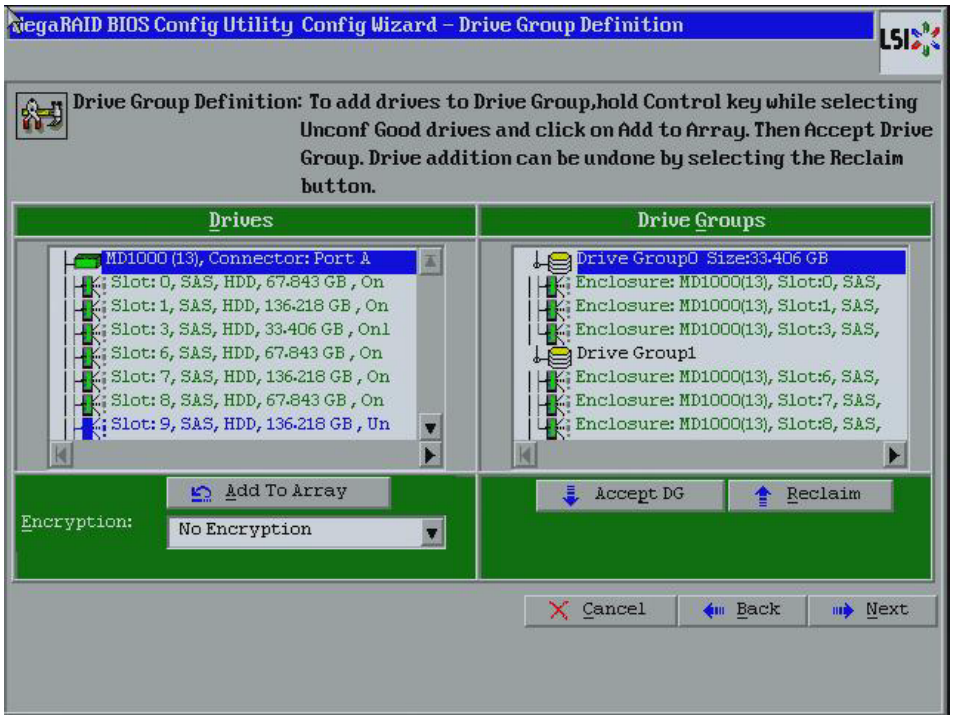


Figure 56: WebBIOS Drive Group Definition Screen

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

   The Span Definition screen appears, as shown in Figure 57.
   This screen displays the drive group holes you can select to add to a span.

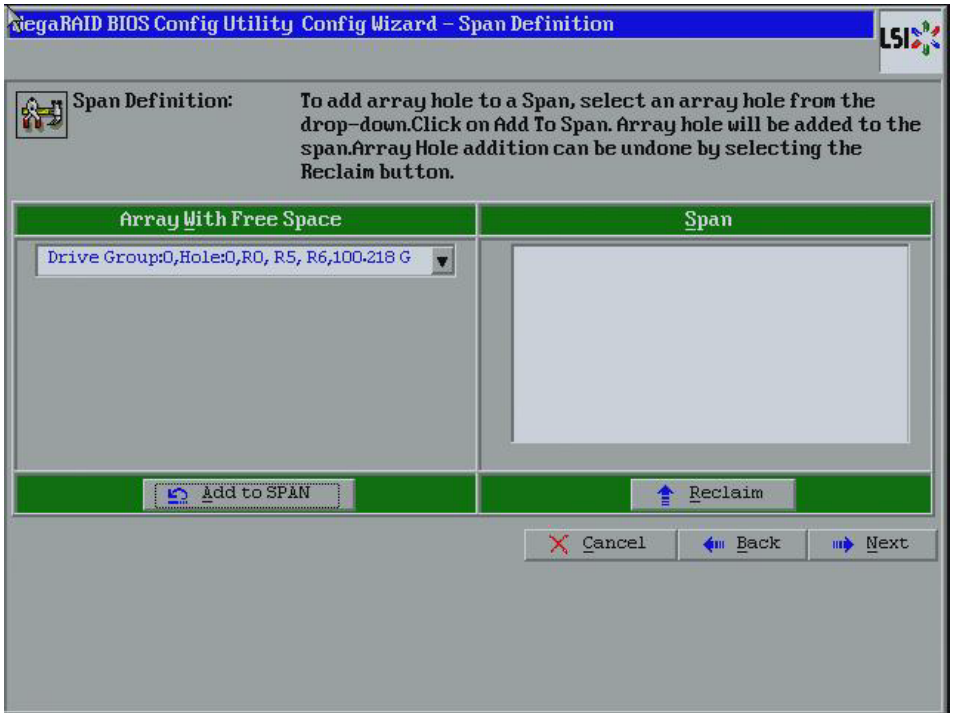


Figure 57: WebBIOS Span Definition Screen

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

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

11. Hold <Ctrl> while you select a second drive group of three or more drives, and click **Add to SPAN**.

    Both drive groups display in the right frame under Span.

12. Click **Next**.

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

13. Hold <Ctrl> while you select two drive groups with at least three drives each in the Configuration window on the right.

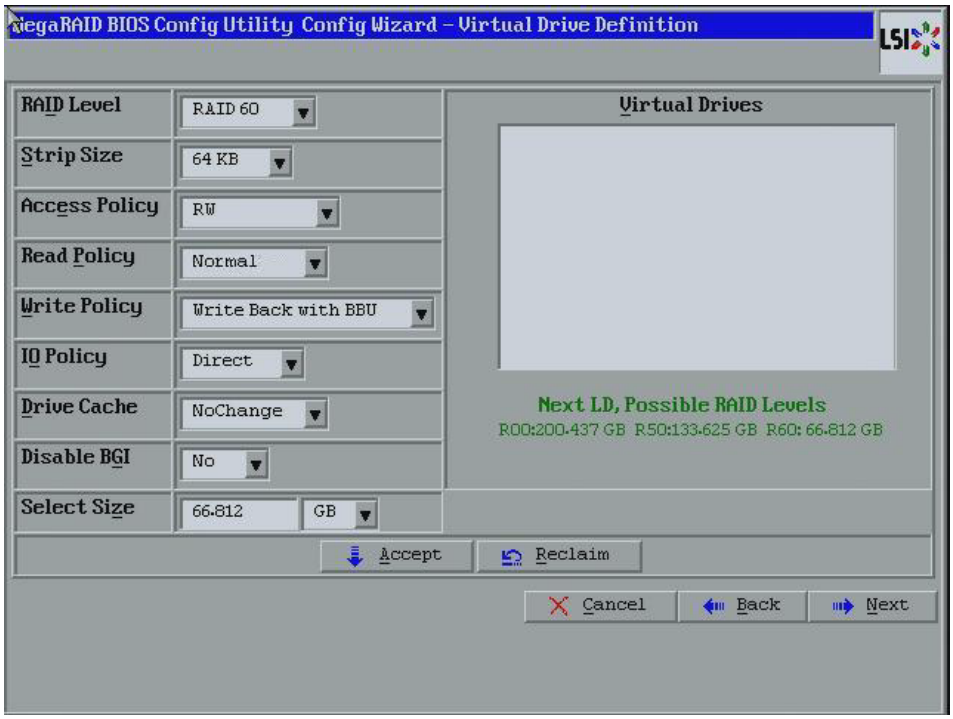


Figure 58: WebBIOS Virtual Drive Definition Screen

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

    Here are brief explanations of the virtual drive options:

    — **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 60.

    — **Stripe Size:** The 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, 16, 32, 64, 128, 256, 512, 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 is the default.

    Blocked: Do not allow access.

    — **Read Policy:** Specify the read policy for this virtual drive:

    Normal: This disables the read ahead capability. This is the default.

Ahead: This enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.

— **Write Policy:** Specify the write policy for this virtual drive:

WBack: In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.

WThru: In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.

Bad BBU: Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

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

— **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.

Direct: In Direct I/O mode, the reads are not automatically cached in controller cache and the data is transferred directly to the host except when 1) the data is buffered in controller memory in order to execute parity; or 2) the data is transferred from the controller cache because of read-ahead algorithms. This is the default.

Cached: In Cached I/O mode, all reads are buffered in the controller's cache memory or in the CacheCade virtual drive.

— **Drive Policy**: Specify the drive cache policy:

Enable: Enable the drive cache.

Disable: Disable the drive cache. This drive policy is the default.

NoChange: Leave the current drive cache policy as is.
This is the default.

— **Disable BGI:** Specify the background initialization status:

No: Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.

Yes: Select *Yes* if you do not want to allow background initializations for configurations on this controller.

— **Select Size:** Specify the size of the virtual drive in megabytes. Normally, this would be the full size for RAID 60 shown in the Configuration panel on the right. You can specify a smaller size if you want to create other virtual drives on the same drive group.

NOTE: WebBIOS does not 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 **Reclaim** button.

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

    The Configuration Preview screen appears, as shown in Figure 59.

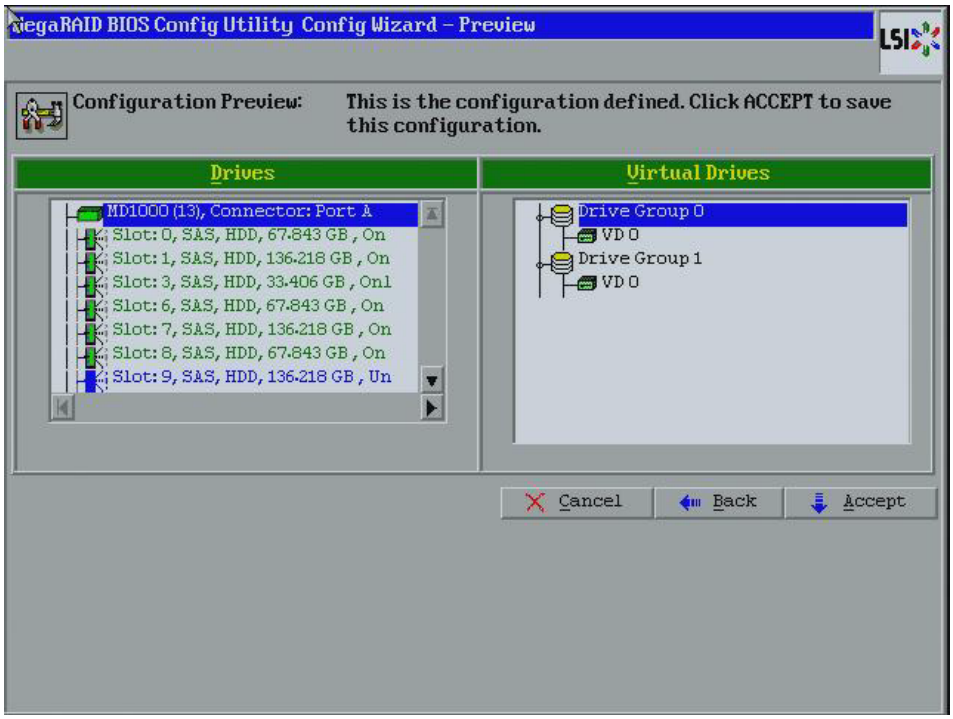


Figure 59: RAID 60 Configuration Preview

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

    The WebBIOS main menu appears.

4.6 Creating a 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. 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 an CacheCade software virtual drive that functions as a secondary tier of cache.

Using CacheCade 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 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.

Follow these steps to create a CacheCade drive group.

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

   The first Configuration Wizard screen appears, as shown in Figure 60.

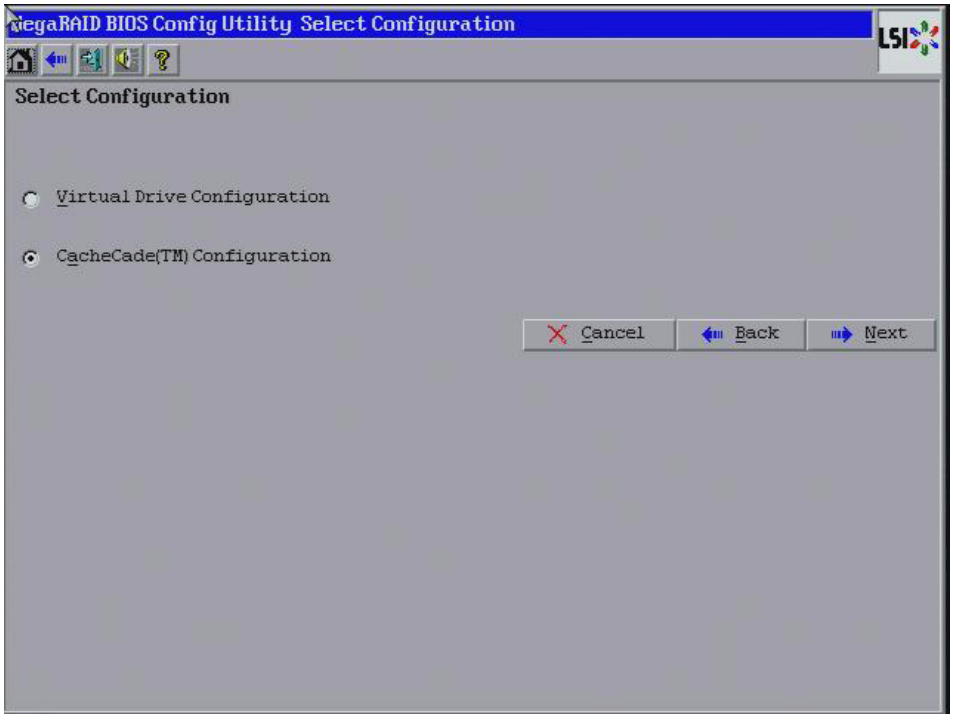


Figure 60: WebBIOS Configuration Wizard Screen

2. Select **CacheCade(TM) Configuration** and click **Next**.

The Span Definition screen appears, as shown in Figure 60.

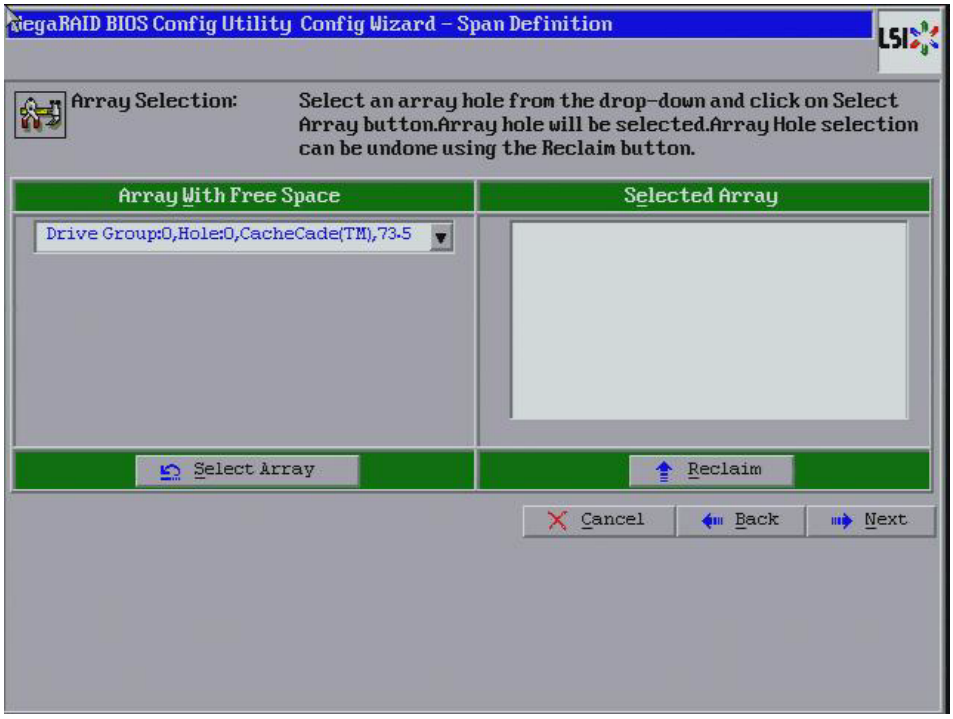


Figure 61: CacheCade Array Selection Screen

3. 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**.

4. Click **Next**.

   The Create CachCade Disk screen appears, as shown in Figure 62.

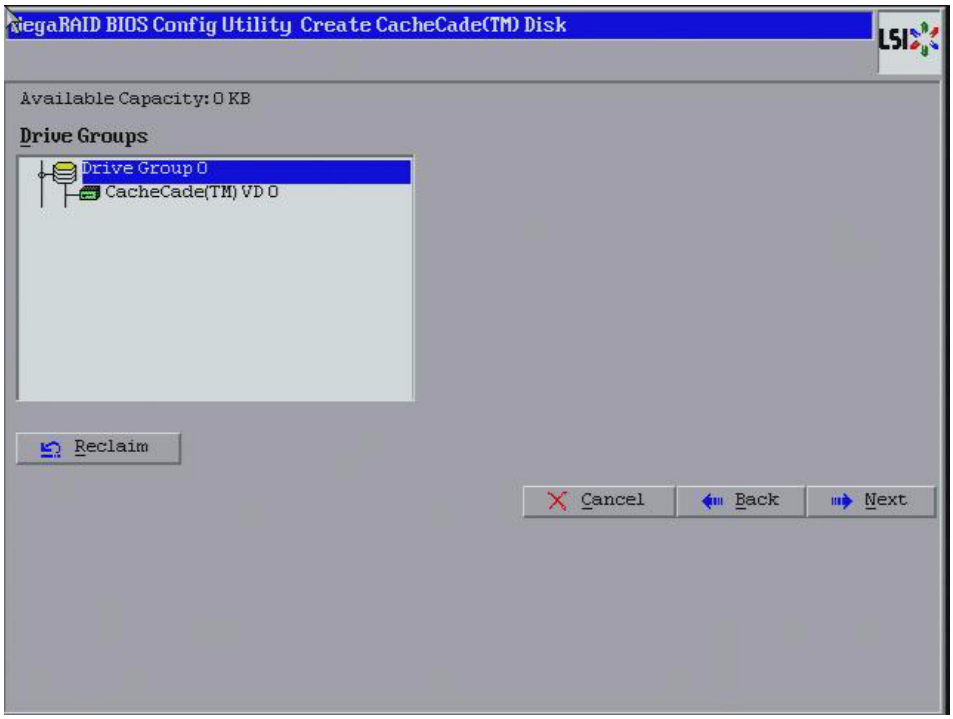


Figure 62: CacheCade Disk Screen

5. Click **Next** to accept the drive group.

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

   The Config Wizard Preview screen appears, as shown in Figure 63.

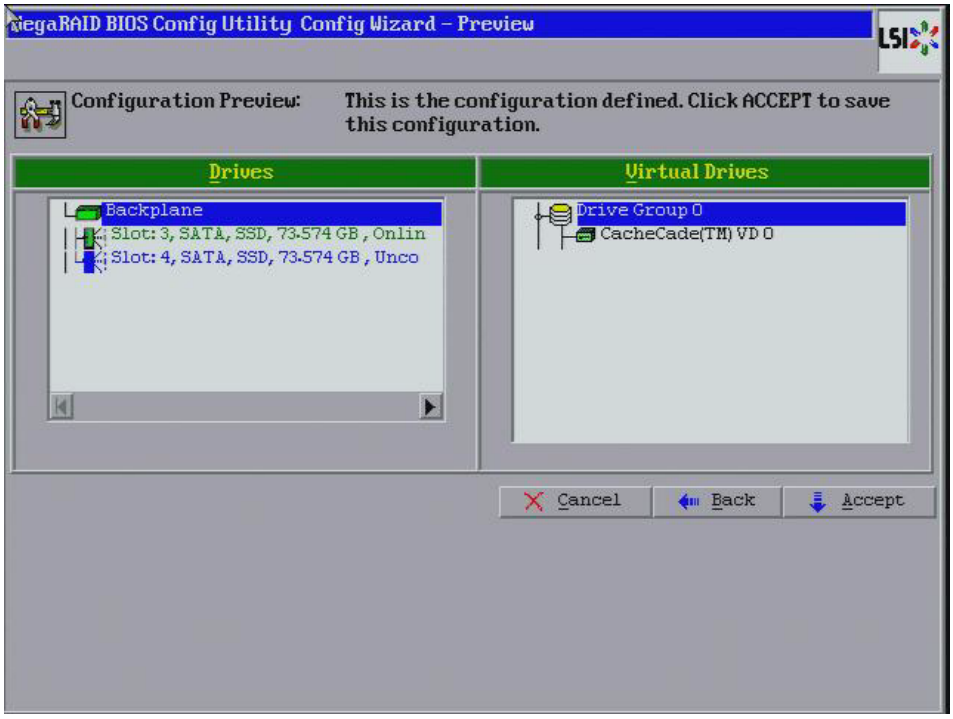


Figure 63: CacheCade Configuration Preview

6. Click **Accept** if the configuration is OK. Otherwise, or click **Back** to return to the previous screens and change the configuration.
7. If you accept the configuration, click **Yes** at the prompt to save the configuration.

   The WebBIOS main menu screen appears, as shown in Figure 64. It shows the CacheCade virtual drive.

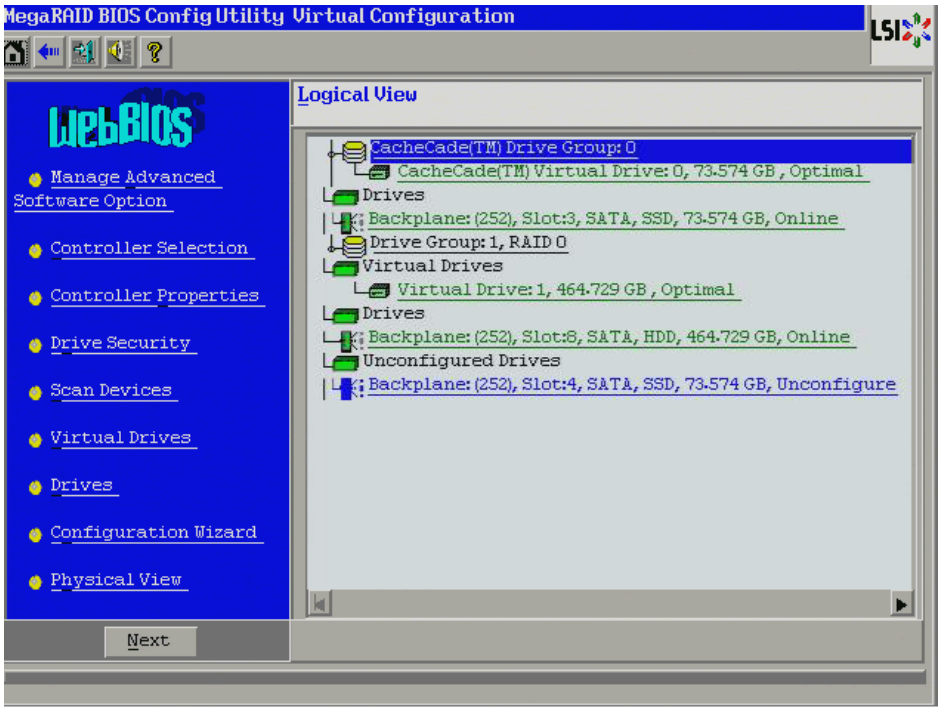


Figure 64: WebBIOS Main Menu with a CacheCade Virtual Drive

4.7 Selecting SafeStore Encryption Services Security Options

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.

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

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 screen.

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

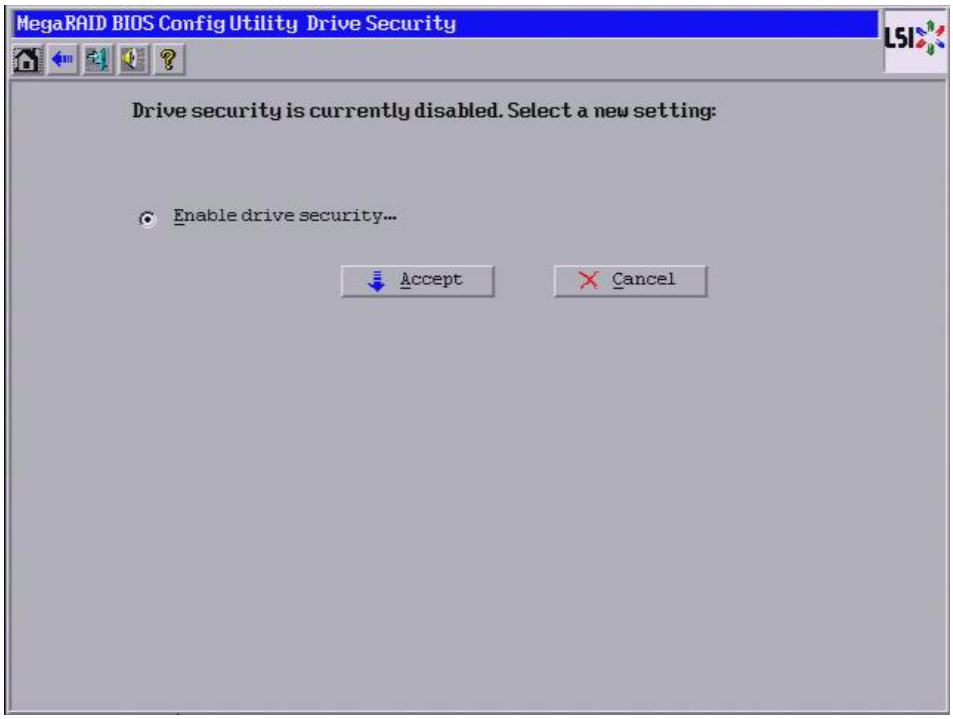


Figure 65: Driver Security Settings Screen

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

   The **Drive Security Choose mode** wizard appears as shown in Figure 66. This wizard allows you to choose the security key management mode.

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

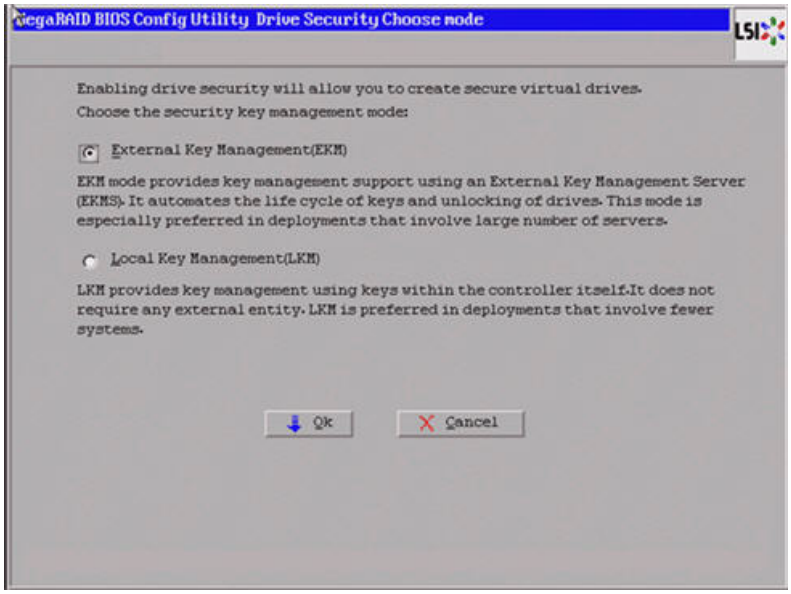


Figure 66: Drive Security Choose mode

3. Select **EKM** in the **Drive Security Choose mode** wizard.
4. Click **OK.**
5. Click **OK**, the system asks you for a restart. The restart message appears as shown in Figure 71.
6. If you select LKM in **Drive Security Choose mode** wizard (Figure 66), the screen used to create a security key identifier appears, as shown in Figure 67.

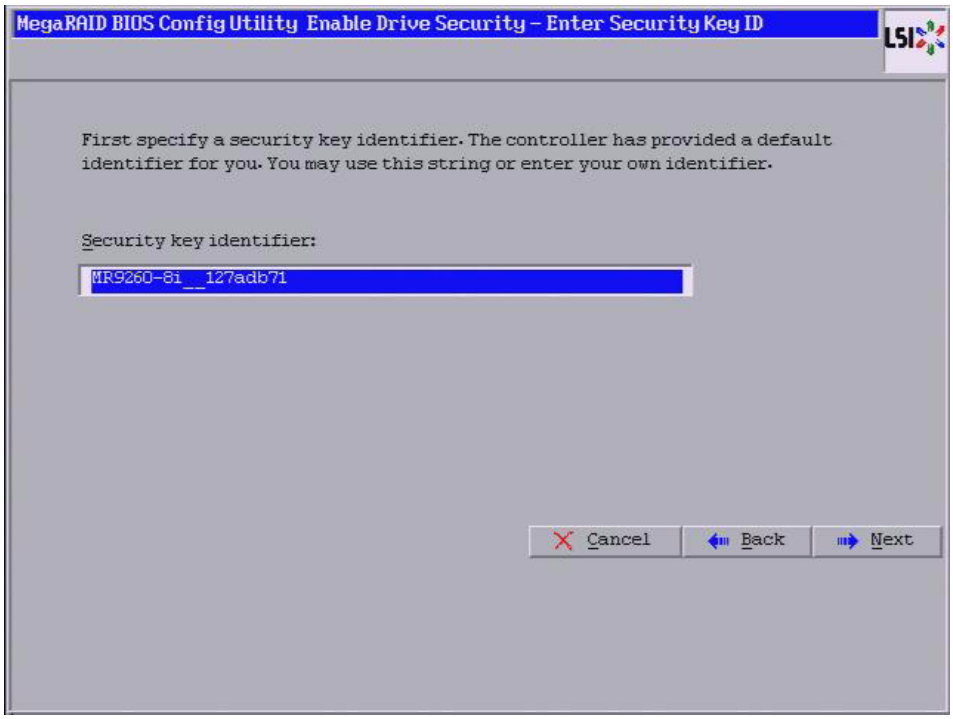


Figure 67: Enable Drive Security – Enter Security Key ID Screen

7. Accept the default security key ID or enter a new security key ID.
8. Click **Next**.

   The Enable Drive Security – Enter Security Key screen appears as shown in Figure 68.

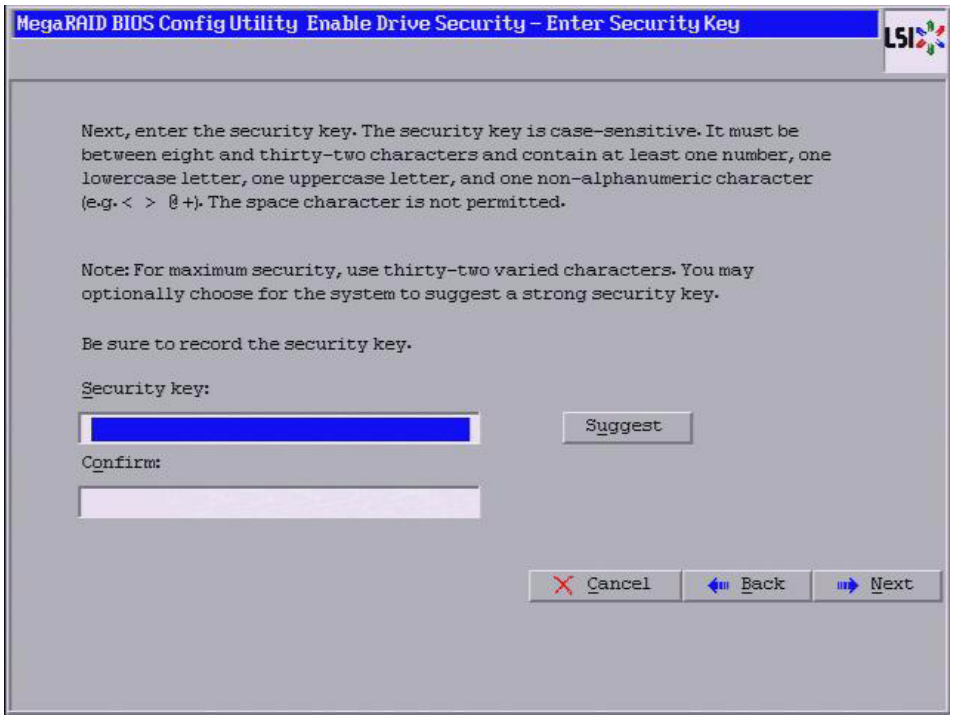


Figure 68: Enable Drive Security – Enter Security Key

9. 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 eight and thirty-two 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.

10. Click **Next**.

    The Enable Drive Security – Enter Pass Phrase screen appears as shown in Figure 69. You have the option to provide a pass phrase for additional security.

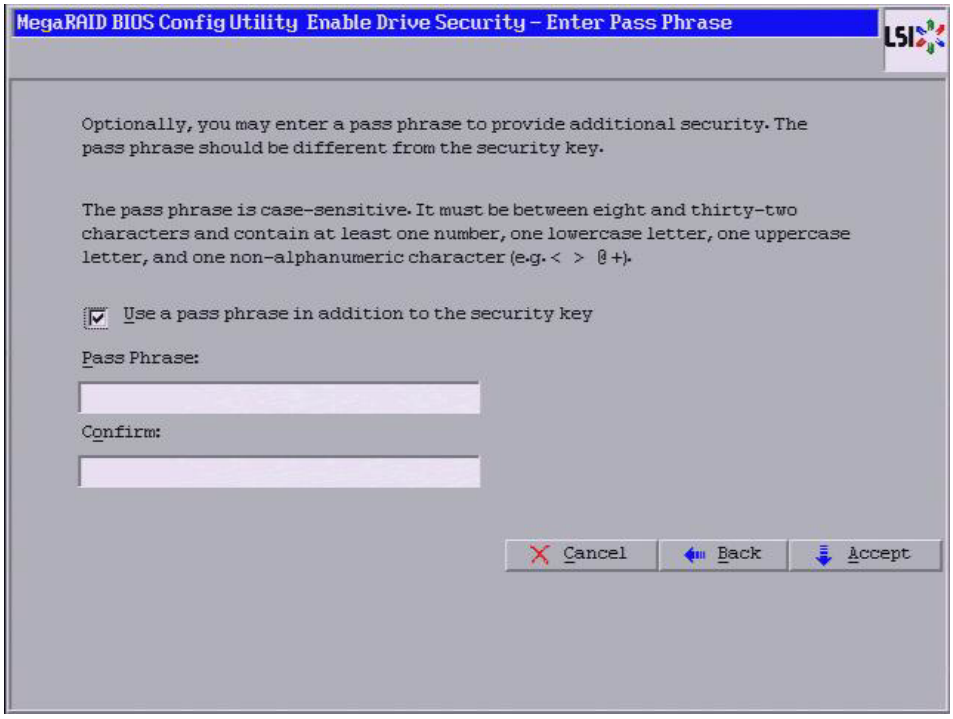


Figure 69: Enable Drive Security – Enter Pass Phrase

11. If you want to use a pass phrase, click the checkbox **Use a pass phrase in addition to the security key**.
12. Enter a new pass phrase and then enter the new pass phrase again to confirm.

    The pass phrase is case-sensitive. It must be between eight and thirty-two 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-US keyboard users must be careful not to enter DBCS characters in the pass phrase field or security key field. Firmware works only with the ASCII character set.
13. Click **Accept**.

    The Confirm Enable Drive Security screen appears, as shown in Figure 70



Figure 70: Confirm Enable Drive Security Screen

14. Click **Yes** on the Confirm Enable Drive Security screen to confirm that you want to enable the drive security settings.

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

When you select EKM in Figure 66, and click **OK**, the following dialog box appears.

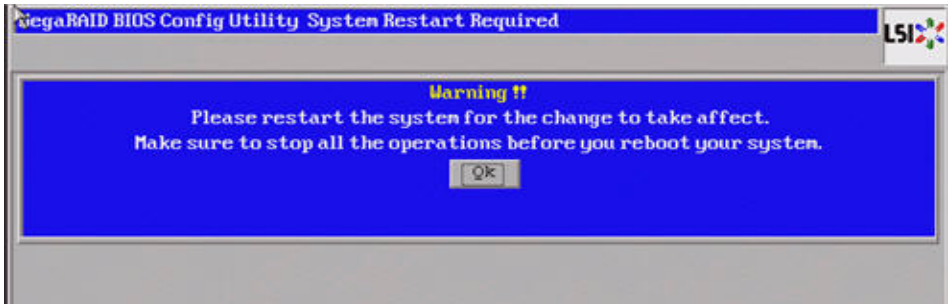


Figure 71: System Restart

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 will be enabled. Perform the following steps to change the encryption settings for the security key identifier, security key, and pass phrase.

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

   The **Change Drive Security** screen appears as shown in Figure 72.

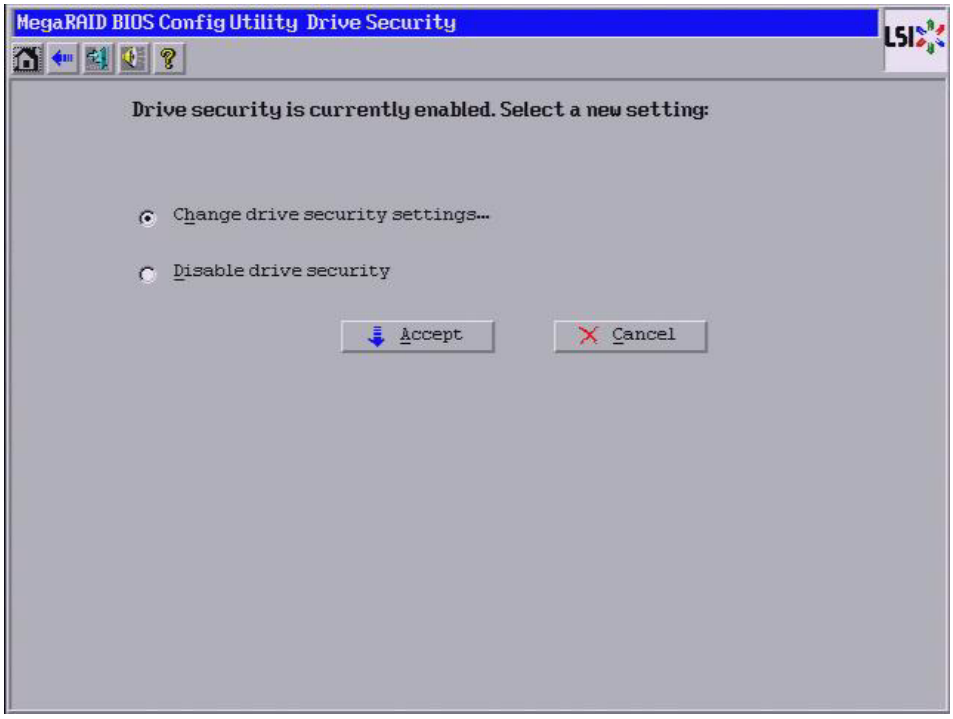


Figure 72: Change Drive Security Settings Screen7

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

4.7.4 Change Security from EKM to LKM

The Change Security Settings – Introduction screen appears as shown in Figure 73.

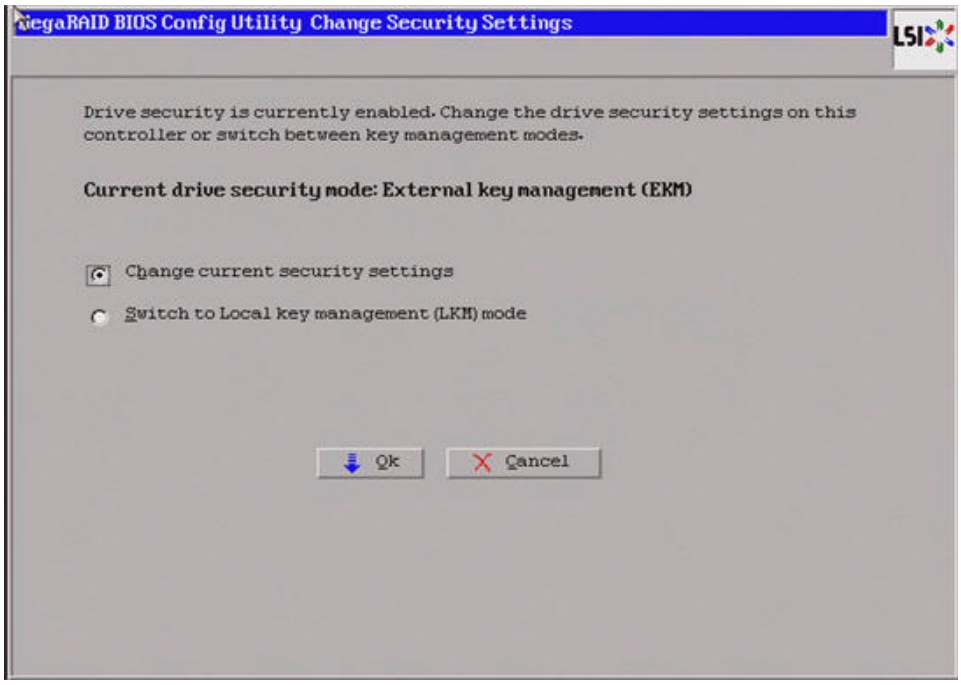


Figure 73: Change Security Settings – Introduction

3. Select **Change current security s ettings**, and click **OK**, the system restart message appears as shown in Figure 71.
4. Select **Switch to Local Key management (LKM) mode**, the Change Security Settings – Security Key ID screen appears as shown in Figure 74.

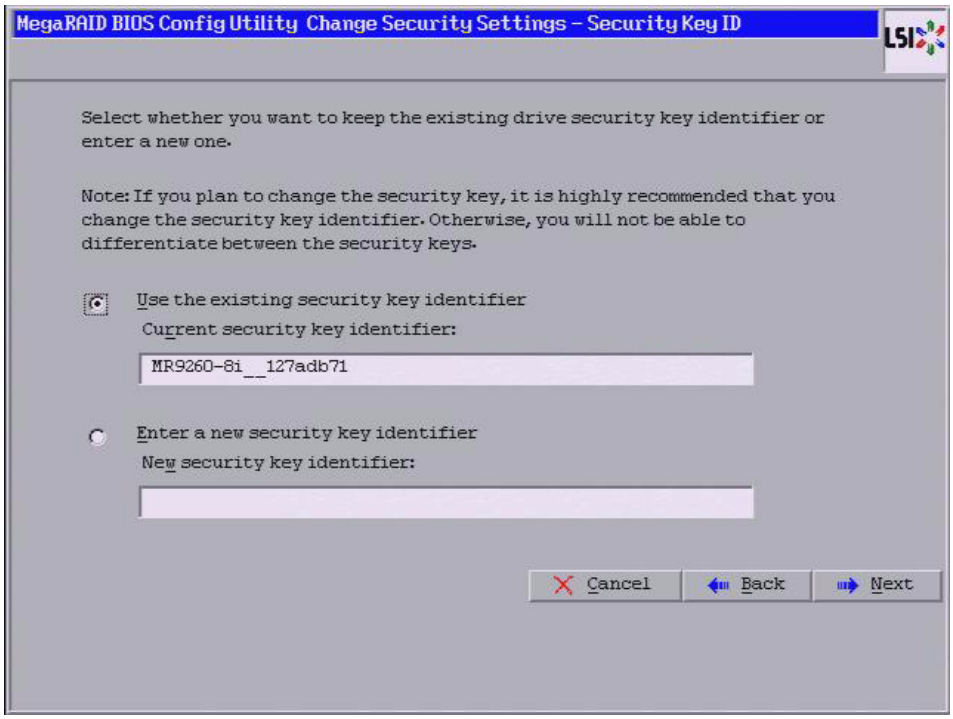


Figure 74: Change Security Settings – Security Key ID

5. Choose whether you want 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).
6. Click **Next**.

   The **Change Security Settings – Security Key** screen appears as shown in Figure 75. You have the option to use the existing security key or enter a new one.

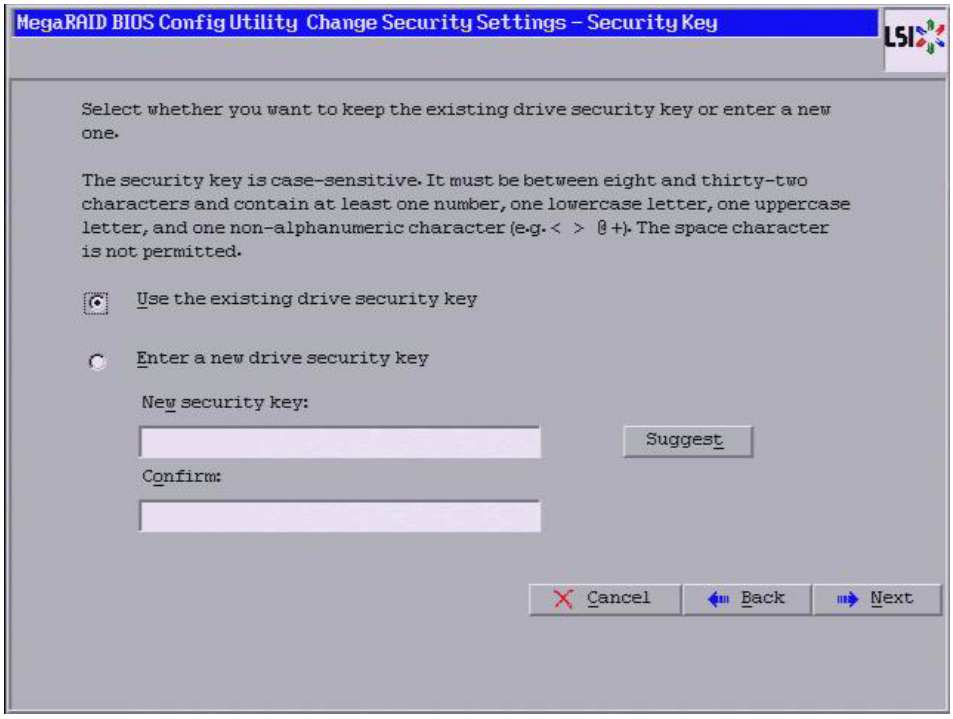


Figure 75: Change Security Settings – Security Key

7. Choose whether you want to use the existing security key or enter a new security key. You have the following options:
   - Use the existing drive security key.
   - Enter a new drive security key.
8. If you want 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.
9. Enter the new drive security key again in the **Confirm** field.

   The security key is case-sensitive. It must be between eight and thirty-two 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.
10. Click **Next**.

    The Change Security Settings – Pass Phrase screen appears as shown in Figure 76.

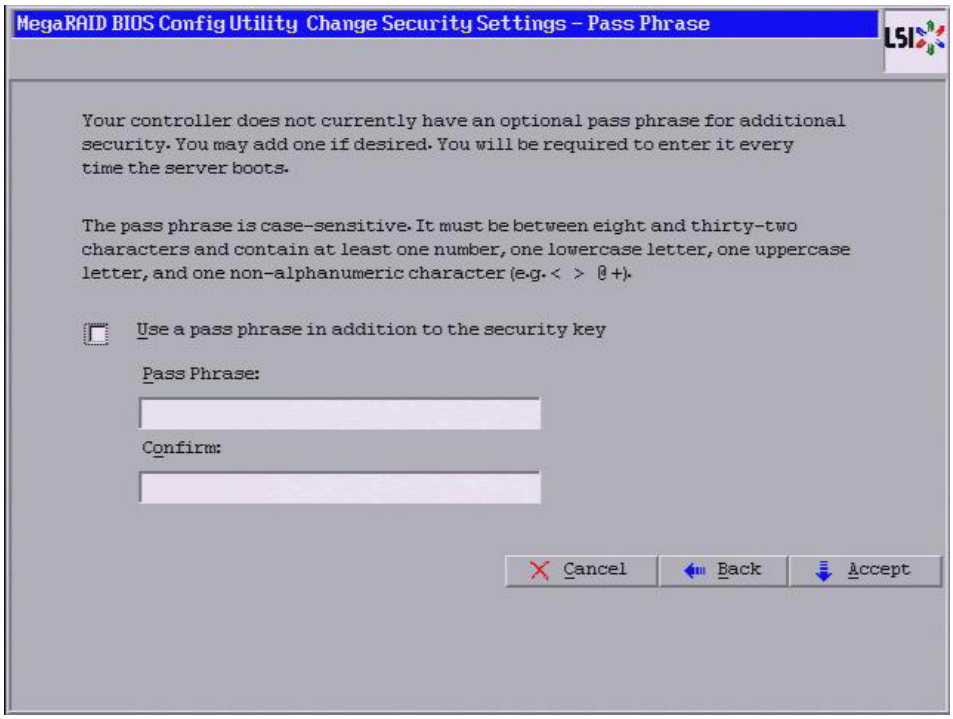


Figure 76: Change Security Settings – Pass Phrase Screen

11. If you want to use a pass phrase, click the checkbox **Use a pass phrase in addition to the security key**.
12. Enter a new pass phrase and then enter the new pass phrase again to confirm.

    The pass phrase is case-sensitive. It must be between 8and 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-US keyboard users must be careful not to enter DBCS characters in the pass phrase field or security key field. Firmware works only with the ASCII character set.
13. Click **Accept**.

    If you entered a new pass phrase, the Authenticate Drive Security Settings screen appears.
14. On the Authenticate Drive Security Settings screen, enter the security key and click **Ok**.

    If you entered a new drive security key, the Authenticate Drive Security Settings screen appears, as shown in Figure 77.

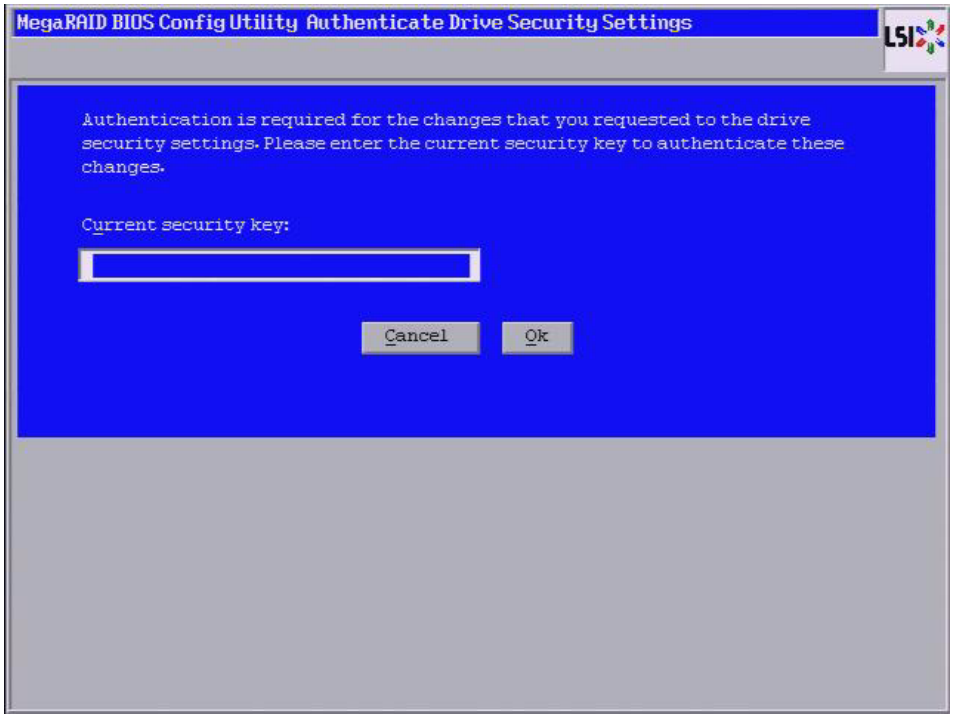


Figure 77: Authenticate Drive Security Settings Screen

15. 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 screen appears, as shown in Figure 78.

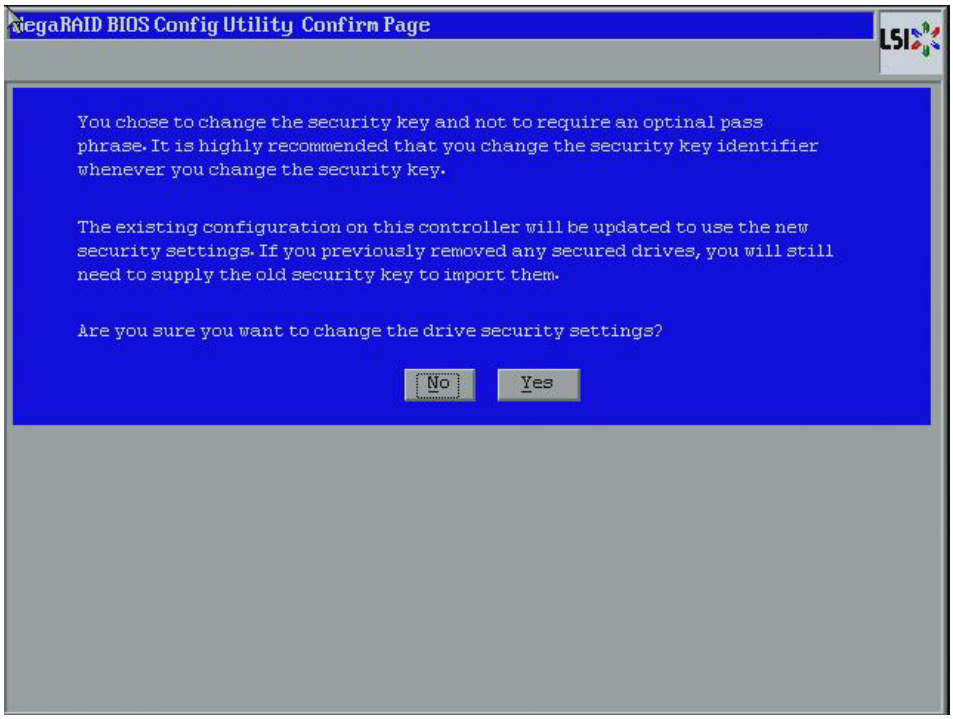


Figure 78: Confirm Screen

16. 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 Authenicate Drive Security Settings screen 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 page.
2. Select **Change Drive security.** The **Change Security Settings** wizard appears as shown in the following figure.

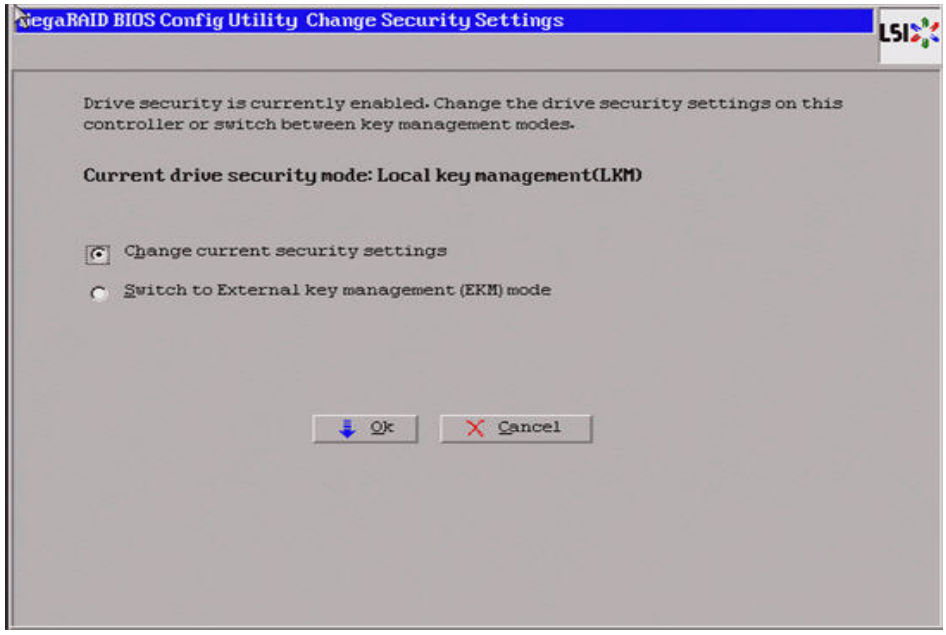


Figure 79: Change Security Settings

3. Select the **Change current security settings** option**, the Change Security Settings – Security Key ID** wizard is displayed **(**Figure 74**).** This procedure continues till Figure 78.
4. Select the **Switch to External Key Management (EKM) mode**, and click **OK,** the **Authenticate Drive Security Settings** screen is displayed(Figure 77).

4.7.6 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.

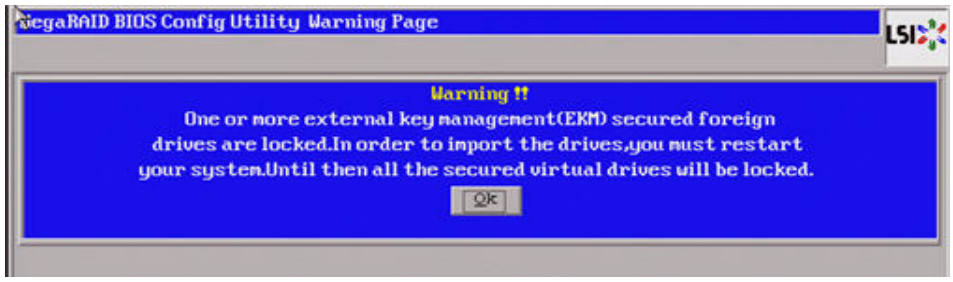


Figure 80: Import Foreign Drives in EKM mode

4.7.7 Import 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 is displayed

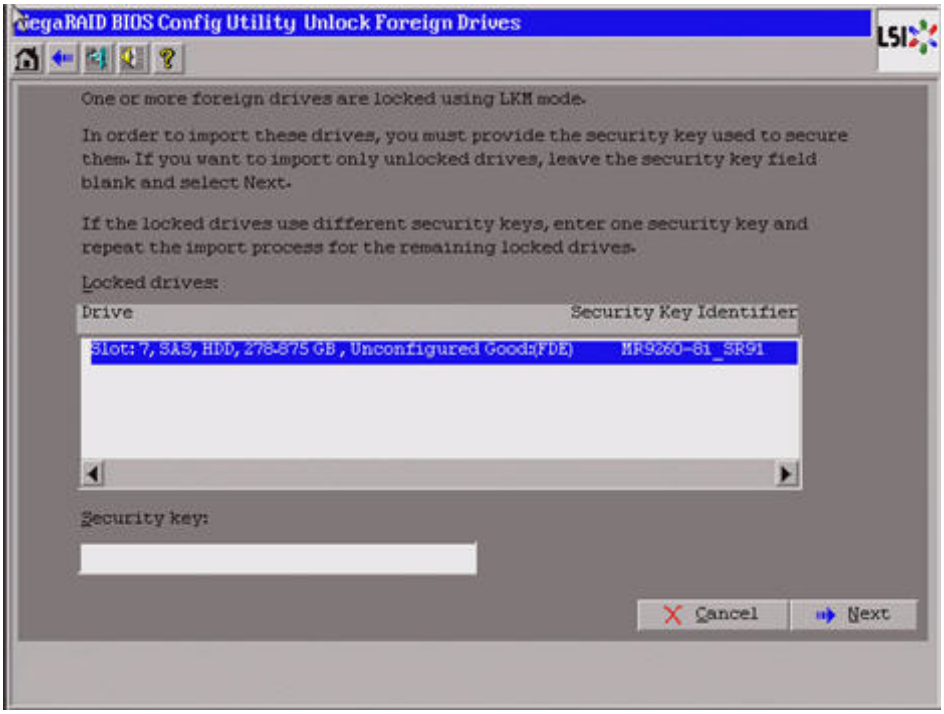


Figure 81: Import Foreign Drives in EKM or LKM Mode

- 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.

4.7.8 Import Foreign Drives in LKM /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 wizard is displayed.

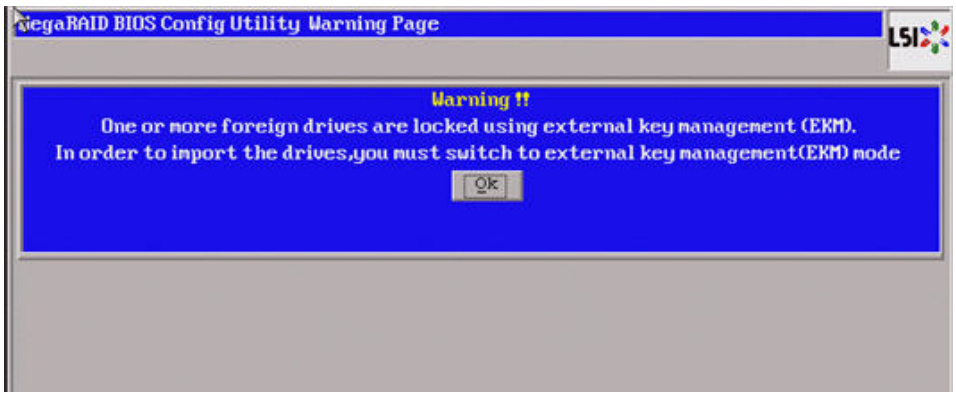


Figure 82: Import Foreign Drives in LKM Mode

4.7.9 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 will still need to enter the security key when you import them.

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

   The Drive Security screen appears, as shown in Figure 83.

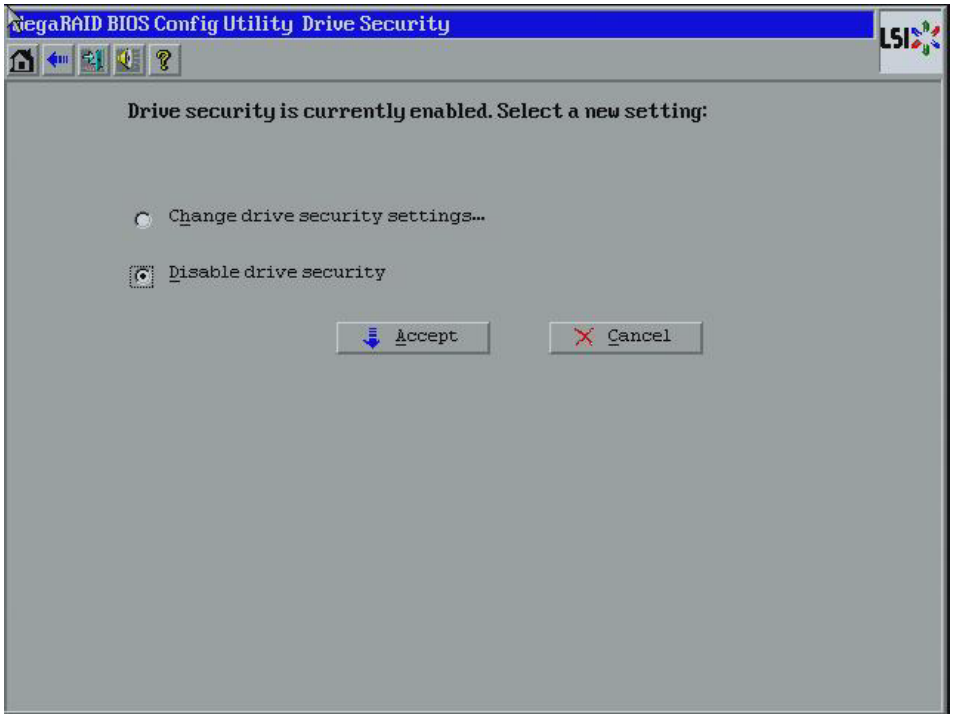


Figure 83: Drive Security

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

   The Confirm Disable Drive Security screen appears as shown in Figure 84.

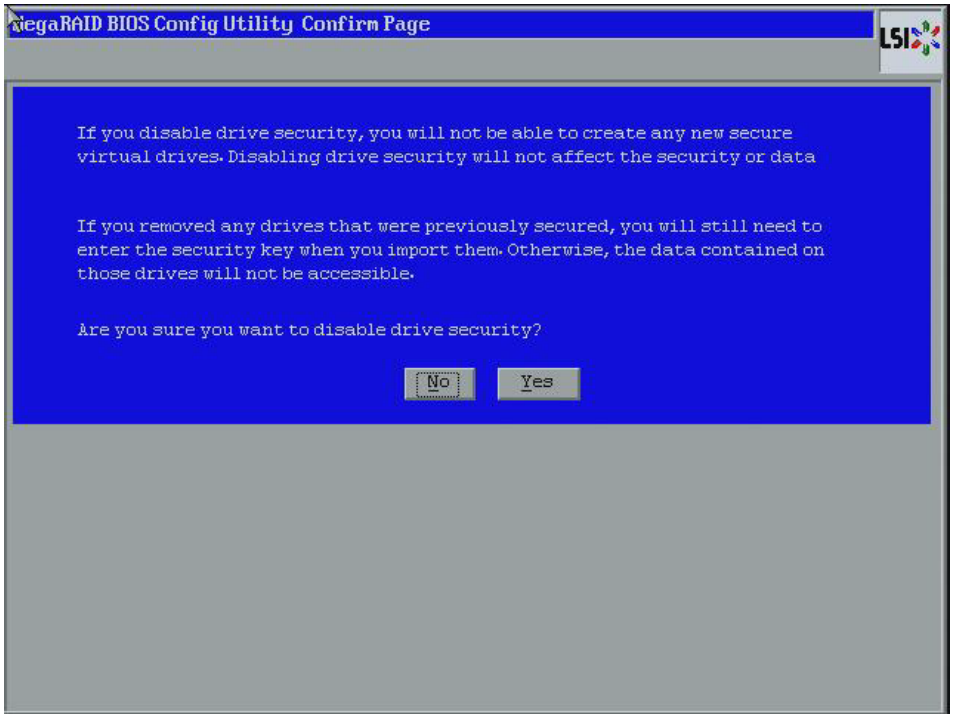


Figure 84: Confirm Disable Drive Security Settings

3. On the Confirm Disable Security Settings screen, click **Yes** to confirm that you want to disable the drive security settings.

   WebBIOS returns you to the MSM main menu.

4.8 Viewing and Changing Device Properties

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

4.8.1 Viewing Controller Properties

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

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

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

   There are three Controller Properties screens. Figure 85 shows the first screen.



Figure 85: First Controller Properties Screen

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

2. Click **Next** to view the second Controller Properties screen, as shown in Figure 86.

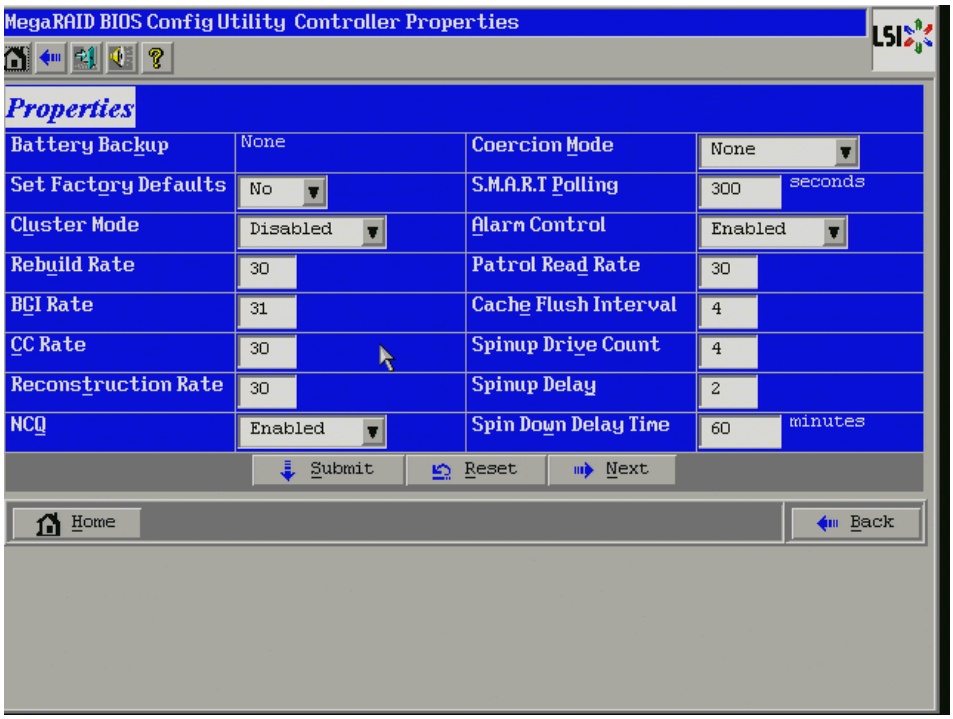


Figure 86: Second Controller Properties Screen

3. Click **Next** to view the third Controller Properties screen, as shown in Figure 87.

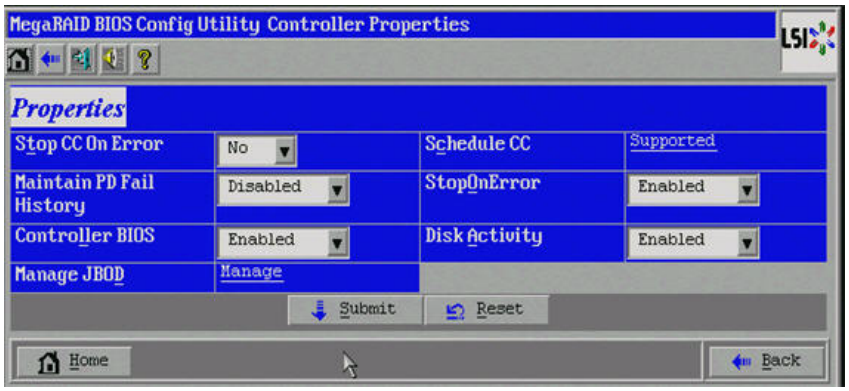


Figure 87: Third Controller Properties Screen

Table 21 describes the entries/options listed on the second and third Controller Properties screens. 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 Properties 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 4.8.4, *Viewing and Changing Battery Backup Unit Information*
Set Factory Defaults	Use this option to load the default MegaRAID® WebBIOS CU 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 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.
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 *None*. The number you choose depends on how much the drives from various vendors vary in their actual size. LSI recommends that you use the 1GB coercion mode option.
S.M.A.R.T. Polling	Use this option to determine how frequently the controller polls for drives reporting a Predictive Drive Failure (S.M.A.R.T.: Self-Monitoring Analysis and Reporting Technology error). The default is 300 seconds (5 minutes).
Alarm Control	Select this option to enable, disable, or silence the onboard alarm tone generator on the controller. The default is *Disabled*.

Table 21: Controller Properties Menu Options (Continued)

Option	Description
Patrol Read Rate	Use this option to select the rate for patrol reads for drives connected to the selected controller. The default is 30 percent. The patrol read rate is the percentage of system resources dedicated to running a patrol read. See Section 5.9, *Patrol Read-Related Controller Properties* for additional information about patrol read.
Cache Flush Interval	Use this option to control the interval (in seconds) at which the contents of the onboard data cache are flushed. The default is 4 seconds.
Spinup Drive Count	Use this option to control the number of drives that spin up simultaneously. The default is 2 drives.
Spinup Delay	Use this option to control the interval (in seconds) between spinup of drives connected to this controller. The delay prevents a drain on the system's power supply that would occur if all drives spun up at the same time. The default is 12 seconds.
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 *Disabled*.
Spin Down Delay Time	Use this option to control the interval (in seconds) between spindown of drives connected to this controller. The delay prevents a drain on the system's power supply that would occur if all drives spun down at the same time. The default is 30 minutes.
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.
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 a enclosure.
Manage JBOD	Converting the multiple JBOD drives to unconfigured drive at single selection.

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

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 CU main screen.

The Virtual Drive screen appears, as shown in Figure 88.

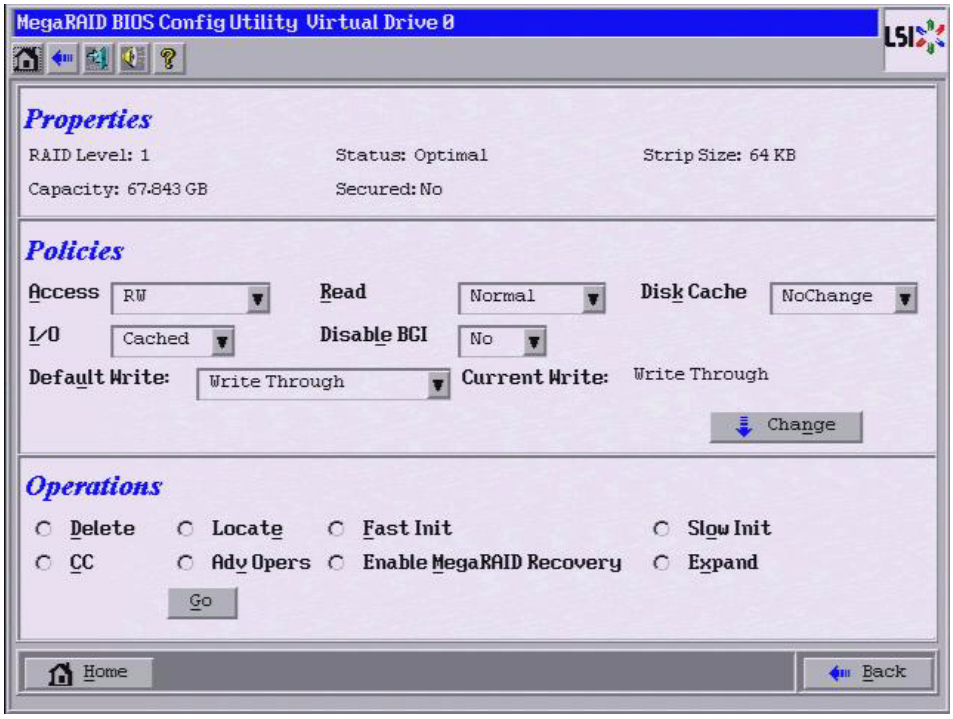


Figure 88: Virtual Drive Screen

The Properties panel of this screen displays the virtual drive's RAID level, state, capacity, strip 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 4.5.4, *Using Manual Configuration* To change any of these policies, make a selection from the drop-down menu and click **Change**.

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

- Select **Delete** to delete this virtual drive. For more information, see Section 4.12.2, *Deleting a Virtual Drive*.
- Select **Locate** to make the LEDs flash on the drives used by this virtual drive. This works only if the drives are installed in a drive enclosure that supports SAFTE (SCSI-Accessed-Fault-Tolerant-Enclosure).
- 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 4.12.1, *Running a Consistency Check*. (This option is not available for RAID 0 virtual drives.)
- Select **AdvOpers** to access screens 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 4.12.5, *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 4.10, *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 4.10, *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. In addition, you can add drives to the virtual drive in order to increase capacity.

  See Section 4.9, *Expanding a Virtual Drive* for the procedure you can use to expand 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 screen displays the properties of a selected drive and enables you to perform operations on the drive. There are two ways to access the Physical Drive screen:

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

Figure 89 shows the Physical Drive screen.

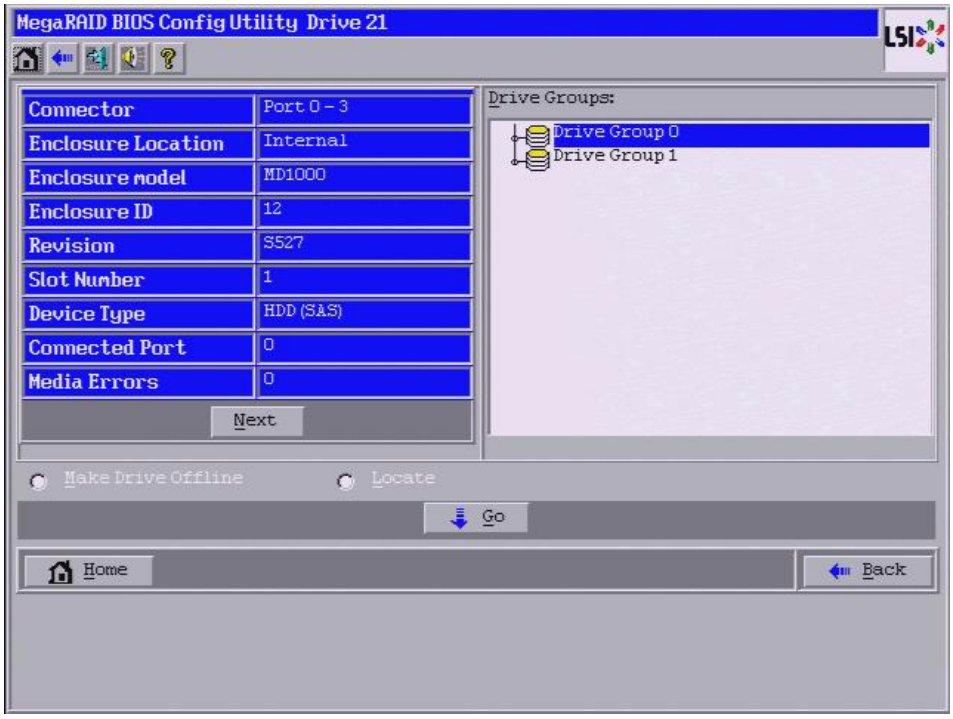


Figure 89: Physical Drive Screen

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

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

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

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

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

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

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

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

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

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

4.8.4 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 do this, follow these steps:

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

   The first Controller Information screen appears, as shown in Figure 90.



Figure 90: First Controller Properties Screen

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

   The second Controller Properties screen appears, as shown in Figure 91. The **Battery Backup** field at the top left of the screen indicates whether the iBBU is present.

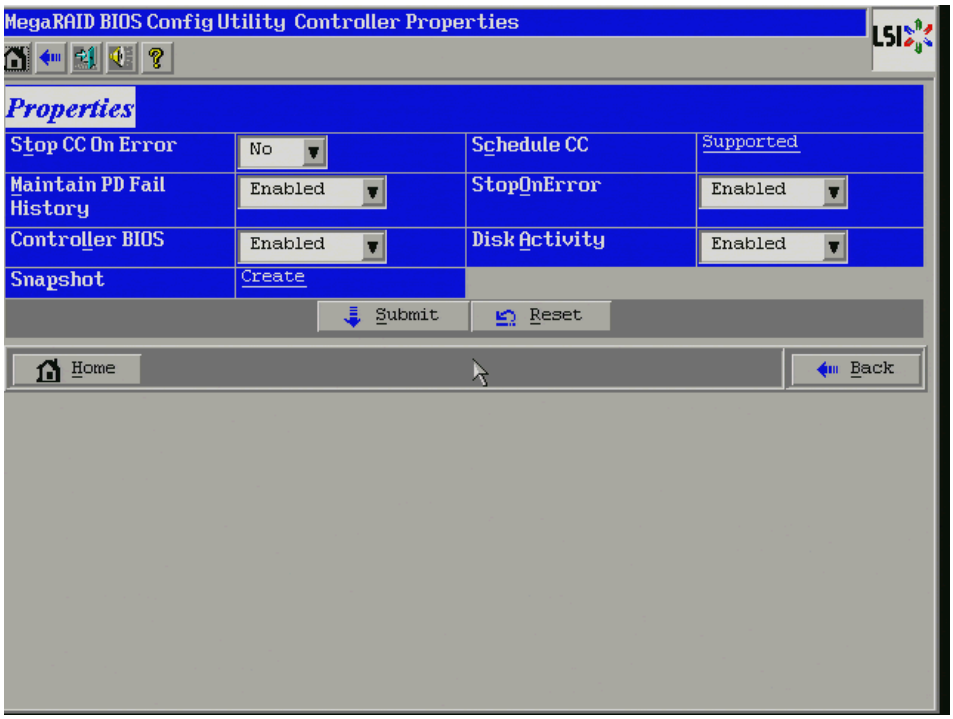


Figure 91: Second Controller Properties Screen

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

   The Battery Module screen appears, as shown in Figure 92. This screen contains the following information:

   — Battery information
   — Design information
   — Capacity information
   — Auto Learn properties and settings

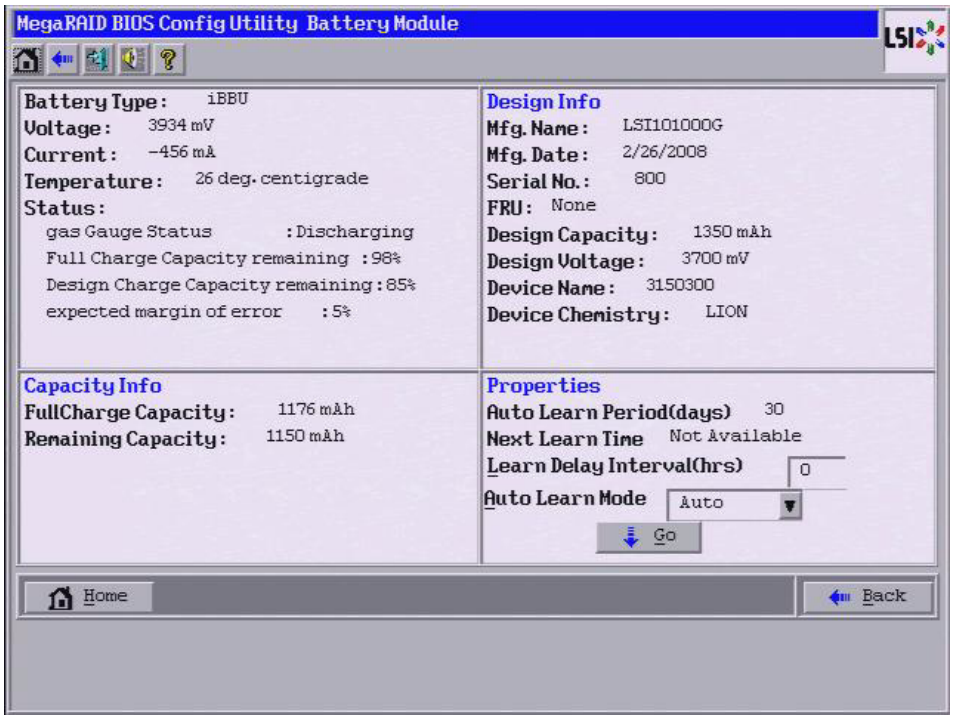


Figure 92: Battery Module Screen

Most of the Battery Module properties are view-only and are self-explanatory.

In the lower right corner of the screen are 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 the learn delay interval and the auto learn mode at their default settings.

4.8.4.1 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 menu in the **Auto Learn Mode** field.
b. Select the learn mode as `Auto` (the default).
   This is so the controller performs the learning cycle automatically.
c. Change the number of hours in the **Learn Delay Interval** field.
   You can delay the start of the learn cycles for up to 168 hours (7 days).
d. Click **Go** to set the interval.

4.8.4.2 Setting the Auto Learn Mode

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

— BBU Auto Learn: Firmware tracks the time since the last learning cycle and performs a learn cycle when due.

— BBU Auto Learn Disabled: Firmware does not monitor or initiate a learning cycle. You can schedule learning cycles manually.

— BBU Auto Learn Warn: Firmware warns about a pending learning cycle. You can initiate a learning cycle manually. After the learning cycle is complete, 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:

a. Open the drop-down menu in the **Auto Learn Mode** field.
b. Select an auto learn mode.
c. Click **Go** to set the auto learn mode.

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

4.8.5 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 must have the same maximum link speed setting. If the phys in a SAS port have different link speeds, the firmware uses the lowest link speed for that SAS port.

Follow these steps to change the link speed.

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

   There are three Controller Properties screens. The first Controller Properties screen appears. See Figure 85 to view this screen.

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

   The second Controller Properties screen appears. See Figure 86 to view this screen.

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

   The third Controller Properties screen appears, as shown in Figure 93.



Figure 93: Third Controller Properties Screen

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

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

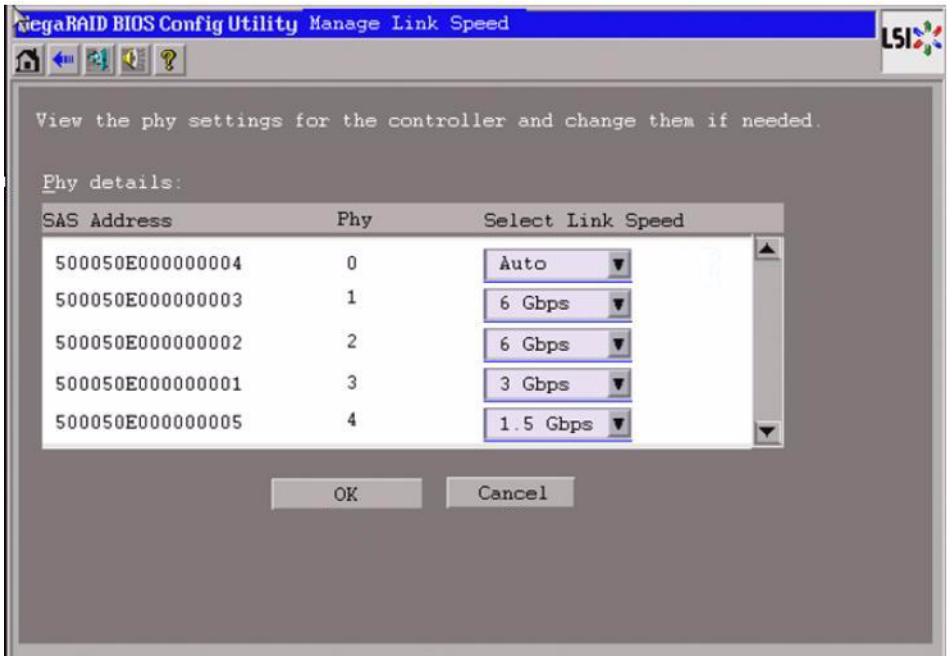


Figure 94: Manage Link Speed Screen

- 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.

5. 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.

6. 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 95: System Restart Required Message Box

4.9 Expanding a Virtual Drive

You can increase the size of a virtual drive to occupy the remaining capacity in a drive group. In addition, you can add drives to the virtual drive in order to increase capacity.

Follow these steps to expand a virtual drive.

1. Access the Virtual Drive screen by clicking a virtual drive icon in the right panel on the WebBIOS CU main screen.

   The Virtual Drive screen appears, as shown in Figure 96.

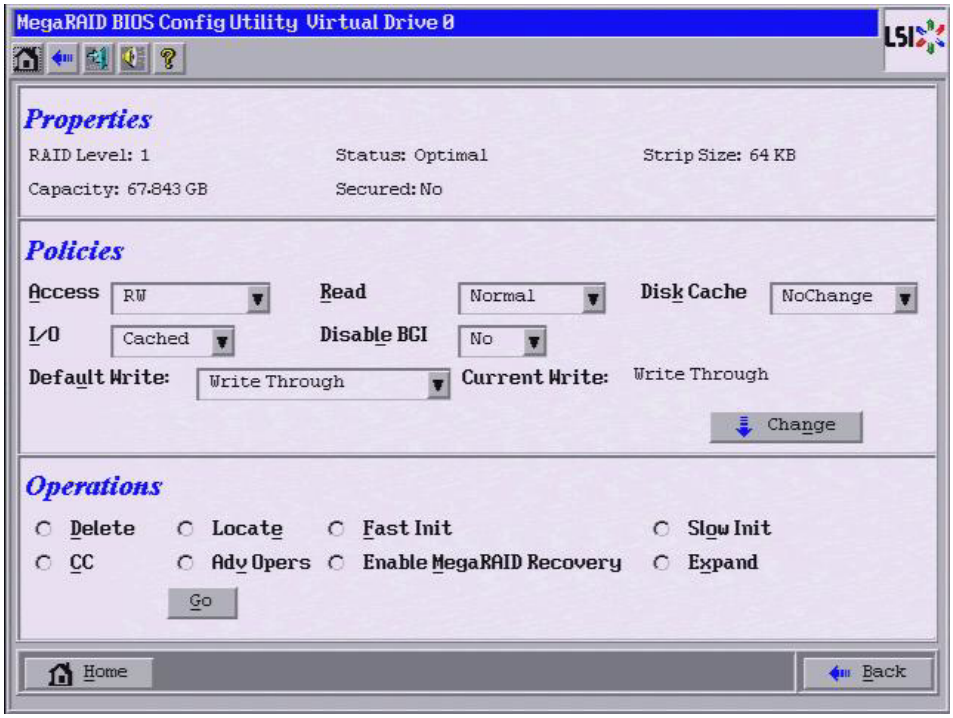


Figure 96: Virtual Drive Screen

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

   The Expand Virtual Drive screen appears, as shown in Figure 97.

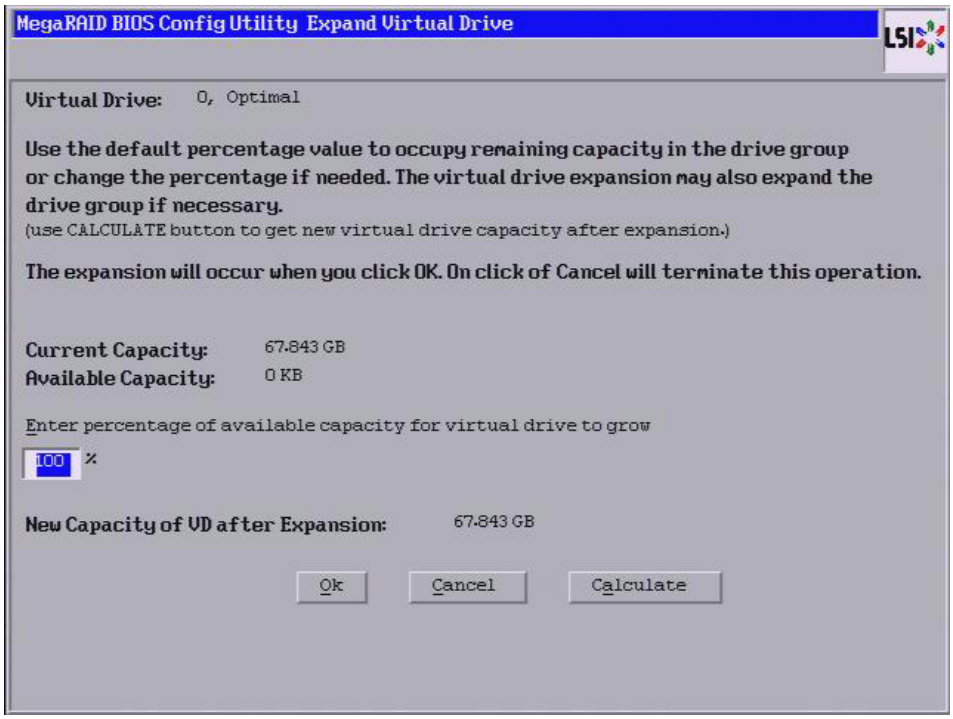


Figure 97: Expand Virtual Drive Screen

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 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) as the Snapshot Repository. Updates to the operating system or operating system crashes could destroy data on that virtual drive.

4.10.1 Recovery Scenarios

There are three primary scenarios in which to use the Recovery feature:

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" af 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 there are corrupt operating system files 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.

4.10.2 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 repository virtual drive.

Follow these steps to enable MegaRAID Recovery.

1. Click on a virtual drive icon in the right panel on the WebBIOS CU main screen to access the Virtual Drive screen.

   The Virtual Drive screen appears, as shown in Figure 98.

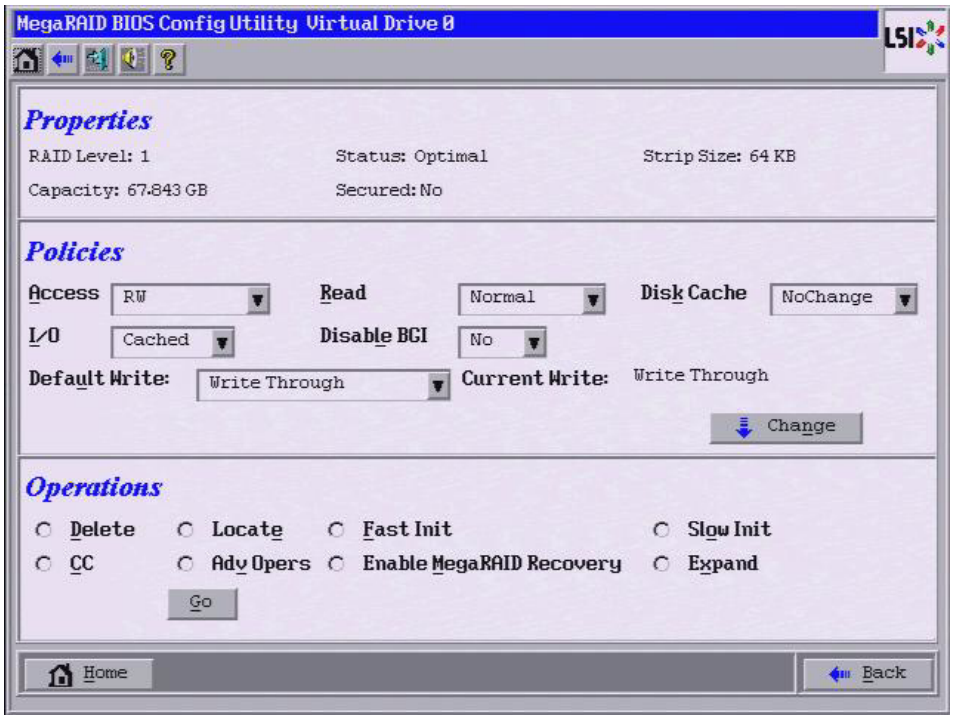


Figure 98: Virtual Drive Screen

2. Click **Enable MegaRAID Recovery** in the Operations panel of the screen.
3. Click **Go** in the Operations panel of the screen.

   The Enable MegaRAID Recovery screen appears, as shown in Figure 99.

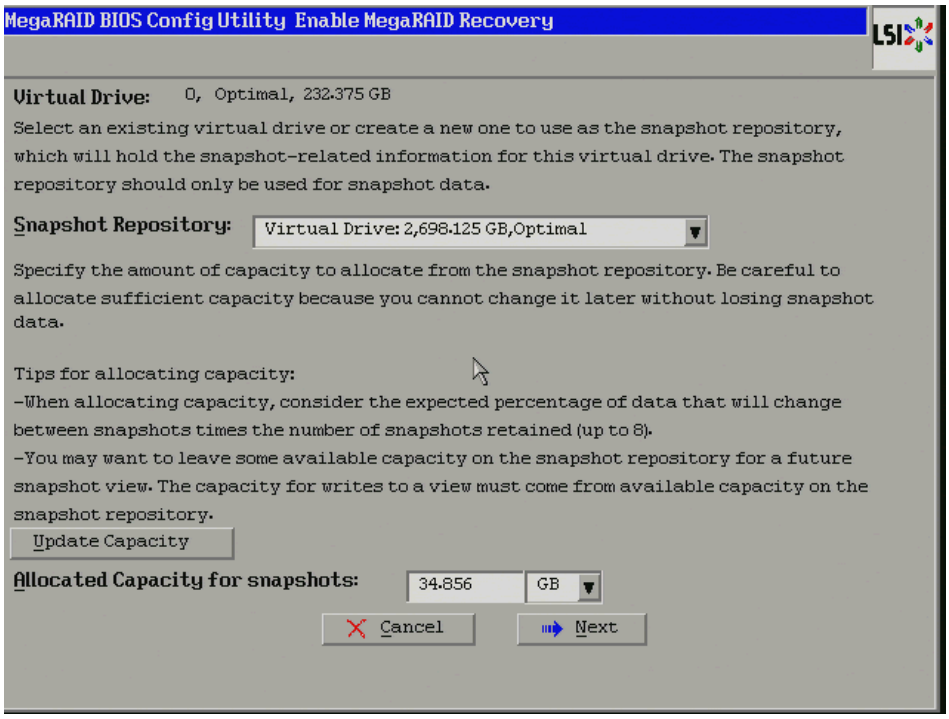


Figure 99: Enable MegaRAID Recovery Screen

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

   This 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 that you are taking snapshots of. 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 Repository virtual drive, there will be insufficient space to create a snapshot and a view.

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

7. Click **Next.**

   The snapshot settings screen appears as shown in Figure 107.

   Click Finish.

   A confirmation dialog box appears.

8. 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.

4.10.3 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 4.10.2, *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 screen to go to the operations for the virtual drive.

   The Virtual Drive screen appears, as shown in Figure 100.

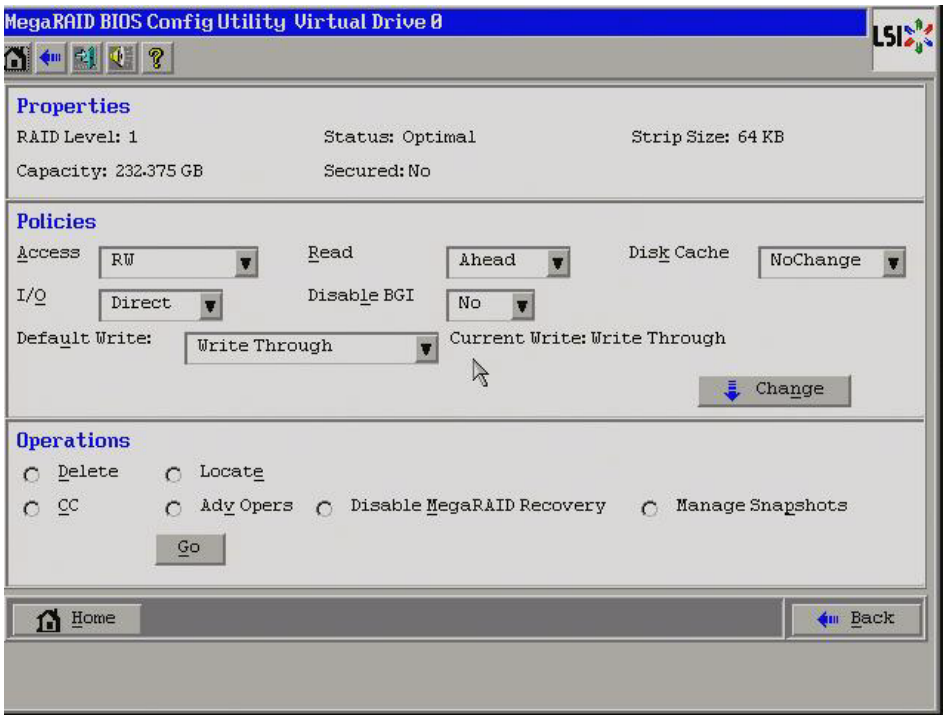


Figure 100: Snapshot Base Virtual Drive Operations

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

   The Virtual Drive Properties screen appears, as shown in Figure 101.

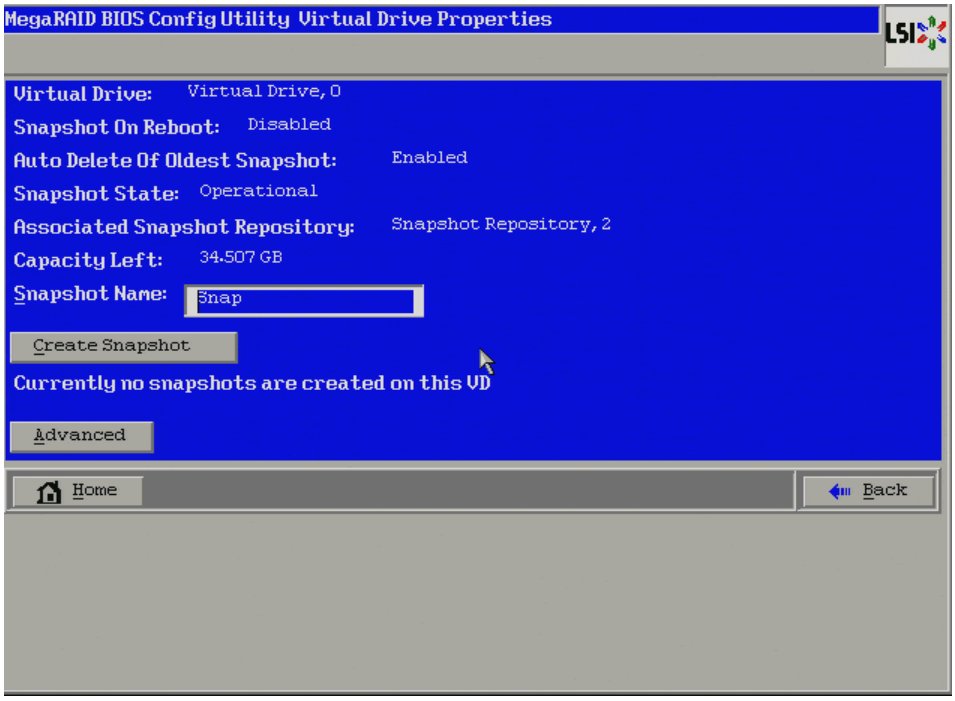


Figure 101: Properties

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

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

5. Click on the link of a specific snapshot.

   The snapshot details appear.

6. Click **Advanced**.

   The Snapshot Settings screen appears as shown in Figure 107.

7. Click **Create View**.

   The Create View screen appears, as shown in Figure 102.

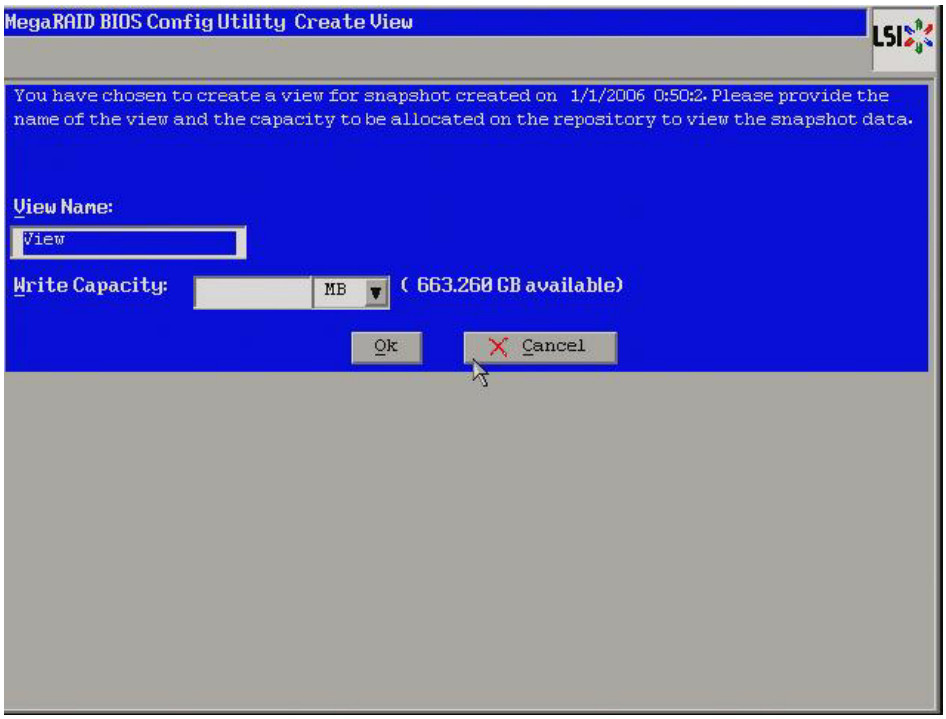


Figure 102: Create View Screen

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

   This 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 103.

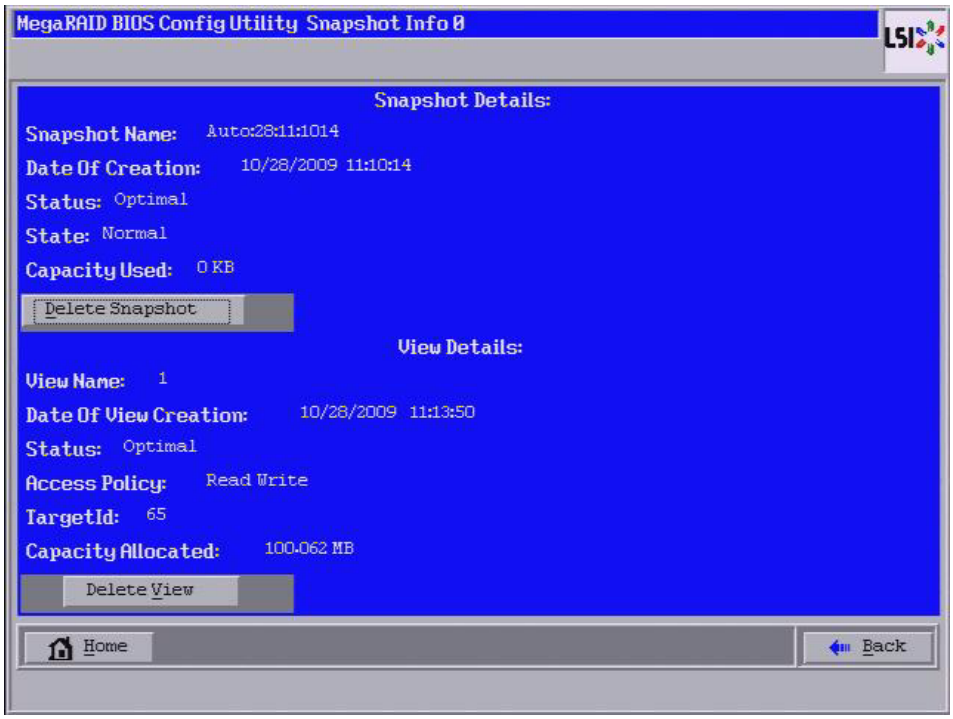


Figure 103: Snapshot and View Details

4.10.4 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 timestamp.

Follow these steps to create concurrent snapshots.

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

   The first Controller Properties screen displays. There are three Controller Properties screen.

2. Click **Next** to on the first two Controller Properties screens to access the third Controller Properties screen, as shown in Figure 104.

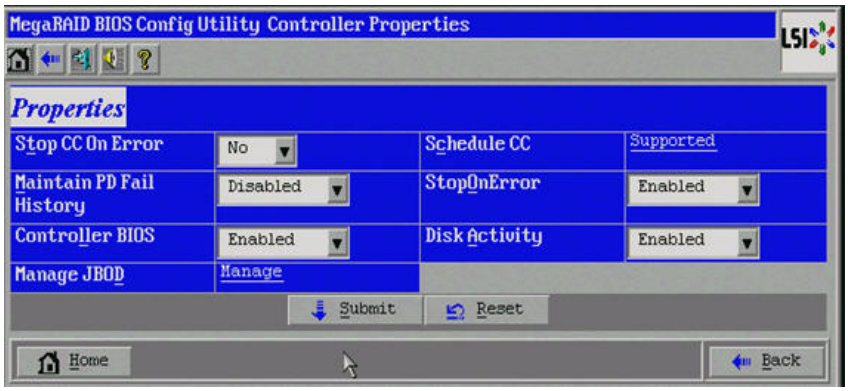


Figure 104: Third Controller Properties Screen

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

   The Snapshot Base Properties screen appears, as shown in Figure 105.

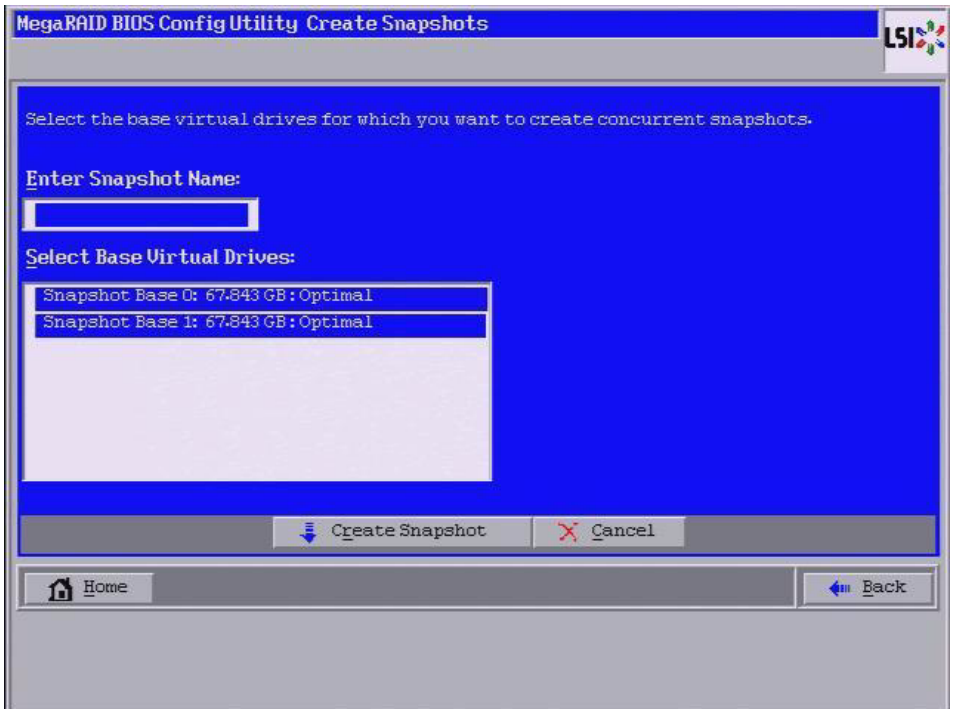


Figure 105: Create Snapshots Screen

4. Enter a snapshot name in the **Enter Snapshot Name** field.
5. Select the Snapshot Base virtual drives on which you want to create concurrent snapshots.
6. Click **Create Snapshot**.

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

4.10.5 Selecting the Snapshot Settings

You can use the Snapshot Settings screen 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 CU main screen.

   The Virtual Drive screen appears, as shown in Figure 106.

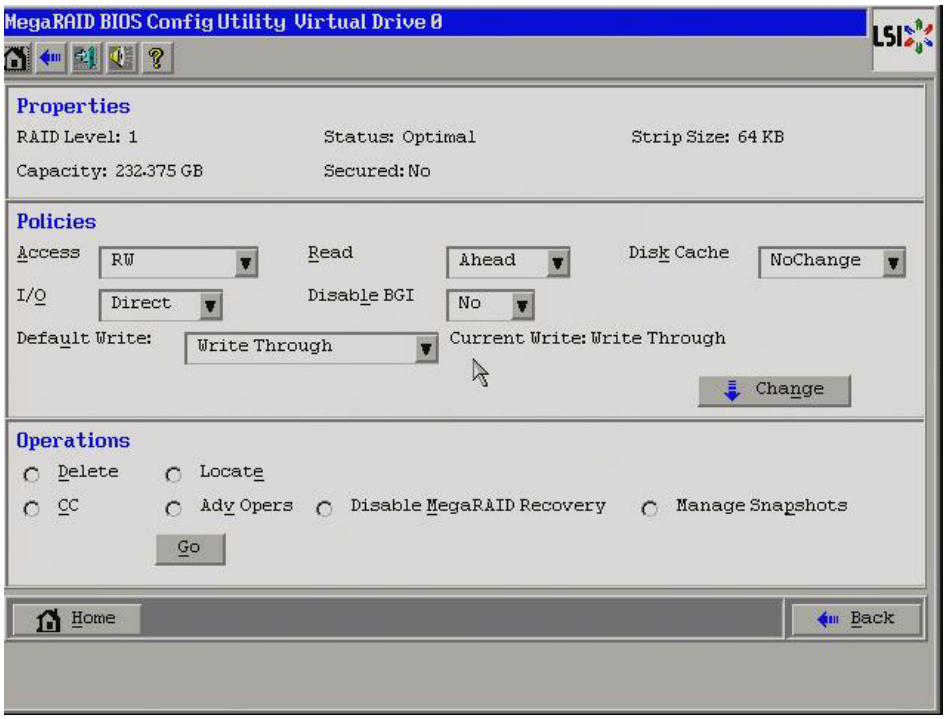


Figure 106: Virtual Drive Screen

2. Click **Manage Snapshots**.

   The Snapshot Settings screen appears, as shown in Figure 107.

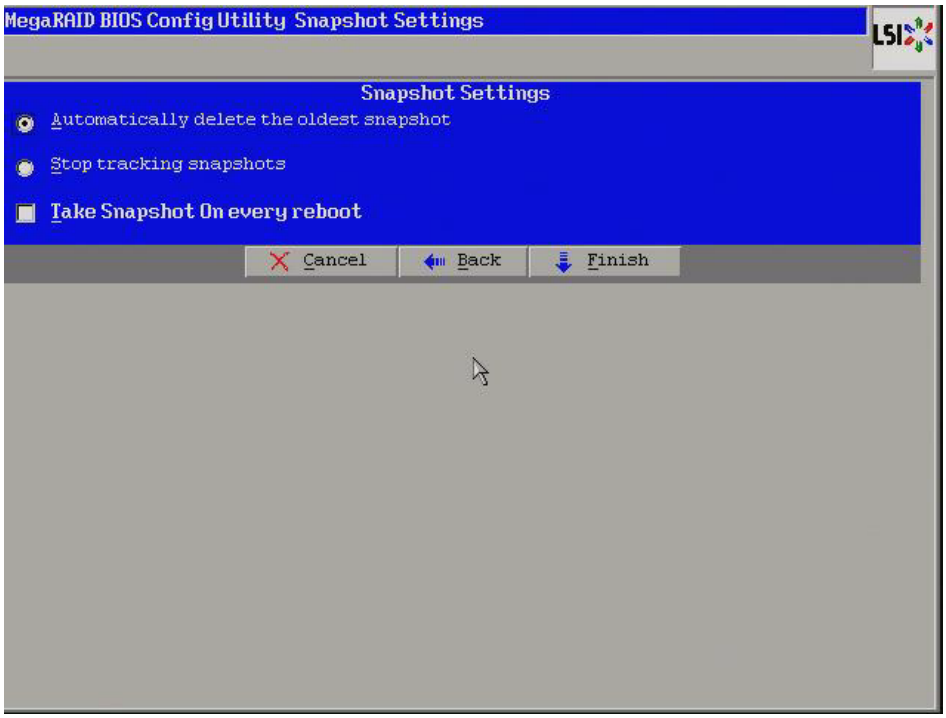


Figure 107: Snapshot Settings Screen

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

4.10.6 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 CU main screen.

   The Virtual Drive screen appears, as shown in Figure 108.

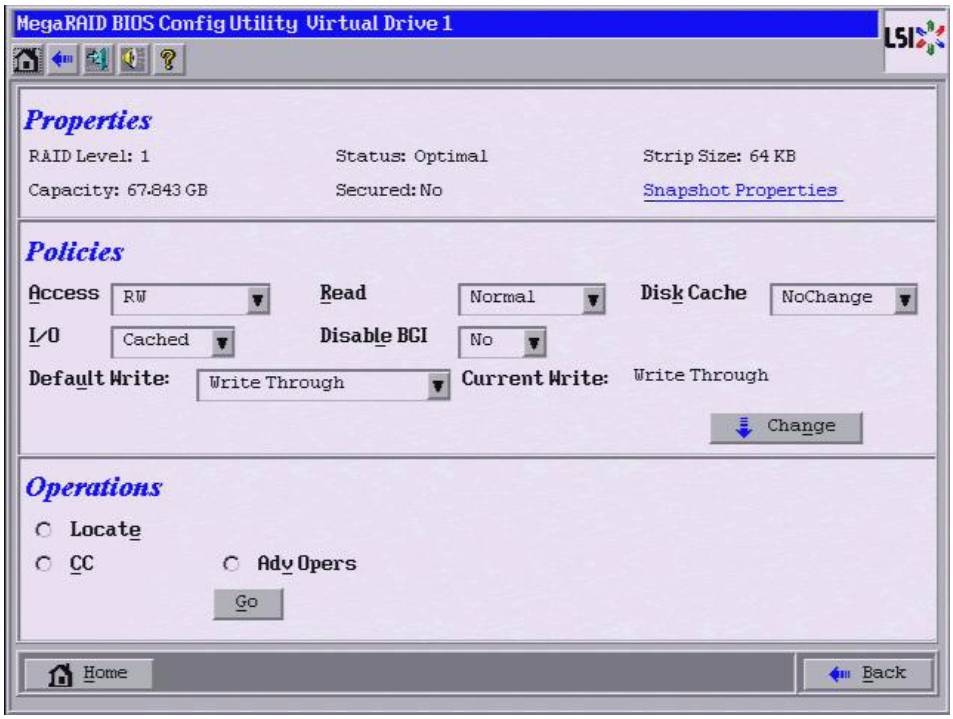


Figure 108: Virtual Drive Screen

2. Click **Snapshot Properties**.

   The Snapshot Repository Properties screen appears, as shown in Figure 109.

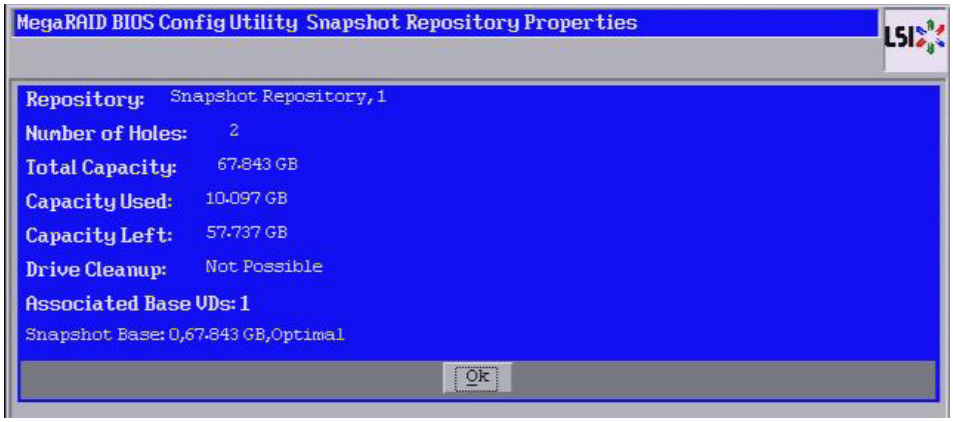


Figure 109: Snapshot Repository Properties

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

4.10.7 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 screen by clicking on a virtual drive icon in the right panel on the WebBIOS CU main screen.

   The Virtual Drive screen appears, as shown in Figure 110.

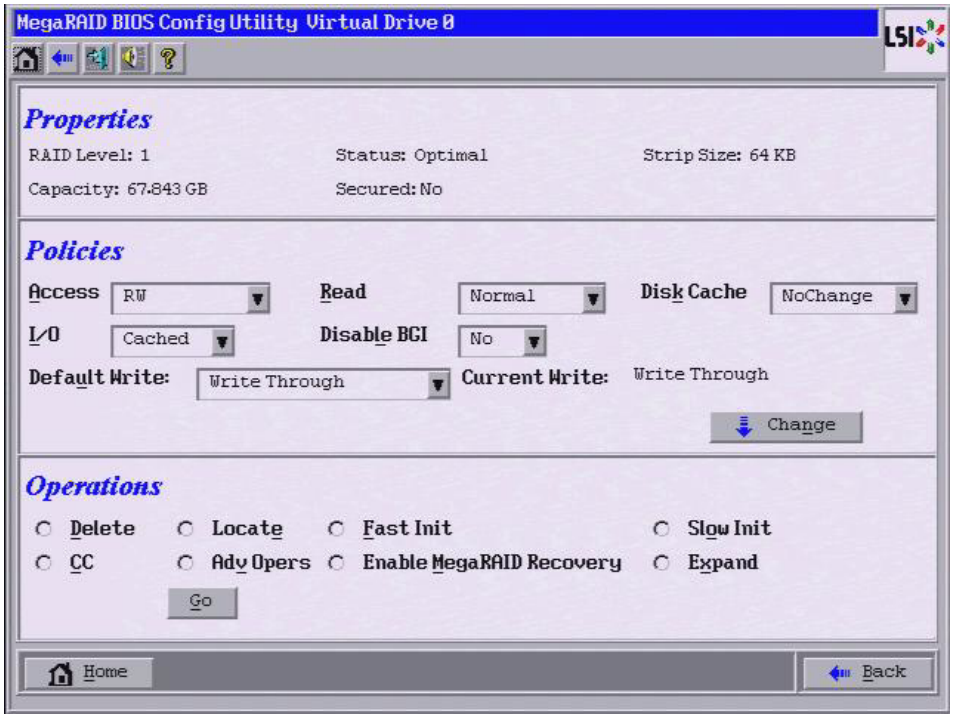


Figure 110: Virtual Drive Screen

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

   The Advanced Operations screen appears, as shown in Figure 111.

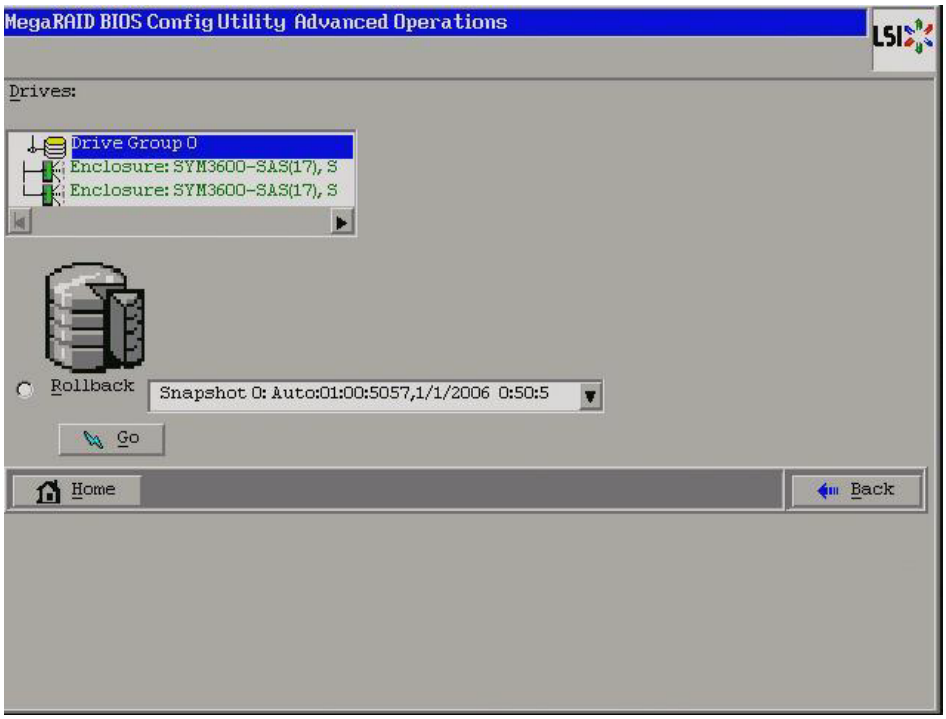


Figure 111: Advanced Operations Screen

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

   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.

4.10.8 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 screen by clicking on a Snapshot Repository virtual drive icon in the right panel on the WebBIOS CU main screen.

   The Virtual Drive screen appears, as shown in Figure 112.

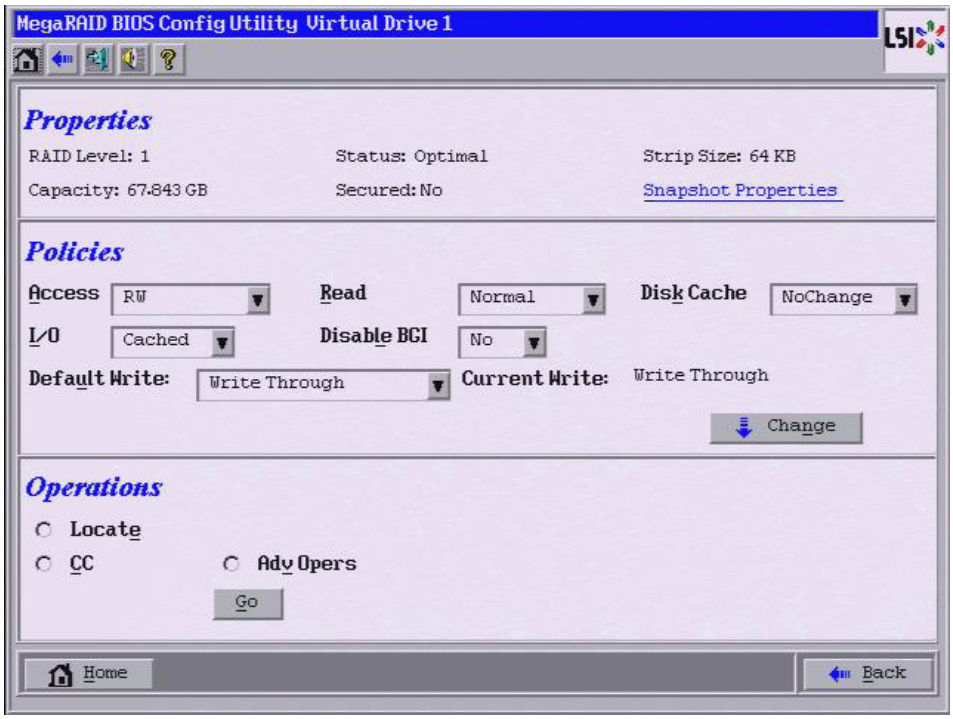


Figure 112: Virtual Drive Screen

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

   The Advanced Operations screen appears, as shown in Figure 113.

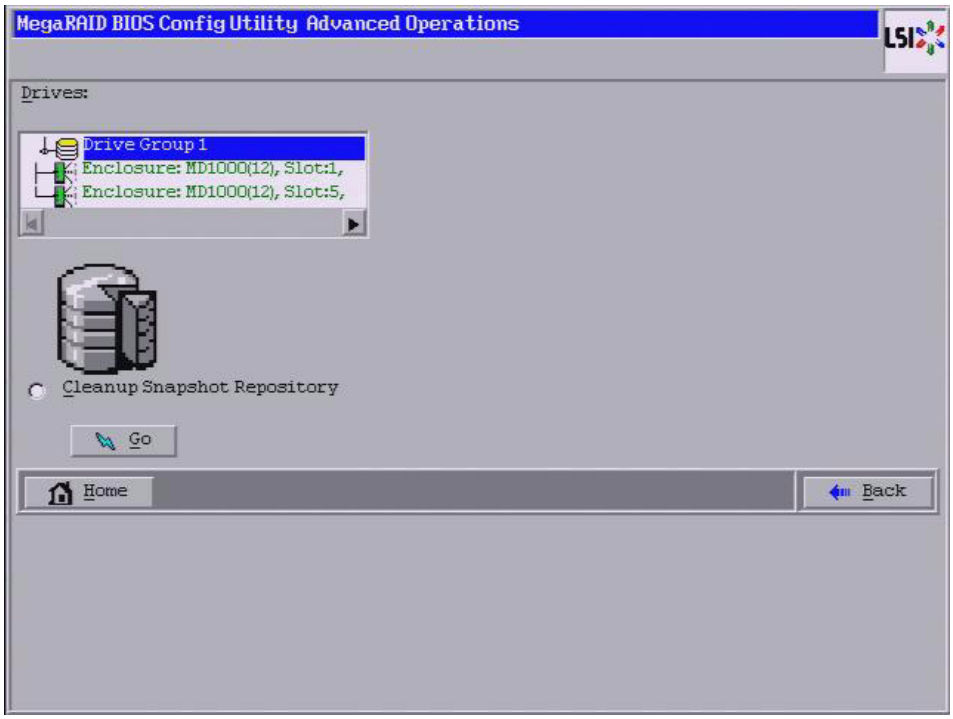


Figure 113: Advanced Operations Screen

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

   This action cleans up the Snapshot Repository.

4.11 Viewing System Event Information

The SAS controller firmware monitors the activity and performance of all storage configurations and devices in the system. When an event occurs (such as the creation of a new virtual drive or the removal of a drive) an event message is generated and is stored in the controller NVRAM.

You can use the WebBIOS CU to view these event messages. To do this, click **Events** on the main WebBIOS CU screen. The Event Information screen appears, as shown in Figure 114.

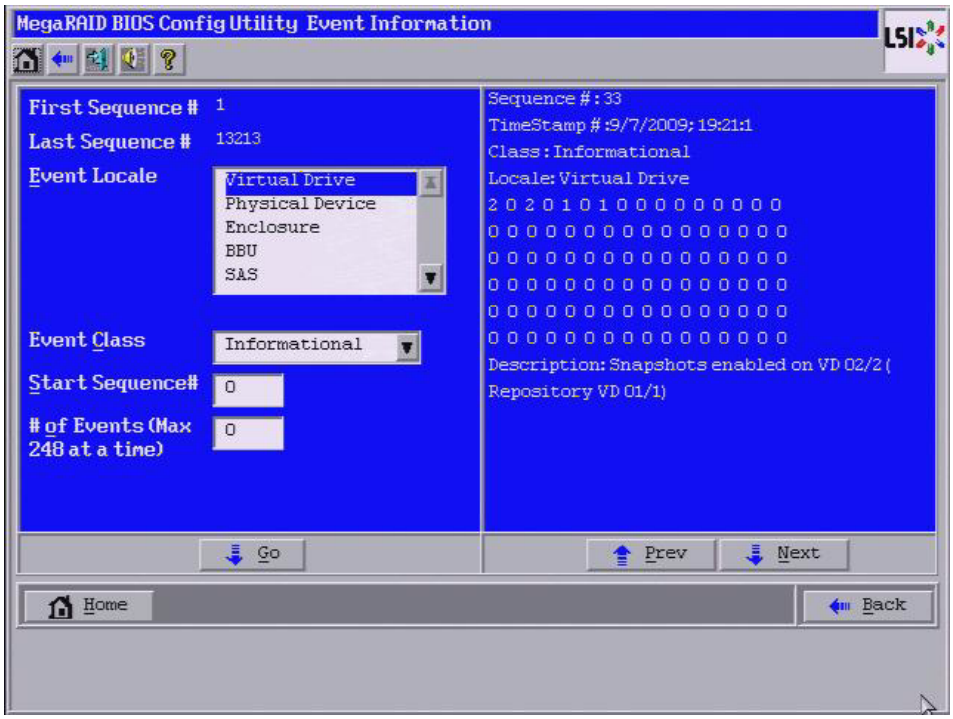


Figure 114: Event Information Screen

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

To view event information, follow these steps:

1. Select an Event Locale from the menu. For example, select **Enclosure** to view events relating to the drive enclosure.
2. Select an Event Class: *Information*, *Warning*, *Critical*, *Fatal*, or *Dead*.
3. Enter a Start Sequence number, between the First Sequence and Last Sequence numbers. The higher the number, the more recent the event.
4. Enter the Number of events of this type that you want to view, and click **Go**.

   The first event in the sequence appears in the right panel.
5. Click **Next** or **Prev** to page forward or backward through the sequence of events.
6. If you want, select different event criteria in the left panel, and click **Go** again to view a different sequence of events.

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

4.12 Managing Configurations

This section includes information about maintaining and managing storage configurations.

4.12.1 Running a Consistency Check

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

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

   The consistency check begins.

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

4.12.2 Deleting a Virtual Drive

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

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 screen by clicking on a virtual drive icon in the right panel on the WebBIOS CU main screen.

   The Virtual Drive screen 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.

4.12.3 Importing or Clearing a Foreign Configuration

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

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

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

If WebBIOS CU detects a foreign configuration, the Foreign Configuration screen appears, as shown in Figure 115.

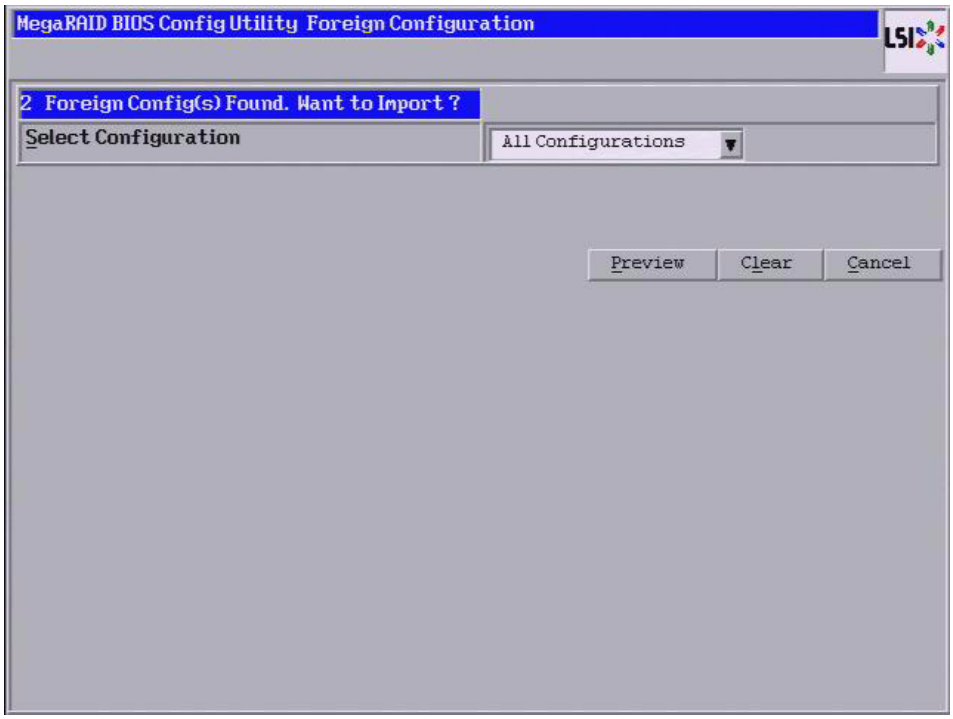


Figure 115: Foreign Configuration Import Screen

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. Select a configuration or **All Configurations**.
3. Perform one of the following steps:
   a. Click **Preview** to preview the foreign configuration(s).
      The Foreign Configuration Preview screen appears, as shown in Figure 116.
   b. Click **Clear** to clear the foreign configuration(s) and reuse the drives for another virtual drive.

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

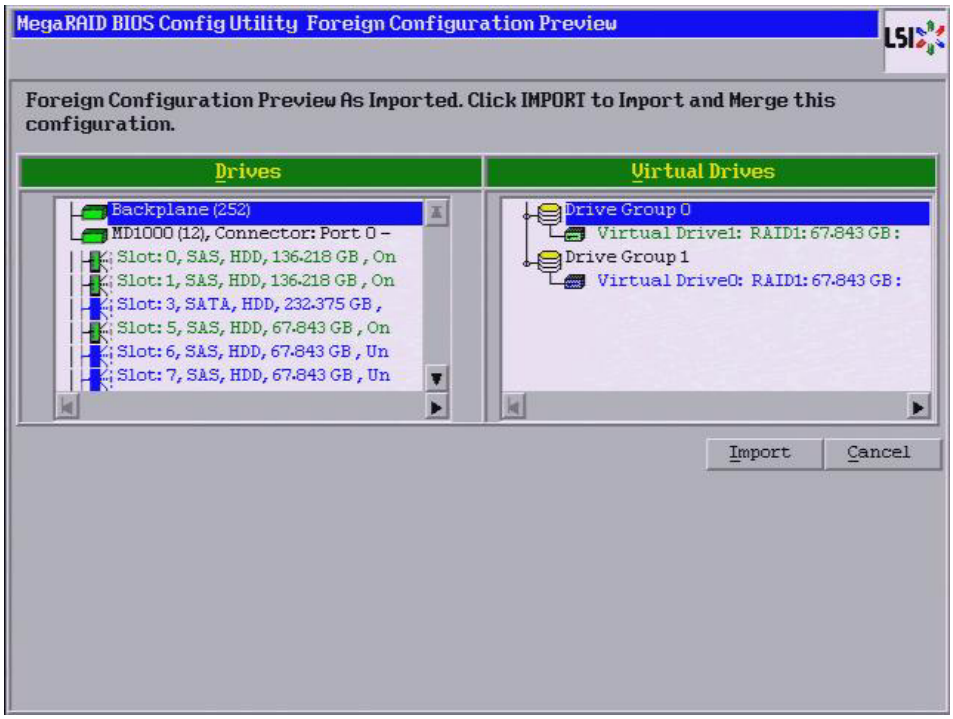


Figure 116: Foreign Configuration Preview Screen

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 this foreign configuration(s) and use it on this controller.

   If you click **Cancel**, you return to Figure 115.

4.12.3.1 Foreign Configurations in Cable Pull and Drive Removal Scenarios

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

Use the **Foreign Configuration Preview** screen to import or clear the foreign configuration in each case. The import procedure and clear procedure are described in Section 4.12.3, *Importing or Clearing a Foreign Configuration*

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

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

- 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 4.12.1, *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 4.12.1, *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.

4.12.3.2 Importing Foreign Configurations from Integrated RAID to MegaRAID

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

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

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

4.12.3.3 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 megadata 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.

In order to import an IR virtual drive into a MegaRAID system, use the **Foreign Configuration Preview** screen to import or clear the foreign configuration. The import procedure and the clear procedure are described in Section 4.12.3, *Importing or Clearing a Foreign Configuration*.

4.12.4 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 utility to import the existing configuration to the RAID controller or clear the configuration so you can create a new one.

See Section 4.12.3, *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.

4.12.5 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 migrate a virtual drive, you can keep the same number of drives, or you can add drives. You can use the WebBIOS CU to migrate the RAID level of an existing virtual drive.

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

Migrations are allowed for the following RAID levels:

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

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

Table 22 lists the number of additional drives required when you change the RAID level of a virtual drive.

Table 22: 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 CU screen, select a virtual drive.
2. Click **Virtual Drives**.

   The Virtual Drive screen appears.
3. Select **AdvOpers** under the **Operations** heading.

   The Advanced Operations screen appears, as shown in Figure 117.

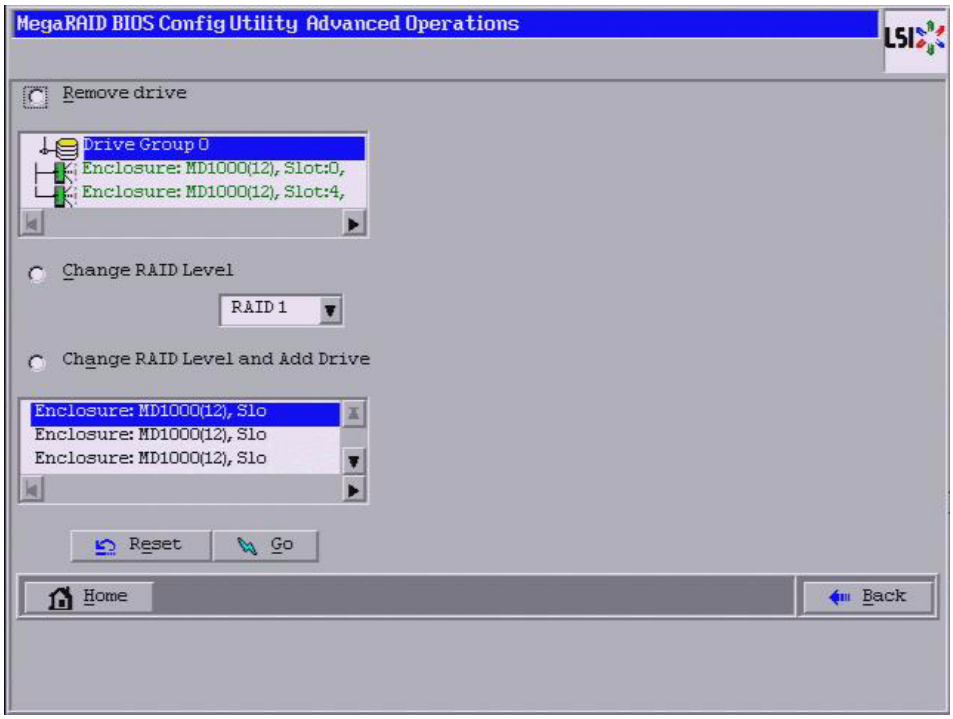


Figure 117: Advanced Operations Screen

4. 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 menu.
   - If you select **Change RAID Level and Add Drive**, change the RAID level from the drop-down menu and then 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.
5. Click **Go**.
6. When the message appears, confirm that you want to migrate the RAID level of the virtual drive.

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

4.12.6 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 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 4.12.3.3, *Troubleshooting Information* for more information about DDF and metadata.

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 MegaRAID SAS RAID controllers and the devices connected to them.

NOTE: The CT supports only 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 Windows is similar to the 32-bit release, so you can follow the 32-bit instructions. 32-bit applications that were validated on an x64 system, such as the Intel Markette system, can use the 32-bit instructions, also.

5.1 Product Overview

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

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

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

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

NOTE: The MegaCLI error messages are listed in Appendix B.

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

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

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

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

For the Solaris operating system, LSI 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

Table 23 lists the abbreviations for the virtual drive parameters used in the following sections.

Table 23: Command Line Abbreviations

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

5.3.2 Conventions

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

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

Table 24 describes the conventions used in the options.

Table 24: 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.

Table 24: Conventions (Continued)

Convention	Description
`-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.
`[E0:S0,E1,S1,…]`	Specifies when one or more physical devices need(s) to be specified in the command line. Each `[E:S]` pair specifies one physical device where `E` means device ID of the enclosure in which a drive resides, and `S` means the slot number of the enclosure. 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 (WriteThrough) is optional. If you enter `WT` at the command line, the application will use WriteThrough write policy for the virtual drive. Otherwise, it uses the default value for the parameter.
`{ }`	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 screen.

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 CLI 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 screen:

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 create CacheCade drives and delete them.

MegaRAID CacheCade improves application performance by expanding the MegaRAID read-caching capacity. 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 CacheCade 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 IO 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 SAS 6Gb/s 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.

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

Use the command in Table 25 to create a cache drive using CacheCade software. You can use that cache as secondary cache. CacheCade software have much greater capacity than HDDs.

Table 25: Create a Solid State Cache Drive to Use as Secondary Cache

Convention	**MegaCli -CfgSscdAdd -Physdrv[E0:S0,...] {-Name LdNamestring} -aN|-a0,1,2|-aALL**
Description	This command is used to create a CacheCade drive group 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 Table 26 to delete a CacheCade software cache drive or multiple CacheCade software cache drives on the selected controller(s).

Table 26: Delete Solid State Cache Drive(s)

Convention	**MegaCli -CfgSscdDel -LX|-L0,2,5...|-LALL -aN|-a0,1,2|-aALL**
Description	Deletes the specified CacheCade software cache drive or drives on the selected controller(s). You can delete multiple CacheCade software cache drives or all of the CacheCade software caches.

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 27: 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
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 (VD) 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 it is housed in, the data on that drive is encrypted and useless to anyone who attempts to access without the 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 Table 28 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 28: 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 drive(s) that you want to perform the Instant Secure Erase on. `-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:** 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.

5.7.2 Secure Data on a Virtual Drive

Use the command in Table 29 to secure data on a virtual drive.

Table 29: 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 Table 30 to destroy the security key.

Table 30: 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 Table 31 to create a security key.

Table 31: Create a Security Key

Convention	**MegaCli -CreateSecurityKey -SecurityKey sssssssssss | [-Passphrase sssssssssss] |[-KeyID kkkkkkkkkkk] -aN**
Description	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. `-CreateSecurityKey`: Creates the security key. `-SecurityKey sssssssssss`: Enters the new security key. The security key is case-sensitive. It must be between eight and thirty-two 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. `[-Passphrase sssssssssss]`: Enters the new passphrase.The pass phrase is case-sensitive. It must be between eight and thirty-two 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.

5.7.5 Drive Security Key

If you want to use the security key using 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 32: Drive Security Key

Convention	MegaCli -CreateSecurityKey useEKMS –aN
Description	Creates security key using EKMS.
Convention	MegaCli -ChangeSecurityKey -SecurityKey sssssssssss [-Passphrase sssssssssss] | [-KeyID kkkkkkkkkkk] –aN
Description	To change the security from EKMS to LKM.
Convention	MegaCli -ChangeSecurityKey useEKMS -OldSecurityKey sssssssssss –aN

Table 32: Drive Security Key

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 Table 33 to change they security key to a new security key.

Table 33: Change the Security Key

Convention	**MegaCli -ChangeSecurityKey -OldSecurityKey sssssssssss | -SecurityKey sssssssssss| [-Passphrase sssssssssss] | [-KeyID kkkkkkkkkkk] -aN**
Description	Changes a security key to a new security key. -ChangeSecurityKey: Changes the security key. -OldSecurityKey sssssssssss: Enters the old security key. The security key is case-sensitive. It must be between eight and thirty-two 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. -SecurityKey sssssssssss: Enters the new security key. The security key is case-sensitive. It must be between eight and thirty-two 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. [-Passphrase sssssssssss]: Enters the new passphrase. The pass phrase is case-sensitive. It must be between eight and thirty-two 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. [-KeyID kkkkkkkkkkk]: 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.

5.7.7 Get the Security Key ID

Use the command in Table 34 to display the security key ID.

Table 34: Get the Security Key ID

Convention	**MegaCli -GetKeyID [-PhysDrv[E0:S0]] -aN**
Description	-GetKeyID: Displays the security key ID.

5.7.8 Set the Security Key ID

Use the command in Table 35 to set the security key ID.

Table 35: Set the Security Key ID

Convention	**MegaCli -SetKeyID -KeyID kkkkkkkkkkk -aN**
Description	-SetKeyID: Set the security key ID. -KeyID kkkkkkkkkkk: 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.

5.7.9 Verify the Security Key

Use the command in Table 36 to verify the security key.

Table 36: Verify the Security Key ID

Convention	**MegaCli -VerifySecurityKey -SecurityKey sssssssssss -aN**
Description	Verifies that the security key is the correct one for the self-encrypted disk. -VerifySecurityKey: Verifies the security key. -SecurityKey sssssssssss: Enters the new security key. The security key is case-sensitive. It must be between eight and thirty-two 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 controller(s), such as the virtual drive parameters and factory defaults.

5.8.1 Display Controller Properties

Use the command in Table 37 to display parameters for the selected controller(s).

Table 37: Controller Parameters

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

5.8.2 Display Number of Controllers Supported

Use the command in Table 38 to display the number of controllers supported on the system.

Table 38: 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 Table 39 to turn automatic rebuild on or off for the selected controller(s). If you have configured hot spares and enabled automatic rebuild, the RAID controller automatically tries to use them to rebuild failed drives. Automatic rebuild also controls whether a rebuild starts when a drive that was part of the drive group is reinserted.

Table 39: 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 controller(s). The `-Dsply` option shows the status of the automatic rebuild state.

5.8.4 Flush Controller Cache

Use the command in Table 40 to flush the controller cache on the selected controller(s). This option sends the contents of cache memory to the virtual drive(s). If the MegaRAID system must be powered down rapidly, you must flush the contents of the cache memory to preserve data integrity.

Table 40: Cache Flush on Selected Controller

Convention	**MegaCli -AdpCacheFlush -aN|-a0,1,2|-aALL**
Description	Flushes the controller cache on the selected controller(s).

5.8.5 Set Controller Properties

This command sets the properties on the selected controller(s).
For example, for `{RebuildRate -val}`, you can enter a percentage between 0 percent and 100 percent as the value for the rebuild rate.
(The rebuild rate is the percentage of the compute cycles dedicated to rebuilding failed drives.) At 0 percent, the rebuild is done only if the system is not doing anything else. At 100 percent, the rebuild has a higher priority than any other system activity.

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

Use the command in Table 41 to display the list of properties you can set for the controller(s).

Table 41: Set Controller Properties

Convention	**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**
Description	Sets the properties on the selected controller(s). The possible settings are: `CacheFlushInterval`: Cache flush interval in seconds. Values: 0 to 255. `RebuildRate`: Rebuild rate. Values: 0 to 100. `PatrolReadRate`: Patrol read rate. Values: 0 to 100. `BgiRate`: Background initilization rate. Values: 0 to 100. `CCRate`: Consistency check rate. Values: 0 to 100. `ReconRate`: Reconstruction rate. Values: 0 to 100. `SpinupDriveCount`: Max number of drives to spin up at one time. Values: 0 to 255. `SpinupDelay`: Number of seconds to delay among spinup groups. Values: 0 to 255. `CoercionMode`: Drive capacity Coercion mode. Values: 0 - None, 1 - 128 MB, 2 - 1 GB. `ClusterEnable`: Cluster is enabled or disabled. Values: 0 - Disabled, 1 - Enabled. `PredFailPollInterval`: Number of seconds between predicted fail polls. Values: 0 to 65535. `BatWarnDsbl`: Disable warnings for missing battery or missing hardware. Values: 0 - Enabled, 1 - Disabled. `EccBucketSize`: Size of ECC single-bit-error bucket. Values: 0 to 255. `EccBucketLeakRate`: Leak rate (in minutes) of ECC single-bit-error bucket. Values: 0 to 65535. AbortCCOnError: `AlarmEnbl`: Set alarm to Enabled. `AlarmDsbl`: Set alarm to Disabled. `AlarmSilence`: Silence an active alarm. `SMARTCpyBkEnbl`: Enable copyback operation on Self-Monitoring Analysis and Reporting Technology (SMART) errors. Copyback is initiated when the first SMART error occurs on a drive that is part of a virtual drive. `AutoDetectBackPlaneDsbl`: Detect automatically if the backplane has been disabled. `CopyBackDsbl`: Disable or enable the copyback operation. `LoadBalanceMode`: Disable or enable the load balancing mode. `NCQEnbl`: Enable the native command queueing. `NCQDsbl`: Disable the native command queueing. `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 an SSD that is part of a virtual drive. `MaintainPdFailHistoryEnbl`: Enable maintenance of the history of a failed drive. `EnblSpinDownUnConfigDrvs`: Enable spindown of unconfigured drives. `AutoEnhancedImportEnbl`: Enable the automatic enhanced import of foreign drives. `AutoEnhancedImportDsbl`: Disable the automatic enhanced import of foreign drives. `UseFDEOnlyEncrypt`: Use encryption on FDE drives only. `PrCorrectUncfgdAreas`:

5.8.6 Display Specified Controller Properties

Use the command in Table 42 to display specified properties on the selected controller(s).

Table 42: 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 | -aN|-a0,1,2|-aALL**
Description	Displays the properties on the selected controller(s). `EccBucketCount`: Count of single-bit ECC errors currently in the bucket. `WBSupport`: Enables support for the WriteBack option as the Write Policy. `DsblSpinDownUnConfigDrvs`: Disable spindown of unconfigured drives. See Table 41 for explanations of the other options.

NOTE: The *tty* log can be saved at the controller level.

5.8.7 Set Factory Defaults

Use the command in Table 43 to set the factory defaults on the selected controller(s).

Table 43: Set Factory Defaults

Convention	**MegaCli –AdpFacDefSet -aN|-a0,1,2|-aALL**
Description	Sets the factory defaults on the selected controller(s).

5.8.8 Set SAS Address

Use the command in Table 44 to set the SAS address on the selected controller(s).

Table 44: Set SAS Address on Controller

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

5.8.9 Set Time and Date on Controller

Use the command in Table 45 to set the time and date on the selected controller(s).

Table 45: Set Time and Date on Controller

Convention	**MegaCli –AdpSetTime yyyymmdd HH:mm:ss -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 Table 47 to display the time and date on the selected controller(s).

Table 46: 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 Table 47 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 MegaRAID SAS 8888ELP RAID controller at this time. You must enable specific ports depending on how you intend to use the controller.

Table 47: 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 Table 48 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 48: 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 Table 49 on the selected controller(s) to set the Patrol Read options.

Table 49: Set Patrol Read Options

Convention	**MegaCli –AdpPR –Dsbl|EnblAuto|EnblMan|Start|Stop|Info | {-SetStartTime yyyymmdd hh} | maxConcurrentPD -aN| -a0,1,2|-aALL**
Description	Sets Patrol Read options on a single controller, multiple controllers, or all controllers: `-Dsbl`: Disables Patrol Read for the selected controller(s). `-EnblAuto`: Enables Patrol Read automatically for the selected controller(s). This means Patrol Read will start automatically after the controller initialization is complete. `-EnblMan`: Enables Patrol Read manually for the selected controller(s). This means that Patrol Read does not start automatically; it has to be started manually by selecting the `Start` command. `-Start`: Starts Patrol Read for the selected controller(s). `-Stop`: Stops Patrol Read for the selected controller(s). `-Info`: Displays the following Patrol Read information for the selected controller(s): • Patrol Read operation mode • Patrol Read execution delay value • Patrol Read status `SetStartTime yyyymmdd hh`: Set the start time for the patrol read in year/month/day format. maxConcurrentPD: Sets the maximum number of concurrent drives that patrol read runs on.

5.9.2 Set Patrol Read Delay Interval

Use the command in Table 50 on the selected controller(s) to set the time between Patrol Read iterations.

Table 50: 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 Table 51 to set or display the ID of the bootable virtual drive.

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

Table 51: 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 from. `-Get`: Displays the bootable virtual drive ID.

5.10.2 Select BIOS Status Options

Use the command in Table 52 to set the options for the BIOS status.

Table 52: Options for BIOS Status

Convention	**MegaCli -AdpBIOS -Enbl|-Dsbl| SOE | BE | EnblAutoSelectBootLd | DsblAutoSelectBootLd |-Dsply| -aN|-a0,1,2|-aALL**
Description	Sets 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 controller(s). `-SOE`: Stops on BIOS errors during POST for selected controller(s). When set to `-SOE`, the BIOS stops in case of a problem with the configuration. This gives you the option to enter the configuration utility to resolve the problem. This is available only when you enable the BIOS status. `-BE`: Bypasses BIOS errors during POST. This is available only when you enable the BIOS status. `-EnblAutoSelectBootLd | DsblAutoSelectBootLd`: Enable or disable automatic selection of the boot virtual drive. `-Dsply`: Displays the BIOS status on selected controller(s).

5.11 Battery Backup Unit-Related Properties

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

5.11.1 Display BBU Information

Use the command in Table 53 to display complete information about the BBU for the selected controller(s).

Table 53: Display BBU Information

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

5.11.2 Display BBU Status Information

Use the command in Table 54 to display complete information about the status of the BBU, such as temperature and voltage, for the selected controller(s).

Table 54: Display BBU Status Information

Convention	**`MegaCli -AdpBbuCmd -GetBbuStatus -aN|-a0,1,2|-aALL`**
Description	Displays complete information about the BBU status, such as the temperature and voltage. The information displays in the following formats: **BBU Status for Adapter: xx** Battery Type: XXXXXX(string) Voltage: xx mV Current: xx mA Temperature: xx C° Firmware Status: xx Battery state: xx **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 Additional status information displays differently for iBBU™ and BBU. **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% **For BBU:** Relative State of Charge: xx Charger Status: xx Remaining Capacity: xx mAh Full Charge Capacity: mAh isSOHGood: Yes/No

5.11.3 Display BBU Capacity

Use the command in Table 55 to display the BBU capacity for the selected controller(s).

Table 55: Display BBU Capacity Information

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

5.11.4 Display BBU Design Parameters

Use the command in Table 56 to display BBU design parameters for the selected controller(s).

Table 56: 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 Table 57 to display the current BBU properties for the selected controller(s).

Table 57: Display Current BBU Properties

Convention	**`MegaCli -AdpBbuCmd -GetBbuProperties -aN|-a0,1,2|-aALL`**
Description	Displays current properties of the BBU. The information displays in the following formats: BBU Properties for Adapter: x Auto Learn Period: xxx Sec Next Learn Time: xxxx Sec Learn Delay Interval=<value>: Value in hours, not greater than 168 hours (7 days) Auto-Learn Mode=<value>: Value can be 0, 1, or 2. **NOTE:** If the battery type is IBBU08, then the BBU mode is displayed as a part of **`GetBbuProperties.`**

5.11.6 Start BBU Learning Cycle

Use the command in Table 58 to start the BBU learning cycle on the selected controller(s). A learning cycle is a battery calibration operation performed by the controller periodically (approximately every three months) to determine the condition of the battery.

Table 58: 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 Table 59 to place the battery into Low-Power Storage mode on the selected controller(s). This saves battery power consumption.

Table 59: 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 Table 60 to set the BBU properties on the selected controller(s) after reading from the file.

Table 60: Set BBU Properties

Convention	**MegaCli -AdpBbuCmd -SetBbuProperties -f<fileName> -aN|-a0,1,2|-aALL**
Description	Sets the BBU properties on the selected controller(s) after reading from the file. The information displays in the following formats: autoLearnPeriod = 1800Sec nextLearnTime = 12345678Sec Seconds past 1/1/2000 learnDelayInterval = 24hours - Not greater than 7 days autoLearnMode = 0 0 – Enabled; 1 - Disabled; 2 – WarnViaEvent. **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 Table 61 to manage the event entries in the event log for the selected controller(s).

Table 61: Event Log Management

Convention	**MegaCli –AdpEventLogInfo –GetEventlogInfo | –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. `Start_entry` 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 controller(s). `-GetCCIncon`: Displays the events relating to inconsistent data found during a consistency check.

5.12.2 Set BBU Terminal Logging

Use the command in Table 62 to set the BBU terminal logging for the selected controller(s).

Table 62: Set BBU Terminal Logging

Convention	**`MegaCli -FwTermLog -Bbuoff |-BbuoffTemp|-Bbuon|-BbuGet| -aN| -a0,1,2|-aALL`**
Description	Sets BBU terminal logging options. The following are the settings you can select on a single controller, multiple controllers, or all controllers: `-Bbuoff`: Turns off the BBU for firmware terminal logging. To turn off the BBU for logging, you have to shut down your system or turn off the power to the system after you run the command. `-BbuoffTemp`: Temporarily turns off the BBU for TTY (firmware terminal) logging. The battery will be turned on at the next reboot. `-Bbuon`: Turns on the BBU for TTY (firmware terminal) logging. TTY logs are cached. `-BbuGet`: Displays the current BBU settings for TTY logging, whether TTY logging is enabled or not.

5.13 Configuration-Related Options

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

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

5.13.1 Create a RAID Drive Group from All Unconfigured Good Drives

Use the command in Table 63 to create one RAID drive group out of all of the unconfigured good drives, and a hot spare, if desired. This is for RAID levels 0, 5, 6, 10, 50, or 60. All free drives are used to create a new drive group and, if desired, one hot spare drive. If it is not possible to use all of the free drives, the command will abort with a related error level. If there are drives of different capacities, the largest drive is used to make the hot spare.

NOTE: 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 63: 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	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. `-Rx[E0:S0,...]`: Specifies the RAID level and the drive enclosure/slot numbers used to construct a drive group. `-WT` (Write through), `WB` (Write back): Selects write policy. `-NORA` (No read ahead), `RA` (Read ahead), `ADRA` (Adaptive read ahead): Selects read policy. `-Direct`, `-Cached`: Selects cache policy. `-CachedBadBBU|NoCachedBadBBU`: Specifies whether to use write cache when the BBU is bad. `Hsp`: Specifies drive to make the hot spare with. `-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. **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.

5.13.2 Add RAID 0, 1, 5, or 6 Configuration

Use the command in Table 64 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 64: 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 [-szYYYYYYYY [...]]] [-strpszM] [-Hsp[E5:S5,...]] [-afterLdX] [-Force] -aN**
Description	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. Note that RAID 1 supports up to 32 drives in a single span of 16 drive groups. RAID 1 requires an even number of drives, as data from one drive is mirrored to the other drive in each RAID 1 drive group. `-Rx[E0:S0,...]`: Specifies the RAID level and the drive enclosure/slot numbers to construct a drive group. `-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 : Specifies whether to use write cache when the BBU is bad. `-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. The configuration of a particular virtual drive will fail if the remaining capacity of the drive group is too small to configure the virtual drive with the specified capacity. This option can also be used to create a configuration on the free space available in the drive group. `-strpszM`: Specifies the stripe size, where the stripe size values are 8, 16, 32, 64, 128, 256, 512, or 1024 KB. `Hsp[E5:S5,...]`: Creates hot spares when you create the configuration. The new hot spares will be dedicated to the virtual drive used in creating the configuration. This option does not allow you to create global hot spares. To create global hot spares, you must use the `-PdHsp` command with proper subcommands. You can also use this option to create a configuration on the free space available in the virtual drive. You can specify which free slot should be used by specifying the `-AfterLdX`: This command is optional. By default, the application uses the first free slot available in the virtual drive. This option is valid only if the virtual drive is already used for configuration.

5.13.3 Add RAID 10, 50, or 60 Configuration

Use the command in Table 65 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 65: 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 [-szYYYYYYYY [...]]] [-strpszM] [–afterLdX] | -Force [FDE|CtrlBased] -aN|-a0,1,2|-aALL**
Description	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. 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. Multiple drive groups are specified using the `–ArrayX[E0:S0,...]` 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. The size option, `–szXXXXXXXX`, can be accepted to allow slicing in the spanned drive groups if the controller supports this feature. The `[–afterLdX]` option is accepted if the size option is accepted. CT exits and does not create a configuration if the size or the `afterLd` option is specified but the controller does not support slicing in the spanned drive groups. **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.

5.13.4 Clear the Existing Configuration

Use the command in Table 66 to clear the existing configuration on the selected controller(s).

Table 66: 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 Table 67 to save the configuration for the selected controller(s) to the given filename.

Table 67: Save Configuration on the Controller

Convention	**MegaCli –CfgSave –f FileName -aN**
Description	Saves the configuration for the selected controller(s) to the given filename.

5.13.6 Restore the Configuration Data from File

Use the command in Table 68 to read the configuration from the file and load it on the selected controller(s). You can restore the read/write properties and RAID configuration using hot spares.

Table 68: Restore Configuration Data from File

Convention	**MegaCli -CfgRestore -f FileName -aN**
Description	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: • 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 will fail if the same device IDs of the drives are not present.

5.13.7 Manage Foreign Configuration Information

Use the command in Table 69 to manage configurations from other controllers, called *foreign configurations*, for the selected controller(s). You can scan, preview, import, and clear foreign configurations.

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

Table 69: Manage Foreign Configuration Information

Convention	**MegaCli -CfgForeign -Scan | [-SecurityKey sssssssssss] | -Dsply [x] | [-SecurityKey sssssssssss] | -Preview [x] | [-SecurityKey sssssssssss] | -Import [x] | [-SecurityKey sssssssssss] | -Clear [x] | [-SecurityKey sssssssssss] -aN|-a0,1,2|-aALL**
Description	Manages foreign configurations. The options for this command are: `-Scan`: Scans and displays available foreign configurations. `-SecurityKey`: This is 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 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 Drive(s)

Use the command in Table 70 to delete one, multiple, or all virtual drives on the selected controller(s).

Table 70: Delete Specified Virtual Drives

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

5.13.9 Display the Free Space

Use the command in Table 71 to display the free space that is available to use for configuration on the selected controller(s).

Table 71: Display Free Space

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

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 Table 72 to display virtual drive information for the selected controller(s).

Table 72: Display Virtual Drive Information

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

5.14.2 Change the Virtual Drive Cache and Access Parameters

Use the command in Table 73 to change the cache policy and access policy for the virtual drive(s) on the selected controller(s).

Table 73: 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 Table 74 to display cache and access parameters for the virtual drive(s) on the selected controller(s).

Table 74: 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 drive(s): -Cache: -Cached, Direct: Displays cache policy. -WT (Write through), WB (Write back): Selects write policy. -NORA (No read ahead), RA (Read ahead), ADRA (Adaptive read ahead): Selects read policy. -Access: -RW, -RO, Blocked: Displays access policy. -DskCache: Displays drive cache policy.

5.14.4 Manage Virtual Drives Initialization

Use the command in Table 75 to manage initialization of the virtual drive(s) on the selected controller(s).

Table 75: 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 drive(s) and displays the progress (this is optional). The fast initialization option initializes the first and last 8 Mbyte areas on the virtual drive. The full option allows you to initialize the entire virtual drive. `-Abort`: Aborts the ongoing initialization on the virtual drive(s). `-ShowProg`: Displays the snapshot of the ongoing initialization, if any. `-ProgDsply`: Displays the progress of the ongoing initialization. The routine continues to display the progress until at least one initialization is completed or a key is pressed.

5.14.5 Manage a Consistency Check

Use the command in Table 76 to manage a data consistency check (CC) on the virtual drives for the selected controller(s).

Table 76: Manage Consistency Check

Convention	**MegaCli –LDCC –Start|-Abort|–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 drive(s), then displays the progress (optional) and time remaining. `-Abort`: Aborts an ongoing CC on the virtual drive(s). `-ShowProg`: Displays a snapshot of an ongoing CC. `-ProgDsply`: Displays ongoing CC progress. The progress displays until at least one CC is completed or a key is pressed.

5.14.6 Schedule a Consistency Check

Use the command in Table 77 to schedule a consistency check (CC) on the virtual drives for the selected controller(s). There are options to set the mode, change the CC start time, set the delay time and display of the CC info.

Table 77: Schedule Consistency Check

Convention	**MegaCli -AdpCcSched -Dsbl|-Info|{-ModeConc | -ModeSeq [-ExcludeLD -LN|-L0,1,2] [-SetStartTime yyyymmdd hh] [-SetDelay val] } -aN|-a0,1,2|-aALL**
Description	Schedules check consistency on the virtual drive of the selected adapter. `Dsbl`: Disables a scheduled CC for the given adapter(s). `Info`: Gets information about a scheduled CC for the given adapter(s). `ModeConc`: The scheduled CC on all of the virtual drives runs concurrently for the given adapter(s). `ModeSeq`: The scheduled CC on all of the virtual drives runs sequentially for the given adapter(s) `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. `StartTime`: Sets the next start time. The date is in the format of `yyyymmdd` in decimal digits and followed by a decimal number for the hour between 0 ~ 23 inclusively. This is optional. `SetDelay`: Sets the execution delay between executions for the given adapter(s). This is optional. `Values`: The value is the length of delay in hours. A value of 0 means continuous execution.

5.14.7 Manage a Background Initialization

Use the command in Table 78 to enable, disable, or suspend background initialization (BGI), as well as display initialization progress on the selected controller(s).

Table 78: Manage Background Initialization

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

5.14.8 Perform a Virtual Drive Reconstruction

Use the command in Table 79 to perform a reconstruction of the virtual drive(s) on the selected controller(s).

Table 79: Virtual Drive Reconstruction

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

5.14.9 Display Information about Virtual Drives and Drives

Use the command in Table 80 to display information about the virtual drives and drives for the selected controller(s), such as the number of virtual drives, RAID level, and drive capacity.

Table 80: Display Virtual Drive and Drive Information

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

5.14.10 Display the Number of Virtual Drives

Use the command in Table 81 to display the number of virtual drives attached to the controller.

Table 81: 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.

5.14.11 Clear the LDBBM Table Entries

Use the command in Table 82 to clear the LDBBM table entries.

Table 82: 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 drive(s) on the selected adapter(s).

5.14.12 Display the List of Virtual Drives with Preserved Cache

Use the command in Table 83 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 83: Display the List of Virtual Drives with Preserved Cache

Convention	`MegaCli -GetPreservedCacheList -aN|-a0,1,2|-aALL`
Description	Display the list of virtual drives that have preserved cache.

5.14.13 Discard the Preserved Cache of a Virtual Drive(s)

Use the command in Table 84 to discard the preserved cache of a virtual drive(s).

Table 84: Discard the Preserved Cache of a Virtual Drive(s)

Convention	`MegaCli -DiscardPreservedCache -Lx|-L0,1,2|-Lall -force -aN|-a0,1,2|-aALL`
Description	Discard the preserved cache of the virtual drive(s).

5.14.14 Expand a Virtual Drive

Use the command in Table 85 to expand a virtual drive.

Table 85: Discard the Preserved Cache of a Virtual Drive(s)

Convention	**MegaCli -LdExpansion -pN -dontExpandArray -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL**
Description	Expands the virtual drive within the existing array or beyond the size of the existing array if you replace the drives with larger drives. `-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. `-dontExpandArray`: Expand a virtual drive within the array, even when there is room to expand the array. For example, you have created a 5 GB RAID 1 virtual drive with two 30 Gbyte 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: 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 Table 86 to display information about the drives on the selected controller(s).

Table 86: 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, and inquiry data. For SAS devices, this includes additional information such as the SAS address of the drive. For SAS expanders, this includes additional information such as the number of devices connected to the expander. `-Physdrv[E0:S0,...]`: Specifies the physical drive enclosure and the slots for the drives to provide information about.

5.15.2 Set the Drive State to Online

Use the command in Table 87 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 87: 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 Table 88 to set the state of a drive to *Offline*.
In the offline state, the virtual drive is not available to the RAID controller.

Table 88: 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 Table 89 to change the state of a drive from *Unconfigured-Bad* to *Unconfigured-Good*.

Table 89: 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 Table 90 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 controller(s).

Table 90: 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 to *ready* (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 Table 91 to manage a drive initialization on the selected controller(s).

Table 91: Drive Initialization

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

5.15.7 Rebuild a Drive

Use the command in Table 92 to start or stop a rebuild on a drive and display the rebuild progress. When a drive in a RAID drive group fails, you can rebuild the drive by recreating the data that was stored on the drive before it failed.

Table 92: Rebuild a Drive

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

5.15.8 Locate the Drive(s) and Activate LED

Use the command in Table 93 to locate the drive(s) for the selected controller(s) and activate the drive activity LED.

Table 93: Locate Drive and Activate LED

Convention	**MegaCli -PDLocate -PhysDrv[E0:S0,E1:S1....] -aN|-a0,1,2|-aALL**
Description	Locates the drive(s) for the selected controller(s) 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 Table 94 to mark the configured drive as missing for the selected controller(s).

Table 94: Mark Configured Drive as Missing

Convention	**MegaCli -PDMarkMissing -PhysDrv[E0:S0,E1:S1....] -aN|-a0,1,2|-aALL**
Description	Marks the configured drive as missing for the selected controller(s). `-Physdrv[E0:S0,...]`: Specifies the physical drive enclosure and the slots for the drives.

5.15.10 Display the Drives in Missing Status

Use the command in Table 95 to mark the configured drive as missing for the selected controller(s).

Table 95: Display Drives in Missing Status

Convention	**MegaCli –PDGetMissing -aN|-a0,1,2|-aALL**
Description	Displays the drive(s) in missing status. The format is: No Row Column SizeExpected(MB) 0 x y zzzzzzzz … 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.

5.15.11 Replace the Configured Drives and Start an Automatic Rebuild

Use the command in Table 96 to replace configured drive(s) and start an automatic rebuild of the drive for the selected controller(s).

Table 96: Replace Configured Drive(s) 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 Table 97 to prepare the unconfigured drive(s) for removal from the selected controller(s).

Table 97: Prepare Unconfigured Drive(s) for Removal

Convention	**MegaCli –PDPrpRmv [-Undo] –PhysDrv[E0:S0,E1:S1....] -aN|-a0,1,2|-aALL**
Description	Prepares unconfigured drive(s) 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 Table 98 to display the total number of drives attached to an controller. Drives can be attached directly or through enclosures.

Table 98: 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 Table 99 to display a list of the physical devices connected to the selected controller(s).

Table 99: Display List of Physical Devices Attached to Controller(s)

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

5.15.15 Download Firmware to the Physical Devices

Use the command in Table 100 to download firmware to the physical devices connected to the selected controller(s).

Table 100: Download Firmware to the Physical Devices

Convention	**MegaCli -PdFwDownload [offline] {[-SataBridge] -PhysDrv[E0:S0,E1:S1....]|{EncdevId[devId1,devId2,...]} -f <filename> -aN|-a0,1,2|-aAll**
Description	Flashes the firmware with the file specified at the command line. 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 has to do error checking. `-SataBridge`: Allows you to download the SATA Bridge firmware in online mode. `-Physdrv[E0:S0,...]`: Specifies the physical drive enclosure and the slots for the drives. `-EncdevId[devId1,devId2,...]`: Lists the device IDs of the enclosure. See Section 5.16.1, Display Enclosure Information for more enclosure information.

5.15.16 Configure All Free Drives into a RAID 0, 1, 5, or 6 Configuration for a Specific Controller

Use the command in Table 101 to download firmware to the physical devices connected to the selected controller(s).

Table 101: Configure All Free Drives into a RAID 0, 1, 5 or 6 Configuration for a Specific Controller

Convention	**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**
Description	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. `Rx[E0:S0,...]`: Specifies the RAID level and the physical drive enclosure/slot numbers to construct a disk group. `WT (Write through), WB (Write back)`: Selects the write policy. `NORA (No read ahead), RA (Read ahead), ADRA (Adaptive read ahead)`: Selects the read policy. `Cached, -Direct`: Selects the cache policy. `[{CachedBadBBU|NoCachedBadBBU }]`: Specifies whether to use write cache when the BBU is bad. `szXXXXXXXX`: Specifies the size for the virtual disk, where XXXX 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. If multiple size options are specified, CT will configure the virtual drives in the order of the options entered in the command line. The configuration of a particular virtual drive will fail if the remaining 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. `strpszM`: Specifies the stripe size, where the stripe size values are 8, 16, 32, 64, 128, 256, 512, or 1024 MB. `Hsp[E5:S5,...]`: 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 `-PdHsp` command with the proper subcommands. You can also use this option to create a configuration on the free space available in the virtual drive. `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 disk is already used for configuration. `FDE|CtrlBased`: If controller support security feature, this option enables FDE/controller-based encryption on virtual disk.

5.15.17 Set the Mapping Mode of the Drives to the Selected Controller(s)

Use the command in Table 102 to set the mapping mode of the physical devices connected to the selected controller(s).

Table 102: 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 controller(s). `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 Perform the Copyback Operation on the Selected Drive

Use the command in Table 103 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 103: Perform the Copyback Operation on the Selected Drive

Convention	**`MegaCli -PDCpyBk -Start | -Stop|-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. `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

The commands in this section are used for enclosures.

5.16.1 Display Enclosure Information

Use the command in Table 104 to display enclosure information for selected controller(s).

Table 104: Display Enclosure Information

Convention	**MegaCli –EncInfo -aN|-a0,1,2|-aALL**
Description	Displays information about the enclosure for the selected controller(s).

5.16.2 Display Enclosure Status

Use the command in Table 105 to display the status of the enclosure for selected controller(s).

Table 105: Display Enclosure Status

Convention	**MegaCli –EncStatus -aN|-a0,1,2|-aALL**
Description	Displays the status of the enclosure for the selected controller(s).

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 Table 106 to flash the firmware with the ROM file specified at the command line for the selected controller(s).

Table 106: Flash Firmware with ROM File

Convention	**MegaCli –AdpFwFlash –f filename [-NoSigChk] [-NoVerChk]-aN| -a0,1,2|-aALL**
Description	Flashes the firmware with the ROM file specified at the command line. 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. If you specify `–NoVerChk`, also, the application flashes the controller firmware without checking the version of the firmware image. The version check applies only to the firmware (`APP.ROM`) version. This command also supports the “Mode 0” flash functionality. For Mode 0 flash, the controller number is not valid. There are two possible methods: • Select which controller to flash after the controllers are detected. • Flash the firmware on all present controllers. XML output data is generated by this option.

5.17.2 Flash the Firmware in Mode 0 with the ROM File

Use the command in Table 107 to flash the firmware in Mode 0 with the ROM file specified at the command line for the selected controller(s). This option is for DOS only.

Table 107: 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 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 Table 108 to display the PHY connection information for physical PHY M on the selected controller(s). Each PHY can form one side of the physical link in a connection with a PHY on a different device. The physical link contains four wires that form two differential signal pairs. One differential pair transmits signals, and the other differential pair receives signals. Both differential pairs operate simultaneously and allow concurrent data transmission in both the receive and the transmit directions. PHYs are contained within ports.

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

Table 108: Display PHY Connection Information

Convention	**MegaCli -PHYInfo -phyM -aN|-a0,1,2|-aALL**
Description	Displays PHY connection information for physical PHY M on the controller(s).

5.19 Diagnostic-Related Options

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

5.19.1 Start Controller Diagnostics

Use the command in Table 109 to start the controller diagnostic for a set amount of time.

Table 109: 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 Start Battery Test

Use the command in Table 110 to start the battery test. This command requires a system reboot.

Table 110: 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 and 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 Recovery to take snapshots of a volume at designated point in time and restore the volume or files from those points in case data is deleted, whether accidentally or maliciously. MegaRAID Recovery supports up to eight snapshots of PiTs for each volume.

5.20.1 Enable the Snapshot Feature

Use the command in Table 111 to enable the snapshot feature on a selected virtual drive..

Table 111: 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 AutoDeletOldestSnapshot 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 Table 112 to enable the snapshot feature on a selected virtual drive.

Table 112: Disable the Snapshot Feature

Convention	**MegaCli -Snapshot -Dsbl -Lx -aN|-a0,1,2|-aALL**
Description	Command disables the snapshot on the source virtual drive. `-Lx`: x specifies the Source LD number on which to disable snapshot.

5.20.3 Take Snapshot of Volume

Use the command in Table 113 to take a snapshot of a volume at designated point in time .

Table 113: 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 \n", appNameP);**
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 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 Parameter, 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 Table 114 to set the snapshot properties..

Table 114: 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 Table 115 to delete a snapshot.

Table 115: Delete a Snapshot

Convention	**MegaCli -Snapshot -DeleteSnapshot [SnapshotTime yyyymmdd hh:mm:ss | -all] [-force|-y] -LN| -L0,1,2 -aN|-a0,1,2|-aALL**
Description	Deletes the snapshot and the associated view if -Force or -Y is specified. `-SnapshotTime yyyymmdd hh:mm:ss`: (Optional) If used, this action deletes the snapshot with the time stamp that is specified in command line, if it is the oldest PIT. `-force`: If specified, this action deletes the snapshot even if it has the view associated with it. `-y`: If specified, this action deletes the snapshot even if it has the view associated with it. `-LN`: N specifies the source LD number for the command. `-L0,1,2`: Specifies the command is for LDs 0, 1, and 2. You can select more than one LD.

5.20.6 Create a View

Use the command in Table 116 to create a view. A view contains the content from the Point-in-Time [PiT] when the snapshot was made..

Table 116: Create a View

Convention	**MegaCli -Snapshot -CreateView -SnapshotTime yyyymmdd hh:mm:ss [-viewName NameString] [-RW|RO|Blocked] [-szXXX] -Lx -aN|-a0,1,2|-aALL \n", appNameP);**
Description	Creates the view on a particular snapshot. `-SnapshotTime yyyymmdd hh:mm:ss`: Creates the view on the snapshot with the time stamp yyyymmdd hh:mm:ss `-viewName NameString`: (Optional) Specifies the name of the view. `-RW|RO|Blocked`: (Optional) Specifies the access policy of the view. `-szXXX`: (Optional) Specifies the size of the view in MB where `XXX` is a decimal number. `-Lx`: x specifies the source LD number for the command.

5.20.7 Delete a View

Use the command in Table 117 to a view..

Table 117: Delete a View

Convention	**MegaCli -Snapshot -DeleteView [-SnapshotTime yyyymmdd hh:mm:ss] -Lx -aN|-a0,1,2|-aALL**
Description	Deletes the view. `-SnapshotTime yyyymmdd hh:mm:ss`: (Optional) If specified, this action deletes the view on the snapshot with the time stamp `yyyymmdd hh:mm:ss`. `-Lx`: x specifies the source LD number for the command.

5.20.8 Rollback to an Old Snapshot

Use the command in Table 118 to roll the virtual drive back to an older snapshot..

Table 118: Rollback to an Old Snapshot

Convention	**MegaCli -Snapshot -Rollback -SnapshotTime yyyymmdd hh:mm:ss [-Force|-Y] -Lx -aN|-a0,1,2|-aALL**
Description	Rolls back the virtual drive to an old snapshot. `-SnapshotTime yyyymmdd hh:mm:ss`: Specifies the snapshot with the time stamp `yyyymmdd hh:mm:ss` to which it has to roll back. `-Force`: If specified, this action overrides the warning message and causes a rollback to an older snapshot. `-Y`: If specified, this action overrides the warning message and causes a rollback to an older snapshot. `-Lx`: x specifies the source LD number for the command.

5.20.9 Display Snapshot and View Information

Use the command in Table 119 to display information about the snapshot and the view.
.

Table 119: Display Snapshot and View Information

Convention	**MegaCli -Snapshot -Info [-SnapshotTime yyyymmdd hh:mm:ss | -ViewTime yyyymmdd hh:mm:ss] -Lx -aN|-a0,1,2|-aALL**
Description	Displays snapshot and view information for the source virtual drive. If the virtual drive is a repository virtual drive, it displays the LD info, the number of source virtual drives mapped and their target ID andthe number of holes. `-SnapshotTime yyyymmdd hh:mm:ss`: (Optional) If specified, this displays the snapshot information for the snapshot with the time stamp `yyyymmdd hh:mm:ss`. `-ViewTime yyyymmdd hh:mm:ss`: (Optional) If specified, this displays the view information for the view with the time stamp `yyyymmdd hh:mm:ss` and the associated snapshot information. `-Lx`: x specifies the source LD number for the command.

5.20.10 Clean the Recoverable Free Space on the Drives in a Virtual Drive

Use the command in Table 120 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 120: Clean the Recoverable Free Space on the Drives in a Virtual Drive

Convention	**MegaCli -Snapshot -Clean -Lx -aN|-a0,1,2|-aALL**
Description	Cleans the recoverable free space on the drives in a snapshot repository virtual drive. `-Lx`: x specifies the LD number for the command. The LD has to be a repository virtual drive.

5.20.11 Display the Information for a Specific View

Use the command in Table 121 to display the information for a specific view if you specify the view target ID..

Table 121: 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 View target ID. 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 Enabling the Snapshot Scheduler

The snapshot scheduler in MSM helps you automate the creation of Point-in-Time (PiT) on one or more VDs. You can schedule the snapshot as monthly, weekly, daily, or hourly. The scheduler does not support biweekly, alternate days, etc.

NOTE: MSM or OEM applications must be running to flush file system buffers and take snapshot when snapshot schedule expires. The system does not support snapshot creation through BMC (sideband) and snapshot on volumes, which are used for VM creation under virtualized environment. To flush the file system buffers, you should interface with hypervisor.

5.21 FastPath-related Options

The command in this section is used to display information about the FastPath option.

MegaRAID FastPath is a high-performance IO accelerator for CacheCade software drive groups connected to a MegaRAID controller card. CacheCade software has a read performance advantage over HDDs and use 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 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 utilizing FastPath technology. Also, FastPath software is faster and more cost-effective than current flash-based adapter card solutions.

5.22 Miscellaneous Options

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

5.22.1 Display the MegaCLI Version

Use the command in Table 122 to display the version number of the MegaCLI utility.

Table 122: Display MegaCLI Version

Convention	`MegaCli -v`
Description	Displays the version number of the MegaCLI utility.

5.22.2 Display Help for MegaCLI

Use the command in Table 123 to display help information for the MegaCLI utility.

Table 123: Display Help for MegaCLI

Convention	`MegaCli -h|-Help|?`
Description	Displays help for the MegaCLI utility.

5.22.3 Summary Information

Use the command in Table 124 to display help information for the MegaCLI utility.

Table 124: Display Help for MegaCLI

Convention	`MegaCli -ShowSummary [-f filename] -aN`
Description	Displays a summary of system information, controller information, drive information, virtual drive information, and enclosure information.

Chapter 6

MegaRAID Storage Manager Overview and Installation

This chapter provides a brief overview of the MegaRAID® Storage Manager™ (MSM) software and explains how to install it on the supported operating systems.

6.1 Overview

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

MegaRAID Storage Manager software enables you to easily configure the controllers, drives, and virtual drives on your workstation or 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 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.

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

MegaRAID Storage Manager software displays the status of controllers, virtual drives, and drives on the workstation or server that you are monitoring. System errors and events are recorded in an event log file and are displayed on the screen. Special device icons appear on the screen to notify you of drive failures and other events that require immediate attention.

6.1.3 Maintaining Storage Configurations

You can use 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 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)
- Drive with at least 400 MB available free space
- Solaris 10 SPARC occupies minimum of 630MB

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, and 5.0
- Installing MegaRAID Storage Manager Software for Solaris 10 X86"
- SUSE Linux/SLES 9, 10, and 11, with latest updates and service packs
- VMWare ESX 4.0i and ESX 4.1i (also known as COSLess, it is an embedded version of VMWare that does not have a console to do configuration)
- Solaris SPARC

Refer to your server documentation and to the operating system documentation for more information on hardware and operating system requirements.

6.3 Prerequisites to Running MSM Remote Administration

MSM 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 MSM 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 MSM 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 MSM. If MSM is already installed, restart MSM framework.

6.4 Installing MegaRAID Storage Manager

This section explains how to install (or reinstall) MegaRAID Storage Manager software on your workstation or server for the supported operating systems: Microsoft Windows, Red Hat Linux, and SUSE Linux.

6.4.1 Prerequisite for MSM Installation

The MSM installation script also installs the LSI SNMP agent RPM (Red Hat Package Manager). 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 MSM.

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 MSM.

6.4.2 Installing MegaRAID Storage Manager Software on Microsoft Windows

Follow these steps if you need to install MegaRAID Storage Manager software on a system running Microsoft Windows Server 2003, Microsoft Windows XP, Microsoft Windows Vista, or Microsoft Windows 7.

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. When the Welcome screen appears, click **Next**.

   If MegaRAID Storage Manager software is already installed on this system, then an upgraded installation occurs.

3. Read the screen text and select **Modify**, **Repair**, or **Remove**.
4. When the next screen appears, read and accept the user license, and click **Next**.

   The Customer Information screen appears, as shown in Figure 118.

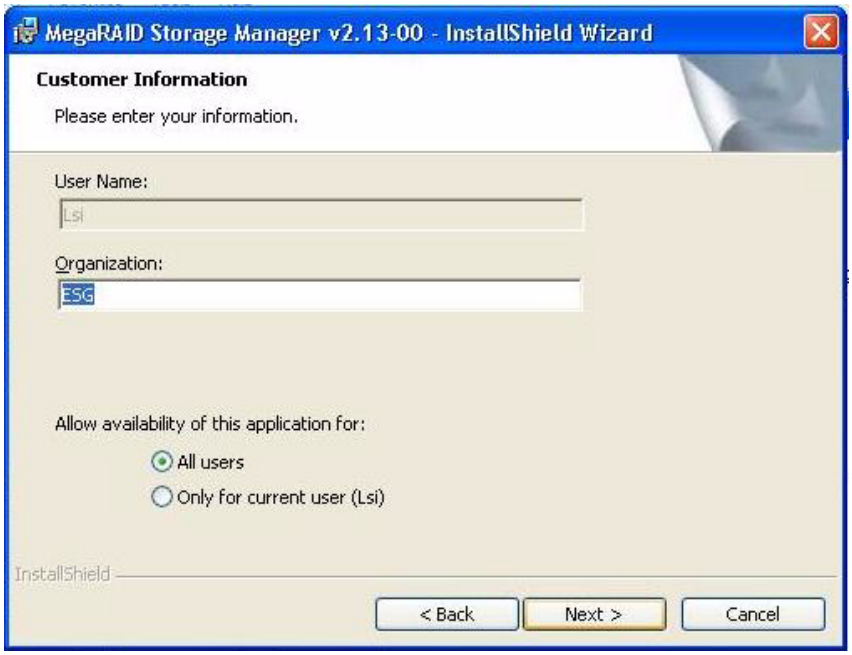


Figure 118: Customer Information Screen

5. Enter your user name and organization name. In the bottom part of the screen, select an installation option:
   - If you select **All users**, any user with administrative privileges can use this version of MegaRAID Storage Manager software to view or change storage configurations.
   - If you select **Only for current user**, the MegaRAID Storage Manager shortcuts and associated icons will be available only to the user with this user name.

6. Click **Next** to continue.
7. On the next screen, accept the default Destination Folder, or click **Change** to select a different destination folder. Click **Next** to continue.

   The Setup Type screen appears, as shown in Figure 119.



Figure 119: Setup Type Screen

8. Select one of the Setup options. The options are fully explained in the screen text.
   - Normally, you would select **Complete** if you are installing MegaRAID Storage Manager software on a server.
   - Select **Custom Installation** 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 screen appears, as shown in Figure 120.

   If you select **Complete** as your setup option, the Installation Wizard is ready to install MSM. To begin installation, click on **Install** on the next screen that appears.

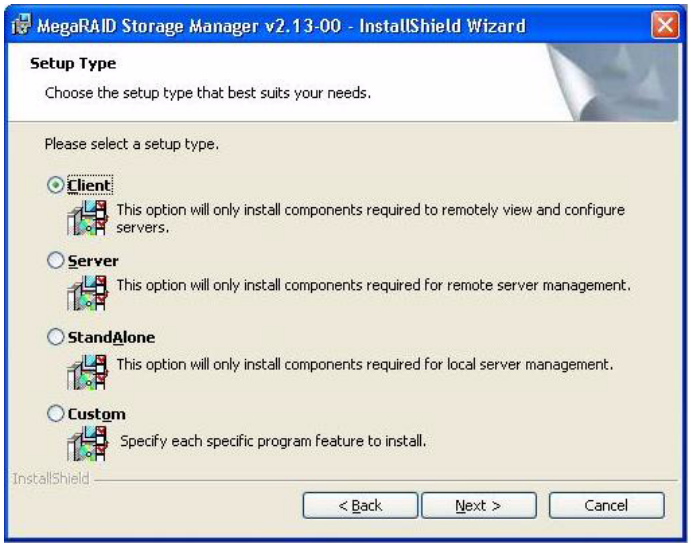


Figure 120: Setup Type Screen

10. Select one of the custom setup options. The options are fully explained in the screen text.
    - Select **Client** if you are installing MegaRAID Storage Manager software on a PC that will be used to view and configure servers over a network. To begin installation, click on **Install** on the next screen that appears.

      In the Client mode of installation, MSM installs only client-related components, such as MSM GUI, and monitor configurator.
      Use this mode when you want to manage and monitor servers remotely. When you install MSM in Client mode on a laptop or a desktop, you can log in to a specific server by providing the IP address.
    - Select **Server** to install only those components required for remote server management. To begin installation, click on **Install** on the next screen that appears.
    - Select **StandAlone** if you will use MegaRAID Storage Manager software to create and manage storage configurations on a standalone workstation. To begin installation, click on **Install** on the next screen that appears.
    - Select **Custom** 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 121. Select the features you want on this screen.

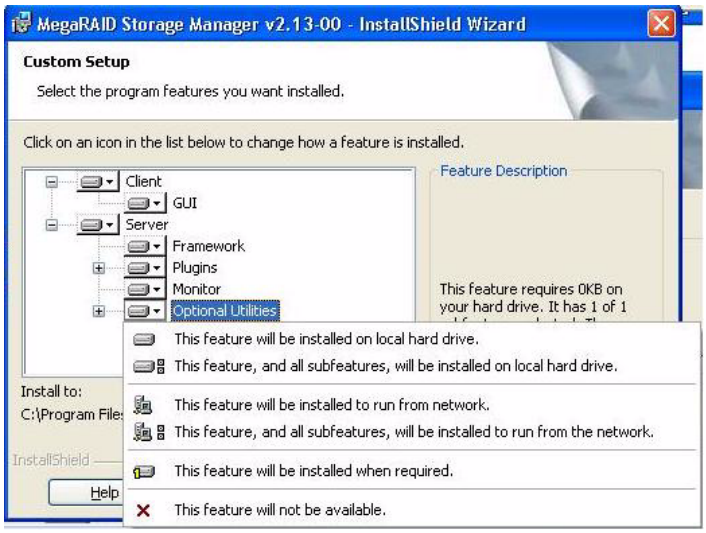


Figure 121: Custom Setup Screen

11. Click **Next** to proceed.
12. Click **Install** to install the program.
13. When the final Configuration Wizard screen appears, click **Finish**.

If you select **Client** installation for a PC used to monitor servers, and if there are no available servers with a registered framework on the local subnet (that is, servers with a complete installation of MegaRAID Storage Manager software), the server screen will appear, as shown in Figure 122. The server screen will not list any servers. You can use this screen to manage systems remotely.

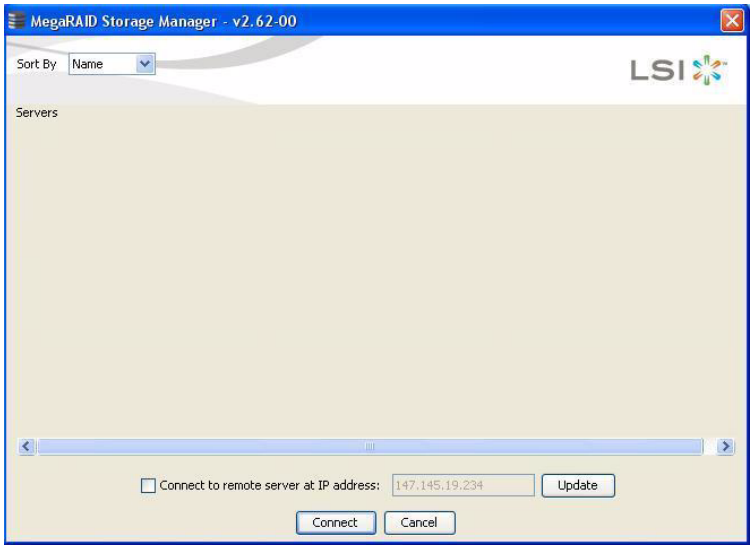


Figure 122: Server Screen

6.4.3 Installing MegaRAID Storage Manager for SPARC

Perform the following steps to install 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, select Y to complete the installation.=

6.4.4 Installing MegaRAID Storage Manager Software for Linux

Follow these steps if you need to install MegaRAID Storage Manager software on a system running Red Hat Linux or SUSE Linux:

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

   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:

   csh install.sh -a

If you select **Client** installation for a PC used to monitor servers, and if there are no available servers with a registered framework on the local subnet (that is, servers with a complete installation of MegaRAID Storage Manager software), the server screen appears. The server screen does not list any servers. You can use this screen to manage systems remotely.

To install the software using 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:

- Complete installation
- Client Component Only
- StandAlone

The `-ru popup` command will remove popup from installation list.

You also can run non-interactive installation using the `RunRPM.sh` command.

The installer offers three types of setup options:

- Complete - This installs all the features of the product.

- Client Components Only - The storelib feature of the product are 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 is not installed in this case, so the resident system is not a part of the network. This means the system cannot browse any other MSM servers in the subnet, and the MSM servers cannot will recognize it as a server.

This installation helps you select any of the setup types, but if you run `RunRPM.sh`, it installs the complete feature.

6.4.5 Prerequisites for Installing MSM on RHEL6.0 x64 Operating System

Before installing MSM on RHEL6.0 x64 system, install the following rpms. Without these files/RPM's MSM may not be install properly or may not work as expected.

1. libstdc++.so.6 -> libstdc++-xx.yy.i686.
2. libXext.so.6 -> libXext-1.1-2.el6.i686.rpm
3. libXtst.so.6 -> libXtst-1.0.99.2-3.el6.i686.rpm

These rpms are available as a part of RHEL6 OS DVD. These RPM's may need additional dependent RPM's as well, and also all the dependent RPM's must be installed on the target system.

6.4.6 Linux Error Messages

The following messages can appear while you are installing MegaRAID Storage Manager software on a Linux 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 MegaRAID Storage Manager software. (This 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 MegaRAID Storage Manager software with the procedure listed previously.

- The version is already installed.

  This message indicates that the version of 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 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.

6.4.7 MSM Customization

You can customize your Logo and Splash screen by by editing the 'msm.properties' file present in the *<installation-directory\MegaRAID Storage Manager>* folder.

In the file 'msm.properties' there are no values for the following keys.

a. CHANNELLOGO=
b. CHANNELSPLASHSCREEN=

There are no default values assigned for the keys mentioned above, therefore the MSM uses the default LSI Logo and Splash screen.

To customize the Logo and Splash screen, you have to enter the Logo and Splash screen file name against these entries.

To enter the filenames follow these steps:

1. Open the file 'msm.properties' 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 the important points.

- File names for both entries should not have any spaces. For example, the valid file name would be: logo_test_1.png or LogoTest1.png or TEST_SPLASH_FILE.jpg
- The logo image dimensions should not exceed 160 * 85 (width * height).
- The splash screen image dimensions should not exceed 390 * 260 (width * height).

After making the changes mentioned above, when you log into MSM, you will be able to view the changes with the new splash screen and logo in MSM.

6.5 MegaRAID Storage Manager Support and Installation on VMWare

This section documents the installation of MegaRAID Storage Manager on VMWare Classic (with console operating system) and on the VMWare ESX 3i operating system.

6.5.1 Installing MegaRAID Storage Manager for VMWare Classic

VMWare does not support any graphics components. In order to install MSM on the VMWare operating system, execute 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 4.0)
- Select the Storelib (Inbox Storelib or Storelib from MSM package)

6.5.2 Uninstalling MegaRAID Storage Manager for VMWare

To uninstall the Server Component of MSM on VMWare, use the `Uninstall` command in the Program menu or run the script `/usr/local/MegaRAID Storage Manager/uninstaller.sh`.

Note the following points:

- A MSM upgrade is supported in this release. This release can be upgraded by future releases.
- To shut down the MSM Framework service, run the following command:

  `/etc/init.d/vivaldiframeworkd stop`

  It is recommended that you stop the Monitor service before you stop the MSM Framework service. To stop the Monitor service run the following command:

  `/etc/init.d/mrmonitor stop`

6.5.3 MegaRAID Storage Manager Support on the VMWare ESXi Operating System

This section outlines the product requirements needed to support the VMWare ESX operating system. Classic VMWare includes a Service Console that is derived from the Linux 2.4 kernel, but with reduced functionality.

The MSM server part cannot be installed directly in VMWare ESX 3i. Management is possible only through Common Information Model (CIM) providers. These CIM providers integrated into the ESX 3i system build an interface between the hardware driver of the LSI MegaRAID controller and remote applications, such as MSM. Management is performed through MSM installed on a remote machine (Linux/Windows). See Section 6.5.4.1, *VMWare ESX 3i Management through CIM and CMPI* for more information.

The Linux installer of MSM works under console with minimal changes. Hardware RAID is currently supported in ESX 3.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.

The network communication is a key element for a proper setup. The communication between the ESXi CIM provider and the LSI management software is an active/passive combination, which requires a highly reliable network. Therefore, we recommend that you install the management on a VM within the ESXi. Follow these steps to install and configure MSM support on the VMWare ESX operating system:

1. **Network Configuration of the ESXi Host:**
   - — Assignment of a ESXi hostname.

     Even if it is not relevant for your network, you need a FQDN (Fully Qualified Domain Name).

     Example: local.lsi.com to be entered using the local ESXi console
   - — Configuration of a virtual network environment:

     You can use the already existing Vswitch, which has a VMkernel port already attached for the communication.

     Alternatively, you can build a new Vswitch without a link to the Host network card.

Which one of the two possibilities to choose depends on your application. It is recommended to choose between both possibilities at a early stage, because the creation of a new Vswitch with VMkernel requires a reboot to make sure a proper communication between the CIM provider and the new interface. For those who want to reach the target as quickly as possible, no change is recommended.

— Configuration of the IP address:

Configure the IP address. The address must be accessible by the VM that will be installed next.

2. **VM Installation:**

Install the operating system as usual, including the VMWare guest tools. The virtual network card should be linked to a Vswitch that has a VMKernel port attached. For a quick installation, no change is recommended.

3. **VM Network Configuration:**

— Case 1: Your network contains a DNS server:

Configure a host entry that belongs to your internal zone and make sure that the FQDN of the ESXi server can be resolved. (Example: local.lsi.com and 192.19.221.186)

— Case 2: Your network does not have a DNS server:

Edit your file `C.\windows\system32\drivers\etc\hosts` and add another entry:

IP of the ESXi Host	FQDN of the ESXi Host
`192.19.221.186`	`local.lsi.com`

4. **Final Steps:**

Reboot the VM and start MSM from a remote Windows/ Linux machine that has MSM installed in complete mode. The ESXi server should now appear in the list of the found hosts. You can now log in with the root account name and password of the ESXi Host.

Host Overview:

Figure 123 lists the host ESXi server and other servers.

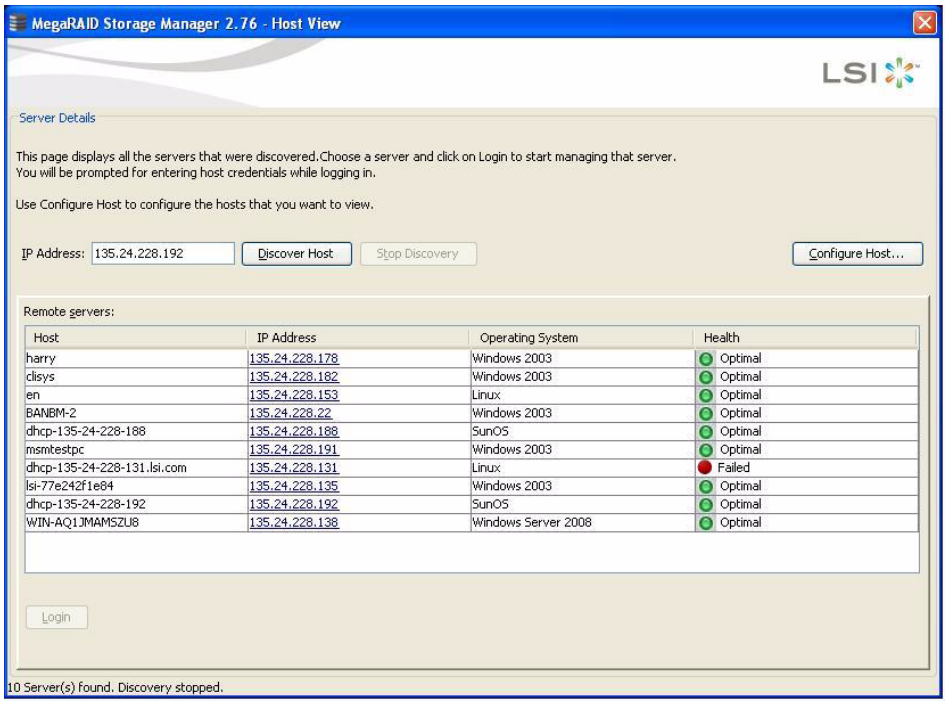


Figure 123: Host ESXi Server Name

You can click Configure Host if you want to select the display preferences for the server. You can choose to display only the local server, systems from a list, or display all of the systems in the network of the local server.

6.5.4 Limitations

The following are the limitations of this installation and configuration:

- There is no active event notification, for example, by popup or email.
- There is no status information for the controller
- Events are collected as long as MSM runs on the Client.
- MSM responds more slowly.

For more details on these limitations, see Section 6.5.4.2, *Differences in MSM for VMware ESXi*.

6.5.4.1 VMWare ESX 3i Management through CIM and CMPI

Management of VMWare ESX 3i is possible only through a Common Information Model (CIM) provider. It is not possible to install anything on the VMWare ESX3i system, so management is performed through MSM installed on a remote machine (Linux/Windows).

VMWare ESX 3i comes with the Small Footprint CIM Broker (CFCB) 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 Manageablity 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).

6.5.4.2 Differences in MSM for VMware ESXi

The following are some of the differences in the MSM 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 display.
  - The operating system type and the operating system architecture do not appear.
  - There is no support for the controller health information.
- Authentication support:
  - MSM 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 MSM Client connected to a VMware ESXi system behaves as follows:

  - There is no support for retrieving initial logs (the events that occurred before a client logs in). Only those events that occur after a client logs in appear in the event logger dialog.
  - System log does not display.
  - The “Save log” feature is not supported; however, the “Save Log as Text” is still supported.
  - The “View Log” option allows you to view the logs saved in a text file on the event logger dialog.
  - The event descriptions might not be identical to a normal MSM Client because the descriptions come from the firmware through the provider.
  - There is no filtering of events, unlike Monitor Service.
  - Refreshing of the MSM GUI after any updates on the firmware is slower for a client connected to VMWare ESXi hosts, compared to one connected to Windows/Linux/Solaris hosts.
- Remote discovery:
  - For networks that do not have DNS configured, the “hosts” file in the machine on which MSM is installed must be edited as follows:

    Add an entry to map the VMWare host’s IP address with the hostname. This is for the discovery to happen correctly. In the absence of this entry, the VMWare host would be discovered as `0.0.0.0`.

    Add an entry to map its own IP address (not the loop back address) with the Hostname. This is to ensure that the Alert Event Notifications (AENs) are delivered correctly.
  - For networks that has DNS configured, the “hosts” file in the machine on which MSM is installed must be edited as follows:

When you do the initial configurations for the VMWare host, provide the correct DNS server IP address.

In the `hosts` file of the machine on which MSM is installed, 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.

- It takes a while to discover the CIMOM servers. If you start the MSM client immediately after you install MSM (or restart Framework service), you will not be able to discover any hosts in the network.
- This is supported only on a full installation of MSM; standalone, client-only, and server-only modes do not support VMWare ESX3i management.
- Supported on the following operating systems:
  - Windows, RHEL and SLES 10
- The following descriptions summarize the status of components related to VMWare ESX3i:
  - MSM client GUI is supported.
  - There is no support for Monitor Configurator; you cannot configure the severity of the AENs.
  - There is no pop-up service support.
  - There is no email and system log support.
  - Monitor service support is not available.

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 ESX3i system, management is possible only through a CIM provider and it is performed through MSM installed on a remote machine running a Linux or Windows operating system.

VMWare ESX3i 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 Manageablity 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 Linux

The following procedure documents how to install and un-install the LSI CIM SAS Storage Provider on a system running on the Linux operating system.

NOTE: Uninstall all the previous versions of LsiSASProvider before you install this version. You can check all of the installed versions of LsiSASProvider by using the command **`rpm -qa | grep LsiSASProvider`**.

- Perform the following step 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 neccessary files and the Managed Object Format (MOF), and it registers the libraries. The 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.

- Perform the following step to un- install 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 cimmof, cimprovider, and wbemexec, should be in PATH variable of `/etc/profile`, and hence, should be defined in all environments of the system.

For Pegasus version 2.5.x, perform the following steps:

1. After you install the LSI SAS Pegasus provider, verify that `libLsiSASProvider.so` and `libLsiSASProvider.so.1` are in `/usr/lib/Pegasus/providers` directory.

   If these files are not present, copy `libLsiSASProvider.so.1` 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 Pegasus CIM Server and LsiServer by performing the following steps:
   — To start the tog-pegasus server, execute the following command:

   `# /etc/init.d/tog-pegasus restart`

   — To start LsiSASSever, execute 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 un-install 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 Provider is now ready to use.

Perform the following steps to uninstall a CIM SAS Storage Provider on a Windows system.

1. Go to **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.

6.7.1 Prerequisite for LSI SNMP Agent RPM Installation

The LSI SNMP agent application depends upon the standard SNMP Utils package. Please 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 rpm's.

Make sure that these RPM's are installed from the operating system media before you install the LSI SNMP agent RPM.

6.7.2 Installing and Configuring an SNMP Agent on Linux

This section explains how to install and configure SAS SNMP Agent for the SUSE Linux and Red Hat Linux operating systems.

To do this, perform the following steps.

NOTE: This procedure requires that you have 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 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_mrdsnmp/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 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 `lsi_mrdsnmpd.conf`. For reference, a sample conf 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/stop the `snmpd` daemon, enter the following command:

   `/etc/init.d/snmpd start/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/stop`

   You can check the status of the SAS SNMP Agent daemon by checked by issuing 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 execute "`sudo snmptrapd -P -F "%02.2h:%02.2j TRAP%w.%q from %A %v\n"`".

NOTE: It is assumed that `snmpd.conf` is located at `/etc/snmp` for Red Hat and `/etc` for SLES. You can change the file location from `/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 SLES 10, 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**.

6.7.3 Installing and Configuring an SNMP Agent on Solaris

This section explains how to install and configure SAS SNMP Agent for the Solaris operating system.

6.7.3.1 Prerequisites

This package requires that you have Solaris System Management Agent installed on the Solaris machine.

6.7.3.2 Installation SNMP on Solaris

To install SNMP for Solaris, perform the following procedure:

1. Unzip the LSI SAS SNMP Agent package.
2. Run the install script by executing the following command:

```
# ./install.sh
```

NOTE: The installation will exit if there are any existing versions of storelib and sassnmp installed on the Solaris machine. Uninstall the existing version by using the following commands:
pkgrm storelib (to uninstall storelib library)
pkgrm sassnmp (to uninstall LSI SAS SNMP Agent)

6.7.3.3 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.

6.7.3.4 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 Managment 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`

6.7.3.5 Configuring snmpd.conf

By default, SNMP queries (walk, get) can be executed 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 executing 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.

6.7.3.6 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 ipaddress 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
```

6.7.3.7 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
```

6.7.4 Installing an SNMP Agent on Windows

This section explains how to install and configure SAS SNMP Agent for the Windows operating system.

6.7.4.1 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 `%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.

6.7.4.2 Installing SNMP Service for Windows

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 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.

6.7.4.3 Configuring SNMP Service on the Server Side

Perform the following steps to configure SNMP Service on the server side.

1. Select **Administrative Tools** from Control Panel.
2. Select **Services** from 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 IPs to which you want the traps to be sent with the community name.

6.8 MegaRAID Storage Manager Support and Installation on Solaris 10

This section documents the installation of MegaRAID Storage Manager on the Solaris 10 U5,U6,U7,U8 x86 and x64 operating system.

6.8.1 Installing MegaRAID Storage Manager Software for Solaris 10

Follow these steps to install MegaRAID Storage Manager software on a system running Solaris 10, update 5:

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.

   When the service is in “maintenance state” , run “svcadm clear MSMFramework” and “svcadm clear MSMMonitor.”
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 Solaris 10

Follow these steps to uninstall MegaRAID Storage Manager software on a system running Solaris 10, update 5:

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 MSM Framework service, run `svcadm disable -t MSMFramework`. It is advisable to stop Monitor service before stopping MSM 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: To check the status of MSM services execute the command `svcs -a | grep -i msm`.

6.9 Prerequisites to Running MSM Remote Administration

MSM requires ports 3071 and 5571 to be open to function. Follow these steps to prepare to run MSM Remote Administration.

1. Configure the system with valid IP address.

   Make sure there is no IP address conflict with in the sub network.

   Ports such as 3071 and 5571 are open and available for MSM 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 MSM to work); if not, create a static route using the following command:

   ```
   Route add 229.111.112.12 dev eth1
   ```

4. Install MSM. If MSM is already installed, restart MSM framework.

Chapter 7

MegaRAID Storage Manager Windows and Menus

This chapter explains how to start MegaRAID Storage Manager software and describes the MegaRAID Storage Manager windows and menus.

7.1 Starting MegaRAID Storage Manager Software

Follow these steps to start MegaRAID Storage Manager software and view the main window:

1. Start the program using the method required for your operating system environment:
   - To start MegaRAID Storage Manager software on a Microsoft Windows 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 MegaRAID Storage Manager software to start. (The Windows Firewall sometimes blocks the operation of programs that use Java.)

   - To start MegaRAID Storage Manager software on a Red Hat Linux system, select **Applications->System Tools->MegaRAID Storage Manager StartupUI**.
   - To start MegaRAID Storage Manager software on a SUSE Linux/SLES system, select **Start->System->More Programs ->MegaRAID Storage Manager**.

   When the program starts, the Select Server window appears, as shown in Figure 124. The remote servers display, along with their IP address, operating system, and health status.

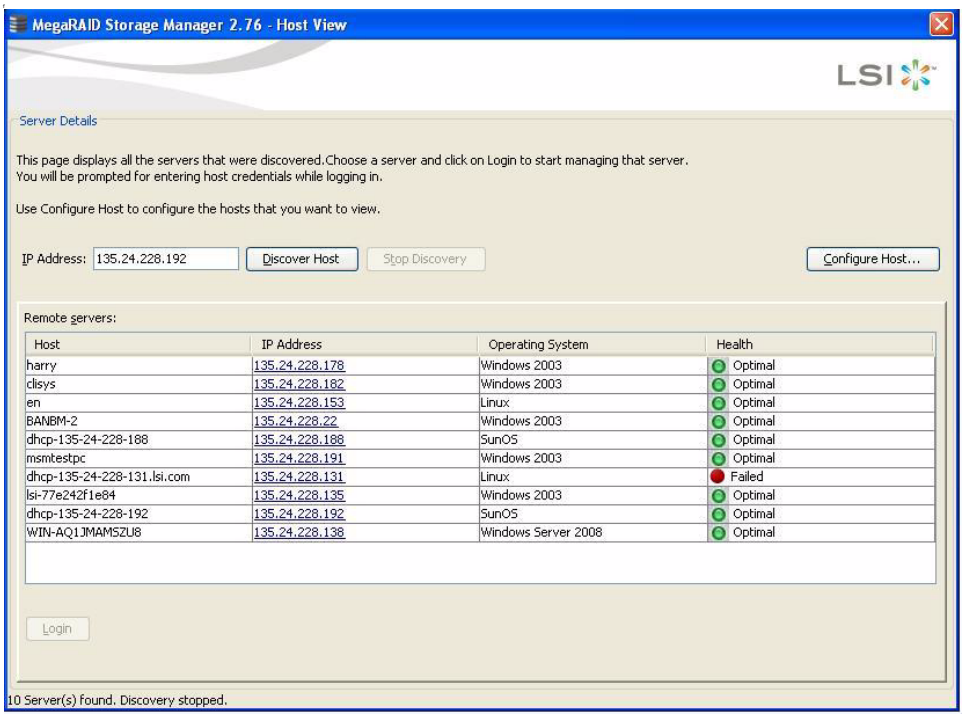


Figure 124: Select Server Window

If the circle in the server icon is orange instead of green, it means that the server is running in a degraded state—for example, because a drive used in a virtual drive has failed. If the circle is red, the storage configuration in the server has failed.

NOTE: To access servers on a different subnet, type in the box at the bottom of the screen the IP address of a server in the desired subnet where the MegaRAID Storage Manager software is running, and click **Update**. If you check the **Connect to remote server at: IP** address box, you can also access a standalone (remote) installation of MegaRAID Storage Manager software, if it has a network connection.

NOTE: For the VMWare CIMOM, the server button does not denote the health of the server. The button 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 Hostname 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 in order 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 Figure 125.



Figure 125: 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.

NOTE: When connected to VMWare 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 is the login to the computer itself, not the network login.

5. Enter the root/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 126.

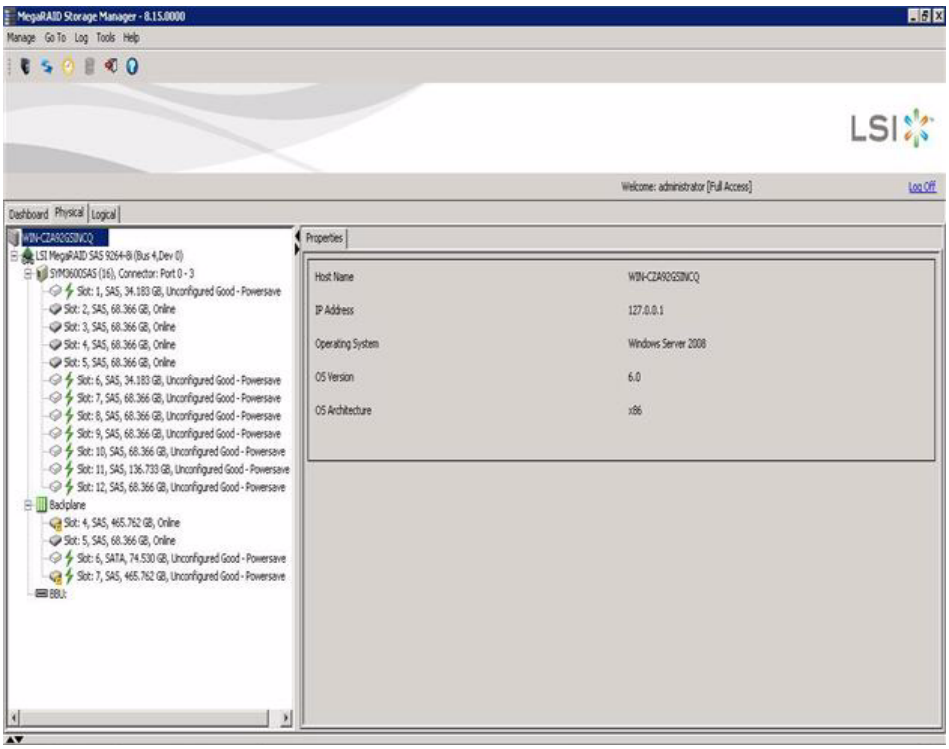


Figure 126: Main MegaRAID Storage Manager Window

The following topics describe the panels and menu options that appear on this screen.

7.2.1 Dashboard/PhysicalView/Logical 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.

7.2.1.1 Dashboard View

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
- MSM 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 127 shows the Dashboard view.

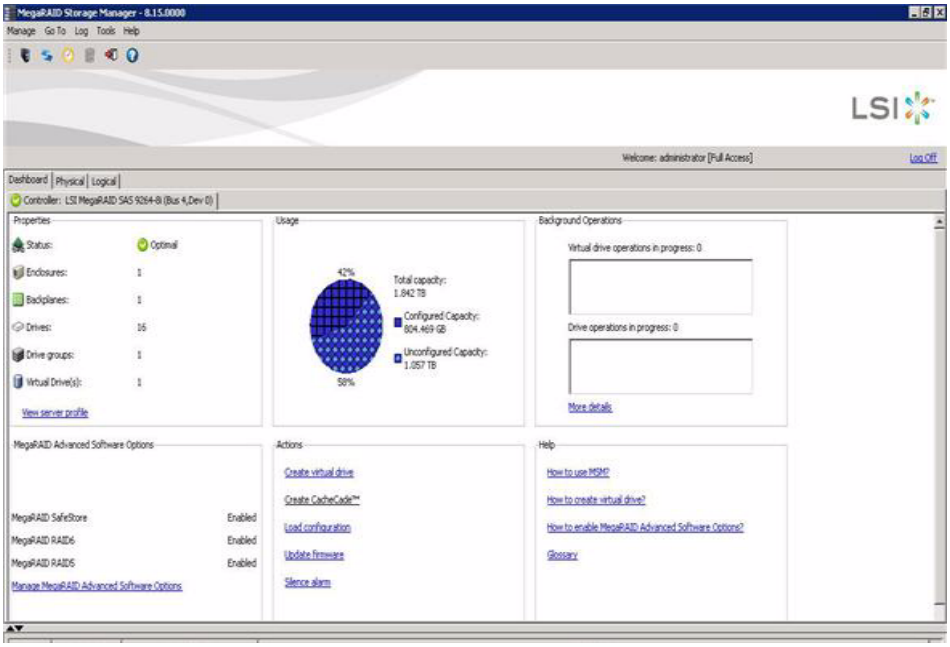


Figure 127: MSM Dashboard View

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 128 shows the Physical view.

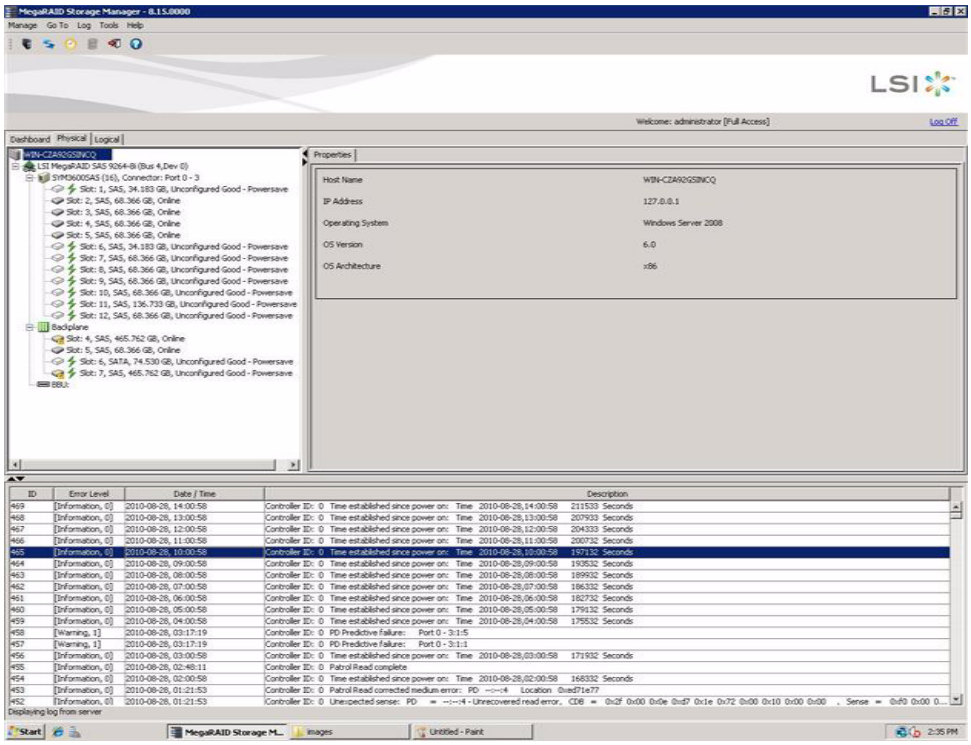


Figure 128: MSM Physical View

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 129 shows the Logical view.

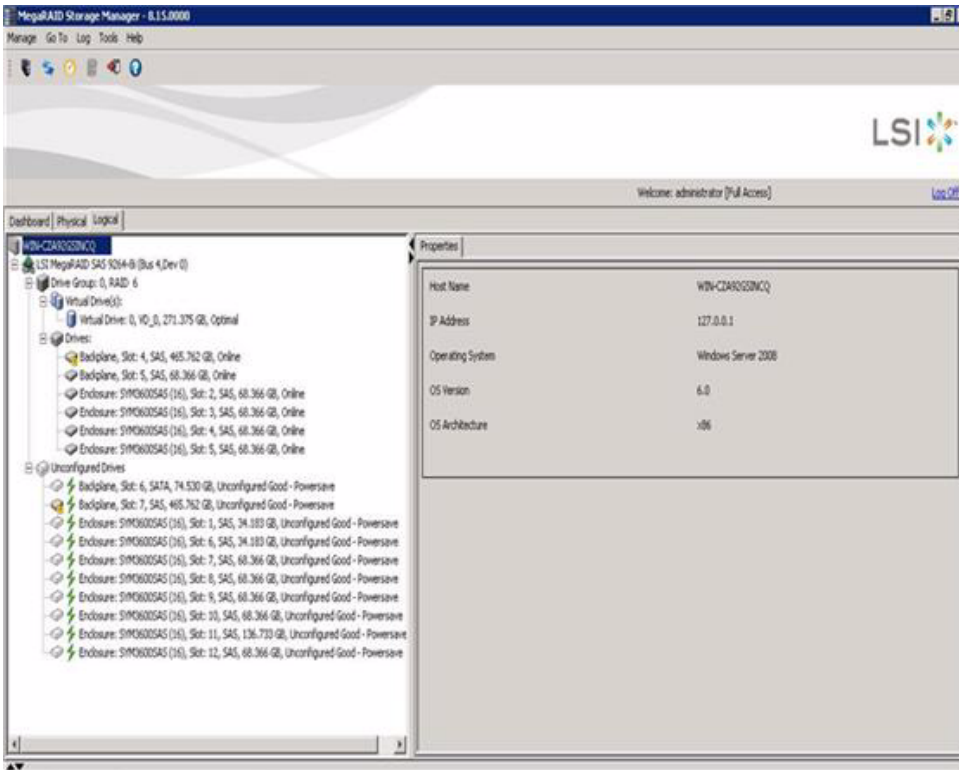


Figure 129: MSM Logical View

The following icons in the left panel represent the controllers, drives, and other devices:

	Status
	System
	Controller
	Backplane
	Enclosure
RAID	Port
	Drive group
	Virtual drive
	Online drive
	Power save mode
	Dedicated hot spare

	Global hot spare
	Battery backup unit (BBU)
	Tape drive
	CD-ROM
	Foreign Drive
	Unconfigured Foreign Drive
	Unconfigured Drive

NOTE: MegaRAID Storage Manager 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 drive has failed:  .

7.2.2 Properties/Graphical View Tabs

The right panel of the MegaRAID Storage Manager window has one tab or two tabs, depending on which kind of device you select in the left panel. Figure 130 shows the MSM main menu.

- 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 9.3, *Monitoring Controllers* Section 9.4, *Monitoring Drives*, and Section 9.6, *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.

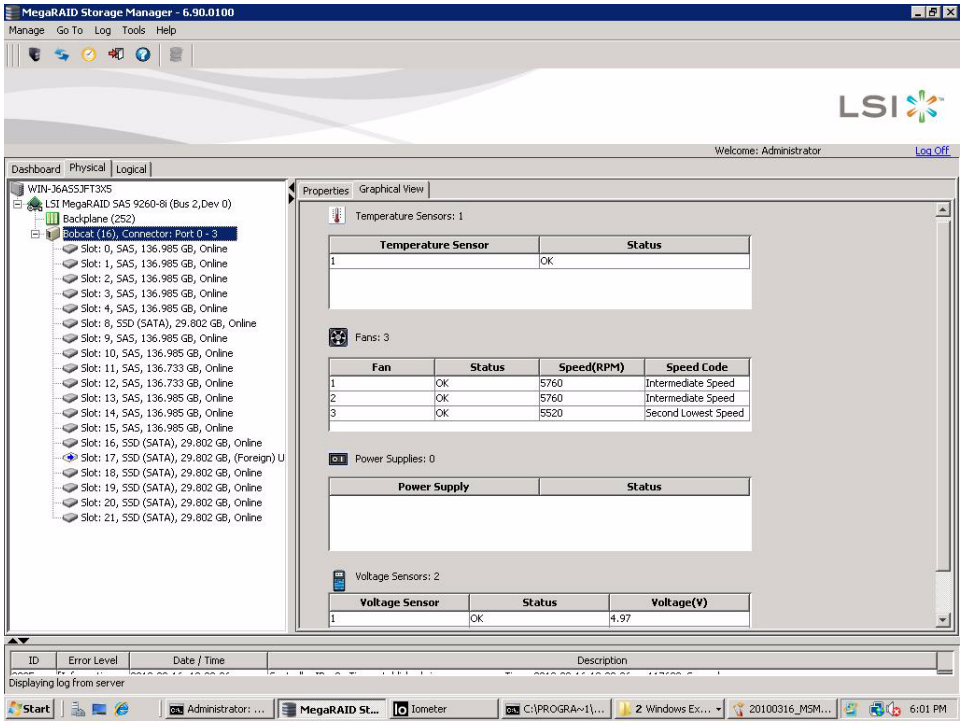


Figure 130: Properties Tab and Graphical View Tab

7.2.3 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 Section 9.1, *Monitoring System Events* For more information about the event log entries, see Appendix A, *Events and Messages*.

7.2.4 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 manual.

7.2.4.1 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 normally updates automatically) and an Exit option to end your session on MegaRAID Storage Manager. The Server menu item 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.

7.2.4.2 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 perform configuration 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**.

7.2.4.3 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*.

7.2.4.4 Tools Menu

On the Tools menu you can select **Tools->Configure Alerts** to access the Configure Alerts screen, which you can use to set the alert delivery rules, event severity levels, exceptions, and email settings. For more information, see Section 9.2, *Configuring Alert Notifications*

7.2.4.5 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® or Mozilla® for the MegaRAID Storage Manager online help to display.

NOTE: When connected to the VMWare server, only the IP address and the hostname information display. The other information, such as the operating system name, version, and architecture do not display.

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 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 explain 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 zeroes 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 can take a long time if the drives are large.

NOTE: 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.

- **Stripe size:** Stripe sizes of 8, 16, 32, 64, 128, 256, 512, 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 speeds up reads for sequential data, but there is little improvement 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 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. It provides a good balance between data protection and performance.

NOTE: The Write Policy depends on the status of the battery backup unit (BBU). If the BBU is not present, is low, is failed, or is being charged, the default Write Policy will switch 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 Cached I/O mode, all reads are buffered in the controller's cache memory or in the CacheCade virtual drive.
  - *Direct IO*: (the default) In Direct I/O mode, the reads are not automatically cached in controller cache and the data is transferred directly to the host except when 1) the data is buffered in controller memory in order to execute parity; or 2) the data is transferred from the controller cache because of read-ahead algorithms.
- **Access policy:** Select the type of data access that is allowed for this virtual drive.
  - *Read/Write*: (the default) Allow read/write access. This is the default.
  - *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.
  - *Disable*d: 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: 64KB

8.1.3 Optimum Controller settings for FastPath

Write Policy: Write Through

IO Policy: Direct IO

Read Policy: No Read Ahead

Stripe Size: 64KB

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 created 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:
   a. Right-click the controller node in the device tree in the left frame of the MegaRAID Storage Manager window and select **Create Virtual Drive**.
   b. Select the controller node and select **Go To>Controller>Create Virtual Drive** in the menu bar, as shown in Figure 131.

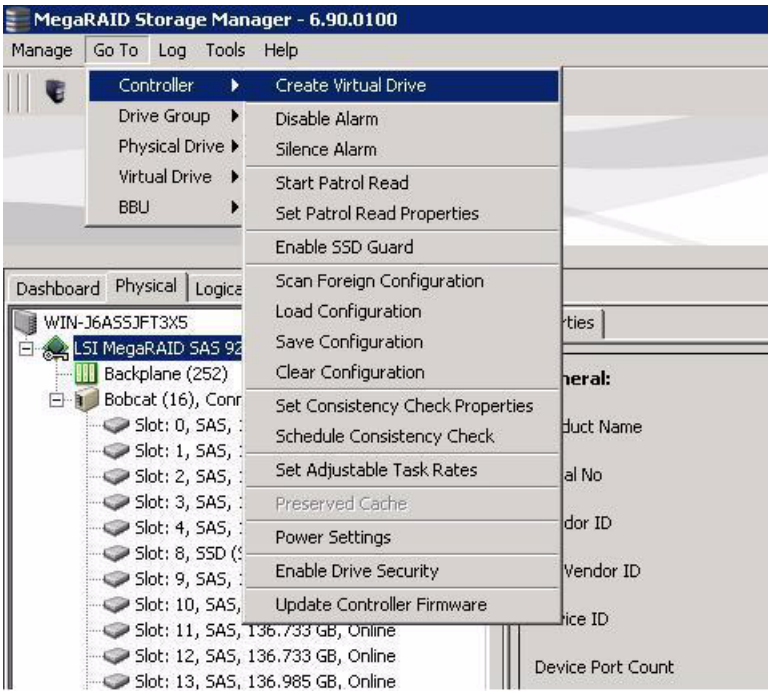


Figure 131: Virtual Drive Creation Menu

The dialog box for the configuration mode (simple or advanced) appears, as shown in Figure 132.

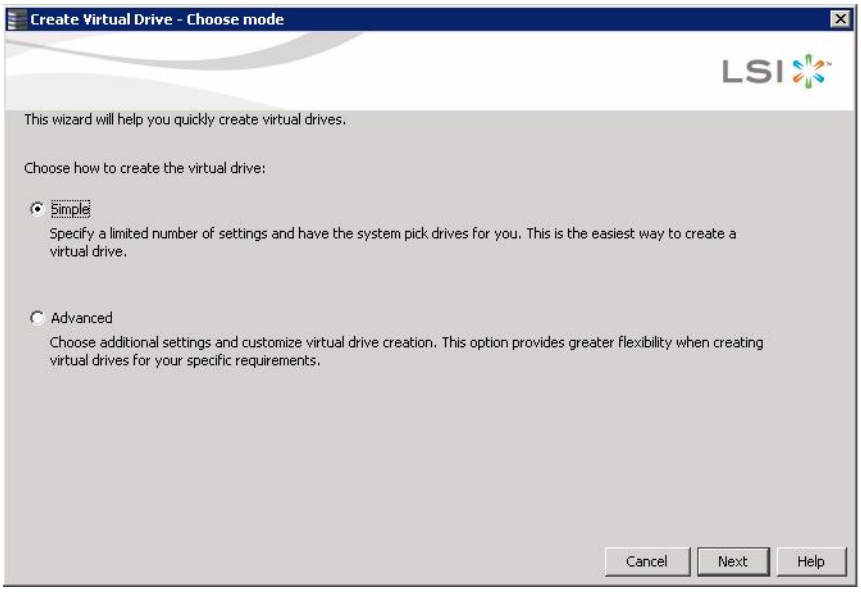


Figure 132: Virtual Drive Simple Configuration Mode

2. Click **Simple** and press **Next**.

If there are no unconfigured drives available, you have the option to use the free capacity of an existing drive group, as shown in Figure 133. The Create Virtual Drive-Summary screen appears as shown in Figure 135.

If unconfigured drives are available, Figure 134 appears, and you can go to step 4.

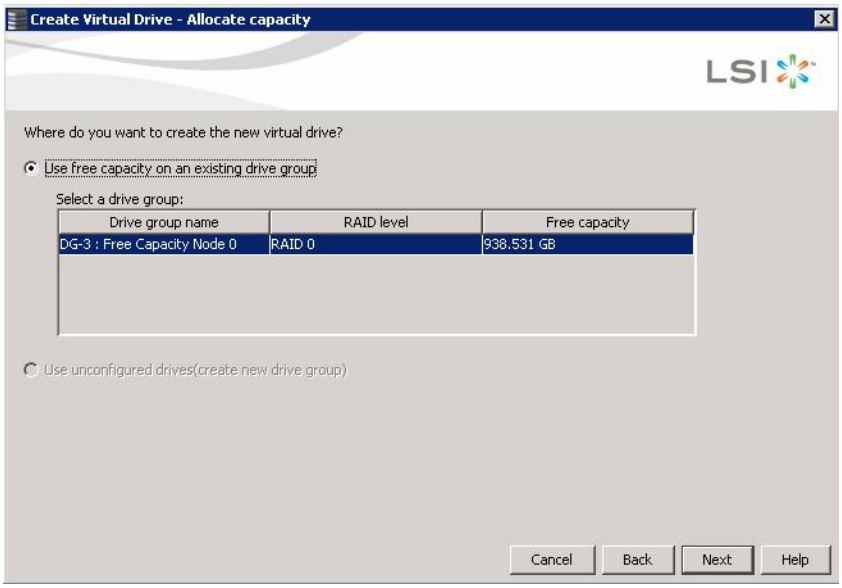


Figure 133: Using the Free Capacity of an Existing Drive Group

3. Check the option **Use Unconfigured drives (create new drive group)** and press **Next**.

   The Create Virtual Drive screen appears, as shown in Figure 134. If there are different types of drives attached to the controller, such as HDD, SDD, SAS, and SATA, there is an option to allow drive type mixing.

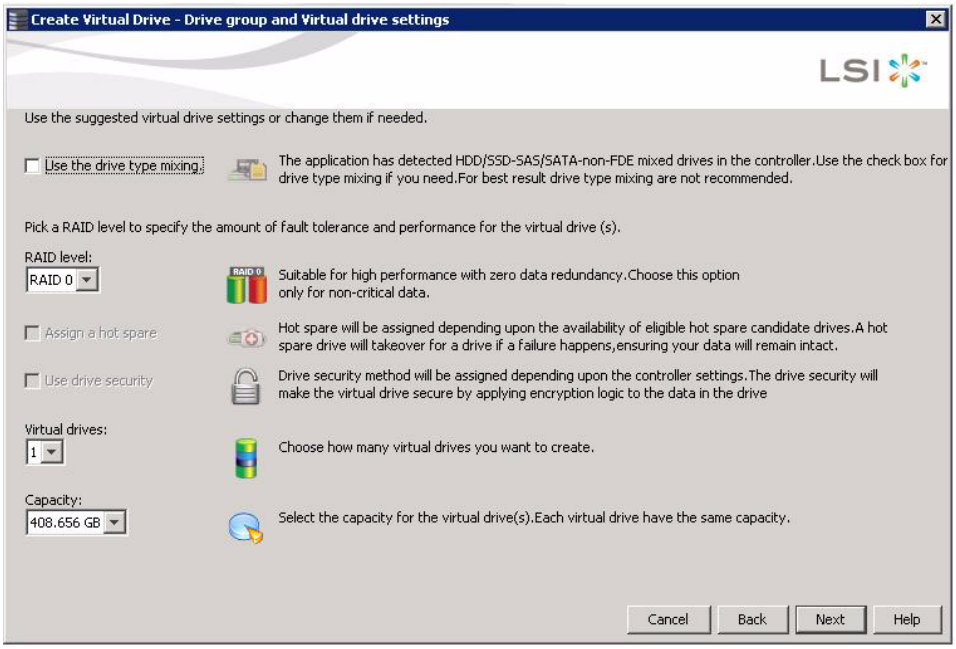


Figure 134: Create Virtual Drive Screen

4. If you want to allow different types of drives in a configuration, click the checkbox **Use the drive type mixing**.

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 screen 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. Click 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. Click 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 menu in the **Virtual drives** field to choose how many virtual drives you want to create.
9. Select the capacity of the virtual drive(s).

Each virtual drive has the same capacity.

10. Click **Next**.

    The **Create Virtual Drive - Summary** window appears, as shown in Figure 135. This window shows the selections you made for simple configuration.

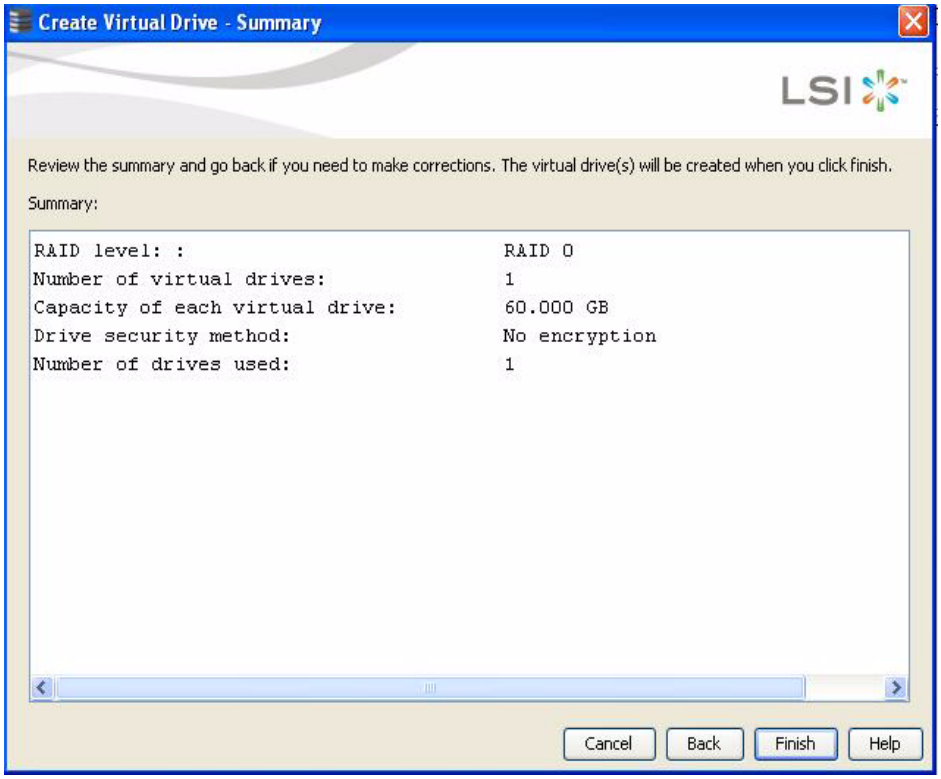


Figure 135: Create Virtual Drive - Summary Window

11. Click **Back** to return to the previous screen 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 powersave 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 box appears that identifies the drive or drives.

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:
   a. Right click on the controller node in the device tree in the left frame of the MegaRAID Storage Manager window and select **Create Virtual Drive.**
   b. Select the controller node and select **Go To->Controller->Create Virtual Drive** in the menu bar, as shown in Figure 136.

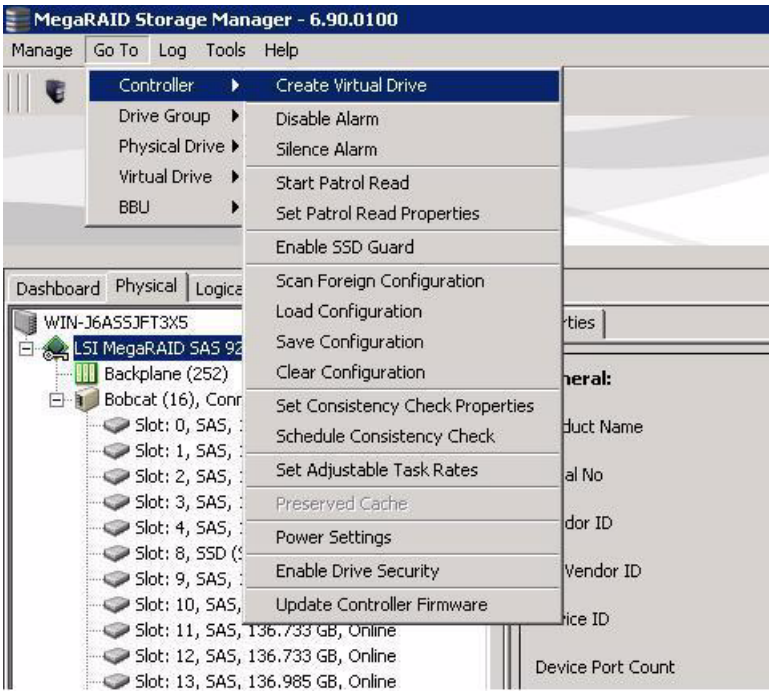


Figure 136: Virtual Drive Creation Menu

The dialog box shown in Figure 137 appears.

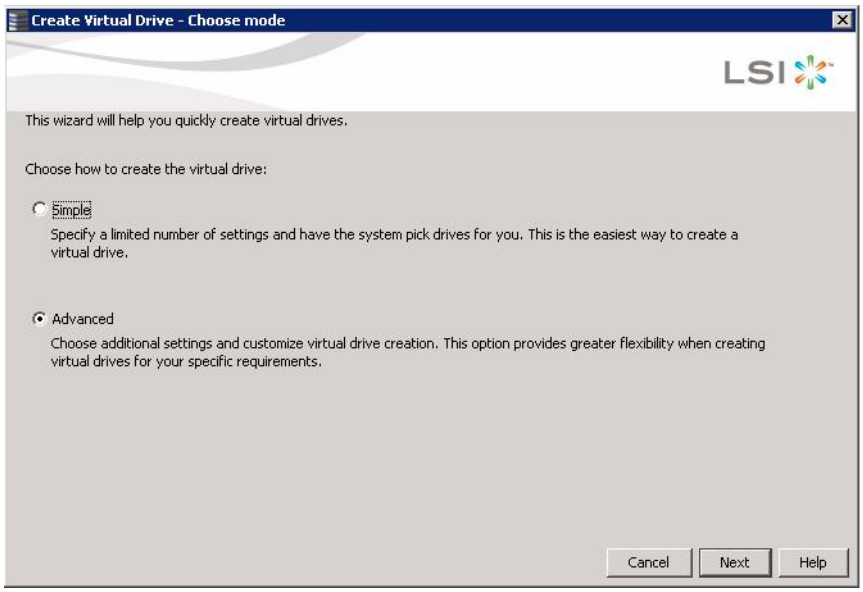


Figure 137: Virtual Drive Advanced Configuration Mode

2. Click **Advanced** and then click **Next**.

   The Create Drive Group Settings screen appears, as shown in Figure 138.

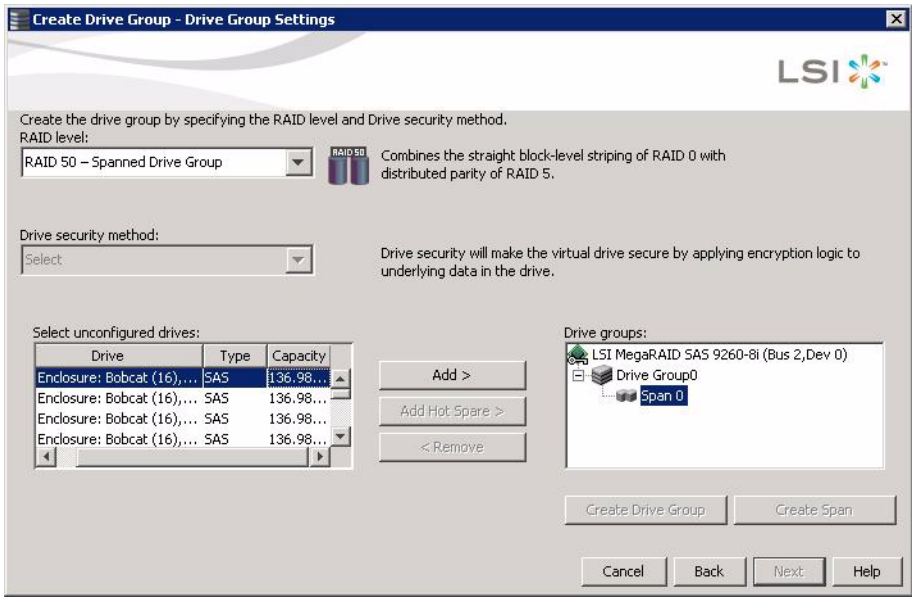


Figure 138: Create Drive Group Settings Screen

3. Select the following items on the Create Drive Group Settings screen:
   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 screen text gives a brief description of the RAID level you select. RAID levels you can choose depend on the number of drives available. 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 Figure 139.

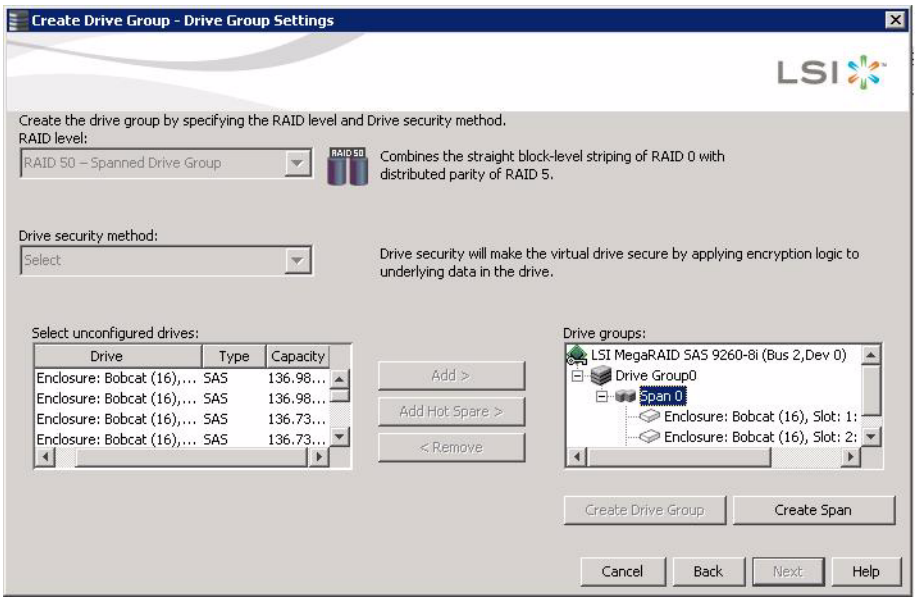


Figure 139: Span 0 of Drive Group 0

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 140.

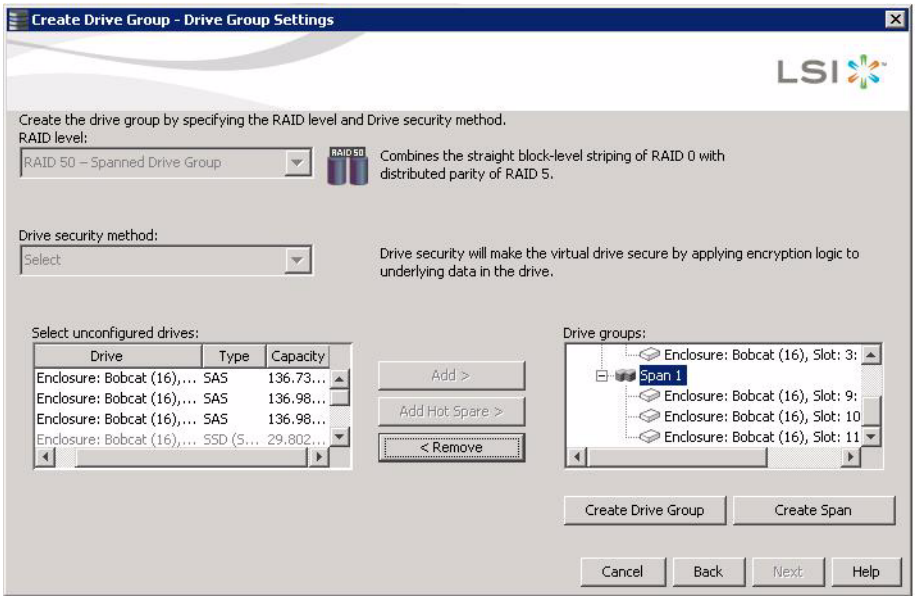


Figure 140: Span 0 and Span 1 of Drive Group 0

f. Click **Create Drive Group** to make a drive group with the spans.

g. Click **Next** to complete this step.

The Virtual drive settings window appears, as shown in Figure 141. 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 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.

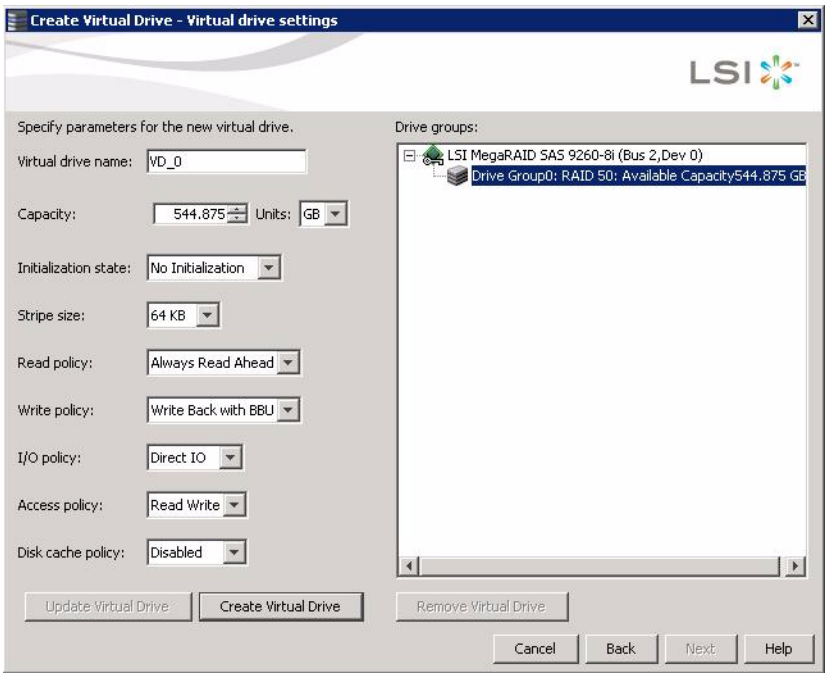


Figure 141: Virtual Drive Settings Window

NOTE: If you select Write Back with BBU as the Write policy, and there is no battery, or the battery is low or failed, or the battery is running through a re-learn cycle, the Write policy switches to Write Through. This eliminates the risk of data loss in case of power failure. A message screen 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 Figure 142. This window shows the selections you made for advanced configuration.

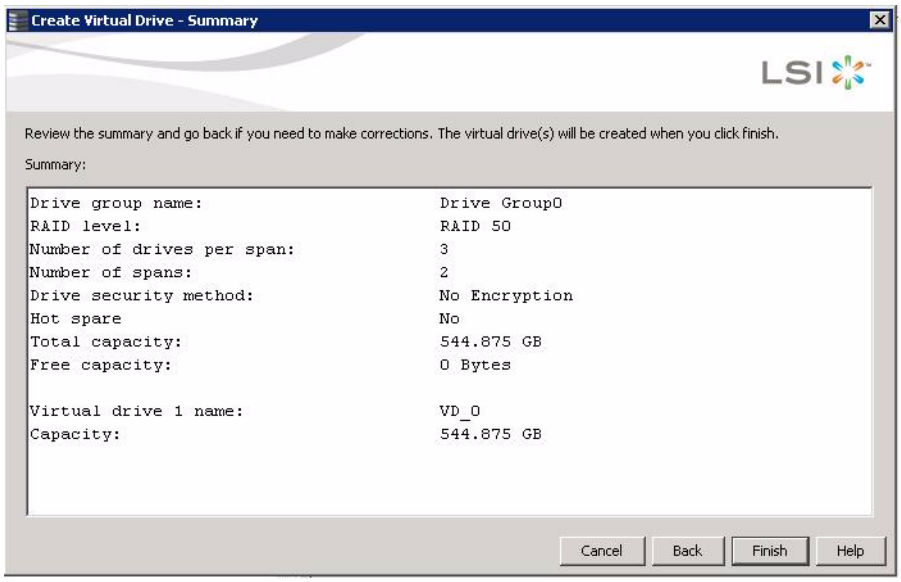


Figure 142: Create Virtual Drive Summary Window

7. Click **Back** to return to the previous screen 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.

NOTE: If you create a large configuration using drives that are in powersave 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 box appears to identify the drive or drives.

   After the configuration is completed, a dialog box notifies you that the virtual drives were created successfully. If more drive capacity exists, the dialog box asks whether you want to create more virtual drives. If no more drive capacity exists, you are prompted to close the configuration session.

8. Select **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 screen, as shown in Figure 134. If you select **No**, the utility asks whether you want to close the wizard.

9. If you selected **No** in the previous step, select **Yes** or **No** to indicate whether you want to close the wizard.

   If you select **Yes**, the configuration procedure closes. If you select **No**, the dialog box 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 a single configuration.

Perform the following steps to configure JBOD to Unconfigured Good drives.

1. Right-click the controller node in the device tree in the left frame of the MegaRAID Storage Manager window and select **Create Virtual Drive**.

-OR-

Select the controller node and select **Go To>Controller>Create Virtual Drive.**

The **Create Virtual Drive - JBOD to Unconfigured Good Conversion** wizard **a**ppears as shown in the following figure.

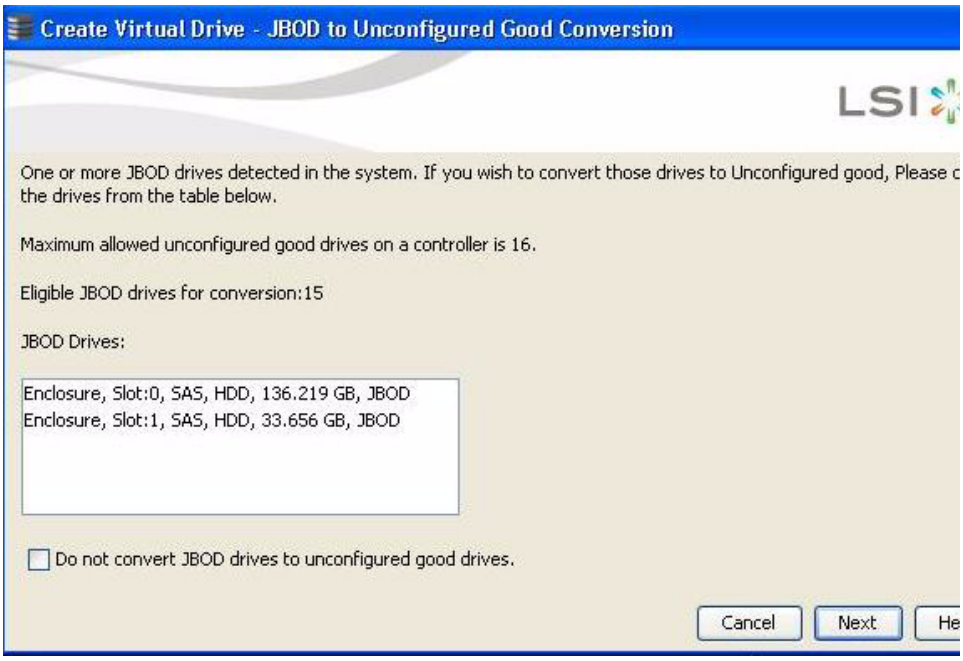


Figure 143: JBOD Drives

The **JBOD drives** field displays the available JBOD drives available in the system.

2. Select the drives which you do not want to be configured as Unconfigured Good, and check the box before "*Do not convert JBOD drives to unconfigured good drives*".

NOTE: If you do not like to make any JBOD as Unconfigured good drive(s), select the check-box before " Do not convert JBOD drives to unconfigured good drives" , then the MSM application skips changing any selected JBOD to unconfigured good drive.

3. Click **Next**.

   The Create Virtual drive Screen appears as shown in Figure 134.

8.2.1 Converting JBOD to Unconfigured Good from the MSM Window

You can convert JBOD to Unconfigured Good by clicking the **Controller** >**Make Configured Good** option from the main MSM window.

The **Make Configured Good** dialog box appears as shown in the following figure.



Figure 144: Make Configured Good

4. Select the JBOD drives to be configured as Unconfigured Good.
5. 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 screen displays the port number, enclosure number, slot number, drive state, drive capacity, and drive manufacturer.
2. Select **Go To>Physical Drive>Assign (G)lobal Hot Spare** or **Go To>Physical Drive>Assign (D)edicated Hot Spare**.
3. If you selected **Assign(De)dicated 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 (G)lobal 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

Follow these steps if you need to change the adjustable rates for rebuilds, and other system tasks that run in the background:

NOTE: LSI recommends 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 value here so you can restore them later, if necessary:
Rebuild Rate: _____
Background Initialization (BGI) Rate: _____
Check Consistency Rate: _____

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 Figure 145.

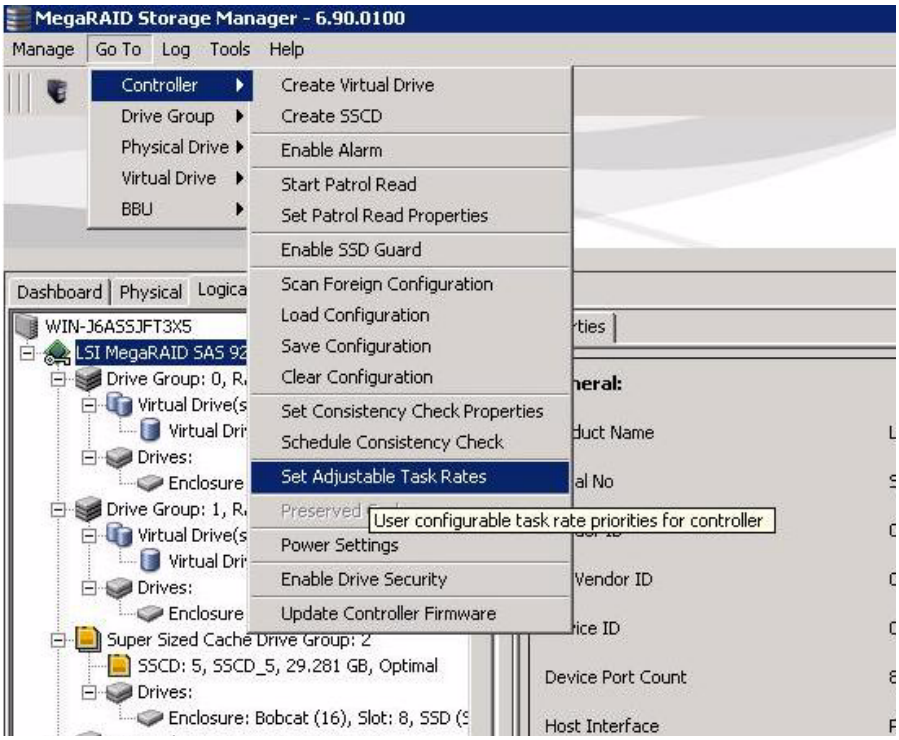


Figure 145: Set Adjustable Task Rates Menu

The Set Adjustable Task Rates screen appears, as shown in Figure 146.

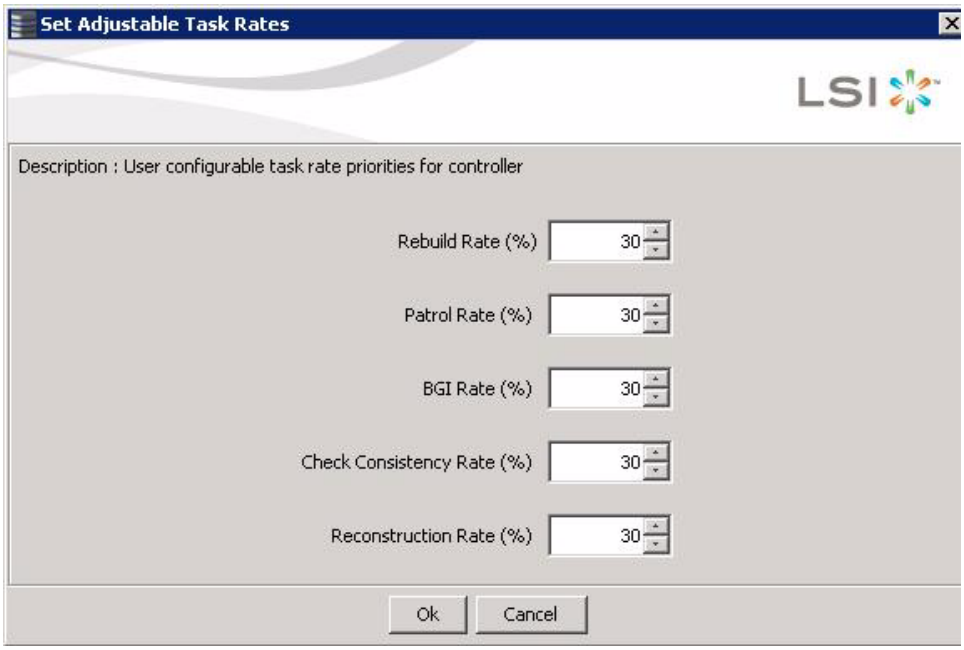


Figure 146: Set Adjustable Task Rates Menu

3. Enter changes, as needed, to the following task rates:
   - Rebuild Rate
   - Patrol Read
   - Background Initialization (BGI) (for fast initialization)
   - Check Consistency (for consistency checks).
   - Reconstruction

   Each task rate can be set from 0 to 100 percent. The higher the number, the faster the activity runs in the background, possibly impacting other system tasks.
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 powersave mode. In powersave 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 MSM to choose whether to allow unconfigured drives or hot spares to enter powersave mode.

NOTE: The Dimmer Switch technology is enabled by default.

When they are in the powersave mode , unconfigured drives and drives configured as hot spares (dedicated or global) can be spun down. When spun down, the drives stay in powersave mode except for periodic maintenance, including:

- Periodic background media scans (Patrol Read) to find and correct media defects to avoid losing data redundancy (hot spare drives only)
- Use of a hot spare to rebuild a degraded drive group (hot spare drives only)
- Update of Disk Data Format (DDF) and other metadata when you make changes to RAID configurations (hot spare drives and unconfigured drives)

NOTE: If your controller does not support this option, the **Power Settings** field does not display.

Follow these steps to change the powersave 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 box displays, as shown in the following figure.



Figure 147: Powersave Mode Checkbox

3. Click the checkbox Allow unconfigured drives to enter powersave mode and then click OK.

   The second Power Settings screen appears, as shown in the following figure.

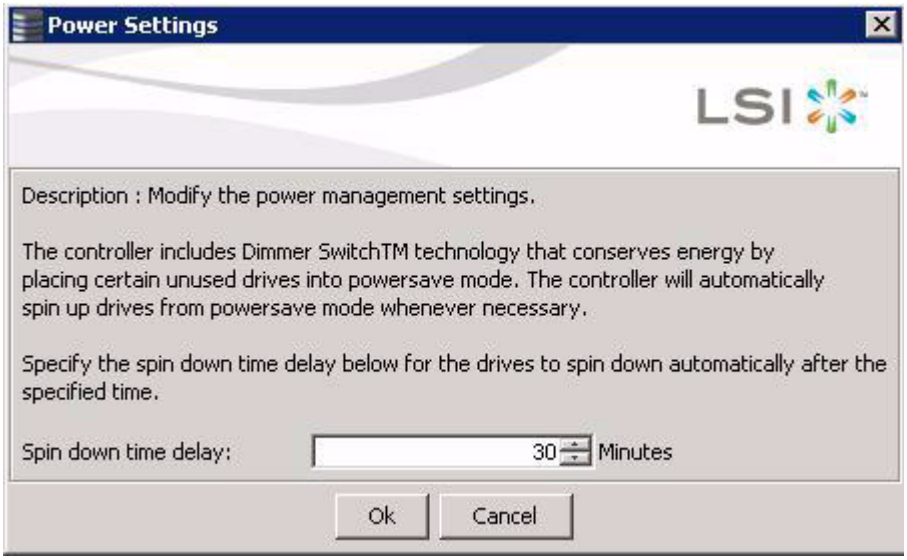


Figure 148: Spin Down Time Delay Setting

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 screen, the nodes for the unconfigured good drives that are spun down appear with - Powersave after their status.

8.5.1 Changing Power Settings Using the Dimmer Switch Enhancement

Follow these steps to change the powersave 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>Manage Power Settings** from the menu bar.

   The Manage Power Settings dialog box appears as shown in the following figure.

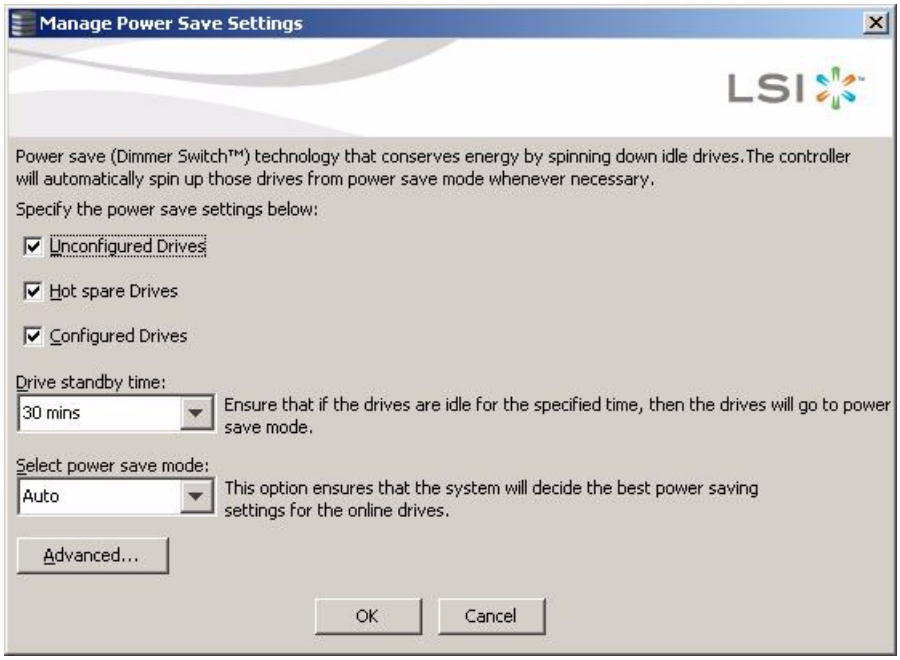


Figure 149: Manage Power Save Settings

3. Select the **Unconfigured Drives** check box; the Unconfigured drives will enter the powersave mode.
4. Select the **Hot spare Drives** check box; the Hot spare drives will enter the powersave mode.
5. Select the **Configured Drives** check box; the Configured drives will enter the power save mode.
6. Select the Drive standby time (alt+d) using the drop-down selector from the **Drive standby time** field. The drive standby time is the time ... write definition here.

NOTE: Until two hours the frequency should be ½ hours and after 2 hours the frequency should be 1 hour.This dropdown will be enabled only if any one of the check boxes above are checked. Else this will be disabled. Refer the descriptions for Drive standby time next to dropdown and Schedule drive active time.

The drive standby time can be 30 Mins, 1 hour, 1.30 Hours, 2.00 Hours ...24 Hours,

7. Select the --- from the **Select power save mode** field. The mode can be Auto, Max or Max without cache (This option will not be shown for DELL and iMR cards depending on the firmware settings)
8. Select the **Advanced** button. The **Power Save Settings - Advanced** screen appears as shown in the following figure.



Figure 150: Power Save Settings - Advanced

9. Select the **Start time** and the **End time** from the **Schedule drive active time** field.
10. Select **Do not schedule drive active time**, if you do not want to schedule the drive active time.
11. Click **OK.**

    The schedule drive active time is done.
12. Right click on **Drive group** and select **Power Settings**.

    The Manage Power Save Settings screen appears, as shown in the following figure.

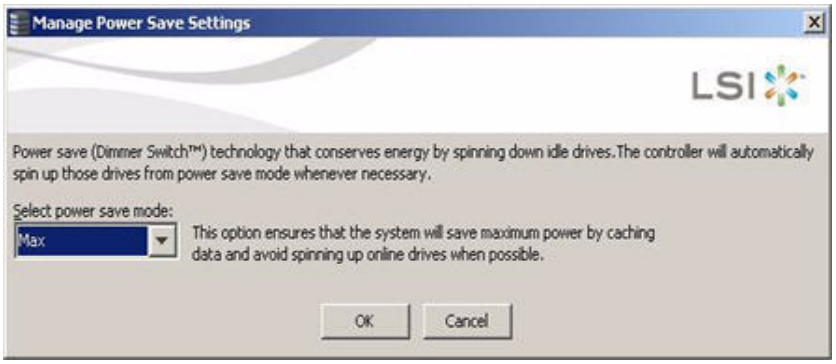


Figure 151: Manage Power Save Settings Screen

13. Select the power save mode from the drop-down selector.

    The values are Max, Max without cache, Auto, None, and controller defined.

8.5.2 Appearance for Controller Defined

This option ensures that the system inherits the controller power save mode for online drives.

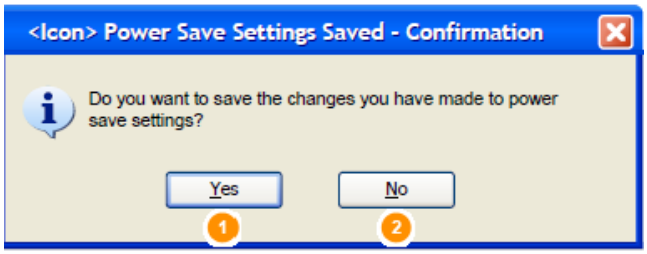


Figure 152: Power Save Settings Saved Confirmation Dialog Box

Click Yes, if you want to save the changes made to power save settings., else click No.

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. Follow these steps to change the virtual drive properties.

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 Figure 153.

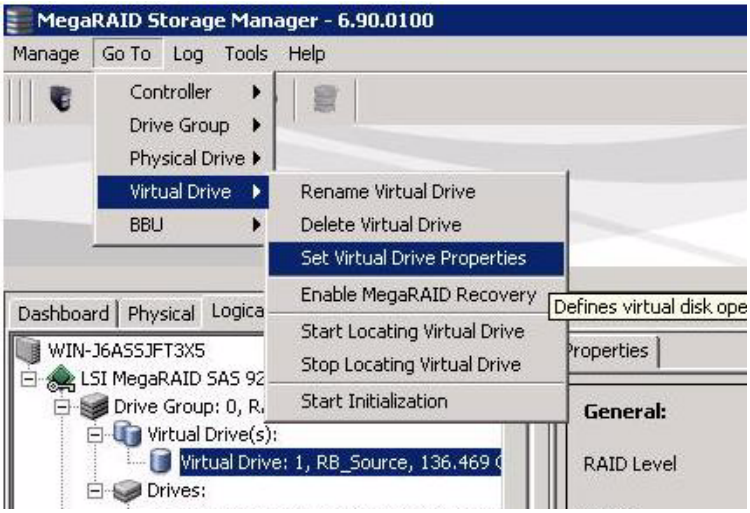


Figure 153: Set Virtual Drive Properties Menu

The Set Virtual Drive Properties dialog box displays, as shown in Figure 8.6.

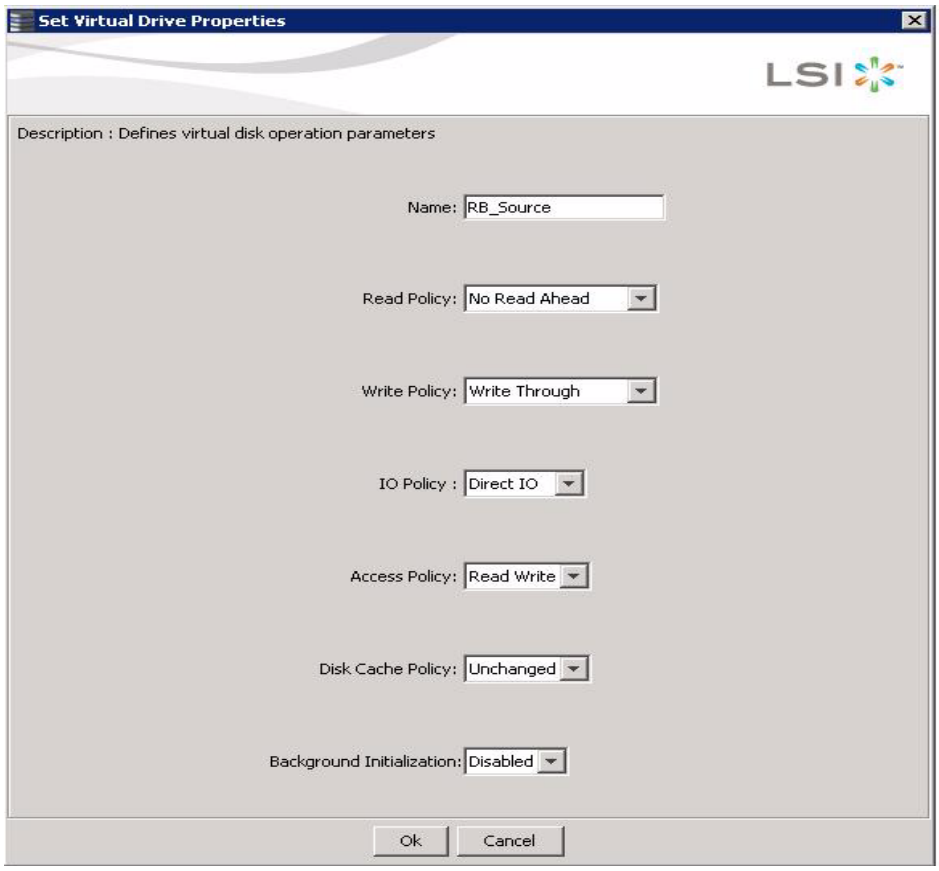


Figure 154: Set Virtual Drive Properties Dialog Box

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 MSM 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, or 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* view 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 screen.
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.

   A 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 155: Reboot Warning Message

4. Select **Confirm** and click **Yes**.

   A warning to back up your data appears, as shown in Figure 156.

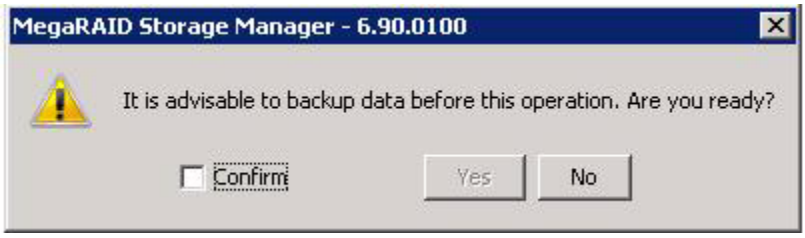


Figure 156: Warning to Back up Data

5. Select **Confirm** and click **Yes**.

   The Modify Drive Group Wizard screen appears, as shown in Figure 157.

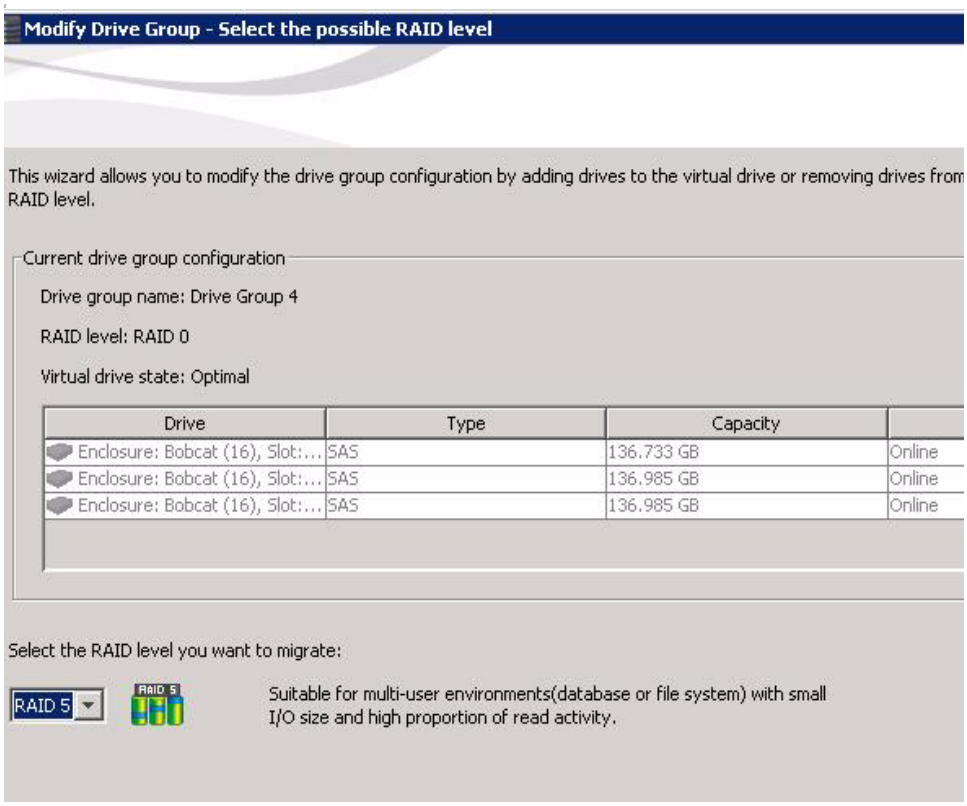


Figure 157: Modify Drive Group Wizard

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. 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.

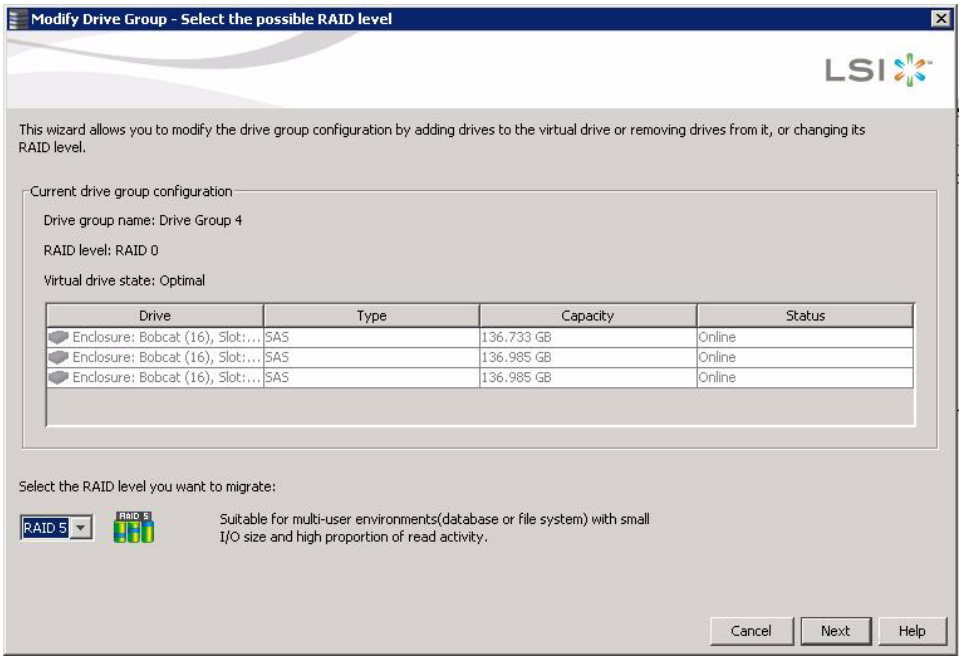


Figure 158: Modify Drive Group Wizard

4. Select the RAID level that you want to change ("migrate") the drive group to and click **Next**.

   The following screen 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.

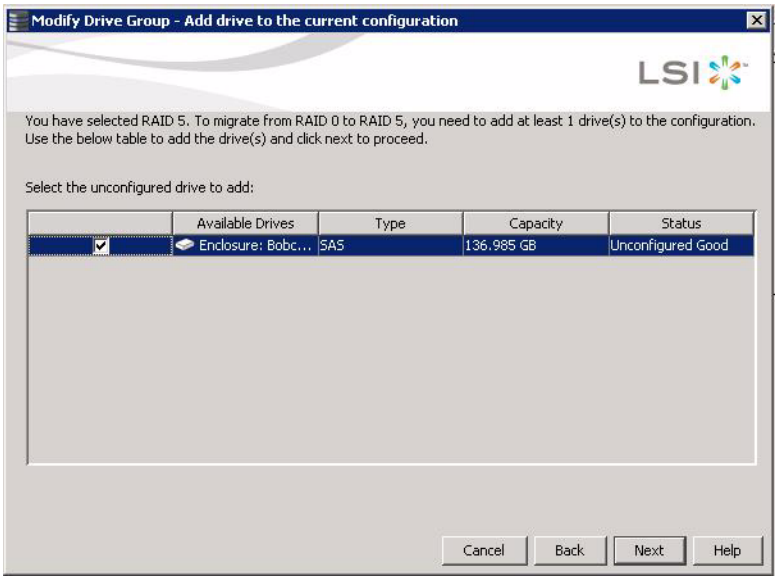


Figure 159: Add Drive(s) to the Current Configuration Screen

5. Click the check box next to any unconfigured drives that you want to add and then click **Next**.

NOTE: The drive(s) 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 Summary screen appears. This screen shows the current settings and what the settings will be after the drives are added.

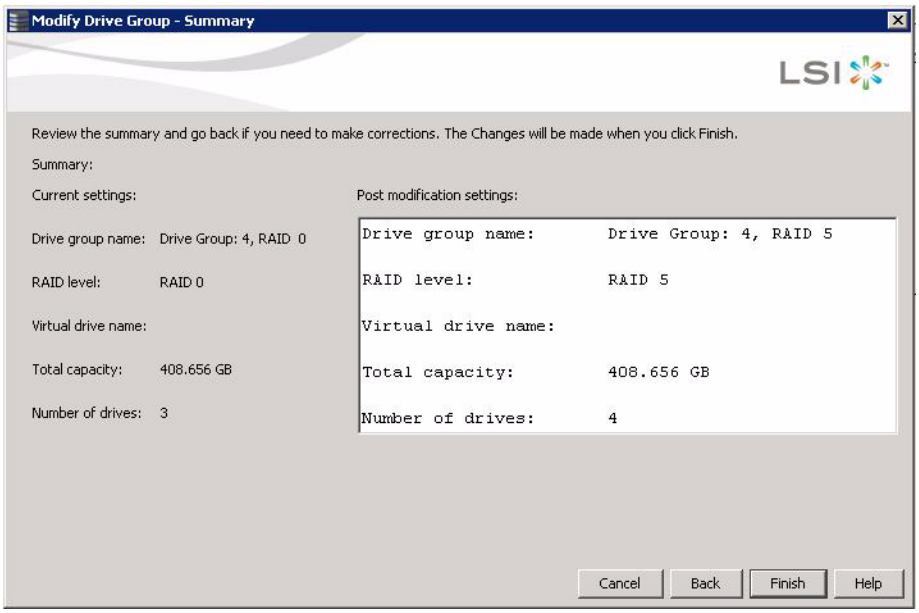


Figure 160: Modify Drive Group Summary Screen

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 main menu screen.
2. Click a drive icon in the left panel of the screen.
3. 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. 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 screen with the replacement drive appears, as shown in Figure 161.

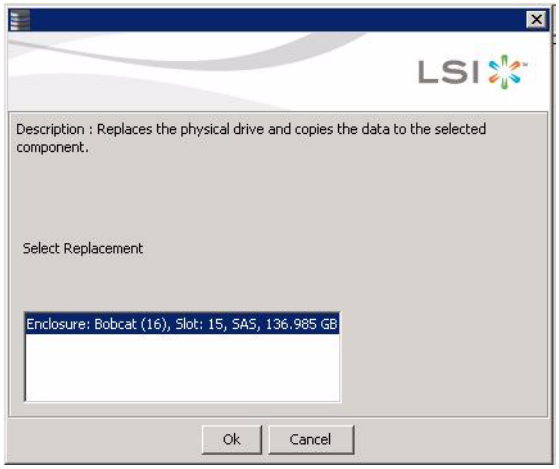


Figure 161: Drive Replacement Window

4. Select a replacement drive.

   A confirmation message appears.

5. Click **Yes**.

   This 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 screen indicates the minimum number of drives you are required to add, if so.

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. 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.

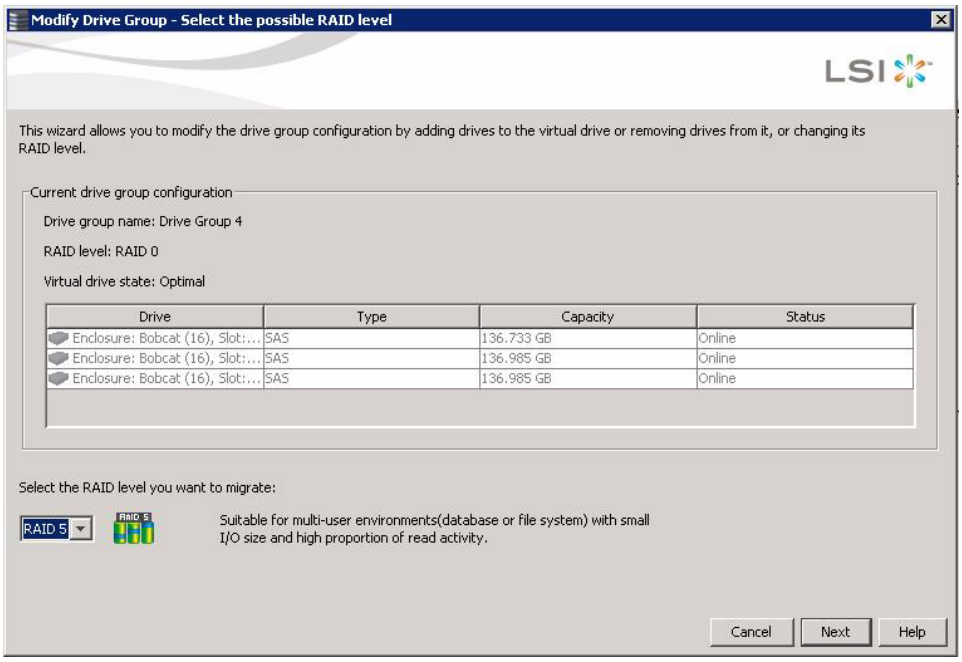


Figure 162: Modify Drive Group Wizard

4. On the Modify Drive Group Wizard screen, select the RAID level that you want to change ("migrate") the drive group to and click **Next**.

   The following screen appears. The screen 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 require more drives.

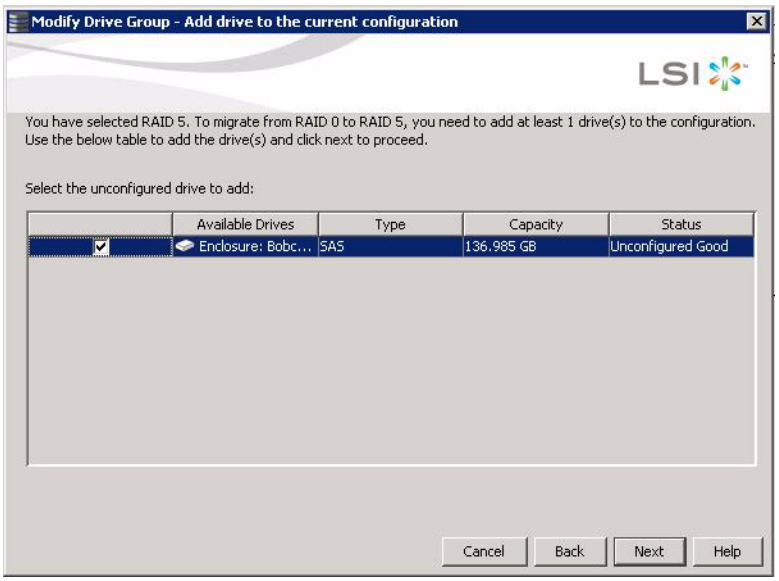


Figure 163: Add Drive(s) to the Current Configuration Screen

5. Select the unconfigured drive or drives to add and then click **Next**.

NOTE: The drive(s) 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 Summary screen appears. This screen shows the current settings and what the settings will be after the drives are added.

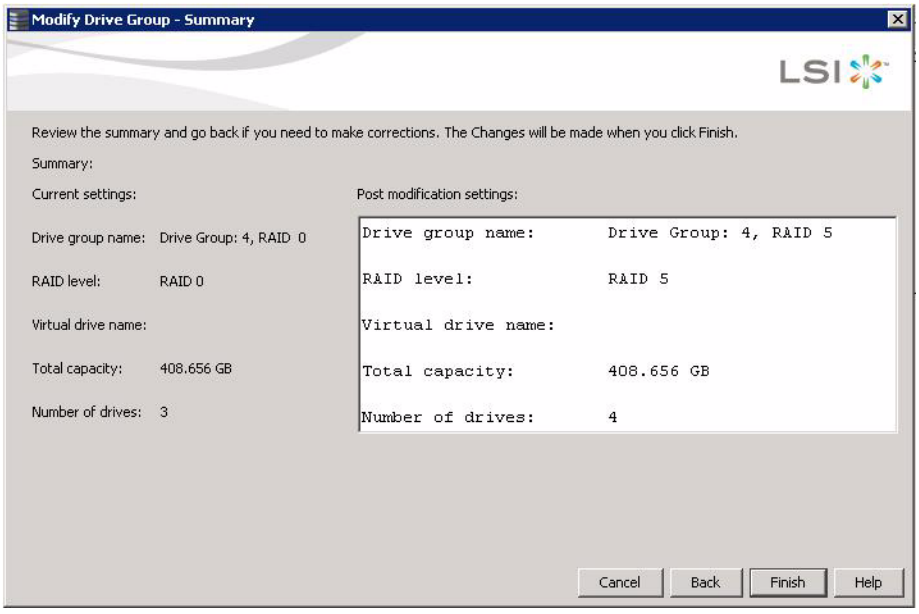


Figure 164: Modify Drive Group Summary Screen

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 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 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 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 4.12.3.3, *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: Be sure to back up the data that is on the virtual drive before you delete it. Be 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 main menu screen, 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 System Events and Storage Devices

This chapter explains how to use MegaRAID Storage Manager software to monitor the status of drives, virtual drives, and other storage devices.

9.1 Monitoring System Events

The MegaRAID Storage Manager utility monitors the activity and performance of all controllers in the system and the storage devices connected to them. When an event occurs (such as the creation of a new virtual drive or the removal of a drive) an event message appears in the log at the bottom of the MegaRAID Storage Manager main menu screen, as shown in Figure 165.

You can use MegaRAID Storage Manager to alert you about events. There are settings are for the delivery of alerts, the severity level of events, exceptions, and email settings.

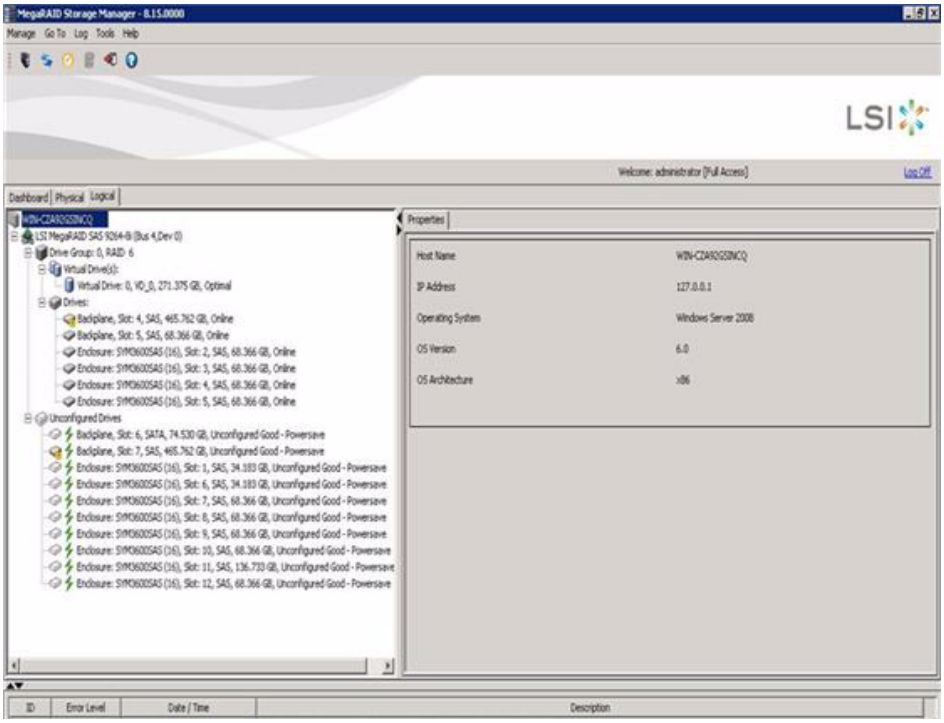


Figure 165: Event Information Window

Each message that appears in the event log has a severity level that indicates the importance of the event, as shown in Table 125, a date and timestamp, and a brief description. You can click an event to display the same information in a window. (For a list of all events, see Appendix A, *Events and Messages*).

Table 125: Event Severity 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.

The Log menu has four options:

- **Save Log**: Saves the current log to a `.log` file.
- **Save Log Text**: Saves the current log in `.txt` format.
- **View Saved Log**: Enables you to load a local `.log` file.
- **Clear Log**: Clears the current log information. You have the option of saving the log first.

9.2 Configuring Alert Notifications

The Alert Notification Configuration feature allows you to control and configure the alerts that MegaRAID Storage Manager software sends when various system events occur.

To access this screen, select **Tools->Configure Alerts** on the main menu screen, as shown in Figure 166.



Figure 166: Alert Notification Configuration Menu

The Alerts Notification Configuration screen appears, as shown in Figure 167. The screen contains three tabs: **Alert Settings**, **Mail Server**, and **Email**. You can use each tab to perform tasks for that topic.

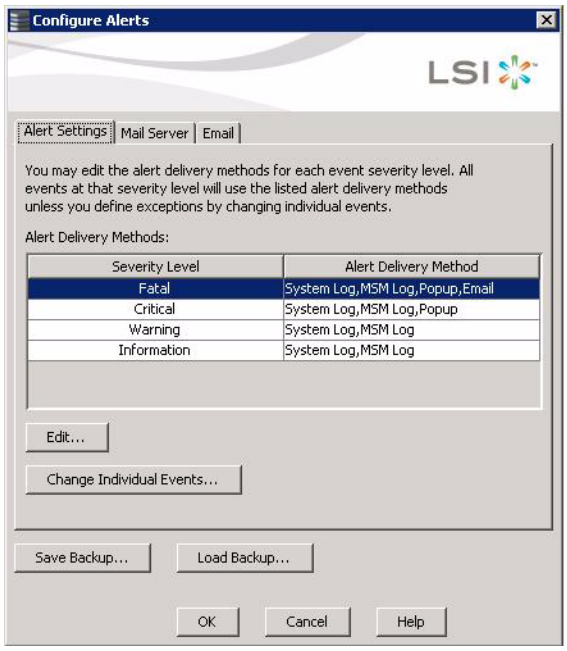


Figure 167: Alerts Notification Configuration Screen

You can select the **Alert Settings** tab to perform the following actions:

- Select the methods for the delivery of alerts.
- Change the severity level of events.
- 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 send to the monitor.

NOTE: When you load a saved backup file, all unsaved changes made in the current session are lost.

You can select the **Mail Server** tab to perform the following actions:

- Enter or edit the sender e-mail address.
- Enter the SMTP server.
- Require authentification of the email server.
- 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 send to the monitor.

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 send to the monitor.

NOTE: When you load a saved backup file, all unsaved changes made in the current session will be lost.

9.2.1 Setting Alert Delivery Methods

You can select the methods used to send alert deliveries, including by popup, email, system log, or MSM log. You can select the alert delivery methods for each event severity level (Information, Warning, Critical and Fatal).

Perform the following steps to select the alert delivery methods:

1. On the Alerts Notification Configuration screen, click the **Alerts Setting** tab.
2. Under the **Alerts Delivery Methods** heading, select one of the severity levels.
3. Click **Edit**.

   The Alert Notification Delivery Methods dialog box appears, as shown in Figure 168.

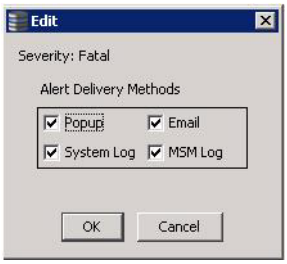


Figure 168: Alert Notification Delivery Methods Dialog Box

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.

9.2.2 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 Alerts Notification Configuration screen, click the **Alerts Setting** tab.

   The the **Alerts Setting** portion of the screen appears, as shown in Figure 167.

2. Click **Change Individual Events**.

   The **Change Individual Events** dialog box appears, as shown in Figure 169. The dialog box shows the events by their ID number, description, and severity level.

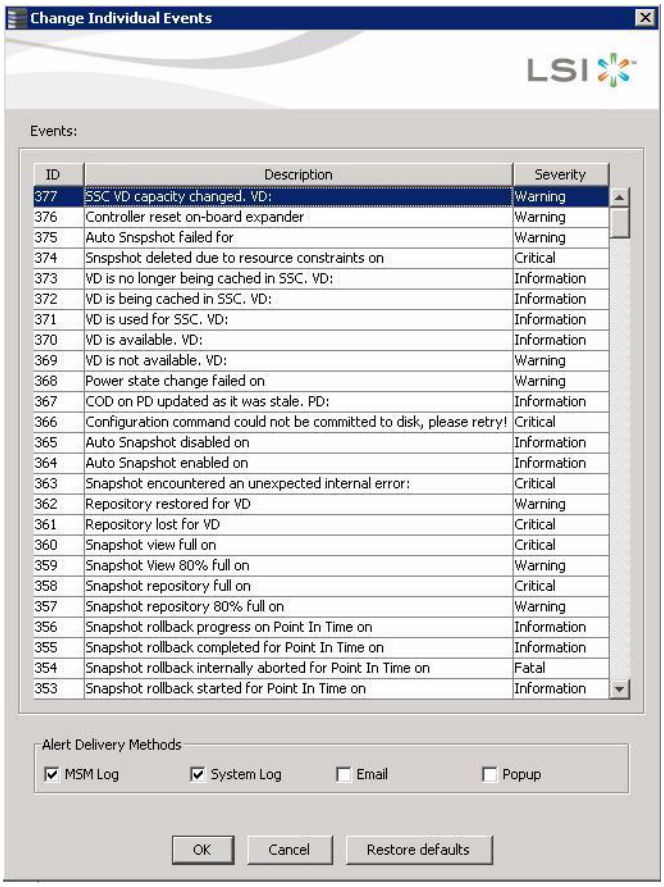


Figure 169: Change Individual Events Dialog Box

3. Click an event in the list to select it.

   The current alert delivery methods appear for the selected event under the **Alert Delivery Methods** heading.

4. Select the desired alert delivery methods for the event.
5. Press ESC to return to the **Alerts Notification Configuration** screen.
6. Click **OK**.

   This saves all of the changes made to the event.

9.2.3 Changing the Severity Level for Individual Events

To change the event severity level for a specific event, perform the following steps:

NOTE: See Table 125 for details about the severity levels.

1. On the Alerts Notification Configuration screen, click the **Alerts Setting** tab.

   The **Alerts Setting** portion of the screen appears.

2. Click **Change Individual Events**.

   The **Change Individual Events** dialog box appears, as shown in Figure 169. The dialog box shows the events by their ID number, description, and severity level.

3. Click an event in the list to select it.

   The current alert delivery methods appear 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 Figure 170.

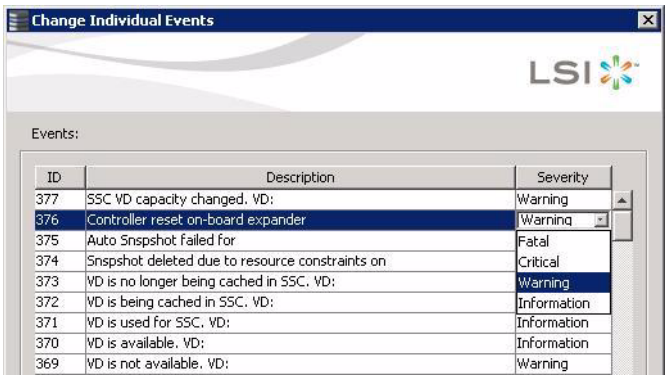


Figure 170: Change Individual Events Severity Level Menu

5. Select a different severity level for the event from the menu.
6. Press ESC to return to the **Alerts Notification Configuration** screen.
7. Click **OK** to save all of the changes made to the events.

9.2.4 Multiple Events Displayed in a Single Pop-up Window

You can view multiple events in a single pop-up window as shown in the following figure.

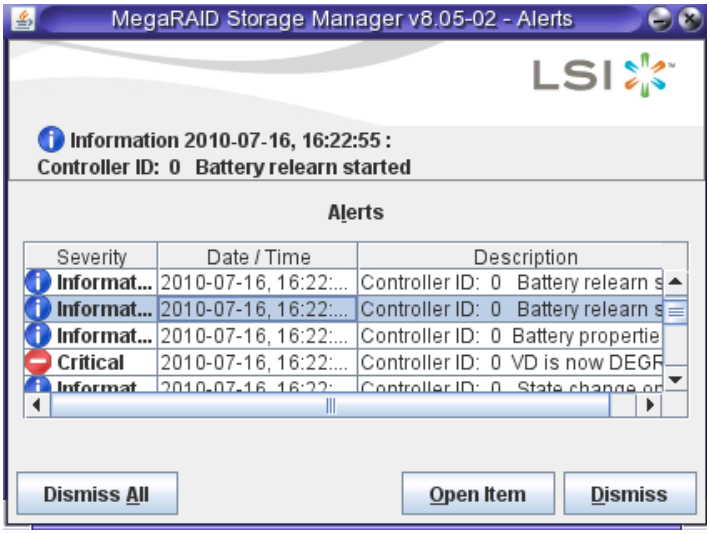


Figure 171: Pop-up for Multiple Events

9.2.5 Entering or Editing the Sender Email Address and SMTP Server

You can use the **Alerts Notification Configuration** screen to enter or edit the sender e-mail address and the SMTP server.

1. On the Alerts Notification Configuration screen, click the **Mail Server** tab.

   The **Mail Server** options appear, as shown in Figure 172.

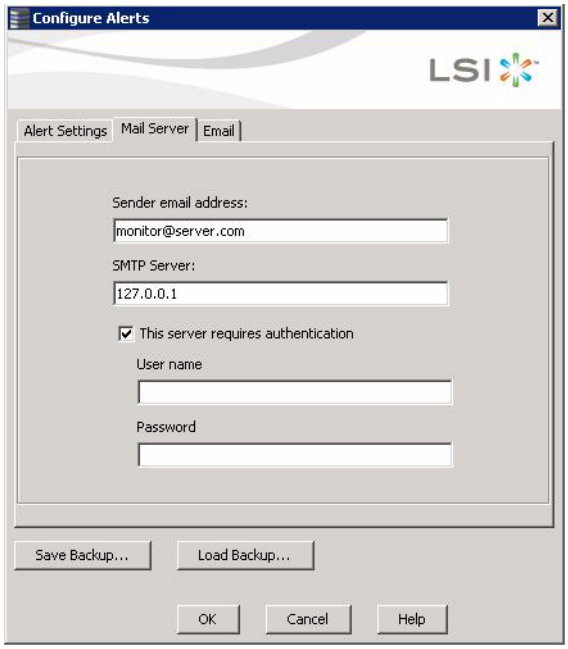


Figure 172: Mail Server Options

2. Enter a new sender email address in the **Sender email address** field or edit the existing sender email address.
3. Click **OK**.

9.2.6 Authenticating a Server

You can use the Alerts Notification Configuration screen to authenticate the SMTP server, providing an extra level of security. The authentication check box enables the **User name** and **Password** fields when selected by default. Clearing the check box disables these fields.

Perform the following steps to enter or edit the address:

1. On the Alerts Notification Configuration screen, click the **Mail Server** tab.

   The **Mail Server** options appears, as shown in Figure 172.
   The authentication check box is selected by default.
2. Enter a user name in the **User name** field.
3. Enter the password in the **Password** field.
4. Click **OK**.

9.2.7 Saving Backup Configurations

You can save an **.xml** backup file of the entire alert configuration.
This includes all the settings on the three tabs.

1. On the Alerts Notification Configuration screen, click the **Alert Setting** tab, **Mail Server** tab, or **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 Alert Notification Configuration screen closes.

9.2.8 Loading Backup Configurations

You can load all of the values from a previously saved backup into the dialog (all tabs) to edit or send to the monitor.

NOTE: If you choose to load a backup configuration and the Configure Alerts dialog currently contains changes that have not yet been sent to the monitor, the changes will be lost. You are prompted to confirm your choice.

1. On the Alerts Notification Configuration screen, click the **Alert Setting** tab, **Mail Server** tab, or **Email** tab.
2. Click **Load Backup**.

   A message warns that when you load a saved backup file, all unsaved changes made in the current session will be lost.

3. Click **Yes**.

   The drive directory appears, from which you can select a backup configuration to load.

4. Select the backup configuration file (it should be in **.xml** format).
5. Click **Open**.

   The drive directory disappears.

6. Click **OK**.

   The backup configuration is laoded and the Alerts Notification Configuration screen closes.

9.2.9 Adding Email Addresses of Recipients of Alert Notifications

The **Email** tab portion of the Alerts Notification Configuration screen shows the email addresses of recipients of the alert notifications. MegaRAID Storage Manager sends alert notifications to those email addresses. Use the screen 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 on the Event Notification Configuration screen.

   The **E-mail** section of the screen appears, as shown in Figure 173.

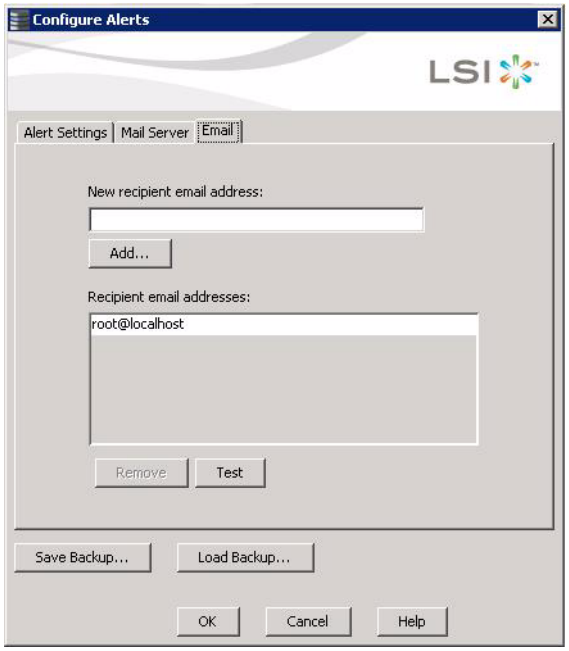


Figure 173: Email Settings

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.

9.2.10 Testing Email Addresses of Recipients of Alert Notifications

Use the **Email** tab portion of the Alerts Notification Configuration screen to send test messages to the email addresses that you added for the recipients of alert notifications.

1. Click the **E-mail** tab on the Event Notification Configuration screen.

   The **E-mail** section of the screen appears, as shown in Figure 173.
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.

   If MegaRAID Storage Manager cannot send an email message to the email address, an error message appears.

9.2.11 Removing Email Addresses of Recipients of Alert Notifications

Use the **Email** tab portion of the Alerts Notification Configuration screen to remove email addresses of the recipients of alert notifications.

1. Click the **E-mail** tab on the Event Notification Configuration screen.

   The **E-mail** section of the screen appears, as shown in Figure 173.
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.

9.3 Monitoring Controllers

When MegaRAID Storage Manager software is running, you can see the status of all controllers in the left panel of the MegaRAID Storage Manager window. If the controller is operating normally, the controller icon looks like this: . If the controller has failed, a small red circle appears to the right of the icon. (See Section 7.2.1, *Dashboard/PhysicalView/Logical View* for a complete list of device icons.)

To display complete controller information, click a controller icon in the left panel of the MegaRAID Storage Manager window. The controller properties display in the right panel, as shown in Figure 174.

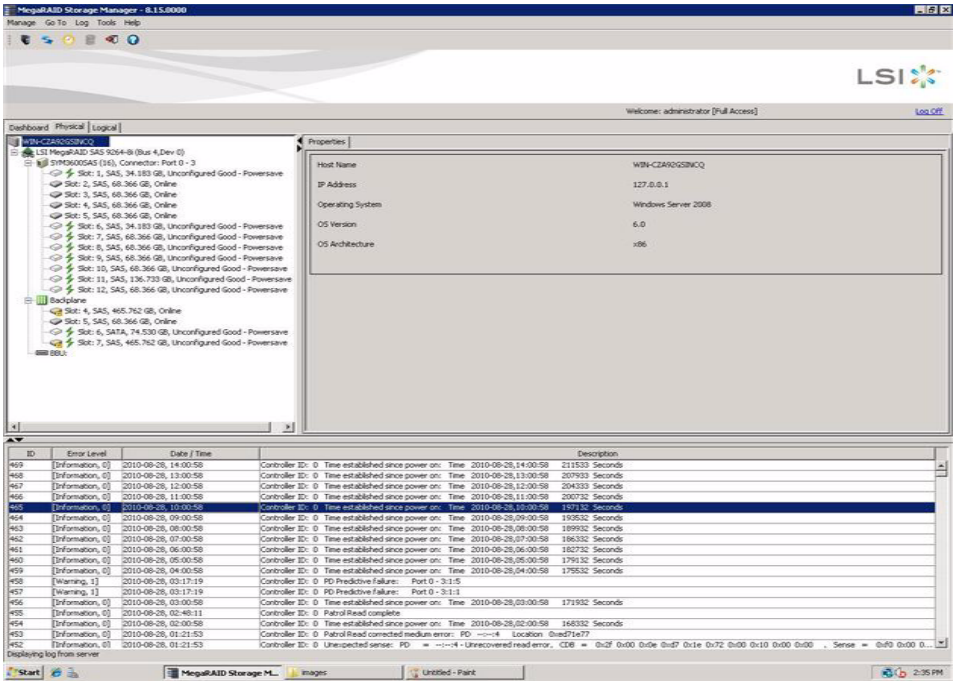


Figure 174: Controller Properties

Most of the information on this screen is self-explanatory. Note the following:

- 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.
- The *Alarm Present* field and the *Alarm Enabled* field indicate whether the controller has an alarm to alert the user with an audible tone when there is an error or problem on the controller.

  The controller properties are defined in Appendix C.

9.4 Monitoring Drives

When MegaRAID Storage Manager software is running, you can see the status of all drives in the left panel of the MegaRAID Storage Manager window. If the drive is operating normally, its icon looks like this: .

If the drive has failed, a small red circle appears to the right of the icon, like this:  . (See Section 7.2.1, *Dashboard/PhysicalView/Logical View* for a complete list of device icons.)

To display complete drive information, click a drive icon in the left panel of the MegaRAID Storage Manager window. The drive properties appear in the right panel, as shown in Figure 175.

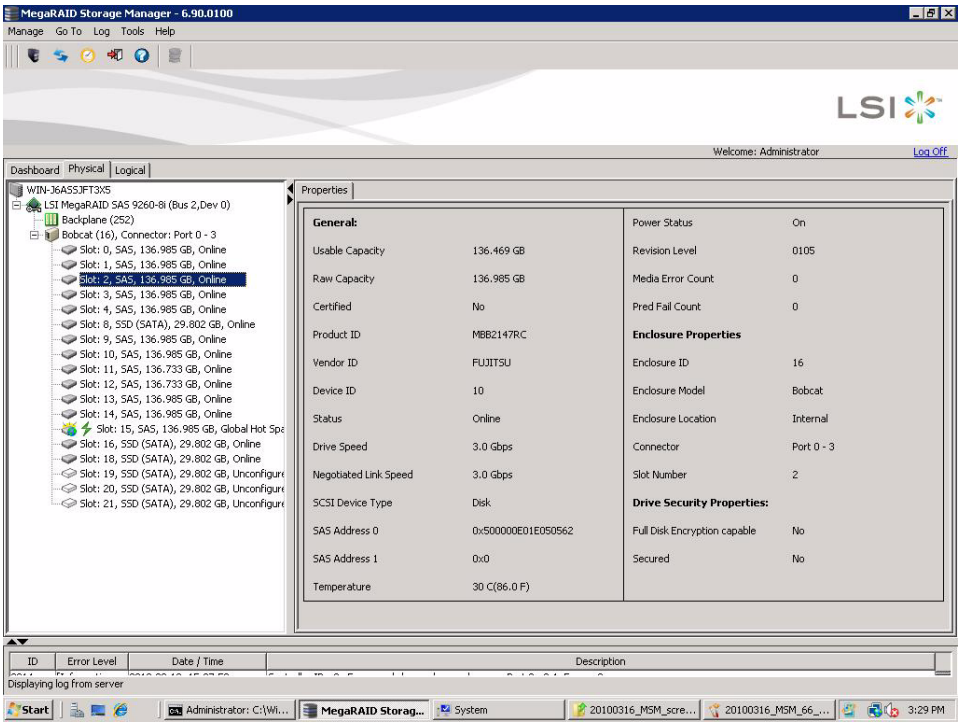


Figure 175: Drive Information

The information on this panel is self-explanatory. There are no user-selectable properties for physical devices. Icons for other storage devices such as CD-ROM drives and DAT drives can also appear in the left panel.

The **Power Status** property shows **On** when a drive is spun up and **Powersave** when a drive is spun down. Note that CacheCade software drives and other drives that never spin down still show **On**.

If the drives are in a drive enclosure, you can identify which drive is represented by each drive LED on the enclosure. Follow these steps to locate the drive:

1. Click the drive icon in the left panel.
2. Click **Go To>Physical Drive>Start Locating Drive**.

   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 LED on the enclosure from blinking, select **Go To>Physical Drive>Stop Locating Drive**.

The drive properties are defined in the Glossary.

To display a graphical view of a drive, click a drive icon in the left panel of the MegaRAID Storage Manager window, and click the **Graphical View** tab. In Graphical View, the drive's storage capacity is color coded according to the legend shown on the screen: configured space is blue, available space is white, and reserved space is red. When you select a virtual drive from the drop-down menu, the drive space used by that virtual drive is displayed in green.

9.5 Running a Patrol Read

A patrol read periodically verifies all sectors of drives connected to a controller, including the system reserved area in the RAID configured drives. A patrol read can be used for all RAID levels and for all hot spare drives. This operation is initiated only when the controller is idle for a defined time 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.

To set the patrol read properties and then start a patrol read, follow these steps:

1. Click a controller icon in the left panel of the MegaRAID Storage Manager main menu.
2. Select **Go To>Controller>Set Patrol Read Properties**.

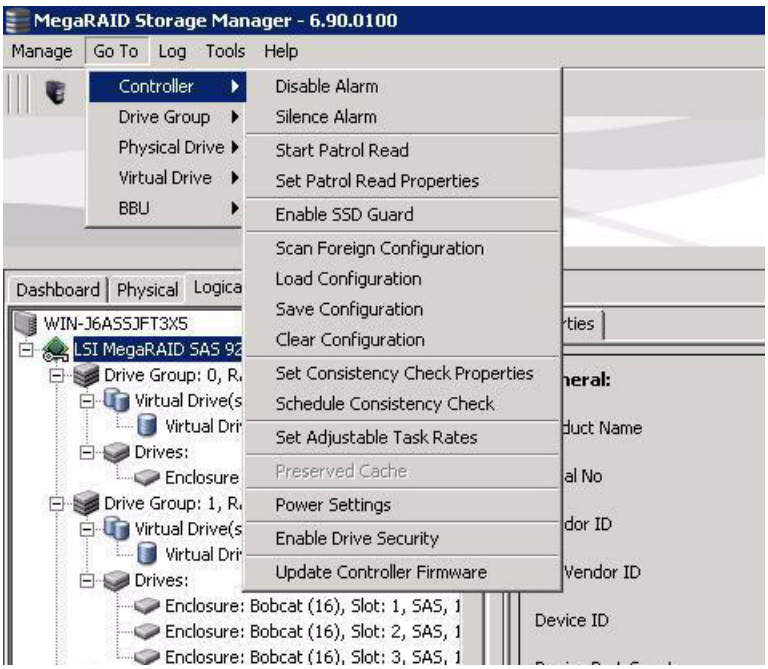


Figure 176: Start Patrol Read Menu

The Patrol Read - Set properties screen displays, as shown in Figure 177.

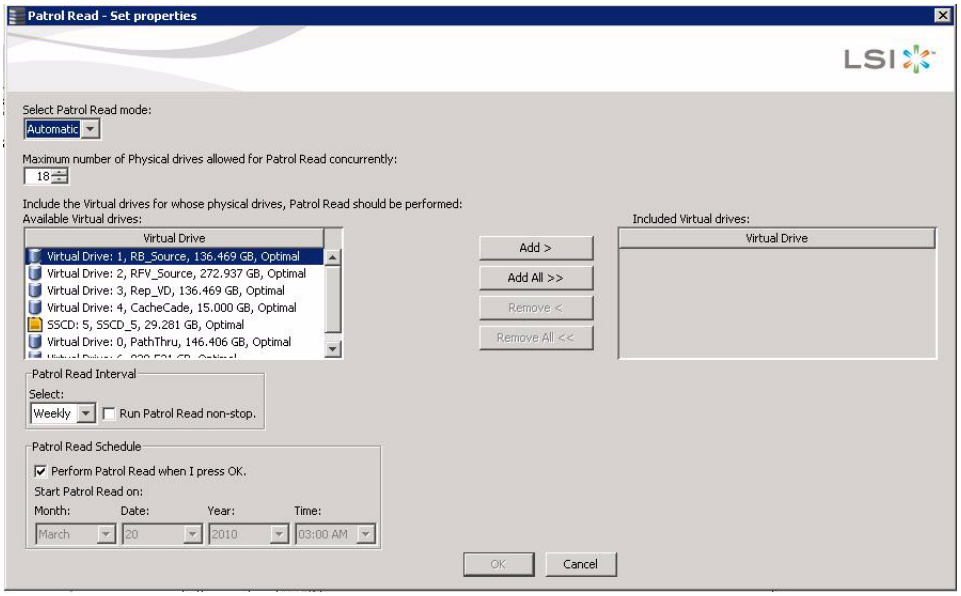


Figure 177: Patrol Read Configuration

3. Select a mode for a patrol read. The options are:
   - **Automatic**: Patrol read runs automatically at the time interval you specify on this screen.
   - **Manual**: Patrol read runs only when you manually start it by selecting **Start Patrol Read** from the controller Options panel.
   - **Disabled**: Patrol read does not run.
4. Specify a maximum count of drives to include in the patrol read. The count must be from 1 to 255.
5. 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.
6. (Optional) Change the frequency at which the patrol read will run.

   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 value here so you can restore them later, if necessary:
Patrol Read Frequency: _____
Continuous Patrolling: Enabled/Disabled
Patrol Read Task Rate: _____

7. (Optional) Set Patrol Read to run at a specific time.

   The default is for the patrol read to start when you click **OK** on this screen. To change the default so that the patrol read starts at a specific time, follow these steps (otherwise, skip this step and proceed to the next step):

a. Uncheck the box **Perform Patrol Read when I click OK**.
b. Select the month, year, day, and time to start patrol read.

8. 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 in the event log only.

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.
3. When prompted, click **Yes** to confirm that you want to start a patrol read.

9.5.1 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, foreground tasks will run more slowly and it might 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 Section 8.4, *Changing Adjustable Task Rates*

9.6 Monitoring Virtual Drives

When MegaRAID Storage Manager is running, you can see the status of all virtual drives. If a virtual drive is operating normally, the icon looks like this:  . If the virtual drive is running in Degraded mode (for example, if a drive has failed), a small yellow circle appears to the right of the icon:  . A red circle indicates that the virtual drive has failed and data has been lost.

When the Logical tab is selected, the left panel of the MegaRAID Storage Manager window 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 a virtual drive icon in the left panel. The properties appear in the right panel. All virtual drive properties are defined in Chapter C, *Glossary*. Figure 178 shows the Properties panel for a virtual drive.

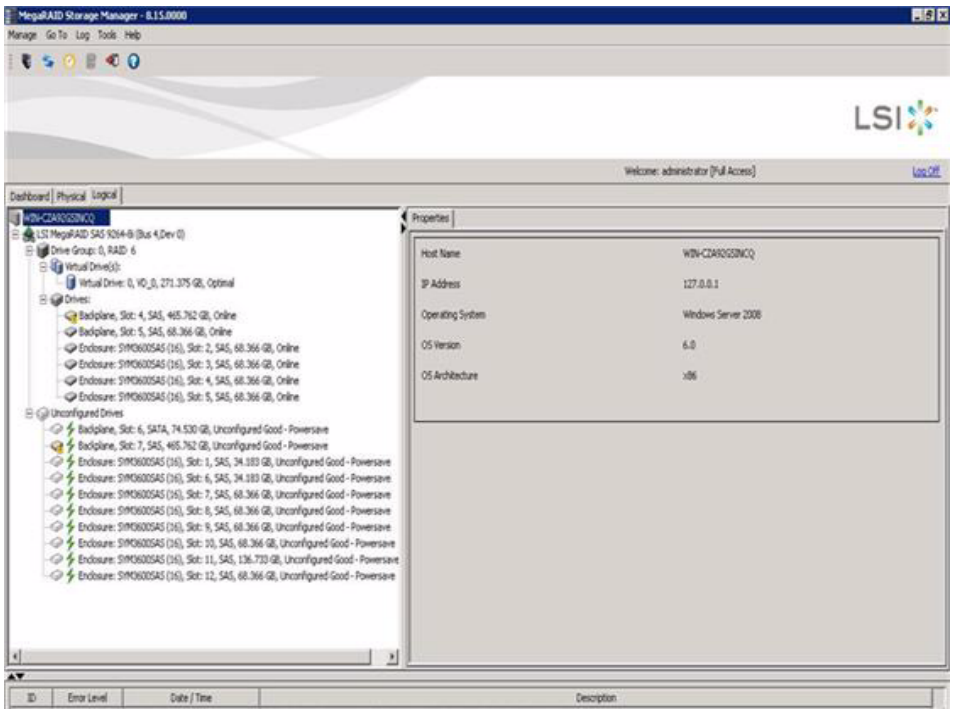


Figure 178: Virtual Drive Properties

The RAID level, stripe size, and access policy of the virtual drive are set when it is configured.

NOTE: You can change the read policy, write policy, and other virtual drive properties. See Section 8.6, *Changing Virtual Drive Properties* for the procedure you can use to change these properties.

If the drives in the virtual drive are in an enclosure, you can identify them by making their LEDs blink. To do this, follow these steps:

1. Click the virtual drive icon in the left panel.
2. Click **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 hot spare drives).
3. To stop the LEDs from blinking, click **Go To>Virtual Drive>Stop Locating Virtual Drive**.

9.7 Monitoring Enclosures

When 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 the enclosure is not functioning normally—for example, if a fan has failed—a small yellow or red circle appears to the right of the icon.

9.8 Monitoring Battery Backup Units

When MegaRAID Storage Manager 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 has failed, 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

9.8.1 Battery Learn Cycle

Learn Cycle is a battery calibration operation 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 cyles, 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.

9.8.1.1 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. Click the **Go To>BBU>Set Learn Cycle Properties**.

   The BBU operations screen appears.
4. On the BBU operations screen, click **Enable automatic learn cycles** and click **Go**.

   You can delay the start of the next learn cycle by up to 7 days (168 hours) using the **Delay next learn cycle** field.
5. To disable automatic learn cycles, click **Disable automatic learn cycles** and then click **Go**.

You can start the learn cycles manually. In addition, you can check the box next to the field **Remind me when to start a learn cycle** to receive a reminder to start a manual learn cycle.

9.8.1.2 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. Click the **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 menu.

9.9 Monitoring Rebuilds and Other Processes

MegaRAID Storage Manager software allows you to monitor the progress of rebuilds and other lengthy operations in the Group Show Progress window.

Follow these steps to monitor the progress of these operations.

1. Open this window, shown in Figure 179, by selecting **Manage>Show Progress** on the menu bar.

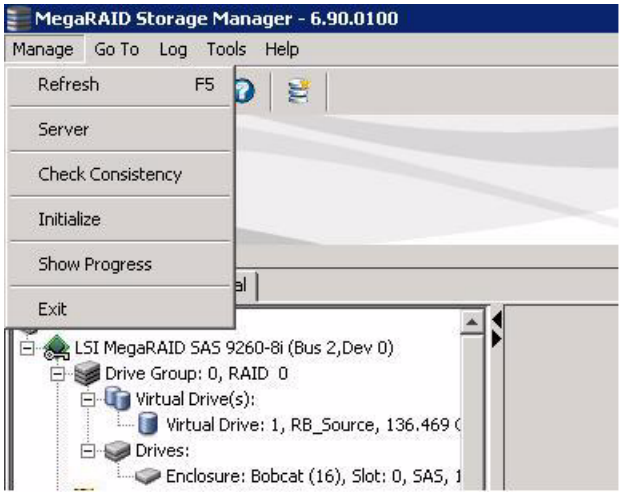


Figure 179: Group Show Progress Menu

The Group Show Progress window appears, as shown in Figure 180.

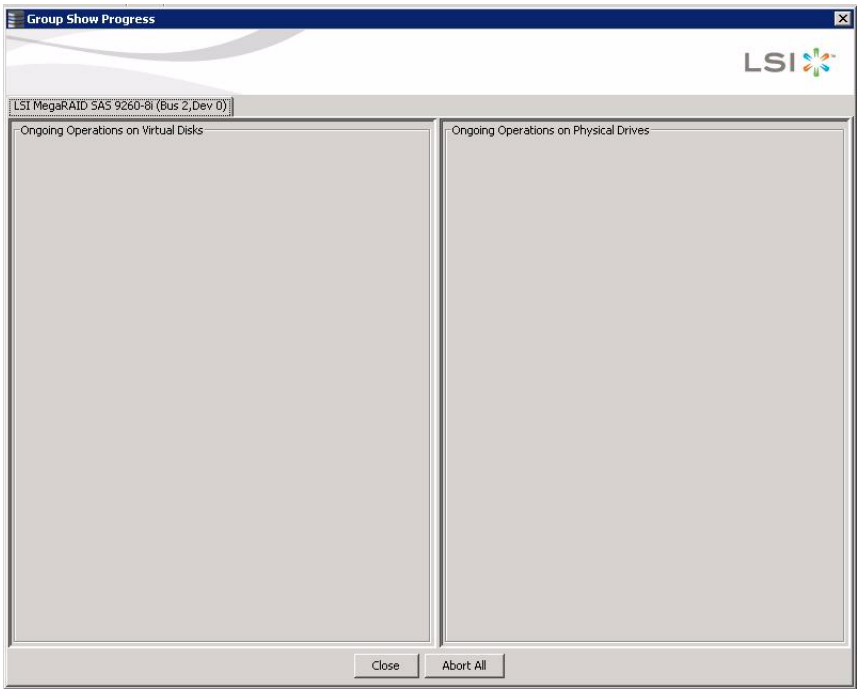


Figure 180: Group Show Progress Window

Operations on virtual drives appear in the left panel of the Group Show Progress window, and operations on drives appear in the right panel. The following operations appear in this window:

- Background or foreground initialization of a virtual drive (see Section 10.1, *Initializing a Virtual Drive*)
- Rebuild (see Section 10.4, *Rebuilding a Drive*)
- Modify Drive Group (see Section 8.7, *Changing a Virtual Drive Configuration*)
- Check Consistency (see Section 10.2, *Running a Consistency Check*)

2. (Optional) Click **Abort All** to abort all ongoing processes.
3. Click **Close** to close the window.

Chapter 10

Maintaining and Managing Storage Configurations

This chapter explains how to use MegaRAID Storage Manager software to maintain and manage storage configurations.

10.1 Initializing a Virtual Drive

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 main menu, and click the icon of the virtual drive that you want to initialize.
2. Select **Go To>Virtual Drive>Start Initialization**.

   The initialize dialog box appears.
3. Select the virtual drive(s) 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. 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 unchecked, MegaRAID Storage Manager software will run a Full Initialization on the virtual drive. See Section 8.1.1, *Selecting Virtual Drive Settings* for more information.
5. Click **Start** to begin the initialization.

   You can monitor the progress of the initialization. See Section 9.9, *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. Click **Manage>Initialize**.

   The Group Consistency Check appears, as shown in Figure 185.

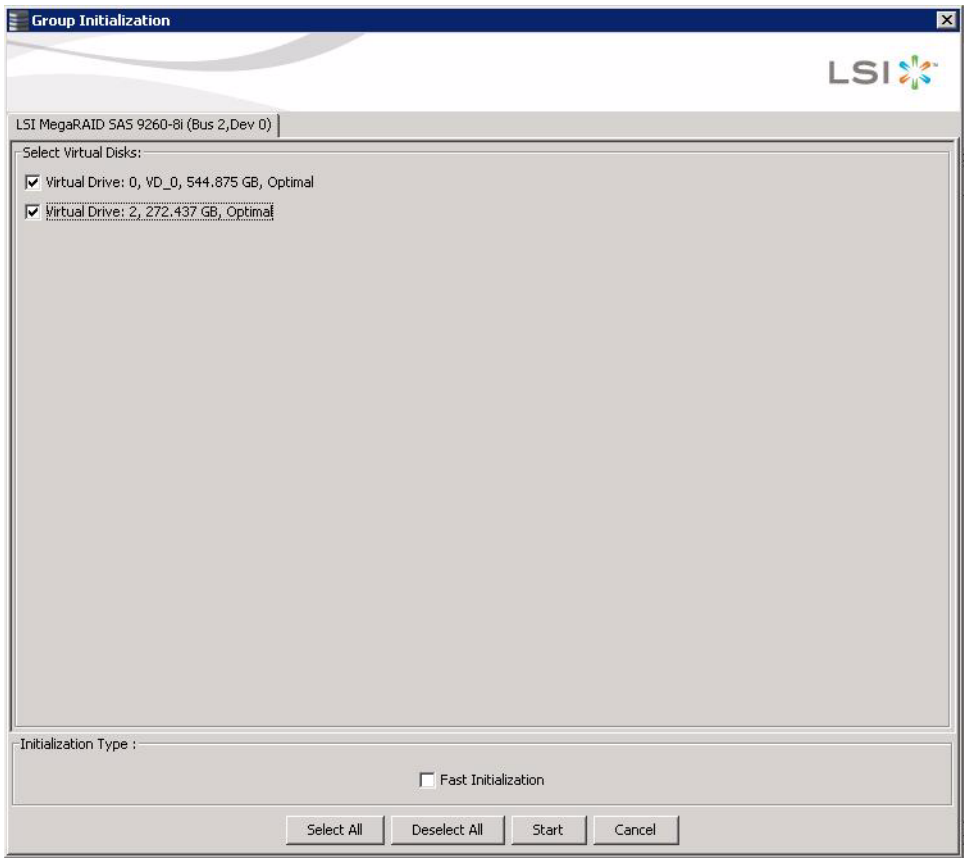


Figure 181: Group Initialization Dialog Box

2. Check the virtual drives to run the initialization on 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 9.9, *Monitoring Rebuilds and Other Processes* for more information.

10.2 Running a 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.

You should run a consistency check on fault-tolerant virtual drives periodically. You must run the consistency check if you suspect that the virtual drive data might be corrupted. Be 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 Logical tab and select the controller.
2. Click **Go To>Controller>Set Consistency Check Properties**.

   Figure 182 shows the consistency check properties menu item.

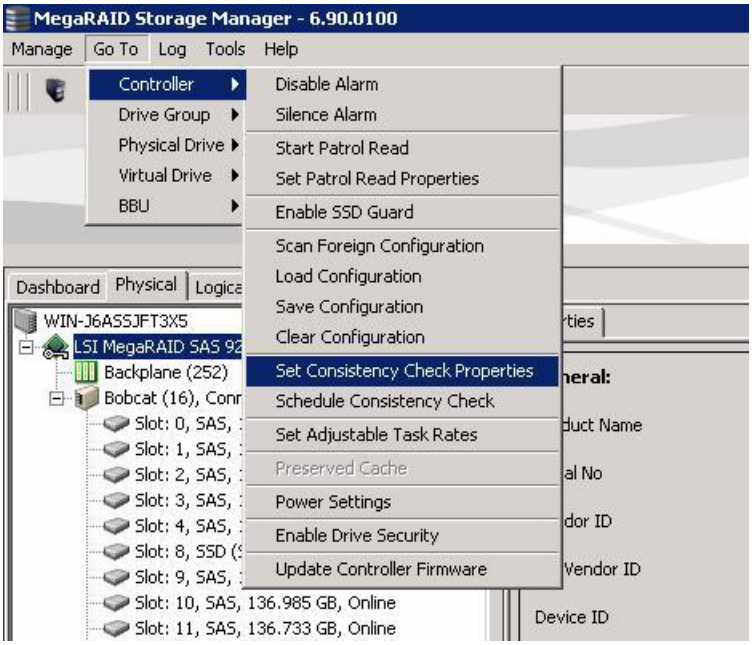


Figure 182: Set Consistency Check Properties Option

The Set Consistency Check Properties dialog box appears, as shown in Figure 183.

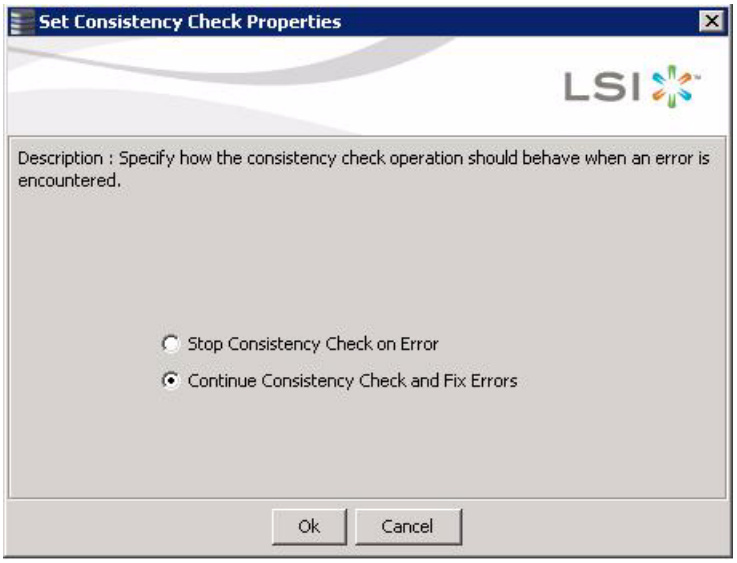


Figure 183: Set Consistency Check Properties Dialog Box

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 and 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 Logical tab and select the controller.
2. Click **Go To>Controller>Schedule Consistency Check**.

   The Schedule Consistency dialog box appears, as shown in Figure 183.

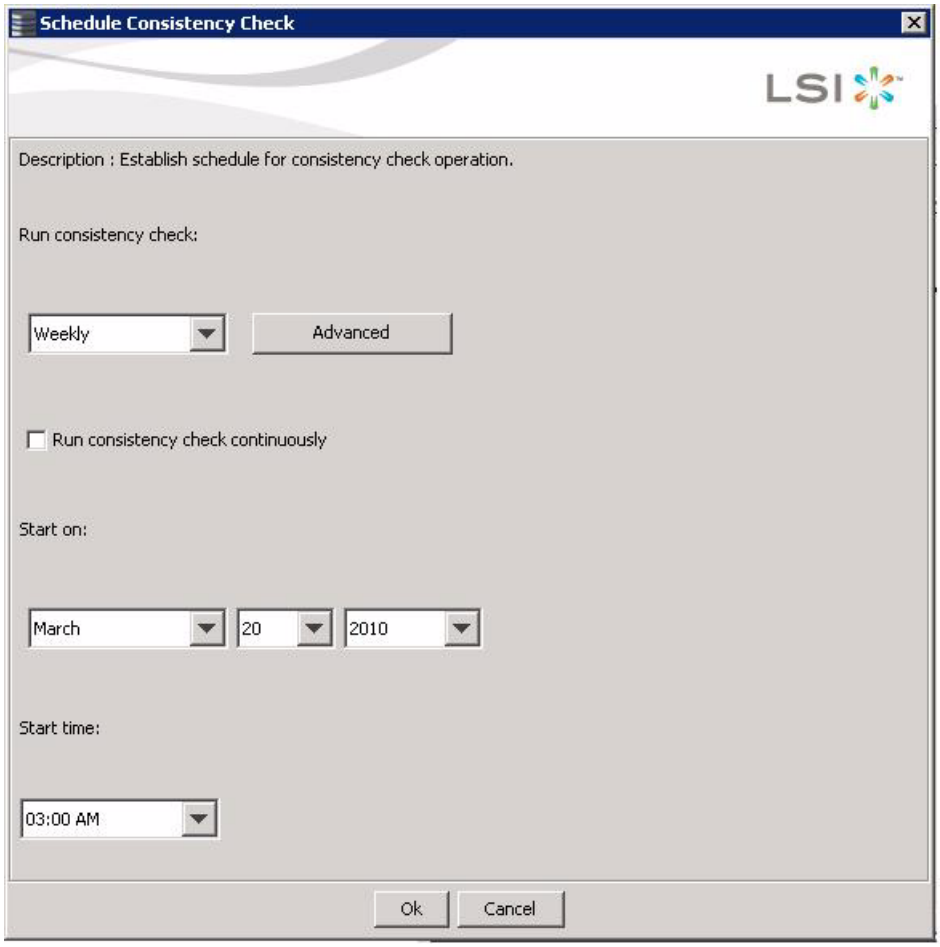


Figure 184: Schedule Consistency Check Dialog Box

3. Perform the following steps to schedule the consistency check:
   a. Select how often to run the consistency check from the drop-down menu. You can click **Advanced** for more detailed date options.
   b. (Optional) Check the **Run consistency check continuously** checkbox.
   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 9.9, *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. Click **Manage>Check Consistency**.

   The Group Consistency Check appears, as shown in Figure 185.

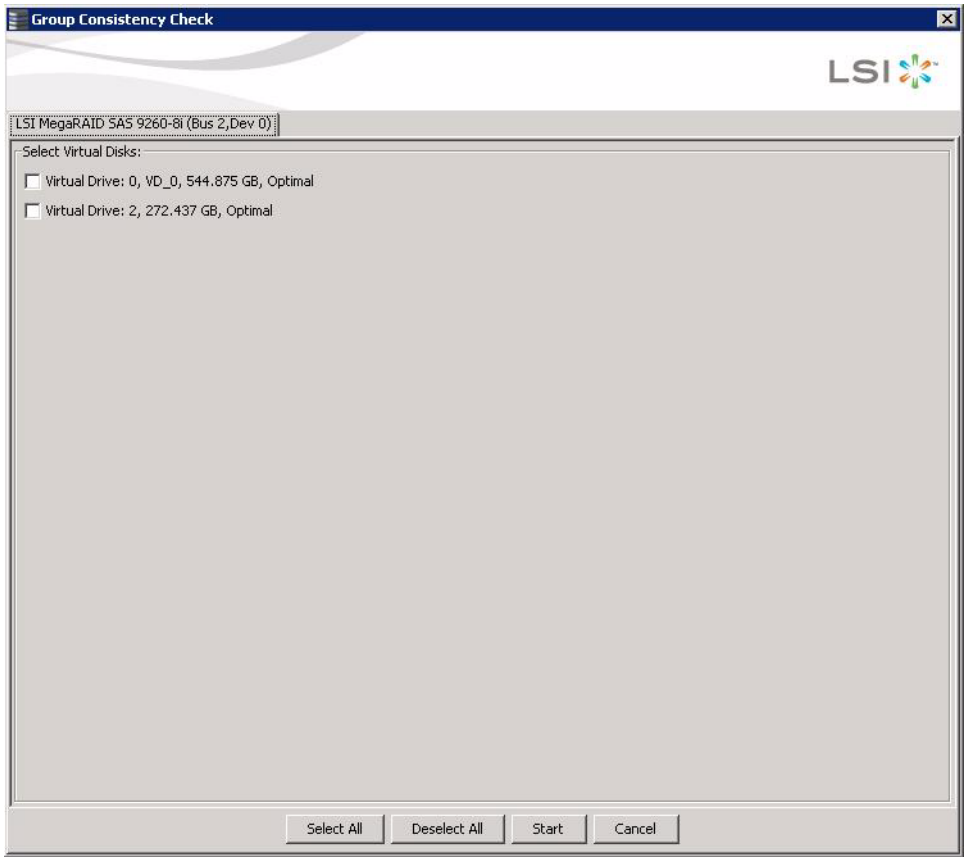


Figure 185: Group Consistency Check Dialog Box

2. Check the virtual drives to run the consistency check on 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 9.9, *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 MegaRAID Storage Manager.

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.

MegaRAID Storage Manager software normally detects newly installed drives and displays icons for them in the MegaRAID Storage Manager window. If for some reason 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 MegaRAID Storage Manager software detects any new drives, it displays a list of them on the screen. If not, it notifies you that no foreign configuration is found.
3. Follow the instructions on the screen to complete the drive detection.

10.4 Rebuilding a Drive

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. If hot spare disks are available, a failed drive is rebuilt automatically without any user intervention. A 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 hot spare drives are available, the failed drive is rebuilt automatically without any user intervention.

If a drive has failed, a red circle appears to the right of the drive icon:  . A small yellow circle appears to the right of the icon of the virtual drive that uses this drive:  . This 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 if you need 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**.

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 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 main menu, click **Go To>Physical Drive>Make Drive (O)ffline**.

   The drive status changes to Offline.

2. Click **Go To>Physical Drive>(M)ark 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 MegaRAID Storage Manager software detects that a drive in a virtual drive has failed, it makes the drive offline. If this happens, 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** and **Rescan** commands.

10.6 Upgrading the Firmware

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 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 MegaRAID Storage Manager main menu, click **Go To>Controller>Update Controller Firmware**.
2. Click **Browse** to locate the `.rom` update file, as shown in Figure 186.

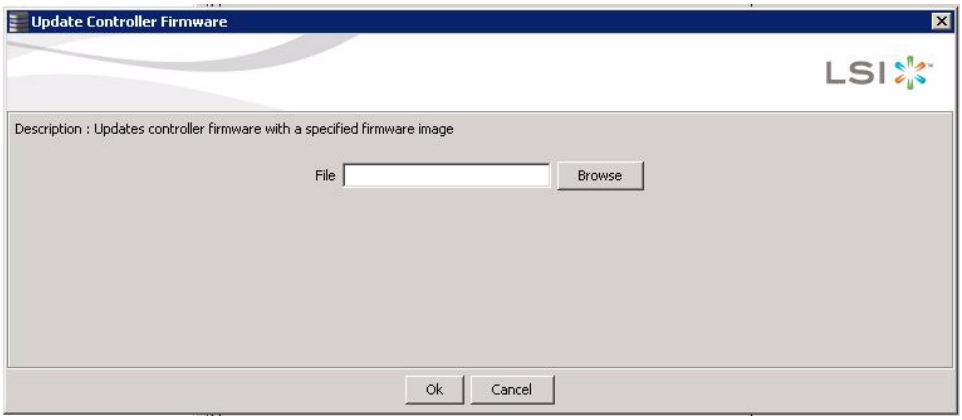


Figure 186: Locate the Controller Firmware File

3. After you locate the file, click **OK**.

   MegaRAID Storage Manager software displays the version of the existing firmware and the version of the new firmware file.

4. 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.

5. 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 MegaRAID Storage Manager (MSM) 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 MegaRAID Storage Manager (MSM) 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.

NOTE: 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
- MegaRAID RAID 6
- MegaRAID RAID 5

11.2.1 MegaRAID Software Licensing

MegaRAID Software licensing authorizes you to enable the MegaRAID advanced software features present in the MSM 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 view** or the **Logical view** tab in the left panel of the MegaRAID Storage Manager window, and click a controller icon.
2. Choose either one of the following options:

- Select **Go To>Controller >Manage MegaRAID Advanced Software Options**,
- Click **Manage MegaRAID Advanced Software Options** from dash board under feature portlet.

The **Manage MegaRAID Advanced Software Options** wizard appears as shown in the following figure.

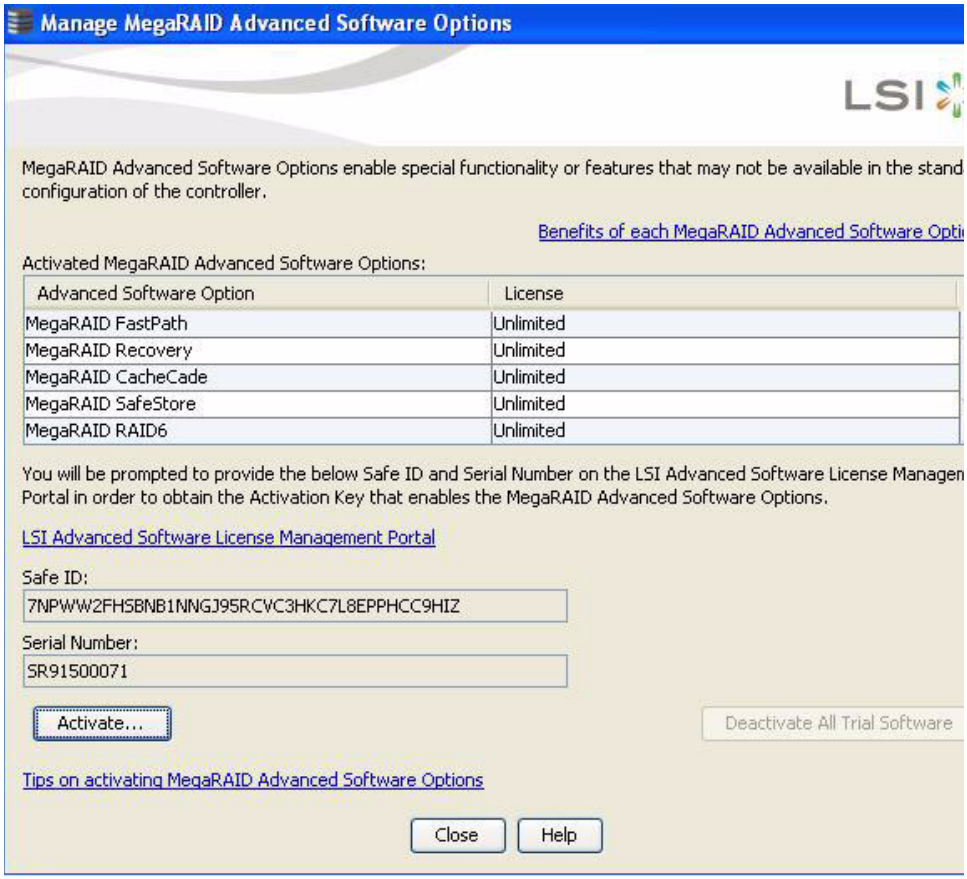


Figure 187: Manage MegaRAID Advanced Software Options

The **Activated MegaRAID Advanced Software** fields consists of **Advanced Software**, **License**, and **Mode** columns.

- The **Advanced Software** column displays the list of advanced softwares present in the controller.
- The **License** column displays the license details for the list of advanced softwares present in the **Advanced Software** 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 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

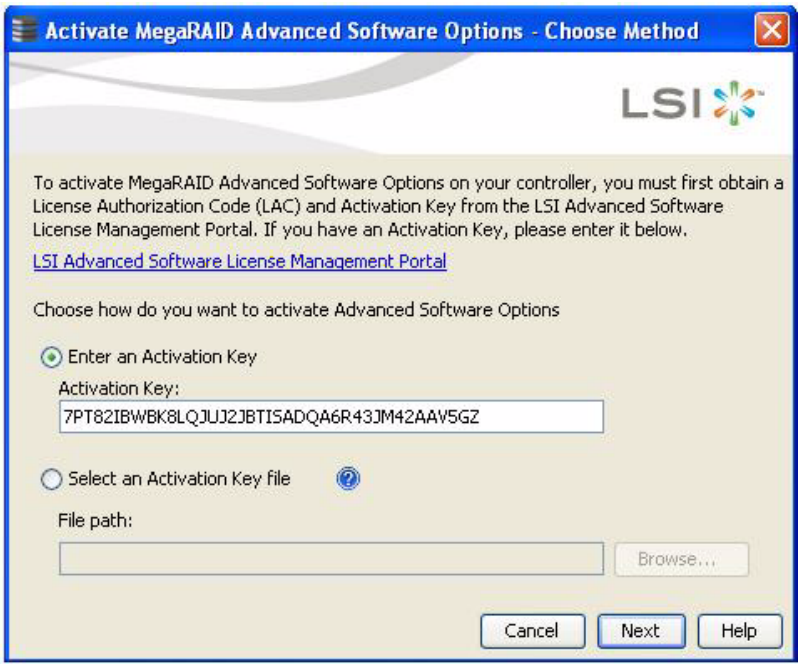


Figure 188: Activate MegaRAID Advanced Software – Choose Method

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:

- Select the **Enter an Activation Key** option, and enter the *activation key* in the text box provided below the **Activation Key** field.
- Select the **Select an Activation Key file,** and click **Browse** to get the path of the *activation Key* file.

3. Click **Next**.

   After you click **Next**, the **Activate MegaRAID Advanced Software – Summary** wizard is displayed as shown in Figure 189.

   ELSE

   Based on some 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** as shown in Figure 188, the **Activate MegaRAID Advanced Software – Summary** wizard displays the list of the advanced softwares along with their *former status* and *new status* in the controller.

- The **Advanced Software** 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.

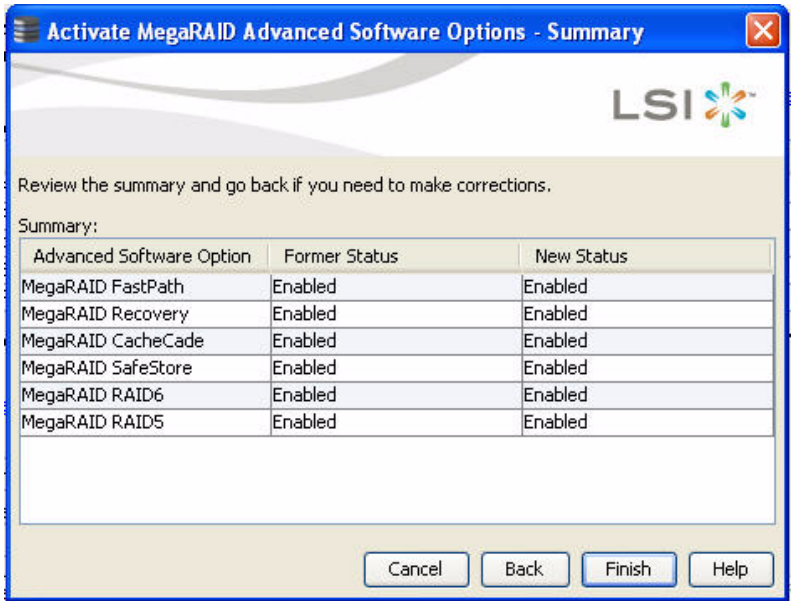


Figure 189: Activate MegaRAID Advanced Software – Summary

4. Click **Finish**.

   The status of the advanced software is enabled, and the advanced features will be secured in the Key Vault.

5. Click **Cancel** to cancel this action.

11.2.5 Application Scenarios and Messages

- Scenario # 1

  If you enter an *invalid* activation key, the following message is displayed.



Figure 190: Invalid Activation Key

- Scenario # 2

  If you enter an *incorrect* activation key file, the following message is displayed.

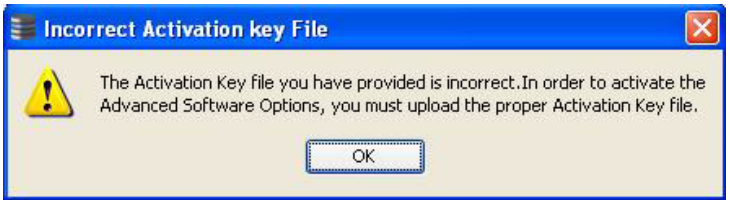


Figure 191: Incorrect Activation Key

- 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 is displayed.



Figure 192: Activation Key Mismatch

NOTE: Entering a space in the **Activation Key** field disables the *Next* button in Figure 188.

If you click **Cancel** in the **Activate MegaRAID Advanced Software – Choose Method,** as shown in Figure 188, the following confirmation dialog box appears.

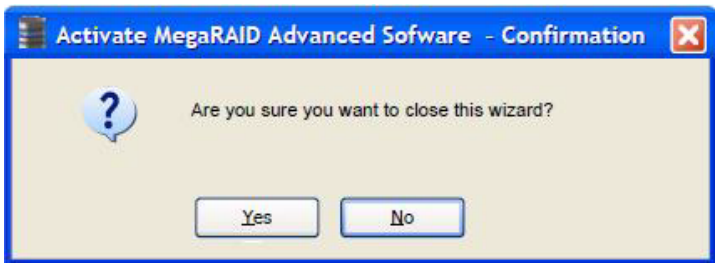


Figure 193: Confirmation Dialog Box

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*" is displayed (indicated in pink text in the following figure).

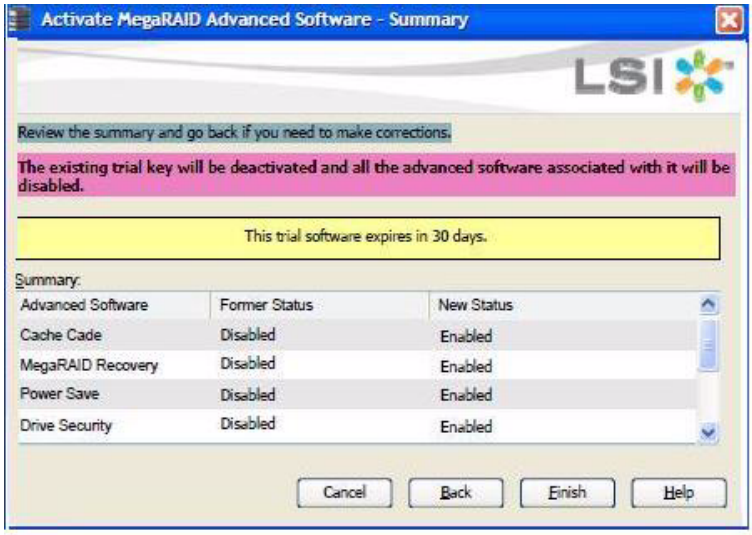


Figure 194: Activating an Unlimited Key over a Trial Key

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).

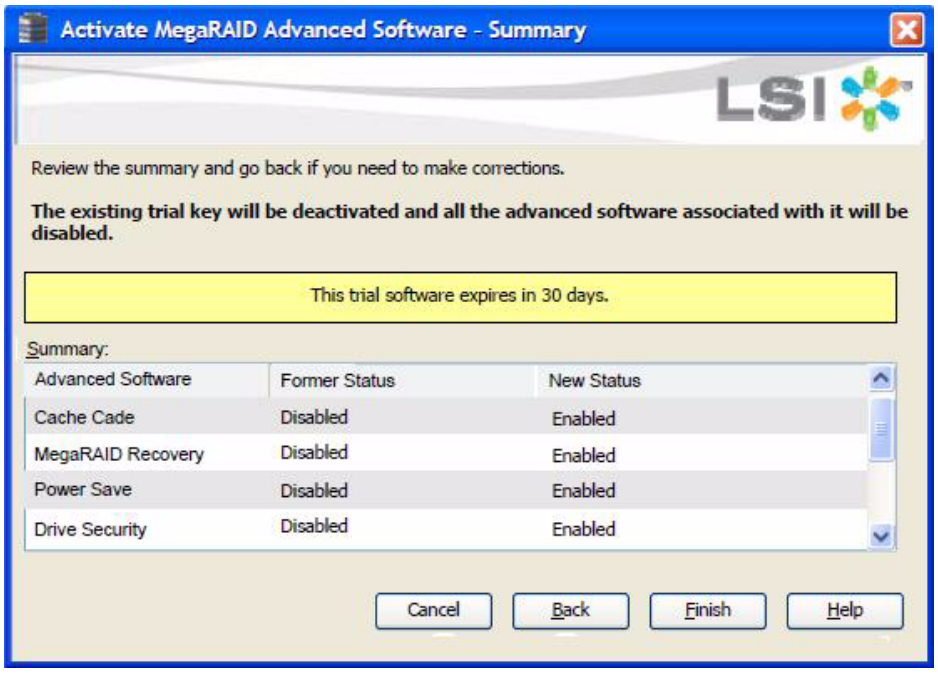


Figure 195: Activating a Trial Software

11.2.6.2 Activating an Unlimited Key

When you activate an unlimited key or a trial key, a message "*Review the summary and go back if you need to make corrections*" is displayed as shown in the following figure.

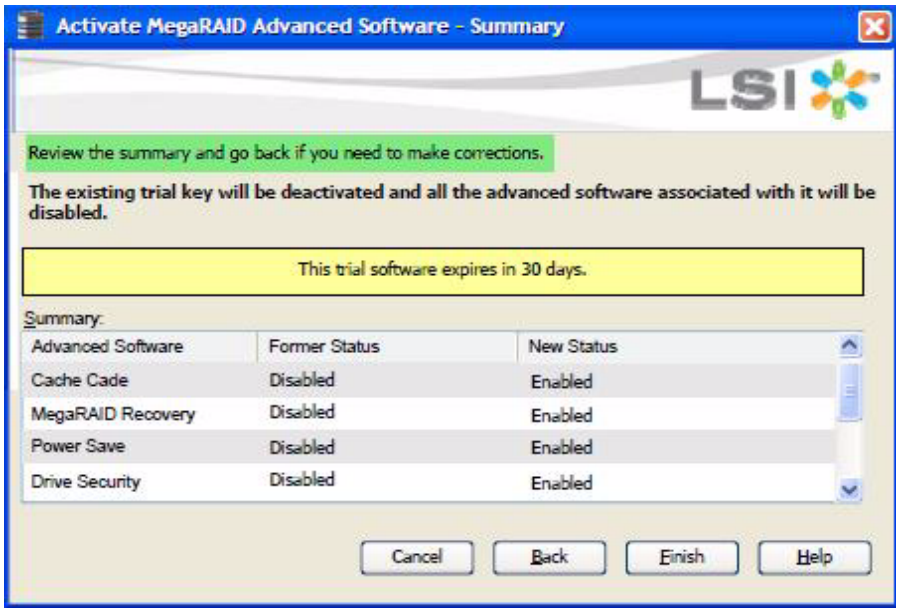


Figure 196: Activating an Unlimited Key

11.2.6.3 Reusing the Activation Key

If you are using an existing activated key, the features have been transferred to the key vault, and the message shown in the following figure is displayed.



Figure 197: Reusing the Activation Key

11.2.6.4 Securing Advanced MegaRAID Software

When you want to transfer the advanced software from the controller to Key vault use the **Securing Advanced MegaRAID Software - Confirmation** wizard. This is a conditional wizard, and will appear only when there is a key vault present and unsecured keys.

Use any one of the following options to view the **Securing Advanced MegaRAID Software - Confirmation** wizard**.**

- Select the **Physical View** tab in the left panel of the MegaRAID Storage Manager main menu, and click a controller icon.
- Select **Go To>Controller>Manage MegaRAID Advanced Software Options** wizard.



Figure 198: Securing Advanced MegaRAID Software

1. Select the check box next to the **Confirm** field, if you want to secure the advanced software.

   After you select the check box, the *Yes* button gets enabled. This implies that the advanced software is secured in the key vault.

If the advanced software is not secured, the **Secure MegaRAID Advanced Software - Confirmation** dialog box appears as shown in Figure 193.

11.2.7 Configuring Key Vault (Re-hosting process)

Rehosting is a process of transferring the advanced software features from one controller to another. To implement the rehosting process, you have to configure the **Configure Key Vault** button in the **Manage MegaRAID Software Options** wizard as shown in Figure 187.

Use 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 from Figure 187.
- Select **Go To>Controller>Manage Premium Feature.**

  The **Configure Key Vault-Confirm Re-Hosting Process** wizard appears as shown in the following figure.

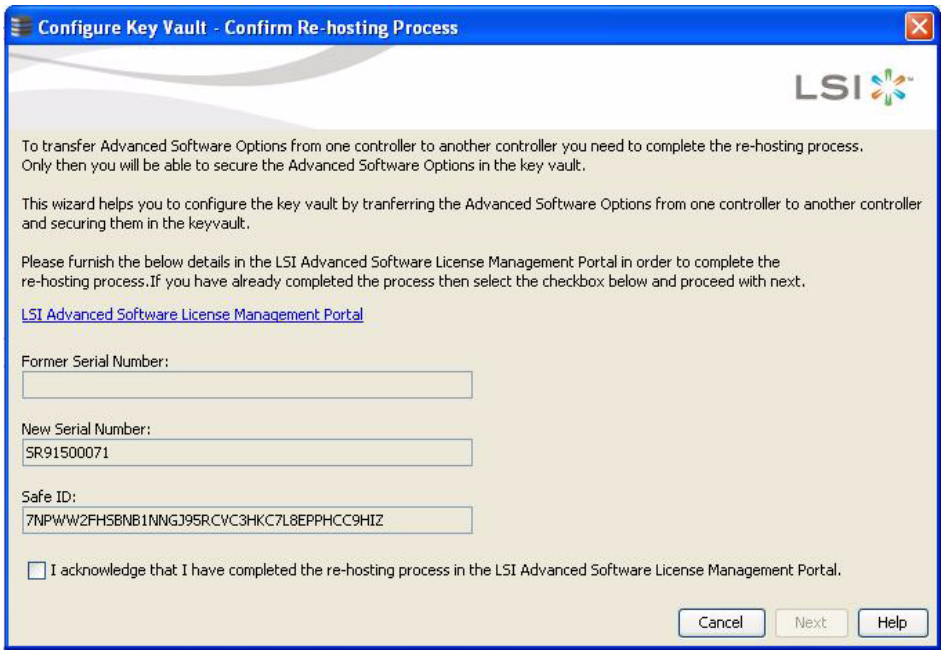


Figure 199: Configure Key Vault

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

   The **Configure key Vault- Secure Advanced Software** Options wizard appears as shown in the following figure.

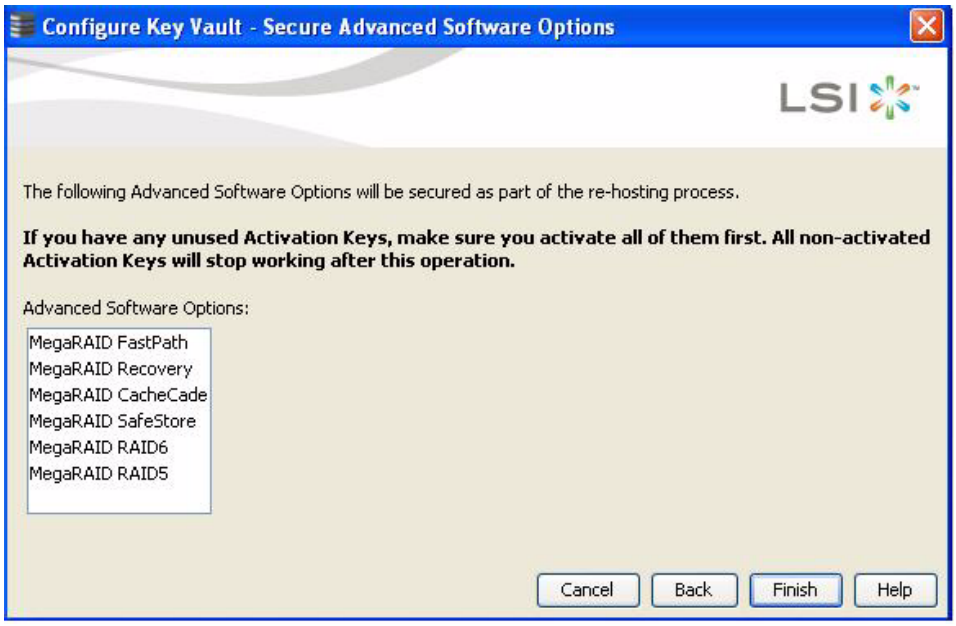


Figure 200: Configure Key Vault - Secure Advanced Software Options

4. Click **Finish**.

   The advanced software options are secured in the key vault.

NOTE: The **Next** button in the Configure Key Vault wizard (See figure 194) will be enabled only when you check the box.

NOTE: This wizard is conditional and will pop-up only, if the rehosting process is necessary when the key vault is present and unsecured keys are also present, both at the same time.

11.2.8 Rehosting Complete

If you want to transfer the advanced software options from one controller to another use the rehosting process. The rehosting process ensures that these options are secured in the key vault. You have to configure the key vault in order to complete the rehosting process.

Use any one of the following options to complete the rehosting 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.

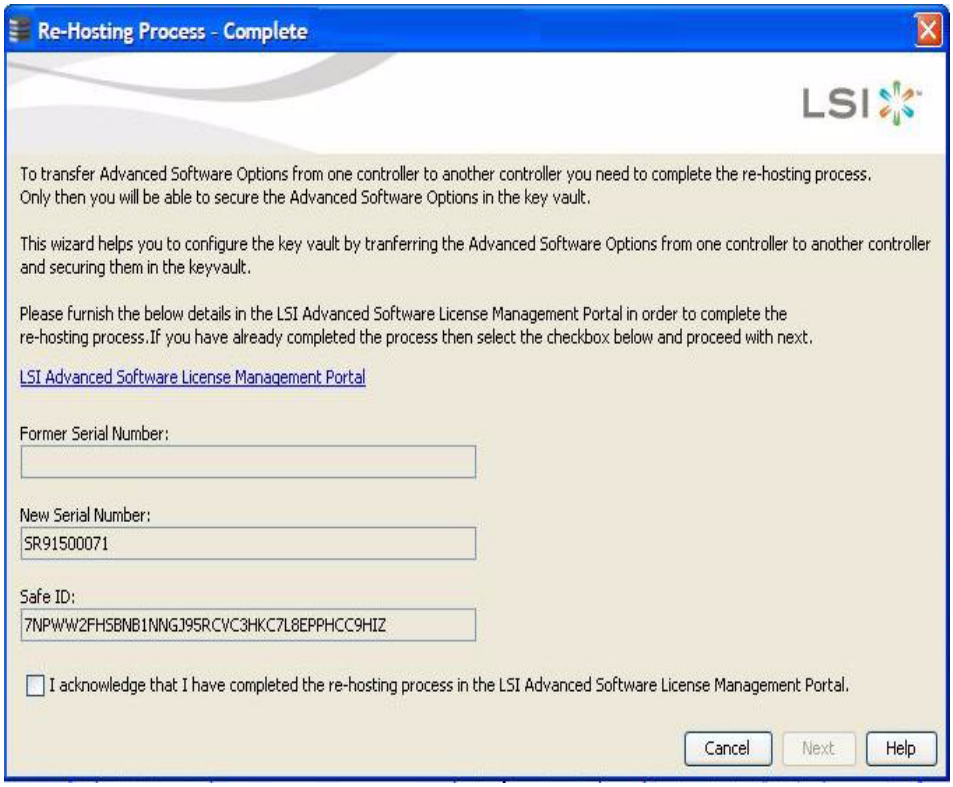


Figure 201: Confirm Rehosting

5. Select the check box next to the message "*I acknowledge that I have completed the re-hosting process in the LSI Advanced Software License Management Portal*" if you want to complete the rehosting process.

   This ensures that the advanced software features are transferred to the controller.

6. Click **Cancel** if you do not want to complete the re-hosting process.

NOTE: Click the LSI Advanced Software License Management Portal to view the www.softwarelicense.lsi.com which will enable you to complete 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 the **Deactivate All Trial Software button** in the **Manage MegaRAID Advanced Software Options** wizard as shown in Figure 187.

   The Deactivate All Trial Software - Confirmation dialog box appears as shown in the following figure.

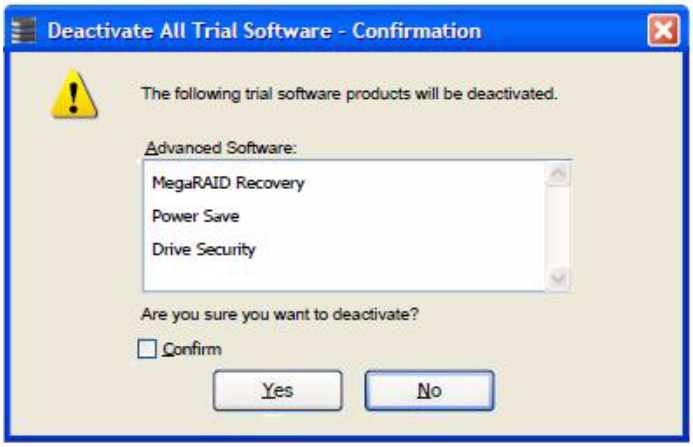


Figure 202: Deactivate All Trial Software

2. Select the check box next to the **Confirm** field, if you want to deactivate the software which are being used with a trial key.
3. Click **Yes** button.

   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 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.

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" af 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 screen.
   c. Drag and drop the earlier version of the file from 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 POST (Power-on Self Test).

b. In the WebBIOS screen, select the corrupted virtual drive and on the next screen that appears, select the option **Adv Opers**.
c. Select **Rollback** and designate the most recent PiT from the drop-down list.
d. Click **Go** and then exit WebBIOS.

   The system reboots.
e. Begin debug/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 application updates/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 MSM. 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.

Follow these steps to enable MegaRAID Recovery.

1. Select the **Logical** tab on the main menu screen 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 Figure 203.

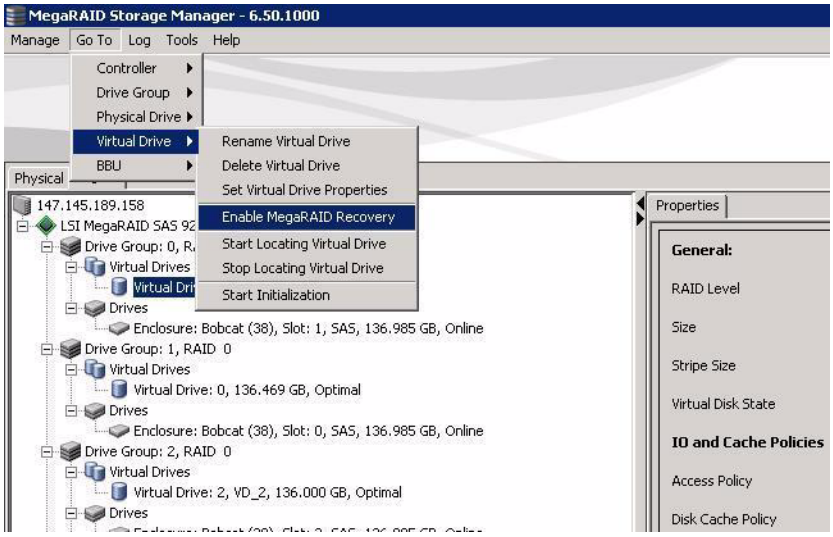


Figure 203: Enable MegaRAID Recovery Menu Option

The Enable MegaRAID Recovery Wizard appears, as shown in Figure 204. This wizard allows you to select the VD to use as the snapshot repository.

11.2.13 Snapshot Repository

You can select an existing VD, or create a new VD to use as the snapshot repository.

Use any one of the options to select or create a new VD.

- Select the virtual drive to use as the Snapshot Repository in the **Snapshot Repository** field from the **Enable MegaRAID Recovery Wizard** wizard as shown in Figure 204.
- Click the **Create New** button to create a new virtual drive to use as the Snapshot Repository.

When you create a new virtual drive, the newly created virtual drive gets appended to the existing rows in the **Snapshot Repository** field.

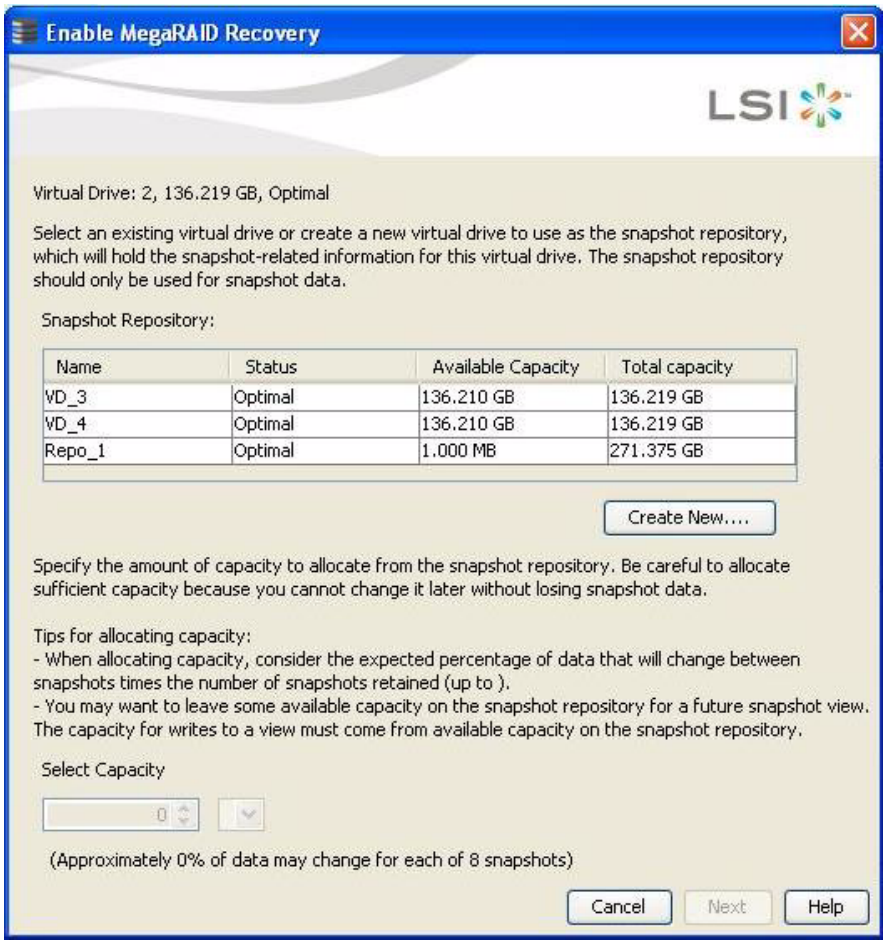


Figure 204: Enable MegaRAID Recovery Wizard

The **Snapshot Repository** field displays the details of the default VDs existing in the system, if there are any.

- The **Name** column displays the name of the VD.
- The **Status** column displays the status of the VD.
- The **Available Capacity** column displays the available capacity on the VD.
- The **Total Capacity** displays the total capacity of the VD.

If the default VDs do not exist in the system the **Snapshot Repository** field columns will be blank.

The status of the VD can be optimal, or degraded, or partially degraded.

4. In the **Select Capacity** field, use the drop-down selector to select the appropriate capacity to use for changes to the base virtual drive.

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

NOTE: Please refer to the tips provided for allocating capacity in Figure 204.

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 wizard- Displaying the selected Virtual Drive** appears as shown in Figure 205. This wizard allows you to select the VD to be used as the snapshot repository.

11.2.14 Selecting the Virtual Drive

You can select the VD to use as the snapshot repository, and also allocate the capacity for the VD from the snapshot repository.

Follow these steps to select the VD in the **Snapshot Repository** field.

1. Select the virtual drive to be used as the snapshot repository to hold snap-shot information. The selected virtual drive gets highlighted as shown in the following figure.

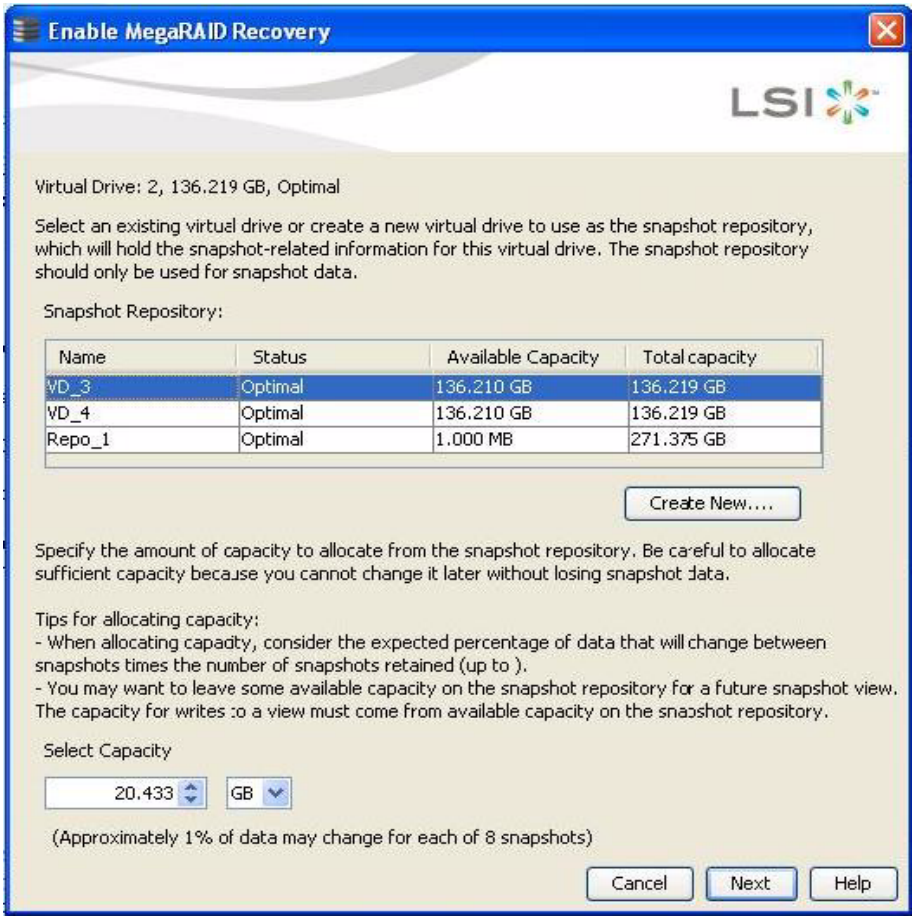


Figure 205: Enable MegaRAID Recovery Wizard- Displaying the selected Virtual Drive

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 wizard - Create Snapshot Schedule** wizard appears as shown in Figure 206. This wizard allows you to schedule snapshots.

11.2.15 Scheduling Snapshots

You can select an existing snapshot schedule or create a new snapshot schedule for the virtual drives.

Follow these steps to schedule snapshots.

1. Select any one of the options shown in the **Enable MegaRAID Recovery - Create Snapshot Schedule** wizard in Figure 206 to schedule snapshots using the radio button.

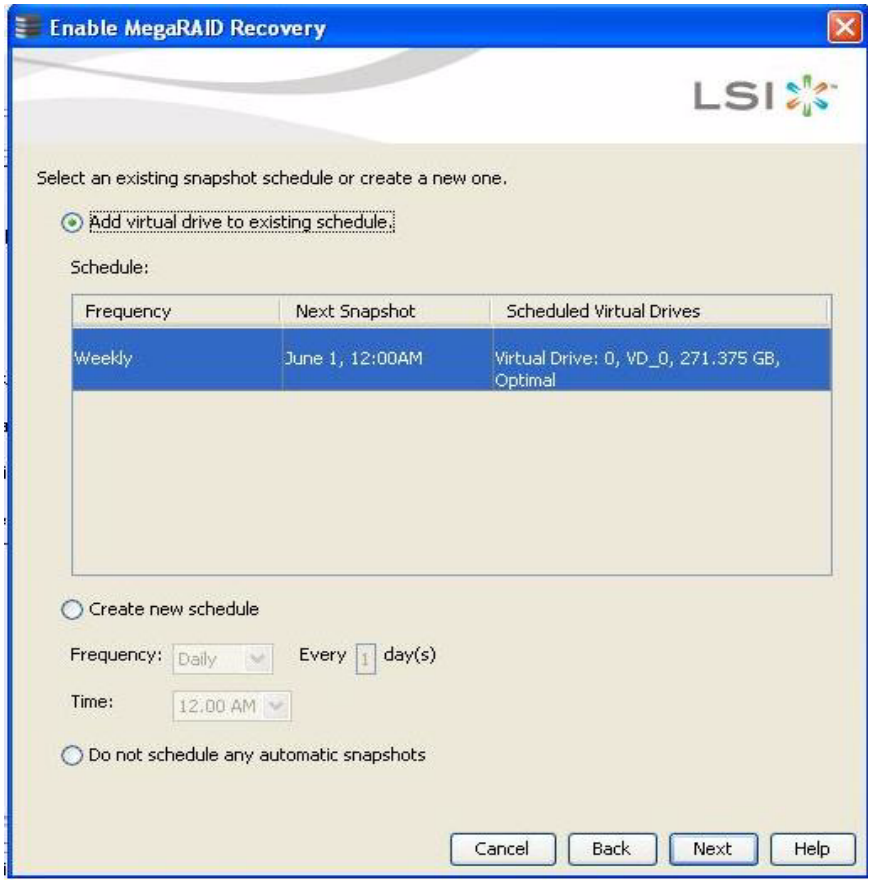


Figure 206: Enable MegaRAID Recovery - Create Snapshot Schedule

- Add virtual drive to existing schedule

  This option allows you to add a VD to a pre-existing schedule already defined in the system.

  The **Schedule** field displays the schedule details of the VD.

  — The **Frequency** column displays the frequency of the snapshot schedule. The frequency of the snapshot can be daily or weekly.

— 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 VD.

  — In the **Frequency** field, use the drop-down selector to select the frequency of the snapshot. It can be daily or weekly.

  — In the **Time** field, use the drop-down selector to select the time when you want 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 207. 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.

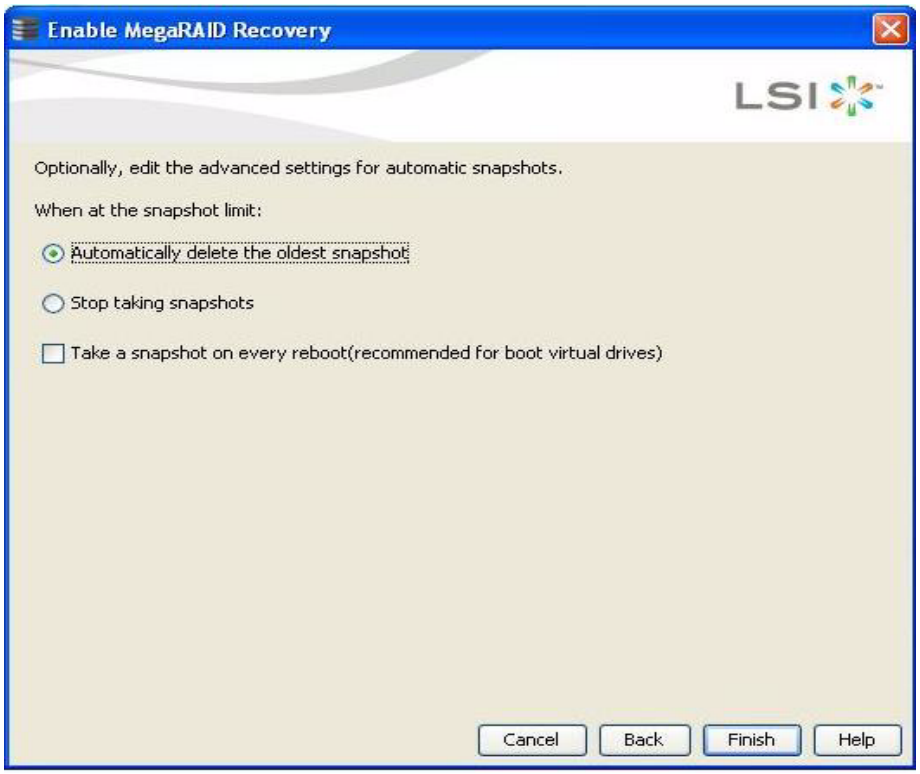


Figure 207: Enable MegaRAID Recovery- Editing Snapshot Properties

Select any one of the options shown in Figure 207 to edit snapshots.

- Automatically delete the oldest snapshot

If you choose this option, the oldest snap shot present in the system automatically get deleted.

- Stop taking snapshots

  If you choose this option, the application stops taking snapshots.

- Take a snapshot on every reboot(recommended for boot virtual drives)

  To use this option, select the check box next to the **Take a snapshot on every reboot (recommended for boot virtual drives)** field.

  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 in case the virtual drive becomes corrupted.

3. Click **Finish**.

   A **Confirm Enable Snapshots** dialog box appears. This dialog box prompts you to make sure whether you want to enable snapshots on VD or not.



Figure 208: Confirm Enable Snapshots

4. Check the **Confirm** box.

   When you select the check box, the **Yes** button gets enabled.

5. Click **Yes**.

   The snapshots get enabled on the virtual drive. 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.

11.2.17 Snapshot Base Details

You can view the snapshot base details of the virtual base drive.

Follow these steps to view the details of the snapshot of the virtual base drive.

1. Select the **Logical** tab on the main menu screen from the MegaRAID Storage Manager.
2. Click a base virtual drive in the left frame.

After you click the base virtual drive, in the right frame you can view the base virtual drive information in the **Properties tab** under **Snapshot Base Details (**highlighted in pink**).**

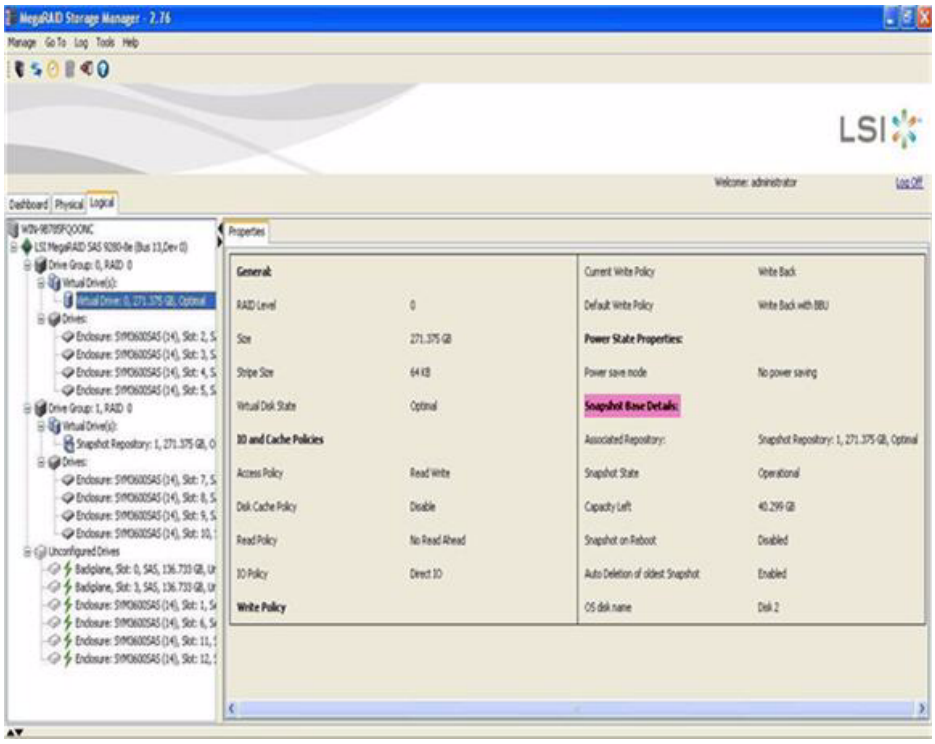


Figure 209: Snapshot Base Details

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 Figure 210.

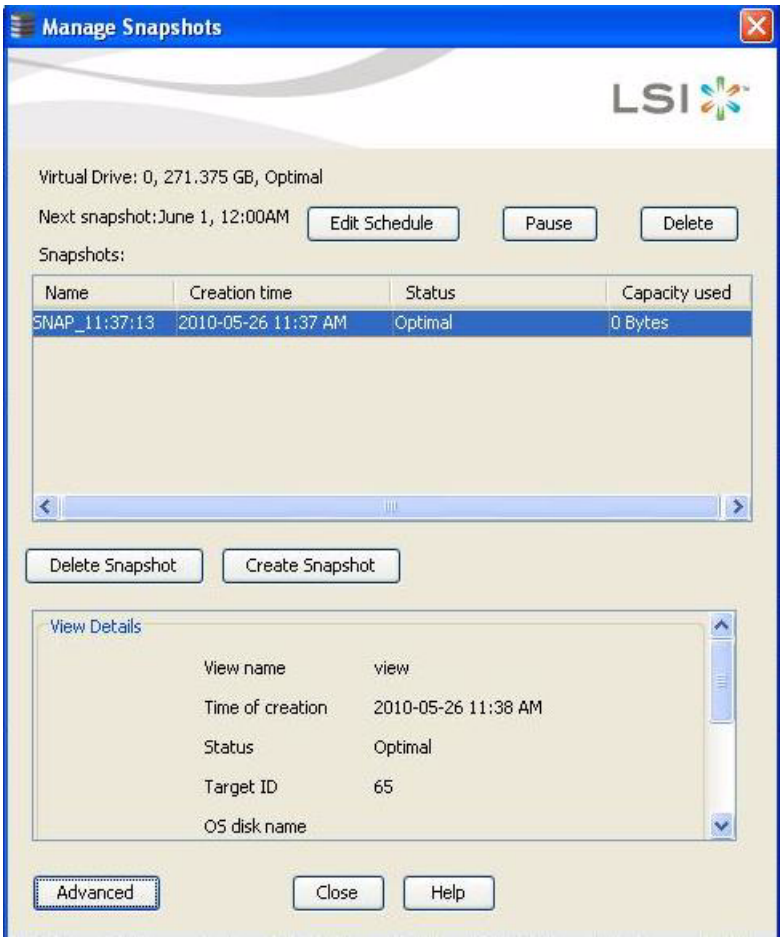


Figure 210: Manage Snapshots

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 Snapshot field 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 button using the **Create Snapshot** button, and delete the snapshot using the Delete Snapshot button.

In the **View Details** field, you can create a view using the **Create View** button, and edit the settings for automatic snapshots using the **Advanced** button.

11.2.19 Editing the Snapshot Schedule

You can edit the schedule using the **Edit Schedule** screen. 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.

Follow these steps to edit the snapshot schedule.

1. Click the **Edit Schedule** button in the Manage Snapshots wizard in Figure 210.

   The **Edit Schedule** screen appears, as shown in Figure 211.

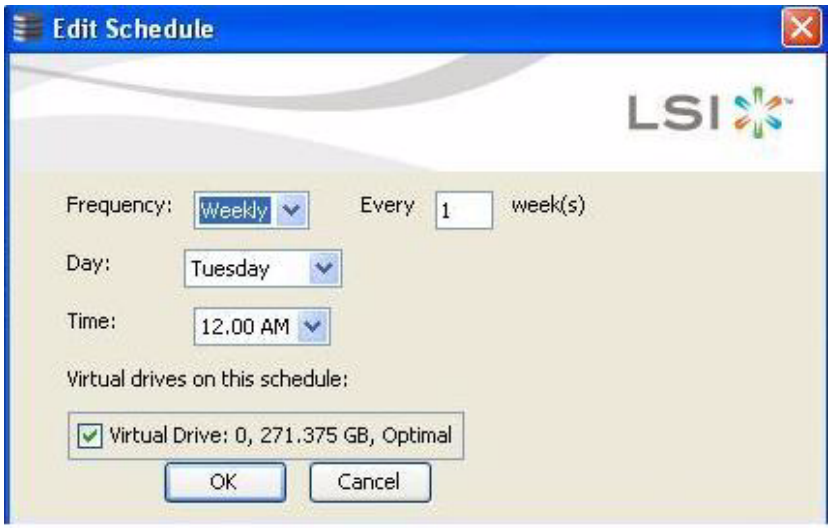


Figure 211: Edit Schedule Dialog Box

- In the **Frequency** field, use the drop-down selector to edit the frequency of the snapshot already taken. The frequency can be daily or weekly.
- In the **Day** field, use the drop-down selector 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 selector to edit the time of the snapshot already taken.

After you select all of the above fields, the virtual drives matching these fields will appear in the **Virtual drives on this schedule** field. Check the box next to this field.

2. 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 210.

   The Advanced screen appears as shown in Figure 212.



Figure 212: Advanced Settings

You can edit the settings by selecting among these options:

- **Auomatically delete the oldest snapshot** option, if you want to delete the oldest snapshot
- **Stop taking snapshots** option, if you want to stop taking snapshots
- **Take a snapshot on every reboot (recommended for boot virtual drive)** option, if you want a snapshot on every reboot.

2. Click **OK**.

   The settings are edited.

11.2.21 Create View

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 VD available at that particular time.

Follow these steps to create views of the snapshots.

1. Click **Create View** button in the **Manage Snapshots** wizard in as shown in Figure 210.

   The **Create View** screen appears as shown in Figure 213.



Figure 213: Create View

2. In the **View name** field, enter the view name. For example, view.
3. In the **Write capacity** field, use the drop-down selector 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 VD, the view details appear under the view details in the **Manage snapshots** wizard as shown in Figure 214.

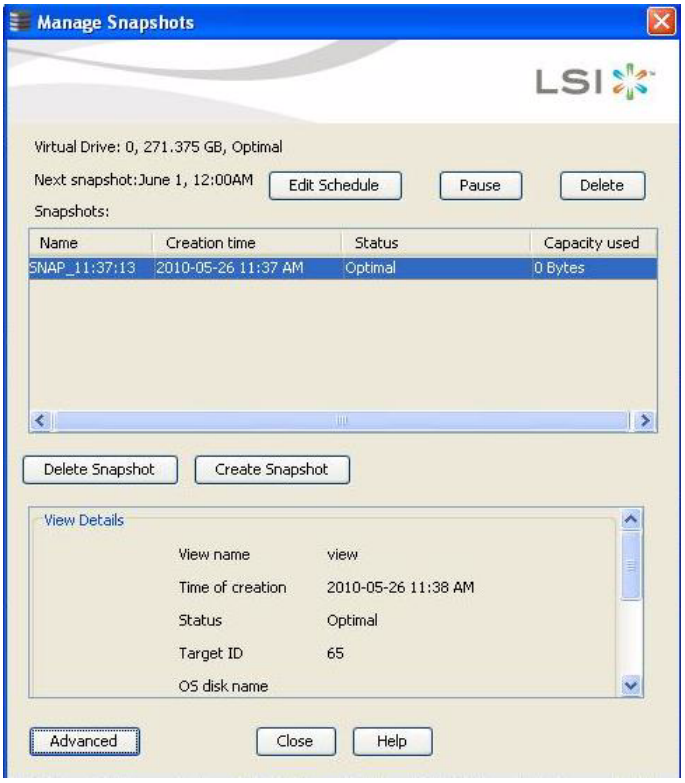


Figure 214: View Details

11.2.23 No View Details for Snapshot

When there is no view for the snapshot, the following message appears in the **View details** area as shown in Figure 215.

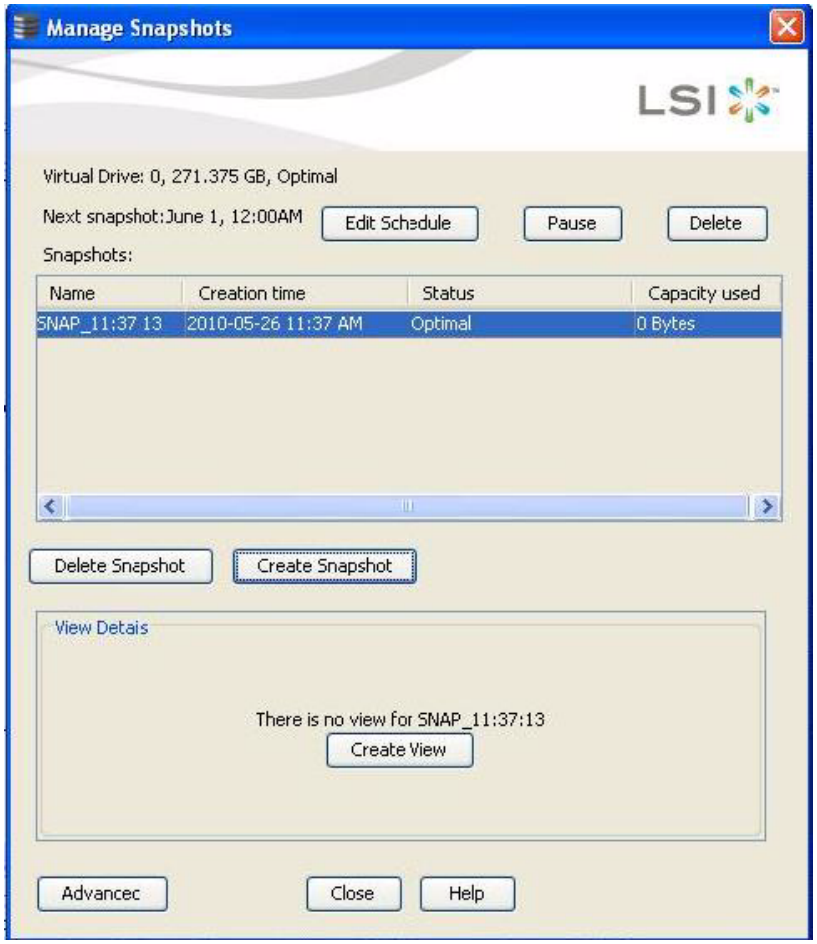


Figure 215: No View Present for the Snapshot

11.2.24 No Snapshot Schedule

When there are virtual drives with no snapshot schedule, the following message appears as shown in Figure 216.

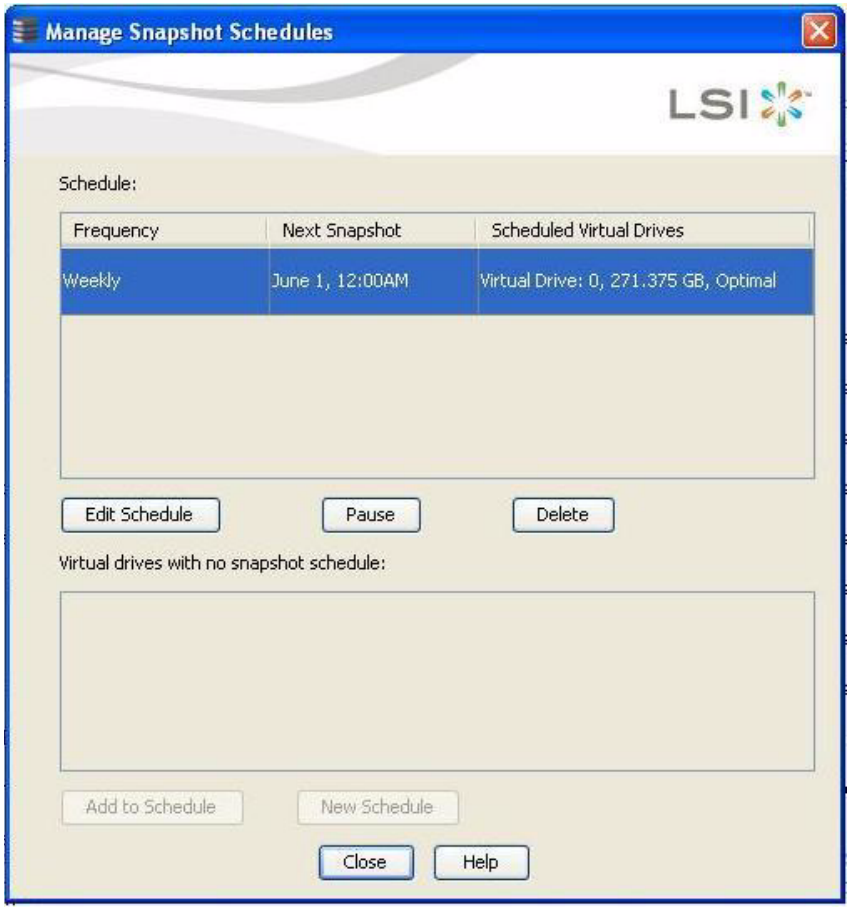


Figure 216: Virtual Drives with No Snapshot Schedule

- Click **Add to Schedule** to add a snapshot schedule.
- Click **New Schedule** to add a new snapshot schedule.

11.2.25 Creating Views

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.

Follow these steps to create views of the snapshots.

1. Click the **Logical** view on the main menu screen.
2. Click the **Snapshot Base** virtual drive in the left frame.
3. Click the **Snapshots** tab in the right frame.

   This screen shows the Snapshot Base details and any existing snapshots.
4. Click **Create View** in the right frame.

   The Create View screen appears, as shown in Figure 217.

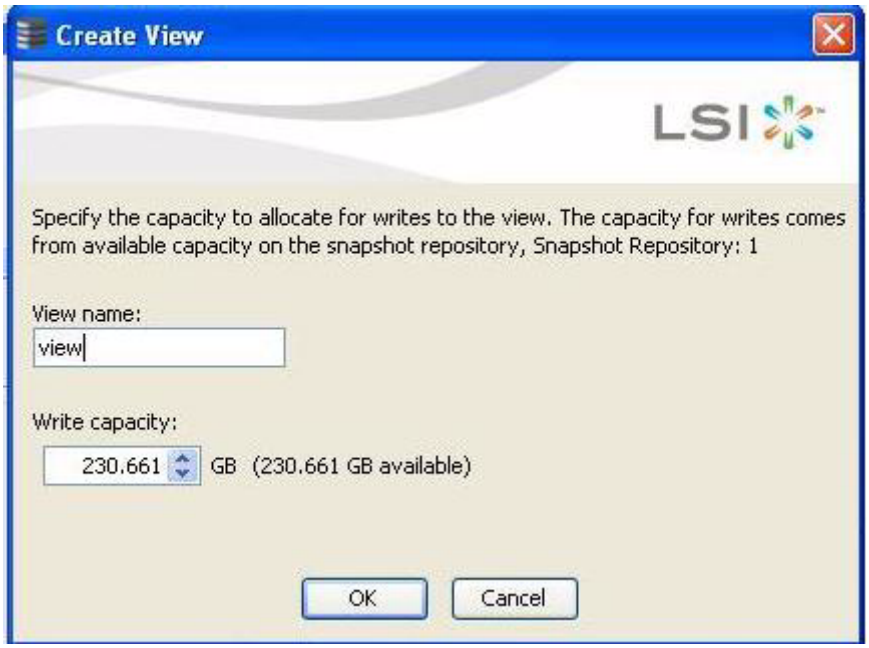


Figure 217: Create View Screen

5. Enter the name of the view in the **Enter View name** field
6. Enter the capacity in the **Available Capacity** field to set aside in the snapshot.

   This is a percentage of the Repository virtual drive capacity.
7. Click **OK**.

   This creates the view of the Point-in-Time snapshot of the volume.

11.2.26 Restoring 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 that snapshot.

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 by pressing Ctrl+H during POST.
2. Click the **Snapshot Base** virtual drive in the right frame.
3. Select **Adv Opers** at the bottom of the screen.
4. Select Rollback and then designate a snapshot PiT from the drop-down list.
5. Click **Go**.
6. Exit WebBIOS and reboot the system.

NOTE: If the operating system is corrupted and cannot boot, go to Section 4.10.7, *Restoring a Virtual Drive by Rolling Back to a Snapshot*, in the WebBIOS Configuration Utility.

11.2.27 Restoring from a View

After you discover that a file, record, or file system is missing or deleted, you can restoring the data from a view. You can mount any PiT to create an instant view of the data at a previous point-in-time. Then you can restore the data in MSM and drag-and-drop the lost data and/or files back into the source data virtual drive volume.

Follow these steps to restore a file from view.

1. Mount a view of the file content from each PiT until you find the missing file.
2. Drag and drop the missing file from Snapshot view back into the online storage volume that was the source of the Snapshot.

11.2.28 Deleting a Snapshot

NOTE: You can delete only the oldest snapshot.

Follow these steps to delete a snapshot.

1. Select the **Logical** tab on the main menu screen for the Logical view.
2. Click the **Snapshot Base** virtual drive in the left frame.
3. Click the **Snapshots** tab in the right frame.

   This screen 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 Join Mirror

The Join Mirror feature uses the mirroring concept implemented by the MSM application.

Perform the following steps to implement the Join Mirror feature.

1. Use any one of the following options to view the Join Mirror wizard.

- Right-click on the drive group from the device tree
- Select **Go To>Drive Group>Join Mirror** in the menu bar.

The Join Mirror wizard appears, as shown in the following figure.

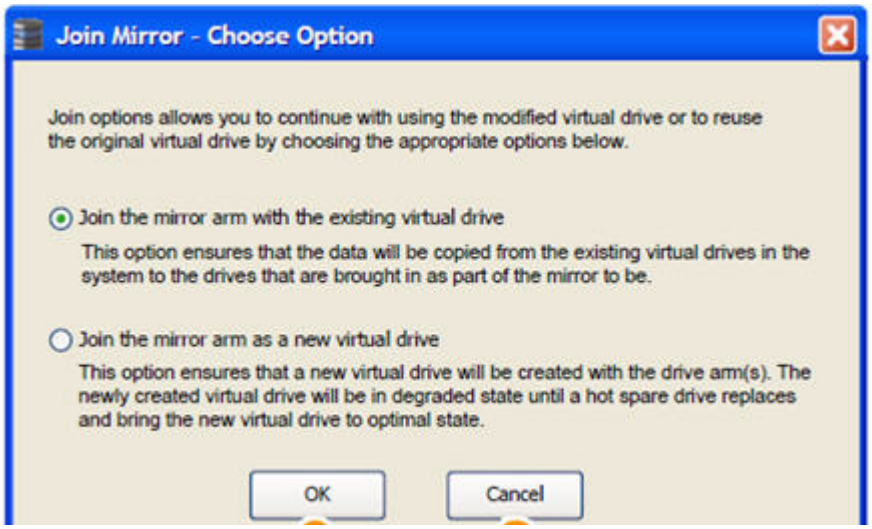


Figure 218: Join Mirror Wizard

TheJoin mirror wizard allows you to choose any of the options based on your requirement.

— Option # 1

2. Select the *Join the mirror arm with the existing virtual drive using* the radio button.

   If you choose this option the data will be copied from the existing virtual drive in the system to the drives that are brought in as part of the mirror to be.

3. Click **OK**.

   The **Join the Mirror with Existing Virtual Drive - Confirmation** dialog box appears, as shown in the following figure.

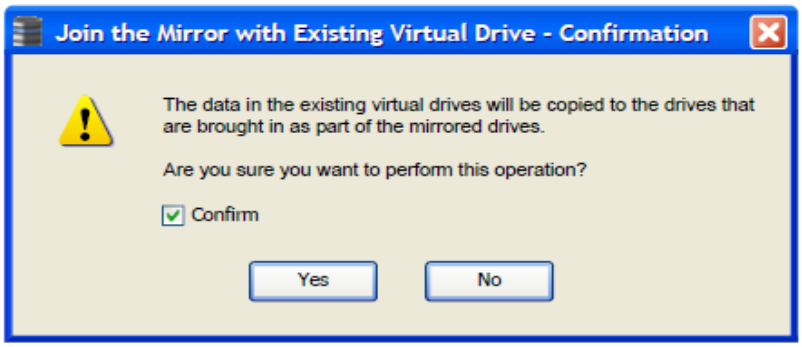


Figure 219: Join Mirror Confirmation

4. Select the check box next to the **Confirm** field if you want the data in the existing virtual drives to be copied to the drives that are part of the mirrored drives.
5. Click **Yes**.

   — Option # 2

6. Select the Join the mirror arm as a new virtual drive.

   If you choose this option, a new virtual drive is created with the drive arms.

7. Click **OK**.

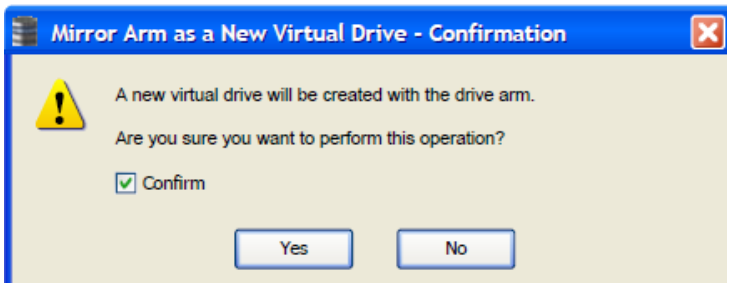


Figure 220: Mirror Arm Confirmation

The Mirror Arm as a New Virtual Drive - Confirmation dialog box appears.

8. Select the check box next to the Confirm field if you want to create a new virtual virtual drive with the drive arm.
9. Click **Yes**.

11.4 CacheCade Advanced Software

MegaRAID CacheCade improves application performance by expanding the MegaRAID read-caching capacity. 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 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 Advanced Software

Perform the following steps to use the CacheCade advanced software.

1. Click a RAID controller icon in the left frame.
2. Click **Go To>Controller>Create CacheCade** on the menu bar, as shown in Figure 221.

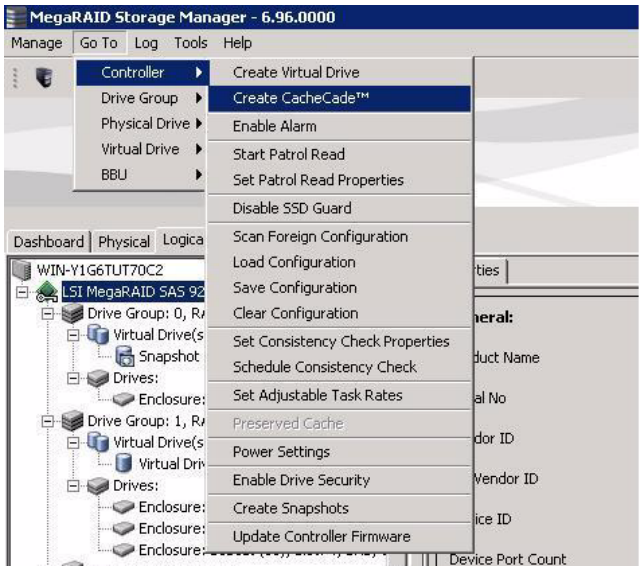


Figure 221: Create CacheCade Menu Option

The Wizard screen appears.

3. Click on unconfigured CacheCade software drives in the left frame to select the drives for the CacheCade drive group, as shown in Figure 222.

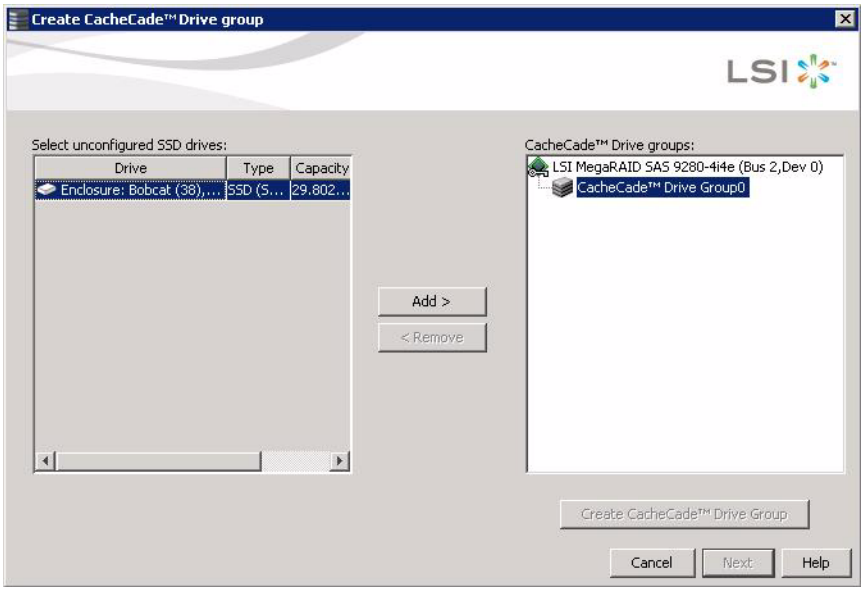


Figure 222: CacheCade Wizard Screen

After you select the unconfigured drives, the **Add >** button is available.

4. Click **Add >** to move the selected drives to the drive group in the right frame, as shown in Figure 223.

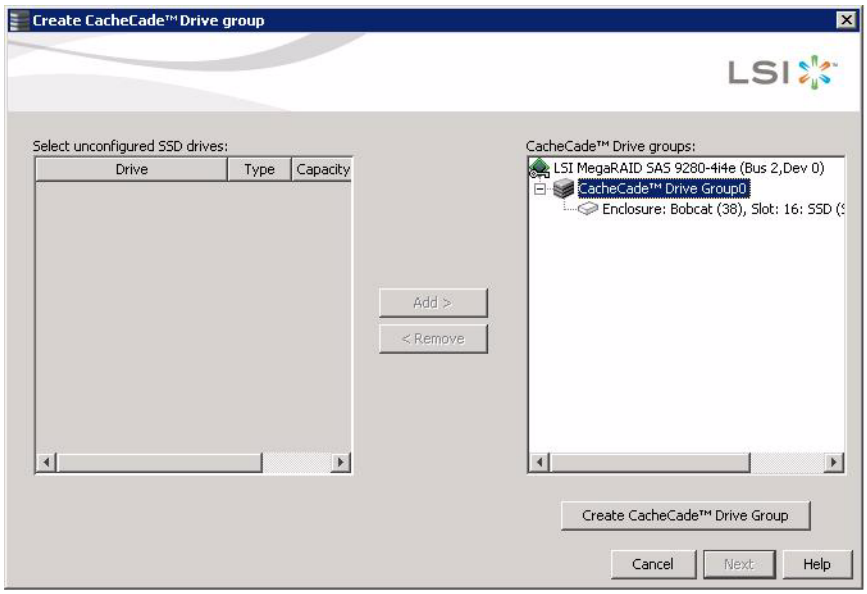


Figure 223: CacheCade Drive Group Screen

After you move the selected drives, the **Create CacheCade™ Drive Group** button is available.

5. Click **Create CacheCade™ Drive Group**.
6. Click **Next**.

   Use the next screen that appears to select parameters for the cache disk.

7. Enter a name for the CacheCade virtual drive in the **CacheCade VD name** field and click **Create**.

   Depending on the number of drives, you might have the option to set the capacity of the CacheCade drive.

   The CacheCade drive group icon appears in the menu screen, as shown in Figure 224.

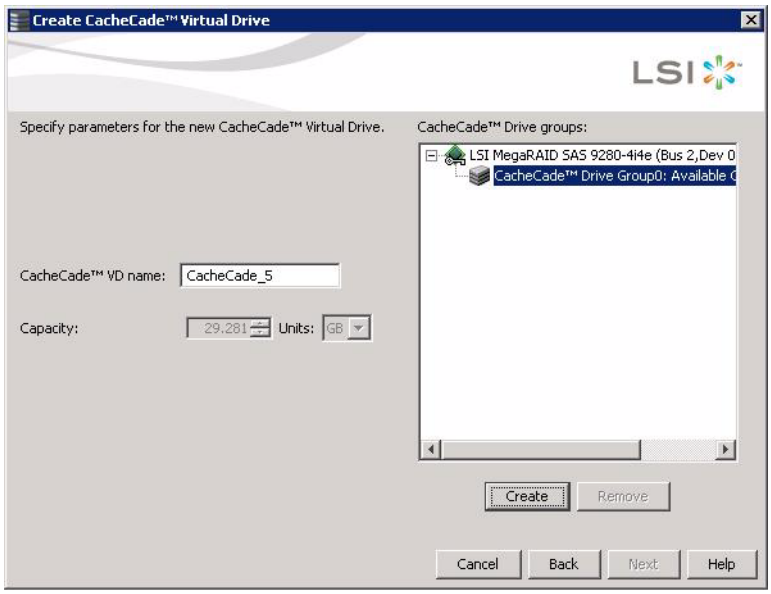


Figure 224: CacheCade Drive Group Icon

8. Click **Next**.

   The summary screen appears, as shown in Figure 225. This screen 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.

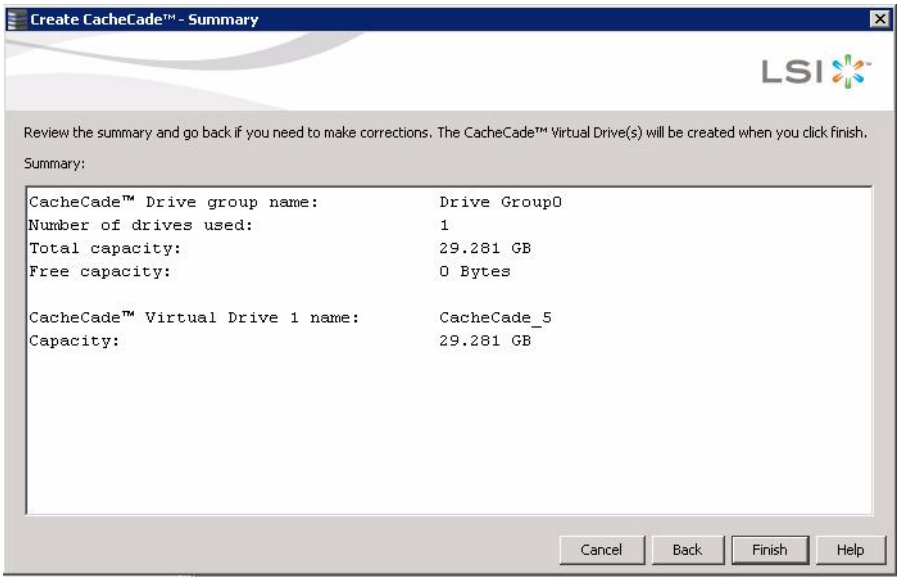


Figure 225: CacheCade Virtual Drive Summary Screen

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 Figure 226.

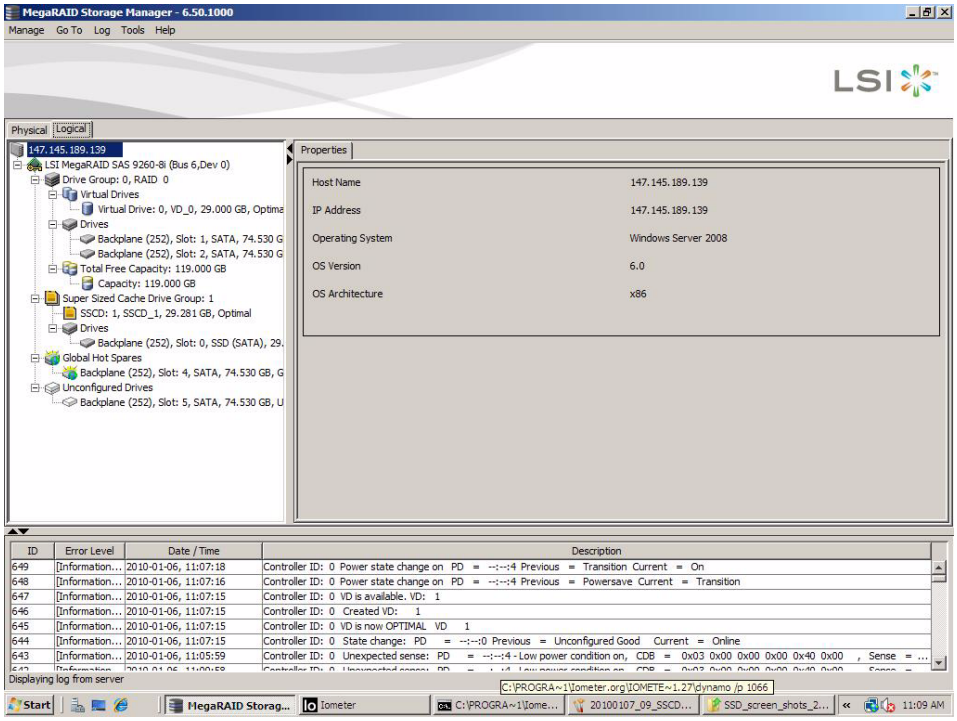


Figure 226: CacheCade Drive Icon in MegaRAID Storage Manager

11.5 FastPath Advanced Software

MegaRAID FastPath is a high-performance IO accelerator for CacheCade software drive groups connected to a MegaRAID controller card. CacheCade software has a read performance advantage over HDDs and use 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 CacheCade 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 utilizing 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 main menu screen for the Logical view.
2. Click a virtual drive icon in the left frame.
3. Click **Virtual Drive>Set Virtual Drive Properties** on the menu bar, as shown in Figure 227.

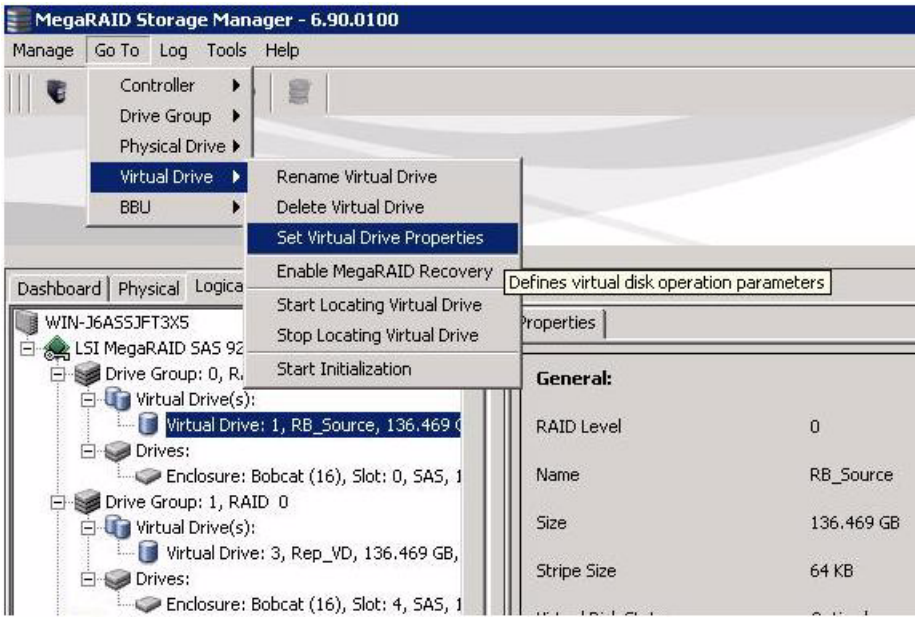


Figure 227: Set Virtual Drive Properties Menu

The **Set Virtual Drive Properties** screen appears. It shows the default settings for the FastPath advanced software. The default settings are:

- Write Policy: Write Thru
- IO Policy: Direct IO
- Read Policy: No Read Ahead
- Dish Cache Policy: Enabled

4. Click **OK**.

   A confirmation screen displays.

5. Click the **Confirm** checkbox and then 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 drives and use 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 it is housed in, the data on that drive is encrypted and useless to anyone who attempts to access without the appropriate security authorization.

This section describes how to enable, change, and disable 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 the **External Key Management** (EKM) and the **Local Key Management** (LKM).

- *Enabling Drive Security using EKM*
- *Enabling Drive Security using 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 also use to unlock configurations.

Yet another important feature of EKM is that you can use EKM 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 VDs
- 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.

You can perform the following configurations to enable the drive security to create secure VD using the EKM mode with the support of EKM servers.

- EKM mode is supported by MSM and EKMS is present
- EKM mode is supported by MSM and EKMS is not present
- Change current security settings or switch between modes
- Change 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 when **EKM mode is supported, and EKMS is present**, the application responds to different behaviors based on the scenarios that take place 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 when **EKM mode is supported, and EKMS is present.** To configure this option, you have to select **EKM** from the **Drive Security Choose** Mode wizard. To arrive at this wizard use any of the following options:

1. Select the **Physical View** tab in the left panel of the MegaRAID Storage Manager window, and click a controller icon.

- Select **Go To>Controller>Enable Drive Security** in the main navigation bar which is present in the top portion of the MSM screen.
- Right - click on the controller icon, and click **Enable Drive Security** menu.The **Drive Security Choose Mode** wizard appears as shown in the following figure.

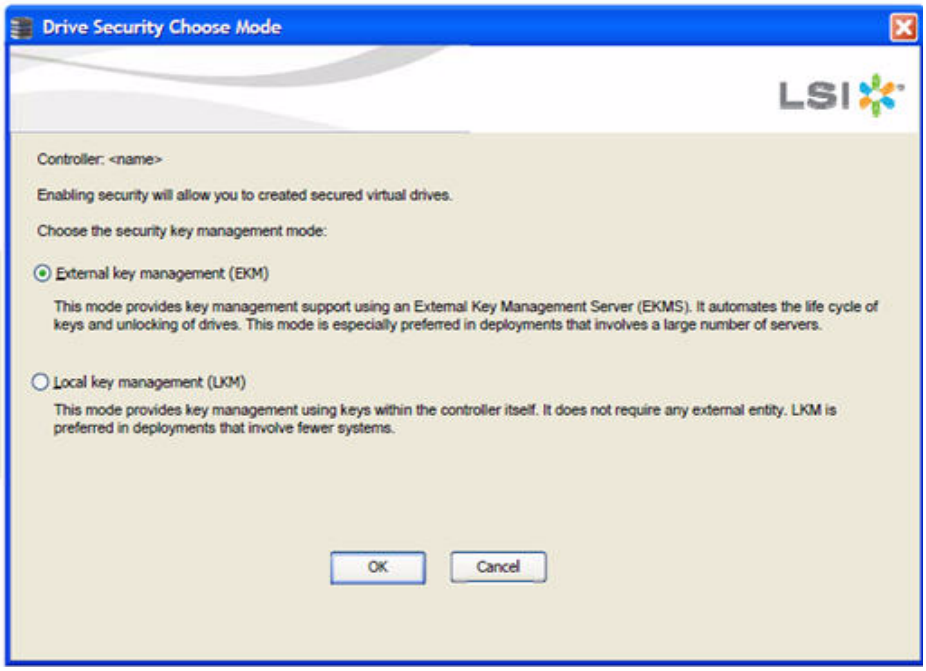


Figure 228: Drive Security Choose Mode

2. Select **External Key management (EKM)**.
3. 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 ensures that security is enabled on the controller using EKMS.



Figure 229: Confirm External Key Management Mode Enabled

NOTE: On the right hand side frame of the controller properties, that is, in **Drive Security Properties**, you can see EKM supported is 'Yes'.

- Scenario # 2

  EKM is selected, and EKMS is not present

  When EKM is selected and EKMS is not present, you have to restart the system. When you restart the system, the system restart message appears as shown in the following figure.

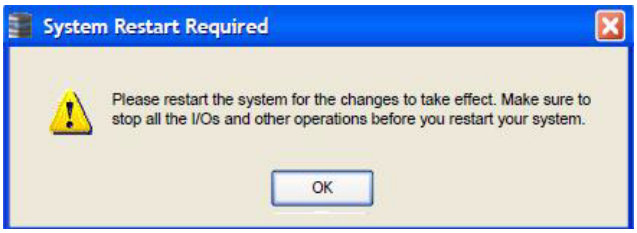


Figure 230: System Restart

When the system restarts, the boot agent generates the security keys for the controller.

NOTE: If the MSM 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 in the **Change Security wizard** as shown in Figure 231, two options are provided for you. If you select the first option, that is, **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.

Use any of the following options to enable the **Change Security Settings** wizard.

1. Select the Physical View tab in the left panel of the MegaRAID Storage Manager window, and click a controller icon.

- Select **Go To>Controller>Change Security Settings** in the main navigation bar which is present in the top portion of the MSM screen.
- Right-click on the controller icon, and click Change Security Settings. The **Change Security Settings** wizard appears as shown in the following figure.

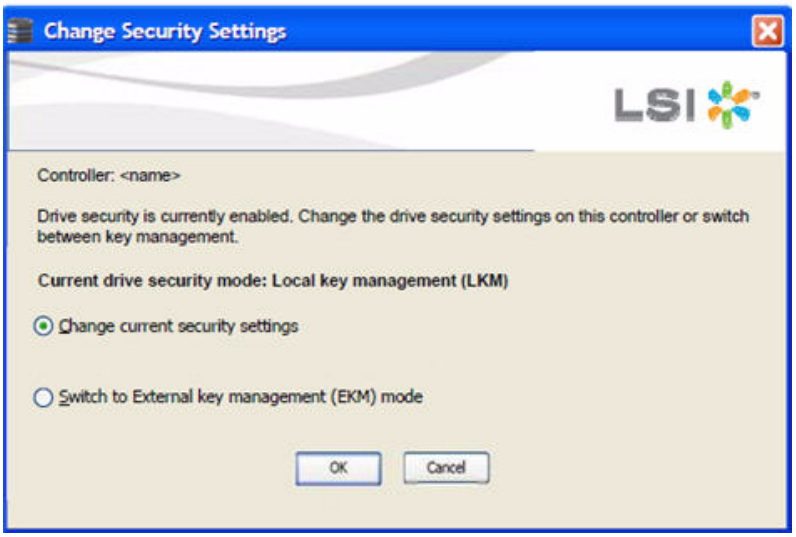


Figure 231: Change Security Settings

2. Select the **Change current security settings** option in Figure 231, if you want to change the drive security settings on the controller using the LKM mode.

   When you select this option, you can view the details of the Change current security settings in section 11.5.9 *Changing the Drive Security Settings*.

3. Select **Switch to External key management (EKM) mode,** if you want to switch over from LKM mode to EKM mode.
4. Click **OK**.

   After you click **OK**, you will get the authentication drive security key. This key is generated internally by the system.

11.6.3.1 Authentication Drive Security Key

After you obtain the authentication drive security key from the EKM mode, you have to enter the authentication drive security key in the **Authentication Drive Security Key** dialog box as shown in the following figure.



Figure 232: Authentication Drive Security Key

If you enter a valid authentication key, you see a confirmation message indicating that EKM is enabled as shown in the following figure.



Figure 233: Confirmation Message When EKM is Enabled

Based on the application's response at that particular point of time, you come across two scenarios, the first scenario is, if **EKMS is selected and EKMS is not present**, and the second scenario is, if you enter a wrong security key in the **Authentication Drive Security Key** dialog box in Figure 232.

- Scenario # 1

EKM is selected and EKMS is not present.

When **EKM is selected and EKMS is not present**, you have to 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.

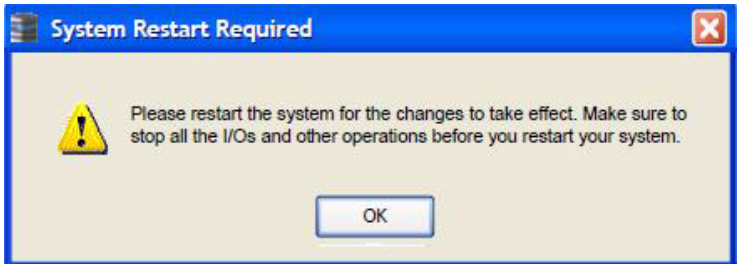


Figure 234: System Restart

- Scenario #2

Entering a wrong security key.

If you enter a wrong security key in the **Authentication Drive Security Key** dialog box in Figure 232, an error message is displayed as shown in the following figure.

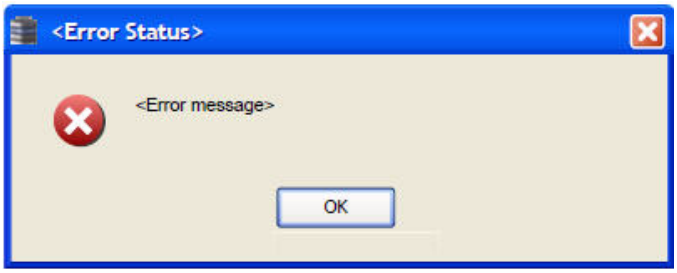


Figure 235: Error Message When You Enter the Wrong Security Key

11.6.4 Change Security Settings - EKM

When you select the **Change Security Settings** in the **Change Security wizard** as shown in Figure 236, two options are provided for you. If you select the first option, that is, **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.

Use any of the following options to enable the **Change Security Settings** wizard.

1. Select the **Physical View** tab in the left panel of the MegaRAID Storage Manager window, and click a controller icon.

- Select **Go To>Controller>Change Security Settings** in the main navigation bar which is present in the top portion of the MSM screen.
- Right-click on the controller icon, and click **Change Security Settings**. The **Change Security Settings** wizard appears as shown in the following figure.



Figure 236: Change Security Settings

When you select the **Change current security settings** option, you will get a different EKM key for the controller.

Based on the application's response two scenarios occur, the first scenario is, when **EKM server is present**, and the second scenario is, when **EKM server is not present**.

- Scenario # 1

EKM server is present.

If EKM server is present, you will get a confirmation message indicating that the drive security is enabled on the controller as shown in Figure 237.



Figure 237: EKMS server is present

- Scenario # 2

EKM server is not present.

When EKM server is not present, you have to 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.

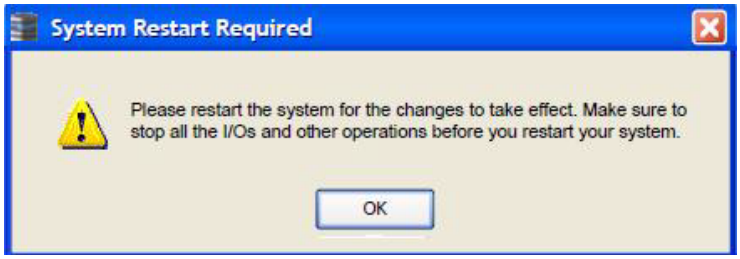


Figure 238: System Restart Required Message

NOTE: If you select the **Switch to Local key Management (LKM)** mode in the **Change Security settings wizard** in Figure 236, you can view the details of LKM in *Enabling Drive Security using LKM*.

11.6.5 Importing Foreign Drives

You can import foreign drives in the LKM or EKM mode. Based 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.

Use any of the following options to import foreign drives.

1. Select the **Physical View** tab in the left panel of the MegaRAID Storage Manager window, and click a controller icon.

- 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 which is present in the top portion of the MSM screen.

  If a Foreign Configuration is detected, the "Foreign Configuration Detected" screen gets displayed. If you choose **Import** and click **OK**, the **Unlock Foreign drives wizard** appears (see Figure 240).

11.6.6 Importing Foreign Drives to LKM

When VD is secured with EKM mode, and if you change the security settings from EKM mode to LKM mode, and want to import those foreign configured VDs, then the following dialog box appears as shown in the following figure.



Figure 239: Importing Foreign Drives- EKM to LKM

11.6.7 Importing Foreign Drives to EKM

When VD is secured with LKM mode, and want to switch from LKM to EKM, and that VD is in Foreign state, and if you want to import those foreign configuration, you need to provide the LKM key for each VD.



Figure 240: Unlock Foreign Drives

- The locked drives field display the drives that are locked with LKM mode.
- The Security Key displays the LKM security key for the VD.

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 screen is displayed. If no drives are unlocked, the Incorrect security Key dialog is displayed.

11.6.8 Importing Foreign Drives to EKM

When VD is secured with EKM mode, and that VD is in Foreign state, and if you switch to LKM mode, and if you want to import to foreign configuration, then the dialog box appears as shown in the following figure.



Figure 241: Switch to EKM Mode

11.6.9 Enabling Drive Security using LKM

This section describes how to enable, change, and disable 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 View** tab in the left panel of the MegaRAID Storage Manager window, and click a controller icon.

2. Select **Go To>Controller>Change Drive Security>Enable**.

   The Enable Drive Security screen appears as shown in the following figure.

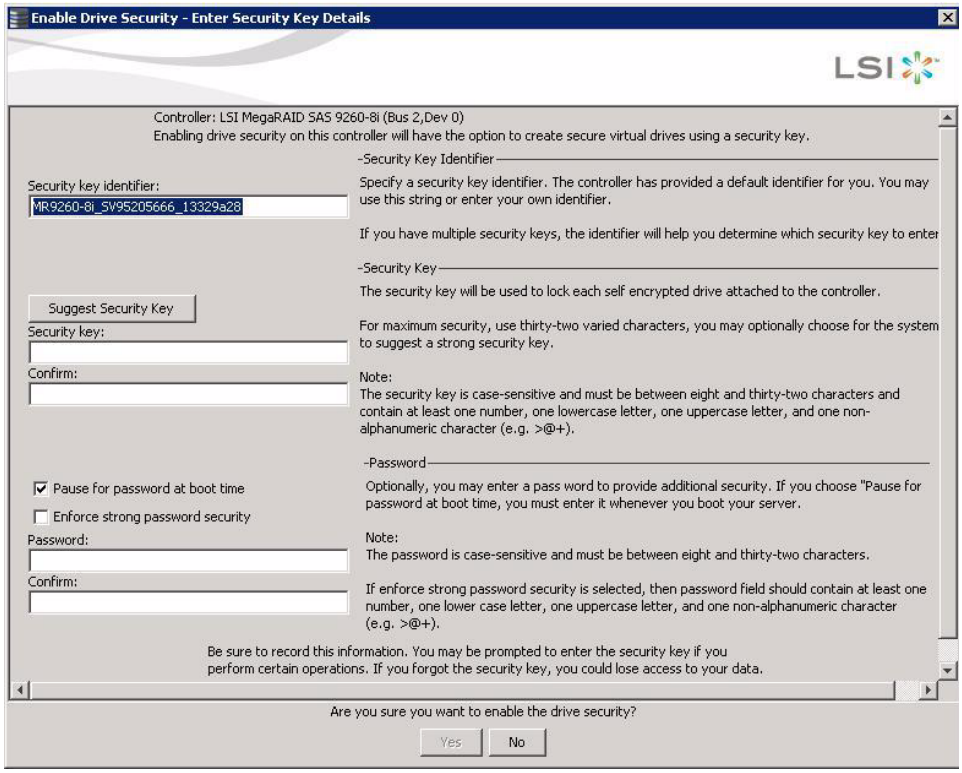


Figure 242: Enable Drive Security - Security Key Identifier

2. 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.

3. Click **Suggest Security Key** to have the system create a security key or you can enter a new security key.
4. Enter the new security key again to confirm, as shown in Figure 243.

NOTE: 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-US keyboard users must be careful not to enter DBCS characters in the security key field. Firmware works with the ASCII character set only.

Figure 243 shows the security key entered and confirmed on this screen.

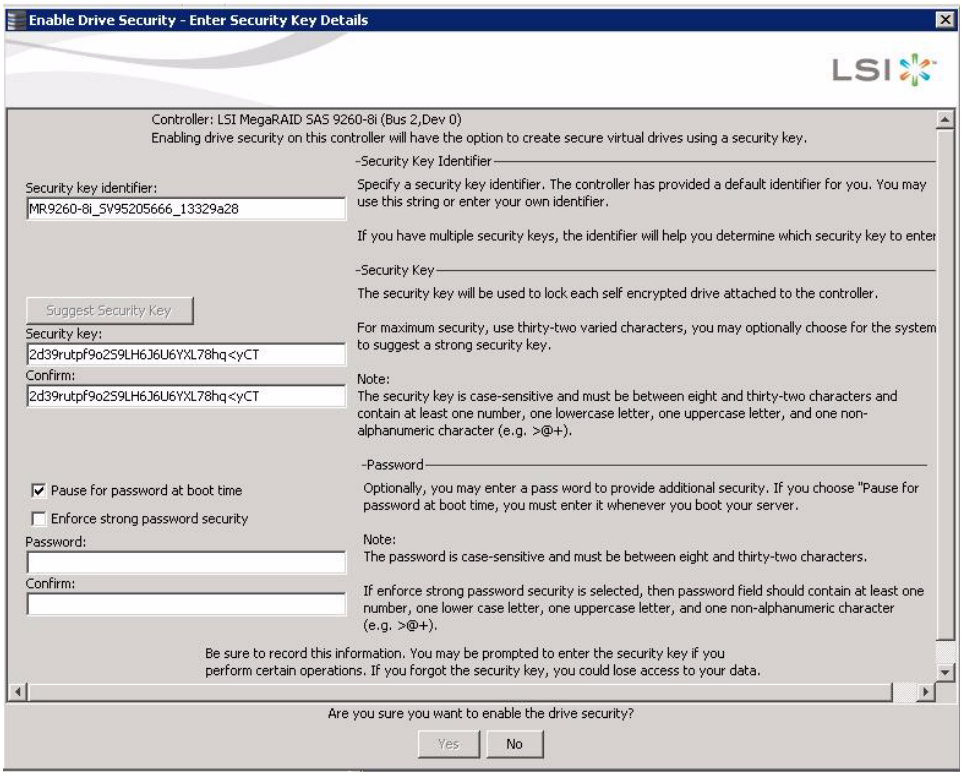


Figure 243: Enable Drive Security - Security Key

5. (Optional) Select the **Pause for password at boot time** check box.

   If you choose this option, you have to enter the password whenever you boot the server.

6. (Optional) Select the **Enforce strong password security** check box.

   If you choose this option, make sure the password is between eight and thirty-two 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.

7. (Optional) Enter a password in the **Password** field and then enter the same password in the **Confirm** field, as shown in Figure 244.

   Warning messages appear if there is a mismatch between the characters entered in the **Password** field and the **Confirm** field, or if there is an invalid character entered.

CAUTION: Be sure to record the password. If you lose the password, you could lose access to your data.

Figure 244 shows the password entered and confirmed on this screen.

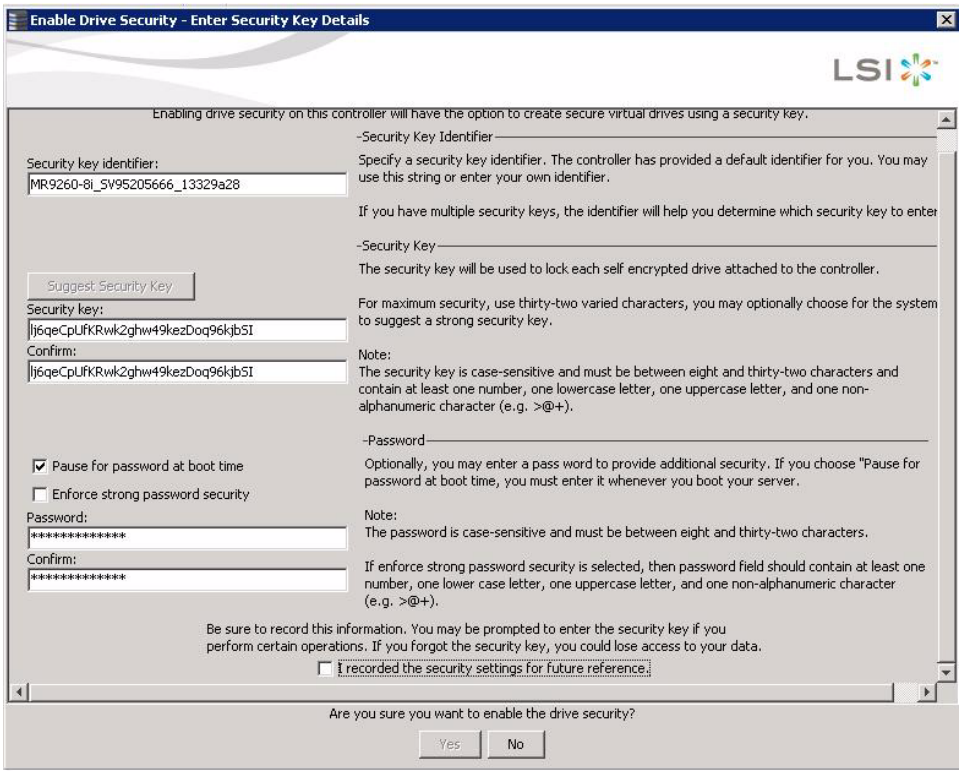


Figure 244: Enable Drive Security - Password

8. Click **Next**.

   The Confirm Enable Drive Security screen appears, as shown in Figure 245, to show the changes requested to the drive security settings.

CAUTION: 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 245: Confirm Create Security Key Screen

9. Click the checkbox **I recorded the security settings for future reference** and then click **Yes** to confirm that you want to enable drive security on this controller and have recorded the security settings for future reference.

   MSM 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 View** tab in the left panel of the MegaRAID Storage Manager main menu, and click a controller icon.
2. Select **Go To>Controller>Change Drive Security**.

   The Change Security Settings – Introduction screen appears. This screen 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 screen appears.
4. Keep the existing security key identifier or enter a new security key identifier.

NOTE: If you change the security key, LSI highly recommendeds that you change the security key identifier. Otherwise, you cannot differentiate between the security keys.

5. Click **Next**.

   The Change Security Settings - Security Key screen appears.

6. Click **Use the existing drive security key** 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 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-US keyboard users must be careful not to enter DBCS characters in the security key field. Firmware works with the ASCII character set only.

7. Click **Next**.

   The Authenticate Drive Security Settings Screen appears. Authentification 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 screen appears.

9. If you choose to, click the option to use a password in addition to the security key.
10. If you chose to 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 passphrase can hold up to 32 characters. The key must be at least eight characters.

    The next screen that appears describes the changes you made and asks you whether you want to confirm these changes.

11. Click the checkbox to confirm that you have recorded the security settings for future reference and then click **Yes** to confirm that you want to change the drive security settings.

    MSM 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 there are any secure drive groups on the controller, you cannot disable drive security. A warning screen 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 View** tab in the left panel of the MegaRAID Storage Manager main menu, and click a controller icon.
2. Select **Go To>Controller>Disable Drive Security**, as shown in Figure 246.

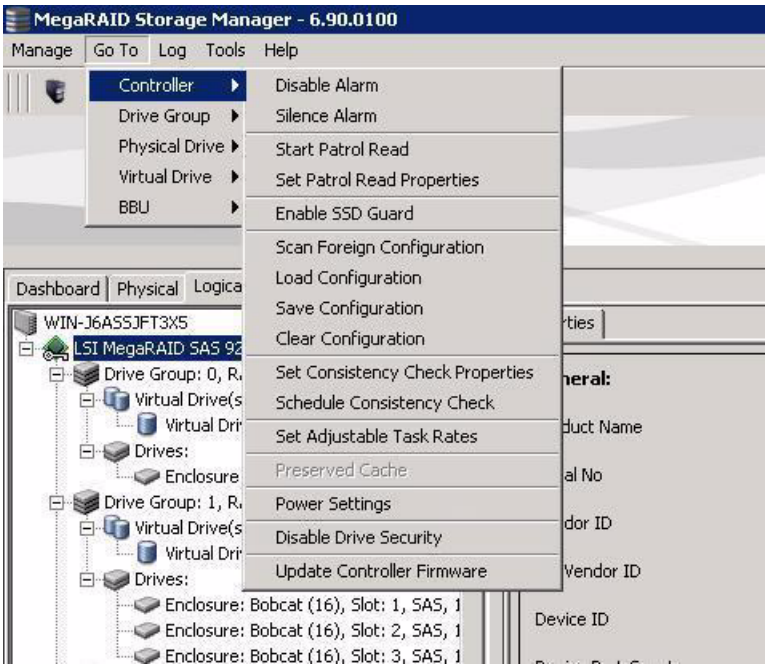


Figure 246: Disable Drive Security Option

The Confirm Disable Drive Security screen appears.

3. To disable drive security, click **Yes**.

MSM 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 will 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 MSM 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 do the following:

- 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 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 the drives are imported, there is no configuration to import.

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

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 used to enable drive security.
2. After you create a security key, right-click the controller icon and click **Scan for Foreign Configuration**.

   If there are locked drives (security is enabled), the Unlock foreign drives dialog box appears.
3. Enter the security key to unlock the configuration.

   The Foreign Configuration Detected screen appears, as shown in Figure 247.



Figure 247: Foreign Configuration Detected Screen

4. Click **Import** to import the foreign configuration from all of the foreign drives, **Clear** to remove the configuration from all foreign drives, or **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 menu.

6. Repeat the import process for any remaining drives.

   Because locked drive can use different security key, you must verify whether there are any remaining drives to imported.

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

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** screen 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 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 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 then 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 must have the same maximum link speed setting. If the phys in a SAS port have different link speeds, the firmware uses the lowest link speed for that SAS port.

Follow these steps to change the link speed.

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 248: 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.

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.

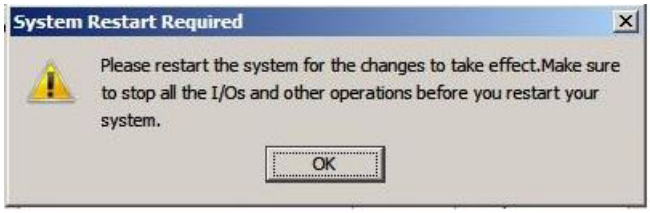


Figure 249: System Restart Required Message Box

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.

A.1 Error Levels

Each message that appears in the event log has an error level that indicates the severity of the event, as shown in Table 126.

Table 126: Event Error Levels

Error 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 127 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 127: Event Messages

Number	Type	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 127: Event Messages (Continued)

Number	Type	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 image(s) 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 127: Event Messages (Continued)

Number	Type	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 127: Event Messages (Continued)

Number	Type	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 127: Event Messages (Continued)

Number	Type	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 127: Event Messages (Continued)

Number	Type	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 127: Event Messages (Continued)

Number	Type	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 127: Event Messages (Continued)

Number	Type	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 127: Event Messages (Continued)

Number	Type	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 127: Event Messages (Continued)

Number	Type	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 127: Event Messages (Continued)

Number	Type	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%lx
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 127: Event Messages (Continued)

Number	Type	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	Snaphot 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(TM) with capacity 0x%lx logical blocks
0x0174	Information	%s is using CacheCade(TM) %s
0x0175	Information	%s is no longer using CacheCade(TM) %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(TM) (%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

Table 127: Event Messages (Continued)

Number	Type	Event Text
0x0187	Information	%s mirror joined
0x0188	Warning	%s link %d failure in wide port
0x0189	Information	%s link %d restored in wide port

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 128.

Table 128: 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 128: Error Messages and Descriptions (Continued)

Number	Event Text
0x17	LD operation not possibe - 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	FW real-time currently not set
0x32	Command issues while FW 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 imcomplete
0x3d	Returned in case if a command is tried on unsupported hardware

Table 128: 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 resert is required
0x5a	No EKM boot agent detected
0x5b	No space on the snapshot repositiry 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 *Read/Write*, *Read Only*, or *Blocked*.

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 *drive group*.

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 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.

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.

current write policy A virtual drive property that indicates whether the virtual drive currently supports Write Back mode or Write Through mode.

- 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 A drive property indicating the status of the drive. A drive can be in one of the following states:

- 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.

- 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 *hole* 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, *PCIX*.

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.

revertible 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 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.

subvendor ID A controller property that lists additional vendor ID information about the controller.

uncorrectable error count 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.

vendor ID A controller property indicating the vendor-assigned ID number of the controller.

vendor info A drive property listing the name of the vendor of the drive.

virtual drive 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.

virtual drive state A virtual drive property indicating the condition of the virtual drive. Examples include Optimal and Degraded.

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

write policy See *Default Write Policy*.

write-through 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.


LSI