12Gb/s MegaRAID® SAS Software

User Guide

Version 1.0

pub-005681

August 08, 2016

For a comprehensive list of changes to this document, see the History of Technical Changes.

Corporate Headquarters

San Jose, CA

Website

www.avagotech.com

Avago Technologies, the A logo, MegaRAID, MegaRAID Storage Manager, CacheCade, SSD Guard, Dimmer Switch, SafeStore, iTBBU, CacheVault, SHIELD, WarpDrive, and 3ware are trademarks of Avago Technologies in the United States and other countries. All other brand and product names may be trademarks of their respective companies.

Data subject to change. Copyright © 2016 Avago Technologies. All Rights Reserved.

Table of Contents

Chapter 1: Overview 13

1.1 SAS Technology 13
1.2 Serial-Attached SCSI Device Interface 14
1.3 Serial ATA III Features 14
1.4 Solid State Drive Features 14
1.4.1 SSD Guard 15
1.5 Dimmer Switch Features 15
1.6 UEFI 2.0 Support 15
1.7 Configuration Scenarios 15
1.7.1 Valid Drive Mix Configurations with HDDs and SSDs 17
1.8 Technical Support 18

Chapter 2: Introduction to RAID 19

2.1 Components and Features 19
2.1.1 Drive Group 19
2.1.2 Virtual Drive 19
2.1.3 Fault Tolerance 20
2.1.3.1 Multipathing 20
2.1.4 Consistency Check 21
2.1.5 Replace 21
2.1.6 Background Initialization 22
2.1.7 Patrol Read 22
2.1.8 Disk Striping 22
2.1.9 Disk Mirroring 23
2.1.10 Parity 23
2.1.11 Disk Spanning 24
2.1.12 Hot Spares 25
2.1.13 Disk Rebuilds 26
2.1.14 Rebuild Rate 27
2.1.15 Hot Swap 27
2.1.16 Drive States 27
2.1.17 Virtual Drive States 27
2.1.18 Beep Codes 28
2.1.19 Enclosure Management 28
2.2 RAID Levels 28
2.2.1 Summary of RAID Levels 29
2.2.2 Selecting a RAID Level 30
2.2.3 RAID 0 Drive Groups 30
2.2.4 RAID 1 Drive Groups 30
2.2.5 RAID 5 Drive Groups 31
2.2.6 RAID 6 Drive Groups 32
2.2.7 RAID 00 Drive Groups 33
2.2.8 RAID 10 34
2.2.9 RAID 50 Drive Groups 35
2.2.10 RAID 60 Drive Groups 36
2.3 RAID Configuration Strategies 37
2.3.1 Maximizing Fault Tolerance 38
2.3.2 Maximizing Performance 39
2.3.3 Maximizing Storage Capacity 40
2.4 RAID Availability 41
2.4.1 RAID Availability Concept 41
2.5 Configuration Planning 42
2.6 Number of Drives 42

Chapter 3: SafeStore Disk Encryption . . . 43

3.1 Terminology . . . 43
3.2 Workflow . . . 44
3.2.1 Enable Security . . . 44
3.2.2 Change Security . . . 45
3.2.3 Create Secure Virtual Drives . . . 45
3.2.4 Import a Foreign Configuration . . . 45
3.3 Instant Secure Erase . . . 46

Chapter 4: Ctrl-R Utility . . . 47

4.1 Overview . . . 47
4.2 Starting the Ctrl-R Utility . . . 47
4.3 Exiting the Ctrl-R Utility . . . 47
4.4 Ctrl-R Utility Keystrokes . . . 48
4.5 Ctrl-R Utility Menus . . . 48
4.5.1 Virtual Drive Management Menu . . . 49
4.5.2 Physical Drive Management Menu . . . 49
4.5.3 Controller Management Menu . . . 50
4.5.4 Properties Menu . . . 52
4.5.5 Foreign View Menu . . . 53
4.6 Managing Software Licensing . . . 54
4.6.1 Managing Advanced Software Options . . . 54
4.6.2 Managing Advanced Software Summary . . . 56
4.6.3 Activating an Unlimited Key over a Trial Key . . . 57
4.6.4 Activating a Trial Software . . . 58
4.6.5 Activating an Unlimited Key . . . 58
4.7 Creating a Storage Configuration . . . 59
4.7.1 Selecting Additional Virtual Drive Properties . . . 61
4.7.2 Creating a CacheCade Virtual Drive . . . 62
4.7.3 Modifying a CacheCade Virtual Drive . . . 64
4.7.4 Creating a CacheCade Pro 2.0 Virtual Drive . . . 65
4.7.5 Modifying a CacheCade Pro 2.0 Virtual Drive . . . 66
4.7.6 Enabling SSD Caching on a Virtual Drive . . . 67
4.7.7 Disabling SSD Caching on a Virtual Drive . . . 68
4.7.8 Enabling or Disabling SSD Caching on Multiple Virtual Drives . . . 68
4.7.9 Deleting a Virtual Drive with SSD Caching Enabled . . . 69
4.8 Clearing the Configuration . . . 70
4.9 Avago SafeStore Encryption Services . . . 70
4.9.1 Enabling Drive Security . . . 70
4.9.2 Changing Security Settings . . . 72
4.9.3 Disabling Drive Security . . . 73
4.9.4 Importing or Clearing a Foreign Configuration . . . 73
4.9.4.1 Foreign Configurations in Cable Pull and Drive Removal Scenarios . . . 75
4.10 Discarding Preserved Cache . . . 76
4.11 Converting JBOD Drives to Unconfigured Good Drives . . . 77
4.12 Converting Unconfigured Good Drives to JBOD Drives . . . 78
4.13 Enabling Security on a JBOD . . . 79
4.14 Viewing and Changing Device Properties . . . 80
4.14.1 Viewing Controller Properties . . . 80
4.14.2 Modifying Controller Properties . . . 80
4.14.3 Viewing and Changing Virtual Drive Properties . . . 82
4.14.4 Deleting a Virtual Drive . . . 84
4.14.5 Deleting a Virtual Drive Group . . . 84
4.14.6 Expanding a Virtual Drive . . . 84
4.14.7 Erasing a Virtual Drive . . . 85
4.14.8 Managing Link Speed . . . 86
4.14.9 Managing Power Save Settings for the Controller . . . 87
4.14.9.1 Setting Advanced Power Save Settings . . . 88

4.14.10 Start Manual Learn Cycle ... 89
4.14.11 Managing Power Save Settings for the Drive Group ... 90
4.14.12 Managing BBU Information ... 90
4.14.13 Managing Dedicated Hot Spares ... 92
4.14.14 Securing a Drive Group ... 92
4.14.15 Setting LED Blinking ... 92
4.14.16 Performing a Break Mirror Operation ... 93
4.14.17 Performing a Join Mirror Operation ... 93
4.14.18 Hiding a Virtual Drive ... 95
4.14.19 Unhiding a Virtual Drive ... 95
4.14.20 Hiding a Drive Group ... 95
4.14.21 Unhiding a Drive Group ... 96
4.14.22 Performing a Transport Ready Operation ... 96
4.14.22.1 Clearing a Transport Ready Operation ... 100
4.15 Managing Storage Configurations ... 101
4.15.1 Initializing a Virtual Drive ... 101
4.15.2 Running a Consistency Check ... 102
4.15.3 Rebuilding a Physical Drive ... 103
4.15.4 Performing a Copyback Operation ... 103
4.15.5 Removing a Physical Drive ... 104
4.15.6 Creating Global Hot Spares ... 104
4.15.7 Removing a Hot Spare Drive ... 105
4.15.8 Making a Drive Offline ... 105
4.15.9 Making a Drive Online ... 105
4.15.10 Instant Secure Erase ... 105
4.15.11 Erasing a Physical Drive ... 105

Chapter 5: HII Configuration Utility ... 107

5.1 Starting the HII Configuration Utility ... 107
5.2 HII Dashboard View ... 108
5.2.1 Main Menu ... 108
5.2.2 HELP ... 109
5.2.3 PROPERTIES ... 109
5.2.4 ACTIONS ... 111
5.2.5 BACKGROUND OPERATIONS ... 112
5.2.6 MegaRAID ADVANCED SOFTWARE OPTIONS ... 112
5.3 Critical Boot Error Message ... 112
5.4 Managing Configurations ... 113
5.4.1 Creating a Virtual Drive from a Profile ... 113
5.4.2 Manually Creating a Virtual Drive ... 117
5.4.3 Creating a CacheCade Virtual Drive ... 121
5.4.4 Viewing Drive Group Properties ... 122
5.4.5 Viewing Global Hot Spare Drives ... 122
5.4.6 Clearing a Configuration ... 123
5.4.7 Make Unconfigured Good, Make JBOD, and Enable Security on JBOD ... 123
5.4.7.1 Make Unconfigured Good ... 124
5.4.7.2 Make JBOD ... 124
5.4.7.3 Enabling Security on JBOD ... 125
5.4.8 Managing Foreign Configurations ... 125
5.4.8.1 Previewing and Importing a Foreign Configuration ... 126
5.4.8.2 Clearing a Foreign Configuration ... 128
5.5 Managing Controllers ... 128
5.5.1 Viewing Advanced Controller Management Options ... 130
5.5.2 Viewing Advanced Controller Properties ... 131
5.5.3 Managing MegaRAID Advanced Software Options ... 133
5.5.4 Scheduling a Consistency Check ... 134
5.5.5 Saving or Clearing Controller Events ... 135
5.5.6 Enabling or Disabling Drive Security ... 136

5.5.7 Changing a Security Key 139
5.5.8 Saving the TTY Log 140
5.5.9 Managing and Changing Link Speeds 141
5.5.10 Setting Cache and Memory Properties 141
5.5.11 Running a Patrol Read 142
5.5.12 Changing Power Save Settings 144
5.5.13 Setting Emergency Spare Properties 145
5.5.14 Changing Task Rates 146
5.5.15 Upgrading the Firmware 147
5.6 Managing Virtual Drives 149
5.6.1 Selecting Virtual Drive Operations 150
5.6.1.1 Locating Physical Drives in a Virtual Drive 150
5.6.1.2 Deleting a Virtual Drive 151
5.6.1.3 Hiding a Virtual Drive 151
5.6.1.4 Unhiding a Virtual Drive 151
5.6.1.5 Hiding a Drive Group 152
5.6.1.6 Unhiding a Drive Group 152
5.6.2 Performing a Transport Ready Operation 152
5.6.2.1 Clearing a Transport Ready Operation 155
5.6.2.2 Reconfiguring a Virtual Drive 156
5.6.2.3 Initializing a Virtual Drive 159
5.6.2.4 Erasing a Virtual Drive 159
5.6.2.5 Enabling and Disabling SSD Caching 159
5.6.2.6 Securing a Virtual Drive 160
5.6.2.7 Running a Consistency Check 160
5.6.2.8 Expanding a Virtual Drive 160
5.6.2.9 Disabling Protection on a Virtual Drive 161
5.6.3 Managing CacheCade Virtual Drives 161
5.6.4 Viewing Associated Drives 162
5.6.5 Viewing and Managing Virtual Drive Properties and Options 162
5.7 Managing Physical Drives 164
5.7.1 Performing Drive Operations 165
5.7.1.1 Locating a Drive 166
5.7.1.2 Making a Drive Unconfigured Bad, Unconfigured Good, or JBOD 166
5.7.1.3 Enabling Security on JBOD 167
5.7.1.4 Replacing a Drive 167
5.7.1.5 Placing a Drive Offline 168
5.7.1.6 Placing a Drive Online 168
5.7.1.7 Marking a Drive Missing 169
5.7.1.8 Replacing a Missing Drive 169
5.7.1.9 Assigning a Global Hot Spare Drive 169
5.7.1.10 Assigning a Dedicated Hot Spare Drive 170
5.7.1.11 Unassigning a Hot Spare Drive 170
5.7.1.12 Initializing or Erasing a Drive 171
5.7.1.13 Rebuilding a Drive 172
5.7.1.14 Securely Erasing a Drive 172
5.7.1.15 Removing a Physical Drive 173
5.7.2 Viewing Advanced Drive Properties 173
5.8 Managing Hardware Components 175
5.8.1 Managing Batteries 176
5.8.1.1 Setting Automatic Learn Cycle Properties 178
5.8.2 Managing Enclosures 179

Chapter 6: StorCLI 181

6.1 Overview 181
6.2 Support for MegaCLI Commands 181
6.3 Devices Supported by the StorCLI Tool 181
6.4 Installation 181

6.4.1 Installing the StorCLI Tool on Microsoft Windows Operating Systems ... 182
6.4.2 Installing the StorCLI Tool on Linux Operating Systems ... 183
6.4.3 Installing the StorCLI Tool on Ubuntu Operating Systems ... 183
6.4.4 Installing the StorCLI Tool on VMware Operating Systems ... 183
6.4.5 Installing the StorCLI Tool on FreeBSD Operating Systems ... 183
6.4.6 Installing the StorCLI Tool on Microsoft EFI ... 183
6.4.7 Installing the StorCLI Tool on Solaris Operating Systems ... 184
6.5 StorCLI Tool Command Syntax ... 184
6.6 Working with the Storage Command Line Interface Tool ... 186
6.6.1 System Commands ... 186
6.6.1.1 System Show Commands ... 186
6.6.2 Controller Commands ... 186
6.6.2.1 Show and Set Controller Properties Commands ... 187
6.6.2.2 Controller Show Commands ... 193
6.6.2.3 Controller Background Tasks Operation Commands ... 193
6.6.2.4 Premium Feature Key Commands ... 197
6.6.2.5 Controller Security Commands ... 197
6.6.2.6 Flashing Controller Firmware Command ... 198
6.6.2.7 Controller Cache Command ... 199
6.6.2.8 Controller Configuration Commands ... 199
6.6.3 Drive Commands ... 199
6.6.3.1 Drive Show Commands ... 200
6.6.3.2 Missing Drives Commands ... 201
6.6.3.3 Set Drive State Commands ... 201
6.6.3.4 Drive Initialization Commands ... 202
6.6.3.5 Drive Firmware Download Commands ... 203
6.6.3.6 Locate Drives Commands ... 203
6.6.3.7 Prepare to Remove Drives Commands ... 204
6.6.3.8 Drive Security Command ... 204
6.6.3.9 Drive Secure Erase Commands ... 204
6.6.3.10 Rebuild Drives Commands ... 205
6.6.3.11 Drive Copyback Commands ... 206
6.6.3.12 Hot Spare Drive Commands ... 207
6.6.3.13 Drive Performance Monitoring Commands ... 208
6.6.4 Virtual Drive Commands ... 209
6.6.4.1 Add Virtual Drives Commands ... 209
6.6.4.2 Delete Virtual Drives Commands ... 212
6.6.4.3 Virtual Drive Show Commands ... 213
6.6.4.4 Preserved Cache Commands ... 214
6.6.4.5 Change Virtual Properties Commands ... 214
6.6.4.6 Virtual Drive Initialization Commands ... 216
6.6.4.7 Virtual Drive Erase Commands ... 217
6.6.4.8 Virtual Drive Migration Commands ... 217
6.6.4.9 Virtual Drive Consistency Check Commands ... 218
6.6.4.10 Background Initialization Commands ... 219
6.6.4.11 Virtual Drive Expansion Commands ... 220
6.6.4.12 Display the Bad Block Table ... 221
6.6.4.13 Clear the LDBBM Table Entires ... 221
6.6.5 Foreign Configurations Commands ... 221
6.6.6 BIOS-Related Commands ... 222
6.6.6.1 OPROM BIOS Commands ... 222
6.6.7 Drive Group Commands ... 223
6.6.7.1 Drive Group Show Commands ... 223
6.6.7.2 Set Transport Ready Drive Group Command ... 224
6.6.8 Dimmer Switch Commands ... 225
6.6.8.1 Change Virtual Drive Power Settings Commands ... 225
6.6.9 BBU Commands ... 226
6.6.10 CacheVault Commands ... 229
6.6.11 Enclosure Commands ... 229

6.6.12 PHY Commands ... 230
6.6.13 Logging Commands ... 231
6.6.14 Automated Physical Drive Caching Commands ... 232
6.7 Frequently Used Tasks ... 233
6.7.1 Showing the Version of the Storage Command Line Interface Tool ... 233
6.7.2 Showing the StorCLI Tool Help ... 233
6.7.3 Showing System Summary Information ... 233
6.7.4 Showing Free Space in a Controller ... 233
6.7.5 Adding Virtual Drives ... 233
6.7.6 Setting the Cache Policy in a Virtual Drive ... 234
6.7.7 Showing Virtual Drive Information ... 234
6.7.8 Deleting Virtual Drives ... 234
6.7.9 Flashing Controller Firmware ... 235

Chapter 7: MegaRAID Storage Manager Overview and Installation ... 236

7.1 Overview ... 236
7.1.1 Creating Storage Configurations ... 236
7.1.2 Monitoring Storage Devices ... 236
7.1.3 Maintaining Storage Configurations ... 236
7.2 Hardware and Software Requirements ... 236
7.3 Installing MegaRAID Storage Manager ... 238
7.3.1 Prerequisite for MegaRAID Storage Manager Installation ... 238
7.3.2 Installing the MegaRAID Storage Manager Software on Microsoft Windows ... 238
7.3.2.1 Setup Options ... 244
7.3.3 Uninstalling the MegaRAID Storage Manager Software on Microsoft Windows ... 244
7.3.3.1 Uninstalling the MegaRAID Storage Manager Software through the Control Panel ... 244
7.3.3.2 Uninstalling the MegaRAID Storage Manager Software Using the Command Prompt ... 244
7.3.3.3 Uninstalling the MegaRAID Storage Manager Software Using the MegaRAID Storage Manager Uninstallation Utility ... 245
7.3.4 Installing and Supporting the MegaRAID Storage Manager Software on Solaris and SPARC Operating Systems ... 245
7.3.4.1 Installing the MegaRAID Storage Manager Software for Solaris 10 x86 ... 245
7.3.4.2 Installing the MegaRAID Storage Manager Software for Solaris 10 SPARC ... 245
7.3.4.3 Installing the MegaRAID Storage Manager Software for Solaris 11 x86 ... 246
7.3.4.4 Installing the MegaRAID Storage Manager Software for Solaris 11 SPARC ... 246
7.3.5 Uninstalling the MegaRAID Storage Manager Software on Solaris 10 (U5, U6, U7, U8, U9, and U10), Solaris 11 (x86 and x64), and Solaris SPARC ... 246
7.3.6 Prerequisites for Installing the MegaRAID Storage Manager Software on RHEL6.x x64 and RHEL7.x x64 ... 247
7.3.7 Installing the MegaRAID Storage Manager Software on RHEL or SLES/SuSE Linux ... 247
7.3.8 Linux Error Messages ... 248
7.3.9 Kernel Upgrade ... 249
7.3.10 Uninstalling the MegaRAID Storage Manager Software on RHEL or SLES or SuSE Linux ... 249
7.3.11 MegaRAID Storage Manager Software Customization ... 249
7.3.12 Updating the Strength of Public and Private RSA keys ... 250
7.3.12.1 Limitations ... 251
7.3.12.2 Updating the Property File and Vivaldikeys ... 252
7.3.13 Stopping the Pop-Up Notification Process ... 252
7.3.13.1 Windows Operating System ... 252
7.3.13.2 Linux, Solaris x86, and Solaris SPARC Operating Systems ... 252
7.3.14 Restarting the Pop-Up Notification Process ... 252
7.4 Installing and Supporting the MegaRAID Storage Manager Software on VMware ... 253
7.4.1 Prerequisites for Installing the MegaRAID Storage Manager for VMware ... 253
7.4.2 Installing the MegaRAID Storage Manager Software on VMware ESX (VMware Classic) ... 253
7.4.3 Uninstalling the MegaRAID Storage Manager Software for VMware ... 253
7.4.4 MegaRAID Storage Manager Support on the VMware ESXi Operating System ... 254
7.4.5 Limitations of Installation and Configuration ... 255
7.4.5.1 Differences in the MegaRAID Storage Manager Software for VMware ESXi ... 255
7.5 Installing and Configuring a CIM Provider ... 256
7.5.1 Installing a CIM SAS Storage Provider on the Linux Operating System ... 256
7.5.2 Running the CIM SAS Storage Provider on Pegasus ... 257

7.5.3 Installing a CIM SAS Storage Provider on Windows ... 257
7.6 Installing and Configuring an SNMP Agent ... 258
7.6.1 Prerequisite for the Avago SNMP Agent RPM Installation ... 258
7.6.2 Installing an SNMP Agent on Windows ... 258
7.6.2.1 Installing an SNMP Agent ... 258
7.6.2.2 Installing SNMP Service for Windows ... 258
7.6.2.3 Configuring SNMP Service on the Server Side ... 259
7.6.2.4 Installing the SNMP Service for the Windows 2008 Operating System ... 259
7.6.2.5 Configuring the SNMP Service on the Server Side for the Windows 2008 Operating System ... 259
7.6.3 Prerequisite for Installing the SNMP Agent on a Linux Server ... 260
7.6.4 Installing and Configuring an SNMP Agent on Linux ... 260
7.6.5 Installing and Configuring an SNMP Agent on Solaris ... 261
7.6.5.1 Prerequisites ... 261
7.6.5.2 Installing the SNMP Agent on Solaris ... 261
7.6.5.3 Avago SAS SNMP MIB Location ... 262
7.6.5.4 Starting, Stopping, and Checking the Status of the Avago SAS SNMP Agent ... 262
7.6.5.5 Configuring the snmpd.conf File ... 262
7.6.5.6 Configuring SNMP Traps ... 264
7.6.5.7 Uninstalling the SNMP Package ... 265
7.7 MegaRAID Storage Manager Remotely Connecting to VMware ESX ... 265
7.8 Prerequisites to Running MegaRAID Storage Manager Remote Administration ... 265
7.9 CLI Packaging Details ... 266

Chapter 8: MegaRAID Storage Manager Window and Menus ... 267

8.1 Starting the MegaRAID Storage Manager Software ... 267
8.2 Discovery and Login ... 267
8.3 High Availability DAS Support ... 271
8.4 LDAP Support ... 273
8.5 Configuring LDAP Support Settings ... 275
8.6 MegaRAID Storage Manager Main Menu ... 276
8.6.1 Dashboard View, Physical View, and Logical View ... 276
8.6.2 Physical Drive Temperature ... 281
8.6.3 Shield State ... 282
8.6.4 Shield State Physical View ... 282
8.6.5 Logical View Shield State ... 283
8.6.6 Viewing the Physical Drive Properties ... 283
8.6.7 Viewing the Server Profile of a Drive in Shield State ... 284
8.6.8 Displaying the Virtual Drive Properties ... 285
8.6.8.1 Parity Size ... 285
8.6.8.2 Mirror Data Size ... 286
8.6.8.3 Metadata Size ... 287
8.6.9 Emergency Spare ... 288
8.6.9.1 Emergency Spare for Physical Drives ... 288
8.6.9.2 Emergency Spare Property for Controllers ... 289
8.6.9.3 Commissioned Hotspare ... 290
8.6.10 SSD Disk Cache Policy ... 291
8.6.10.1 Virtual Drive Settings ... 292
8.6.10.2 Set the Virtual Drive Properties ... 293
8.6.11 Non-SED Secure Erase Support ... 294
8.6.11.1 Group Show Progress for Drive Erase ... 296
8.6.11.2 Virtual Drive Erase ... 297
8.6.11.3 Group Show Progress for Virtual Drive Erase ... 299
8.6.12 Rebuild Write Cache ... 300
8.6.13 Background Suspend/Resume Support ... 300
8.6.14 Enclosure Properties ... 302
8.6.15 Expander Properties ... 302
8.7 GUI Elements in the MegaRAID Storage Manager Window and Menus ... 303
8.7.1 Device Icons ... 303

8.7.2 Properties and Graphical View Tabs ... 304
8.7.3 Event Log Panel ... 305
8.7.4 Menu Bar ... 305

Chapter 9: Configuration ... 307

9.1 Creating a New Configuration ... 307
9.1.1 Selecting Virtual Drive Settings ... 307
9.1.2 Optimum Controller Settings for CacheCade ... 309
9.1.3 Optimum Controller Settings for Fast Path ... 309
9.1.4 Creating a Virtual Drive Using Simple Configuration ... 309
9.1.5 Creating a Virtual Drive Using Advanced Configuration ... 313
9.2 Converting JBOD Drives to Unconfigured Good ... 320
9.2.1 Converting JBOD to Unconfigured Good from the MegaRAID Storage Manager Main Menu ... 321
9.2.2 Removing a JBOD Drive ... 322
9.3 Enabling Security on JBOD ... 322
9.4 Creating Hot Spare Drives ... 323
9.5 Changing Adjustable Task Rates ... 323
9.6 Changing Power Settings ... 325
9.7 Recovering and Clearing Punctured Block Entries ... 326
9.8 Changing Virtual Drive Properties ... 327
9.9 Changing a Virtual Drive Configuration ... 328
9.9.1 Accessing the Modify Drive Group Wizard ... 329
9.9.2 Adding a Drive or Drives to a Configuration ... 330
9.9.3 Removing a Drive from a Configuration ... 333
9.9.4 Replacing a Drive ... 334
9.9.5 Migrating the RAID Level of a Virtual Drive ... 335
9.10 Deleting a Virtual Drive ... 338
9.11 Performing a Join Mirror Operation ... 339
9.11.0.1 Configuration ... 340
9.12 Hiding and Unhiding a Virtual Drive or a Drive Group ... 371
9.12.1 Hiding a Virtual Drive ... 371
9.12.2 Unhiding a Virtual Drive ... 371
9.12.3 Hiding a Drive Group ... 371
9.12.4 Unhiding a Drive Group ... 372
9.13 Performing a Transport Ready Operation ... 372
9.13.1 Clearing a Transport Ready Operation ... 378

Chapter 10: Monitoring Controllers and Their Attached Devices ... 379

10.1 Alert Delivery Methods ... 379
10.1.1 Vivaldi Log/MegaRAID Storage Manager Log ... 379
10.1.2 System Log ... 381
10.1.2.1 Setting Up the Custom Facility Level in the System Log File for the Solaris x86 Operating System ... 381
10.1.3 Pop-up Notification ... 382
10.1.4 Email Notification ... 382
10.2 Configuring Alert Notifications ... 383
10.3 Editing Alert Delivery Methods ... 385
10.4 Changing Alert Delivery Methods for Individual Events ... 386
10.5 Changing the Severity Level for Individual Events ... 387
10.6 Roll Back to the Default Individual Event Configuration ... 387
10.7 Entering or Editing the Sender Email Address and SMTP Server ... 388
10.8 Authenticating the SMTP Server ... 389
10.9 Adding Email Addresses of Recipients of Alert Notifications ... 389
10.10 Testing Email Addresses of Recipients of Alert Notifications ... 390
10.11 Removing Email Addresses of Recipients of Alert Notifications ... 390
10.12 Saving Backup Configurations ... 391
10.13 Loading Backup Configurations ... 391
10.14 Monitoring Server Events ... 391
10.15 Monitoring Controllers ... 392

10.16 Monitoring Drives ... 393
10.17 Running a Patrol Read ... 394
10.17.1 Patrol Read Task Rates ... 396
10.18 Monitoring Virtual Drives ... 396
10.19 Monitoring Enclosures ... 397
10.20 Monitoring Battery Backup Units ... 398
10.21 Battery Learn Cycle ... 399
10.21.1 Setting Automatic Learn Cycle Properties ... 399
10.21.2 Starting a Learn Cycle Manually ... 400
10.22 Monitoring Rebuilds and Other Processes ... 401
10.23 Managing Link Speed ... 402

Chapter 11: Maintaining and Managing Storage Configurations ... 404

11.1 Initializing a Virtual Drive ... 404
11.1.1 Running a Group Initialization ... 404
11.2 Running a Consistency Check ... 405
11.2.1 Setting the Consistency Check Settings ... 406
11.2.2 Scheduling a Consistency Check ... 406
11.2.3 Running a Group Consistency Check ... 407
11.3 Scanning for New Drives ... 408
11.4 Rebuilding a Drive ... 409
11.4.1 New Drives Attached to a MegaRAID Controller ... 410
11.5 Making a Drive Offline or Missing ... 410
11.6 Removing a Drive ... 410
11.7 Upgrading Firmware ... 411
11.7.1 Upgrading the CPLD Version ... 412

Chapter 12: Using MegaRAID Advanced Software ... 413

12.1 MegaRAID Advanced Software ... 413
12.2 MegaRAID Software Licensing ... 413
12.3 Managing MegaRAID Advanced Software ... 413
12.4 Activation Key ... 416
12.5 Advanced MegaRAID Software Status Summary ... 416
12.6 Application Scenarios and Messages ... 417
12.7 Activating an Unlimited Key over a Trial Key ... 418
12.7.1 Activating a Trial Software ... 419
12.7.2 Activating an Unlimited Key ... 420
12.7.3 Reusing the Activation Key ... 421
12.7.4 Securing Advanced MegaRAID Software ... 421
12.8 Configuring Key Vault (Re-hosting Process) ... 422
12.9 Re-hosting Complete ... 423
12.10 Deactivate Trial Software ... 424
12.11 Using the MegaRAID CacheCade Advanced Software ... 425
12.12 Using the MegaRAID CacheCade Pro 2.0 Software ... 429
12.12.1 Modifying the CacheCade Virtual Drive Properties ... 432
12.12.2 Enabling SSD Caching on a Virtual Drive ... 433
12.12.3 Disabling SSD Caching on a Virtual Drive ... 434
12.12.4 Enabling or Disabling SSD Caching on Multiple Virtual Drives ... 434
12.12.5 Modifying a CacheCade Drive Group ... 435
12.12.6 Clearing Configuration on CacheCade Pro 2.0 Virtual Drives ... 435
12.12.7 Removing Blocked Access ... 436
12.12.8 Deleting a Virtual Drive with SSD Caching Enabled ... 437
12.13 Fast Path Advanced Software ... 438
12.13.1 Setting Fast Path Options ... 438
12.14 Avago MegaRAID SafeStore Encryption Services ... 439
12.14.1 Enabling Drive Security ... 439
12.14.2 Enabling Drive Security ... 441
12.14.3 Changing Security Settings ... 444

12.14.4 Disabling Drive Security ... 445
12.14.5 Importing or Clearing a Foreign Configuration ... 446
12.14.5.1 Foreign Configurations in Cable Pull and Drive Removal Scenarios ... 447

Appendix A: Events and Messages ... 449
A.1 Error Levels ... 449
A.2 Event Messages ... 449

Appendix B: 3ware CLI Commands to StorCLI Command Conversion ... 468
B.1 System Commands ... 468
B.2 Controller Commands ... 468
B.3 Alarm Commands ... 471
B.4 Patrol Read and Consistency Check Commands ... 471
B.5 BBU Commands ... 472
B.6 Virtual Drive Commands ... 473
B.7 Physical Drive Commands ... 475
B.8 Enclosure Commands ... 476
B.9 Events and Logs ... 477
B.10 Miscellaneous Commands ... 477

Appendix C: MegaCLI Commands to StorCLI Command Conversion ... 478
C.1 System Commands ... 478
C.2 Controller Commands ... 478
C.3 Patrol Read Commands ... 481
C.4 Consistency Check Commands ... 482
C.5 OPROM BIOS Commands ... 482
C.6 Battery Commands ... 483
C.7 RAID Configuration Commands ... 484
C.8 Security Commands ... 485
C.9 Virtual Drive Commands ... 486
C.10 Physical Drive Commands ... 487
C.11 Enclosure Commands ... 489
C.12 PHY Commands ... 489
C.13 Alarm Commands ... 490
C.14 Event Log Properties Commands ... 490
C.15 Premium Feature Key Commands ... 490

Appendix D: Unsupported Commands in Embedded MegaRAID ... 492

Appendix E: CLI Error Messages ... 494
E.1 Error Messages and Descriptions ... 494

Appendix F: 240 Virtual Drive Feature Limitations ... 498
F.1 Host Software Utility Support ... 498
F.2 BIOS Known Limitations ... 498

Appendix G: Online Firmware Upgrade and Downgrade ... 499
G.1 Online Firmware Upgrade Support ... 499

Appendix H: Boot Messages and BIOS Error Messages ... 501
H.1 Displaying Boot Messages ... 501
H.2 Differences in the System Boot Mode ... 502

Appendix I: Transport Ready Pre-requisites ... 528

Appendix J: Glossary ... 531

History of Technical Changes ... 540

Chapter 1: Overview

This chapter provides an overview of this guide, which documents the utilities used to configure, monitor, and maintain MegaRAID® Serial-attached SCSI (SAS) RAID controllers with RAID control capabilities and the storage-related devices connected to them.

This guide describes how to use the MegaRAID Storage Manager™ software, the Ctrl- R utility, the StorCLI tool software and the Avago® MegaRAID Human Interface Infrastructure (HII) configuration utility.

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

1.1 SAS Technology

The MegaRAID 12Gb/s SAS RAID controllers are high-performance intelligent PCI Express®-to-SAS/Serial ATA II controllers with RAID control capabilities. The MegaRAID 12Gb/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. The MegaRAID 12Gb/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 and SATA III devices, and CacheCade SSD Read Caching software devices within the same storage infrastructure. These devices bring individual characteristics that make each of these more suitable choice depending on your storage needs. MegaRAID gives you the flexibility to combine these two similar technologies on the same controller, within the same enclosure, and in the same virtual drive.

NOTE Carefully assess any decision to combine SAS drives and SATA drives within the same virtual drives. Avoid mixing drives; this applies to both HDDs and CacheCade SSD Read Caching software.

The MegaRAID 12Gb/s SAS RAID controllers are based on the Avago 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. Avago offers a family of MegaRAID SAS RAID controllers addressing the needs for both internal and external solutions.

The SAS controllers support the ANSI® *Serial Attached SCSI standard, version 2.1*. In addition, the controller supports the SATA II protocol defined by the *Serial ATA specification, version 3.0*. Supporting both the SAS and SATA II interfaces, the SAS controller is a versatile controller that provides the backbone of both server environments and high-end workstation environments.

Each port on the SAS RAID controller supports SAS devices or SATA III devices using the following protocols:

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

1.2 Serial-Attached SCSI Device Interface

SAS is a serial, point-to-point, enterprise-level device interface that leverages the proven SCSI protocol set. SAS is a convergence of the advantages of SATA, 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 protocols use a thin, 7-wire connector instead of the 68-wire SCSI cable or 26-wire ATA cable. The SAS/SATA connector and cable are easier to manipulate, allow connections to smaller devices, and do not inhibit airflow. The point-to-point SATA architecture eliminates inherent difficulties created by the legacy ATA master-slave architecture, while maintaining compatibility with existing ATA firmware.

1.3 Serial ATA III Features

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

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

- Supports SATA III data transfers of 12Gb/s
- Supports STP data transfers of 12Gb/s
- Provides a serial, point-to-point storage interface
- Simplifies cabling between devices
- Eliminates the master-slave construction used in parallel ATA
- Allows addressing of multiple SATA II targets through an expander
- Allows multiple initiators to address a single target (in a fail-over configuration) through an expander

1.4 Solid State Drive Features

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

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

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

You can choose whether to allow a virtual drive to consist of both CacheCade 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 uses 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.

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

1.4.1 SSD Guard

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

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

1.5 Dimmer Switch Features

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

Dimmer Switch 1 – Spin down unconfigured disks. This feature is configurable and can be disabled.

Dimmer Switch 2 – Spin down Hot Spares. This feature is configurable and can be disabled.

1.6 UEFI 2.0 Support

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

1.7 Configuration Scenarios

You can use the SAS RAID controllers in three scenarios:

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

  In these configurations, use the RAID controller as a high-end SATA II-compatible controller that connects up to 8 disks. These configurations are mostly for low-end or entry servers. Enclosure management is provided through out-of-band Inter-IC (I^2C) bus. Side bands of both types of internal SAS connectors support the SFF-8485 (SGPIO) interface.

- **Midrange Internal SAS Configurations**

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

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

  These configurations are for both internal connectivity and external connectivity, using SATA drives, SAS drives, or both. External enclosure management is supported through in-band, SCSI-enclosed storage. The configuration must support STP and SMP.

The following figure shows a direct-connect configuration. The I^2C interface communicates with peripherals. The external memory bus provides a 32-bit memory bus, parity checking, and chip select signals for pipelined burst static random access memory (PBSRAM), 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 Avago SAS Direct-Connect Application

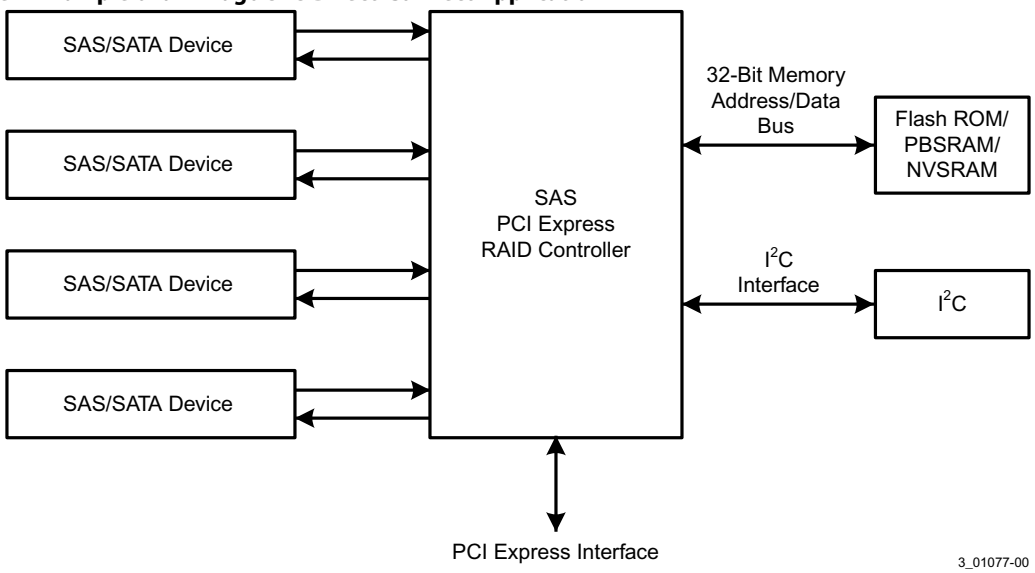

The following figure shows an example of a SAS RAID controller configured with an LSISASx12 expander that is connected to SAS disks, SATA II disks, or both.

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

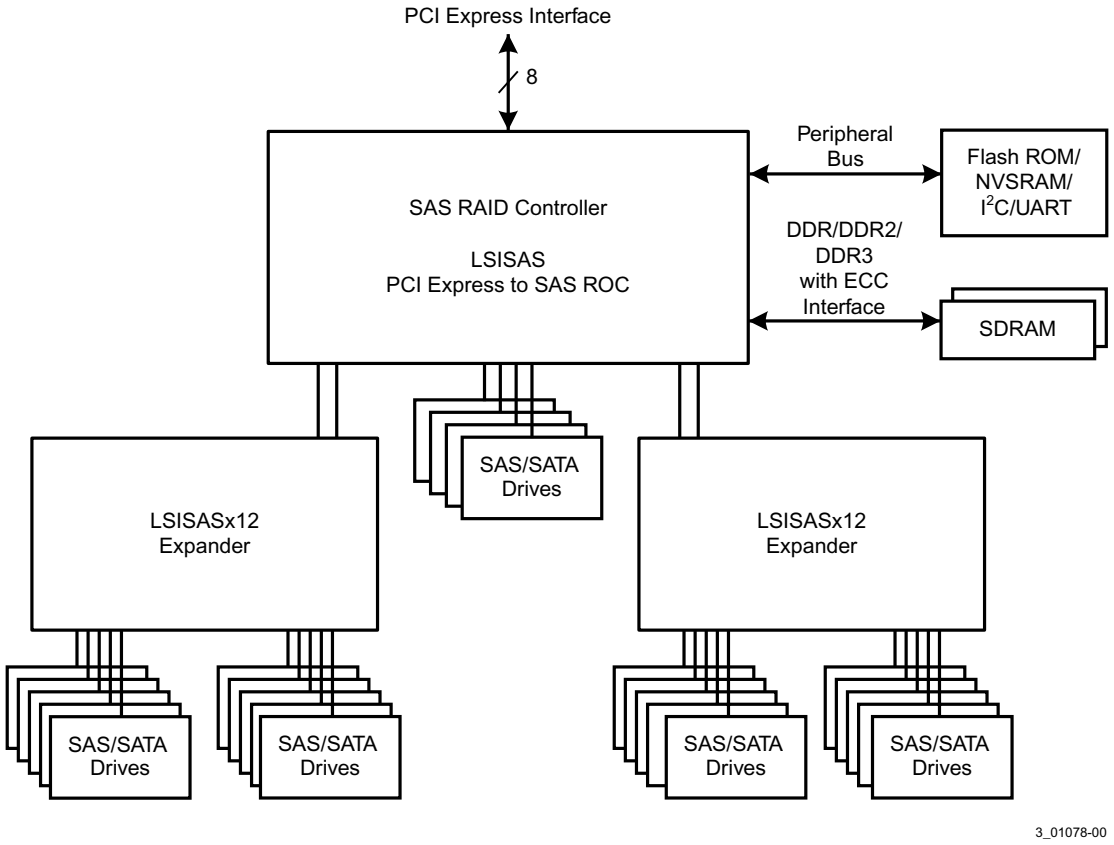

1.7.1 Valid Drive Mix Configurations with HDDs and SSDs

You can allow a virtual drive to consist of both Solid State Drives and Hard Disk Drives. For virtual drives that have both SSDs 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 HDD drives, you can choose whether to mix SAS and SATA HDD drives with SAS and SATA CacheCade software devices in various combinations.

The following table 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 SSD (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)

Table 1 Valid Drive Mix Configurations (Continued)

#	Valid Drive Mix Configurations
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 the above table is allowed based on your controller card manufacturing settings.

NOTE The valid drive mix also applies to hot spares. For information on hot spares, see Hot Spares.

1.8 Technical Support

For assistance with installing, configuring, or running your MegaRAID SAS RAID controllers, contact an Avago Technical Support representative.

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

http://www.avagotech.com/support/submit-storage-support-request

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

E-mail:

http://www.avagotech.com/support/email/megaraid

Phone Support:

http://www.avagotech.com/support/call-us

1-800-633-4545 (North America)

00-800-5745-6442 (International)

+ 49 (0) 8941 352 0123 (Germany)

Chapter 2: Introduction to RAID

This chapter describes a 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.

RAID Description

A Redundant Array of Independent Disks 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. An I/O transaction is expedited because several drives can be accessed simultaneously.

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.

RAID Functions

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

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

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

2.1 Components and Features

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

2.1.1 Drive Group

A drive group is a group of physical drives. These drives are managed in partitions known as virtual drives.

2.1.2 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 these components:

- An entire drive group
- More than one entire drive group

- A part of a drive group
- Parts of more than one drive group
- A combination of any two of these conditions

2.1.3 Fault Tolerance

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

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

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

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

Fault tolerance is often associated with system availability because it allows the system to be available during the failures. However, fault tolerance means that it is also important for the system to be available during the repair of the problem.

A hot spare is an unused drive. You can use a hot spare to rebuild the data and re-establish redundancy in case of a disk failure in a redundant RAID drive group. 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.

2.1.3.1 Multipathing

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

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

Multipathing provides the following features:

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

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

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

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

2.1.4 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 calculating 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.

2.1.5 Replace

The Replace operation lets you copy data from a source drive into a destination drive that is not a part of the virtual drive. The Replace operation often creates or restores a specific physical configuration for a drive group (for example, a specific arrangement of drive group members on the device I/O buses). You can run a Replace operation 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 Replace operation runs as a background activity, and the virtual drive is still available online to the host.

A Replace operation is also initiated when the first 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 that has the SMART error is marked as *failed* only after the successful completion of the Replace operation. This situation avoids putting the drive group in Degraded status.

NOTE During a Replace operation, if the drive group involved in the Replace operation is deleted because of a virtual drive deletion, the destination drive reverts to an Unconfigured Good state or Hot Spare state.

NOTE When a Replace operation is enabled, the alarm continues to beep even after a rebuild is complete; the alarm stops beeping only when the Replace operation is completed.

Order of Precedence

In the following scenarios, a rebuild takes precedence over a Replace operation:

- If a Replace operation is already taking place to a hot spare drive, and any virtual drive on the controller degrades, the Replace operation aborts, and a rebuild starts. A Rebuild operation changes the virtual drive to the Optimal state.
- The Rebuild operation takes precedence over the Replace operation when the conditions exist to start both operations. Consider the following examples:
  - Hot spare is not configured (or unavailable) in the system.
  - Two drives (both members of virtual drives) exist, with one drive exceeding the SMART error threshold, and the other failed.

— If you add a hot spare (assume a global hot spare) during a Replace operation, the Replace operation is ended abruptly, and a Rebuild operation starts on the hot spare.

2.1.6 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 fewer drives exist, the background initialization does not start. The background initialization needs to be started manually. The following number of drives are required:

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

The default and recommended background initialization rate is 30 percent. Before you change the rebuild rate, you must stop the background initialization or the rate change will not affect the background initialization rate. After you stop background initialization and change the rebuild rate, the rate change takes effect when you restart background initialization.

2.1.7 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 Storage Manager software to select the patrol read options, which you can use to set automatic or manual operation, or disable patrol read. See Running a Patrol Read.

2.1.8 Disk Striping

Disk striping lets you 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 a minimum of 64 KB to 1 MB for MegaRAID controllers and 64 KB for Integrated MegaRAID controllers. The LSISAS2108 controller allows stripe size from 8 KB to 1 MB. 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)

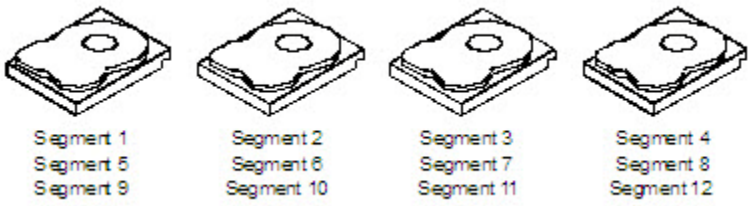

Stripe Width

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.

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 1 MB of drive space and has 64 KB of data residing on each drive in the stripe. In this case, the stripe size is 1 MB and the strip size is 64 KB.

Strip Size

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

2.1.9 Disk Mirroring

With disk 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 run the system and reconstruct the failed disk.

Disk mirroring provides 100 percent redundancy, but it is expensive because each drive in the system must be duplicated. The following figure shows an example of disk mirroring.

Figure 4 Example of Disk Mirroring (RAID 1)

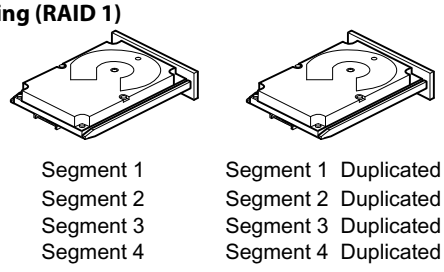

2.1.10 Parity

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 a RAID drive group, 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 the following table.

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.

A RAID 5 drive group 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 the following figure. A RAID 5 drive group uses parity to provide redundancy for one drive failure without duplicating the contents of entire drives. A RAID 6 drive group also uses distributed parity and disk striping, but adds a second set of parity data so that it can survive up to two drive failures.

Figure 5 Example of Distributed Parity (RAID 5 Drive Group)

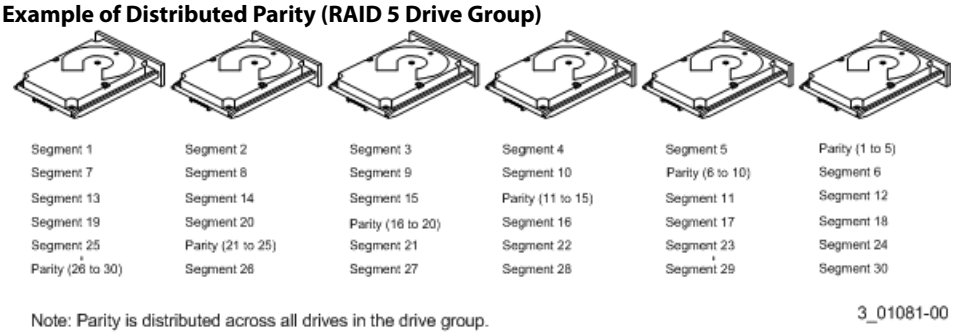

2.1.11 Disk Spanning

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 the following figure, RAID 1 drive groups are turned into a RAID 10 drive group.

Figure 6 Example of Disk Spanning

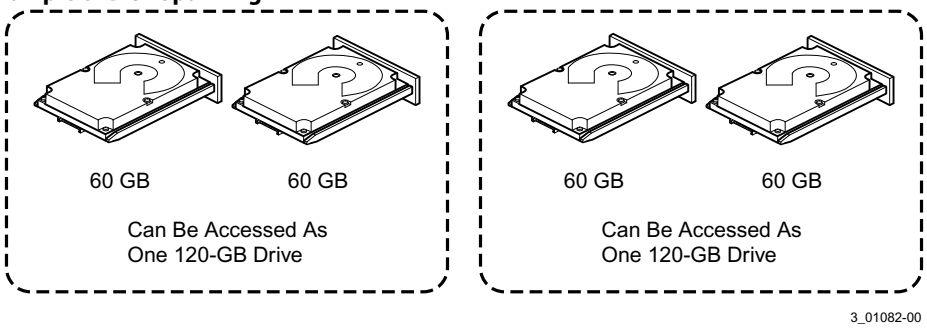

Spanning two contiguous RAID 0 virtual drives does not produce a new RAID level or add fault tolerance. It does increase the capacity of the virtual drive and improves performance by doubling the number of spindles.

Spanning for RAID 00, RAID 10, RAID 50, and RAID 60 Drive Groups

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

See 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 Drive Groups

Level	Description
00	Configure a RAID 00 by spanning two or more contiguous RAID 0 virtual drives, up to the maximum number of supported devices for the controller.
10	Configure RAID 10 by spanning two or more contiguous RAID 1 virtual drives, up to the maximum number of supported devices for the controller. A RAID 10 drive group supports a maximum of 8 spans. You must use an even number of drives in each RAID virtual drive in the span. The RAID 1 virtual drives must have the same stripe size.
50	Configure a RAID 50 drive group by spanning two or more contiguous RAID 5 virtual drives. The RAID 5 virtual drives must have the same stripe size.
60	Configure a RAID 60 drive group by spanning two or more contiguous RAID 6 virtual drives. The RAID 6 virtual drives must have the same stripe size.

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

2.1.12 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 let you replace failed drives without system shutdown or user intervention. The MegaRAID SAS RAID controllers can implement automatic and transparent rebuilds of failed drives using hot spare drives, which provide a high degree of fault tolerance and zero downtime.

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

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

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

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

The hot spare can be of two types:

- Global hot spare
- Dedicated hot spare

Global Hot Spare

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

Dedicated Hot Spare

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

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

Observe the following parameters when using hot spares:

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

  For example, to replace a 500-GB drive, the hot spare must be 500-GB or larger.

2.1.13 Disk Rebuilds

When a drive in a RAID drive group fails, you can rebuild the drive by re-creating the data that was stored on the drive before it failed. The RAID controller re-creates the data using the data stored on the other drives in the drive group. Rebuilding can be performed 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 operation 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 operation to a hot spare begins. If the system goes down during a Rebuild operation, the RAID controller automatically resumes the rebuild after the system reboots.

NOTE When the Rebuild operation to a hot spare begins, the failed drive is often removed from the virtual drive before management applications detect the failed drive. When this removal occurs, the event 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 operation begins to a hot spare. If a source drive fails during a rebuild to a hot spare, the Rebuild operation 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 operation 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 operation will not start if you replace a drive during a RAID-level migration. The Rebuild operation 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.)

2.1.14 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 assigns priority to rebuilding the failed drives.

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

2.1.15 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 operation occurs automatically if these situation occurs:

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

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

2.1.16 Drive States

A drive state is a property indicating the status of the drive. The drive states are described in the following table.

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.
Shield State	An interim state of physical drive for diagnostic operations.
Copyback	A drive that has replaced the failed drive in the RAID configuration.

2.1.17 Virtual Drive States

The virtual drive states are described in the following table.

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. A RAID 6 drive group can tolerate up to two drive failures.
Failed	The virtual drive has failed.
Offline	The virtual drive is not available to the RAID controller.

2.1.18 Beep Codes

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

Table 6 Beep Codes, Events, and Virtual Drive States

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

2.1.19 Enclosure Management

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

2.2 RAID Levels

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

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

2.2.1 Summary of RAID Levels

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

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

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

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

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

A RAID 10 drive group, a combination of RAID 0 and RAID 1 drive groups, 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. A RAID 10 drive group allows a maximum of 8 spans. You must use an even number of drives in each RAID virtual drive in the span. The RAID 1 virtual drives must have the same stripe size. A RAID 10 drive group provides high data throughput and complete data redundancy but uses a larger number of spans.

A RAID 50 drive group, a combination of RAID 0 and RAID 5 drive groups, 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. A RAID 50 drive group 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 Level 0 and RAID Level 5, in the same drive group is not allowed. For example, if an existing RAID 5 virtual drive is created out of partial space in an array, the next virtual drive in the array has to be RAID Level 5 only.

A RAID 60 drive group, a combination of RAID level 0 and RAID Level 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. A RAID 60 drive group works best with data that requires high reliability, high request rates, high data transfers, and medium-to-large capacity.

NOTE The MegaSR controller supports the standard RAID levels – RAID 0, RAID 1, RAID 5, and RAID 10. The MegaSR controller comes in two variants, SCU and AHCI, both supporting a maximum of eight physical drives. A maximum of eight virtual drives can be created (using RAID 0, RAID 1, RAID 5, and RAID 10 only) and controlled by the MegaSR controller. One virtual drive can be created on an array (a maximum of eight if no other virtual drives are already created on the MegaSR controller), or you can create eight arrays with one virtual drive each. However, on a RAID 10 drive group, you can create only one virtual drive on a particular array.

2.2.2 Selecting a RAID Level

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

2.2.3 RAID 0 Drive Groups

A RAID 0 drive group provides disk striping across all drives in the RAID drive group. A RAID 0 drive group does not provide any data redundancy, but the RAID 0 drive group offers the best performance of any RAID level. The RAID 0 drive group 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. A RAID 0 drive group offers high bandwidth.

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

By breaking up a large file into smaller segments, the RAID controller can use both SAS drives and SATA drives to read or write the file faster. A RAID 0 drive group involves no parity calculations to complicate the write operation. This situation makes the RAID 0 drive group ideal for applications that require high bandwidth but do not require fault tolerance. The following table provides an overview of the RAID 0 drive group. The following figure provides a graphic example of a RAID 0 drive group.

Table 7 RAID 0 Drive Group Overview

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

Figure 7 RAID 0 Drive Group Example with Two Drives

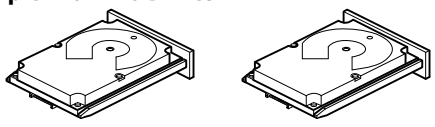

Segment 1 Segment 2
Segment 3 Segment 4
Segment 5 Segment 6
Segment 7 Segment 8

3_01083-00

2.2.4 RAID 1 Drive Groups

In RAID 1 drive groups, the RAID controller duplicates all data from one drive to a second drive in the drive group. A RAID 1 drive group supports an even number of drives from 2 through 32 in a single span. The RAID 1 drive group provides complete data redundancy, but at the cost of doubling the required data storage capacity. The following

table provides an overview of a RAID 1 drive group. The following figure provides a graphic example of a RAID 1 drive group.

Table 8 RAID 1 Drive Group Overview

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

Figure 8 RAID 1 Drive Group

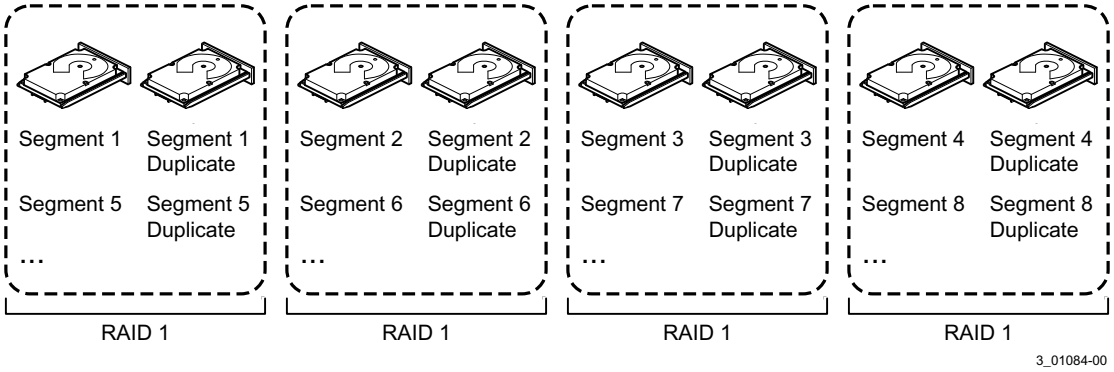

2.2.5 RAID 5 Drive Groups

A RAID 5 drive group 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 drive groups, the parity information is written to all drives. A RAID 5 drive group is best suited for networks that perform a lot of small input/output (I/O) transactions simultaneously.

The following table provides an overview of a RAID 5 drive group. The following figure provides a graphic example of a RAID 5 drive group.

Table 9 RAID 5 Drive Group Overview

Uses	Provides high data throughput, especially for large files. Use RAID 5 drive groups for transaction processing applications because each drive can read and write independently. If a drive fails, the RAID controller uses the parity drive to re-create all missing information. Use also for online customer service that requires fault tolerance. Use for any application that has high read request rates but random 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 lots of small writes or small block write operations. Suffers more impact if no cache is used. Drive performance is reduced if a drive is being rebuilt. Environments with few processes do not perform as well because the RAID drive group overhead is not offset by the performance gains in handling simultaneous processes.
Number of drives in this RAID level	3 through 32

Figure 9 RAID 5 Drive Group with Six Drives

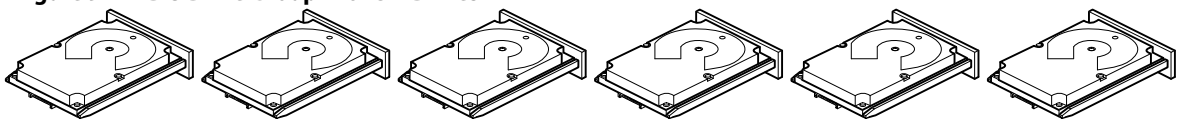

Segment 1	Segment 2	Segment 3	Segment 4	Segment 5	Parity (1–5)
Segment 7	Segment 8	Segment 9	Segment 10	Parity (6–10)	Segment 6
Segment 13	Segment 14	Segment 15	Parity (11–15)	Segment 11	Segment 12
Segment 19	Segment 20	Parity (16–20)	Segment 16	Segment 17	Segment 18
Segment 25	Parity (21–25)	Segment 21	Segment 22	Segment 23	Segment 24
Parity (26–30)	Segment 26	Segment 27	Segment 28	Segment 29	Segment 30

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

3_01085-00

2.2.6 RAID 6 Drive Groups

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

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

The following table provides an overview of a RAID 6 drive group.

Table 10 RAID 6 Drive Group Overview

Uses	Use for any application that has high read request rates but low random or small block write 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. Performance is similar to that of a RAID 5 drive group.
Weak points	Not well-suited to tasks requiring a lot of small and/or random write operations. A RAID 6 virtual drive must generate two sets of parity data for each write operation, which results in a significant decrease in performance during write operations. Drive performance is reduced during a drive Rebuild operation. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes. A RAID 6 drive group costs more because of the extra capacity required by using two parity blocks per stripe.
Drives	3 through 32.

The following figure shows a RAID 6 drive group data layout. The second set of parity drives is denoted by *Q*. The *P* drives follow the RAID 5 drive group parity scheme.

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

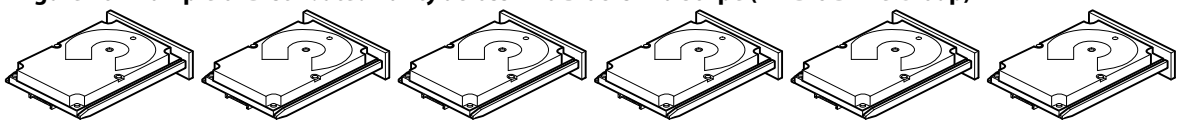

Segment 1	Segment 2	Segment 3	Segment 4	Parity (P1–P4)	Parity (Q1–Q4)
Segment 6	Segment 7	Segment 8	Parity (P5–P8)	Parity (Q5–Q8)	Segment 5
Segment 11	Segment 12	Parity (P9–P12)	Parity (Q9–Q12)	Segment 9	Segment 10
Segment 16	Parity (P13–P16)	Parity (Q13–Q16)	Segment 13	Segment 14	Segment 15
Parity (P17–P20)	Parity (Q17–Q20)	Segment 17	Segment 18	Segment 19	Segment 20

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

3_01086-00

2.2.7 RAID 00 Drive Groups

A RAID 00 drive group is a spanned drive group that creates a striped set from a series of RAID 0 drive groups. A RAID 00 drive group does not provide any data redundancy, but, along with the RAID 0 drive group, does offer the best performance of any RAID level. A RAID 00 drive group 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. A RAID 00 drive group offers high bandwidth.

NOTE RAID level 00 is not fault tolerant. If a drive in a RAID 0 drive group fails, the entire virtual drive (all drives associated with the virtual drive) fails.

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

tolerance. The following table provides an overview of the RAID 00 drive group. The following figure provides a graphic example of a RAID 00 drive group.

Table 11 RAID 00 Drive Group Overview

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

Figure 11 RAID 00 Drive Group Example with Two Drives

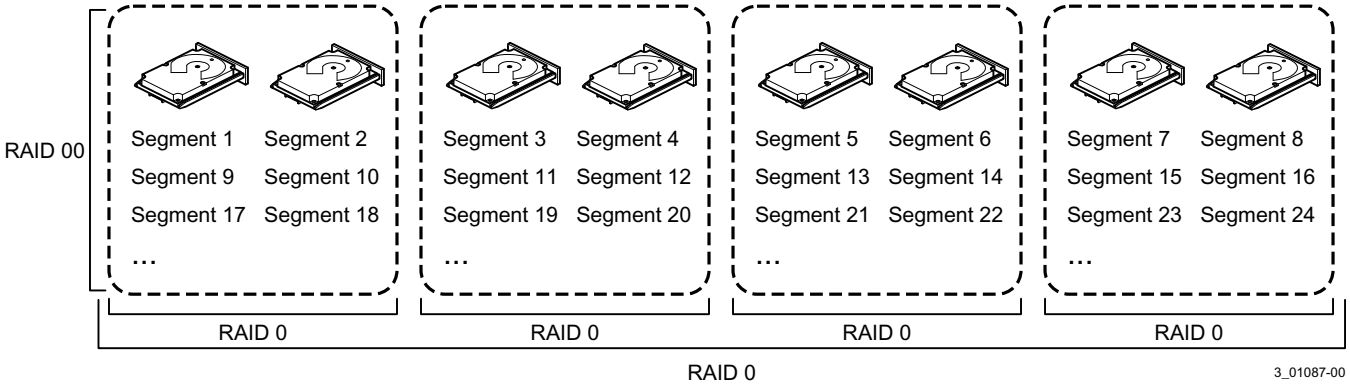

2.2.8 RAID 10

A RAID 10 drive group is a combination of RAID level 0 and RAID level 1, and it consists of stripes across mirrored drives. A RAID 10 drive group 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 level 1 drive groups are referred to as RAID level 10, (1+0). Data is striped across drive groups to increase performance by enabling access to multiple drive groups simultaneously.

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

Configure RAID 10 drive groups by spanning two contiguous RAID 1 virtual drives, up to the maximum number of supported devices for the controller. A RAID 10 drive group supports a maximum of 8 spans, with a maximum of 32 drives per span. You must use an even number of drives in each RAID 10 virtual drive in the span.

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

The following table provides an overview of a RAID 10 drive group.

Table 12 RAID 10 Drive Group 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.) A RAID 10 drive group 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 in RAID 1 drive groups.
Drives	4 to 32 in multiples of 4 — The maximum number of drives supported by the controller (using an even number of drives in each RAID 10 virtual drive in the span).

In the following figure, virtual drive 0 is created by distributing data across four drive groups (drive groups 0 through 3).

Figure 12 RAID 10 Level Virtual Drive

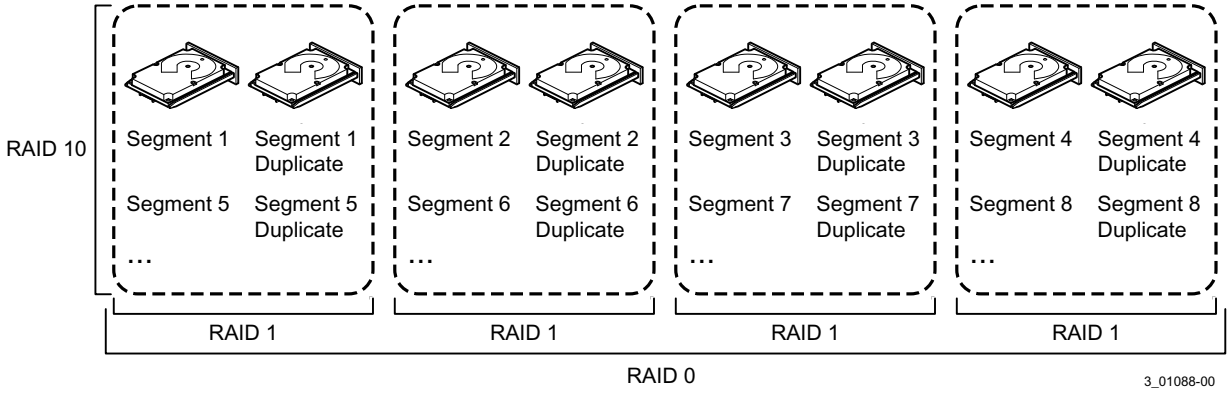

2.2.9 RAID 50 Drive Groups

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

A RAID 50 drive group breaks up data into smaller blocks and then stripes the blocks of data to each RAID 5 disk set. A RAID 5 drive group breaks up data into smaller blocks, calculates parity by performing an exclusive OR operation on the blocks, and then performs write operations to 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.

A RAID level 50 drive group 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.

The following table provides an overview of a RAID 50 drive group.

Table 13 RAID 50 Drive Group Overview

Uses	Appropriate when used with data that requires high reliability, high request rates, high data transfer, and medium-to-large capacity. Also used when a virtual drive of greater than 32 drives is needed.
Strong points	Provides high data throughput, data redundancy, and very good performance.
Weak points	Requires two times to eight times as many parity drives as a RAID 5 drive group.
Drives	Eight spans of RAID 5 drive groups that contain 3 to 32 drives each (limited by the maximum number of devices supported by the controller)

Figure 13 RAID 50 Level Virtual Drive

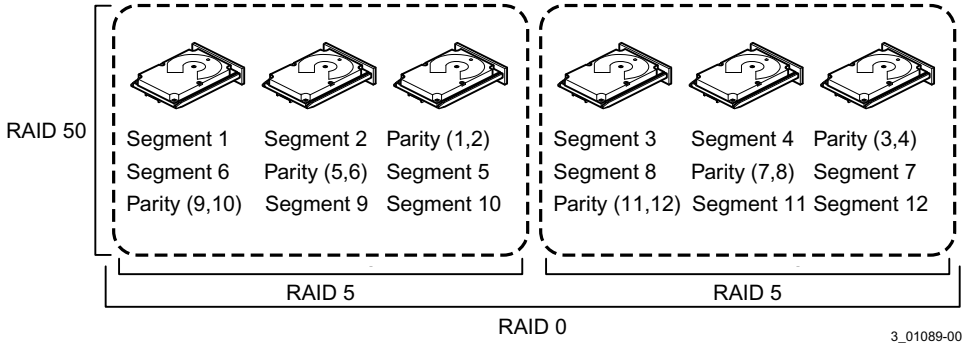

2.2.10 RAID 60 Drive Groups

A RAID 60 drive group provides the features of both RAID 0 and RAID 6 drive groups, and includes both parity and disk striping across multiple drive groups. A RAID 6 drive group 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 drive group sets without losing data. A RAID 60 drive group is best implemented on two RAID 6 drive groups with data striped across both drive groups.

A RAID 60 drive group breaks up data into smaller blocks and then stripes the blocks of data to each RAID 6 disk set. A RAID 6 drive group breaks up data into smaller blocks, calculates parity by performing an exclusive-OR operation on the blocks, and then performs write operations to the blocks of data and writes the 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.

A RAID 60 drive group 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 Drive Group Overview

Uses	Provides a high level of data protection through the use of a second parity block in each stripe. Use a RAID 60 drive group 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 re-create all of the missing information. If two drives in a RAID 6 set in a RAID 60 virtual drive fail, two drive Rebuild operations are required, one for each drive. These Rebuild operations can occur at the same time. Use for online customer service that requires fault tolerance. Use for any application that has high read request rates but low write request rates. Also used when a virtual drive of greater than 32 drives is needed.
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.
Weak points	Not well-suited for small block write or random write operations. A RAID 60 virtual drive must generate two sets of parity data for each write operation, which results in a significant decrease in performance during write operations. Drive performance is reduced during a drive Rebuild operation. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes. A RAID 6 drive group costs more because of the extra capacity required by using two parity blocks per stripe.
Drives	A minimum of 6.

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

Figure 14 RAID 60 Level Virtual Drive

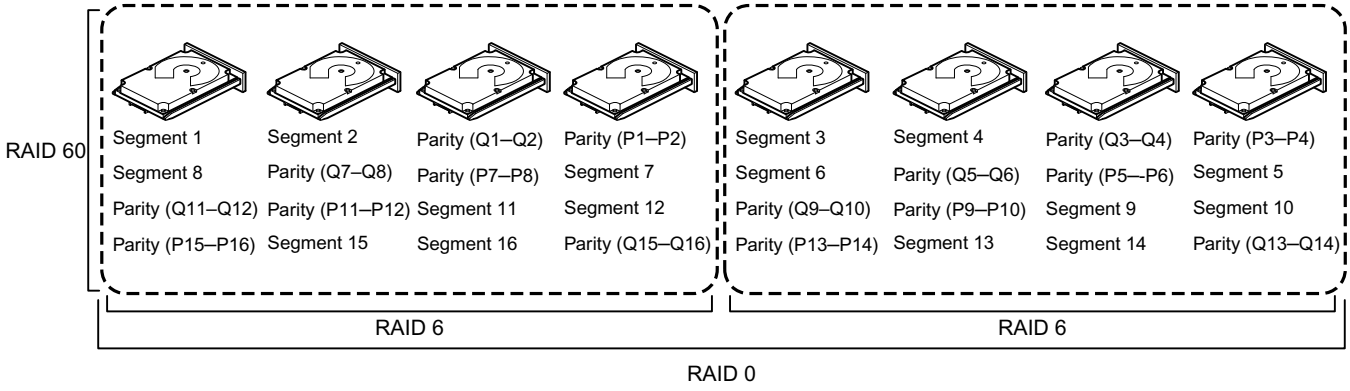

Note: Parity is distributed across all drives in the drive group.

3_01090-00

2.3 RAID Configuration Strategies

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

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

You cannot configure a virtual drive that optimizes all three factors, but it is easy to choose a virtual drive configuration that maximizes one factor at the expense of another factor. For example, RAID 1 (mirroring) provides excellent fault tolerance, but requires a redundant drive.

The following subsections describe how to use the RAID levels to maximize virtual drive availability (fault tolerance), virtual drive performance, and virtual drive capacity.

2.3.1 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 operation occurs.

A *hot swap* is the manual substitution of a replacement unit in a disk subsystem for a defective one, where the substitution can be performed while the subsystem is running hot swap drives. An Auto-Rebuild feature 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 operation occurs, which provides 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. A RAID 0 drive group is ideal for applications that require high performance but do not require fault tolerance.
1	Provides complete data redundancy. If one drive fails, the contents of the other drive in the drive group can be used to run the system and reconstruct the failed drive. The primary advantage of disk mirroring is that it provides 100 percent data redundancy. Because the contents of the drive are completely written to a second drive, no data is lost if one of the drives fails. Both drives contain the same data at all times. A RAID 1 drive group 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 a RAID 5 drive group, this method is applied to entire drives or stripes across all drives in a drive group. Using distributed parity, a RAID 5 drive group offers fault tolerance with limited overhead.
6	Combines distributed parity with disk striping. A RAID 6 drive group 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 a RAID 6 drive group, this method is applied to entire drives or stripes across all of the drives in a drive group. Using distributed parity, a RAID 6 drive group 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. A RAID 00 drive group 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. A RAID 10 drive group works well for any environment that requires the 100 percent redundancy offered by mirrored drive groups. A RAID 10 drive group can sustain a drive failure in each mirrored drive group and maintain data integrity.
50	Provides data redundancy using distributed parity across spanned RAID 5 drive groups. A RAID 50 drive group includes both parity and disk striping across multiple drives. If a drive fails, the RAID controller uses the parity data to re-create all missing information. A RAID 50 drive group 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. A RAID 60 drive group 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. A RAID 60 drive group includes both parity and disk striping across multiple drives. If a drive fails, the RAID controller uses the parity data to re-create all missing information.

2.3.2 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. The I/O performs faster because drives can be accessed simultaneously. The following table 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 a minimum of 64 KB to 1 MB for MegaRAID controllers and only 64 KB for Integrated MegaRAID controllers. The LSISAS2108 controller allows strip size from 8 KB to 1 MB. These stripes are interleaved in a repeated sequential manner. Disk striping enhances performance because multiple drives are accessed simultaneously.
1	With a RAID 1 (mirroring) drive group, each drive in the system must be duplicated, which requires more time and resources than striping. Performance is impaired during drive Rebuild operations.
5	A RAID 5 drive group provides high data throughput, especially for large files. Use this RAID level for any application that requires high read request rates, but low write request rates, such as transaction processing applications, because each drive can read and write independently. Because each drive contains both data and parity, numerous write operations can take place concurrently. In addition, robust caching algorithms and hardware-based exclusive-or assist make RAID 5 drive group performance exceptional in many different environments. Parity generation can slow the write process, making write performance significantly lower for RAID 5 drive group than for RAID 0 or RAID 1 drive groups. 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	A RAID 6 drive group 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, a RAID 6 drive group is not well suited to tasks requiring a lot of write operations. A RAID 6 virtual drive must generate two sets of parity data for each write operation, which results in a significant decrease in performance during write operations. 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	A RAID 00 drive group (striping in a spanned drive group) offers excellent performance. A RAID 00 drive group 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 a minimum of 64 KB to 1 MB for MegaRAID controllers and only 64 KB for Integrated MegaRAID controllers. The LSISAS2108 controller allows strip size from 8 KB to 1 MB. 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	A RAID 10 drive group works best for data storage that need the enhanced I/O performance of a RAID 0 drive group (striped drive groups), which provides high data transfer rates. Spanning increases the capacity of the virtual drive and improves performance by doubling the number of spindles. The system performance improves as the number of spans increases. (The maximum number of spans is 8.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans, and RAID performance degrades to that of a RAID 1 or RAID 5 drive group.
50	A RAID 50 drive group works best when used with data that requires high reliability, high request rates, and high data transfer. It provides high data throughput, data redundancy, and very good performance. Spanning increases the capacity of the virtual drive and improves performance by doubling the number of spindles. The system performance improves as the number of spans increases. (The maximum number of spans is 8.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans and RAID drive group performance degrades to that of a RAID 1 or RAID 5 drive group.
60	A RAID 60 drive group works best when used with data that requires high reliability, high request rates, and high data transfer. It provides high data throughput, data redundancy, and very good performance. Spanning increases the capacity of the virtual drive and improves performance by doubling the number of spindles. The system performance improves as the number of spans increases. (The maximum number of spans is 8.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans, and RAID performance degrades to that of a RAID 1 or RAID 6 drive group. A RAID 60 drive group is not well suited to tasks requiring a lot of writes. A RAID 60 virtual drive must generate two sets of parity data for each write operation, which results in a significant decrease in performance during write operations. 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.

2.3.3 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 drive group) or distributed parity (RAID 5 or RAID 6 drive group). A RAID 5 drive group, which provides redundancy for one drive failure without duplicating the contents of entire drives, requires less space than a RAID 1 drive group. The following table explains the effects of the RAID levels on storage capacity.

Table 17 RAID Levels and Capacity

RAID Level	Capacity
0	A RAID 0 drive group (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. A RAID 0 drive group provides maximum storage capacity for a given set of drives. The usable capacity of a RAID 0 array is equal to the number of drives in the array into the capacity of the smallest drive in the array.
1	With a RAID 1 drive group (mirroring), data written to one drive is simultaneously written to another drive, which doubles the required data storage capacity. This situation is expensive because each drive in the system must be duplicated. The usable capacity of a RAID 1 array is equal to the capacity of the smaller of the two drives in the array.
5	A RAID 5 drive group provides redundancy for one drive failure without duplicating the contents of entire drives. The RAID 5 drive group breaks up data into smaller blocks, calculates parity by performing an exclusive-or on the blocks and then writes the blocks of data and parity to each drive in the drive group. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set. The usable capacity of a RAID 5 array is equal to the number of drives in the array, minus one, into the capacity of the smallest drive in the array.
6	A RAID 6 drive group 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 a RAID 60 drive group more expensive to implement. The usable capacity of a RAID 6 array is equal to the number of drives in the array, minus two, into the capacity of the smallest drive in the array.

Table 17 RAID Levels and Capacity (Continued)

RAID Level	Capacity
00	A RAID 00 drive group (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. A RAID 00 drive group provides maximum storage capacity for a given set of drives.
10	A RAID 10 drive group requires twice as many drives as all other RAID levels except RAID level 1. A RAID 10 drive group works well for medium-sized databases or any environment that requires a higher degree of fault tolerance and moderate-to-medium capacity. Disk spanning allows multiple drives to function like one large drive. Spanning overcomes lack of disk space and simplifies storage management by combining existing resources or adding relatively inexpensive resources.
50	A RAID 50 drive group requires two to four times as many parity drives as a RAID 5 drive group. This RAID level works best when used with data that requires medium to large capacity.
60	A RAID 60 drive group provides redundancy for two drive failures in each RAID set without duplicating the contents of entire drives. However, it requires extra capacity because a RAID 60 virtual drive has to generate two sets of parity data for each write operation. This situation makes a RAID 60 drive group more expensive to implement.

2.4 RAID Availability

2.4.1 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.

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 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.

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. A manual rebuild is necessary if hot spares with enough capacity to rebuild the failed drives are not available. You must insert a drive with enough storage into the subsystem before rebuilding the failed drive.

2.5 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 to optimize 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.

2.6 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.

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, RAID 6, RAID 10, RAID 50, or RAID 60.
- Does this drive group support any software system that must be available 24 hours per day?
  Use RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, or RAID 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 RAID 00.
- Will this drive group contain data from an imaging system?
  Use RAID 0, RAID 00, or RAID 10.

Fill out the following table to help you plan the drive group configuration. Rank the requirements for your drive group, such as storage space and data redundancy, in order of importance, and then review the suggested RAID levels.

Table 18 Factors to Consider for Drive Group Configuration

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

Chapter 3: SafeStore Disk Encryption

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

Overview

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

With the SafeStore encryption service, data is encrypted by the drives. You can designate which data to encrypt at the individual virtual drive (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 and the MegaRAID Storage Manager software offer procedures that you can use to manage the security settings for the drives.

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.

3.1 Terminology

The following table describes the terminology related to the SafeStore encryption feature.

Table 19 Terminology Used in the SafeStore Encryption Feature

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.
Key backup	You need to provide the controller with a lock key if the controller is replaced or if you choose to migrate secure virtual disks. To do this task, you must back up the security key.
Re-provisioning	Re-provisioning disables the security system of a device. For a controller, it involves destroying the security key. For SafeStore encrypted drives, when the drive lock key is deleted, the drive is unlocked and any user data on the drive is securely deleted. This situation does not apply to controller-encrypted drives, because deleting the virtual disk destroys the encryption keys and causes a secure erase. See 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. 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.

3.2 Workflow

Overview

The SafeStore workflow follows:

1. Activate the SafeStore key in the software.
2. Enable SafeStore on the controller.
3. Use a compatible SED drive.
4. Enable encryption when the virtual drive is created with the SED drives.
5. Create a security key that conforms to the security requirements.
6. You can configure the system with the desired password.
7. After the system is booted, you need not enter the password again to access the virtual drives.
8. If the virtual drive is moved to a different controller, then the controller to which the virtual drive is moved, in order to access the data must have the following features:
   - SafeStore enabled.
   - Encryption enabled.
   - The security key must be entered.

3.2.1 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)

Create the Security Key Identifier

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

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. The security key must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (for example, < > @ +).

ATTENTION If you forget the security key, you lose access to the data if you are prompted for the security key again.

Create a Password

Password creation is optional. If you create a password, (referred to as a *passphrase* in StorCLI) it causes the controller to stop during POST and requests a password. If the correct password is not provided, the data on that virtual drive is not accessible. If the virtual drive is a boot device, booting is not possible. The password (*passphrase*) can be the same as the security key. The security key must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (for example, < > @ +).

ATTENTION If you forget the password and you reboot, you will lose access to your data.

3.2.2 Change Security

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

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

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

See Avago MegaRAID SafeStore Encryption Services for the procedures used to change security options in the MegaRAID Storage Manager software.

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.

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.

Add or Change the Password

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

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

3.2.3 Create Secure Virtual Drives

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

Simple Configuration

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

See Creating a Virtual Drive Using Simple Configuration, for the procedures used to select the redundancy type and drive security method for a configuration.

Advanced Configuration

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

See 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.

3.2.4 Import a Foreign Configuration

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

Configuration Utility and the MegaRAID Storage Manager software allows you to import the existing configuration to the RAID controller or clear the configuration so you can create a new one.

See Importing or Clearing a Foreign Configuration, for the procedure in the MegaRAID Storage Manager software.

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

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

3.3 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 read operations) and from the host to the drive cache (on write operations) 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 must manage a password if they are locked. Even if you do not lock the drives, a benefit still exists to using encrypted disks.

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

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

Consider the following reasons for using instant secure erase.

To repurpose the hard drive for a different application

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

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.

To return a disk for warranty activity

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

Chapter 4: Ctrl-R Utility

This chapter describes the Ctrl-R Utility, a BIOS configuration utility, that lets you create and manage RAID configurations on Avago SAS controllers. You can configure the drive groups and drives on the system before the operating system has been installed.

4.1 Overview

The Ctrl-R Utility resides in the SAS controller BIOS and operates independently of the operating system.

You can use the Ctrl-R Utility to perform tasks such as these:

- Create drive groups and virtual drives for storage configurations
- View controller, physical drive, virtual drive, enclosure, and battery backup unit (BBU) properties, and change parameters
- Delete virtual drives
- Modify power settings
- Import and clear foreign configurations
- Initialize virtual drives
- Check configurations for data consistency
- Create CacheCade virtual drives

4.2 Starting the Ctrl-R Utility

When you boot the system, perform the following steps to start the Ctrl-R Utility:

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

   ```
   Copyright© AVAGO Technologies
   Press <Ctrl><R> for Ctrl-R
   ```

2. Based on the controllers on the system, one of the two following scenarios occurs:
   - If the system has multiple SAS controllers, a controller selection dialog appears. Select a controller and press Enter. The Ctrl-R Utility main menu screen appears.
   - If the system has only one SAS controller, the Ctrl-R Utility main menu screen appears.

4.3 Exiting the Ctrl-R Utility

To exit the Ctrl-R Utility, perform these steps:

1. Perform one of these actions:
   - If you are not in a dialog, press Esc once.
   - If you are in a dialog, press Esc twice (once to exit the dialog, and the second time to exit the utility).

   A confirmation message box appears.
2. Press **OK** to exit the utility.

4.4 Ctrl-R Utility Keystrokes

The following table lists the keystrokes that you can use in the Ctrl-R Utility to navigate between the screens.

Table 20 Ctrl-R Utility Keystrokes

Keystroke	Action
F1	Displays help for the particular screen that you are in.
F2	Displays a list of commands that can be performed for the selected device. This key stroke is available only in the VD Mgmt, the PD Mgmt, and the Foreign View menus. The commands that are enabled are highlighted in white and the disabled commands are highlighted in black. **NOTE** Based on the configurations that you make, commands are enabled or disabled.
F5	Refreshes the screen that you currently are in.
F11	Switches between controllers.
F12	Displays a list of all the available controllers. You can also scroll to the next controller.
<Ctrl><N>	Displays the next menu screen.
<Ctrl><P>	Displays the previous menu screen
<Ctrl><S>	shortcut key for the Apply button in the Controller Settings screens.
<Tab>	Moves the cursor to the next control.
<Shift><Tab>	Moves the cursor to the previous control on a screen or a dialog.
<Enter>	Lets you to select a menu item, a button, a check box and values in a list box.
<Esc>	Closes a screen or a window. Press Esc twice to exit from the Ctrl-R Utility.
Up Arrow	Moves the cursor to the next menu selection.
Down Arrow	Moves the cursor to the lower menu items or to a lower level menu.
Right Arrow	Opens a submenu, moves from a menu heading to the first submenu, or moves to the first item in a submenu. The right arrow also closes a menu list in a popup window.
Left Arrow	Closes a submenu, moves from a menu item to the menu heading or moves from a sub menu to a higher level menu.
Spacebar	Lets you select a menu item, a button and a check box.

4.5 Ctrl-R Utility Menus

The Ctrl-R Utility contains the following menus:

- **VD Mgmt**
- **PD Mgmt**
- **Ctrl Mgmt**
- **Properties**
- **Foreign View**

4.5.1 Virtual Drive Management Menu

The **VD Mgmt** tab is the first menu screen that appears when you start the Ctrl-R Utility.

Figure 15 Virtual Drive Management Screen

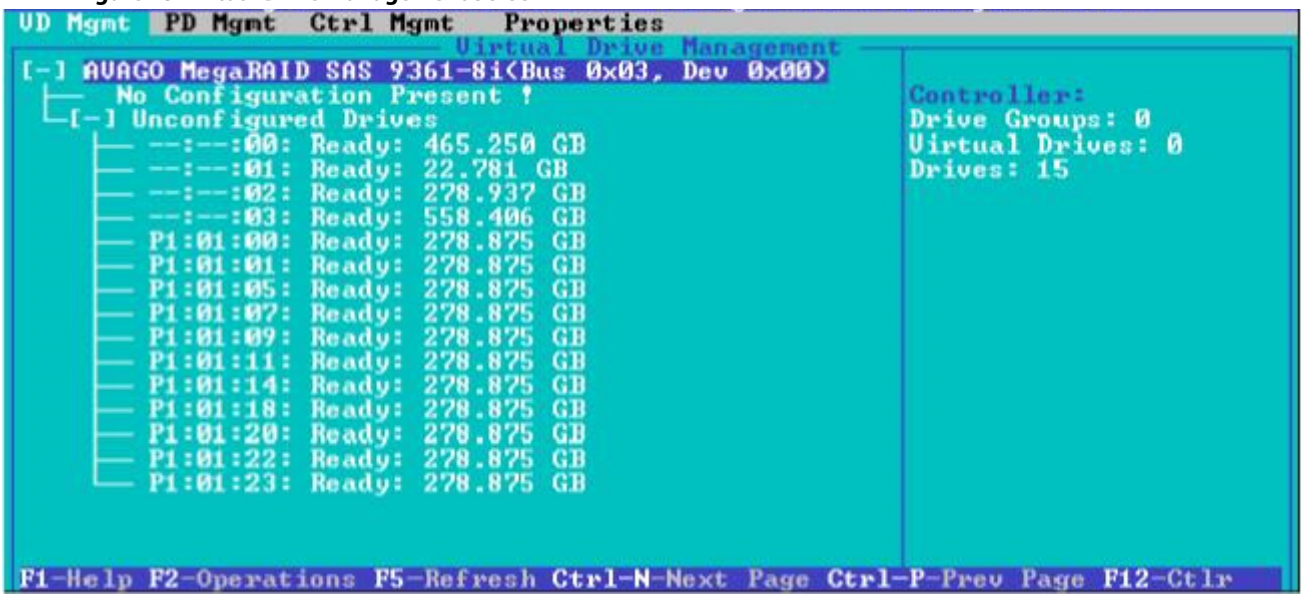

This screen shows information on the configuration of controllers, drive groups, and virtual drives.

The right panel of the screen shows attributes of the selected device.

In the **Virtual Drive Management** screen, you can perform tasks, such as creating and initializing virtual drives; performing a consistency check; deleting, expanding, and erasing virtual drives; importing or clearing foreign configurations; and creating CacheCade virtual drives.

NOTE Based on the controller settings that you make, options will be enabled or disabled.

4.5.2 Physical Drive Management Menu

The **PD Mgmt** tab shows information about all the physical drives connected to the selected controller. This menu also shows information about enclosures, the number of physical drives in an enclosure, and all of the direct-attached drives under a backplane node.

Figure 16 Physical Drive Management Screen

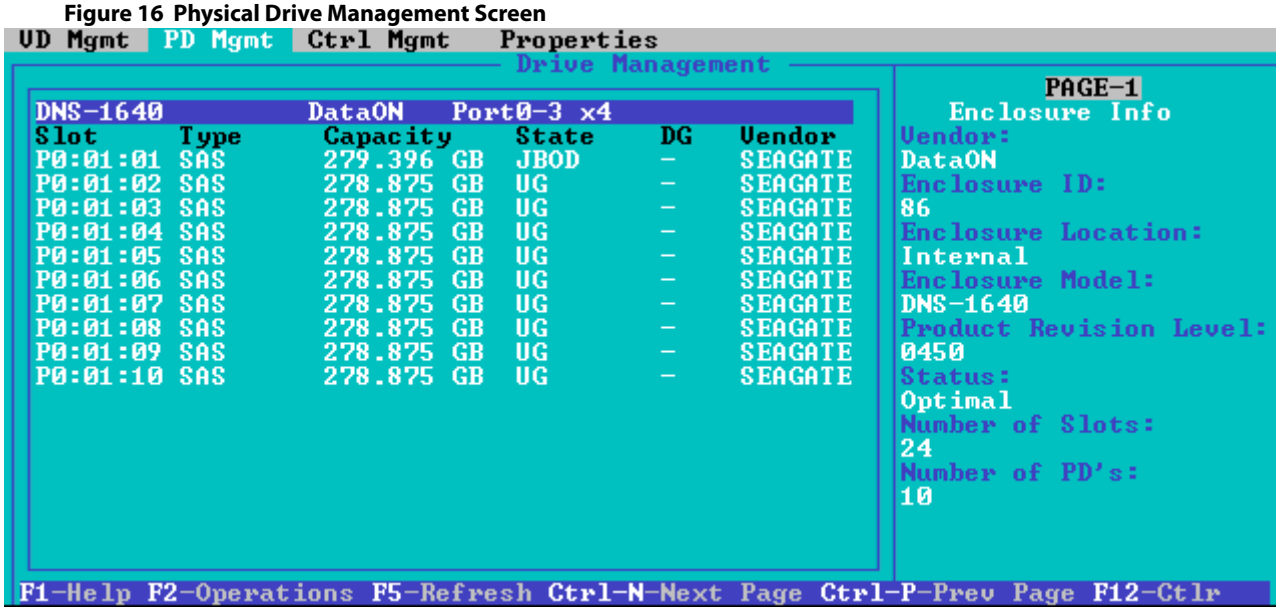

The right panel of the screen shows additional attributes of the selected device.

In the **Drive Management** screen, you can perform tasks, such as rebuilding a failed drive, making a drive offline, or making it a global hot spare drive.

4.5.3 Controller Management Menu

The **Ctrl Mgmt** tab lets you change the settings of the selected controller. The **Ctrl Mgmt** menu consists of two screens.

In the first **Controller Settings** screen (as shown in the following figure), you can change controller options, such as **Maintain PD Fail History**, **Enable Controller BIOS**, **Enable Stop CC on Error**, **Auto Enhanced Import**, and **Enable JBOD**. You also can perform tasks, such as enabling or silencing an alarm, entering values for Rebuild Rate and Patrol Rate, and enabling or disabling the JBOD mode. If you enable the JBOD mode, the drive comes up as JBOD; otherwise, the drive comes up as Unconfigured Good.

NOTE When you disable the JBOD mode, if one or more selected JBODs have an operating system or a file system, a warning message appears indicating that the JBODs contain an operating system or a file system. If you want to proceed, click **Yes**. Otherwise, click **No** to return to the previous screen.

Figure 17 Controller Settings – First Screen

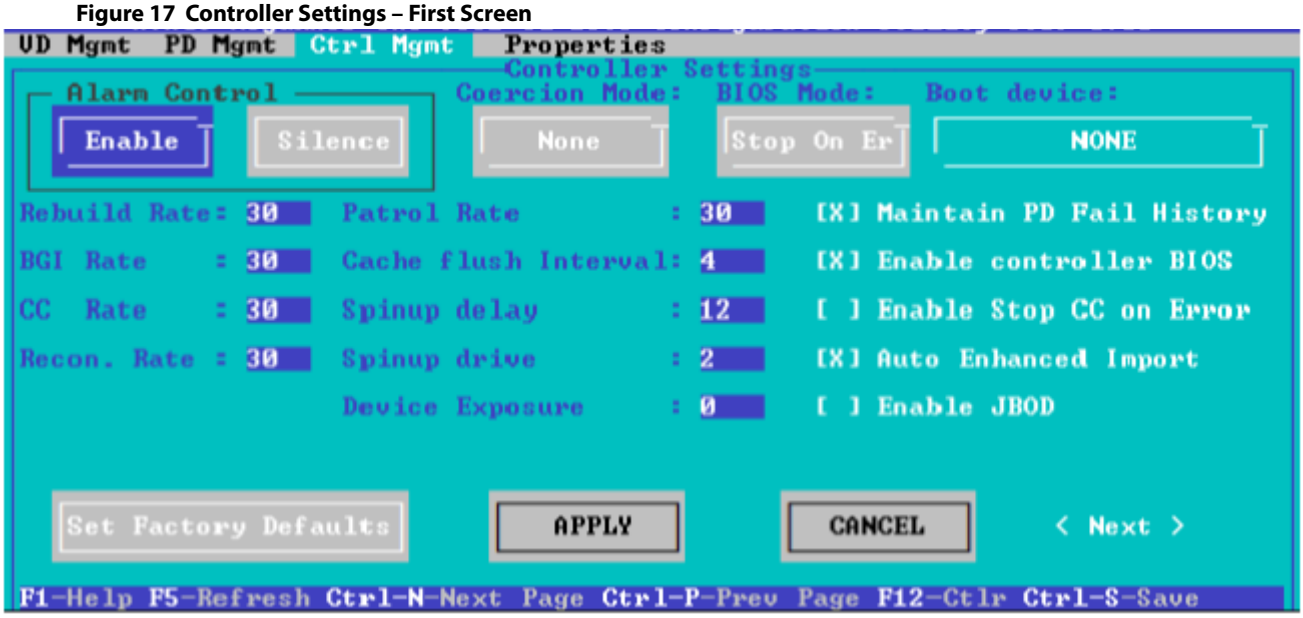

Click **Next** to open the **Controller Settings** screen (as shown in the following figure). You can manage the Link Speed, Power Save, battery settings, begin a Start Manual Learn Cycle, enable or disable Write Verify, and enable or disable large I/O support.

You can enable the **Write Verify** option to verify if the data was written correctly to the cache before flushing the controller cache.

Figure 18 Controller Settings – Second Screen

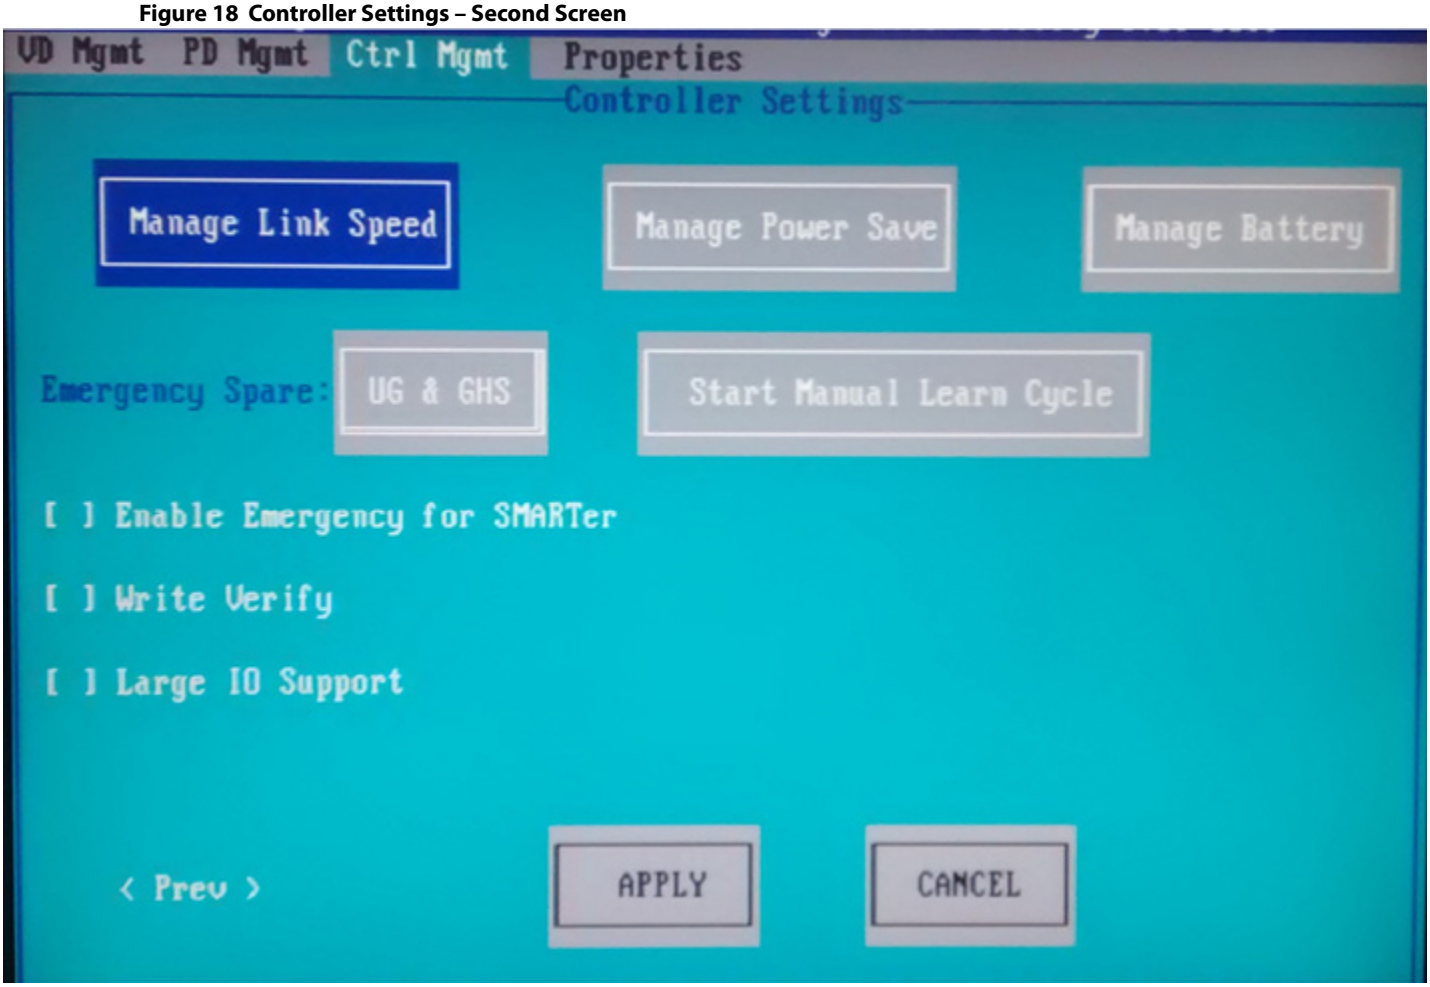

4.5.4 Properties Menu

The **Properties** menu shows all of the properties of the active controller. The **Properties** menu consists of two screens. The information shown in these screens is read only.

In the first **Properties** screen (as shown in the following figure), you can view properties, such as controller status, firmware version, BIOS version, and metadata size.

Figure 19 Properties

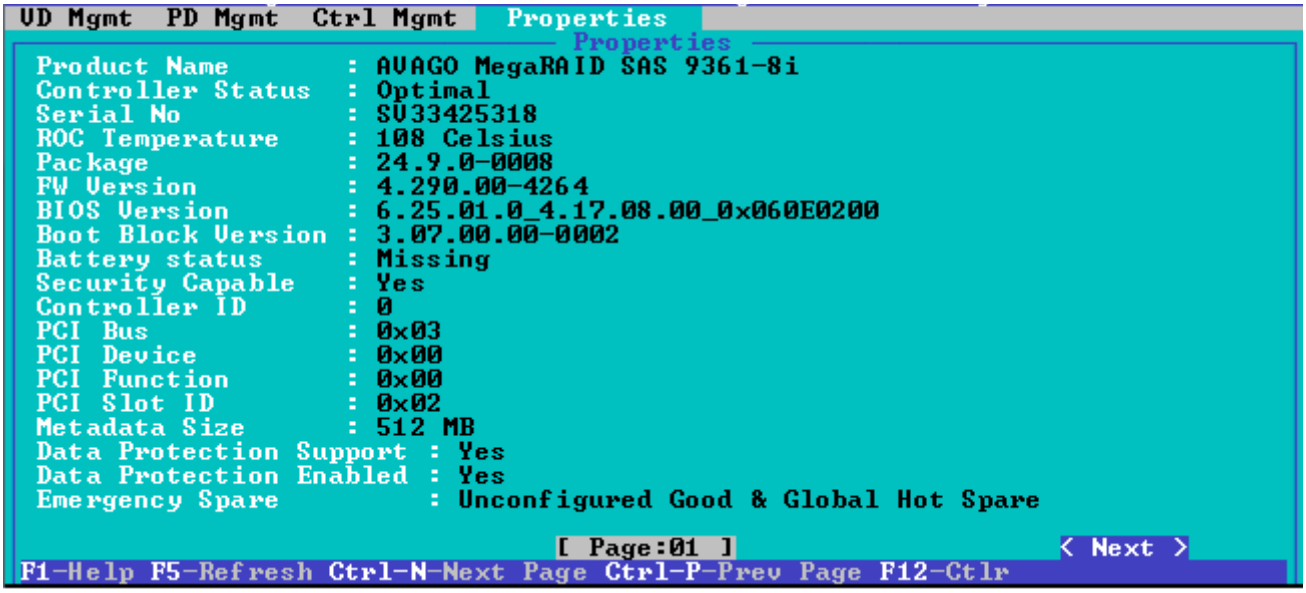

To view additional properties, you can navigate to **Next** and press Enter. The second **Properties** screen shows information, such as maximum cache size, drive standby time, battery status, and power saving properties.

To go back to the previous **Properties** screen, navigate to **Prev**, and press Enter.

4.5.5 Foreign View Menu

If one or more physical drives in a configuration are removed and reinserted, the controller considers the drives as foreign configurations.

The **Foreign View** tab is shown only when the controller detects a foreign configuration. If no foreign configurations exists, the **Foreign View** tab is not shown.

Figure 20 Foreign View Menu

You can use the **Foreign Config View** screen to view information about the foreign configuration, such as drive groups, virtual drives, physical drives, and hot spares.

The **Foreign Config View** screen lets you import foreign configurations to the RAID controller or clear the foreign configurations.

4.6 Managing Software Licensing

The MegaRAID advanced software offers the software license key feature to enable the advanced options in the Ctrl-R Utility. The license key is also known as the activation key.

You need to configure the Advanced Software options present in the Ctrl-R Utility to use the advanced features present in the controller.

4.6.1 Managing Advanced Software Options

Perform the following steps to configure the Advanced Software options by using the activation key.

1. In the **VD Mgmt** screen, navigate to the controller and press the F2 key.
2. Navigate to **Advanced Software Options**, and press Enter.
   The **Manage MegaRAID Advanced Software Options** dialog appears, as shown in the following figure.

Figure 21 Manage MegaRAID Advanced Software Options

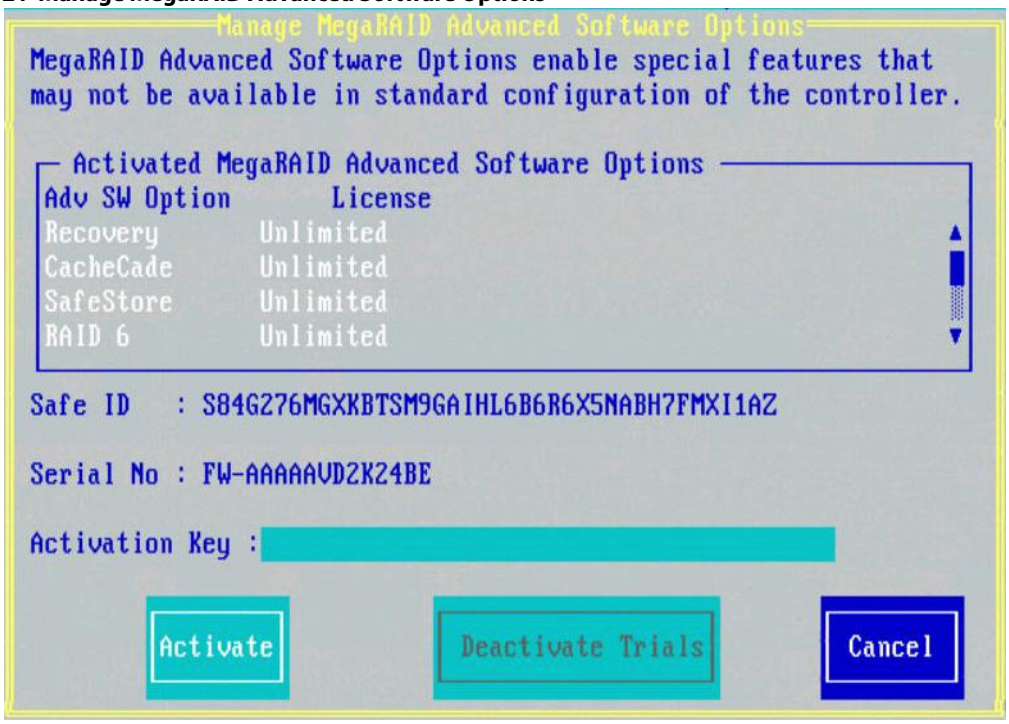

The **Activated MegaRAID Advanced Software Options** box contains the **Adv SW Option** and **License** columns.

- The **Adv SW Option** column shows the list of advanced software features available in the controller.
- The **License** column shows the license details for the list of advanced software options present in the **Adv SW Option** column. The license details validates if the software is under trial period, or whether it can be used without any trial period (Unlimited).

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

3. Enter a valid activation key in the **Activation Key** field.

4. Click **Activate**.
   The **Advanced Software Options Summary** dialog appears, as shown in Figure 26 on page 57.
5. Click **Deactivate Trials**.
   The **Deactivate Trial** dialog appears, as shown in the following figure.

Figure 22 Deactivate Trial Dialog

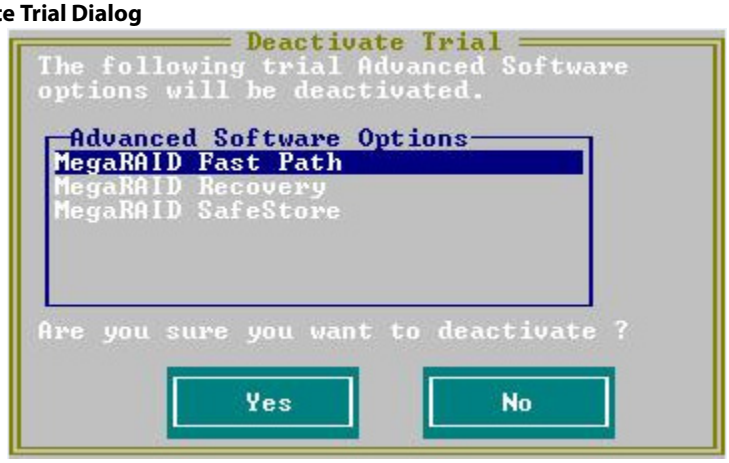

6. Perform one of these actions:
   — If you want to *deactivate* the software that is being used with a trial key, press **Yes**.
   — If you do not want to deactivate the software, press **No**.

If the activation key entered in the **Activation Key** field is incorrect, the following scenario messages appear:

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

Figure 23 Invalid Activation Key Message

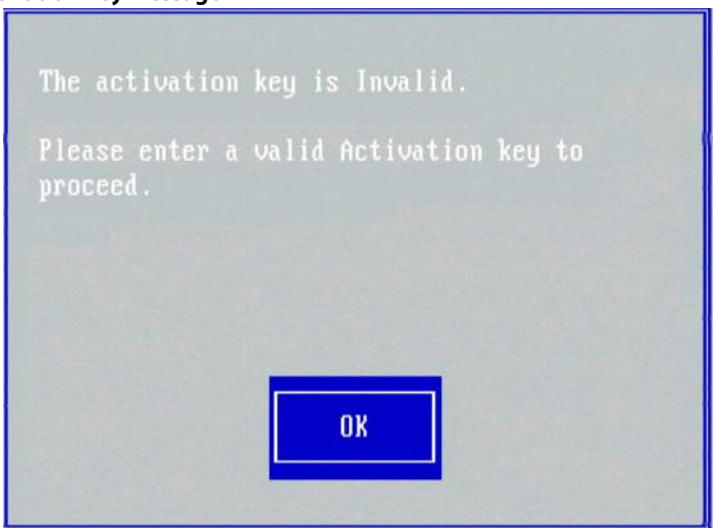

- Scenario 2
  If you leave the **Activation Key** field *blank* or enter *space* characters, the following message appears.

Figure 24 Activation Key Left Blank

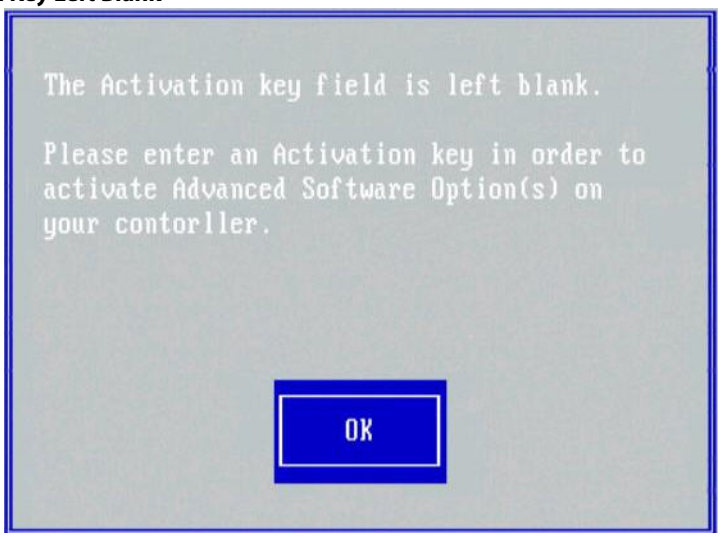

- Scenario 3

  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 25 Activation Key Mismatch Message

4.6.2 Managing Advanced Software Summary

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

Figure 26 Advanced Software Options Summary

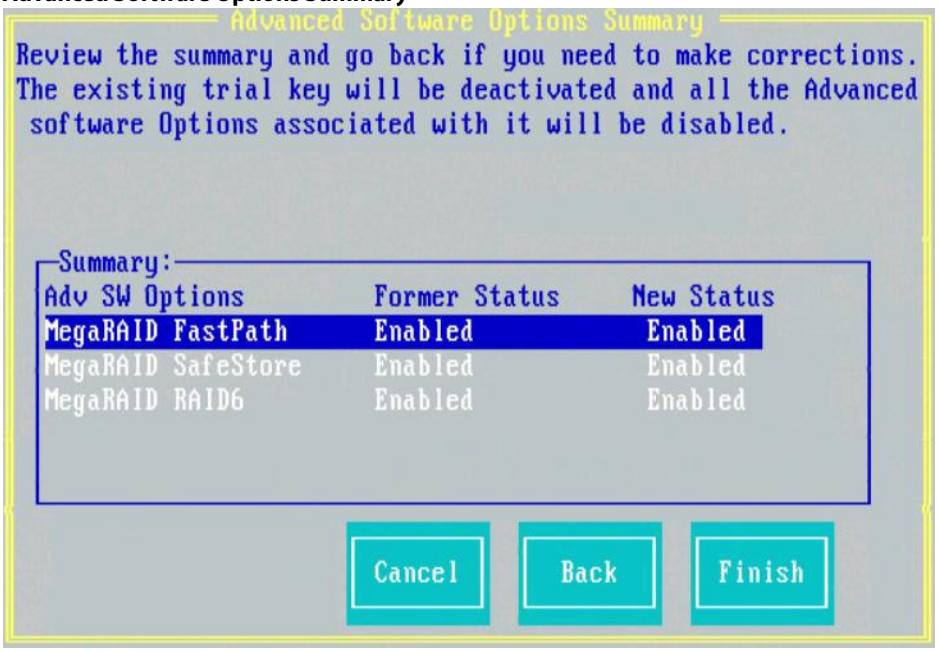

The **Summary** box shows the list of the advanced software options along with their former status and new status.

- The **Advanced SW Options** column shows the currently available software in the controller.
- The **Former Status** column shows the status of the available advanced software before you enter the activation key.
- The **New Status** column shows the status of the available advanced software, after you enter the activation key.

4.6.3 Activating an Unlimited Key over a Trial Key

When you activate an unlimited key over a trial key, the following dialog appears.

Figure 27 Activating an Unlimited Key over a Trial Key

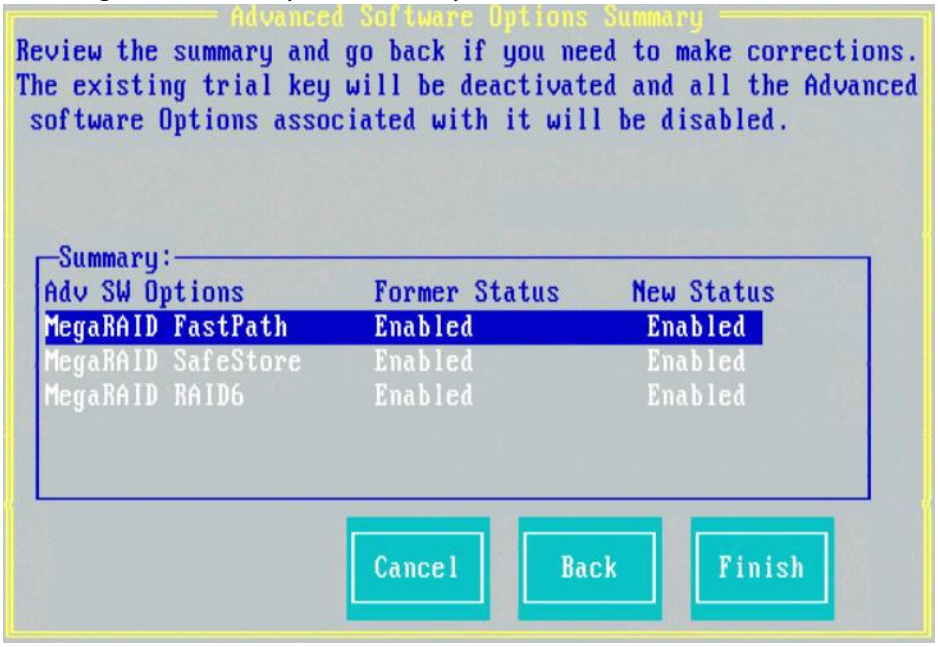

4.6.4 Activating a Trial Software

When you activate a trial software, the following dialog appears.

Figure 28 Activating a Trial Software Application

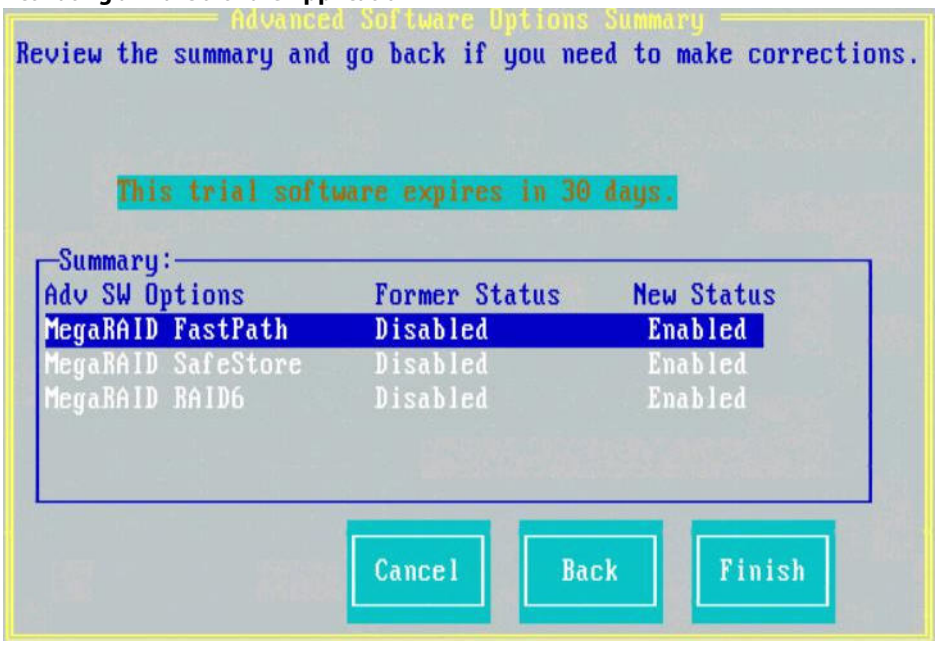

4.6.5 Activating an Unlimited Key

When you activate an unlimited key, the following dialog appears.

Figure 29 Activating an Unlimited Key

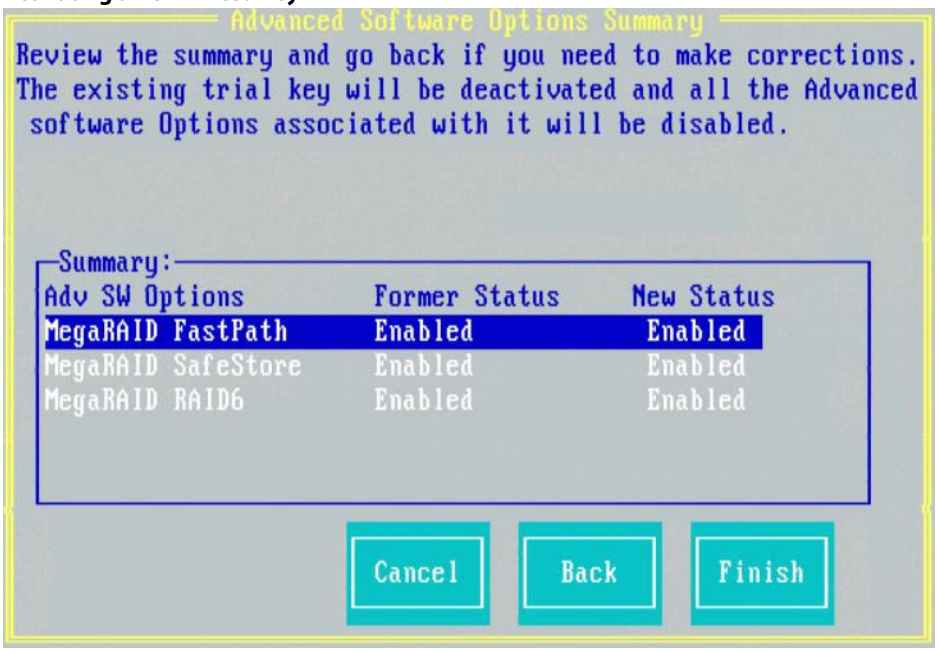

4.7 Creating a Storage Configuration

You can use the Ctrl-R Utility to configure RAID drive groups and virtual drives to create storage configurations on systems with Avago SAS controllers.

NOTE The Ctrl-R utility supports 240 VD creation. For more information, see the 240 Virtual Drive Feature Limitations appendix.

Table 21 RAID Levels

Level	Description
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 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 00	Is a spanned drive group that creates a striped set from a series of RAID 0 drive groups to provide high data throughput, especially for large files.
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.

1. In the **VD Mgmt** screen, navigate to the controller and press the F2 key.
2. Press Enter.

   The **Create New VD** screen appears.

NOTE You can use the **Create New VD** dialog to create virtual drives for Unconfigured Good drives. To create virtual drives for existing drive groups, navigate to a drive group and press the F2 key to view the **Add VD in Drive Group** dialog. The fields in the **Add VD in Drive Group** dialog are the same as in the **Create New VD** dialog.

Figure 30 Create a New Virtual Drive

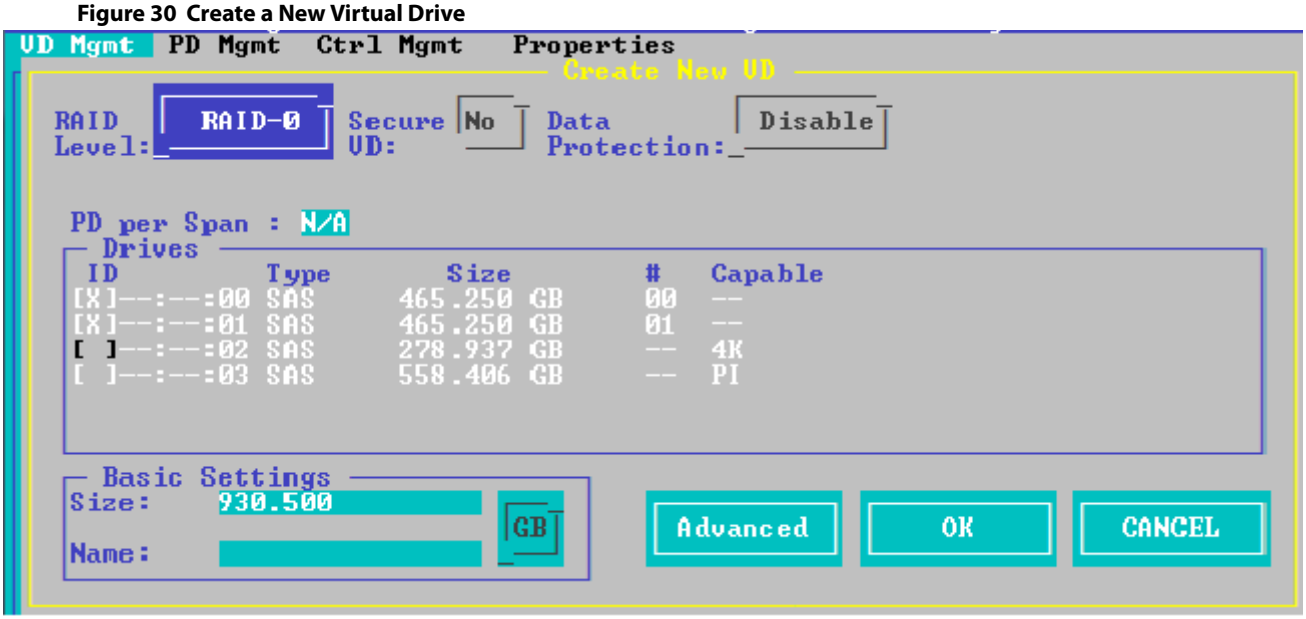

NOTE If your system detects any JBODs, the **Convert JBOD to Unconfigured Good** dialog (Figure 48 on page 78) appears before the **Create New VD** dialog. The **Convert JBOD to Unconfigured Good** dialog lets you convert the JBOD drives to Unconfigured Good.

3. Select a RAID level for the drive group from the **RAID Level** field. For more information, refer to Table 21, RAID Levels
4. Select a power save mode for the drive group from the **Power save mode** field.

   The options available are **Auto**, **Max**, and **Controller defined**.

   This field is enabled only if power saving on configured drives is supported on the controller.

   Power Save (Dimmer Switch feature) is a technology that conserves energy by placing certain unused drives into a Power Save mode. In Power-Save mode, the drives use less energy. The fan and the enclosure require less energy to cool and house the drives, respectively. Also, this technology helps avoid application time-outs caused by spin-up delays and drive wear caused by excessive spin-up/down cycles.
5. You can encrypt data and use drive-based key management for your data security solution.

   This option protects the data in the event of theft or loss of drives. Select a value from the **Secure VD** field. The options available are **Yes** and **No**.
6. You can choose whether you want to use the data protection feature on the newly created virtual drive.

   Select a value from the **Data Protection** field. The options available are **Yes** and **No**. The **Data Protection** field is enabled only if the controller has data protection physical drives connected to it.
7. You can change the sequence of the physical drives in the **Drives** box.

   All the available unconfigured good drives appear in the **Drives** box. Select the physical drives in the sequence that you prefer. Based on your selection, the sequence number appears in the **#** column. The **Type** column shows the drive type; for example, SAS, SATA, IDE, and so on. The **Capable** column shows the capability of the drive.
8. You can select a size lesser than the maximum size of the drive group, if you want to create other virtual drives on the same drive group.

   The maximum size of the drive group appears in the **Size** field. Select either MB, GB, or TB from the drop-down menu.

NOTE Drive group size of floating data type up to three decimal places is supported. Some of the screens in this chapter may not show this feature.

9. Enter a name for the virtual drive in the **Name** field.

   The name given to the virtual drive cannot exceed 15 characters.

   You may press **Advanced** to set additional properties for the newly created virtual drive. For more information, see Selecting Additional Virtual Drive Properties.

10. Press **OK**.

    A dialog appears, asking you whether you want to initialize the virtual drive you just created.

11. To initialize the virtual drive, press **OK**.

    The **Create New VD** dialog appears again.

12. Press **Advanced**.

    The **Create Virtual Drive – Advanced** dialog appears.

Figure 31 Create Virtual Drive – Advanced

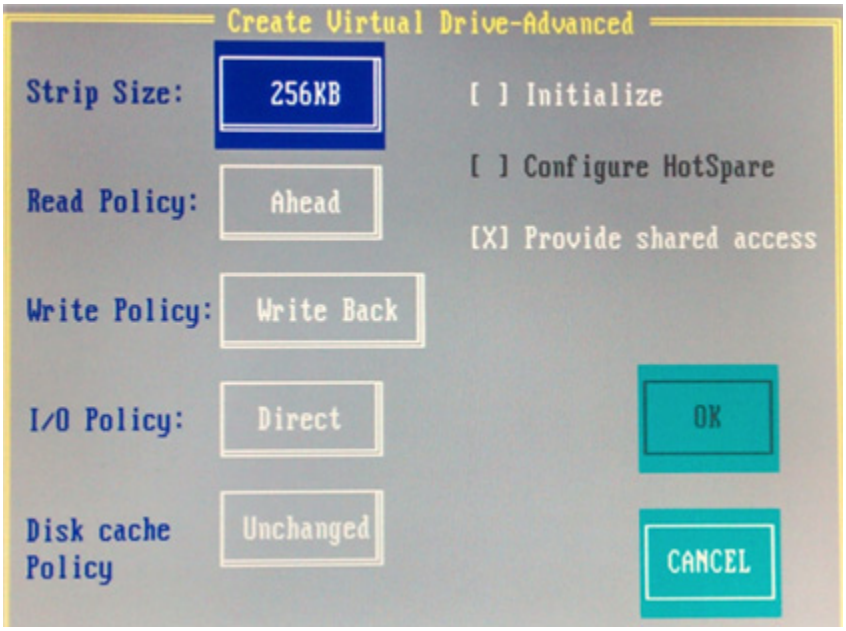

NOTE The **Provide shared access** check box appears only if the controller supports High Availability DAS.

13. Select **Initialize**, and press **OK**.

    The new virtual drive is created and initialized.

4.7.1 Selecting Additional Virtual Drive Properties

This section describes the following additional virtual drive properties that you can select while you create virtual drives. Change these parameters only if you have a specific reason for doing so. It is usually best to keep them at their default settings.

- **Strip Size** – The strip size is the portion of the stripe that resides on a single virtual drive in the drive group. Strip sizes of 64 KB, 128 KB, 256 KB, 512 KB, or 1 MB are supported.

  NOTE The Integrated MegaRAID controller supports only64 KB stripe size.

- **Read Policy** – Specify one of the following options to specify the read policy for this virtual drive:
  - **Normal** – Read ahead capability lets the controller read sequentially ahead of requested data and to store the additional data in cache memory, thereby anticipating that the data will be needed soon. This process speeds up reads for sequential data, but there is little improvement when the computer accesses random data.
  - **Ahead** – Disables the read ahead capability.
- **Write Policy** – Select one of the following options to specify the write policy for this virtual drive:
  - **Write Back** – In this mode, the controller sends a data transfer completion signal to the host when the controller cache receives all of the data in a transaction. If you select the Write Back policy and the battery is absent, the firmware disables the Write Back policy and defaults to the Write Through policy.
  - **Write Through** – In this mode, the controller sends a data transfer completion signal to the host when the drive subsystem receives all the data in a transaction.
  - **Always Write Back** – In this mode, the controller sends a data transfer completion signal to the host when the controller cache receives all the data in a transaction. If you select the Always Write Back policy and the battery is absent, the firmware is forced to use the Write Back policy.
- **I/O Policy** – The I/O policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.
  - **Cached** – In this mode, all reads are buffered in cache memory. **Cached I/O** provides faster processing.
  - **Direct** – In this mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. **Direct I/O** makes sure that the cache and the host contain the same data.
- **Disk cache policy** – Select a cache setting for this virtual drive:
  - **Enable** – Enable the drive cache.
  - **Disable** – Disable the drive cache.
  - **Unchanged** – Updating the drive cache policy to **Unchanged** may enable /disable the drive cache based on the WCE (Write Cache Policy) bit of the save mode page of the drive.
- **Emulation** – Lets you to set the emulation type on a virtual drive to default or none. The force option forces the emulation to be set on a controller even when MFC settings do not support it. The possible options are **Default**, **Disabled**, or **Forced**.
- **Initialize** – Select to initialize the virtual drive. Initialization prepares the storage medium for use. Fast initialization will be performed on the virtual drive.
- **Configure Hot Spare** – Select to configure physical drives as hot spares for the newly created virtual drive.

  This option is enabled only if there are additional drives and if they are eligible to be configured as hot spares. This option is not applicable for RAID 0 or RAID 00. If you select this option and after the Virtual drive is created, a dialog appears. The dialog asks you to choose the physical drives that you want to configure as hot spares.
- **Provide shared access**– Select this option if you want the virtual drive to be shared between the servers in a cluster. This option appears only if the controller supports High Availability DAS.

4.7.2 Creating a CacheCade Virtual Drive

The MegaRAID CacheCade software provides you with read caching capability.

Perform the following steps to create a CacheCade virtual drive:

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Navigate to **Create CacheCade Virtual Drive**, and press Enter.

   The **Create CacheCade Virtual Drive** dialog appears.

Figure 32 Create CacheCade Virtual Drive

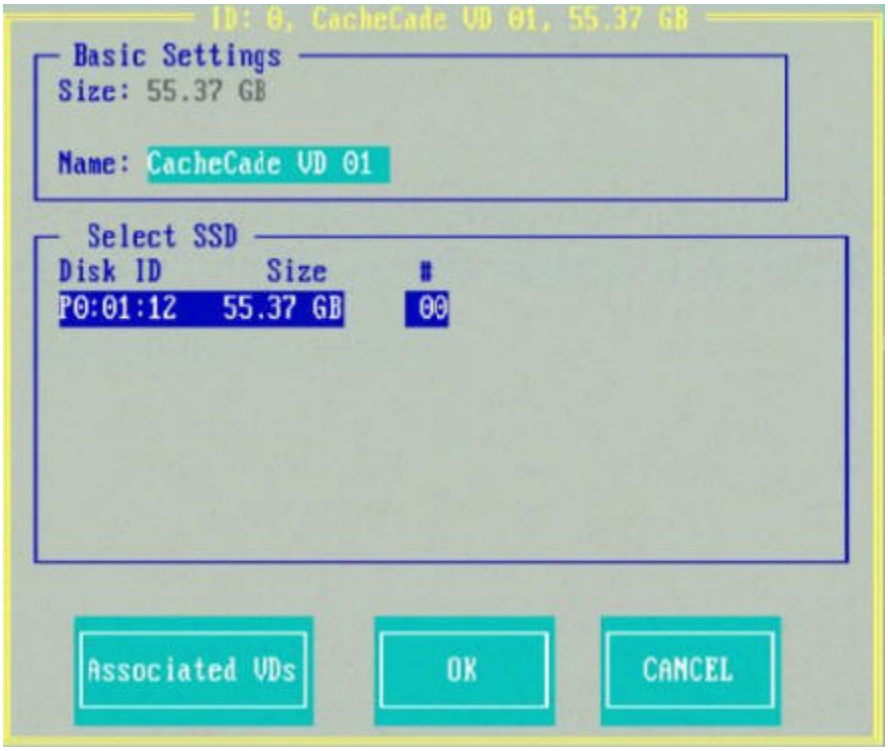

3. Enter a name for the CacheCade virtual drive in the **Name** field.
4. Select a SSD from the **Select SSD** box.

   The size of the SSD is reflected in the **Size** field (in the **Basic Settings** box).
5. Click **OK**.

   A message appears, stating that the CacheCade virtual drive has been created.

   To view the virtual drives associated with this CacheCade virtual drive, click **Associated VDs** in the **Create CacheCade Virtual Drive** dialog.

   The **Associated Virtual Drives** dialog appears.

Figure 33 Associated Virtual Drives

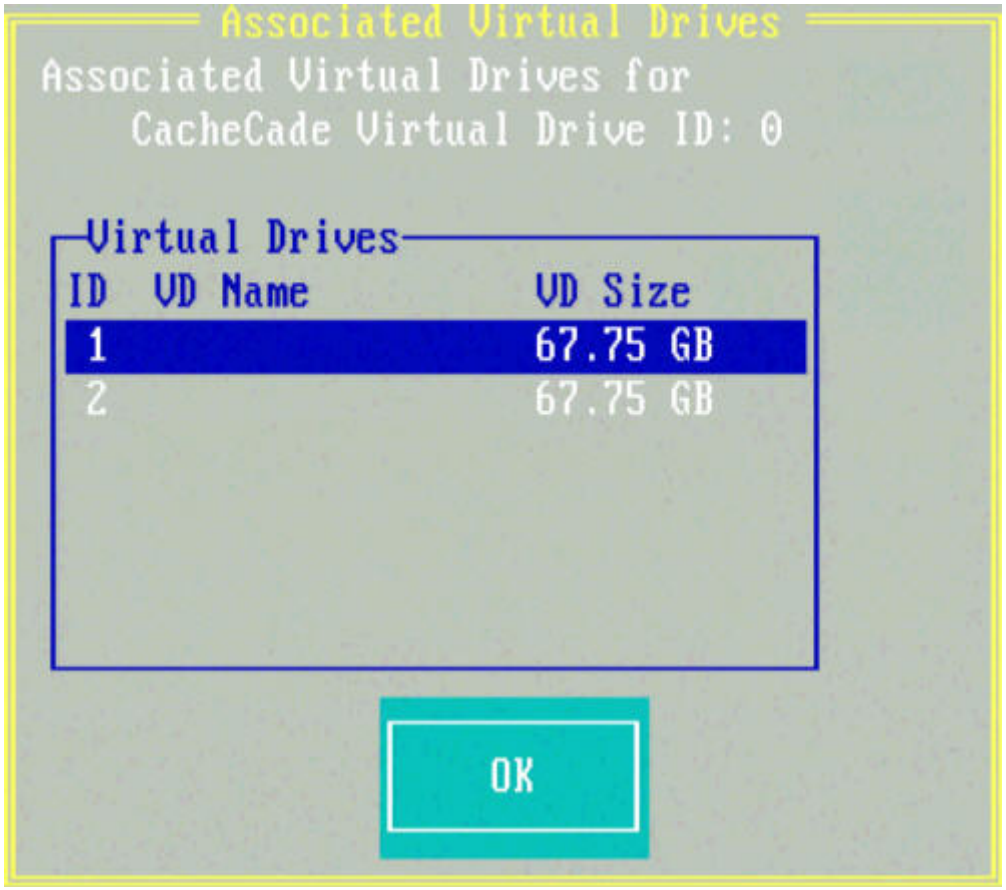

You can view the ID, the name, and the size of the associated virtual drives.

4.7.3 Modifying a CacheCade Virtual Drive

You can modify an existing CacheCade virtual drive by renaming it.

Perform the following steps to modify the CacheCade virtual drive:

1. In the **VD Mgmt** screen, navigate to the CacheCade virtual drive. and press the F2 key.
2. Navigate to **Properties**, and press Enter.
   The following dialog appears.

Figure 34 Modifying CacheCade Virtual Drive

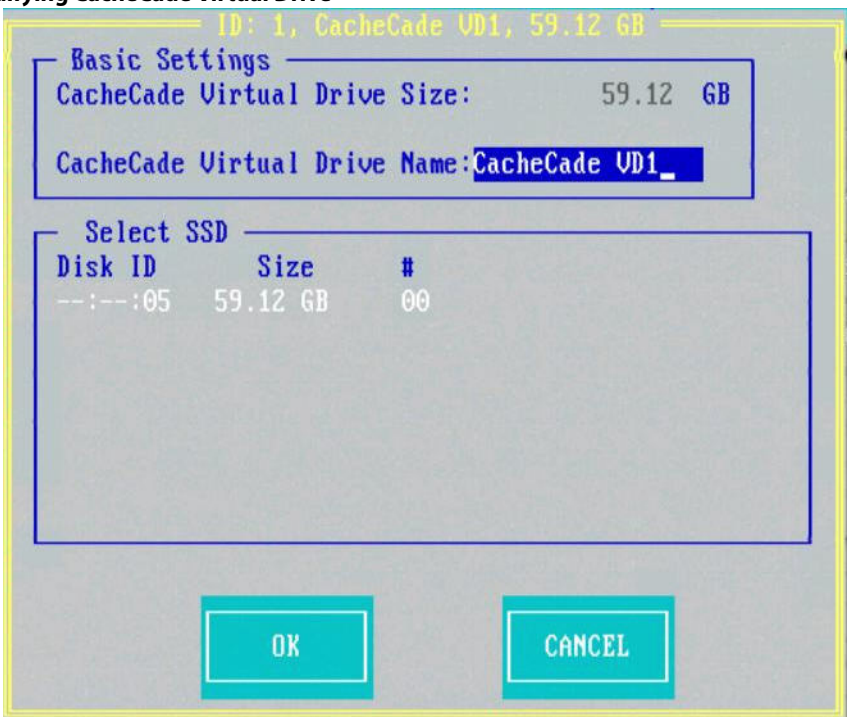

3. You can rename a CacheCade virtual drive in the **CacheCade Virtual Drive Name** field.
4. Press **OK**.

4.7.4 Creating a CacheCade Pro 2.0 Virtual Drive

The MegaRAID CacheCade Pro 2.0 provides you with read and write capability.

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

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Navigate to **Create CacheCade Virtual Drive**, and press Enter.
   The **Create CacheCade Virtual Drive** dialog appears.

Figure 35 Create CacheCade Virtual Drive

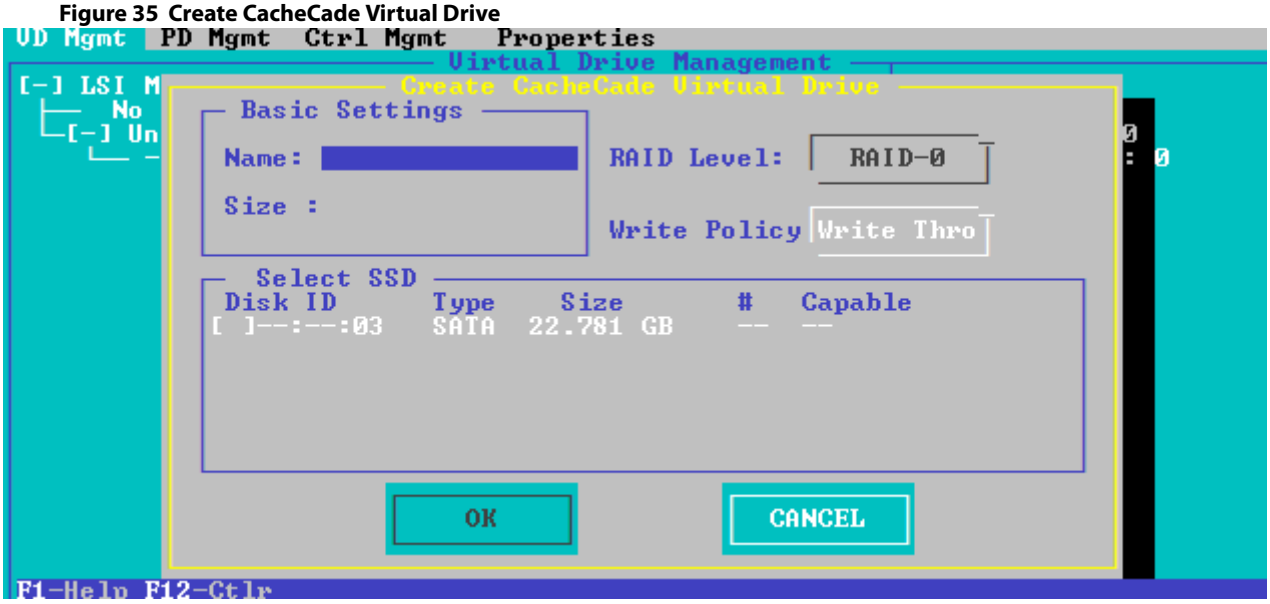

3. Enter a name for the CacheCade virtual drive in the **Name** field.
4. Select a SSD from the **Select SSD** box.
5. Press **OK**.

   A message appears, stating that the CacheCade virtual drive has been created.

4.7.5 Modifying a CacheCade Pro 2.0 Virtual Drive

You can modify the name and the write policy of an existing CacheCade virtual drive any time after a CacheCade virtual drive is created.

Perform the following steps to modify the CacheCade virtual drive:

1. In the **VD Mgmt** screen, navigate to the CacheCade virtual drive. and press the F2 key.
2. Navigate to **Properties**, and press Enter.

   The following dialog appears.

Figure 36 Modifying CacheCade Virtual Drive

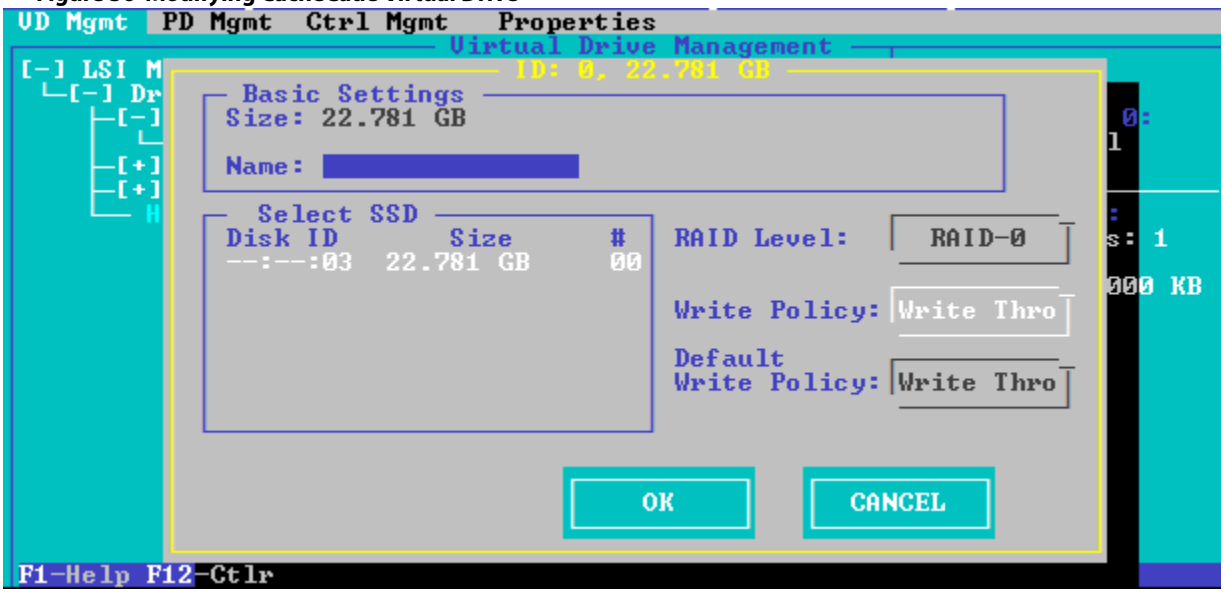

3. You can rename a CacheCade virtual drive in the **CacheCade Virtual Drive Name** field.
4. You can also modify the write policy by selecting one from the **Write Policy** field.
5. Press **OK**.

4.7.6 Enabling SSD Caching on a Virtual Drive

You can enable SSD caching on a virtual drive. When you enable SSD caching on a virtual drive, that virtual drive becomes associated with an existing or with a future CacheCade SSD Caching virtual drive. This option is only available when the virtual drive's caching is currently disabled.

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

1. In the **VD Mgmt** screen, navigate to a virtual drive. and press the F2 key.
2. Select **Enable Caching** and press Enter.

   The following message dialog appears.

Figure 37 Message Box for Enabling SSD Caching

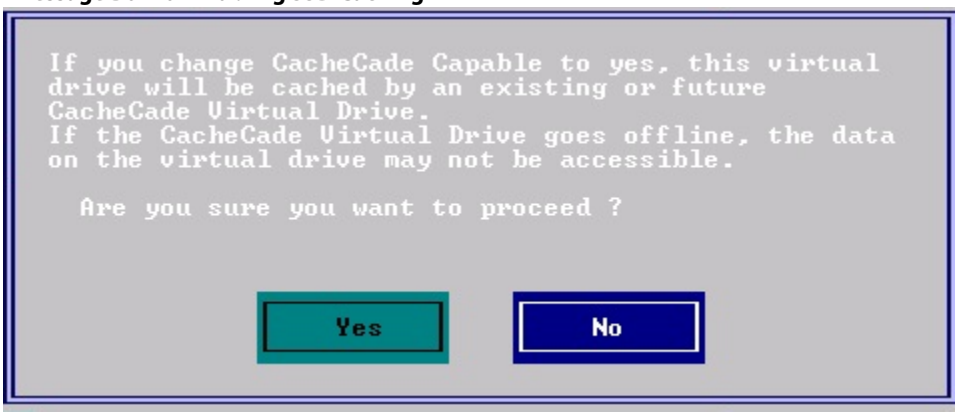

   The virtual drives that have SSD caching enabled, have the check boxes next to them selected. The virtual drives that have SSD caching disabled, have deselected check boxes.
3. Click **Yes** to enable caching for that virtual drive.

4.7.7 Disabling SSD Caching on a Virtual Drive

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

Perform the following steps to enable SSD Caching on a virtual drive:

1. In the **VD Mgmt** screen, navigate to a virtual drive. and press the F2 key.
2. Select **Disable Caching** and press Enter.
   The following message dialog appears.

Figure 38 Message Box for Disabling SSD Caching

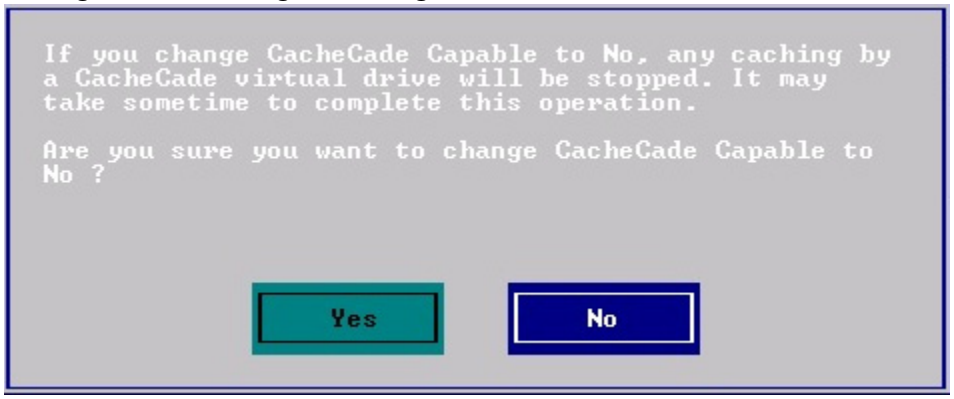

3. Click **Yes** to disable caching for that virtual drive.

4.7.8 Enabling or Disabling SSD Caching on Multiple Virtual Drives

You can enable or disable SSD caching on multiple virtual drives at one go.

Perform the following steps to enable or disable SSD caching on multiple drives:

1. In the **VD Mgmt** screen, navigate to a virtual drive and press the F2 key.
2. Select **Manage SSD Caching** and press Enter.
   The **Manage SSD Caching** dialog appears.

Figure 39 Manage SSD Caching

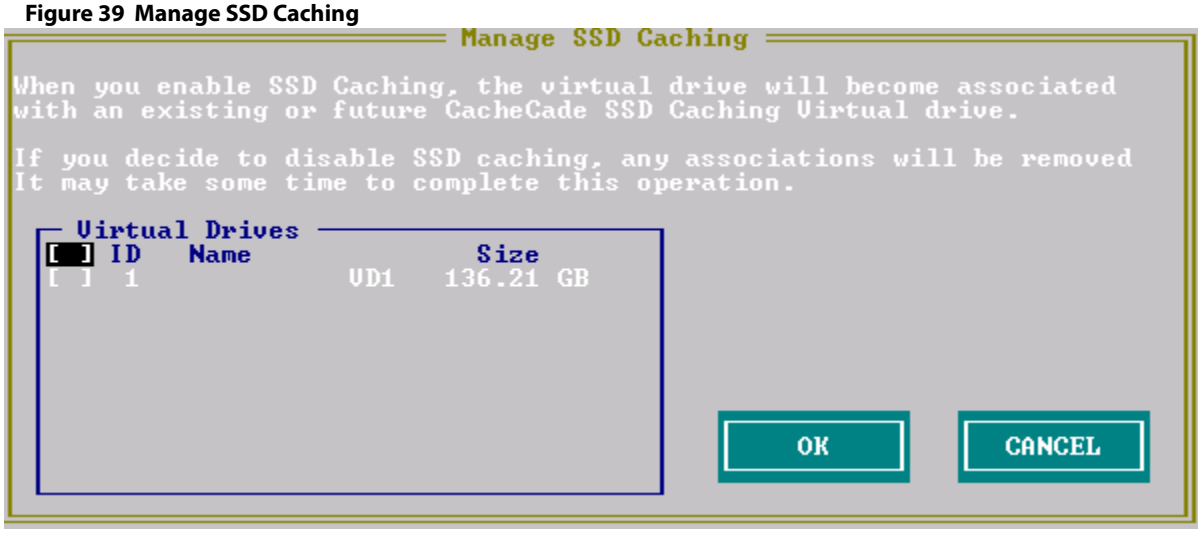

The virtual drives that have SSD caching enabled have the check boxes next to them selected. The virtual drives that have SSD caching disabled have deselected check boxes.

3. Select or deselect a check box to change the current setting of a virtual drive.
4. Click **OK** to enable or disable SSD caching on the selected virtual drives.

4.7.9 Deleting a Virtual Drive with SSD Caching Enabled

You can delete a virtual drive that has SSD caching enabled on it.

Perform the following steps to delete the virtual drive:

1. In the **VD Mgmt** screen, navigate to a virtual drive. and press the F2 key.
2. Select **Delete VD**, and click **Yes**.
   The following message dialog appears.

Figure 40 Message Box for Deleting Virtual Drive

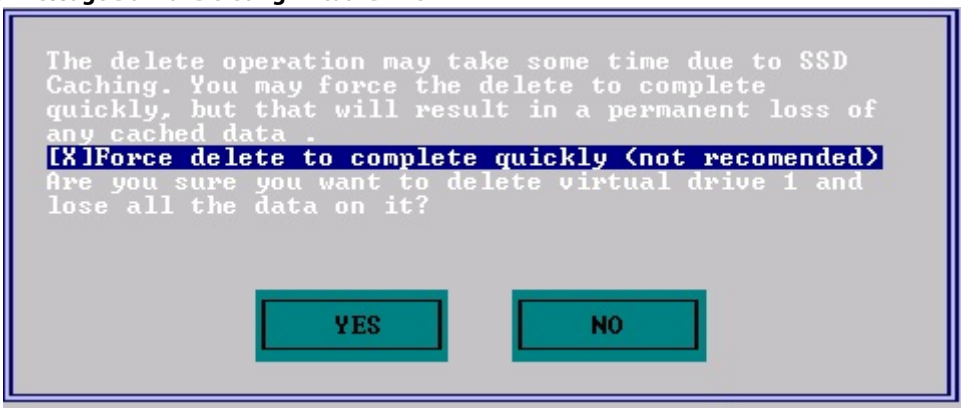

NOTE If you select the **Force delete to complete quickly** check box to delete the virtual drive, the data is not flushed before deleting the virtual drive. In this scenario, if you create this virtual drive after deleting it, there will be no data available.

3. Press **Yes** to proceed with the delete operation.

4.8 Clearing the Configuration

You can clear all the existing configuration on virtual drives by deleting the virtual drives.

Perform the following steps to clear configuration:

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Navigate to **Clear Configuration** and press Enter.
   The following dialog appears.

Figure 41 Clear Configuration

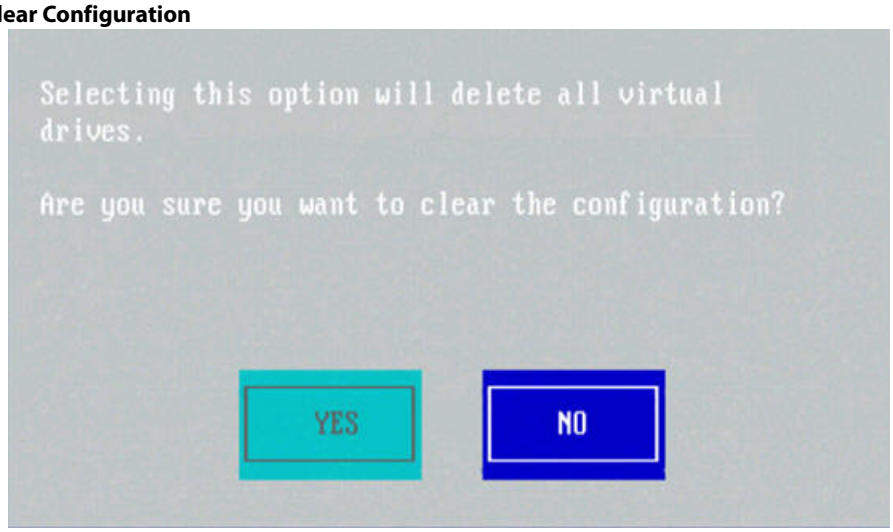

3. Press **Yes** to delete all the virtual drives.

4.9 Avago SafeStore Encryption Services

The Avago SafeStore Encryption Services can encrypt data on the drives and use the drive-based key management to provide data security. This solution protects data in the event of theft or loss of physical drives. If you remove a self-encrypting drive from its storage system or the server in which it resides, the data on that drive is encrypted, and becomes useless to anyone who attempts to access it without the appropriate security authorization.

4.9.1 Enabling Drive Security

This section describes how to enable, change, and disable the drive security, and how to import a foreign configuration by 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.
- Enter a security key.
  After you create a security key, you can create secure virtual drives by using the key. You must use the security key to perform certain operations.

You can improve security by entering a password. To provide additional security, you can request for the password whenever anyone boots the server.

Perform the following steps to enable drive security.

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Navigate to **Drive Security**, and press Enter.
3. Navigate to **Enable Security**, and press Enter.
   The **Create Security Key** dialog appears.

Figure 42 Create Security Key

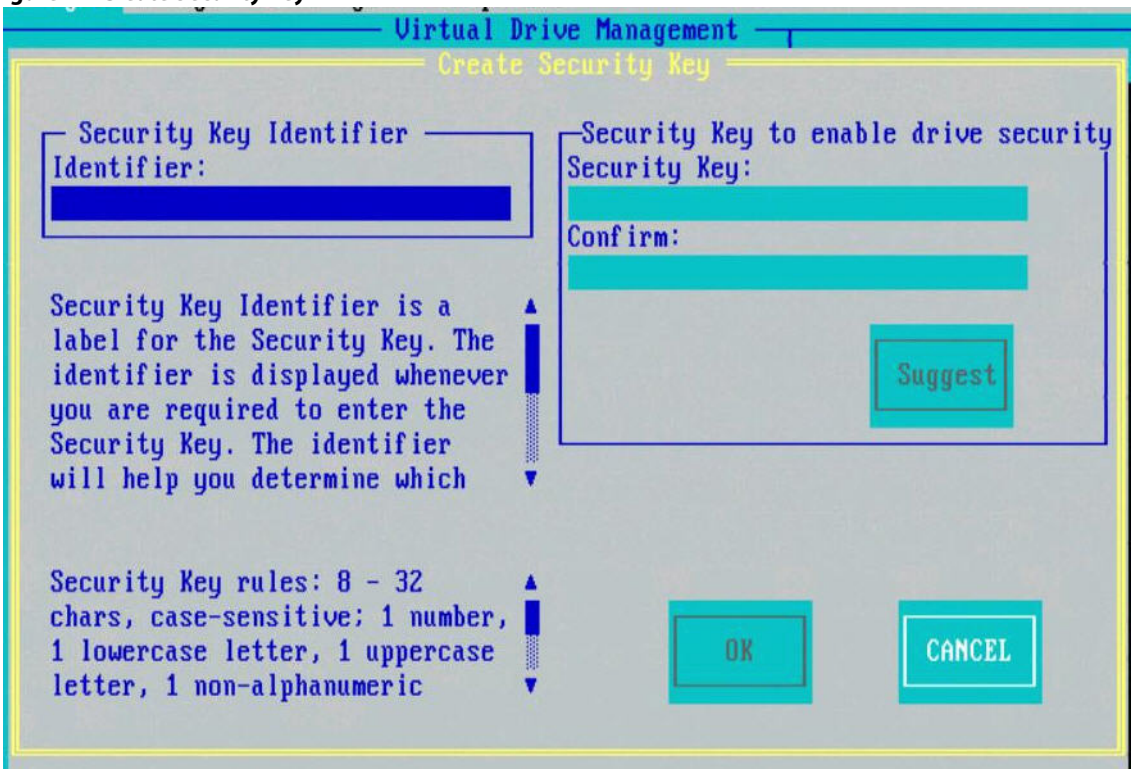

4. Either use the default security key identifier, or enter a new security key identifier.

   NOTE After you create a security key, the **Enable Security** option is disabled. This option is re-enabled only after you delete the existing key.

5. Either click **Suggest** to ask the system to create a security key, or you can enter a new security key.
6. Reenter the new security key to confirm it.

   ATTENTION **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 (a symbol, for example, < > @ +). The space character is not permitted.

   NOTE Non-U.S. keyboard users must be careful not to enter double-byte character set (DBCS) characters in the security key field. The firmware works with the ASCII character set only.

4.9.2 Changing Security Settings

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

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Navigate to **Drive Security**, and press Enter.
3. Select **Change Security Settings**, and press Enter.

   The **Change Security Key** dialog appears.

Figure 43 Change Security Key

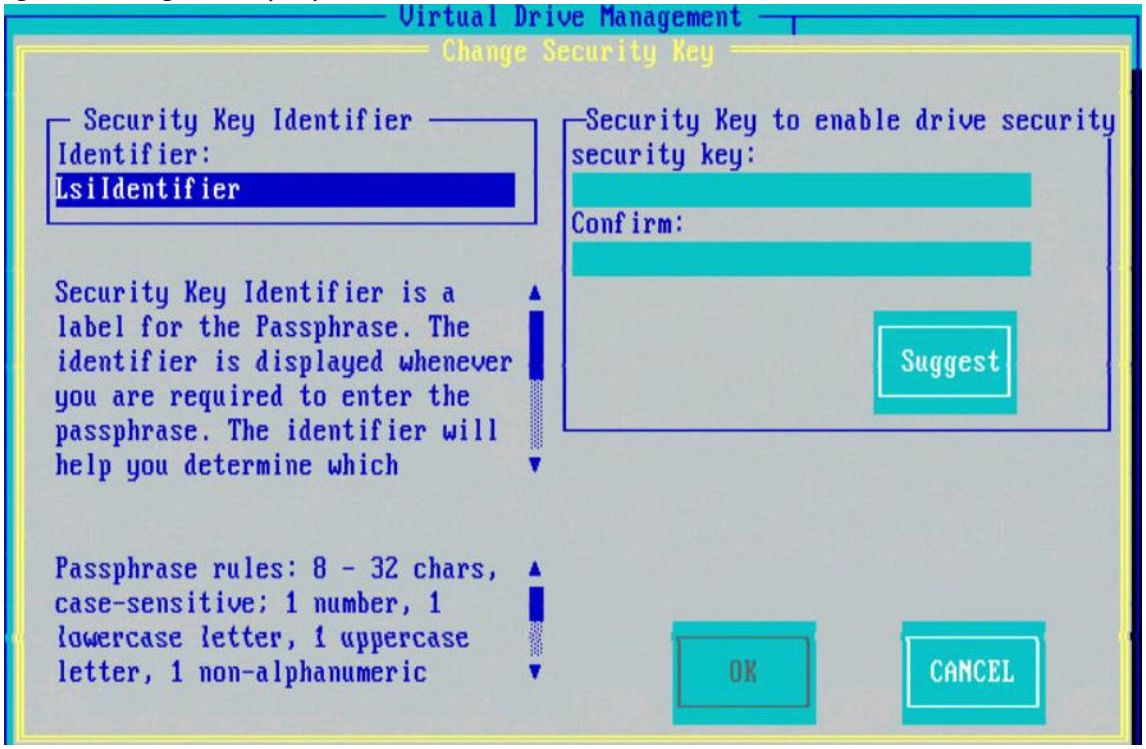

4. Either keep the existing security key identifier, or enter a new security key identifier.

   NOTE If you change the security key, you need to change the security key identifier. Otherwise, you cannot differentiate between the security keys.

5. Either click **Suggest** to ask the system to create a security key, or you can enter a new security key.
6. Re-enter the new security key to confirm it.

   ATTENTION **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 (for example, < > @ +). The space character is not permitted.

   NOTE Non-U.S. keyboard users must be careful not to enter DBCS characters in the Security Key field. The firmware works with the ASCII character set only.

4.9.3 Disabling Drive Security

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 data security 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.

If there are any secure drive groups on the controller, you cannot disable drive security. A warning dialog appears if you attempt to do so. To disable drive security, you first must delete the virtual drives on all the secure drive groups.

Perform the following steps to disable drive security:

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Navigate to **Drive Security**, and press Enter.
3. Select **Disable Security**.

   A message box appears.
4. To disable drive security, click **Yes** to delete the security key.

ATTENTION If you disable drive security, you cannot create any new encrypted virtual drives and the data on all encrypted unconfigured drives will be erased. Disabling drive security does not affect the security or data of foreign drives.

4.9.4 Importing or Clearing a Foreign Configuration

A foreign configuration is a RAID configuration that already exists on a replacement set of drives that you install in a computer system. You can use the Ctrl-R Utility to import the foreign configuration to the RAID controller or to clear the foreign configuration so that you can create a new configuration by using these drives.

To import a foreign configuration, you must perform the following tasks:

- Enable security to permit importation of locked foreign configurations. You can import unsecured or unlocked configurations when security is disabled.
- If a locked foreign configuration is present and security is enabled, enter the security key, and unlock the configuration.
- Import the foreign configuration.

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

Verify whether any drives are left to import because the locked drives can use different security keys. If any drives remain, repeat the import process for the remaining drives. After all the drives are imported, there is no configuration to import.

NOTE When you create a new configuration, the Ctrl-R Utility shows only the unconfigured drives. Drives that have existing configurations, including foreign configurations, do not appear. To use drives with existing configurations, you first must clear the configuration on those drives.

You can import or clear a foreign configuration from the **VD Mgmt** menu or from the **Foreign View** menu.

Perform the following steps to import or clear a foreign configuration from the **VD Mgmt** menu:

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.

2. Navigate to **Foreign Config**, and press Enter.
   The foreign configuration options **Import** and **Clear** appear.

Figure 44 Foreign Configuration Options

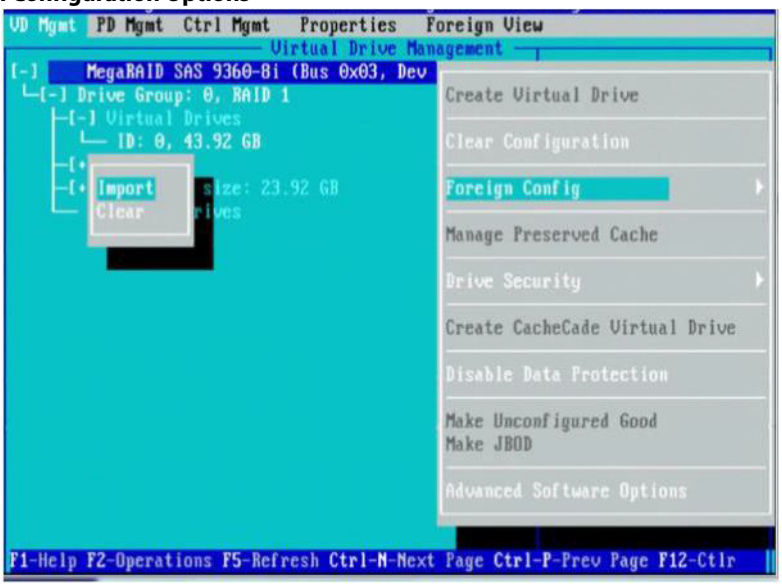

3. Navigate to the command you want to perform.
   — To import a foreign configuration, go to step 4.
   — To clear a foreign configuration, go to step 6.

4. To import a foreign configuration, select **Import**, and press Enter.
   The following dialog appears.

Figure 45 Foreign Configuration – Import

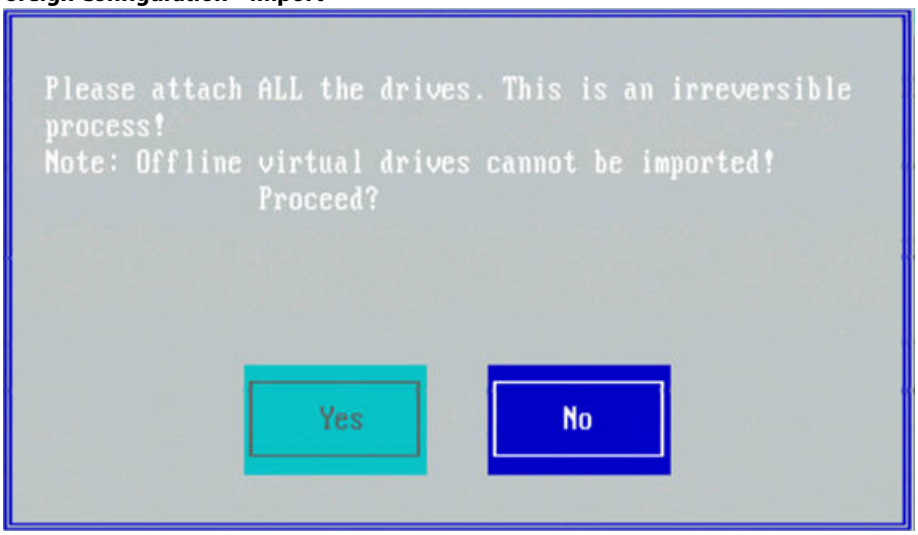

5. Press **Yes** to import the foreign configuration from all the foreign drives. Repeat the import process for any remaining drives.
   Because locked drives can use different security keys, you must verify whether there are any remaining drives to be imported.

NOTE When you create a new configuration, the Ctrl-R Utility shows only the unconfigured drives. Drives that have existing configurations, including foreign configurations, do not appear. To use drives with existing configurations, you first must clear the configuration on those drives.

6. To clear a foreign configuration, select **Clear**, and press Enter.

   The following dialog appears.

Figure 46 Foreign Configuration – Clear

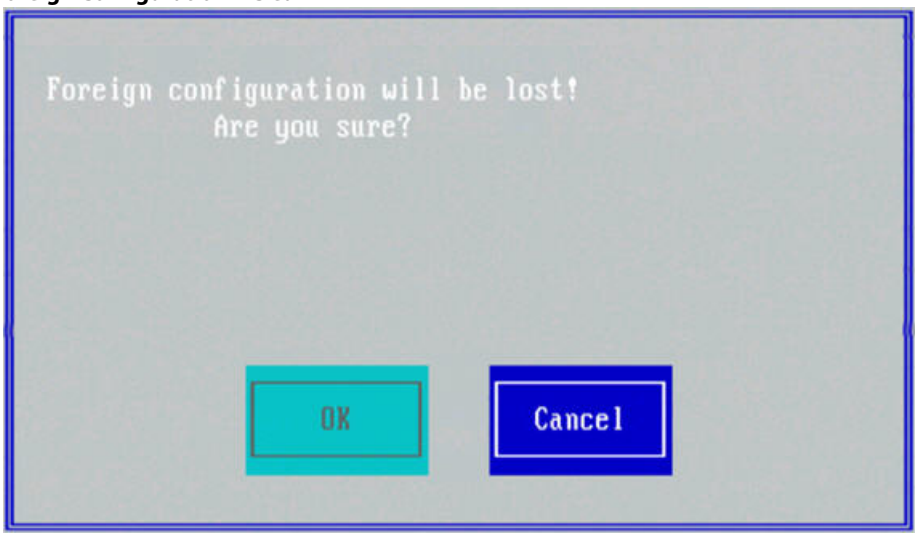

7. Press **OK** to clear a foreign configuration.

NOTE The operation cannot be reversed after it is started. Imported drives appear as Online in the Ctrl-R Utility.

4.9.4.1 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.

NOTE To import the foreign configuration in any of the following scenarios, you must have all the drives in the enclosure before you perform the import operation.

- Scenario 1: If all the drives in a configuration are removed and reinserted, 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.

- Scenario 2: If some of the drives in a configuration are removed and reinserted, 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.

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

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

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

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

4.10 Discarding Preserved Cache

If the controller loses access to one or more virtual drives, the controller preserves the data from the virtual drive. This preserved cache, is preserved until you import the virtual drive or discard the cache.

Certain operations, such as creating a new virtual drive, cannot be performed if preserved cache exists.

CAUTION If there are any foreign configurations, import the foreign configuration before you discard the preserved cache. Otherwise, you might lose data that belongs to the foreign configuration.

Perform the following steps to discard the preserved cache:

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Navigate to **Manage Preserved Cache**, and press Enter.

   The **Manage Preserved Cache** dialog appears.

Figure 47 Manage Preserved Cache

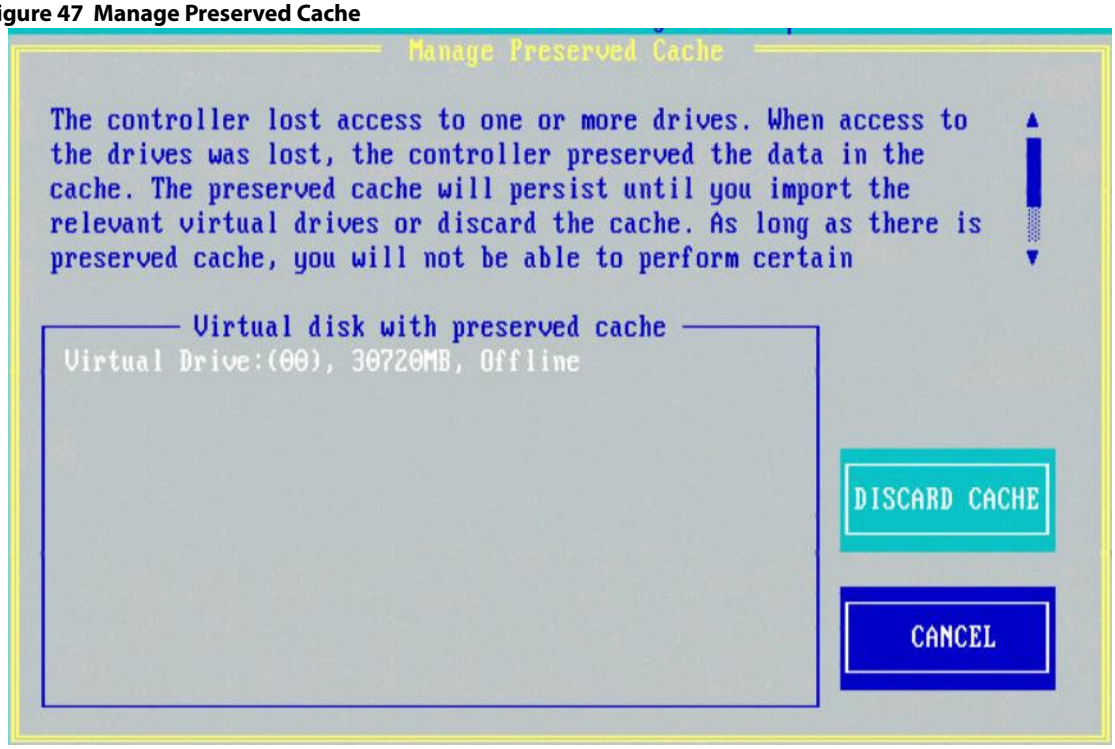

3. Click **Discard Cache** to discard the preserved cache from the virtual drive. A message box appears, asking you to confirm your choice.
4. Click **OK** to continue.

4.11 Converting JBOD Drives to Unconfigured Good Drives

You can convert multiple JBODs to Unconfigured Good drives (from the **VD Mgmt** screen), or you can convert a particular JBOD drive to an Unconfigured Good drive (from the **Drive Management** screen).

NOTE The MegaRAID SAS 9240-4i and the MegaRAID SAS 9240-8i controllers support JBOD.

Perform the following steps to convert multiple JBODs to Unconfigured Good drives:

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Navigate to **Make Unconfigured Good**, and press Enter.

   The **Convert JBOD to Unconfigured Good** dialog appears, which shows all JBODs available in the system.

Figure 48 Convert JBOD to Unconfigured Good

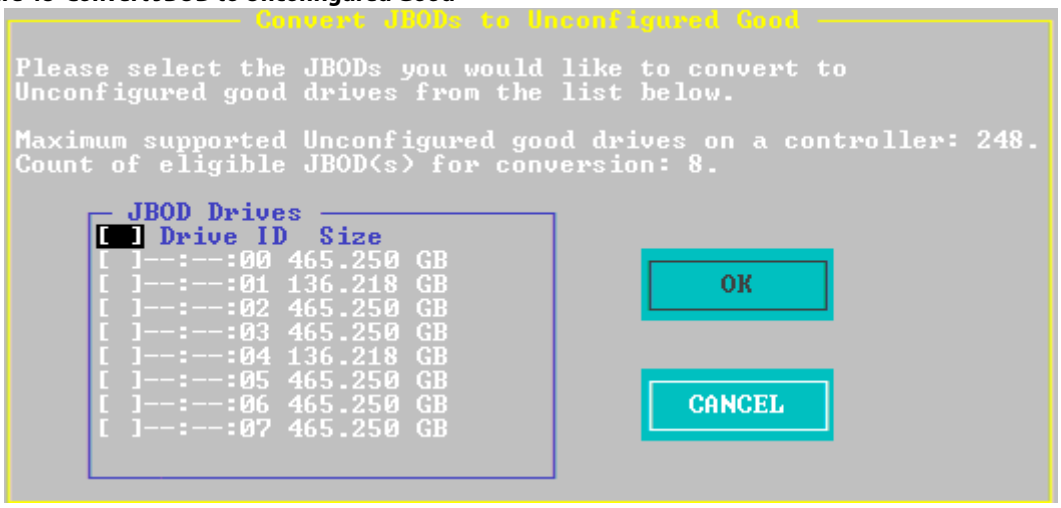

3. Select the JBODs that you want configured as Unconfigured Good drives.

   To select or deselect all the JBODs at one go, select the top most square brackets in the **JBOD Drives** box.

   NOTE If the selected JBODs have an operating system or a file system, a warning message appears indicating that the listed JBODs contain an operating system or a file system, and any existing data on the drives would be lost if you proceed with the conversion. If you want to proceed with the conversion, click **Yes**. Else, click **No** to return to the previous screen and unselect those JBODs that have the OS or the file system installed on them.

4. Click **OK**.

   The selected JBODS are converted to Unconfigured Good drives.

Perform the following steps to convert a particular JBOD drive to an Unconfigured Good drive:

1. In the **Drive Management** screen, navigate to a JBOD drive, and press the F2 key.
2. Navigate to **Make Unconfigured Good**, and press Enter.

   NOTE If the JBOD has an operating system or a file system, a warning message appears indicating that the JBOD contains an operating system or a file system, and any existing data on the drive would be lost if you proceed with the conversion. If you want to proceed with the conversion, click **Yes**. Else, click **No** to return to the previous screen.

4.12 Converting Unconfigured Good Drives to JBOD Drives

You can convert a bunch of Unconfigured Good drives to JBOD drives (from the **VD Mgmt** screen), or you can convert a particular Unconfigured Good drive to a JBOD drive (from the **Drive Management** screen).

Perform the following steps to convert a bunch of Unconfigured Good drives to JBOD drives:

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Navigate to **Make JBOD**, and press Enter.

   The **Convert Unconfigured Good to JBOD** dialog appears, which shows all Unconfigured Good drives available in the system.

Figure 49 Convert Unconfigured Good to JBOD

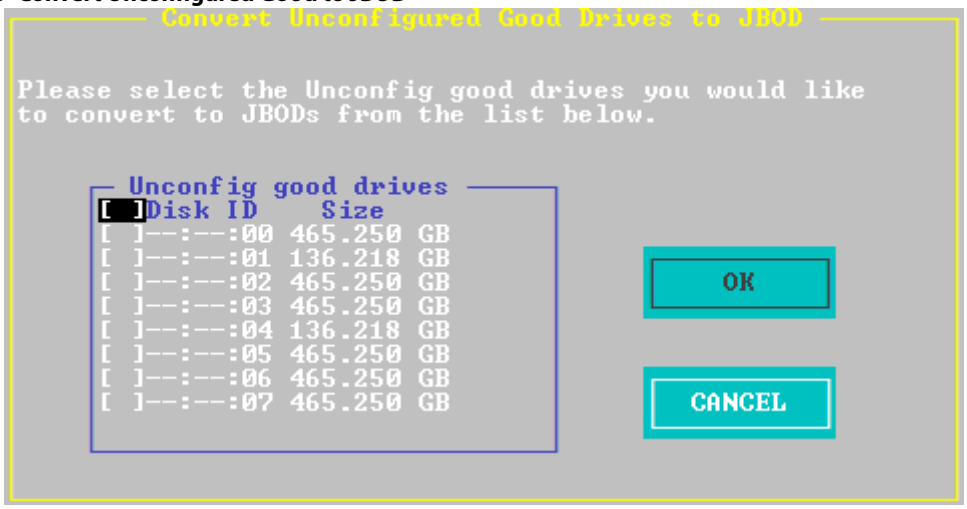

3. Select the Unconfigured Good drives that you want configured as JBODs.

   To select or deselect all the Unconfigured Good drives at one go, select the top most square brackets in the **Unconfig good drives** box.
4. Click **OK**.

   The selected Unconfigured Good drives are converted to JBOD drives.

Perform the following steps to convert a particular Unconfigured Good drive to a JBOD drive:

1. In the **Drive Management** screen, navigate to a Unconfigured Good drive, and press the F2 key.
2. Navigate to **Make JBOD**, and press Enter.
3. Click **OK** in the message confirmation box to continue.

4.13 Enabling Security on a JBOD

You can enable security on the JBOD drives (from the **VD Mgmt** screen or the **Drive Management** screen). The following are the prerequisites for enabling security on the JBOD drives:

- The drive must be an SED capable drive.
- The controller must support Security feature.
- The controller must support JBOD functionality.

Perform the following steps to convert a bunch of Unconfigured Good drives to JBOD drives:

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Navigate to **Enable Security on JBOD**, and press Enter.

   The **Enable Security on JBOD** dialog appears, which shows all of the SED-enabled JBOD drives available in the system.

Figure 50 Enable Security on JBOD

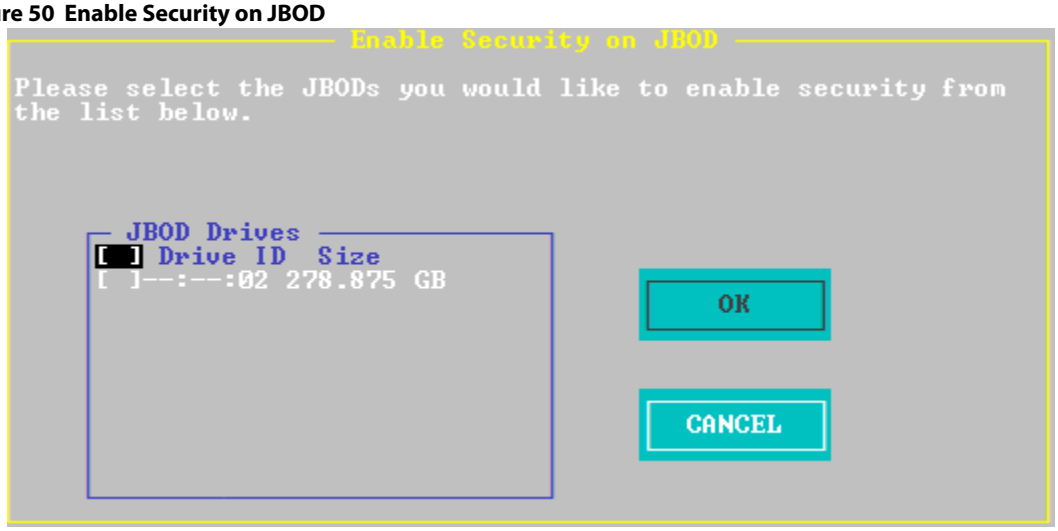

3. Select the JBOD drives for which you want to enable security.
   To select or deselect all the JBOD drives at one go, select the top most square brackets in the **JBOD drives** box.
4. Click **OK**.
   The security is enabled on all of the selected JBOD drives.

Perform the following steps to enable security on a JBOD drive from the **Drive Management** screen:

1. In the **Drive Management** screen, navigate to a JBOD drive, and press the F2 key.
2. Navigate to **Enable Security on JBOD**, and press Enter.
3. Click **OK** in the message confirmation box to continue.

4.14 Viewing and Changing Device Properties

This section explains how you can use the Ctrl-R Utility to view and change the properties for controllers, virtual drives, drive groups, physical drives, and BBUs.

4.14.1 Viewing Controller Properties

The Ctrl-R Utility shows information for one Avago SAS controller at a time. If your system contains multiple Avago SAS controllers, you can view information for a different controller by pressing the F12 key and selecting a controller from the list.

Navigate to the **Properties** menu to view the properties of the active controller.

The information in the **Properties** screen (Figure 19 on page 53) is read only. Most of this information is self-explanatory. To view additional properties, navigate to **Next**, and press Enter.

4.14.2 Modifying Controller Properties

You can change the properties of the controller in the **Ctrl Mgmt** menu.

Perform the following steps to change the controller properties:

1. Navigate to the **Ctrl Mgmt** menu to view the first **Controller Settings** screen.

2. You can change the values of the properties for the editable fields.

   To change additional properties, such as link speed, battery properties, and power settings, Write Verify properties, and large I/O support, click **Next** to go to the second **Controller Settings** screen.

3. Click **Apply**.

The following table describes all entries and options listed on both the **Controller Settings** screen. Leave these options at their default settings to achieve the best performance, unless you have a specific reason for changing them.

Table 22 Controller Settings

Options	Descriptions
Alarm Control	Select this option to enable, disable, or silence the onboard alarm tone generator on the controller.
Coercion Mode	Use this option to force 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 number you choose depends on how much the drives from various vendors vary in their actual size.
BIOS Mode	Specifies the following options to set the BIOS boot mode: ■ **Stop on Error**: Shows the errors encountered during boot up and waits for your input. The firmware does not proceed with the boot process until you take some action. ■ **Ignore Error**: Ignores errors and the firmware proceeds with boot. ■ **Pause on Error**: The firmware might halt because of hardware faults. If the firmware encounters no hardware faults, the boot up continues. ■ **SafeMode Error**: Boots the controller to run on safe mode.
Boot Device	Use this option to select the boot device from the list of virtual drives and JBODs.
Rebuild Rate	Use this option to select the rebuild rate for drives connected to the selected controller. The rebuild rate is the percentage of system resources dedicated to rebuilding a failed drive. The higher the number, the more system resources that are devoted to a rebuild. The range of rebuild rate is between 0 and 100 percent.
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 range of background initialization (BGI) rate is between 0 and 100 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 range of Consistency Check (CC) rate is between 0 and100 percent.
Recon. Rate	Use this option to select the amount of system resources dedicated to reconstruction of drives connected to the selected controller. The range of Recon rate is between 0 and100 percent.
Patrol Rate	Use this option to select the rate for patrol reads for drives connected to the selected controller. The patrol read rate is the percentage of system resources dedicated to running a patrol read. The range of patrol read is between 0 to 100 percent.
Cache Flush Interval	Use this option to control the interval at which the contents of the onboard data cache are flushed. The range of Cache Flush Interval is between 0 to100 seconds.
Spinup Delay	Use this option to control the interval (in seconds) between the spin-up of drives connected to this controller. The delay prevents a drain on the system's power supply that would occur if all drives spun up at the same time. The range of the Spin-up Delay is between 0 to 255 seconds.
Spinup Drive	Use this option to control the interval at which the contents of the onboard data cache are flushed. The range of Spin-up Drive is between 0 to 255 seconds.
Maintain PD Fail History	Use this option to maintain the history of all drive failures.

Table 22 Controller Settings (Continued)

Options	Descriptions
Device Exposure	Displays the actual number of devices to be exposed to the host. You can assign the following values: ■ 0 and 1 = Exposes all drives to the host ■ 2 to 255 = The actual number of devices to be exposed. For example, 4 = 4 devices, 10 = 10 devices exposed, 100 = 100 devices exposed and so on.
Enable Controller BIOS	Use this option to enable or disable the BIOS for the selected controller. If the boot device is on the selected controller, the BIOS must be enabled. Otherwise, the BIOS should be disabled, or you might be unable to use a boot device elsewhere.
Enable Stop CC on Error	Use this option to stop a consistency check when the controller BIOS encounters an error.
Auto Enhanced Import	Use this option to import automatically at boot time.
Set Factory Defaults	Use this option to load the default Ctrl-R Utility settings.
Manage Link Speed	Use this option to change the link speed between the controller and the expander, or between a controller and a drive that is directly connected to the controller.
Manage Power Save	Use this option to reduce the power consumption of drives that are not in use, by spinning down the unconfigured good drives, hot spares, and configured drives.
Start Manual Learn Cycle	The manual learn cycle re-calibrates the battery integrated circuit so that the controller can determine whether the battery can maintain the controller cache for the prescribed period of time in the event of a power loss.
Manage Battery	Use this option to view information about the BBU, if the selected controller has a BBU.
Emergency Spare	Use this option to commission unconfigured good drives or global hot spares as emergency spare drives. You can select from the options **None**, **UG** (Unconfigured Good), **GHS** (Global Hot spare), or **UG and GHS** (Unconfigured Good and Global Hot spare).
Enable Emergency for SMARTer	Use this option to commission emergency hot spare drives for predictive failure analysis events.
Write Verify	Use this option to verify if the data was written correctly to the cache before flushing the controller cache.
Large I/O Support	Use this option to enable or disable large I/O support feature. By default, large I/O support is disabled. A reboot is required if this property is changed. When this property is changed, `The controller property change has been performed successfully. Reboot the machine for the change to take effect` message is displayed.

4.14.3 Viewing and Changing Virtual Drive Properties

The Ctrl-R Utility shows the properties, policies, and the operations for virtual drives.

To view these items for the currently selected virtual drive and to change some of these settings, perform the following steps:

1. In the **VD Mgmt** screen, navigate to a virtual drive, and press the F2 key.
2. Press Enter.

   The **Virtual Drive Properties** dialog appears.

Figure 51 Virtual Drive Properties

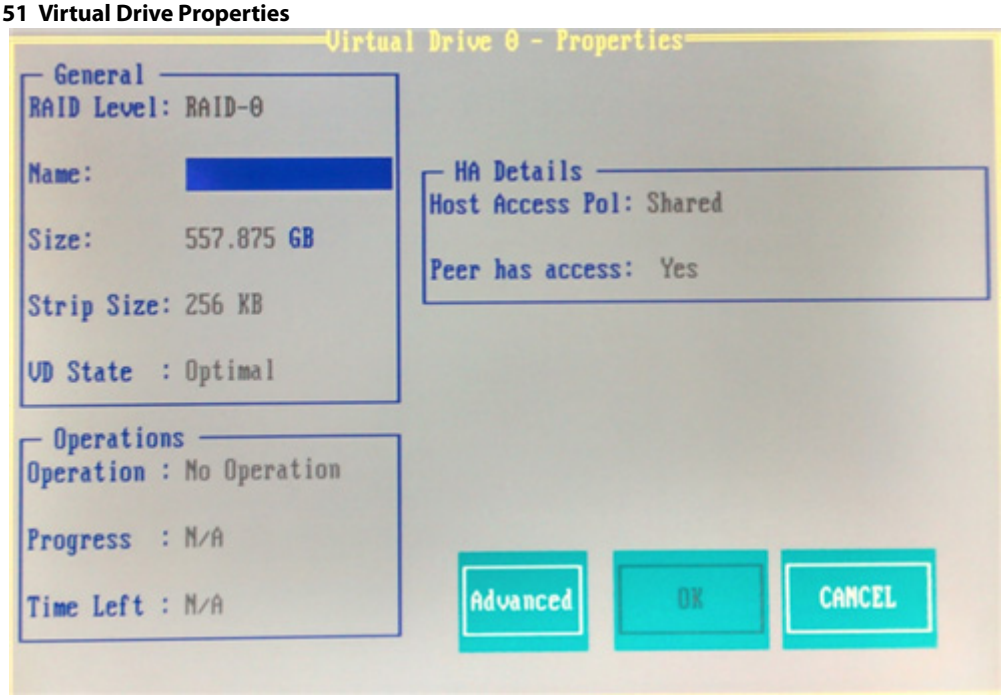

The **General** box shows the virtual drive's RAID level, name, state, size, and strip size.

The **Operations** box lists any operation (performed on the virtual drive) in progress, along with its progress status and the time remaining for the operation to be completed.

If High Availability DAS is supported on the controller, the **HA Details box** lists additional virtual drive properties; **Host access policy** and **Peer has access** appear on the **Properties** page.

— **Host access policy**

Indicates whether the virtual drive is shared between the servers in a cluster. The values for this property are **Shared** and **Exclusive**.

— **Peer has access**

Indicates whether the peer controller has access to the shared virtual drive. This property appears only if the virtual drive is shared.

3. Change the settings for the fields that are enabled in this dialog.

ATTENTION Before you change a virtual drive configuration, back up any data on the virtual drive that you want to save, or you might lose access to that data.

4. Click **OK** to save your changes.
5. Click **Advanced** to view additional virtual drive properties.

The **Advanced Properties** dialog appears.

Figure 52 Virtual Drive Management – Advanced Properties

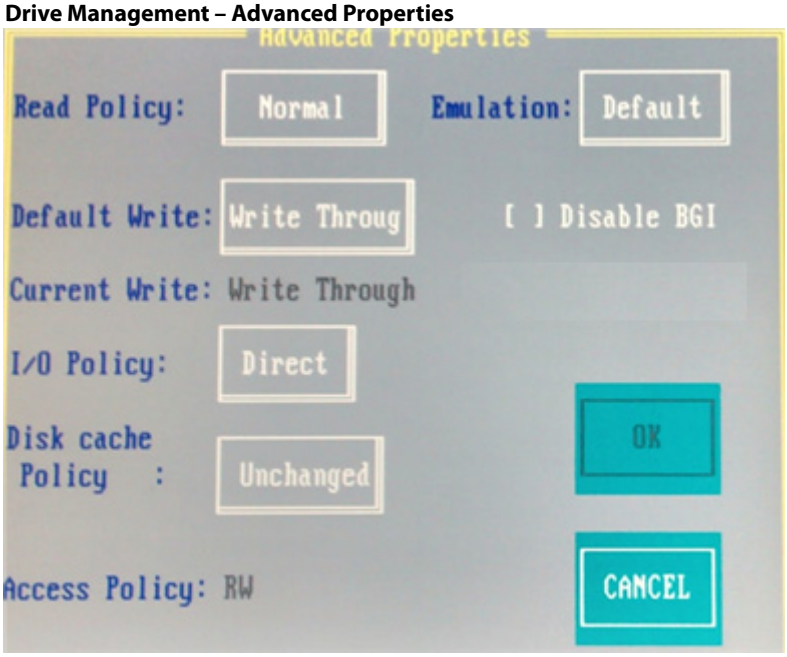

You can view the virtual drive policies that were defined when the storage configuration was created.

4.14.4 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 Ctrl-R Utility lists configurable drive groups where there is space to configure. If multiple virtual drives are defined on a single drive group, you can delete a virtual drive without deleting the entire drive group.

ATTENTION Back up any data that you want to keep before you delete a virtual drive.

Perform the following steps to delete a virtual drive:

1. In the **VD Mgmt** screen, navigate to the virtual drive, and press the F2 key.
2. Navigate to **Delete VD**, and press Enter.
   A message box appears, asking you to confirm the deletion.
3. Click **OK** to delete the virtual drive.

4.14.5 Deleting a Virtual Drive Group

You can delete a virtual drive group. On deleting a drive group, all the virtual drives in that drive group also are deleted.

Perform the following steps to delete a drive group:

1. In the **VD Mgmt** screen, navigate to a drive group, and press the F2 key.
2. Navigate to **Delete Drive Group**, and press Enter.
   The drive group is deleted and is removed from the **VD Mgmt** screen.

4.14.6 Expanding a Virtual Drive

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

Perform the following steps to expand the size of a virtual drive:

1. In the **VD Mgmt** screen, select the virtual drive whose size you want to expand and press the F2 key.
2. Navigate to **Expand VD**, and press Enter.

   The **Expand Virtual Drive** dialog appears.

Figure 53 Expand Virtual Drive

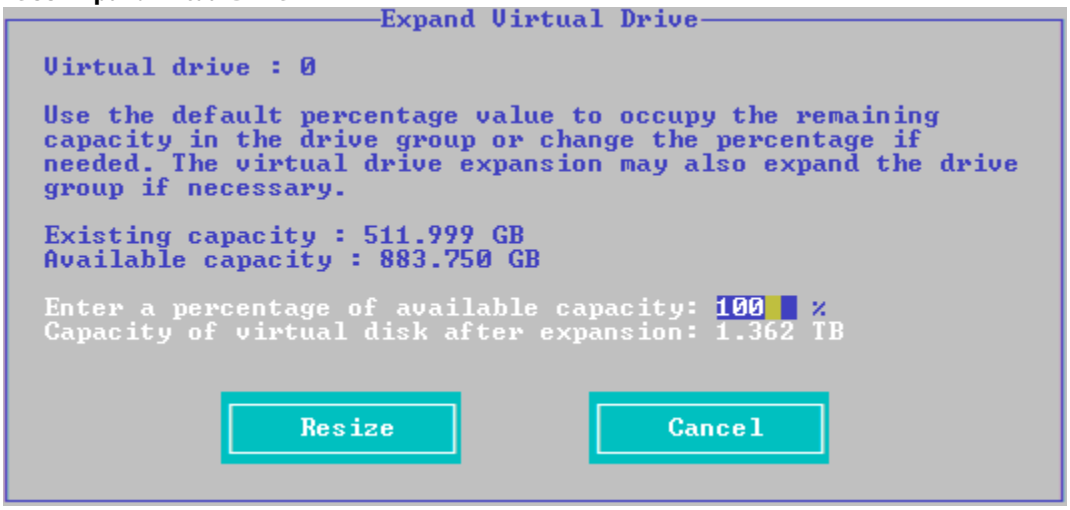

3. Enter the percentage of the available capacity that you want the virtual drive to use.

   For example, if 100 GB of capacity is available and you want to increase the size of the virtual drive by 30 GB, select 30 percent.
4. Click **Resize** to determine the capacity of the virtual drive after expansion.

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

4.14.7 Erasing a Virtual Drive

Virtual drive erase operates on a specified virtual drive and overwrites all user-accessible locations. It supports nonzero patterns and multiple passes. Virtual drive erase optionally deletes the virtual drive and erases the data within the virtual drive's logical base address range. Virtual drive erase is a background operation that posts events to notify users of their progress.

Perform the following steps to perform the virtual drive erase operation.

1. In the **VD Mgmt** screen, select a virtual drive, and press the F2 key.

2. Navigate to **Erase VD**, and press Enter.

   A menu appears displaying the following modes:

   — **Simple**

   Specifies a single-pass erase operation that writes pattern A to the virtual drive.

   — **Normal**

   Specifies a three-pass erase operation that first overwrites the virtual drive content with random values, then overwrites it with pattern A, and then overwrites it with pattern B.

   — **Thorough**

   Specifies a nine-pass erase operation that repeats the **Normal** erase three times.

   — **Stop Erase**

   Stops the erase operation that has already been started. This option is disabled at first. After the erase operation begins, this option is enabled.

3. Select a mode and press Enter.

   A message box appears.

Figure 54 Erase Virtual Drive

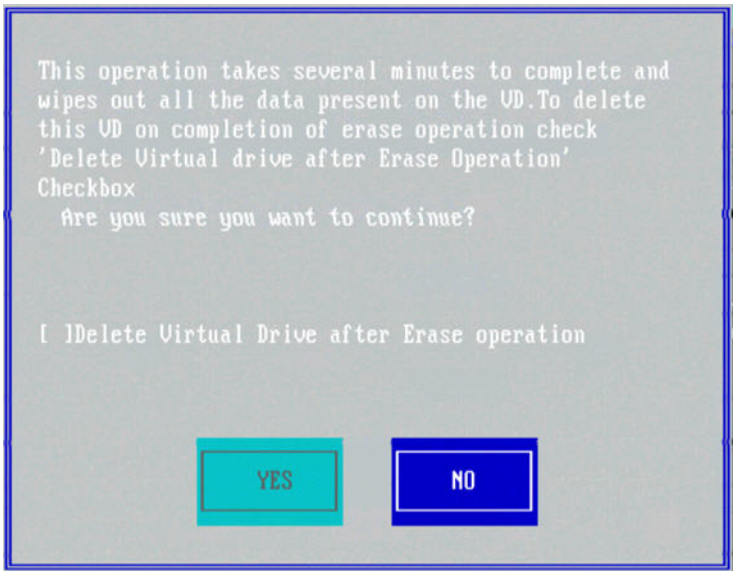

4. To delete the virtual drive after the erase operation has been completed, select the **Delete Virtual Drive after Erase operation** check box.
5. Click **Yes** for the erase operation to start.

   After the Drive Erase operation has started, the **Simple**, **Normal**, and **Thorough** options are disabled and the **Stop Erase** option is enabled.

4.14.8 Managing Link Speed

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

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

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

Perform the following steps to change the link speed:

1. In the **Controller Settings** screen, click **Next**.
   The second **Controller Settings** screen appears.
2. Click **Manage Link Speed**.
   The **Manage Link Speed** dialog appears.

Figure 55 Manage Link Speed

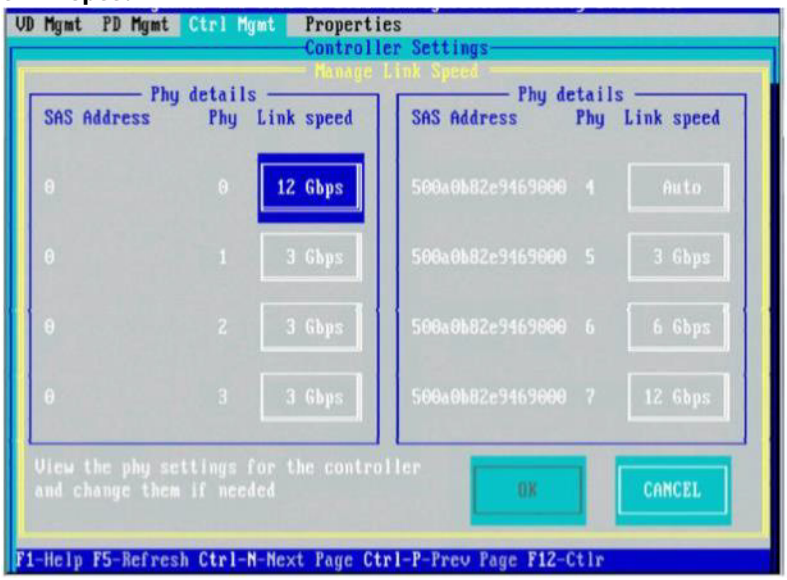

   — The **SAS Address** column shows the SAS address that uniquely identifies a device in the SAS domain.
   — The **Phy** column shows the system-supported phy link values. The phy link values are from 0 through 7.
   — The **Link Speed** column shows the phy link speeds.
3. Select the desired link speed by using the drop-down list.
   The link speed values are Auto,1.5Gb/s, 3Gb/s, 6Gb/s, or 12Gb/s.

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

4. Click **OK**.
   A message box appears, asking you to restart your system for the changes to take effect.
5. Click **OK**.
   The link speed value is now reset. The change takes place after you restart the system.

4.14.9 Managing Power Save Settings for the Controller

You can change the controller's power-save settings by using the Dimmer Switch enhancement (Power-Save mode).

Perform the following steps to change the power save settings:

1. Navigate to the second **Controller Settings** screen.
2. Navigate to **Manage Power Save**, and press Enter.
   The **Manage Power Save** dialog appears.

Figure 56 Manage Power Save

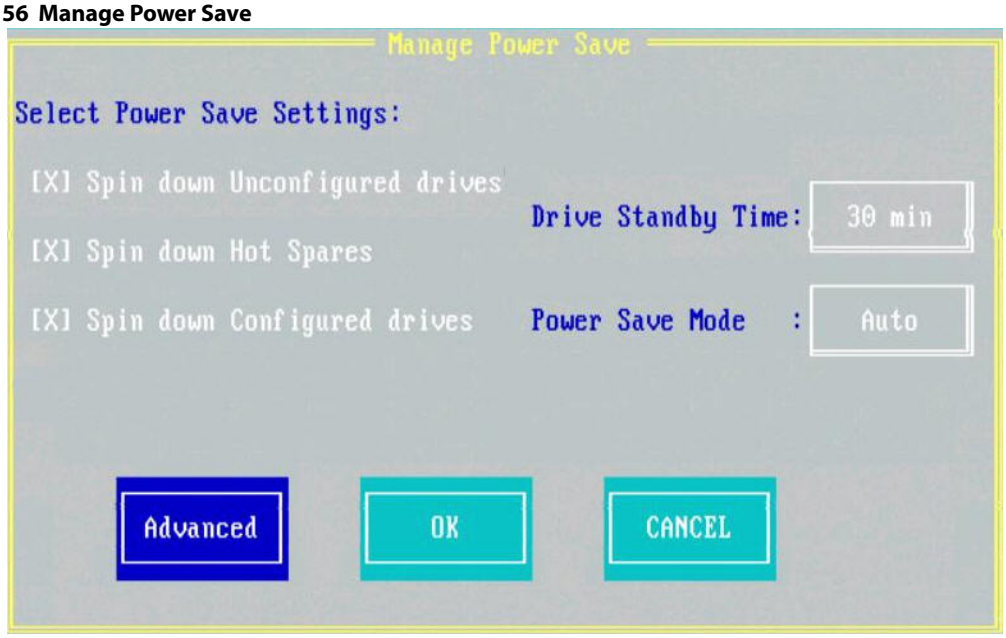

3. Select the **Spin down Unconfigured drives** check box to let the controller enable the unconfigured drives to enter the Power-Save mode.
4. Select the **Spin down Hot Spares** check box to let the controller enable the Hot spare drives to enter the Power-Save mode.
5. Select the **Spin down Configured Drives** check box to let the controller enable the Configured drives to enter the Power-Save mode.
6. Select the drive standby time from the **Drive Standby Time** drop-down list.

   NOTE The **Drive Standby Time** drop-down list is enabled only if any of the preceding check boxes are checked. The drive standby time can be 30 minutes, 1 hour, 90 minutes, or 2 hours through 24 hours.

7. Select the power save mode from the **Power- Save Mode** drop-down list.

   The mode can be **Auto** or **Max**.

   NOTE The **Power Save Mode** drop-down list is enabled only if the **Spin down Configured drives** check box is selected.

   You can click **Advanced** to set additional power save settings. This button is enabled only if the **Spin down Configured drives** check box is selected. The **Manage Power Save – Advanced** dialog appears, as shown in Figure 57.
8. Click **OK**.

   A message box appears, asking you to save the power-save settings.
9. Click **Yes** to save the settings.

4.14.9.1 Setting Advanced Power Save Settings

You can set additional power save properties in the **Manage Power Save – Advanced** dialog.

Perform the following steps to set advanced power save settings:

1. In the **Manage Power Save** dialog, click **Advanced**.

   The **Manage Power Save – Advanced** dialog appears.

Figure 57 Manage Power Save – Advanced

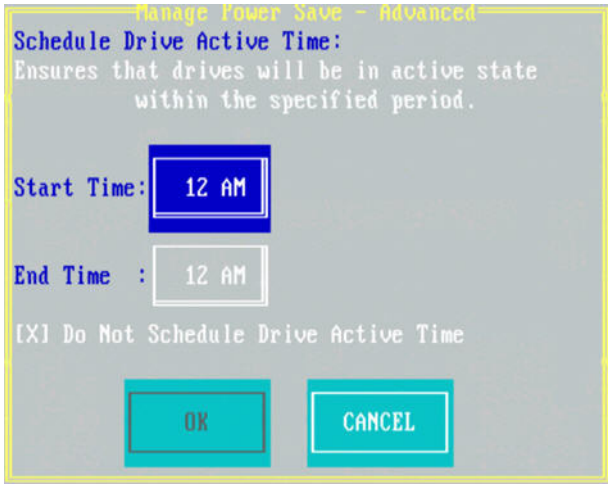

2. Perform one of these two actions:
   - Specify a start and end time, for the drives to be in active state, in the **Start Time** and **End Time** fields, respectively.
   - Select the **Do not Schedule Drive Active Time** check box to disable the **Start Time** and **End Time** fields.
3. Click **OK**.

4.14.10 Start Manual Learn Cycle

You can launch a cycle re-calibration of the battery integrated circuit so that the controller can determine whether the battery can maintain the controller cache for the prescribed period of time in the event of a power loss.

Re-calibrate the battery integrated circuit using the following steps:

1. Navigate to the second **Controller Settings** screen.
   The **Ctrl Mgmt – Controller Settings** dialog appears.

Figure 58 Controller Settings – Second Screen

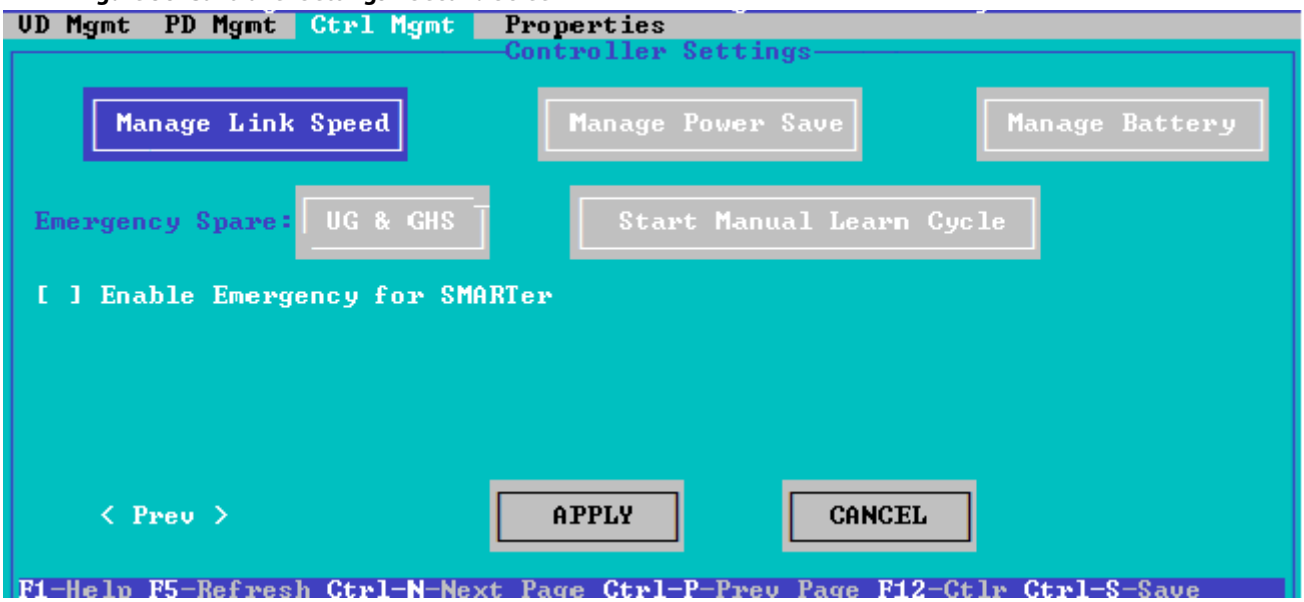

2. Click **Start Manual Learn Cycle**.
   An information box appears stating that the battery learn cycle will take a few minutes.

Figure 59 Manual Learn Cycle Warning

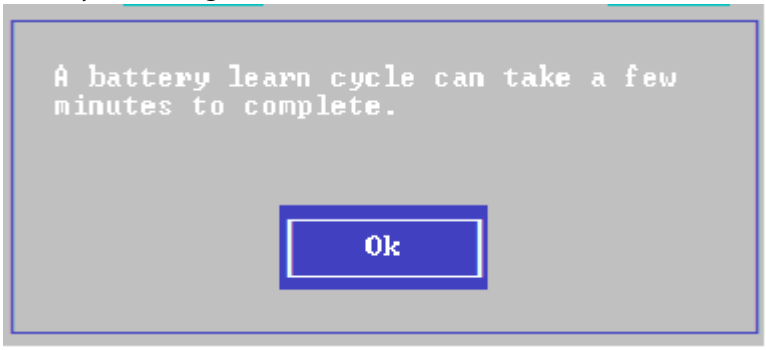

3. Click **Ok** to continue.

4.14.11 Managing Power Save Settings for the Drive Group

You can change the power save settings for a selected drive group.

Perform the following steps to change the power save settings for a drive group:

1. Navigate to a drive group in the **VD Mgmt** screen, and press the F2 key.
2. Navigate to **Manage Power Save Settings** and press Enter.
   The **Manage Power Save Settings** dialog appears.

Figure 60 Manage Power Save Settings – Drive Group

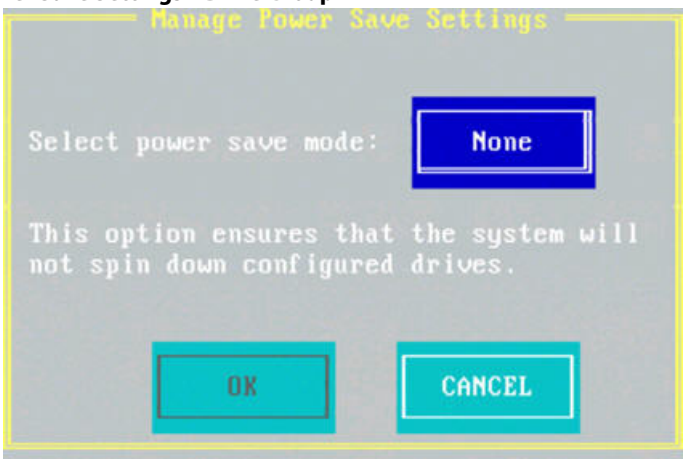

3. Select a power save mode from the **Select power save mode** drop-down list.
   A description of the selected mode appears in the dialog.
4. Click **OK**.

4.14.12 Managing BBU Information

If your SAS controller has a BBU, you can view information about it and change some settings.

A learning cycle is a battery calibration operation that the controller performs periodically to determine the condition of the battery.

Perform the following steps to view and change the battery settings:

1. Navigate to the second **Controller Settings** screen and select **Manage Battery**.
   The **Battery Properties** dialog appears. Most of the battery properties are read only.

Figure 61 Battery Properties

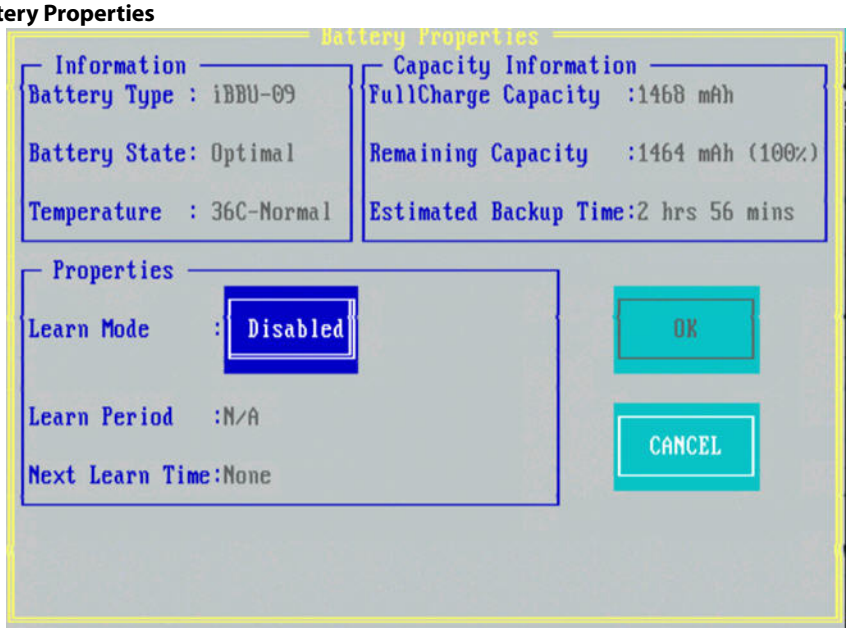

If the **Battery State** field has a value other than **Optimal**, the **Non-Optimal Reason** field appears at the bottom of the **Battery Properties** dialog. The **Non-Optimal Reason** field is a read-only field and states a reason for the non optimal state of the battery.

2. Select a battery learn mode from the **Learn Mode** drop-down list.
   The values in the drop-down list differ based on whether the battery supports transparent learn cycles.
   - If the battery supports transparent learn, the following values appear in the **Learn Mode** drop-down list:
   - **Transparent**
     The firmware tracks the time since the last learning cycle and performs a learn cycle when it is due.
   - **Disabled**
     The firmware does not monitor or initiate a learning cycle. You can schedule learning cycles manually.
   - **Unknown**
     The firmware warns about a pending learning cycle. You can start a learning cycle manually. After the learning cycle completes, the firmware resets the counter and warns you when the next learning cycle time is reached.
   - If the battery does not support transparent learn, the following values appear in the **Learn Mode** drop-down list:
   - **Automatic**
     The firmware tracks the time since the last learning cycle and performs a learn cycle when due. Write caching need not be disabled.
   - **Disabled**
     The firmware does not monitor or initiate a learning cycle. You can schedule learning cycles manually.
   - **Disabled (Warning Only)**
     The firmware never initiates a battery learn cycle but notifies you through events when a learn cycle is needed.
3. Click **OK** to change the learn mode.

4.14.13 Managing Dedicated Hot Spares

A dedicated hot spare is used to replace failed drives only in a selected drive group that contains the hot spare. You can create or delete dedicated hot spares in the **Virtual Drive Management** screen.

Perform the following steps to create or delete dedicated hot spares:

1. Navigate to a drive group in the **VD Mgmt** screen, and press the F2 key.
2. Navigate to **Manage Dedicated Hotspare**, and press Enter.

   The **Dedicated Hotspare** dialog appears, which shows a list of all hot spares that are available to create dedicated hot spares.

Figure 62 Dedicated Hotspare

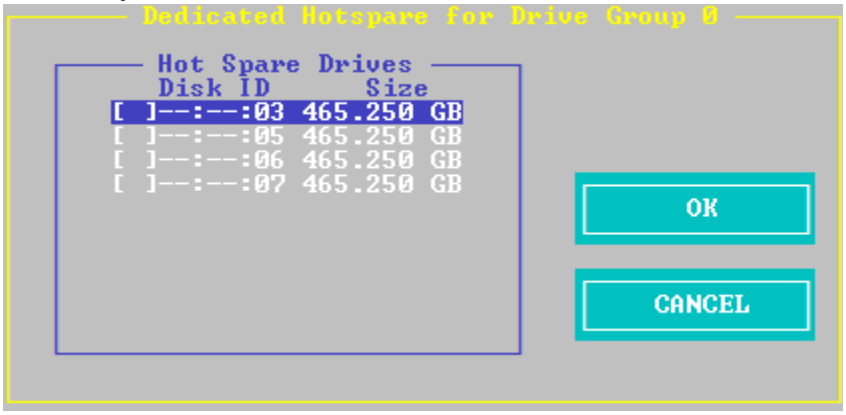

3. Perform one of these steps:
   - To create a dedicated hot spare, select a drive and click **OK**.
   - To delete a dedicated hot spare, deselect the hot spare and click **OK**.

4.14.14 Securing a Drive Group

If a drive group is created with FDE drives (security enabled drives) and at the time of creation, the security is set to **No**; later, you can secure that drive group using encryption.

Perform the following steps to secure a drive group:

1. Navigate to the **VD Mgmt** screen, navigate to the drive group that you want to secure, and press the F2 key.
2. Navigate to **Secure Drive Group**, and press Enter.

   A message box appears asking for your confirmation.
3. Click **Yes** to secure the drive group.

NOTE After a virtual drive is secured, you will not be able to remove the encryption without deleting the virtual drive.

4.14.15 Setting LED Blinking

You can use the **Locate** option to make the LEDs blink on the physical drives used by a virtual drive. You can choose to start or stop the LED blinking.

Perform the following steps to start or stop LED blinking:

1. Navigate to the **Drive Management** screen (in the **PD Mgmt** menu).

2. Select a physical drive, and press the F2 key.
3. Navigate to **Locate**, and press Enter.

   The **Start** and the **Stop** options appear.
4. Perform one of these actions:
   - Select **Start**, and press Enter to start LED blinking.
   - Select **Stop**, and press Enter to stop LED blinking.

NOTE Both the **Start** and **Stop** options of **Locate** only work if the drive is installed in a drive enclosure.

4.14.16 Performing a Break Mirror Operation

You can perform a Break Mirror operation on a drive group. The Break Mirror operation enables a RAID 1 configured drive group to be broken into two volumes. You can use one of the volumes in another system and replicate it without making a copy of the virtual drive.

Perform the following steps to perform a break mirror operation:

1. Navigate to the **VD Mgmt** screen, navigate to a drive group on which you want to perform the break mirror operation, and press the F2 key.
2. Navigate to **Break Mirror**, and press Enter.

   The following message box appears, asking for your confirmation.

Figure 63 Break Mirror

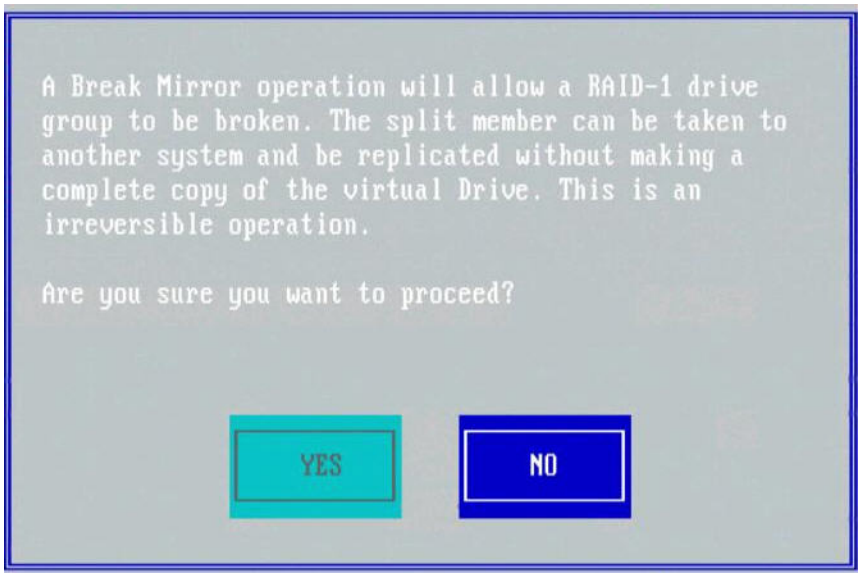

3. Click **Yes** to proceed.

4.14.17 Performing a Join Mirror Operation

You can perform a join mirror operation on a drive group to continue using the modified virtual drive or to reuse the original virtual drive.

Perform the following steps to perform a join mirror operation:

1. Navigate to the **VD Mgmt** screen, navigate to a drive group on which you want to perform the join mirror operation, and press the F2 key.

2. Navigate to **Join Mirror**, and press Enter.
   The following dialog appears.

Figure 64 Join Mirror – Choose Option

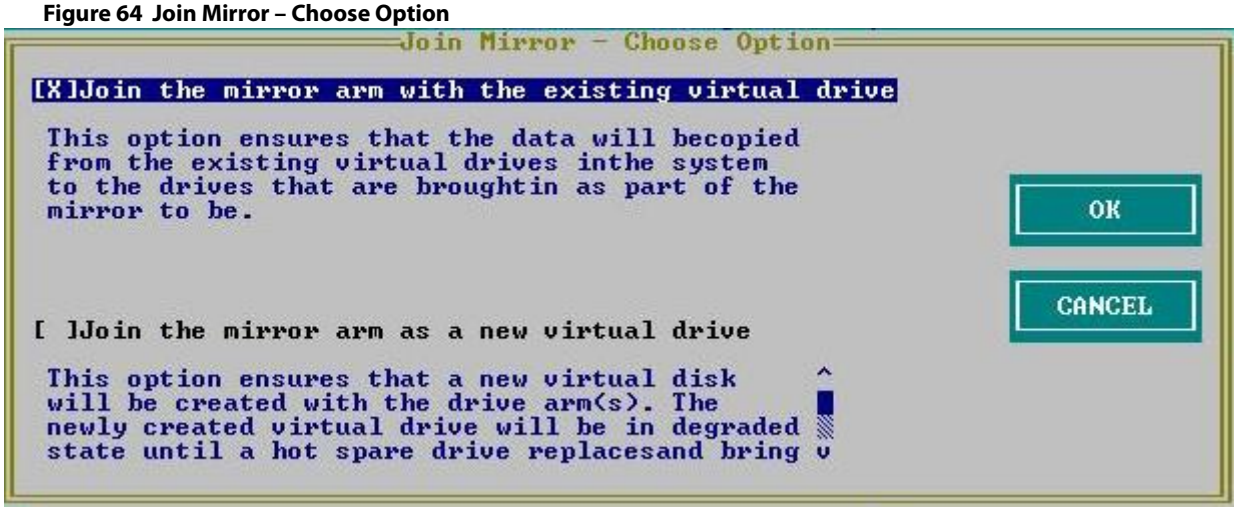

3. Select one of the options and click **OK**.
   If you select **Join the mirror arm with the existing virtual drive**, the following confirmation dialog appears.

Figure 65 Confirmation Message

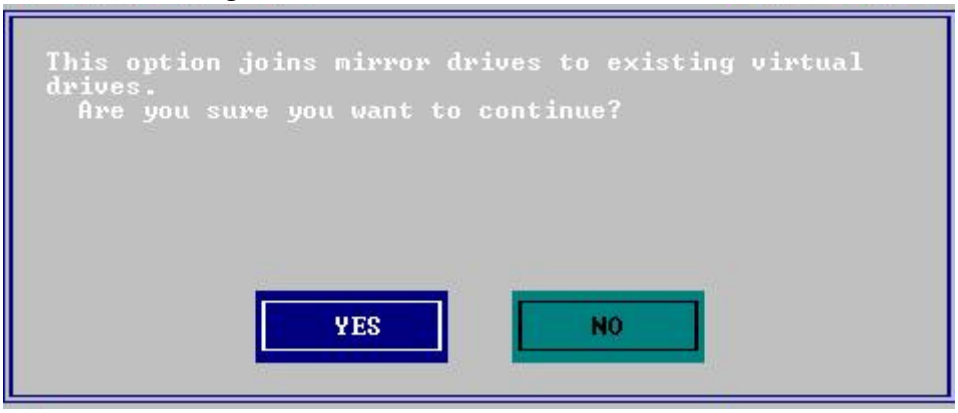

If you select **Join the mirror arm as a new virtual drive**, the following confirmation dialog appears.

Figure 66 Confirmation Message

4. Click **Yes** to proceed.
   The following dialog appears.

Figure 67 Join Mirror – Choose Option

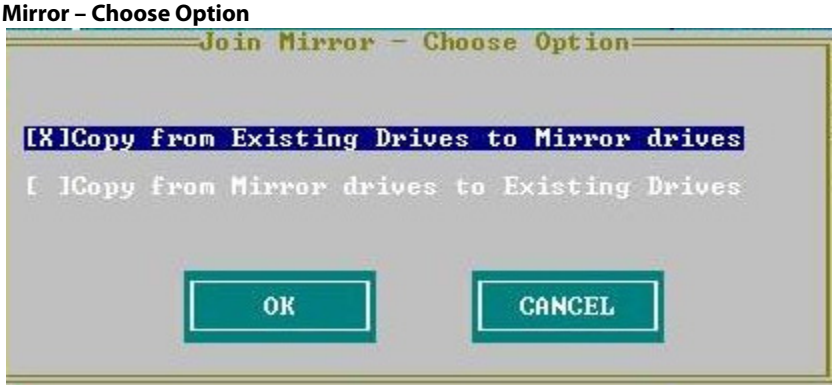

5. Select one of the options and click **OK**.

4.14.18 Hiding a Virtual Drive

You can hide a virtual drive on the controller.

Perform the following steps to hide a virtual drive:

1. In the **VD Mgmt** screen, select a virtual drive, and press the F2 key.
2. Navigate to **Hide VD**, and press Enter.
   A message box appears, asking you to confirm the operation.
3. Click **OK** to hide the virtual drive.

4.14.19 Unhiding a Virtual Drive

You can unhide a virtual drive on the controller.

Perform the following steps to unhide a virtual drive:

1. In the **VD Mgmt** screen, select a virtual drive, and press the F2 key.
2. Navigate to **Unhide VD**, and press Enter.
   A message box appears, asking you to confirm the operation.
3. Click **OK** to unhide the virtual drive.

4.14.20 Hiding a Drive Group

You can hide a drive group on the controller. If you hide a drive group, all of the virtual drives that are a part of this drive group become hidden.

Perform the following steps to hide a drive group:

1. In the **VD Mgmt** screen, select a drive group, and press the F2 key.
2. Navigate to **Hide Drive Group**, and press Enter.
   A message box appears, asking you to confirm the operation.
3. Click **OK** to hide the drive group.

4.14.21 Unhiding a Drive Group

You can unhide a drive group on the controller. If you unhide a drive group, all of the virtual drives that are a part of this drive group become unhidden.

Perform the following steps to unhide a drive group:

1. In the **VD Mgmt** screen, select a drive group, and press the F2 key.
2. Navigate to **Unhide Drive Group**, and press Enter.

   A message box appears, asking you to confirm the operation.
3. Click **OK** to unhide the drive group.

4.14.22 Performing a Transport Ready Operation

Transport Ready operation allows you to move a drive group and all the shared or exclusive dedicated hot spare drives associated with that drive group from one server to another server non-disruptively, in a real time.

NOTE To know more about the Transport Ready feature, see Transport Ready Pre-requisites section.

Follow these steps to perform a transport ready operation:

1. In the Ctrl-R **Ctrl Mgmt**, navigate to the **VD Mgmt** screen.
2. Navigate to the drive group on which want to perform transport ready and press the F2 Key.
3. Select **Transport Ready**.

Figure 68 Transport Ready

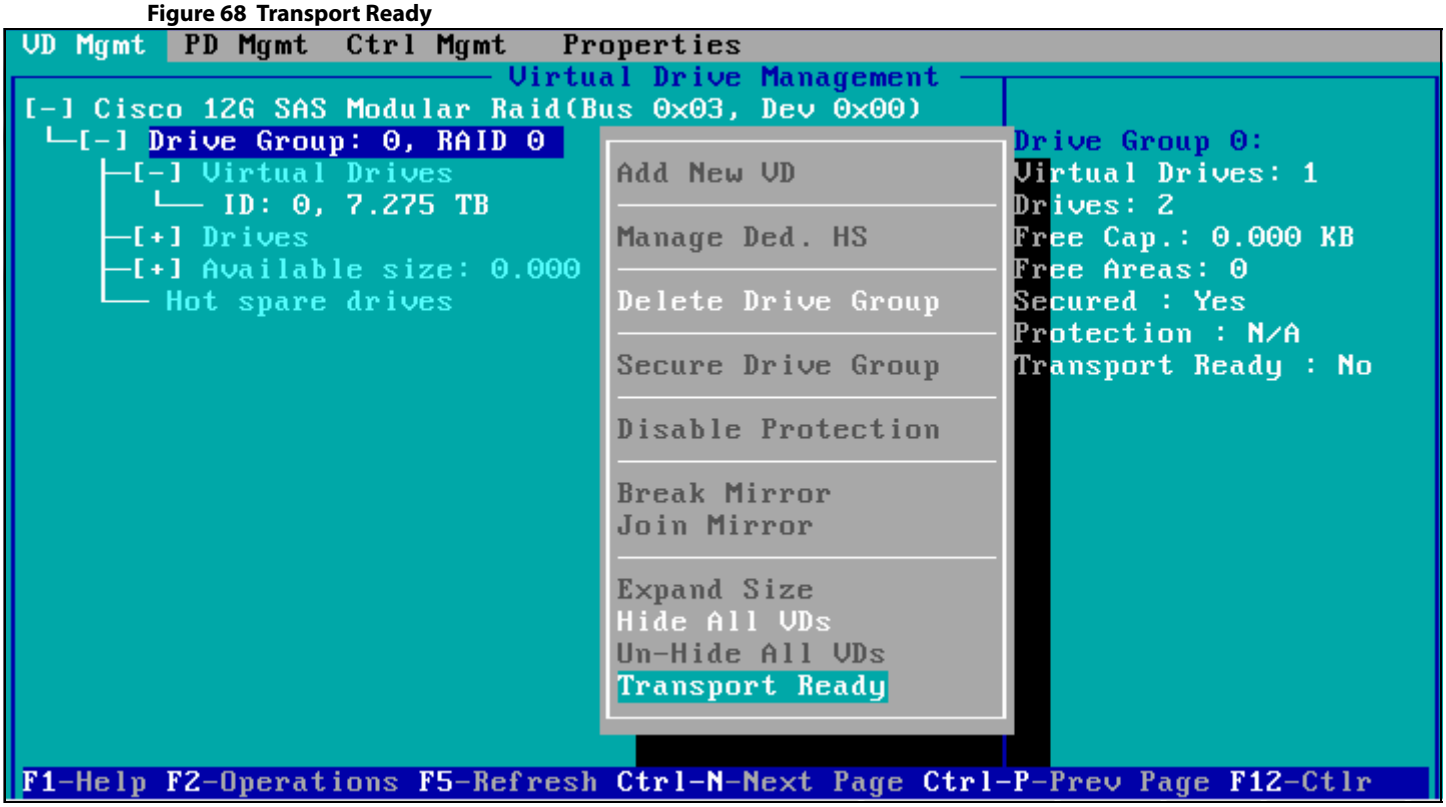

A warning message is displayed requesting you to confirm your action.

4. Select **Include Exclusive Dedicated Hot Spare Drive(s)**
5. Click **Yes**.

   If you have any `exclusive dedicated hot spare drives` associated with the drive group on which you are performing the Transport Ready operation, a warning message is displayed requesting you to confirm your action.

Figure 69 Transport Ready for Exclusive Dedicated Hot Spare Drives

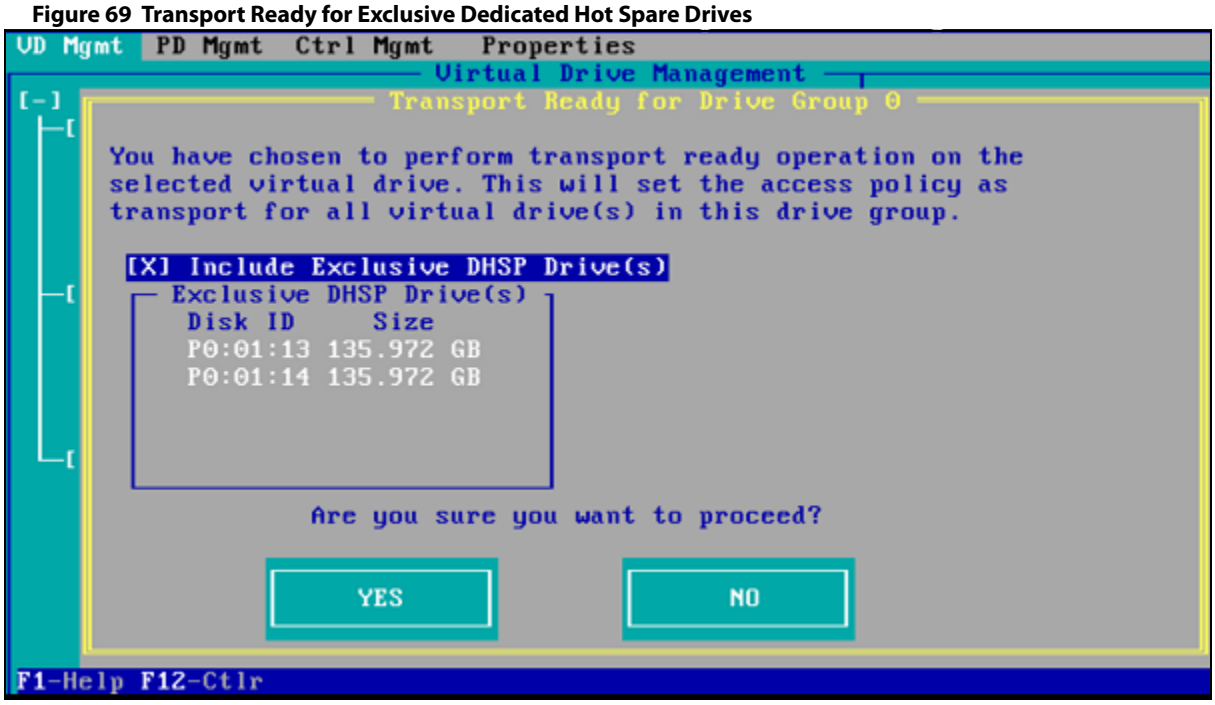

   a. Select the **Include Shared Dedicated Hot Spare Drive(s)** option.
   b. Click **`Yes`**.

   If you have any `shared dedicated hot spare drives` associated with the drive group on which you are performing the Transport Ready operation, a warning message is displayed requesting you to confirm your action.

Figure 70 Transport Ready for Shared Dedicated Hot Spare Drive(s)

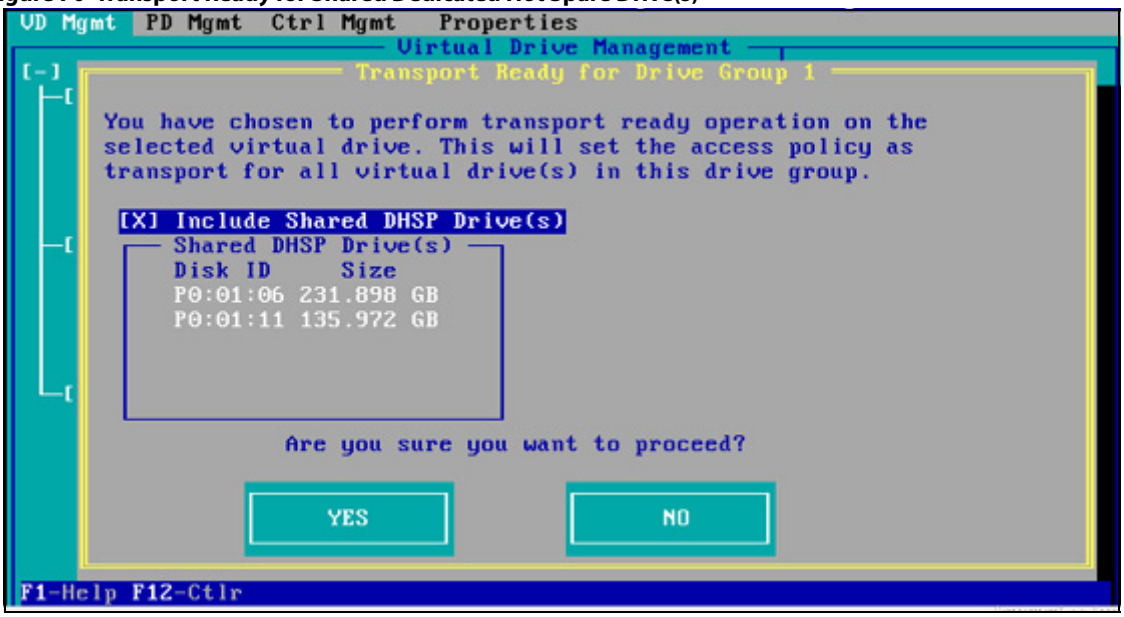

a. Select **Include Shared Dedicated Hot Spare Drive(s)** option.

b. Click **Yes**.

A warning message is displayed requesting you to confirm your action.

Figure 71 Transport Ready Confirmation Message

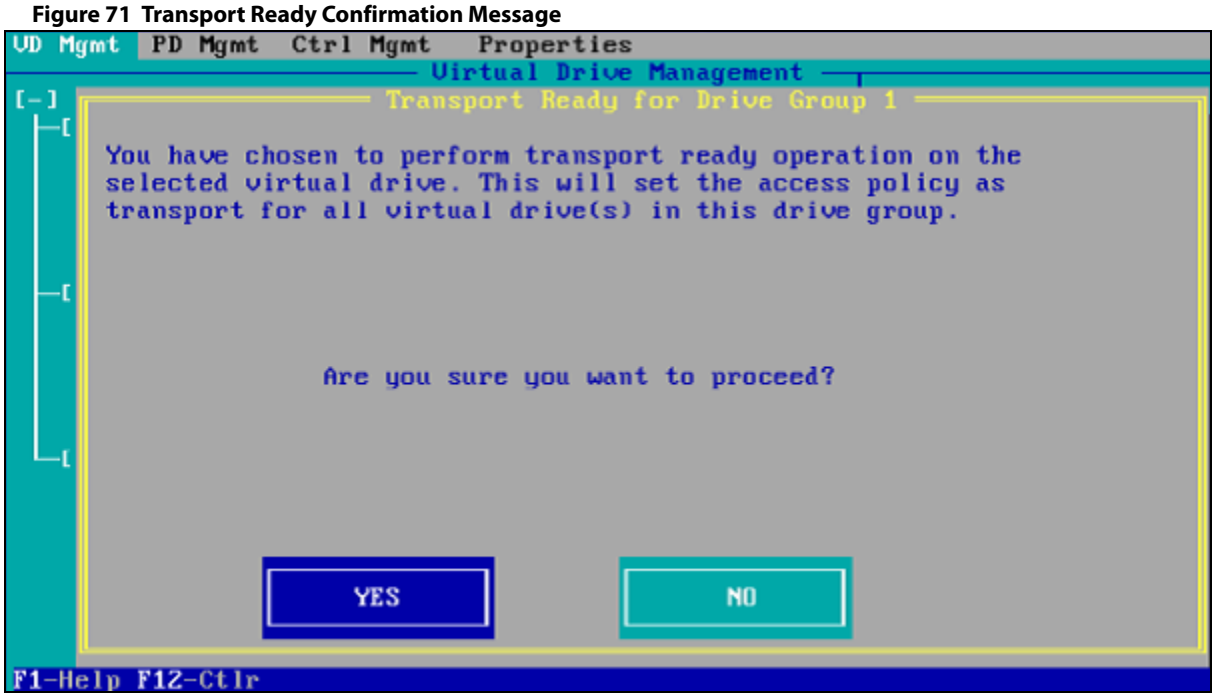

If you have both exclusive and shared dedicated host spare drives associated with the drive group on which you are performing the Transport Ready operation, a warning message is displayed requesting you to confirm your action.

Figure 72 Transport Ready for Shared Dedicated Hot Spare Drive(s) and Exclusive Dedicated Hot Spare Drive(s)

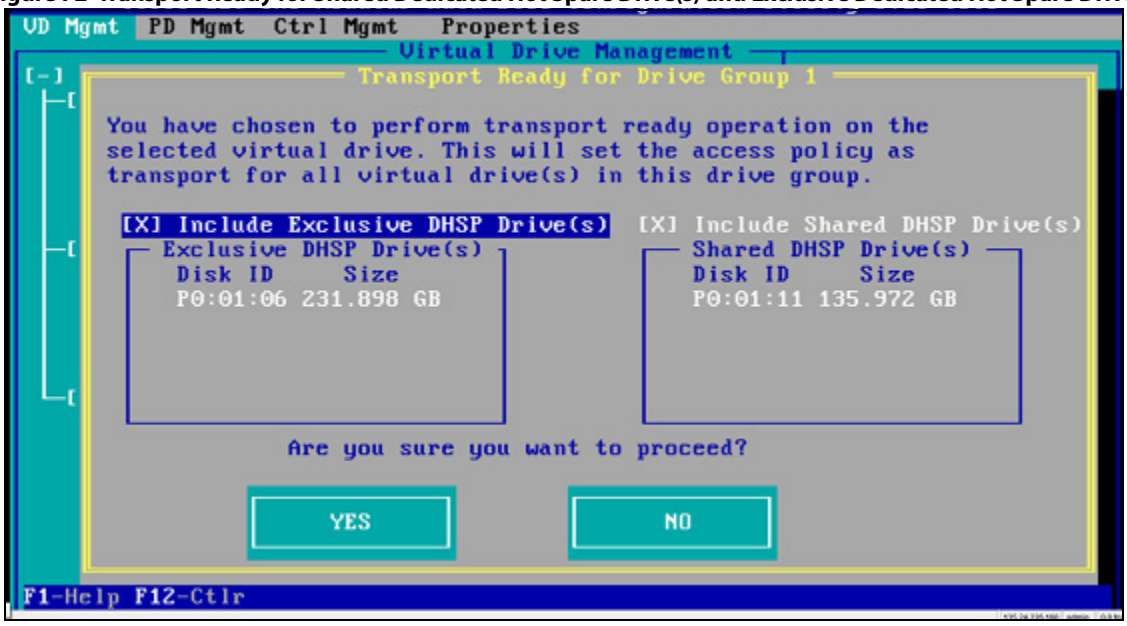

The drive group on which you have performed the Transport Ready operation is appended and a Transport Ready property is displayed on the drive group. Once the Transport Ready property is displayed on the drive group, you can remove the drive group

Figure 73 Access Policy set as Transport Ready

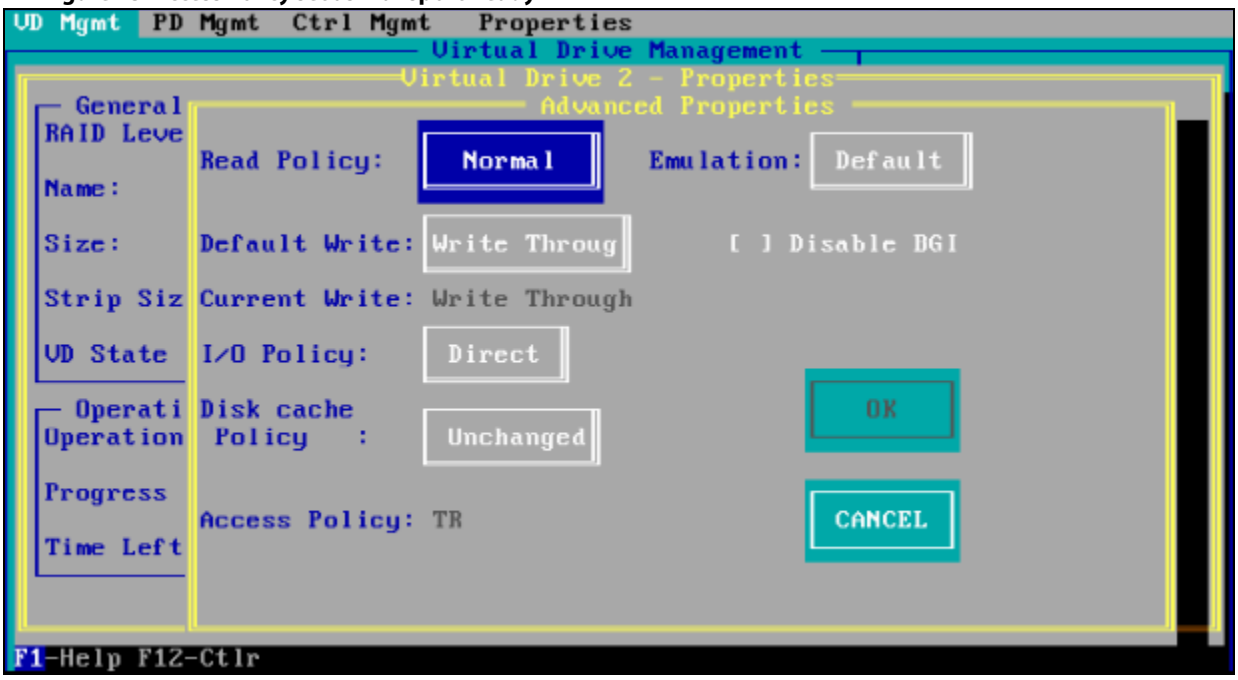

After the Transport Ready operation is performed successfully, the access policy of the virtual drive associated with the drive group on which the Transport Ready operation has been performed, is defaulted to Transport Ready. You will not have any other access policy options (such as Read Write, Read Only, and so on) available.

4.14.22.1 Clearing a Transport Ready Operation

If you have performed the Transport Ready operation on a drive group and if you want to reverse it (undo your previous action), you can do so using the Clear Transport Ready option.

Follow these steps to clear a Transport Ready operation:

1. In the **VD Mgmt** screen, select the drive group and press F2 Key.
2. Select **Clear Transport Ready**.

Figure 74 Clear Transport Ready

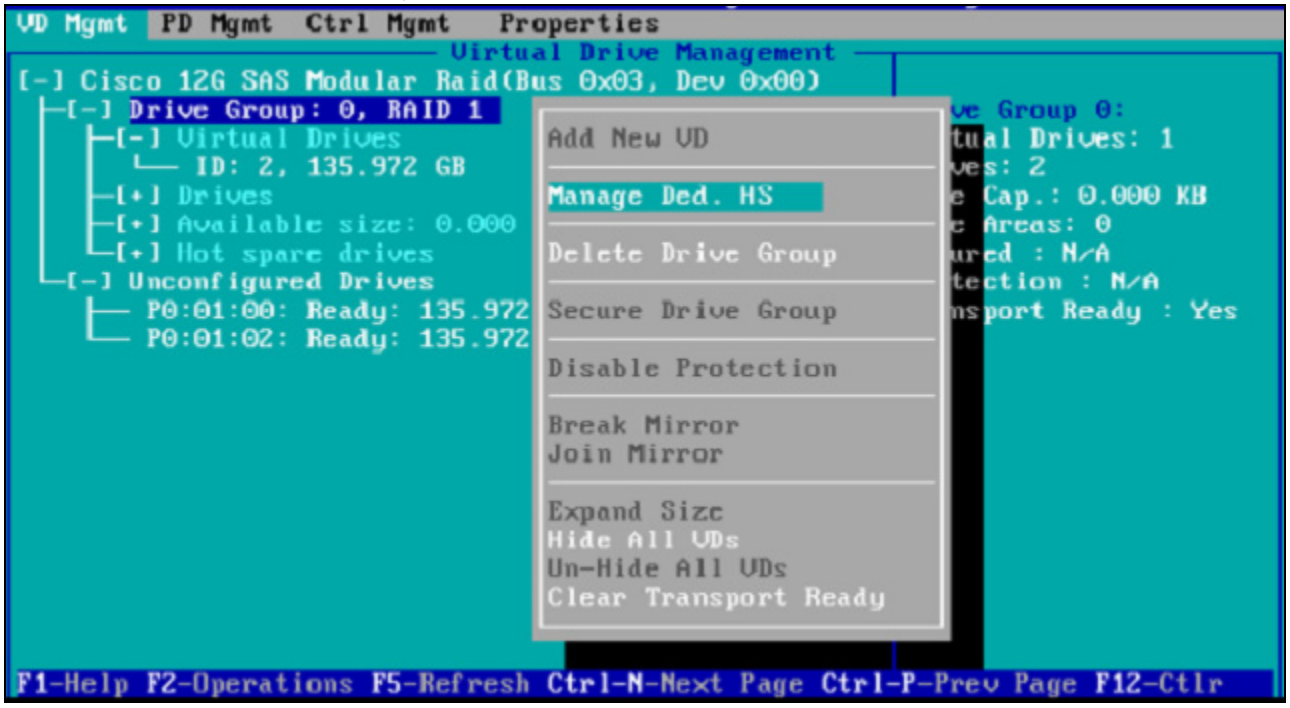

A warning message is displayed requesting you to confirm your action.

Figure 75 Confirm Clear Transport Ready

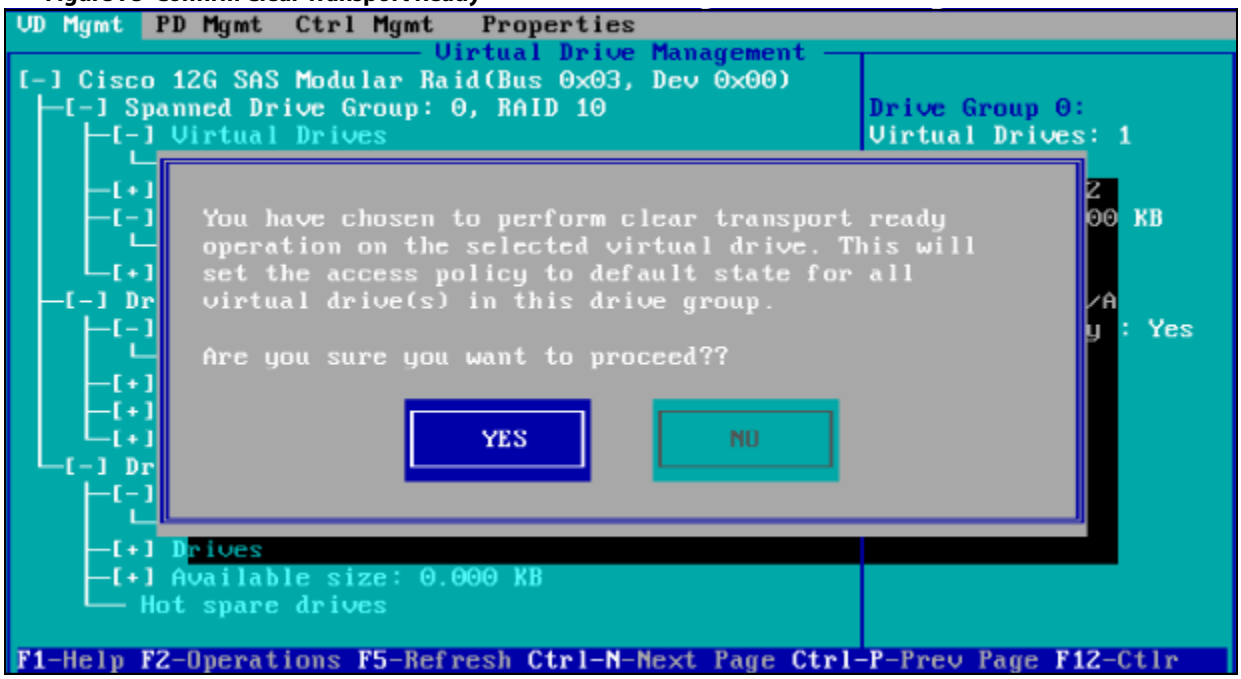

3. Select `Yes` and press Enter.

4.15 Managing Storage Configurations

This section describes how to use the Ctrl-R Utility to maintain and manage storage configurations.

4.15.1 Initializing a Virtual Drive

When you create a new virtual drive, the Ctrl-R Utility asks whether you would like to initialize the virtual drive. If you do not want to initialize the virtual drive at that stage, you can initialize the drive later.

Perform the following steps to initialize a virtual drive:

1. Navigate to the **VD Mgmt** screen, navigate to a virtual drive, and press the F2 key.
2. Select **Initialization**, and press Enter.
   The two initialization options, **Fast Init** and **Slow Init**, appear.
3. Select one of the two options, and press Enter.
   A confirmation dialog appears.

Figure 76 Initialize a Virtual Drive

4. Click **Yes** to begin initialization.

CAUTION Initialization erases all data on the virtual drive. Make sure to back up any data you want to keep before you initialize a virtual drive. Make sure the operating system is not installed on the virtual drive you are initializing.

4.15.2 Running a Consistency Check

You should periodically run a consistency check on fault-tolerant virtual drives (RAID 1, 5, 6, 10, 50, or 60 configurations; RAID 0 and RAID 00 do not provide data redundancy). A consistency check scans the virtual drive to determine whether the data has become corrupted and needs to be restored.

For example, in a system with parity, checking consistency means computing the data on one drive and comparing the results with the contents of the parity drive. You must run a consistency check if you suspect that the data on the virtual drive might be corrupted.

ATTENTION Make sure to back up the data before you run a consistency check, if you think the data might be corrupted.

Perform the following steps to run a consistency check:

1. Navigate to a virtual drive in the **VD Mgmt** screen, and press the F2 key.
2. Navigate to **Consistency Check**, and press Enter.
3. Navigate to **Start**, and press Enter.

   The consistency check starts and checks the redundant data in the virtual drive.

   If you attempt to run a consistency check on a virtual drive that has not been initialized, a confirmation dialog appears, asking for your confirmation.

Figure 77 Consistency Check

4. Click **Yes** to run the consistency check.

4.15.3 Rebuilding a Physical Drive

If a drive in a redundant virtual drive (RAID 1, 5, 6, 10, 50, or 60) fails, you must rebuild that drive on a hot spare drive to prevent data loss.

Perform the following steps to rebuild a physical drive:

1. Navigate to the **Drive Management** screen (in the **PD Mgmt** menu), and press the F2 key.
2. Select **Rebuild**, and press Enter.
   The rebuild operation starts.

4.15.4 Performing a Copyback Operation

You can perform a copyback operation on a selected drive.

The copyback operation copies 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 often creates or restores a specific physical configuration for a drive group (for example, a specific arrangement of drive group members on the device I/O buses).

Perform the following steps to perform the copyback operation:

1. Navigate to the **Drive Management** screen, navigate to a physical drive, and press the F2 key.
2. Navigate to **Copyback**, and press Enter.
   The following dialog appears.

Figure 78 Copyback Operation

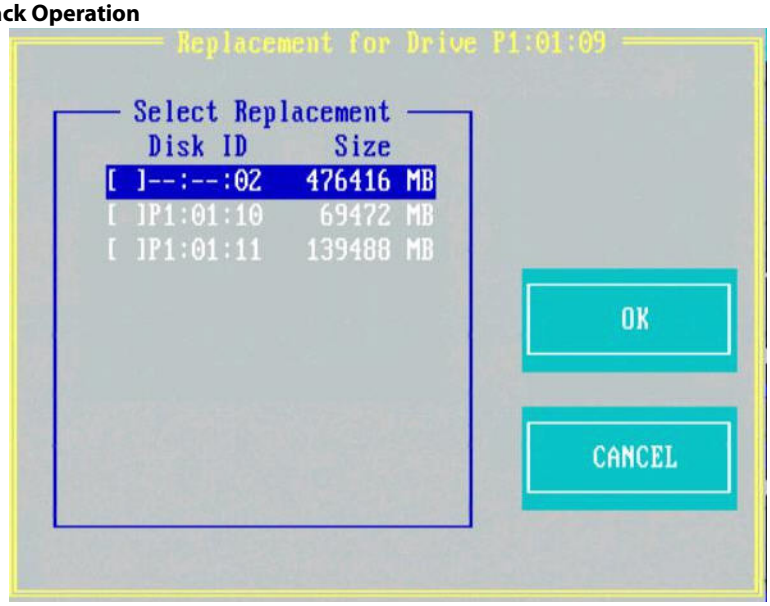

3. Select the replacement drive to which you want the data copied.
4. Click **OK**.

   The copyback operation is performed on the selected drive.

4.15.5 Removing a Physical Drive

You might sometimes need to remove a non-failed drive that is connected to the controller. Preparing a physical drive for removal spins the drive into a power save mode.

Perform the following steps to prepare a physical drive for removal:

1. Navigate to the **Drive Management** screen, and press the F2 key.
2. Select **Prepare for Removal**, and press Enter.

   The physical drive is now in a power save mode.

   If you change your mind and do not want to remove the drive, navigate to **Undo Removal**, and press Enter.

4.15.6 Creating Global Hot Spares

A global hot spare is used to replace a failed physical drive in any redundant array, as long as the capacity of the global hot spare is equal to or larger than the coerced capacity of the failed physical drive.

You can designate the hot spare to have enclosure affinity. In an enclosure affinity, if drive failures are present on a split backplane configuration, the hot spare first is used on the backplane in which it resides.

Perform the following steps to create global hot spares:

1. Navigate to the **Drive Management** screen, navigate to a physical drive that you want to change to a hot spare, and press the F2 key.
2. Select **Make Global HS**, and press Enter.

   The physical drive is changed to a global hot spare. The status of the physical drive as a global hot spare appears in the **Drive Management** screen.

4.15.7 Removing a Hot Spare Drive

Perform these steps to remove a hot spare drive:

1. Navigate to the **Drive Management** screen, navigate to a hot spare drive that you want to remove, and press the F2 key.
2. Select **Remove Hot Spare drive**, and press Enter.
   The hot spare drive is removed.

4.15.8 Making a Drive Offline

If a drive is part of a redundant configuration and you want to use it in another configuration, you can remove the drive from the first configuration and change the drive state to Unconfigured Good.

ATTENTION After you perform this procedure, all data on that drive is lost.

Perform the following steps to remove the drive from the configuration without harming the data on the virtual drive:

1. Navigate to the **Drive Management** screen, select a physical drive, and press the F2 key.
2. Navigate to **Place Drive Offline**, and press Enter.
   The drive status changes to Unconfigured Good.

ATTENTION After you perform this step, the data on this drive is no longer valid.

4.15.9 Making a Drive Online

You can change the state of a physical drive to online. In an online state, the physical drive works normally and is a part of a configured virtual drive.

Perform the following steps to make a physical drive online:

1. Navigate to the **Drive Management** screen, select a physical drive, and press the F2 key.
2. Navigate to **Place Drive Online**, and press Enter.
   The state of the physical drive changes to Online.

4.15.10 Instant Secure Erase

You can erase data on SED drives by using the **Instant Secure Erase** option in the **PD Mgmt** menu.

Perform the following steps to erase data on SED drives:

1. Navigate to the **Drive Management** screen, select a physical drive and press the F2 key.
2. Navigate to **Instant Secure Erase**, and press Enter.
   A confirmation dialog appears, asking whether you would like to proceed.
3. Click **Yes** to proceed.

4.15.11 Erasing a Physical Drive

You can securely erase data on Non SEDs (normal HDDs) by using the **Drive Erase** option in the **PD Mgmt** menu.

For Non–SEDs, the erase operation consists of a series of write operations to a drive that overwrites every user-accessible sector of the drive with specified patterns. It can be repeated in multiple passes using different data patterns for enhanced security. The erase operation is performed as a background task.

Perform the following steps to erase data on Non-SEDs:

1. Navigate to the **Drive Management** screen, select a physical drive and press the F2 key.
2. Navigate to **Drive Erase**, and press Enter.

   A menu appears displaying the following modes:

   — **Simple**

   Specifies a single pass operation that writes pattern A to the physical drive.

   — **Normal**

   Species a three pass erase operation that first overwrites the physical drive content with random values, then overwrites it with pattern A and then overwrites it with pattern B.

   — **Thorough**

   Specifies a nine pass erase operation that repeats the **Normal** erase operation three more times.

   — **Stop Erase**

   This option is disabled. This option is disabled at first. After the erase operation begins, this options is enabled.
3. Select a mode and press Enter.

   When you select **Simple**, **Normal**, or **Thorough**, a confirmation dialog appears.
4. Click **Yes** on the confirmation dialog to proceed with the drive erase operation.

   After the Drive Erase operation has started, you are intimated with the progress of the operation. Also, the **Simple**, **Normal**, and **Thorough** modes are disabled and the **Stop Erase** mode is enabled.

Chapter 5: HII Configuration Utility

The Avago MegaRAID Human Interface Infrastructure (HII) configuration utility is a tool used to configure controllers, physical disks, and virtual disks, and to perform other configuration tasks in a pre-boot, Unified Extensible Firmware Interface (UEFI) environment.

5.1 Starting the HII Configuration Utility

Follow these steps to start the HII configuration utility and to access the Dashboard View.

1. Boot the computer and press the appropriate key to start the setup utility during bootup.

   NOTE The startup key might be F2 or F1 or some other key, depending on the system implementation. Refer to the on-screen text or the vendor-specific documentation for more information.

2. When the initial window appears, highlight **System Settings** and press Enter.

   The **System Settings** dialog appears.

3. Highlight **Storage** and press Enter.

   The **Controller Selection** menu appears.

   The **Controller Selection** menu dialog lists the Avago MegaRAID controllers installed in your computer system. Use the PCI slot number to differentiate between controllers of the same type.

4. Use the arrow keys to highlight the controller you want to configure and press Enter.

   The **Dashboard View** appears as shown in the following figure. The **Dashboard View** shows an overview of the system. You can manage configurations, controllers, virtual drives, drive groups, and other hardware components from the **Dashboard View**.

Figure 79 Dashboard View

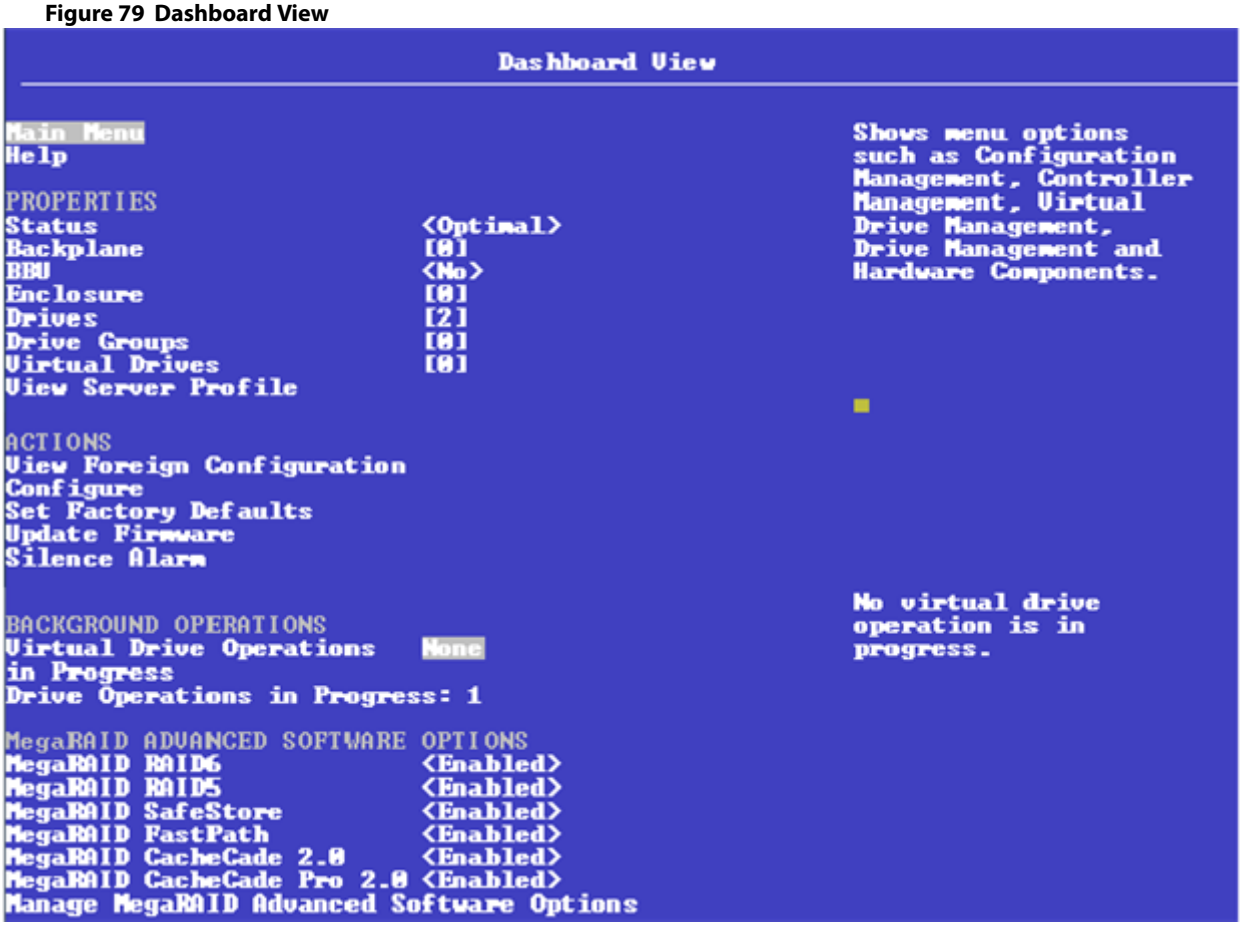

5.2 HII Dashboard View

The following sections describe the **Dashboard View**.

5.2.1 Main Menu

When you select the **Main Menu** option in the **Dashboard View**, the **Main Menu** dialog appears. The **Main Menu** provides various menu options to configure and manage controllers, virtual drives, drive groups, and hardware components. When the controller is running in Safe Mode, the **Main Menu** includes the warning message as shown in the following figure.

Figure 80 Main Menu – Safe Mode

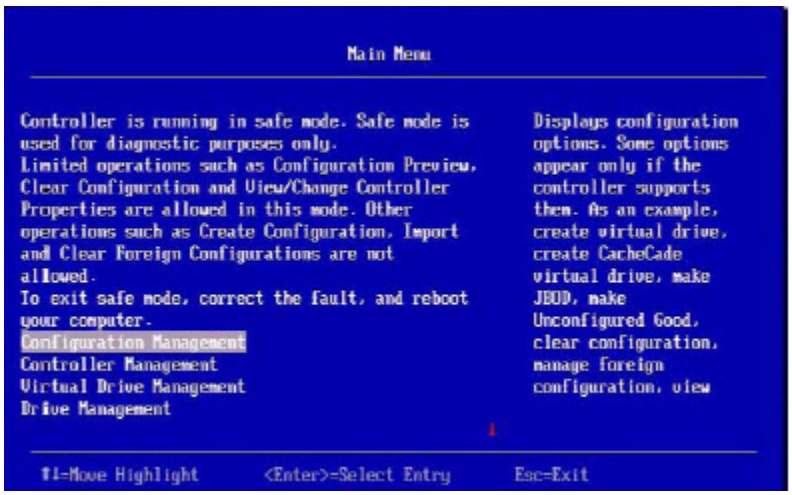

1. Select one of the following menu options:
   - Select **Configuration Management** to perform tasks, such as creating virtual drives, viewing drive group properties, viewing hot spare information, and clearing a configuration. For more information, see Managing Configurations.
   - Select **Controller Management** to view and manage controller properties and to perform tasks, such as clearing configurations, scheduling and running controller events, and running patrol reads. For more information, see Managing Controllers.
   - Select **Virtual Drive Management** to perform tasks, such as viewing virtual drive properties, locating virtual drives, and running a consistency check. For more information, see Managing Virtual Drives.
   - Select **Drive Management** to view physical drive properties and to perform tasks, such as locating drives, initializing drives, and rebuilding a drive after a drive failure. For more information, see Managing Physical Drives.
   - Select **Hardware Components** to view battery properties, manage batteries, and manage enclosures. For more information, see Managing Hardware Components.

5.2.2 HELP

The **HELP** section displays the HII utility context-sensitive help. It displays help strings for the following functions:

- Discard Preserved Cache
- Foreign Configuration
- Configure
- Silence Alarm

NOTE The help strings are displayed for the Discard Preserved Cache function only if pinned cache is present, and the help strings are displayed for the Foreign Configuration function only if the foreign configuration is present.

5.2.3 PROPERTIES

The **PROPERTIES** section displays the following information.

Figure 81 Dashboard View – PROPERTIES

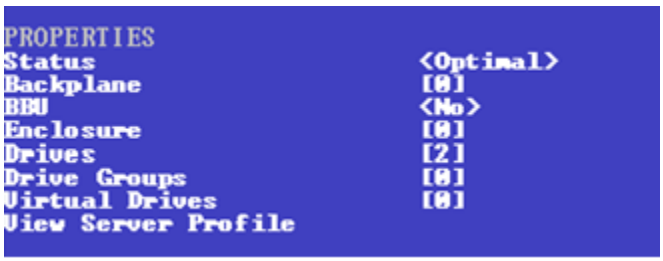

- **Status**

  Displays the status of the controller.
- **Backplanes**

  Displays the total number of backplanes connected to the controller.
- **BBU**

  Displays whether the battery backup unit is present.
- **Enclosures**

  Displays the total number of enclosures connected to the controller.
- **Drives**

  Displays the total number of drives connected to the controller.
- **Drive Groups**

  Displays the number of drives groups.
- **Virtual Drives**

  Displays the number of virtual drives.
- **View Server Profile**

  Displays the UEFI specification version that the system supports and the following menu options, as shown in the following figure.

Figure 82 Dashboard View – PROPERTIES – Server Profile

— Select **Controller Management** to view and manage controller properties and to perform tasks, such as clearing configurations, scheduling and running controller events, and running patrol reads.

For more information, see Managing Controllers.

— **Hardware Components** to view battery properties, manage batteries, and manage enclosures.

For more information, see Managing Hardware Components.

— **Drive Management** to view physical drive properties and to perform tasks, such as locating drives, initializing drives, and rebuilding a drive after a drive failure.

For more information, see Managing Physical Drives.

— **Virtual Drive Management** to perform tasks, such as viewing virtual drive properties, locating virtual drives, and running a consistency check.

For more information, see Managing Virtual Drives.

5.2.4 ACTIONS

The **ACTIONS** section displays some actions that you can perform on the controller:

Figure 83 Dashboard View – ACTIONS

- **Discard Preserved Cache**

  To discard the preserved cache for the selected controller, highlight **Discard Preserved Cache**, press Enter.

  ATTENTION If any foreign configurations exist, import them before discarding the preserved cache. Otherwise, you might lose data that belongs with the foreign configuration.

  NOTE The **Discard Preserved Cache** option is displayed only if pinned cache is present on the controller.

- **View Foreign Configuration**

  Helps you to preview and import a foreign configuration and clear a foreign configuration. It also displays the final configuration before the foreign configuration is imported or cleared. See Managing Foreign Configurations.

  NOTE If there are secured virtual drives, make sure you enter the pass-phrase.

- **Configure**

  Displays configuration options. See Managing Configurations.

- **Set Factory Defaults**

  Resets the controller to its factory settings.

- **Update Firmware**

  To update the controller's firmware, highlight **Update Firmware** and press Enter. The **Controller Firmware Update** window appears. See Upgrading the Firmware.

- **Silence Alarm**

  To silence the alarm on the controller, highlight **Silence Alarm** and press Enter.

  NOTE This option is disabled if the Alarm Control is disabled.

5.2.5 BACKGROUND OPERATIONS

This section displays the total number of background operations in progress for the virtual drives and the drives. If no background operations are in progress, it displays **None**.

When background operations for the virtual drives or drives are in progress, you can click the numbers to navigate to the **Virtual Drive Management** dialog or the **Drive Management** dialog, respectively. From these dialogs, you can click a specific virtual drive or a drive to view the progress of the operation and stop or suspend the operation. You can also view the basic properties and advanced properties of the virtual drives or drives.

Figure 84 Dashboard View – BACKGROUND OPERATIONS

5.2.6 MegaRAID ADVANCED SOFTWARE OPTIONS

This section displays the enabled advanced software options, such as the RAID levels, MegaRAID SafeStore, MegaRAID FastPath, MegaRAID CacheCade 2.0, and MegaRAID CacheCade Pro 2.0. This section also allows you to configure and use the advanced features. See Managing MegaRAID Advanced Software Options.

Figure 85 Dashboard View – MegaRAID ADVANCED SOFTWARE OPTIONS

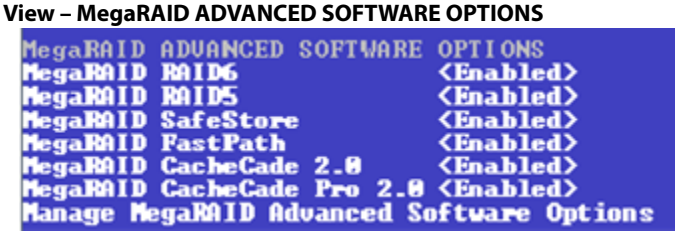

5.3 Critical Boot Error Message

The HII Configuration Utility shows an error screen with the title **Critical Message**, if preserved cache related to a missing drive in a virtual drive exists. This message can occur if a drive has failed or accidentally disconnected from the system, or for any other reason the drive is not visible to the system. This message appears pre-POST and must be addressed to continue a boot.

NOTE Some of the error messages that appear in the **Critical Message** screen might have spaces in them. This is a known limitation.

If this message appears when the system is started, perform these steps to resolve the problem:

1. Check the cabling that connects all of the drives to the system.

   Make sure that all of the cables are well connected and that the host bus adapter (if applicable) is securely seated in its slot.
2. If your system has activity LEDs, make sure that all of the LEDs do not show a fault.
3. If a cabling or connection issue does not exist with the physical drives, the problem might be the driver.

   Press C or Y in the input field when prompted by the critical boot error screen until no more screens appear. Then press Esc to exit, and the driver installs.
4. If these steps do not fix the problem, contact the Avago Customer Support team for further assistance.

5.4 Managing Configurations

When you select **Configuration Management** from the **Main Menu** or the **Configure** options in the **Dashboard View**, the **Configuration Management** dialog appears, as shown in the following figure.

Figure 86 Configuration Management

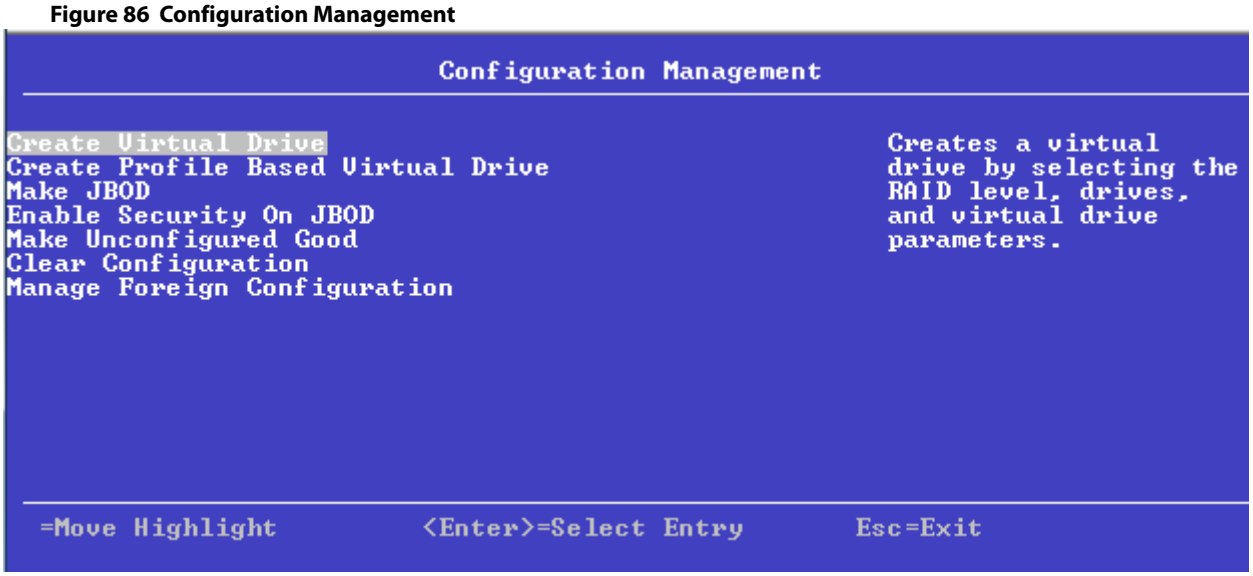

Only the first two menu options appear if no virtual drives have been created on this controller. The Make JBOD, Enable Security on JBOD, and Make Unconfigured Good options are included for some controllers. (See Make Unconfigured Good, Make JBOD, and Enable Security on JBOD.) You can enable security on the JBOD drives either from the Configuration Management screen or the Drive Management Screen. The following are the prerequisites for enabling security on JBOD drives:

- The JBOD drive must be an SED-capable drive.
- The controller must support the security feature.
- The controller must support the JBOD functionality.

The Manage Foreign Configuration option is included for some configurations. (See Managing Foreign Configurations.)

The HII utility supports 240 VD creation. For more information see the 240 Virtual Drive Feature Limitations appendix.

5.4.1 Creating a Virtual Drive from a Profile

To create a virtual drive from a profile, perform the following steps:

1. Select **Configuration Management** from the **Main Menu**.
2. Select **Create Profile Based Virtual Drive** from the **Configuration Management** menu.
3. Select a RAID level from the **Create Virtual Drive** menu. For example, select **Generic RAID 0**. The available RAID levels are: Generic RAID 0, Generic RAID 1, Generic RAID 5, and Generic RAID 6.

   The **Generic R0** dialog appears if you select Generic RAID 0 profile.

   The small red arrow at the bottom of the dialog indicates that you can scroll down to view more information.

NOTE The red arrow appears when there is too much information to display in one window. The amount of information that can be displayed in

one window depends on the capabilities of the HII browser. The **Save Configuration** option is not displayed in the previous figure.

4. Choose an option from the **Drive Selection Criteria** field (if more than one option exists).
5. Select **Save Configuration** to create the chosen profile.
6. Highlight **Confirm** and press the spacebar, then highlight **Yes** and press Enter.

You can create a virtual drive by using the profile shown in the previous figure. The following table describes the profile options.

Table 23 Virtual Drive Creation Profile Options

Option	Description
Drive Selection Criteria	You need to select one of the various combinations of options that exist. If only one option is possible, only one option appears.
Profile Parameters:	
Virtual Drive Name	Displays the name of the virtual drive.
RAID Level	Displays the RAID level based on the profile selected. For example, if the profile selected is Generic RAID 0, **RAID 0** is displayed.
Virtual Drive Size	Displays the amount of virtual drive storage space. By default, the maximum capacity available for the virtual drive is displayed. **NOTE** Virtual drive size of floating data type up to three decimal places is supported. Some of the screens in this chapter may not reflect this feature.
Power Save Mode	Displays the selected Power Save Mode of the five available options: **None**, **Auto**, **Max**, **Max without Cache**, and **Controller Defined**.
Strip Size	Displays the strip element size for the virtual drive. Drive Stripping involves partitioning each physical drive storage space in strips of the following sizes: **64 KB**, **128 KB**, **256 KB**, **512 KB**, **1 MB**.
Read Policy	Displays the read cache policy for the virtual drive. For any profile, if the drive is an SSD drive, the **No Read Ahead** option is displayed. Otherwise, the **Default** option is displayed. The possible options follow: ■ **Default** A virtual drive property that indicates whether the default read policy is **Read Ahead** or **No Read Ahead**. ■ **Read Ahead** Permits the controller to read requested data and store the additional data in cache memory, anticipating that the data is required soon. ■ **No Read Ahead** Specifies that the controller does not use Read Ahead for the current virtual drive.

Table 23 Virtual Drive Creation Profile Options (Continued)

Option	Description
Write Policy	Displays the write cache policy for the virtual drive. For any profile, if the drive is an SSD drive, the **Write Through** option is displayed. Otherwise, the **Always Write Back** option is displayed. The possible options follow: ■ **Write Back** The controller sends a data transfer completion signal to the host when the controller cache receives all of the data in a transaction. If you select the Write Back policy and the battery is absent, the firmware disables the Write Back policy and defaults to the Write Through policy. ■ **Write Through** The controller sends a data transfer completion signal to the host when the drive subsystem receives all the data in a transaction. ■ **Always Write Back** The controller sends a data transfer completion signal to the host when the controller cache receives all the data in a transaction. If you select the Always Write Back policy and the battery is absent, the firmware is forced to use the Write Back policy.
I/O Policy	Displays the Input/Output policy for the virtual drive. For any profile, if the drive is an SSD drive, the **Direct** option is displayed. The possible options follow: ■ **Default** A virtual drive property that indicates whether the default I/O policy is **Direct IO** or **Cached IO**. ■ **Direct IO** Data read operations are not buffered in the cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from the cache memory. (The I/O policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.) ■ **Cached IO** All read operations are buffered in cache.
Access Policy	The access policy for the virtual drive. The options are **Read/Write** and **Read Only**.
Disk Cache Policy	Displays the virtual drive cache setting. The possible options are **Unchanged**, **Enable**, and **Disable**.
Default Initialization	Displays the virtual drive initialization setting. The Default Initialization displays the following options: ■ **No** Do not initialize the virtual drive. ■ **Fast** Initializes the first 100 MB on the virtual drive. ■ **Full** Initializes the entire virtual drive.
Save Configuration	Saves the configuration that the wizard created.

The profile based virtual drive creation method has special requirements. The following table describes these requirements.

Table 24 Profile Based Virtual Drive Creation Requirements

Properties	Generic RAID0	Generic RAID1	Generic RAID5	Generic RAID6
HDD	Supported	Supported	Supported	Supported
SSD	Supported	Supported	Supported	Supported
SAS	Supported	Supported	Supported	Supported

Table 24 Profile Based Virtual Drive Creation Requirements (Continued)

Properties	Generic RAID0	Generic RAID1	Generic RAID5	Generic RAID6
SATA	Supported	Supported	Supported	Supported
PCIe	Supported	Supported	Supported	Not Supported
SED	Supported	Supported	Supported	Supported
NonSED	Supported	Supported	Supported	Supported
Protected Information (PI)	Supported	Supported	Supported	Supported
NonProtected Information (NonPI)	Supported	Supported	Supported	Supported
Sector Size (logical block format size) – 4 KB	Supported	Supported	Supported	Supported
Sector Size (logical block format size) – 512 B	Supported	Supported	Supported	Supported
Link speed – 3Gb/s	Supported	Supported	Supported	Supported
Link speed – 6Gb/s	Supported	Supported	Supported	Supported
Link speed – 12Gb/s	Supported	Supported	Supported	Supported
Direct attached	Supported	Supported	Supported	Supported
Backplane	Supported	Supported	Supported	Supported
Enclosure	Supported	Supported	Supported	Supported
Minimum number of PDs	1	2	3	4
Maximum number of PDs	0xFF	2	0xFF	0xFF
Power-save mode	Controller-defined	Controller-defined	Controller-defined	Controller-defined
Strip Size	256 KB	256 KB	256 KB	256 KB
Read Policy	If the drive is an SSD drive, then the No Read Ahead option appears. Else, the Default options appears.	If the drive is an SSD drive, then the No Read Ahead option appears. Else, the Default option appears.	If the drive is an SSD drive, then the No Read Ahead option appears. Else, the Default option appears.	If the drive is an SSD drive, then the No Read Ahead option appears. Else, the Default option appears.
Write Policy	If the drive is an SSD drive, then the Write Through option appears. Else, the Write Back option appears.	If the drive is an SSD drive, then the Write Through option appears. Else, the Write Back option appears.	If the drive is an SSD drive, then the Write Through option appears. Else, the Write Back option appears.	If the drive is an SSD drive, then the Write Through option appears. Else, the Write Back option appears.
IO Policy	If the drive is an SSD drive, then the Direct IO option appears. Else, the Default options appears.	If the drive is an SSD drive, then the Direct IO option appears. Else, the Default options appears.	If the drive is an SSD drive, then the Direct IO option appears. Else, the Default options appears.	If the drive is an SSD drive, then the Direct IO option appears. Else, the Default options appears.
Access policy	Read/Write	Read/Write	Read/Write	Read/Write
Disk Cache Policy	Enable	Unchanged	Unchanged	Unchanged
Initialization	Fast	Fast	Full	Full
Dedicated Hot Spare	Not Supported	Supported	Supported	Supported

Table 24 Profile Based Virtual Drive Creation Requirements (Continued)

Properties	Generic RAID0	Generic RAID1	Generic RAID5	Generic RAID6
Mixing of Media HDD and SSD drives	Not Supported	Not Supported	Not Supported	Not Supported
Mixing of Interface Type SAS and SATA drives	Not Supported	Not Supported	Not Supported	Not Supported
Mixing of PI and NonPI drives	Not Supported	Not Supported	Not Supported	Not Supported
Mixing SED and NonSED drives	Not Supported	Not Supported	Not Supported	Not Supported
Mixing of 1.5Gb/s, 3Gb/s, 6Gb/s, and 12Gb/s link speeds	Not Supported	Not Supported	Not Supported	Not Supported

5.4.2 Manually Creating a Virtual Drive

The following dialog appears when you select **Create Virtual Drive** from the **Configuration Management** menu.

Figure 87 Create Configuration Window

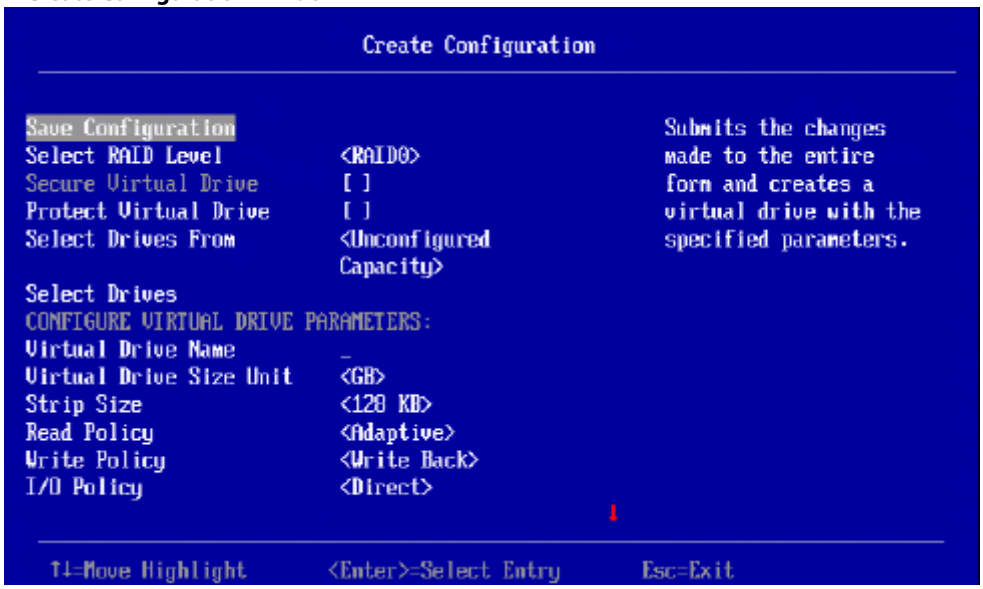

The small red arrow at the bottom of the window indicates that you can scroll down to view more information.

NOTE If your system detects any JBODs, the **Make Unconfigured Good** dialog appears before the **Create Configuration** window. The **Make Unconfigured Good** dialog lets you convert the JBOD drives to Unconfigured Good. See Make Unconfigured Good.

NOTE The red arrow appears when there is too much information to display in one dialog. The amount of information that can be displayed in one dialog depends upon the capabilities of the HII browser.

Perform these steps to select options for a new configuration (that is, a new virtual drive) on the controller.

1. Highlight the **Select RAID Level** field and press **Enter**.

2. Select a RAID level for the virtual drive from the pop-up menu.

   The available RAID levels are listed in the help text of the **Create Configuration** dialog. Some system configurations do not support all these RAID levels. See Table 26 for brief descriptions of the RAID levels.
3. To view the **Secure Virtual Drive** field, enable security and attach an FDE drive. If either is missing, the field is grayed out.
4. To view the **Protect Virtual Drive** field, enable protection and attach a protected drive.

   If either is missing, the field is grayed out.
5. If the security key is enabled, highlight the **Secure Virtual Drive** field to secure the new virtual drive.

   This field is not available unless the security feature is already enabled.
6. If protection is enabled, highlight the **Protect Virtual Drive**.

   This field is not available unless the protection feature is already supported by the controller.
7. Highlight the **Select Drives From** field, press Enter, and select **Unconfigured Capacity** or **Free Capacity**.

   Free capacity means the new virtual drive is created from unused (free) drive capacity that is already part of a virtual drive. *Unconfigured capacity* means the new virtual drive is created on previously unconfigured drives.
8. Highlight the **Virtual Drive Name** field, press Enter, and enter a name for the new virtual drive.
9. (Optional) Change the **Virtual Drive Size Unit** value by highlighting this field, pressing Enter, and selecting a value from the pop-up menu.

   The options are `MB`, `GB`, and `TB`.
10. (Optional) Change the default values for **Strip Size**, **Read Policy**, **Write Policy**, **I/O Policy**, **Access Policy**, **Drive Cache**, **Disable Background Initialization**, and **Default Initialization**.

    See Table 25 for descriptions of these options.
11. Highlight **Select Drives** and press Enter.

    The following dialog appears.

Figure 88 Select Drives Window

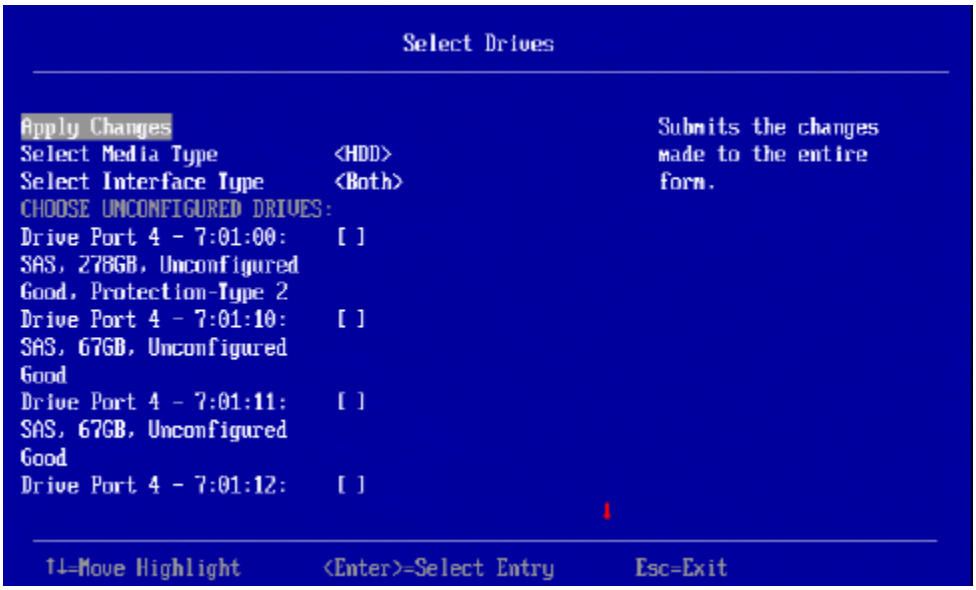

Follow these steps to select physical drives for the new virtual drive.

1. (Optional) Change the default **Select Media Type** by highlighting this field, pressing Enter, and selecting an option from the pop-up menu.

   The choices are `HDD` and `SSD`. Combining HDDs and SSDs in a single virtual drive is not supported.

2. (Optional) Change the default **Select Interface Type** by highlighting this field, pressing Enter, and selecting an option from the pop-up menu.

   The choices are **SAS**, **SATA**, and **Both**. Depending on the configuration of your system, combining SAS and SATA drives in a virtual drive might not be supported.

3. Select physical drives for the virtual drive by highlighting each drive and pressing the spacebar to select it.

   A small red arrow at the bottom of the window indicates you can scroll down to view more drives.

   NOTE The red arrow appears when there is too much information to display in one dialog. The amount of information that can be displayed in one dialog depends on the capabilities of the HII browser.

   Alternatively, use the **Select All** and **Deselect All** options at the bottom of the list of drives to select or deselect all available drives. If you select drives of varying sizes, the usable space on each drive is restricted to the size of the smallest selected drive.

   NOTE Be sure to select the number of drives required by the specified RAID level, or the HII utility will return you to the root menu when you try to create the virtual drive. For example, RAID 1 virtual drives use exactly two drives, and RAID 5 virtual drives use three or more virtual drives. See Table 26 for more information.

4. When you have selected all of the drives for the new virtual drive, highlight **Apply Changes** and press Enter to create the virtual drive.

   NOTE If you selected drives of varying sizes, the HII utility shows a message warning stating that the remaining free capacity on the larger drives would be unusable.

5. If the warning message about different size capacities appears, press the spacebar to confirm the configuration, then highlight **Yes** and press Enter.

   The HII utility returns you to the **Create Configuration** dialog.

6. Highlight **Save Configuration** and press Enter to create the virtual drive.

   A message appears confirming that the configuration is being created.

7. Highlight **OK** and press Enter to acknowledge the confirmation message.

The following table describes the policies that you can change when creating a virtual drive.

Table 25 Virtual Drive Policies

Property	Description
Strip Size	The virtual drive strip size per DDF. The possible values are as follows: ■ **7**: 64 KB ■ **8**: 1 MB
Read Policy	The read cache policy for the virtual drive. The possible values are as follows: ■ **Ahead** 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 process speeds up reads for sequential data, but provides little improvement when accessing random data. ■ **Normal** Read-ahead capability is disabled.
Write Policy	The write cache policy for the virtual drive. The possible values are as follows: ■ **Write Back** The controller sends a data transfer completion signal to the host when the controller cache receives all of the data in a transaction. If you select the Write Back policy and the battery is absent, the firmware disables the Write Back policy and defaults to the Write Through policy. ■ **Write Through** The controller sends a data transfer completion signal to the host when the drive subsystem receives all the data in a transaction. ■ **Always Write Back** The controller sends a data transfer completion signal to the host when the controller cache receives all the data in a transaction. If you select the Always Write Back policy and the battery is absent, the firmware is forced to use the Write Back policy.
I/O Policy	The I/O policy for the virtual drive. The possible values are as follows: ■ **Direct** Data 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 I/O policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.) ■ **Cached** All reads are buffered in cache.
Access Policy	The access policy for the virtual drive. The options are **Read/Write**, **Read Only**, and **Blocked**.
Drive Cache	The disk cache policy for the virtual drive. The possible values are **Unchanged**, **Enable**, and **Disable**.
Disable Background Initialization (BGI)	Specifies whether background initialization is enabled or disabled. When BGI is enabled, the firmware runs the initialization process in the background. When BGI is disabled, the initialization process does not start automatically and does not run in the background.
Default Initialization	Allows choice of virtual drive initialization option. The possible options are **No**, **Fast**, and **Slow**.

The following table describes the RAID levels that you can select when creating a new virtual drive. Some system configurations do not support RAID 6 and RAID 60.

Table 26 RAID Levels

Level	Description
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 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 00	Is a spanned drive group that creates a striped set from a series of RAID 0 drive groups to provide high data throughput, especially for large files.
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.

5.4.3 Creating a CacheCade Virtual Drive

A CacheCade virtual drive is a software virtual drive that enables SSDs to be configured as a secondary tier of cache to maximize transactional I/O performance for read-intensive applications. The following window appears when you select **Create CacheCade Virtual Drive** from the **Virtual Drive Management** window.

Figure 89 Create CacheCade Virtual Drive Window

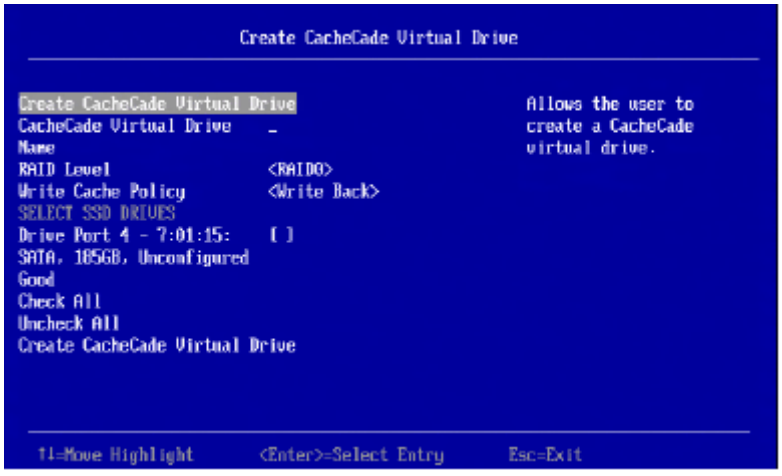

Follow these steps to create a CacheCade virtual drive.

1. Highlight **CacheCade Virtual Drive Name**, press Enter, and enter a name for the virtual drive.
2. Highlight the **RAID Level** field and press Enter.
3. Select a RAID level for the CacheCade virtual drive from the pop-up menu.

   The available RAID levels are listed in the help text of the **Create Configuration** window. Some system configurations do not support all these RAID levels. See Table 26 for brief descriptions of the RAID levels.

4. Highlight the **Write Cache Policy** field and press Enter.
5. Select a write cache policy from the popup menu. The choices are as follows:
   - **Write Through**: The controller sends a data transfer completion signal to the host when the virtual drive has received all of the data and has completed the write transaction to the drive.
   - **Write Back**: 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 virtual drive in accordance with policies set up by the controller. These policies include the amount of dirty and clean cache lines, the number of cache lines available, and the elapsed time from the last cache flush.
   - **Force Write Back**.
6. Highlight the available SSD drives listed in the window and press the spacebar to select them.

   Alternatively, highlight **Select All** and press Enter to select all available SSD drives for the virtual drive.
7. When you have selected all the SSD drives, highlight **Create CacheCade Virtual Drive** and press Enter to create the virtual drive.

5.4.4 Viewing Drive Group Properties

The following window appears when you select **View Drive Group Properties** from the **Virtual Drive Management** menu.

Figure 90 View Drive Group Properties Window

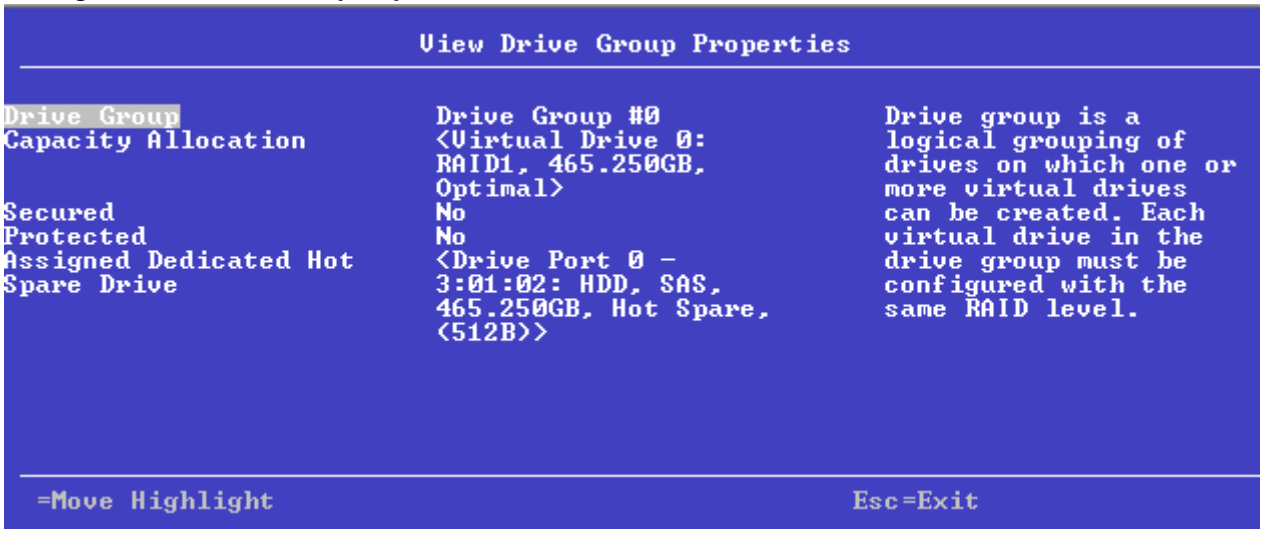

A drive group is a logical grouping of drives attached to a RAID controller on which one or more virtual drives can be created. Each virtual drive in the drive group must be configured with the same RAID level. This figure shows information for one drive group.

In this window, the Capacity Allocation entry for each drive group displays associated virtual drives for the drive group. The window also indicates whether the drive group is secured and protected. To see how much free space is available in the drive group, highlight a **Capacity Allocation** field and press Enter. The information appears in a pop-up window.

The **Assigned Dedicated Hot Spare Drive** field provides information about the dedicated hot spare drives that are assigned to this drive group. You can assign more than one dedicated Hot Spare drive to single drive group.

5.4.5 Viewing Global Hot Spare Drives

To view all the assigned global hot spare drives on the controller, select **View Global HotSpares** on the **Configuration Management** menu. The following figure shows a sample of the **View Global Hot Spare Drives** dialog.

Figure 91 View Global Hot Spare Drives

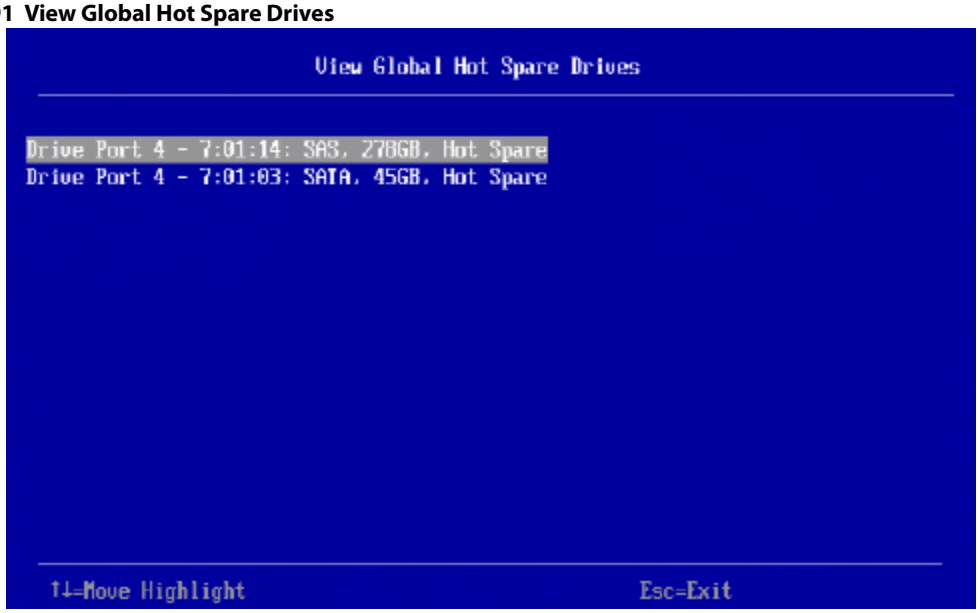

Press Esc to exit this window when you are finished viewing information.

5.4.6 Clearing a Configuration

A warning message dialog appears when you select **Clear Configuration** from the **Configuration Management** menu.

As stated in the warning text, this command deletes all virtual drives and hot spare drives attached to the controller.

ATTENTION All data on the virtual drives is erased. If you want to keep this data, be sure you back it up before using this command.

To complete the command, follow these steps:

1. Highlight the brackets next to **Confirm** and press the spacebar.
   An X appears in the brackets.
2. Highlight **Yes** and press Enter.
   A success message appears.
3. Highlight **OK** and press Enter.
   The HII Utility clears the configuration and returns you to the **Configuration Management** menu.

5.4.7 Make Unconfigured Good, Make JBOD, and Enable Security on JBOD

When you power down a controller and insert a new physical drive, if the inserted drive does not contain valid DDF metadata, the drive status is listed as JBOD (Just a Bunch of Disks) when you power the system again. When you power down a controller and insert a new physical drive, if the drive contains valid DDF metadata, its drive state is **Unconfigured Good**. A new drive in the JBOD drive state is exposed to the host operating system as a stand-alone drive. You cannot use JBOD drives to create a RAID configuration, because they do not have valid DDF records. First, the drives must be converted to an Unconfigured Good state.

If the controller supports JBOD drives, the **Configuration Management** menu of the HII utility includes options for converting JBOD drives to Unconfigured Good, or vice versa. You can also enable security on the JBOD drives.

NOTE If the controller supports JBOD drives, you can also change the status of JBOD drives to Unconfigured Good when you create a new configuration using the **Create Configuration** option.

5.4.7.1 Make Unconfigured Good

Follow these steps to change the status of JBOD drives to Unconfigured Good.

1. Highlight **Make Unconfigured Good** on the **Configuration Management** menu and press Enter.
   The following dialog appears, which lists information about the JBOD drives currently connected to the controller.

Figure 92 Make Unconfigured Good

   Scroll down, if necessary, to view other drives listed in the dialog.
2. Highlight each JBOD drive you want to make Unconfigured Good and press the spacebar to select it.

   ATTENTION If one or more JBOD drives that you have selected have an operating system (OS) or a file system on them, a warning message appears indicating that the listed JBOD drives have an operating system or a file system and any data on them would be lost if you proceed with the conversion. If you want to proceed, highlight **Confirm** and press the spacebar, then highlight **Yes** and press Enter. Otherwise, highlight **No** and press Enter to return to the previous screen and unselect those JBOD drives that have an OS or a file system installed on them.
3. Highlight **OK** (at the bottom of the JBOD drive list) and press Enter to convert the JBOD drives to Unconfigured Good status.

5.4.7.2 Make JBOD

Perform these steps to change the status of Unconfigured Good drives to JBOD.

1. Highlight **Make JBOD** on the **Configuration Management** menu and press Enter.
   The **Make JBOD** dialog appears listing the Unconfigured Good drives currently connected to the controller.
2. Highlight each drive you want to convert to JBOD status and press the spacebar to select it.
3. Highlight **OK** and press Enter to convert the Unconfigured Good drives to JBOD status.

5.4.7.3 Enabling Security on JBOD

If you have SED-enabled JBOD drive that meets the perquisites mentioned in Managing Configurations, you can enable security on it. Follow these steps to enable the security on a JBOD.

ATTENTION All of the data on the drive is lost when you enable security on it. Therefore, back up any data that you want to keep.

1. Highlight **Enable Security on JBOD** on the **Configuration Management** menu and press Enter.

   The **Enable Security on JBOD** dialog appears listing the SED-enabled JBOD drives currently connected to the controller.

Figure 93 Enable Security on JBOD

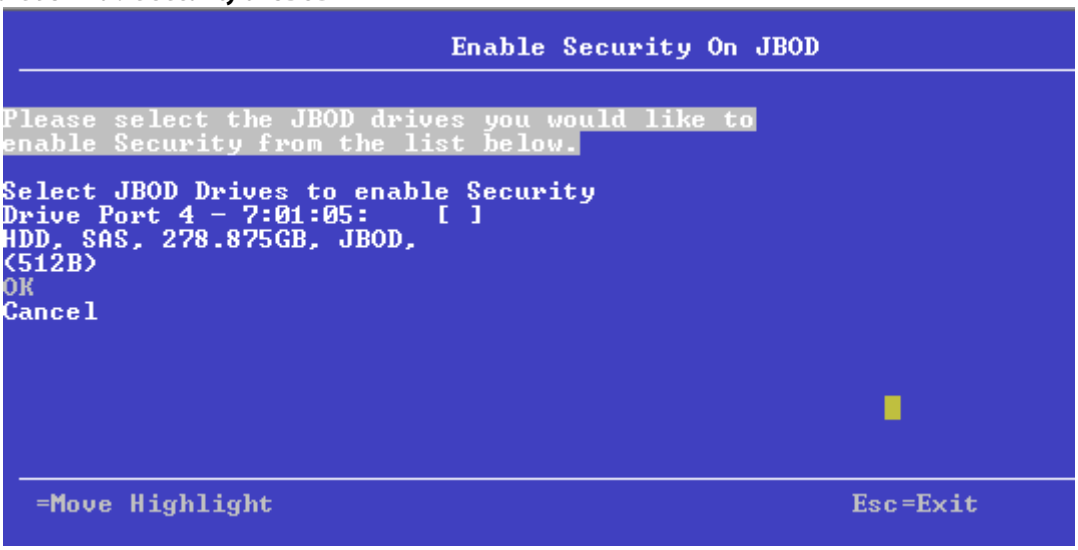

2. Highlight each JBOD drive to enable security on it and press the spacebar to select it.
3. Highlight **OK** and press Enter to enable security on the JBOD drive.

   A message appears stating that the existing data in the drive would be lost if you proceed and prompting for your confirmation.
4. Highlight **Confirm** and press the spacebar, then highlight **Yes** and press Enter.

   A success message appears.
5. Highlight **OK** and press Enter.

   The HII Utility enables security on the JBOD drive and returns you to the Configuration Management menu.

5.4.8 Managing Foreign Configurations

The following dialog appears when you select **Manage Foreign Configuration** from the **Dashboard View** or the **Configuration Management** menu.

Figure 94 Manage Foreign Configuration

A *foreign configuration* is a virtual disk that was created on another controller, and whose member drives have been moved to this controller.

The following sections explain how to preview and import a foreign configuration and how to clear a foreign configuration.

5.4.8.1 Previewing and Importing a Foreign Configuration

You can preview a foreign configuration before importing it or clearing it. Importing a foreign configuration means activating an inactive virtual drive that you physically transferred to the controller from another system. You might be unable to import a foreign configuration if any of the following conditions exist:

- The volume state is not INACTIVE.
- The volume state is either FAILED or MISSING.
- The volume uses incompatible Gen1 metadata.
- The maximum number of two RAID volumes already exist on this controller.
- The maximum number of supported physical drives are already in use in active volumes on this controller. Global hot spares also count because they must be activated along with other drives in the foreign volume.

HII displays the following message if you try to import a foreign configuration that is locked, and if drive security is disabled on the controller.

Figure 95 Enter Security Key for Locked Drives

To successfully import the foreign configuration, follow the directions in the message.

Perform these steps to preview and import a foreign configuration.

1. Highlight **Preview Foreign Configuration** on the **Manage Foreign Configuration** menu and press Enter. The following dialog appears, listing information about the physical drives in the foreign configuration.

Figure 96 Preview Configuration Window 1

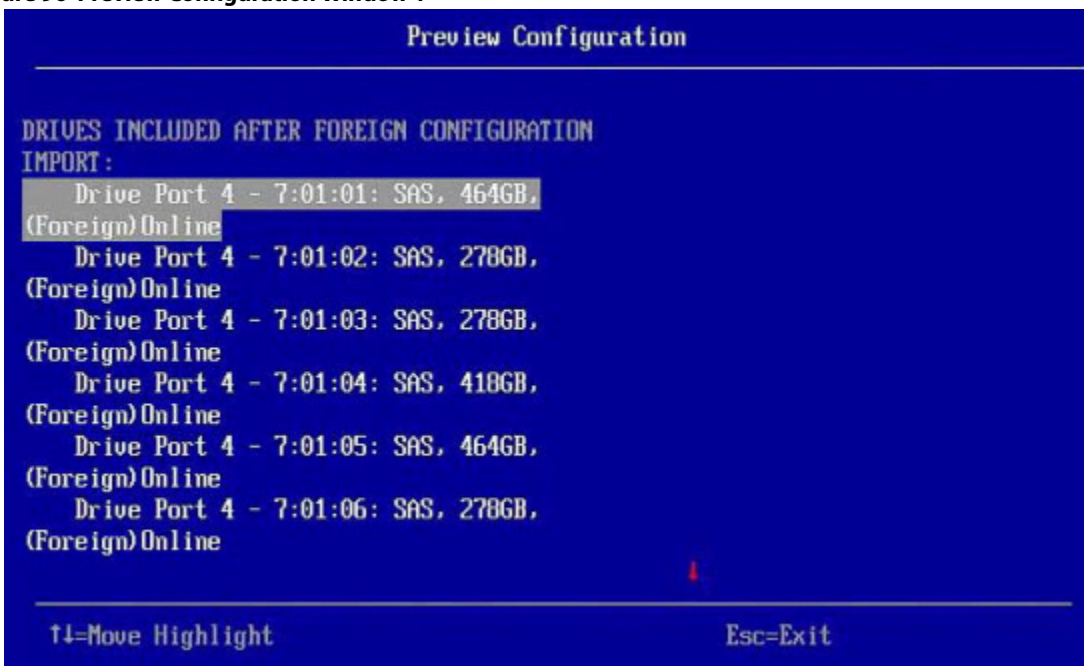

2. Scroll down, if needed, to view more information about the drives in the foreign configuration, as shown in the following figure.

Figure 97 Preview Configuration Window 2

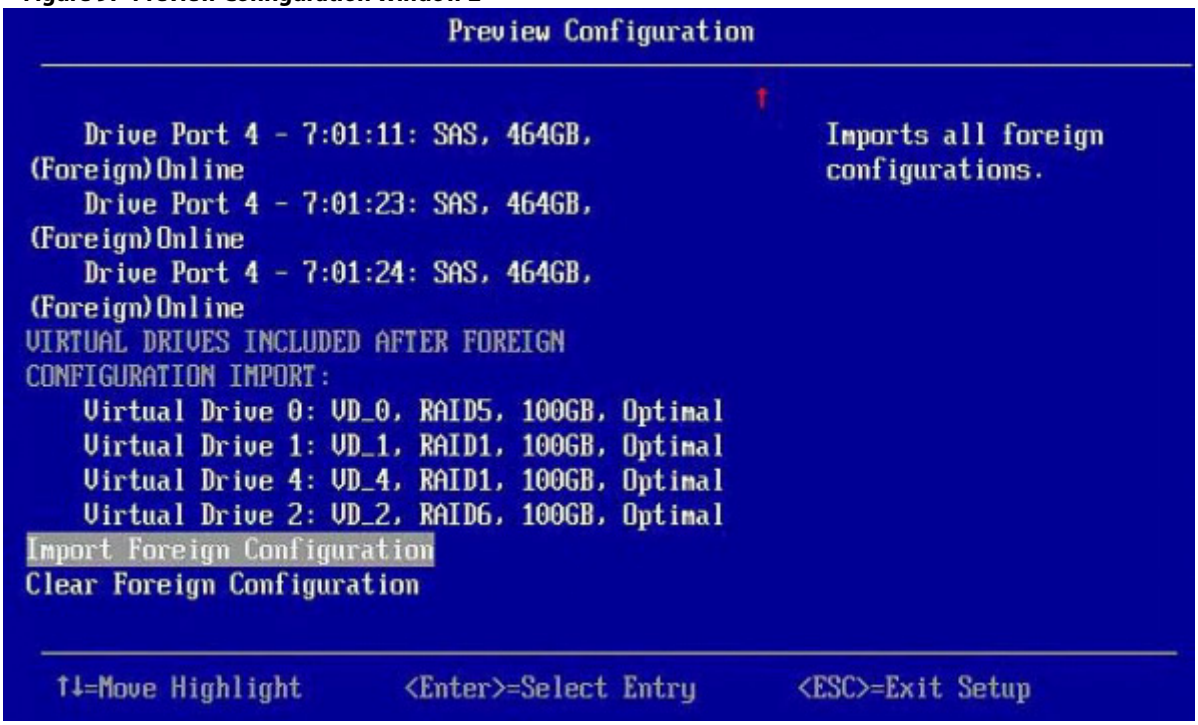

3. Review the information listed on the window.
4. Highlight **Import Foreign Configuration** and press Enter.
   A warning message appears that indicates the foreign configuration from the physical drives will merge with the existing configuration.
5. To confirm the import, highlight **Confirm** and press the spacebar.
6. Highlight **Yes** and press Enter.
   The foreign configuration is imported.

5.4.8.2 Clearing a Foreign Configuration

Perform these steps to clear a foreign configuration.

1. Highlight **Clear Foreign Configuration** on the **Manage Foreign Configuration** menu and press Enter.
   A warning message appears that indicates all of the foreign VDs will be deleted.
2. To confirm clearing the foreign configuration, highlight **Confirm** and press the spacebar.
3. Highlight **Yes** and press Enter.
   The foreign configuration is deleted.

NOTE You can also delete (clear) a foreign configuration after you preview the configuration.

5.5 Managing Controllers

When you select **Controller Management** from the **Main Menu** or from the **View Server Profile**, the **Controller Management** dialog appears, as shown in the following figure.

The top-level **Controller Management** dialog lists some actions that you can perform on the controller.

- To view additional controller management properties, in the **Basic Properties** section, highlight **Advanced Controller Management** and press Enter.

  For more information, see Viewing Advanced Controller Management Options.
- To view additional controller properties, in the **Basic Properties** section, highlight **Advanced Controller Properties**.

  For more information, see Viewing Advanced Controller Properties.

Figure 98 Controller Management

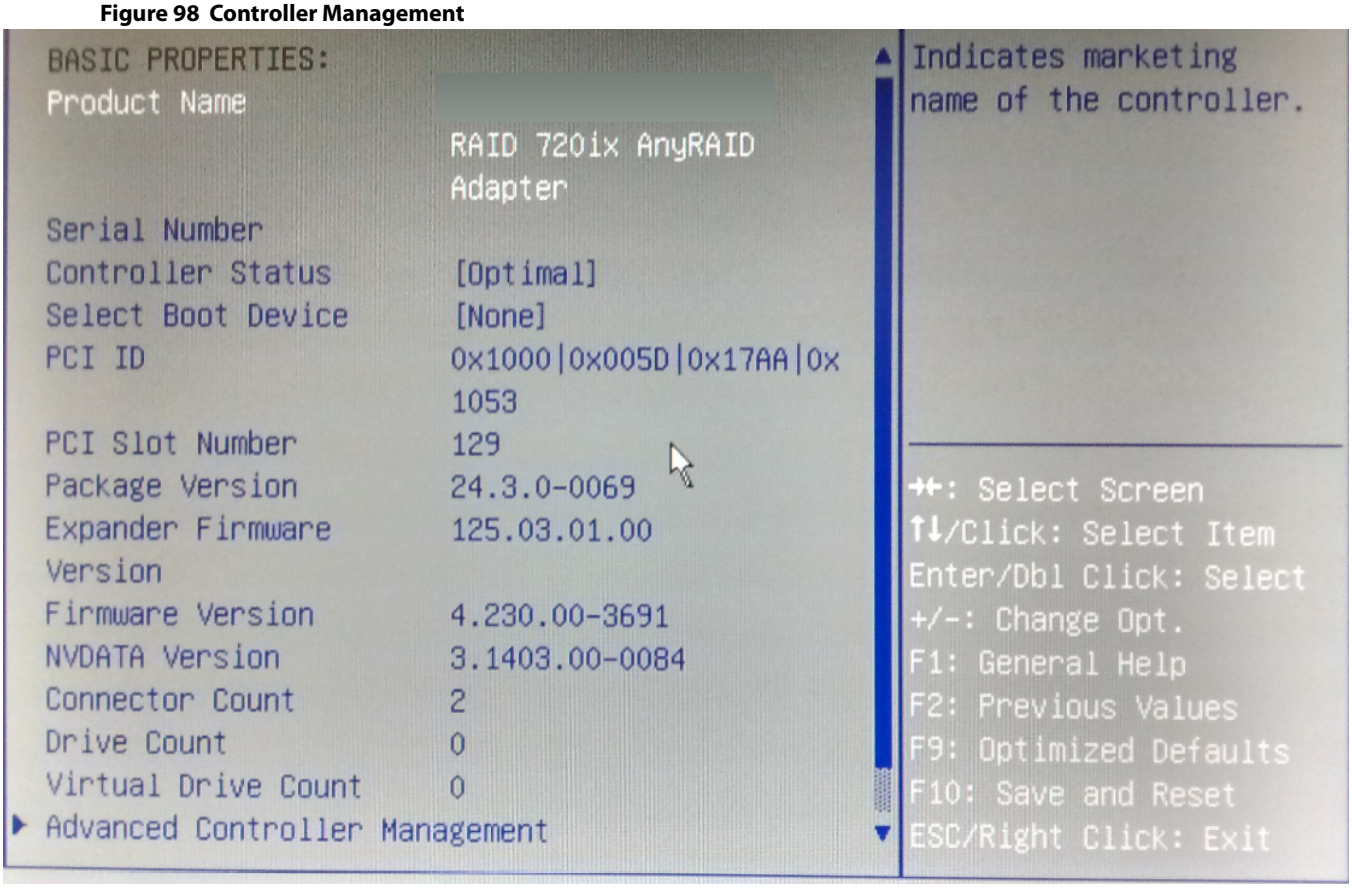

The **Controller Management** dialog lists the following basic controller properties.

Table 27 Basic Controller Properties

Property	Description
Product Name	The marketing name of the controller.
Serial Number	The serial number of the controller.
Controller Status	The cumulative status of virtual drives and physical drives connected to the controller, plus the backup battery, the enclosure and the NVDATA. The status is one of the following: ■ **Optimal**, if all components are operating normally. ■ **Needs Attention**, if any component needs attention. ■ **Safe Mode**, if the controller encountered critical errors. Most features are disabled and the controller requires user attention.
Select Boot Device	This field selects the primary boot device. **NOTE** This property is applicable for legacy BIOS.
PCI ID	The PCI ID of the controller.

Table 27 Basic Controller Properties (Continued)

Property	Description
PCI Slot Number	The slot ID number of the PCI slot where the controller is installed.
Package Version	The version number of the package.
Expander Firmware Version	This field shows the firmware version of the expander that is connected to the controller. **NOTE** This field only appears when an expander is connected to the controller.?
Firmware Version	The version number of the controller firmware.
NVDATA Version	The version number of the controller NVDATA.
Connector Count	Number of host data ports, connectors, or both currently in use on this controller.
Drive Count	Number of physical drives attached to this controller.
Virtual Drive Count	Number of virtual drives defined on this controller

5.5.1 Viewing Advanced Controller Management Options

The **Advanced Controller Management** dialog lists all the controller management properties and also includes options for performing various actions on the controller.

Figure 99 Advanced Controller Management

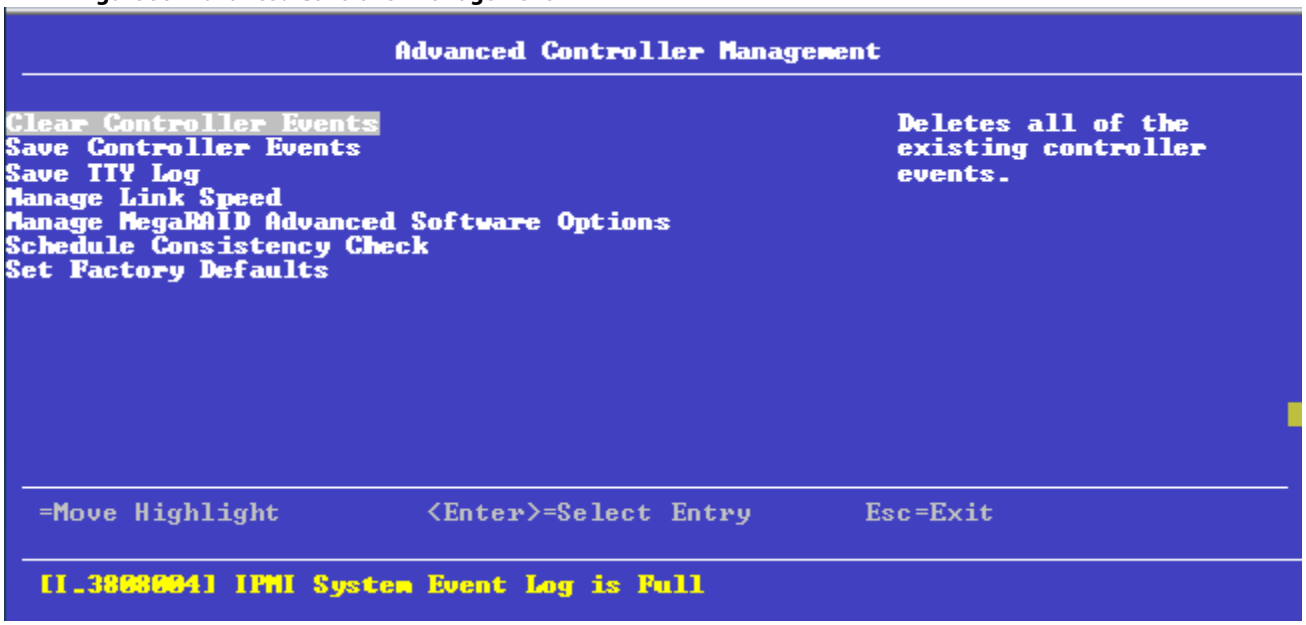

The following table describes all of the entries on the **Advanced Controller Management** dialog, including the ones that are not visible.

Table 28 Controller Management Options

Property	Description
Clear Controller Events	Clears entries from the log.
Save Controller Events	Saves the controller log entries to a file.
Save TTY Log	Saves a copy of the firmware's terminal log entries for the controller.
Enable Drive Security	Enables drive security to protect the data on your system from unauthorized access or use.

Table 28 Controller Management Options (Continued)

Property	Description
Disable Drive Security	Disables drive security.
Change Security Key	Changes the security key or switch between drive security modes on the controller.
Manage Link Speed	Enables 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. For more information, see Managing and Changing Link Speeds.
Manage MegaRAID Advanced Software Options	DIsplays the activated MegaRAID Advanced Software Options on the controller and lets you configure these options to use the advanced features in the controller. You need to activate the activation key to use the advanced features. **NOTE** The MegaRAID Advanced Software Options are displayed only if the controller supports MegaRAID software licensing.
Schedule Consistency Check	Schedules a consistency check operation to verify and correct the mirror and parity data for fault tolerant virtual drives.
Set Factory Defaults	Resets the controller to its factory settings.

5.5.2 Viewing Advanced Controller Properties

The **Advanced Controller Properties** dialog lists all the controller properties and also includes options for performing various actions on the controller.

The top-level of the **Advanced Controller Management** dialog lists some actions that you can perform on the controller.

- To view and modify the controller cache, highlight **Cache and Memory** and press Enter.
  For more information, see Setting Cache and Memory Properties.
- To view and set patrol read properties, highlight **Patrol Read**, press Enter.
  For more information, see Running a Patrol Read.
- To view and modify physical drive power settings, highlight **Power Settings** and press Enter.
  For more information, see Changing Power Save Settings.
- To view and modify properties related to replacing a drive, an emergency spare, or a hot spare, highlight **Spare** and press Enter.
  For more information, see Setting Emergency Spare Properties.
- To modify the rebuild rate and other task rates for a controller, highlight **Task Rates**.
  For more information, see Changing Task Rates.

Figure 100 Advanced Controller Properties

This dialog lists various properties, all of them cannot b e shown in one dialog. Scroll down to view all of the options.

NOTE The red arrow appears when there is too much information to display in one dialog. The amount of information that can be displayed in one dialog depends on the capabilities of the HII browser.

Many of the entries in this dialog are view-only, but some are selectable and configurable. Perform these steps to change any user-configurable option on this dialog.

1. Move the highlight to the value for any option and press Enter.

   A pop-up menu of the available options appears.

2. Highlight the value you want and press Enter. For options, such as **SMART Polling** that require a number, use the + and – keys on the keypad to increase or decrease the number, and press Enter.

   NOTE Some systems permit you to enter numeric values directly, without using the + and – keys.

3. When you finish changing the controller properties, scrolling up and down on the menu as needed, move the highlight to **Apply Changes** and press Enter.

   The changes to the controller properties are applied, and a success message appears.

The following table describes all the controller properties listed in the **Advanced Controller Properties** dialog, including the ones that are not visible.

Table 29 Advanced Controller Properties

Property	Description
Alarm Control	Enables or disables the controller alarm.
Auto Import Foreign Configuration	Enables or disables the automatic import of foreign configurations without any user intervention.
Boot Mode	Specifies the option to handle errors that the firmware might encounter during the boot process. The errors might require you to take action or to acknowledge the error and permit the boot process to continue. The options are *Stop on error*, *Pause on error*, *Ignore errors*, and *Safe mode*.
Controller BIOS	Enables or disables the controller BIOS. The controller BIOS should be enabled if the boot device is connected to the selected RAID controller.

Table 29 Advanced Controller Properties (Continued)

Property	Description
Controller Temperature	Indicates the temperature of the controller.
ROC Temperature	Current temperature of the RAID-on-a-chip (ROC) on the controller, in degrees Celsius.
Shield State Supported	Indicates whether the controller supports shield state.
Drive Security	Indicates the drive security (encryption) feature status on the controller.
T10-PI	Indicates the status of the data protection feature on the controller.
Maintain Drive Fail History	Enables or disables the option to track bad physical drives through a reboot.
SMART Polling	Determines the interval, in seconds, at which the controller polls for drives reporting a Predictive Drive Failure. The default is 300 seconds. To change the value, use the + and – keys on the keypad. **NOTE** Some systems let you edit the numeric value directly, without using the + and – keys.
Stop Consistency Check on Error	Enables or disables the option of stopping a consistency check operation on a redundant virtual drive if a data inconsistency is detected.
JBOD Mode	Enables or disables the JBOD mode. **NOTE** When the JBOD mode is enabled, the drive comes up as a JBOD; otherwise, it comes up as an Unconfigured Good drive. **NOTE** When the JBOD mode is disabled, if one or more selected JBODs contain an operating system or a file system, a warning message appears indicating that the listed JBOD drives have an operating system or a file system and any data on them would be lost if you proceed. If you want to disable the JBOD mode, highlight **Confirm** and press the spacebar, then highlight **Yes** and press Enter. Else, highlight **No**.
Write Verify	Enables or disables the write verify feature during controller cache flush. This feature verifies if the data was written correctly to the cache before flushing the cache.
Large IO Support	Enables or disables the large I/O support feature. By default, large I/O support is disabled. A reboot is required if this property is changed. When this property is changed,`The controller property change has been performed successfully. Reboot the machine for the change to take effect` message is displayed.

5.5.3 Managing MegaRAID Advanced Software Options

The **Manage MegaRAID Advanced Software Options** dialog lists all the activated advance software options on the controller. You can configure the MegaRAID advanced software options to use the advanced software features.

Follow these steps to enable the activation key in order to use the advanced software features:

1. In the **Dashboard View** dialog or the **Advanced Controller Management** dialog, highlight **Manage MegaRAID Advanced Software Options** and press Enter.

   The **Manage MegaRAID Advanced Software Options** dialog appears, as shown in the following figure.

Figure 101 Manage MegaRAID Advanced Software Options

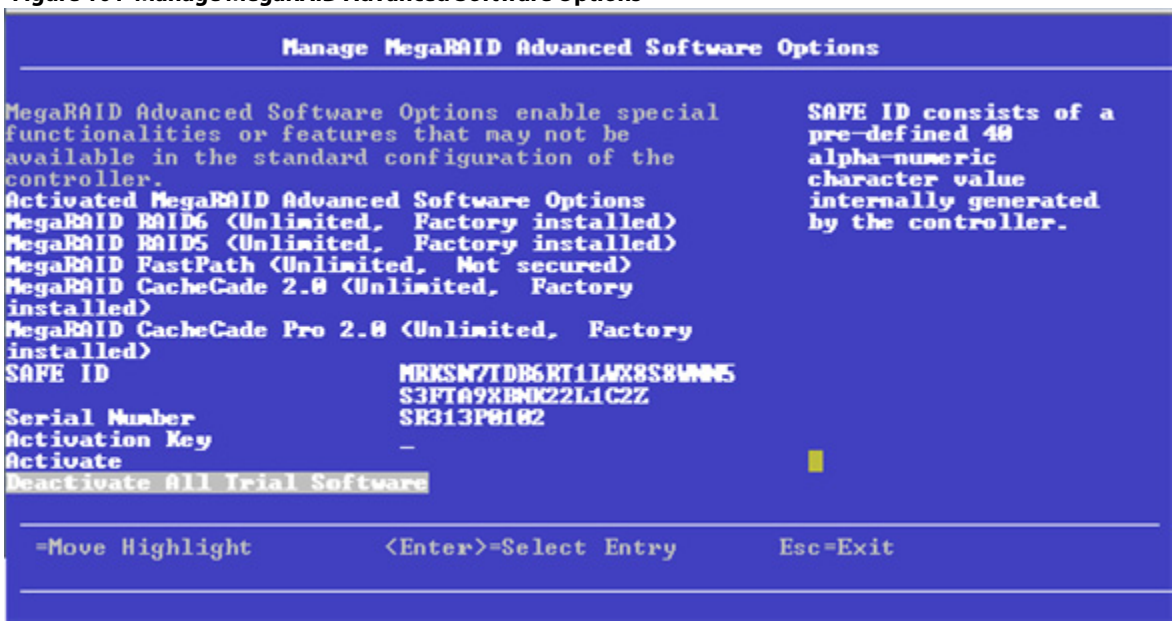

This dialog lists fields that cannot all be shown in one dialog. Scroll down to view all of the fields.

NOTE The red arrow appears when there is too much information to display in one dialog. The amount of information that can be displayed in one dialog depends on the capabilities of the HII browser.

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

2. Highlight **Activation Key** and press Enter. Enter the activation key and press Enter.
3. Click **Activate**.

   The activation key is activated. You can now use the advanced software features.

5.5.4 Scheduling a Consistency Check

The **Schedule Consistency Check** dialog appears when you select **Schedule Consistency Check** from the **Advanced Controller Management** menu.

Use this dialog to schedule consistency checks on the redundant virtual drives configured on the controller. The nonselectable entries in the **Consistency Check Start** fields indicate the date and time of the next scheduled consistency check.

Follow these steps to change the consistency check settings.

1. Highlight the **Consistency Check Frequency** field and press Enter.

   A selectable popup menu appears.

Figure 102 Scheduling a Consistency Check

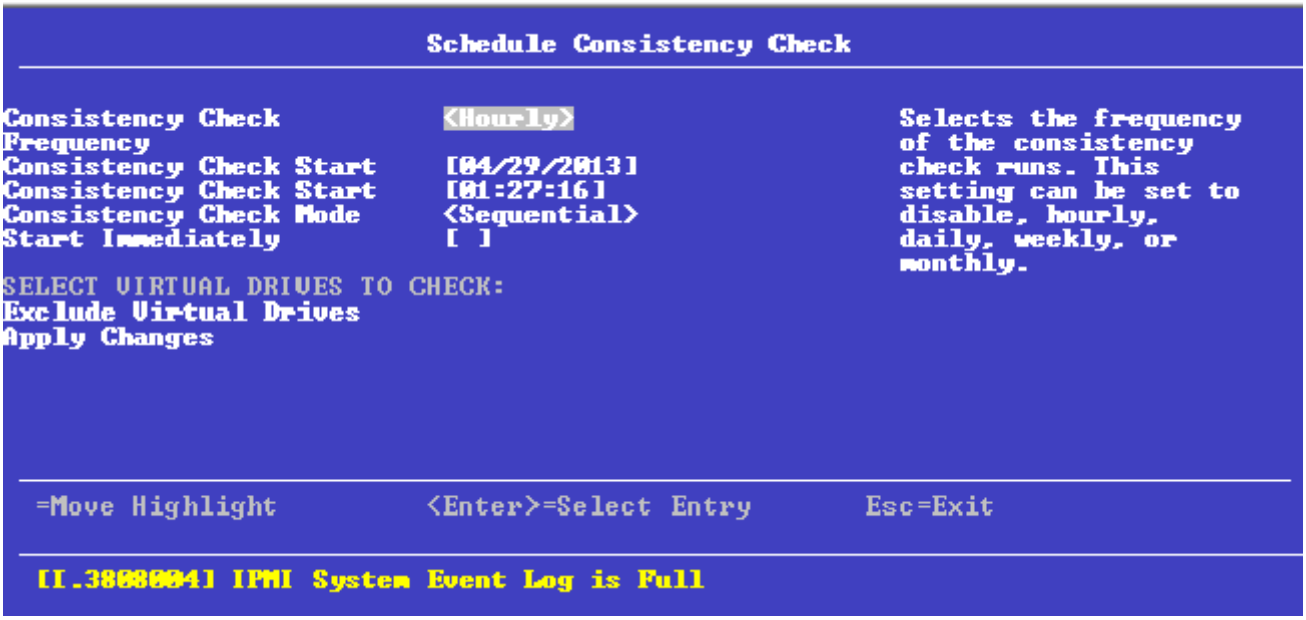

2. Select the desired interval at which to run consistency checks.

   The choices are **Hourly**, **Daily**, **Weekly**, or **Monthly**. You can also choose to disable consistency checks, which is not recommended because it reduces the level of protection for your system.
3. To change the mode of operation, highlight the **Consistency Check Mode** field and press Enter.

   A selectable pop-up menu appears.
4. Select **Concurrent** to run consistency checks concurrently on all virtual drives, or select **Sequential** to run consistency checks on one virtual drive at a time.
5. Check the **Start Immediately** checkbox to run consistency checks immediately on all virtual drives that are *not* excluded, not just on a single virtual drive.
6. (Optional) To exclude specified virtual drives from consistency checks, highlight the **Exclude Virtual Drives** field and press Enter.

   The **Exclude Virtual Drives** dialog appears, listing the virtual drives defined on this controller.

   You might want to exclude a virtual drive from a consistency check if, for example, you are running some operation on the drive and you do not want it to be interrupted by a consistency check.
7. To exclude a virtual drive from the consistency check, highlight the field to the right of the drive name and press the spacebar.

   An X in this field means the virtual drive does not undergo a consistency check.
8. Highlight the **Select Entry** field and press Enter.

   The program returns you to the **Schedule Consistency Check** dialog.
9. Highlight the **Select Entry** field on the **Schedule Consistency Check** dialog and press Enter.

   The consistency check changes are now registered.

5.5.5 Saving or Clearing Controller Events

The following window appears when you select **Save Controller Events** from the **Advanced Controller Management** menu.

NOTE An error message appears if the controller events log is empty.

Figure 103 Save Controller Events

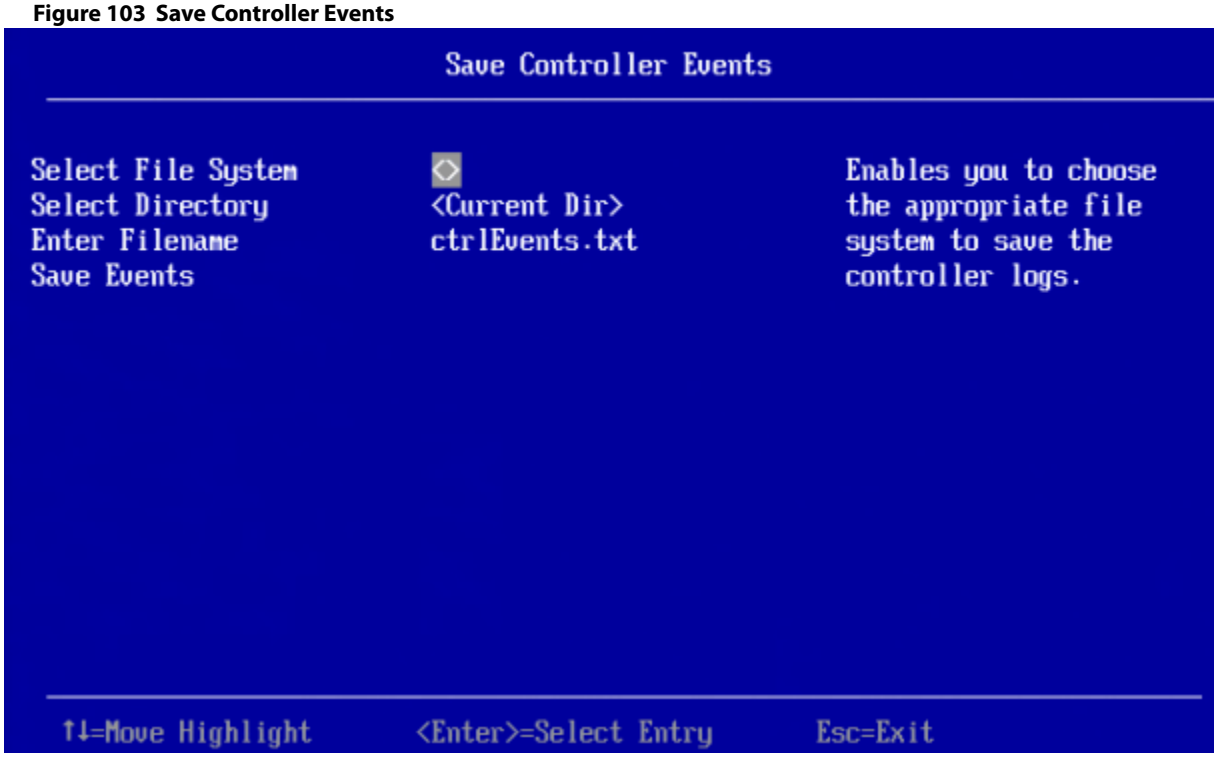

Perform these steps to save controller event log entries to a file.

1. To select a different file system from the one listed in the **Select File System** field, highlight the current file system name and press Enter.

   An error message appears if there is no file system.
2. Select a file system from the pop-up menu and press Enter.
3. To save the controller events file to a different directory from the one listed in the **Select Directory** field, highlight the current directory name and press Enter.
4. Select a directory name from the pop-up menu and press Enter.
5. To enter a different name for the controller event log file, highlight the current file name and press Enter.
6. Type the new file name in the pop-up dialog and press Enter.
7. Highlight **Save Events**, and press Enter to save the event log entries to the file.

To clear controller events, highlight **Clear Controller Events** in the **Advanced Controller Management** dialog. When the confirmation message appears, highlight **OK** and press Enter.

5.5.6 Enabling or Disabling Drive Security

The following dialog appears when you select **Enable Drive Security** from the **Advanced Controller Management** menu.

Figure 104 Enable Drive Security (Choose Drive Security Mode)

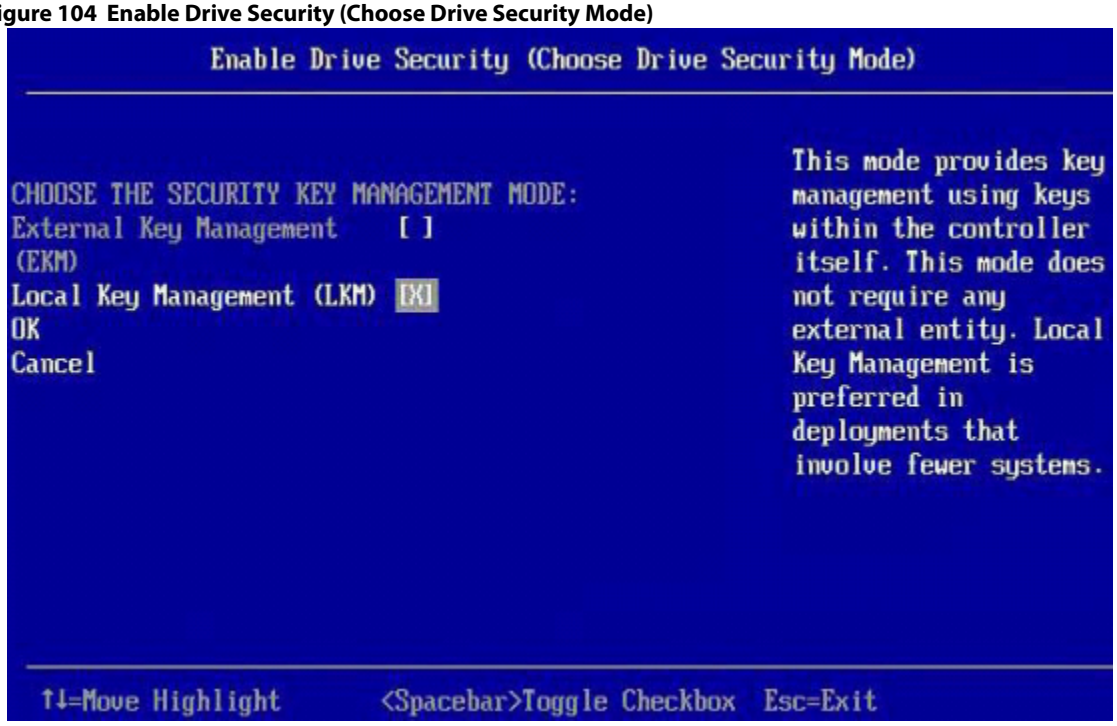

Enable drive security to protect the data on your system from unauthorized access or use. Local Key Management (LKM) is the method that the HII utility provides to manage drive security. LKM uses security keys within the controller and does not require any external entity to implement. Therefore, it is the preferred security mode for configurations that involve a smaller number of computer systems.

Follow these steps to enable LKM security on your configuration.

1. Highlight the **Local Key Management (LKM)** field and, if required, press the spacebar to enter an X in this field.
2. Highlight **OK** and press Enter.

   The following dialog appears.

Figure 105 Enable Drive Security

The highlighted field is the security key identifier, which appears whenever you need to enter the security key. If you have more than one security key, the identifier helps you determine which security key to enter.

3. To change the security key identifier, press Enter and enter the new identifier in the popup window.

4. To request the controller to suggest a drive security key, highlight **Suggest Security Key** and press Enter.
5. To enter your own security key, highlight the **Security Key** field, press Enter, and type the security key.

   The **Security Key** field is case-sensitive. The security key must be between eight and thirty-two characters and must contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (for example, > @ +).
6. After entering the security key, highlight **Confirm** and press Enter. Enter the security key again to confirm it.

   The security key must match exactly the characters you entered in the **Security Key** field.
7. If you do not want the controller to require a password at boot time, deselect the **Pause for Password at Boot** option by highlighting it and pressing the spacebar.

   This option is selected by default.
8. To enforce strong password restrictions, highlight **Enforce Strong Password Security** and press the spacebar.

   A strong password must be between eight and thirty-two characters and must contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (for example, > @ +).
9. Highlight the **Password** field, press Enter, and type the boot time password.
10. Highlight **Confirm** and re-enter the password.

    The password must match exactly the characters you entered in the **Password** field.
11. Record the drive security information and store it in a safe place.
12. Highlight the **I Recorded The Security Settings**... field and press the spacebar to select it.
13. Highlight **Enable Drive Security** and press Enter.
14. When the popup window appears, confirm that you want to enable drive security and select **Yes**.

    Drive security is enabled for the drives connected to this controller.

Follow these steps to disable LKM drive security:

1. Select **Disable Drive Security** from the **Advanced Controller Management** menu.

   The following warning appears.

Figure 106 Disable Drive Security Warning

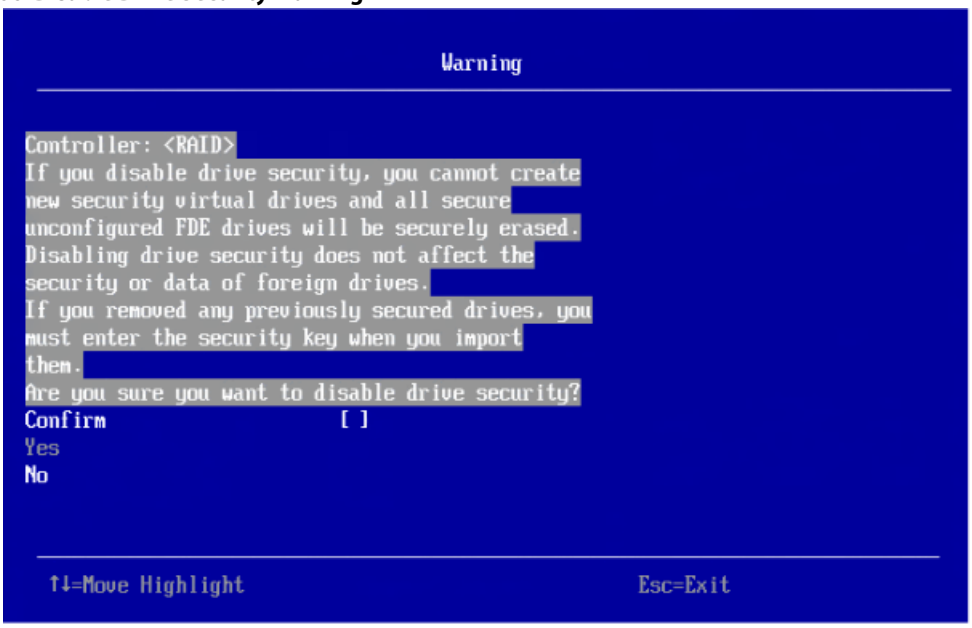

2. Read the warning and be sure you understand what will happen if you disable the drive security.
3. Highlight **Confirm** and press the spacebar to select it.

4. Highlight **Yes** and press Enter.

   Drive security is disabled.

5.5.7 Changing a Security Key

The **Change Security Key** dialog appears when you select **Change Security Key** from the **Advanced Controller Management** menu.

Perform these steps to change the security key settings.

1. Highlight **OK** and press Enter.

   The following dialog appears.

Figure 107 Change Security Key

   By default, the same security key identifier is retained.

2. To change the security key identifier, press the spacebar to deselect **Use the Existing Security Key Identifier**.
3. Highlight the **Enter a New Security Key Identifier** field, press Enter, and enter the new security key identifier in the popup window.
4. Highlight the **Enter Existing Security Key** field and press Enter.

   You are required to enter the security key to prevent unauthorized changes to the security settings.
5. Type the current security key in the popup window and press Enter.
6. Highlight **Suggest Security Key** and press Enter to have the system create a new security key.
7. To enter your own new security key, highlight the **Security Key** field, press Enter, and type the new security key.

   The **Security Key** field is case-sensitive. The security key must be between eight and thirty-two characters and must contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (for example, > @ +).
8. After entering the new security key, highlight **Confirm** and press Enter. Enter the security key again to confirm it.

   The security key must match exactly the characters you entered in the **Security Key** field.
9. If you do not want the controller to require a password at boot time, deselect the **Pause for Password at Boot** option by highlighting it and pressing the spacebar.

   This option is selected by default.
10. To enforce strong password restrictions, highlight **Enforce Strong Password Security** and press the spacebar.

    A strong password must be between eight and thirty-two characters and must contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (for example, > @ +).

11. Highlight the **Password** field, press Enter, and type the new boot time password.
12. Highlight **Confirm** and reenter the new password.
    The password must match exactly the characters you entered in the **Password** field.
13. Record the drive security information and store it in a safe place.
14. Highlight the **I Recorded The Security Settings**... field and press the spacebar to select it.
15. Highlight **Change Security Key** and press Enter.
16. When the popup window appears, confirm that you want to change the security settings and select **Yes**.
    The security changes are entered for the drives connected to this controller.

5.5.8 Saving the TTY Log

The following dialog appears when you select **Save TTY Log** from the **Advanced Controller Management** menu.

Figure 108 Save TTY Log

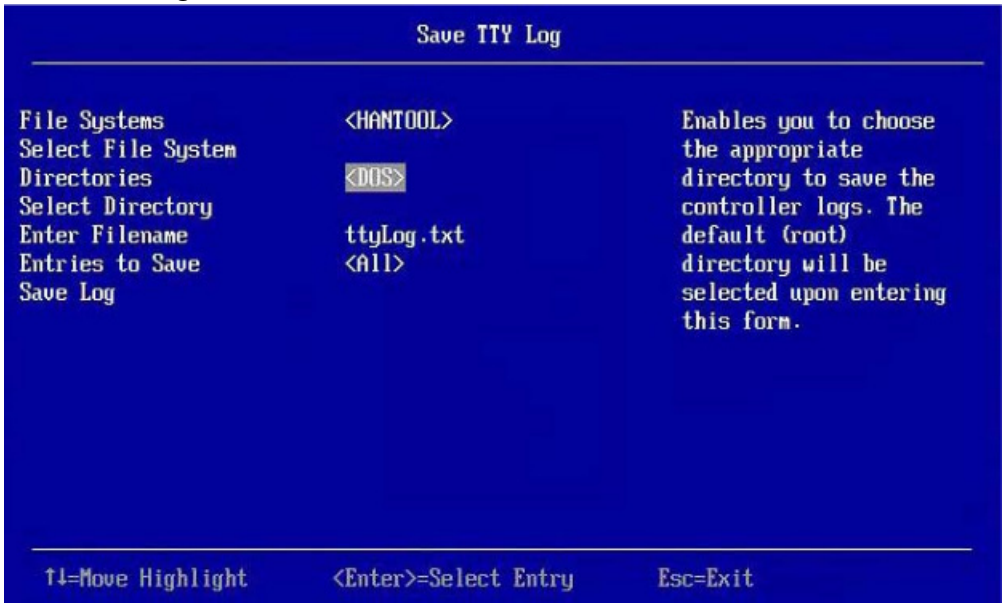

Follow these steps to save the TTY log entries to a file.

1. To select a different file system from the one listed in the **File Systems** field, highlight the current file system name, and press Enter.
   An error message appears if there is no file system.
2. Select a file system from the popup menu, and press Enter.
3. Highlight **Select File System** and press Enter.
4. To save the TTY log events file to a different directory from the one listed in the **Directories** field, highlight the current directory name, and press Enter.
5. Select a directory name from the pop-up menu, and press Enter.
6. Highlight **Select Directory**, and press Enter.
7. To enter a different name for the TTY log file, highlight the current file name, and press Enter.
8. Type the new file name in the pop-up window, and press Enter.
9. To select how many TTY log entries to save, highlight the **Entries to Save** field, and press Enter.

10. Select an option from the popup menu, and press Enter.
    Your choices are **2 KB**, **4 KB**, **8 KB**, **16 KB**, or **All**.
11. Highlight **Save Log**, and press Enter to save the log entries to the file.

5.5.9 Managing and Changing Link Speeds

The Manage Link Speed feature lets you change the link speed between the controller and an expander or between the controller and a drive that is directly connected to the controller. The following dialog appears when you select **Manage Link Speed** on the **Advanced Controller Management** dialog. The default settings for all phys is **Auto**.

Figure 109 Manage Link Speed

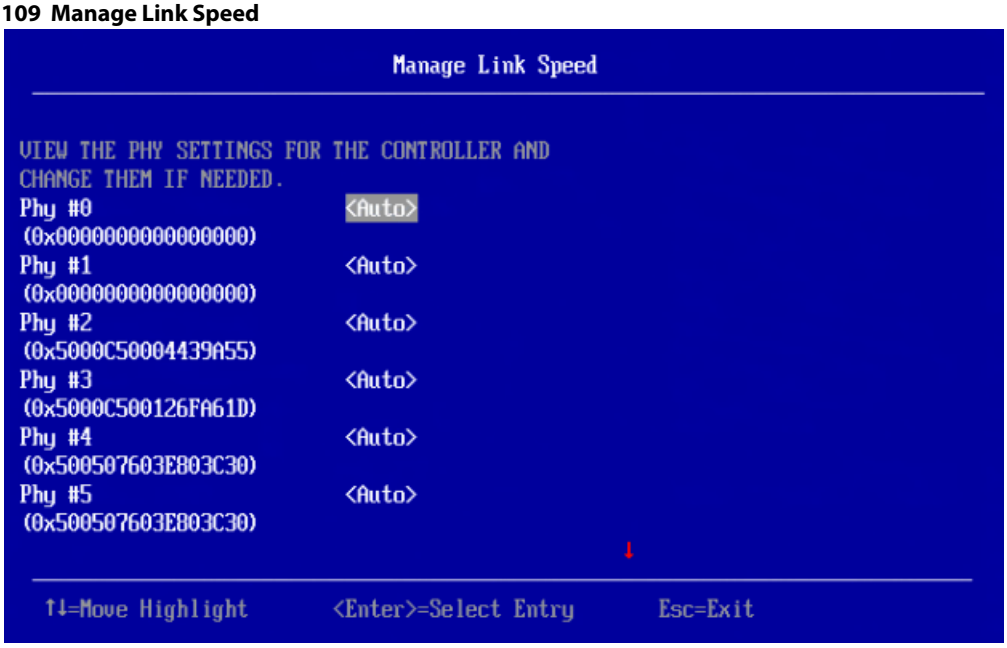

Follow these steps to change the link speed for one or more phys:

1. Highlight the field to the right of the phy number and press Enter.
2. Select an option from the pop-up menu.
   The link speed values are Auto,1.5Gb/s, 3Gb/s, or 6Gb/s.
3. Scroll to the bottom of the phy list, highlight **OK**, and press Enter.

5.5.10 Setting Cache and Memory Properties

The following dialog appears when you select **Cache and Memory** from the **Advanced Controller Properties** dialog.

Figure 110 Cache and Memory

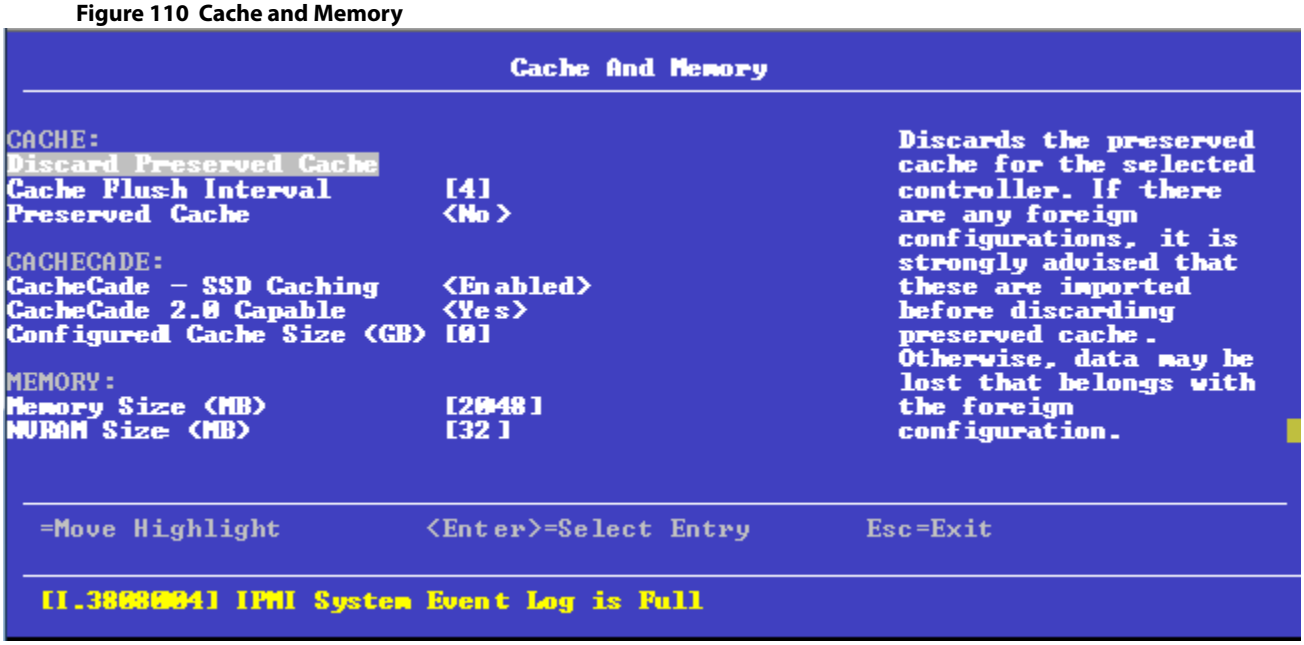

Follow these steps to set cache and memory properties:

1. To discard the preserved cache for the controller, highlight **Discard Preserved Cache** and press Enter.

   NOTE If any foreign configurations exist, import them before discarding the preserved cache. Otherwise, you might lose data that belongs with the foreign configuration.

2. To change the interval, in seconds, at which the contents of the onboard data cache are flushed, highlight **Cache Flush Interval** and press Enter. Specify a numeric value and press Enter.
3. If you want the controller to preserve cache because of missing or offline virtual drives (the cache is preserved until the virtual drive is imported or the cache is discarded), highlight **Preserved Cache**, and press Enter. Select either **Yes** or **No** and press Enter.
4. Highlight **Apply Changes** and press Enter.

   The new settings are saved in the controller properties.

5.5.11 Running a Patrol Read

The following dialog appears when you select **Patrol Read** from the **Advanced Controller Properties** dialog.

Figure 111 Patrol Read

A patrol read operation scans and resolves potential problems on configured physical drives.

You can set the patrol read properties and start the patrol read operation, or you can start the patrol read without changing the properties:

Follow these steps to set patrol read properties:

NOTE You can only view the properties/options supported by your controller.

1. To select a mode for the patrol read operation, highlight **Mode** and press Enter. Select any of the following modes and press Enter.
   - **Auto**: Patrol read runs continuously on the controller based on a schedule. You do not need to start it manually.
   - **Manual**: Patrol read can be started or stopped manually.
   - **Disabled**: Patrol read does not run.
2. To specify a rate for the percentage of system resources dedicated to perform a patrol read operation on configured drives, highlight **Rate**, specify a rate as a numeric value and press Enter.

   100 is the maximum numeric value that you can enter as the rate.
3. To select a patrol read setting for unconfigured space, highlight **Setting for Unconfigured Space**, and press Enter. Select either **Enabled** or **Disabled** and press Enter.
4. Highlight **Apply Changes** and press Enter.

   The new settings are saved in the controller properties.

To start a patrol read without changing the patrol read properties, follow these steps:

1. Highlight **Start** in the **Patrol Read** dialog and press Enter.
2. A message box appears stating that the operation has been successful. Click **OK** to return to the **Patrol Read** dialog.

   Suspend and **Stop** are now active.

5.5.12 Changing Power Save Settings

The following dialog appears when you select **Power Save Settings** from the **Advanced Controller Properties** dialog.

Figure 112 Power Save Settings

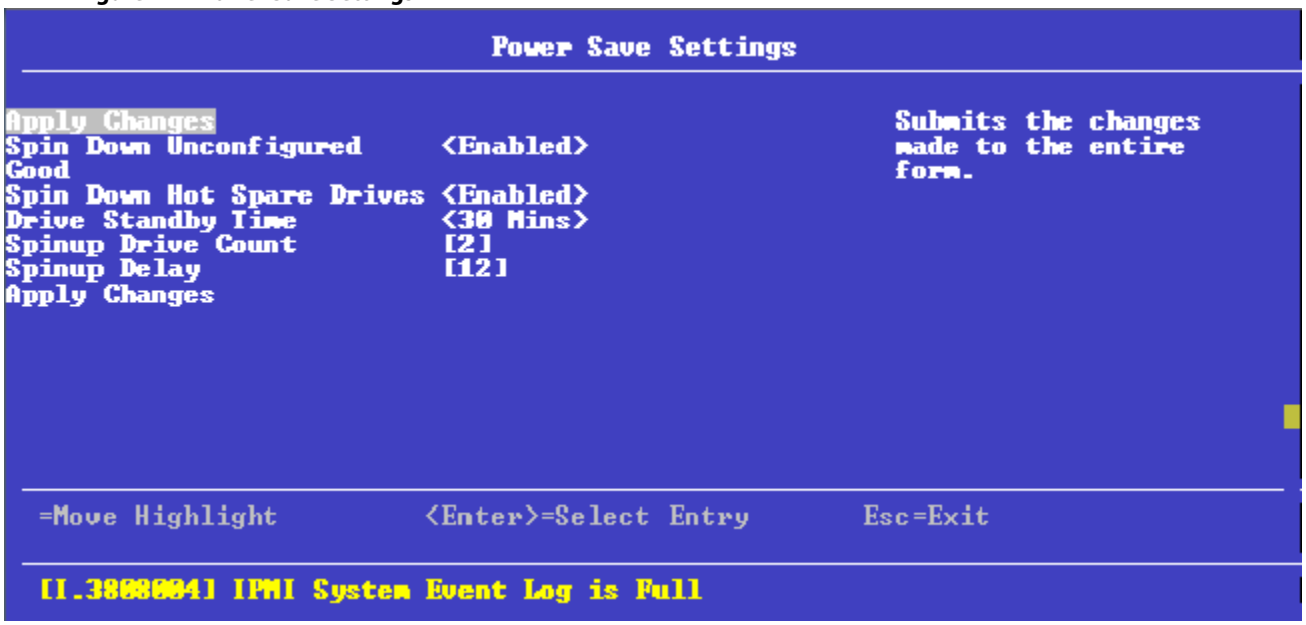

The above dialog lets you choose if you want unconfigured drives, hot spares, and configured drives to enter the power-save mode. When the unconfigured drives, hot spares, and configured drives are in power-save mode, they can be spun down.

Follow these steps to change the power-save settings:

NOTE You can only view the properties/options supported by your controller.

1. To enable or disable spinning down of unconfigured good drives, highlight **Spin Down Unconfigured Good** and press Enter. Select **Enable** or **Disable** and press Enter.
2. To enable or disable spinning down of hot spares, highlight **Spin Down Hot Spare Drives** and press Enter. Select **Enable** or **Disable** and press Enter.
3. To enable or disable spinning down of configured drives, highlight **Spin Down Configured Drives** and press Enter. Select **Enabled** or **Disabled** and press Enter.
4. To specify a drive's idle time, after which the drive goes into the power save mode, highlight **Drive Standby Time** and press Enter. Specify the time duration and press Enter.

   The drive standby time can be 30 minutes, 1 hour, 1.30 hours, or 2 hours through 24 hours.
5. To select the desired power-save mode, highlight **Power Save Mode** and press Enter. Select a mode (**None**, **Auto**, **Max**, and **Max without Cache**) and press Enter.
6. To specify the maximum number of drives that spin up simultaneously, highlight **Spinup Drive Count** and press Enter. Specify a numeric value and press Enter.
7. To control the interval (in seconds) between spin up of drives connected to the controller, highlight **Spinup Delay** and press Enter. Specify the time in seconds and press Enter.

   The delay prevents a drain on the system's power supply that would occur if all drives spun up at the same time.

8. If you do not want to schedule the drive active time, highlight **Do Not Schedule Drive Active Time** and press Enter.
9. To specify the Quality of Service window start time, highlight **Qos Window Start Time** and press Enter. Specify a start time and press Enter.
10. To specify the Quality of Service window end time, highlight **Qos Window End Time** and press Enter. Specify a end time and press Enter.
11. Highlight **Apply Changes** and press Enter.

    The new settings are saved in the controller properties.

5.5.13 Setting Emergency Spare Properties

The following dialog appears when you select **Spare** from the **Advanced Controller Properties** dialog.

Figure 113 Spare

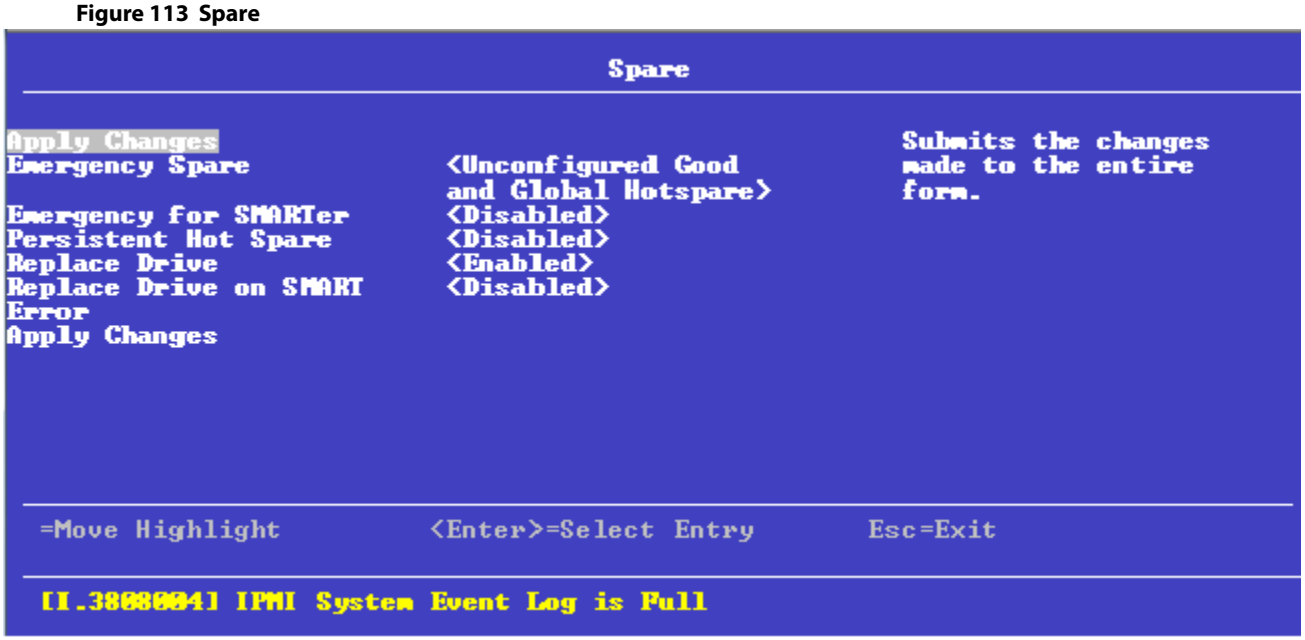

When a drive within a redundant virtual drive fails or is removed, the MegaRAID firmware automatically rebuilds the redundancy of the virtual drive by providing a emergency spare drive, even if no commissionable dedicated drive or global hot spare drive is present.

Follow these steps to set emergency spare properties:

1. To specify whether it is acceptable to commission otherwise incompatible global hot spare drive and/or unconfigured good drives as emergency hot spare drives, highlight **Emergency Spare** and press Enter. Select any of the following modes and press Enter.
   - **Global Hotspare**
   - **Unconfigured Good**
   - **Unconfigured Good and Global Hotspare**
   - **None**
2. To specify whether it is acceptable to commission emergency hot spare drives for PFA events, highlight **Emergency for SMARTer** and press Enter. Select an option (**Enabled** or **Disabled**) and press Enter.

3. To enable or disable the ability to have drive slots in the system backplane or in a storage enclosure dedicated as hot spare slots, highlight **Persistent Hot Spare** and press Enter. Select either **Enabled** or **Disabled** and press Enter.

   If enabled, replacement of a hot spare drive in the same slot automatically configures the drive as a hot spare.
4. To enable or disable the option to copy data back from a hot spare drive to a physical drive, highlight **Replace Drive** and press Enter. Select either **Enabled** or **Disabled** and press Enter.
5. To enable or disable the option to start a Drive Replace operation, if a Self-Monitoring Analysis and Report Technology (SMART) error is detected on a physical drive, highlight **Replace Drive on SMART Error** and press Enter. Select either **Enabled** or **Disabled** and press Enter.
6. Highlight **Apply Changes** and press Enter.

   The new settings are saved in the controller properties.

5.5.14 Changing Task Rates

The following dialog appears when you select **Task Rates** from the **Advanced Controller Properties** dialog.

Figure 114 Task Rates

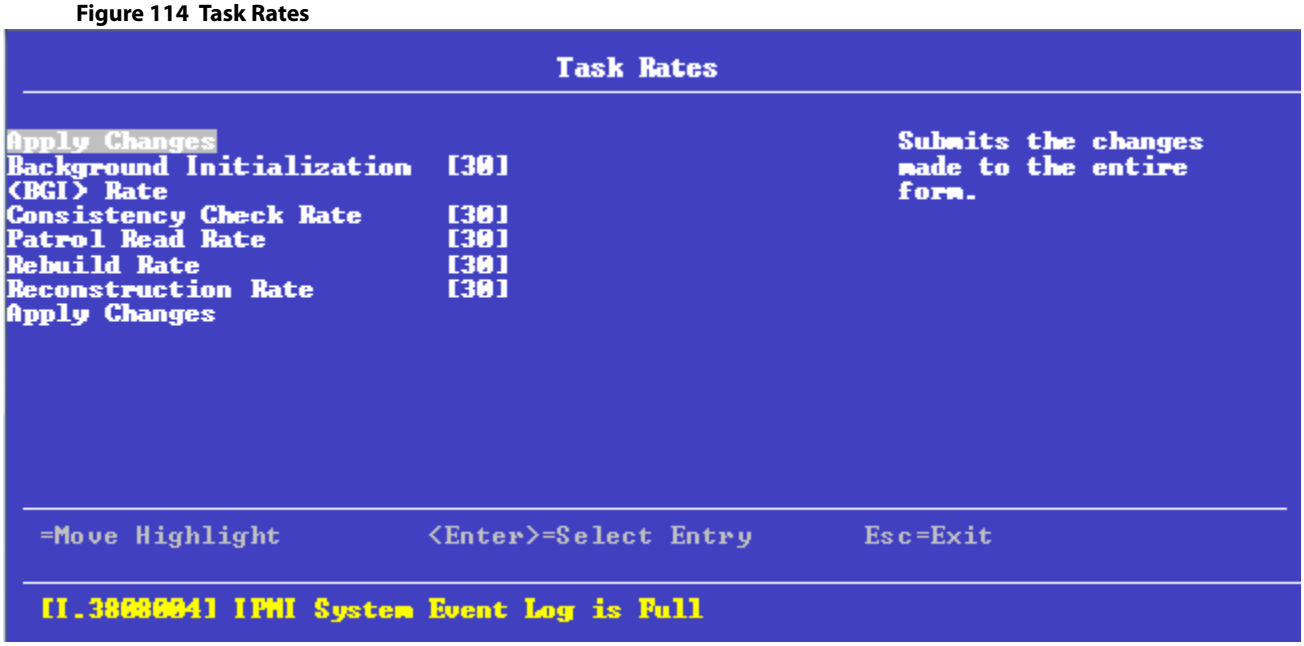

You can change the Rebuild rate and other task rates for a controller in the above dialog.

Follow these steps to change the task rates:

NOTE You can only view the properties/options supported by your controller.

1. To change the percentage of system resources dedicated to performing a BGI on a redundant virtual drive, highlight **Background Initialization <BGI> Rate** and press Enter. Specify a number from 0 to 100 and press Enter.

   The BGI rate is the percentage of the compute cycles dedicated to running a background initialization of drives on this controller. You can configure the BGI rate between 0 percent and100 percent. At 0 percent, the initialization operation runs only if the firmware is not doing anything else. At 100 percent, the initialization operation has a higher priority than I/O requests from the operating system. For best performance, use an initialization rate of approximately 30 percent.

2. To specify a rate for the percentage of system resources dedicated to performing a consistency check operation on a redundant virtual drive, highlight **Consistency Check Rate**, and press Enter. Specify a number from 0 to 100 and press Enter.

   The consistency check rate is the percentage of the compute cycles dedicated to running a consistency check on drives on this controller. You can configure the consistency check rate between 0 percent and 100 percent. At 0 percent, the consistency check operation runs only if the firmware is not doing anything else. At 100 percent, the consistency check operation has a higher priority than I/O requests from the operating system. For best performance, use a consistency check rate of approximately 30 percent.

3. To specify a rate for the percentage of system resources dedicated to performing a patrol read operation on configured physical drives, highlight **Patrol Read Rate** and press Enter. Specify a number from 0 to 100 and press Enter.

   The patrol read rate is the percentage of the compute cycles dedicated to running a patrol read on drives on this controller. You can configure the patrol read rate between 0 percent and 100 percent. At 0 percent, the patrol read runs only if the firmware is not doing anything else. At 100 percent, the patrol read has a higher priority than I/O requests from the operating system. For best performance, use a patrol read rate of approximately 30 percent.

4. To specify a rate for the percentage of system resources dedicated to rebuilding data on a new drive after a storage configuration drive has failed, highlight **Rebuild Rate** and press Enter. Specify a number from 0 to 100 and press Enter.

   The rebuild rate is the percentage of the compute cycles dedicated to rebuilding failed drives in virtual drives on this controller. You can configure the rebuild rate between 0 percent and 100 percent. At 0 percent, the rebuild runs only if the firmware is not doing anything else. At 100 percent, the rebuild operation has a higher priority than I/O requests from the operating system. For best performance, use a rebuild rate of approximately 30 percent.

5. To specify a rate for the percentage of system resources dedicated to performing a RAID Level Migration (RLM) or an Online Capacity Expansion (OCE) on a virtual drive, highlight **Reconstruction Rate** and press Enter. Specify a number from 0 to 100 and press Enter.

   The reconstruction rate is the percentage of the compute cycles dedicated to reconstructing data on drives on this controller. You can configure the reconstruction rate between 0 percent and100 percent. At 0 percent, the reconstruction operation runs only if the firmware is not doing anything else. At 100 percent, the reconstruction operation has a higher priority than I/O requests from the operating system. For best performance, use a reconstruction rate of approximately 30 percent.

6. Highlight **Apply Changes** and press Enter.

   The new settings are saved in the controller properties.

5.5.15 Upgrading the Firmware

The following dialog appears when you select **Update Firmware** from the **Dashboard View**. For a list of limitations, see Online Firmware Upgrade Support.

Figure 115 Controller Firmware Update

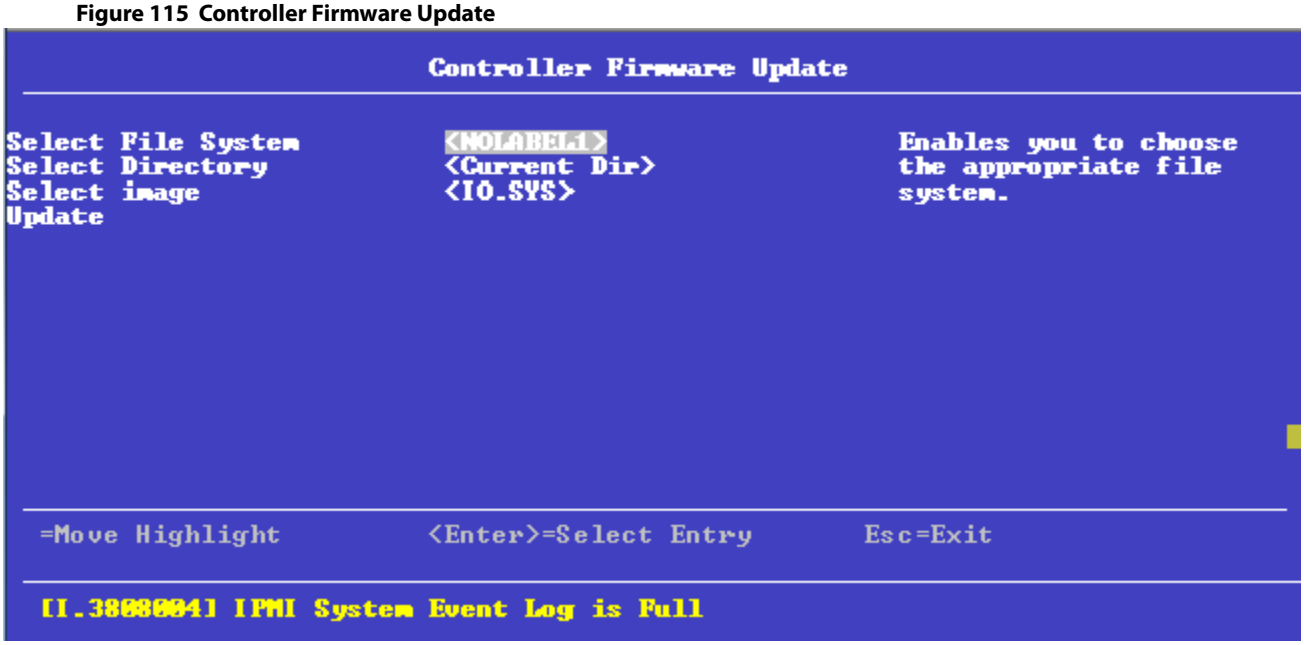

Follow these steps to upgrade the firmware:

1. To specify the file system where the `.rom` update file resides, highlight **Select File System** and press Enter. Select the file system and press Enter.
2. To specify the directory where the `.rom` file resides, highlight **Select Directory** and press Enter. Browse to the required the directory and press Enter.

   The current directory is normally highlighted. You can browse to only one level higher or one level lower.
3. To specify the `.rom` file, highlight **Select Image** and press Enter. Select the `.rom` file and press Enter.
4. Highlight **Update** and press Enter.

   The following Warning dialog appears.

Figure 116 Firmware Update Warning

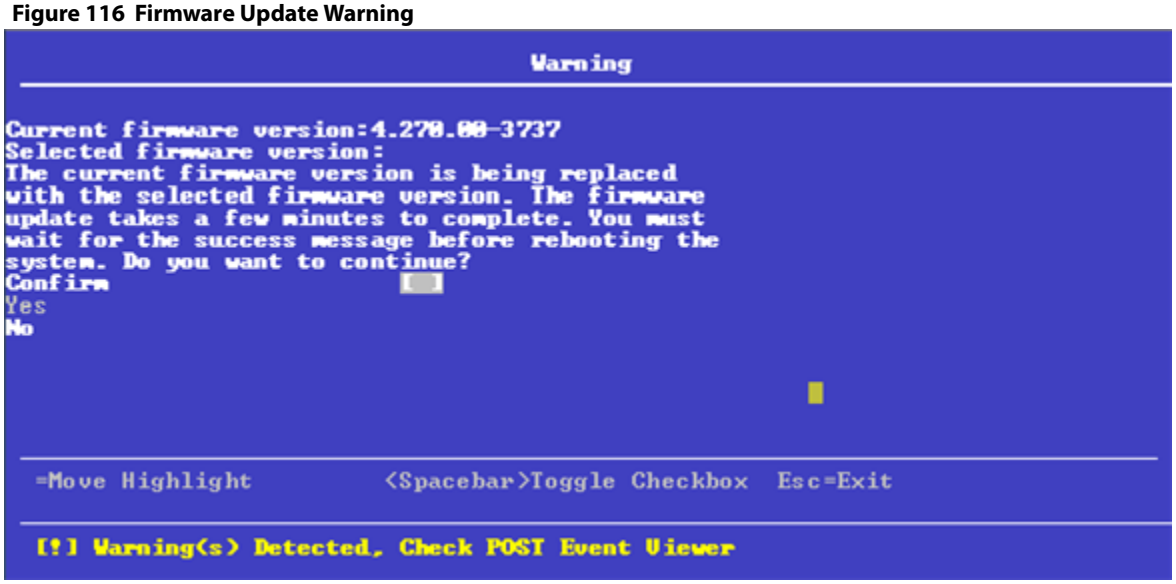

5. Highlight the **Confirm** check box and press the spacebar to select the check box.

6. Click **Yes** to continue with the firmware update.

   After the controller is successfully updated with the new firmware code, a message box appears stating the same. Highlight **OK** and click **Enter** in the message box to return to the **Controller Management** dialog.

5.6 Managing Virtual Drives

When you select **Virtual Drive Management** on the **Main Menu**, the **Virtual Drive Management** dialog appears, as shown in the following figure.

Figure 117 Virtual Drive Management

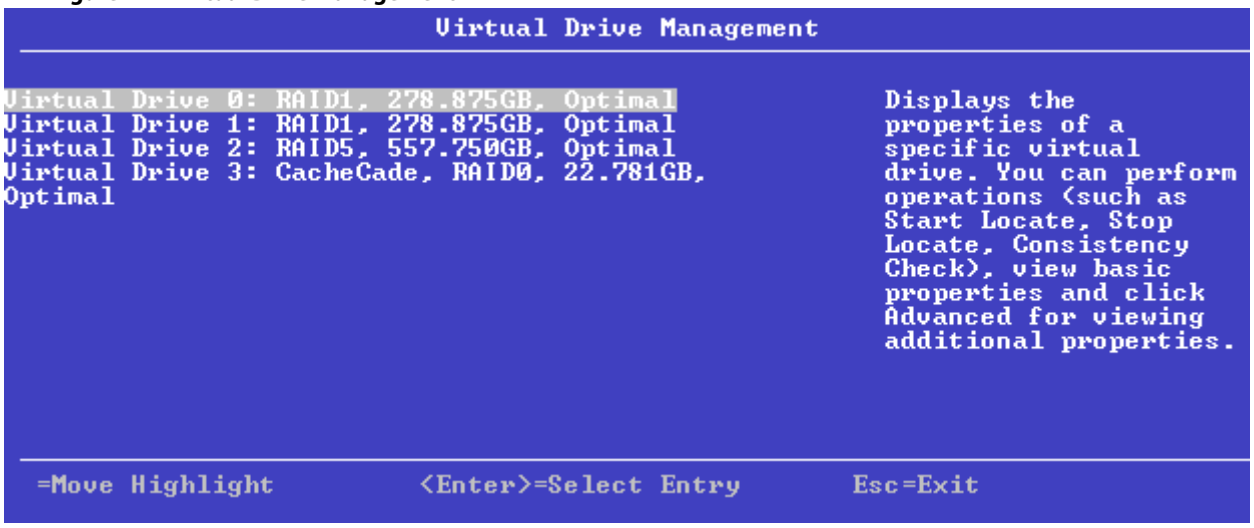

The menu lists all the virtual drives that currently exist on the controller. Highlight the virtual drive you want to manage and press Enter. The following dialog appears.

Figure 118 Virtual Drive Management

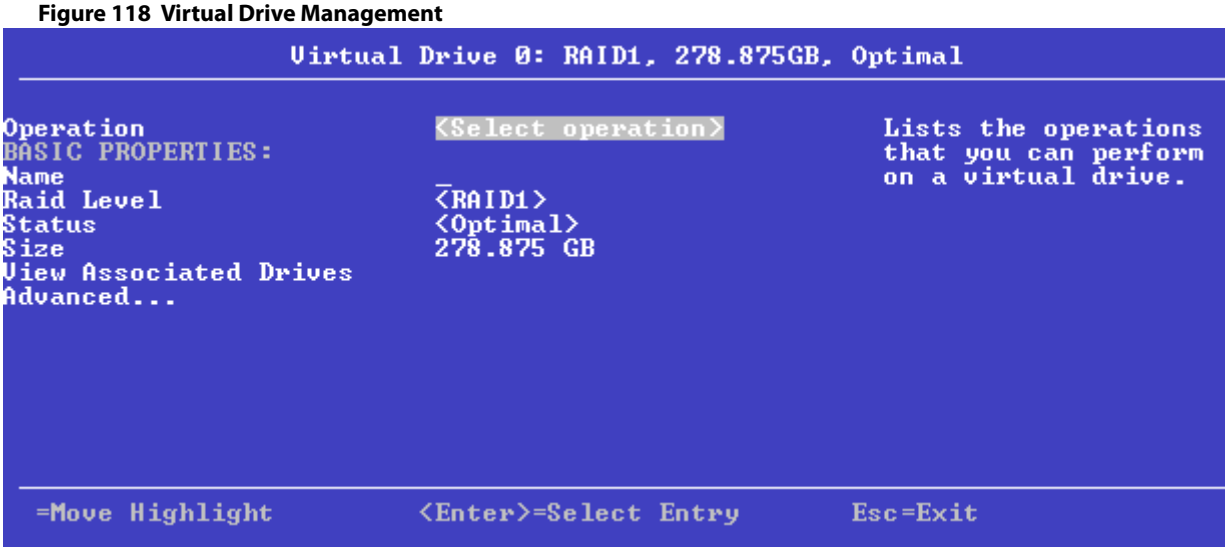

This dialog lists the following basic virtual drive properties.

Table 30 Basic Virtual Drive Properties

Property	Description
Name	The name assigned to the virtual drive. To assign a name or to change the name, highlight the field, press Enter, and type the new name in the popup window.
RAID Level	The RAID level of the virtual drive.
Status	The current status of the virtual drive.
Size	The capacity of the virtual drive, in MB or GB. **NOTE** Virtual drive size of floating data type up to three decimal places is supported. Some of the screens in this chapter may not show this feature.

For information on how to perform virtual drive operations, see Selecting Virtual Drive Operations.

For information on how to view the physical drives associated with the virtual drive, see Viewing Associated Drives.

For information on how to view and change advanced virtual drive settings, see Viewing and Managing Virtual Drive Properties and Options.

5.6.1 Selecting Virtual Drive Operations

The following popup menu appears when you highlight **Operation** in the **Virtual Drive** window and press Enter.

Figure 119 Virtual Drive Operations Menu

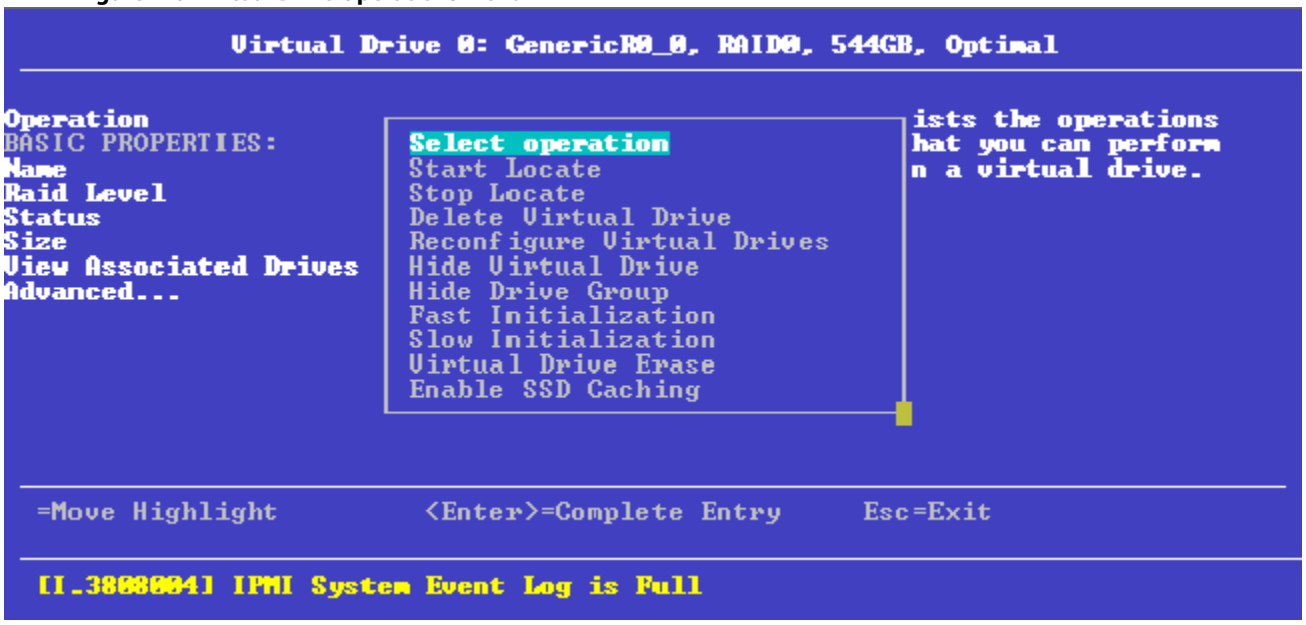

Other options, such as **Enable/Disable SSD Caching**, **Secure Virtual Drive**, **Check Consistency**, and **Expand Virtual Drive**, might also appear, depending on the current configuration of the system.

Highlight the operation you want to select and press Enter. Then highlight the word **Go** that appears beneath **Operation** and press Enter to start the operation for the currently selected virtual drive.

The following sections explain how to run the operations.

5.6.1.1 Locating Physical Drives in a Virtual Drive

To locate the physical drives in a virtual drive by flashing their LEDs, perform these steps:

1. Highlight **Start Locate** on the pop-up menu and press Enter.
2. Highlight the word **Go** that appears beneath **Operation** and press Enter.
   A Success message appears.
3. Highlight **OK** and press Enter to return to the **Virtual Drive** dialog.
   The LEDs on the physical drives start flashing, if the drive firmware supports this feature.
4. Observe the location of the drives with the flashing LEDs.
5. To stop the LEDs from flashing, access the popup menu again, highlight **Stop Locate**, and press Enter.
6. Highlight the word **Go** that appears beneath **Operation** and press Enter.
   A Success message appears.
7. Highlight **OK** and press Enter to return to the **Virtual Drive** dialog.
   The LEDs on the physical drives stop flashing.

5.6.1.2 Deleting a Virtual Drive

CAUTION All data on a virtual drive is lost when you delete it. Back up data you want to keep before you delete a virtual drive.

The delete virtual drive action is performed on the currently selected virtual drive. To select a different virtual drive for deletion, press Esc to return to the **Virtual Drive Selection** dialog and select the virtual drive.

To delete a virtual drive, perform these steps:

1. Highlight **Delete Virtual Drive** on the pop-up menu and press Enter.
2. Highlight the word **Go** that appears beneath **Operation** and press Enter.
   The **Delete Virtual Drive** warning message appears.
3. Highlight **Confirm** and press the spacebar to confirm the deletion, then highlight **Yes** and press Enter.
   The virtual drive is deleted.

NOTE The group initialization process is time-consuming when it is performed simultaneously on multiple drives when I/O transactions are in progress. You cannot close the **Group Initialization** dialog and perform any other operation on the MegaRAID Storage Manager application until this process completes.

5.6.1.3 Hiding a Virtual Drive

To hide a virtual drive, perform these steps:

1. Highlight **Hide Virtual Drive** on the pop-up menu and press Enter.
2. Highlight the word **Go** that appears beneath **Operation** and press Enter.
   The **Hide Virtual Drive** warning message appears.
3. Highlight **Confirm** and press the spacebar to confirm the deletion, and then highlight **Yes** and press Enter.
   The virtual drive is hidden.

5.6.1.4 Unhiding a Virtual Drive

To unhide a virtual drive, perform these steps:

1. Highlight **Un-Hide Virtual Drive** on the pop-up menu and press Enter.
2. Highlight the word **Go** that appears beneath **Operation** and press Enter.
   The **Un-Hide Virtual Drive** warning message appears.

3. Highlight **Confirm** and press the spacebar to confirm the deletion, and then highlight **Yes** and press Enter.
   The virtual drive is unhidden.

5.6.1.5 Hiding a Drive Group

To hide a drive group to which the virtual drive is associated, perform these steps:

1. Highlight **Hide Drive Group** on the pop-up menu and press Enter.
2. Highlight the word **Go** that appears beneath **Operation** and press Enter.
   The **Hide Drive Group** warning message appears.
3. Highlight **Confirm** and press the spacebar to confirm the deletion, and then highlight **Yes** and press Enter.
   The drive group is hidden.

5.6.1.6 Unhiding a Drive Group

To unhide a drive group to which the virtual drive is associated, perform these steps:

1. Highlight **Un-Hide Drive Group** on the pop-up menu and press Enter.
2. Highlight the word **Go** that appears beneath **Operation** and press Enter.
   The **Un-Hide Drive Group** warning message appears.
3. Highlight **Confirm** and press the spacebar to confirm the deletion, and then highlight **Yes** and press Enter.
   The drive group is unhidden.

5.6.2 Performing a Transport Ready Operation

Transport Ready operation allows you to move a drive group and all the shared or exclusive dedicated hot spare drives associated with that drive group from one server to another server non-disruptively, in a real time.

NOTE To know more about the Transport Ready feature, see Transport Ready Pre-requisites section.

Follow these steps to perform a transport ready operation:

1. In the **VD Mgmt** screen, navigate to **Operation** and select **Transport Ready**.

Figure 120 Transport Ready

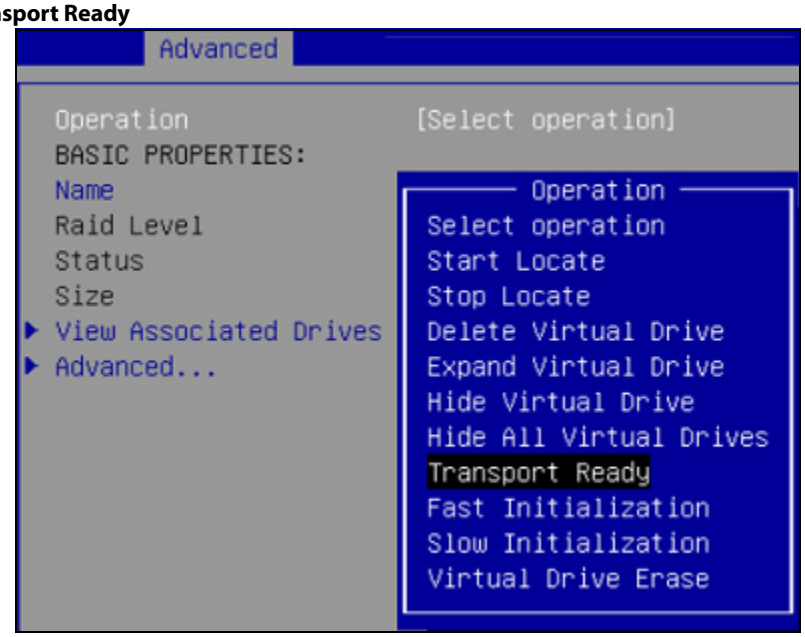

2. In the **Operation** screen, type `Go`
3. Select **Include Exclusive Dedicated Hot Spare Drive(s)** option
4. Click `Yes`.

   If you have any `exclusive dedicated host spare drives` associated with the drive group on which you are performing the Transport Ready operation, a warning message is displayed requesting you to confirm your action.

Figure 121 Transport Ready for Exclusive Dedicated Hot Spare

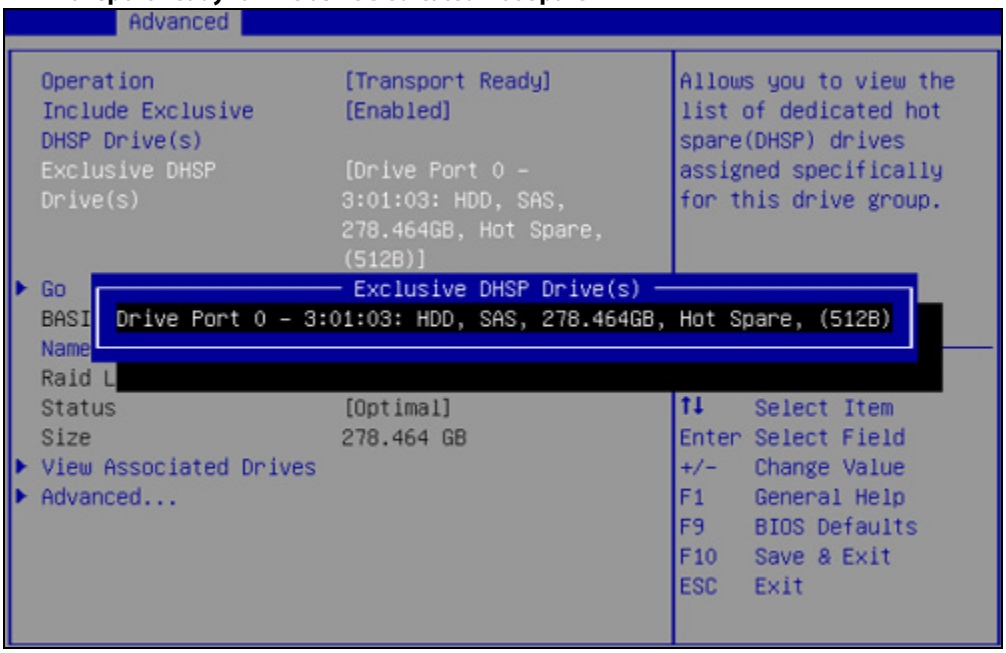

5. Select **Include Shared Dedicated Hot Spare Drive(s)** option
6. Click `Yes`.

   If you have any Shared `dedicated host spare drives` associated with the drive group on which you are performing the Transport Ready operation, a warning message is displayed requesting you to confirm your action.

Figure 122 Transport Ready for Shared Dedicated Hot Spare

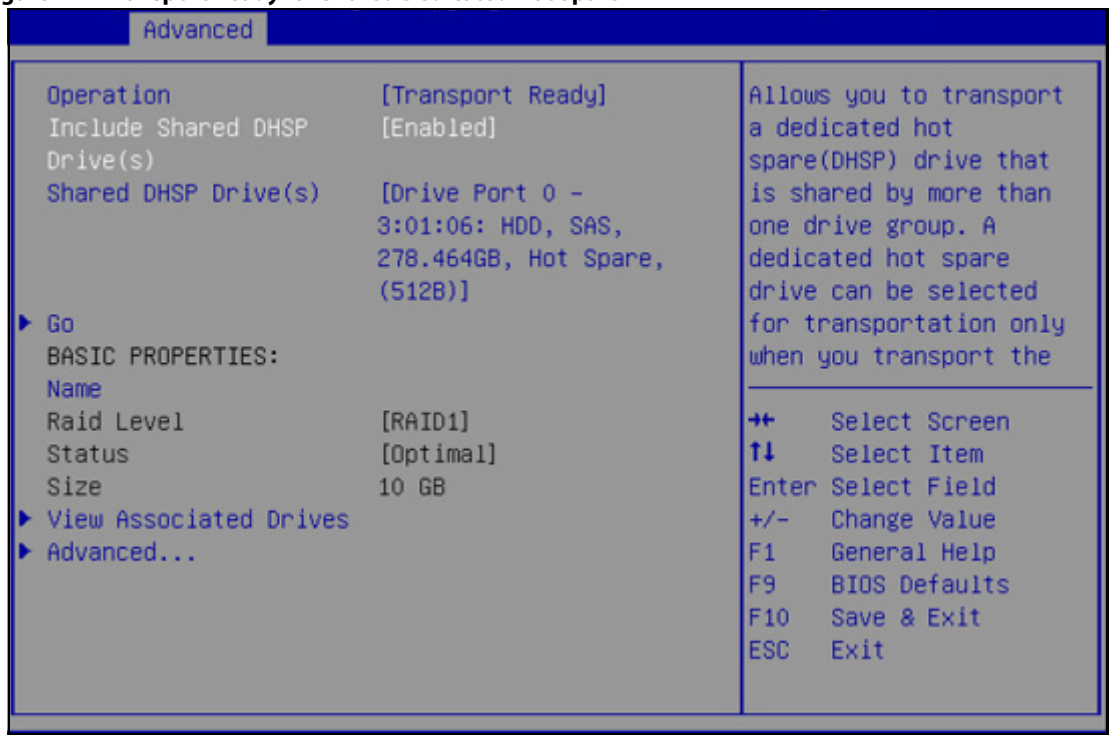

If you have both exclusive and shared dedicated host spare drives associated with the drive group on which you are performing the Transport Ready operation, a warning message is displayed requesting you to confirm your action.

Figure 123 Transport Ready with Exclusive Dedicated Hot Spare and Shared Dedicated Hot Spare

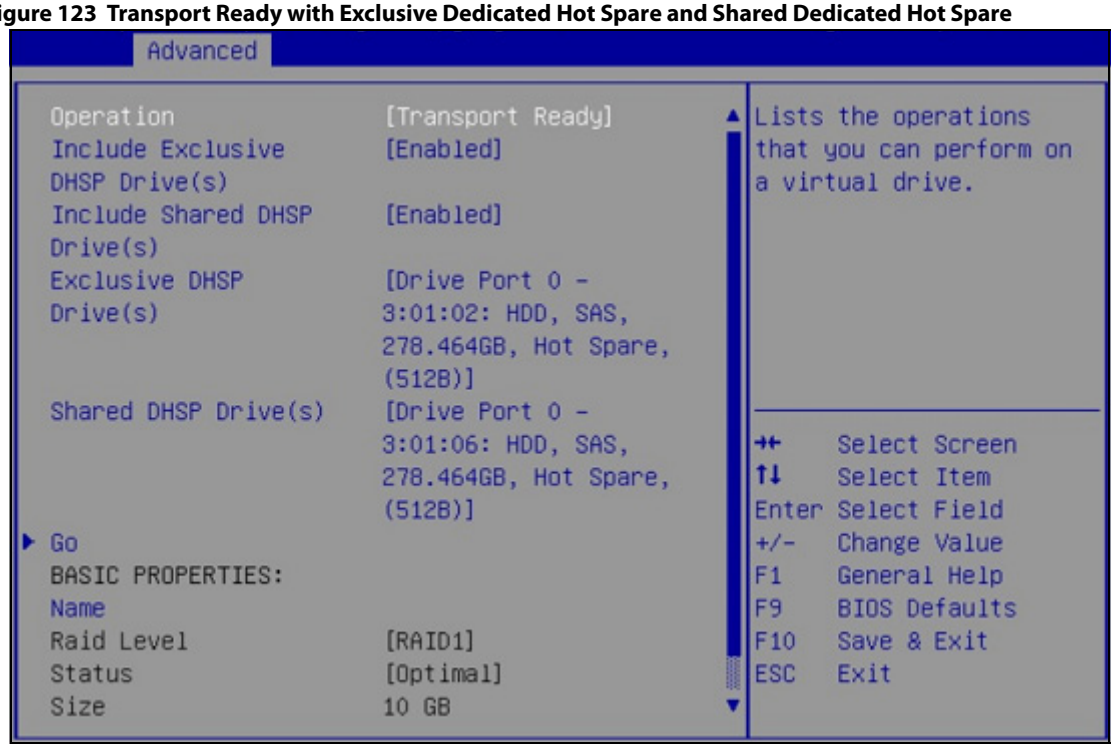

A warning message is displayed requesting you to confirm your action.

Figure 124 Confirm Transport Ready

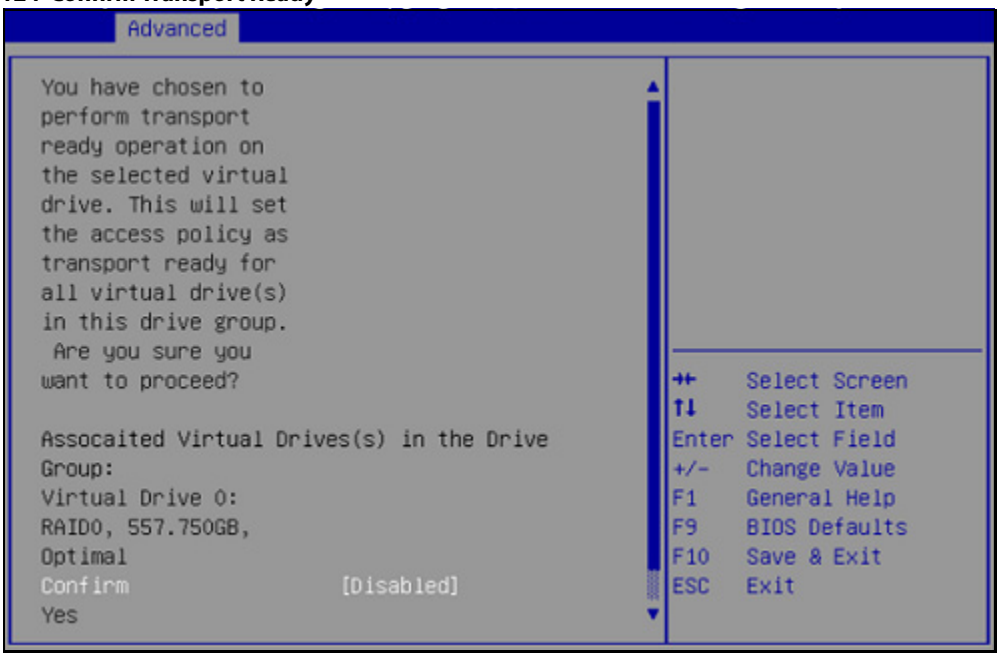

7. Click `Yes`.

   The drive group on which you have performed the Transport Ready operation is appended and a Transport Ready property is displayed on the drive group. Once the Transport Ready property is displayed on the drive group, you can remove the drive group

Figure 125 Access Policy set as Transport Ready

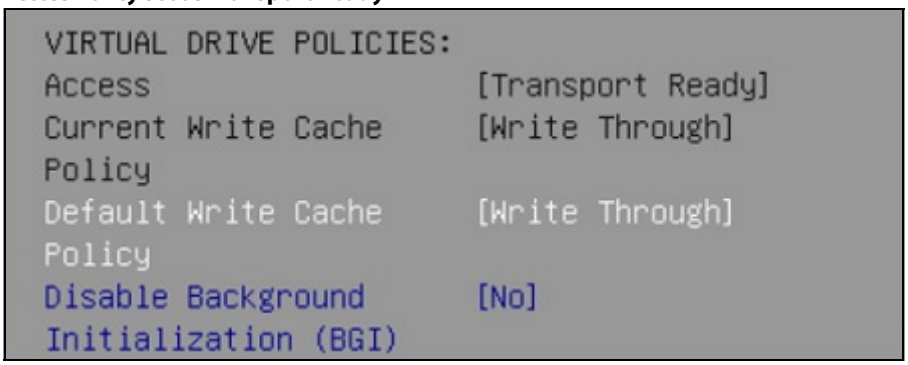

After the Transport Ready operation is performed successfully, the access policy of the virtual drive associated with the drive group on which the Transport Ready operation has been performed, is defaulted to Transport Ready. You will not have any other access policy options (such as Read Write, Read Only, and so on) available.

5.6.2.1 Clearing a Transport Ready Operation

If you have performed the Transport Ready operation on a drive group and if you want to reverse it (undo your previous action), you can do so using the Clear Transport Ready option.

Follow these steps to clear a Transport Ready operation:

1. In the **VD Mgmt** screen, navigate to the controller, and press F2 Key.
2. Select the drive group on which want to perform transport ready and select **Clear Transport Ready**.

Figure 126 Clear Transport Ready

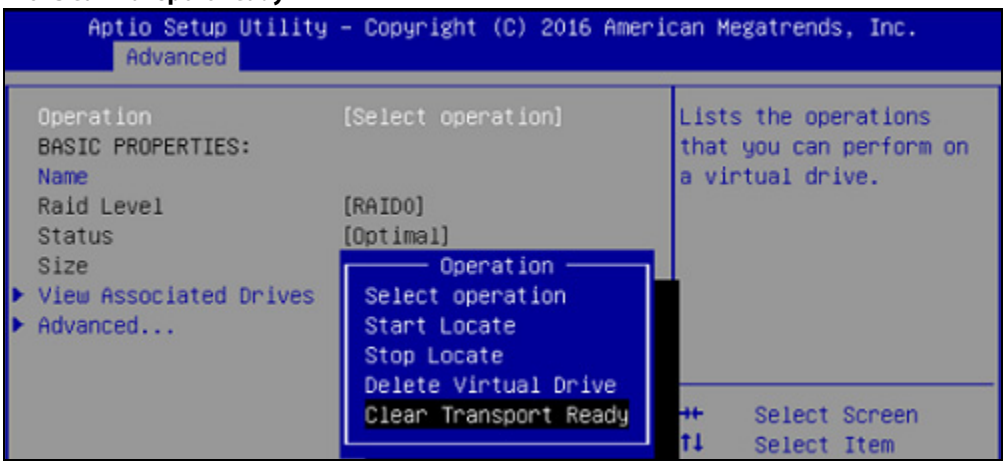

A warning message is displayed requesting you to confirm your action.

Figure 127 Confirm Clear Transport Ready

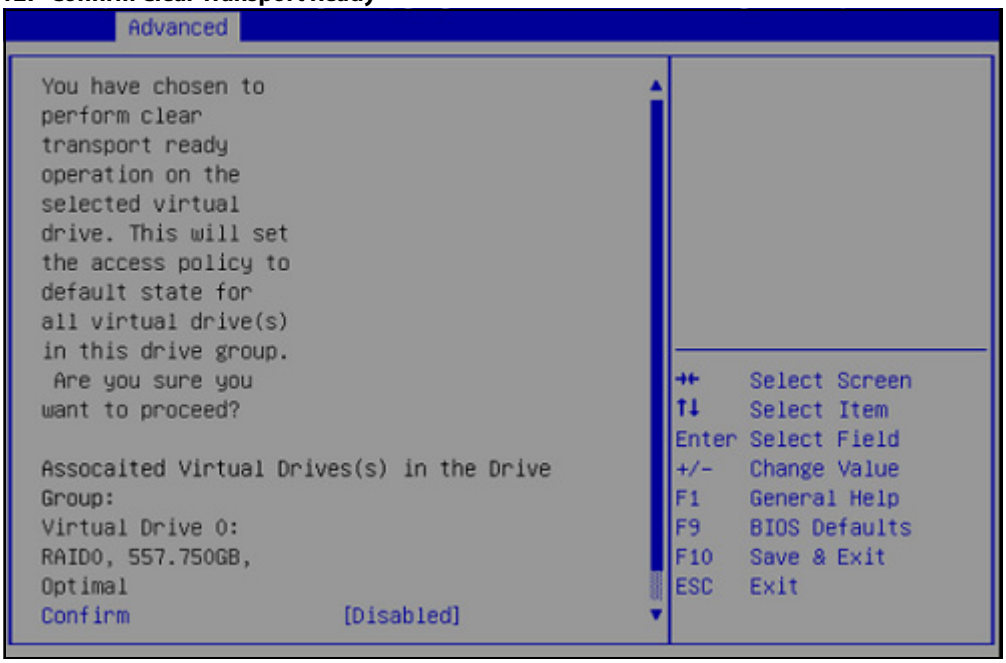

3. Select `Yes` and press Enter.

5.6.2.2 Reconfiguring a Virtual Drive

You can reconfigure a virtual drive by changing its RAID level, or by adding physical drives to it, or by doing both of these actions. When performing these changes, however, you must observe the maximum drive and minimum drive restrictions for the various RAID levels. See Table 26 for more information.

To reconfigure a virtual drive, perform these step:

1. Highlight **Reconfigure Virtual Drive** on the pop-up menu and press Enter.
2. Highlight the word **Go** that appears beneath **Operation** and press Enter.

   The following dialog appears.

Figure 128 Reconfigure Virtual Drives

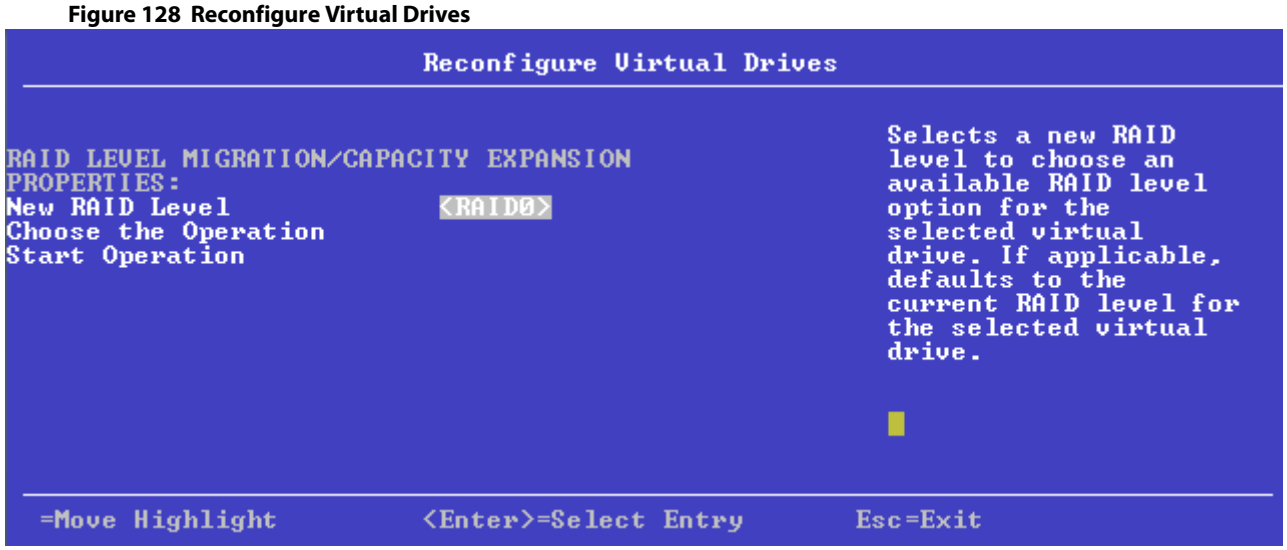

3. To change the RAID level of the selected virtual drive, highlight **New RAID Level** and press Enter.
4. Select a RAID level from the pop-up menu.
5. Depending on the source and the target RAID levels, you can either add drives or remove drives. Highlight **Choose the Operation** and press Enter.
6. Choose either **Add Drives** or **Remove Drives**

5.6.2.2.1 Adding Drives to a Configuration

Perform the following steps to add unconfigured drives to a configuration while reconfiguring a virtual drive.

1. If you select the **Add Drives** option and press Enter, the following dialog appears.

Figure 129 Select Drives – Add Drives

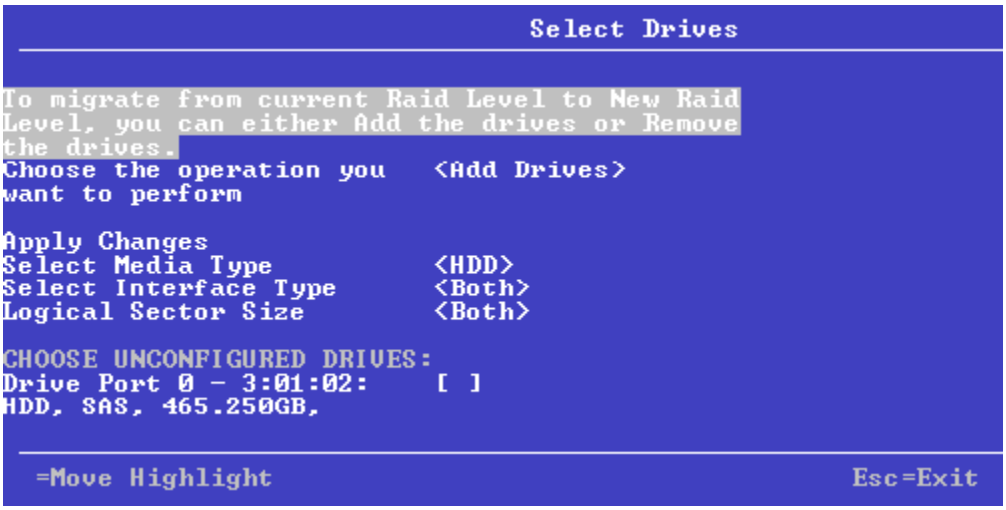

2. (Optional) To change the default **Select Media Type** value, highlight this field, press Enter, and select an option from the pop-up menu.

   The choices are **HDD** and **SSD**. Combining HDDs and SSDs in a virtual drive is not supported.
3. (Optional) To change the default **Select Interface Type** value, highlight this field, press Enter, and select an option from the pop-up menu.

   The choices are **SAS**, **SATA**, and **Both**. Depending on the configuration of your system, combining SAS and SATA drives in a virtual drive might not be supported.

4. To select unconfigured drives to add to the configuration, highlight the drives and press the spacebar. A small red arrow at the bottom of the dialog indicates you can scroll down to view more drives.

   NOTE The red arrow appears when there is too much information to display in one dialog. The amount of information that can be displayed in one dialog depends on the capabilities of the HII browser.

   Alternatively, use the **Check All** and **Uncheck All** options at the bottom of the list of drives to select or deselect all available drives.

   NOTE Be sure to select the number of drives required by the specified RAID level; otherwise, the HII utility displays an error message when you try to create the virtual drive. For example, RAID 1 virtual drives use exactly two drives and RAID 5 virtual drives use three or more drives. See Table 26 for more information.

5. When you have selected the unconfigured drives to add, highlight **Apply Changes** and press Enter.

   NOTE If you have selected drives of varying sizes, the HII utility displays a message warning you that the remaining free capacity on the larger drives will be unusable.

   The HII utility returns you to the **Reconfigure Virtual Drives** dialog.

5.6.2.2.2 Removing Drives from a Configuration

Perform the following steps to remove drives from a configuration while reconfiguring a virtual drive.

1. If you select the **Remove Drives** option and press Enter, the following dialog appears.

Figure 130 Select Drives – Remove Drives

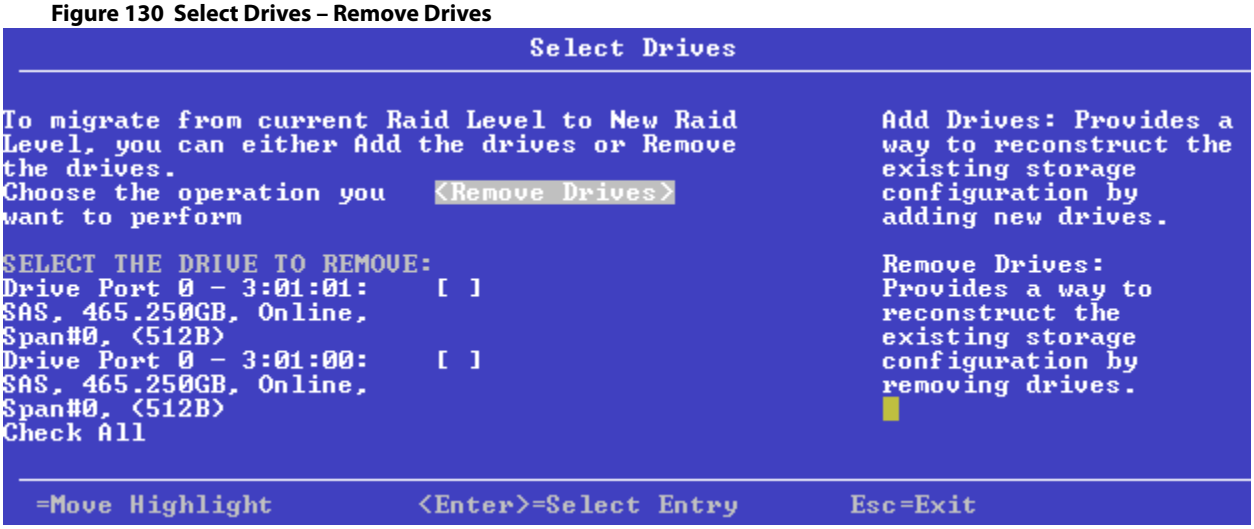

2. To select the drives to remove from the configuration, highlight the drives and press the spacebar. A small red arrow at the bottom of the dialog indicates you can scroll down to view more drives.

   NOTE The red arrow appears when there is too much information to display in one dialog. The amount of information that can be displayed in one dialog depends on the capabilities of the HII browser.

   Alternatively, use the **Check All** and **Uncheck All** options at the bottom of the list of drives to select or deselect all available drives.

3. When you have selected the drives to remove, highlight **Apply Changes** and press Enter.

   The HII utility returns you to the **Reconfigure Virtual Drives** dialog.

5.6.2.3 Initializing a Virtual Drive

To initialize a virtual drive, perform these steps:

ATTENTION All data on the virtual drive is lost when you initialize it. Before you start this operation, back up any data that you want to keep.

1. Highlight **Fast Initialization** or **Slow Initialization** on the pop-up menu and press Enter.

   A fast initialization overwrites the first and last 8 MB of the virtual drive, clearing any boot records or partition information. A slow (full) initialization overwrites all blocks and destroys all data on the virtual drive.
2. Highlight the word **Go** that appears beneath **Operation** and press Enter.

   The **Initialize Virtual Drive Warning** dialog appears.
3. Highlight **Confirm** and press the spacebar to confirm the operation, then highlight **Yes** and press Enter.

   A progress indicator shows the percentage completion of the initialization process. This indicator refreshes automatically.

5.6.2.4 Erasing a Virtual Drive

To erase data on a virtual drive, perform these steps:

ATTENTION All data on the virtual drive is lost when you erase it. Before you start this operation, back up any data that you want to keep.

NOTE After the data is erased, you have the option to keep the blank virtual drive, which you can use to store other data, or to delete the virtual drive completely.

1. Highlight **Virtual Drive Erase** on the pop-up menu and press Enter.

   Two additional fields appear.
2. Highlight **Erase Mode** and press Enter.
3. Select **Simple**, **Normal**, or **Thorough** from the pop-up menu.

   A Simple erase writes a pattern to the virtual drive in a single pass. The other erase modes make additional passes to erase the data more thoroughly.
4. (Optional) Highlight **Delete After Erase** and press the spacebar to select it.
5. Highlight **Go** and press Enter.

   The **Virtual Drive Erase** warning message appears.
6. Highlight **Confirm** and press the spacebar to confirm the operation, then highlight **Yes** and press Enter.

   A progress indicator shows the percentage completion of the operation. This indicator refreshes automatically. After the completion of the operation, the virtual drive is erased.

5.6.2.5 Enabling and Disabling SSD Caching

When you enable SSD caching, the selected virtual drive becomes associated with an existing or future CacheCade SSD caching virtual drive. When you disable SSD caching, this association is deleted. Follow these steps to enable or disable SSD caching for a virtual drive.

1. Highlight **Enable/Disable SSD Caching** on the pop-up menu and press Enter.
2. Highlight **Go** and press Enter.

   The **Enable SSD Caching Warning** message appears.

3. Highlight **Confirm** and press the spacebar to confirm the operation, then highlight **Yes** and press Enter.
   SSD caching is enabled for this virtual drive.

The warning is similar when you disable SSD caching.

5.6.2.6 Securing a Virtual Drive

A Secure Virtual Drive operation enables security on a virtual drive. You can only disable the security by deleting the virtual drive. Perform these steps to secure a virtual drive.

1. Highlight **Secure Virtual Drive** on the pop-up menu and press Enter.
   The **Secure Virtual Drive** warning appears.
2. Highlight **Confirm** and press the spacebar to confirm the operation, then highlight **Yes** and press Enter.
   The virtual drive is secured.

5.6.2.7 Running a Consistency Check

Follow these steps to run a consistency check on the currently selected redundant virtual drive.

1. Highlight **Check Consistency** on the pop-up menu and press Enter.

   NOTE The **Check Consistency** option does not appear on the menu if the currently selected virtual drive is either RAID 0 or RAID 00 (nonredundant).

2. Highlight **Go** and press Enter.
   The **Consistency Check Success** dialog appears.
   As the message indicates, the consistency check is now running.
3. Highlight **OK** and press Enter.
   The Progress indicator in the dialog shows the percentage progress of the consistency check. To refresh the indicator, exit the dialog and re-enter it.
4. To stop or suspend the consistency check, highlight **Stop** or **Suspend** and press Enter.
5. To resume a suspended consistency check, highlight **Resume** and press Enter.
   A progress indicator shows the percentage completion of the operation. This indicator refreshes automatically.

For more information about consistency checks, see Scheduling a Consistency Check.

5.6.2.8 Expanding a Virtual Drive

Expanding a virtual drive means increasing its capacity. Existing data on the virtual drive is not impacted by the expansion. Follow these steps to expand the currently selected virtual drive.

1. Select **Expand Virtual Drive** from the pop-up menu.
   The **Expand Virtual Drive** dialog appears.
   The dialog shows the current capacity of the selected virtual drive, the available capacity that can be added to it, and the capacity of the expanded virtual drive, if all available capacity is added.
2. To change the amount of available capacity, highlight the **Enter a Percentage of Available Capacity** field and use the minus key (–) on the keypad to reduce percentage.

   NOTE Some systems permit you to enter numeric values directly, without using the + and – keys.

3. When you have set the capacity to the desired level, highlight **OK** and press Enter.
   The capacity of the virtual drive is expanded.

5.6.2.9 Disabling Protection on a Virtual Drive

To disable data protection on virtual drives, perform these steps:

1. Highlight **Disable Protection** on the pop-up menu and press Enter.
2. Highlight the word **Go** that appears beneath Operation and press Enter.
   Data protection is disabled on virtual drives.

5.6.3 Managing CacheCade Virtual Drives

After you create a CacheCade virtual drive, as described in Creating a CacheCade Virtual Drive, you can select it on the **Virtual Drive Management** menu, run operations on it, and manage it in other ways.

The following window appears when you select a CacheCade virtual drive in the **Virtual Drive Management** menu.

Figure 131 Manage CacheCade Virtual Drive Window

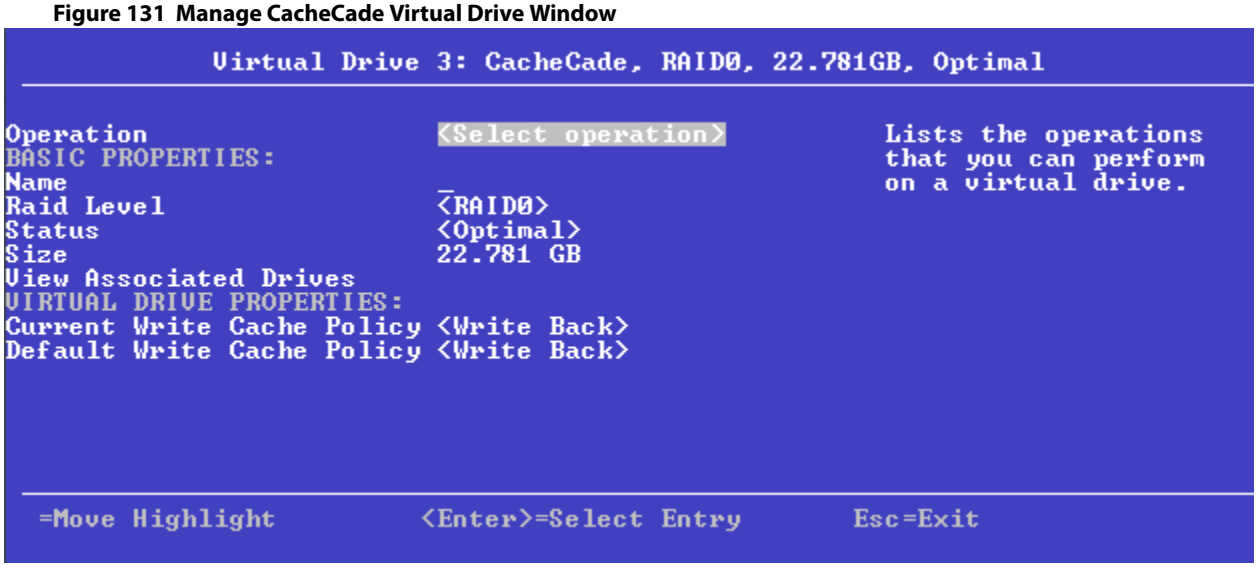

This window lists basic information about the CacheCade virtual drive, including name, RAID level, status, and size.

You can select and run the following operations on a CacheCade virtual drive:

- **Start Locate/Stop Locate**
  Use this option to flash the light on the SSD used for the CacheCade virtual drive. For more information, see Locating Physical Drives in a Virtual Drive.
- **Delete Virtual Drive**
  Use this option to delete the CacheCade virtual drive. For more information, see Deleting a Virtual Drive.

To assign a name to the CacheCade virtual drive, highlight **Name**, press Enter, type the name, and press Enter again.

To change the default write cache policy, highlight **Default Write Cache Policy**, press Enter, and select an option from the popup menu. Options are **Write Through**, **Write Back**, and **Force Write Back**.

To view the drives associated with the CacheCade virtual drive, highlight **View Associated Drives** and press Enter. For more information, see Enabling and Disabling SSD Caching.

5.6.4 Viewing Associated Drives

The **View Associated Drives** dialog appears when you select **View Associated Drives** at the bottom of the **Virtual Drive** window.

The dialog lists all the physical drives associated with the currently selected virtual drive. Follow these steps to view information about the associated drives.

1. To select a different virtual drive, highlight **Selected Virtual Drive**, press Enter, and select an entry from the popup menu.
2. Highlight one of the associated drives and press the spacebar to select it.
3. Highlight **View Drive Properties** and press Enter.

   The **View Drive Properties** window for the drive appears.
4. View the information on the **View Drive Properties** window.

   For more information, see Viewing Advanced Drive Properties.

5.6.5 Viewing and Managing Virtual Drive Properties and Options

The following dialog appears when you select **Advanced** from the **Virtual Drive** dialog. (The second dialog shows the rest of the options that are visible when you scroll down.)

NOTE The properties and options shown in the dialog apply to the currently selected virtual drive. To manage properties for a different virtual drive, press Esc until you return to the **Virtual Drive Selection** menu, select the desired virtual drive, and navigate back to this dialog.

Figure 132 Advanced Virtual Drive Properties 1

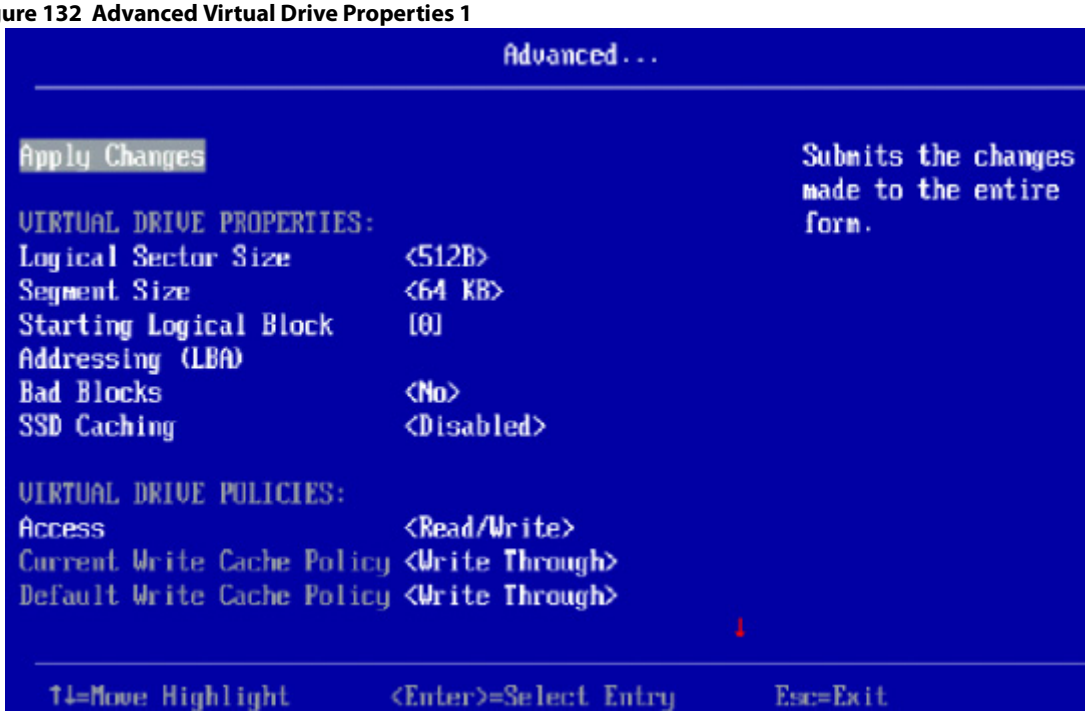

The small red arrow at the bottom of the dialog indicates that you can scroll down to view more virtual drive properties and virtual drive policies, as shown in the preceding figure.

Figure 133 Advanced Virtual Drive Properties 2

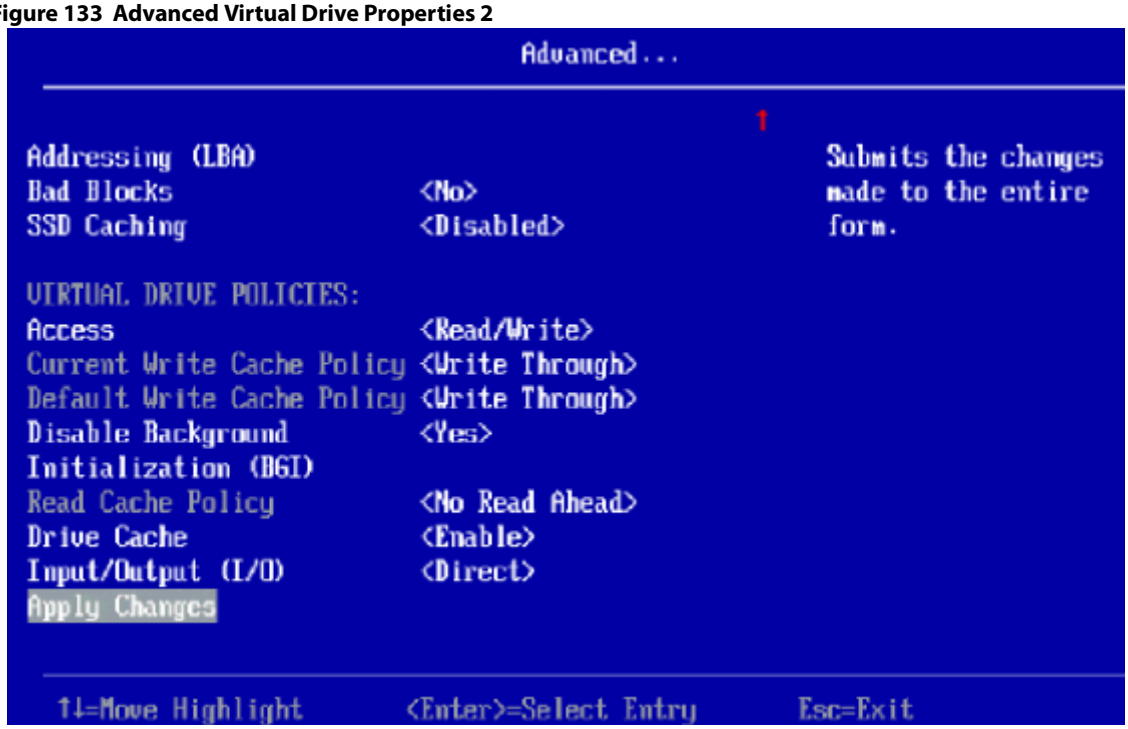

NOTE The red arrow appears when there is too much information to display in one dialog. The amount of information that can be displayed in one dialog depends on the capabilities of the HII browser.

The following table describes all of the virtual drive properties listed in this dialog.

Table 31 Virtual Drive Properties

Property	Description
Logical Sector Size	The logical sector size of this virtual drive. The possible options are **4 KB** and **512 B**.
Segment Size	The segment size used on this virtual drive.
Starting Logical Block	The address of the first location of a block of data stored on the virtual drive.
Addressing (LBA)	Indicates whether the virtual drive is secured.
Bad Blocks	Indicates whether the virtual drive has bad blocks.
SSD Caching	Indicates whether solid-state disk (SSD) caching is enabled on this virtual drive.

Following the virtual drive properties listed in the dialog are virtual drive policies that you can select and change. To change any policy, highlight the field, press Enter, and select a value from the pop-up menu. When you finish changing policy settings, highlight **Apply Changes** at the top or the bottom of the selections and press Enter.

The following table describes the virtual drive policies.

Table 32 Virtual Drive Policies

Property	Description
Access	The access policy for the virtual drive. The options are **Read/Write**, **Read Only**, and **Blocked**.
Current Write Cache Policy	Displays the current write cache policy. The possible values are as follows: ■ **Write-Through** (**WThru**) The controller sends a data transfer completion signal to the host when the virtual drive has received all of the data and has completed the write transaction to the drive. ■ **Write-Back** (**WBack**) 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 virtual drive in accordance with policies set up by the controller. These policies include the amount of dirty and clean cache lines, the number of cache lines available, and the elapsed time from the last cache flush. ■ **Force Write Back**.
Default Write Cache Policy	Displays the default write cache policy of the virtual drive.
Disable Background Initialization (BGI)	Specifies whether background initialization is enabled or disabled. When BGI is enabled, the firmware runs the initialization process in the background. When BGI is disabled, the initialization process does not start automatically and does not run in the background.
Read Cache Policy	The read cache policy for the virtual drive. The possible values are as follows: ■ **Ahead**: 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 process speeds up reads for sequential data, but provides little improvement when accessing random data. ■ **Normal**: Read-ahead capability is disabled.
Drive Cache	The disk cache policy for the virtual drive. The possible values are **Unchanged**, **Enable**, and **Disable**.
Input/Output (I/O)	The I/O policy for the virtual drive. The possible values are as follows: ■ **Direct**: Data reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. (The I/O policy applies to reads on a specific virtual drive. It does not affect the read- ahead cache.) ■ **Cached**: All reads are buffered in cache.

5.7 Managing Physical Drives

When you select **Drive Management** on the **Main Menu**, the **Drive Management Selection** dialog appears.

The menu lists all the physical drives that are connected to the controller. Highlight the drive you want to manage and press Enter. The following dialog appears.

Figure 134 Drive Management

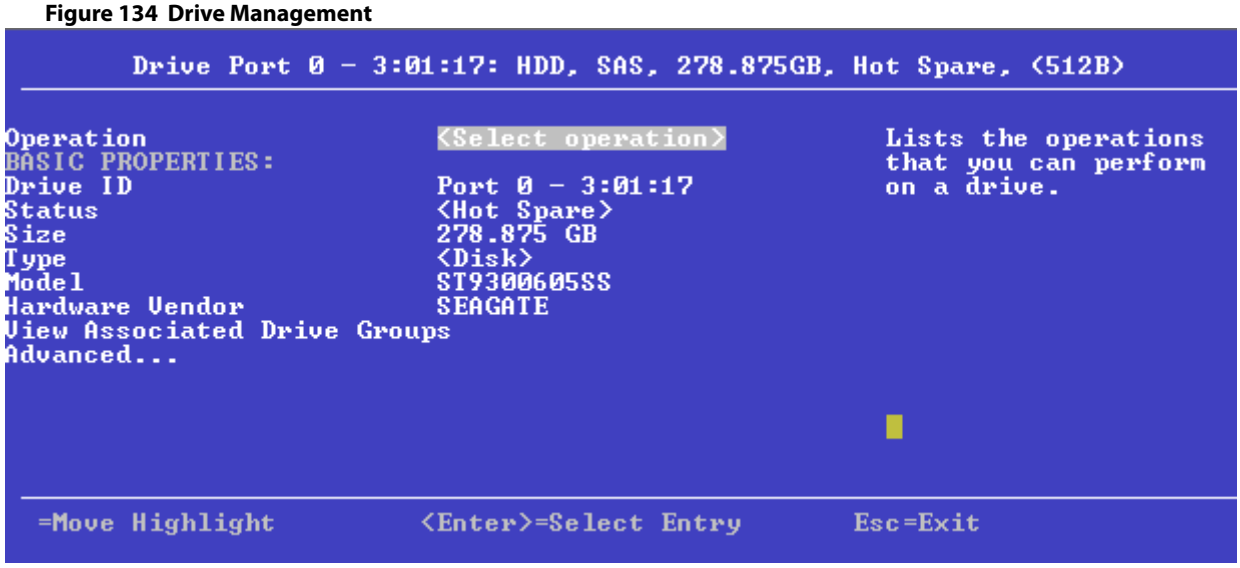

The preceding dialog lists the following basic drive properties for the selected drive:

Table 33 Basic Physical Drive Properties

Property	Description
Drive ID	The ID of the currently selected drive. The format of the ID is **Connector: Port wired order: Slot**. If the drive is not installed in an enclosure, the format of the ID is **Connector: Port wired order**.
Status	The status of the drive, such as **Online**, **Ready**, **Available**, or **Failed**.
Size	The drive capacity, in GB. Drive size of floating data type up to three decimal places is supported. Some of the screens in this chapter may not show this feature.
Type	The device type of the drive, which is normally **Disk**.
Model	The model number of the drive.
Hardware Vendor	The hardware vendor of the drive.
Associated Virtual Drive	If this physical drive is currently used in a virtual drive, this field lists information about the virtual drive. Highlight this field and press Enter to view a popup window with additional information about the virtual drive.
Associated Drive Groups	If this physical drive is associated with drive groups, this field lists information about the drive groups. Highlight this field and press Enter to view a popup window with a list of associated drive groups. Highlight a drive from the list and press Enter to view additional information about the drive group, such as associated virtual drives, the capacity allocation, and the assigned dedicated hot spare drives, if any.

For information on performing drive operations, see Performing Drive Operations.

For information on viewing and changing drive settings and properties, see Viewing Advanced Drive Properties.

5.7.1 Performing Drive Operations

When you highlight the **Select operation** field, press Enter and a pop-up drive operations menu appears.

Figure 135 Select Drive Operations Menu

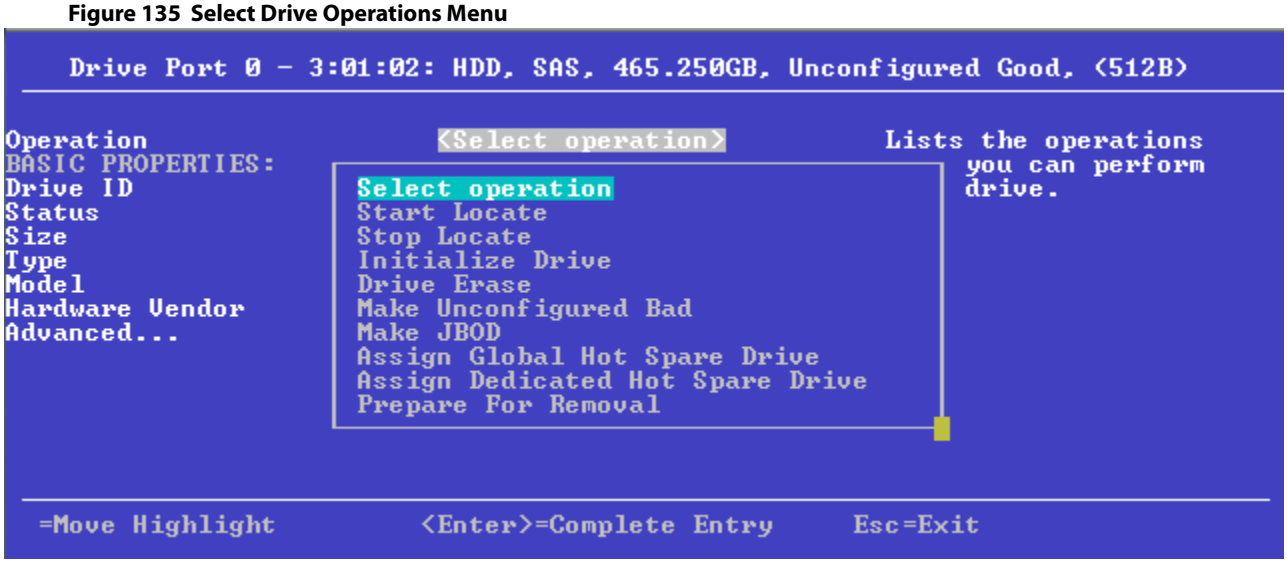

Start Locate and **Stop Locate** are the available options for any selected drive. The other menu options vary based on the status of the drive, which can be **Online**, **Offline**, **JBOD**, **Unconfigured Good**, **Unconfigured Bad**, **Global Hot Spare**, and **Dedicated Hot Spare**.

The following sections describe the available drive operations.

NOTE The drive operations run on the currently selected drive. To run an operation on a different drive, press Esc to return to the **Drive Selection** menu, highlight the drive you want to select, press Enter to select it, and return to this dialog.

5.7.1.1 Locating a Drive

Perform these steps to locate a physical drive by flashing its LED.

1. Open the pop-up drive operations menu, highlight **Start Locate**, and press Enter.
2. Highlight **Go**, which appears beneath **Operation**, and press Enter.
   A success message appears.
3. Highlight **OK** on the success message and press Enter.
   The LED on the selected drive starts flashing, if the drive firmware supports this feature.
4. Observe the location of the drive with the flashing LED.
5. To stop the LED from flashing, highlight **Stop Locate** on the popup menu and press Enter.
6. Highlight **Go**, which appears beneath **Operation**, and press Enter.
   A success message appears.
7. Highlight **OK** on the success message and press Enter, to exit the message dialog.

5.7.1.2 Making a Drive Unconfigured Bad, Unconfigured Good, or JBOD

When you force a drive offline, it enters the *Unconfigured Bad* state.

When you power down a controller and insert a new physical drive, if the inserted drive does not contain valid DDF metadata, the drive status is Just a Bunch of Disks (*JBOD*) when you power the system again. A new drive in the JBOD drive state is exposed to the host operating system as a stand-alone drive. You cannot use the JBOD drives to create a RAID configuration, because they do not have valid DDF records. You must first convert the drives into *Unconfigured Good*.

If a drive contains valid DDF metadata, its drive state is *Unconfigured Good*.

A drive must be in *Unconfigured Good* status before you can use it as a hot spare or use it as a member of a virtual drive. Follow these steps to change the status of an Unconfigured Bad, or Unconfigured Good, or JBOD drive.

1. Open the pop-up drive operations menu, highlight **Make Unconfigured Good**, **Make Unconfigured Bad**, or **Make JBOD**, and press Enter.
2. Highlight **Go**, which appears beneath **Operation**, and press Enter.

   ATTENTION If you have selected the Make Unconfigured Good operation, and if the JBOD that you have selected has an operating system or a file system on it, a warning message appears indicating that the JBOD has an operating system or a file system and any data on it would be lost if you proceed with the conversion. If you want to proceed, highlight **Confirm** and press the spacebar, then highlight **Yes** and press Enter. Otherwise, highlight **No** and press Enter to return to the previous screen. To run this operation on a different drive, press Esc to return to the **Drive Selection** menu and select another drive.

   A message appears indicating that the operation was successful.
3. Highlight **OK** on the success message and press Enter.

   NOTE To refresh the status of the drive displayed in the dialog, exit back to the **Main Menu**, then re-enter the **Drive Management** dialog.

5.7.1.3 Enabling Security on JBOD

If you have SED-enable JBOD that meets the perquisites mentioned in Managing Configurations, you can enable security on it. Follow these steps:

1. Open the pop-up drive operations menu, highlight **Enable Security on JBOD** and press Enter.
2. Highlight **Go**, which appears beneath **Operation**, and press Enter.

   A success message appears.
3. Highlight **OK** and press Enter.

5.7.1.4 Replacing a Drive

You might want to replace a drive that is a member of a redundant virtual drive connected to the controller if the drive shows signs of failing. Before you start this operation, be sure that an available Unconfigured Good replacement drive is available. The replacement drive must have at least as much capacity as the drive you are replacing.

Follow these steps to replace a drive.

1. Open the pop-up drive operations menu, highlight **Replace Drive**, and press Enter.
2. Highlight **Go**, which appears beneath **Operation**, and press Enter.

   The following dialog appears.

Figure 136 Replace Drive Window

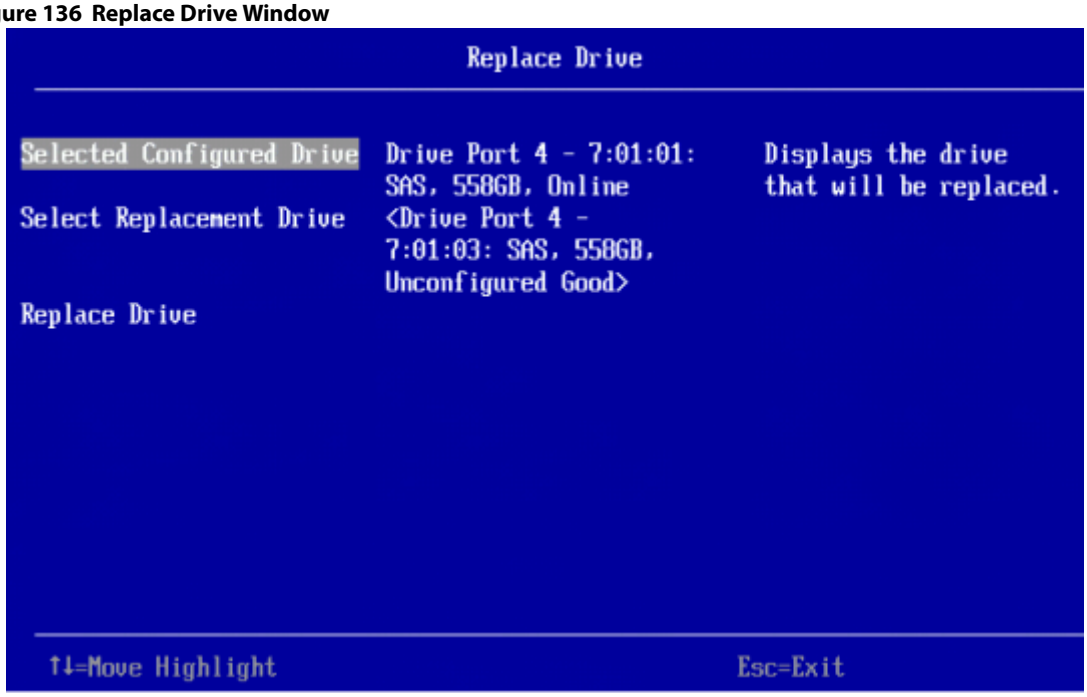

3. Highlight **Select Replacement Drive** and press Enter.

   A pop-up list of available replacement drives appears. In this example, only one replacement drive is available.
4. Select the replacement drive and press Enter.
5. Highlight **Replace Drive** and press Enter.

   A success message appears, and the replacement process begins as the data on the drive is rebuilt on the replacement drive.
6. Click **OK**.

   You are returned to the **Drive Management** menu. The status of the drive changes from **Online** to **Replacing**. You can perform other tasks in the HII utility while the replacement operation runs.

5.7.1.5 Placing a Drive Offline

Perform these steps to force a physical drive offline. If you perform this operation on a good drive that is part of a redundant virtual drive with a hot spare, the drive rebuilds to the hot spare drive. The drive you force offline goes into the Unconfigured Bad state.

1. Open the pop-up drive operations menu, highlight **Place Drive Offline**, and press Enter.
2. Highlight **Go**, which appears beneath **Operation**, and press Enter.

   The Place Drive Offline message appears.
3. Highlight **Confirm**, and press the spacebar to confirm the operation.
4. Highlight **Yes**, and press Enter.

   The selected drive is forced offline.

5.7.1.6 Placing a Drive Online

Perform these steps to force a selected member drive of a virtual drive online after it been forced offline.

1. Open the pop-up drive operations menu, highlight **Place Drive Online**, and press Enter.
2. Highlight **Go** and press Enter.

   The **Place Drive Online** warning appears.

ATTENTION Forcing a drive online that is part of a redundant array is *not* recommended.

3. Highlight **Confirm** and press the spacebar to confirm the operation.
4. Highlight **Yes** and press Enter.
   A message appears indicating that the action has been completed.
5. Highlight **Yes** and press Enter to return to the previous dialog.
   The drive is now online.

5.7.1.7 Marking a Drive Missing

Perform the following steps to mark a drive missing.

NOTE To set a drive that is part of an array as missing, you must first set it as offline. After the drive is set to offline, you can then mark the drive as missing.

1. Open the pop-up drive operations menu, highlight **Mark Drive as Missing**, and press Enter.
2. Highlight **Go** and press Enter.
   A warning message appears.
3. Highlight **Confirm** and press the space bar to confirm the operation.
4. Highlight **Yes** and press Enter.
   A message appears indicating that the action has been completed.
5. Highlight **OK** and press Enter to return to the previous dialog.
   The drive is marked as missing.

5.7.1.8 Replacing a Missing Drive

Perform the following steps to replace the drive that is marked as missing.

1. Open the pop-up drive operations menu, highlight **Replace Missing Drive**, and press Enter.
2. Highlight **Go** and press Enter.
   A warning message appears.
3. Highlight **Confirm** and press the space bar to confirm the operation.
4. Highlight **Yes** and press Enter.
   A message appears indicating that the action has been completed.
5. Highlight **OK** and press Enter to return to the previous dialog.
   The drive that was marked as missing is replaced.

5.7.1.9 Assigning a Global Hot Spare Drive

Global hot spare drives provide protection to redundant virtual drives on the controller. If you select an Unconfigured Good drive, you have the option to assign it as a global hot spare drive. Perform these steps to assign a global hot spare.

1. Open the pop-up drive operations menu, highlight **Assign Hot Spare Drive**, and press Enter.
2. Highlight **Go**, which appears beneath **Operation**, and press Enter.
   The hot spare selection dialog appears.

3. Highlight **Assign Global Hot Spare Drive** and press Enter.
   The status of the selected drive changes to hot spare.

NOTE To refresh the status of the drive displayed in the dialog, exit back to the **Main Menu**, then re-enter the **Drive Management** dialog.

5.7.1.10 Assigning a Dedicated Hot Spare Drive

Dedicated hot spare drives provide protection to one or more specified redundant virtual drives on the controller. If you select an Unconfigured Good drive, you have the option to assign it as a dedicated spare drive. Perform these steps to assign a dedicated hot spare.

1. Open the pop-up drive operations menu, highlight **Assign Dedicated Spare Drive**, and press Enter.
2. Highlight **Go**, which appears beneath **Operation**, and press Enter.
   The following dialog appears.

Figure 137 Associate Virtual Drives to the Dedicated Hot Spare Drive

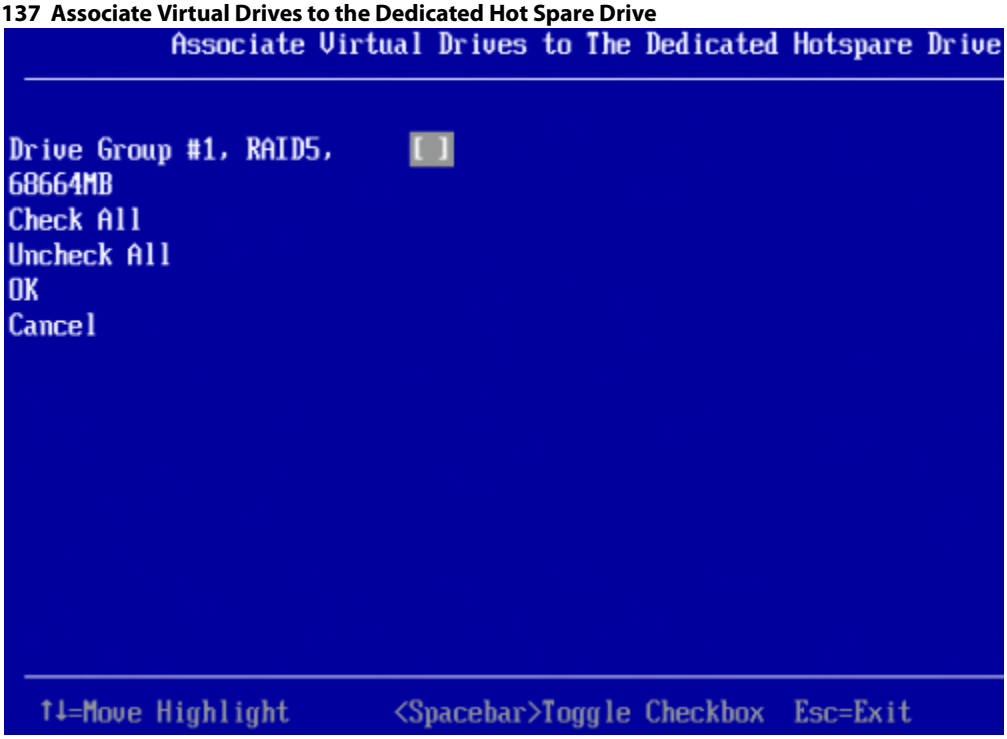

The preceding figure lists a single entry for each existing drive group. If you create a partial virtual drive on the same drive group, you can view a single entry with the cumulative size.

3. Select the drive groups to which this hot spare drive is dedicated, by highlighting each drive group and by pressing the spacebar.
   Alternatively, use the **Check All** or **Uncheck All** commands to select or deselect all of the drive groups.
4. When your selection is complete, highlight **OK**, and press Enter.
   When you return to the previous dialog, the status of the selected drive changes to hot spare.

NOTE To refresh the status of the drive displayed in the dialog, exit back to the **Main Menu** and then re-enter the **Drive Management** dialog.

5.7.1.11 Unassigning a Hot Spare Drive

If the currently selected drive is a hot spare drive, you can unassign it and return it to Unconfigured Good status.

Perform these steps to unassign a hot spare drive.

ATTENTION If you unassign a global hot spare drive or a dedicated hot spare drive, you reduce the protection level of the data on the VDs.

1. Open the pop-up drive operations menu, highlight **Unassign Hot Spare Drive**, and press Enter.
2. Highlight **Go**, which appears beneath **Operation**, and press Enter.
   The **Unassign Hotspare Drive** warning message appears.
3. Highlight **Confirm** and press the spacebar to confirm the operation.
4. Highlight **Yes** and press Enter.
   A confirmation message appears.
5. Click **OK** to return to the **Drive Management** menu.
   The drive that was formerly a hot spare now appears as Unconfigured Good.

NOTE To refresh the status of the drive displayed in the dialog, exit back to the **Main Menu** and then re-enter the **Drive Management** dialog.

5.7.1.12 Initializing or Erasing a Drive

Follow these steps to initialize or erase the currently selected drive. An initialize operation fills the drive with zeroes. An erase operation initializes the drive with a pattern of zeros and ones.

ATTENTION All data on the drive is lost when you initialize it or erase it. Before starting these operations, back up any data that you want to keep.

1. Open the pop-up drive operations menu, highlight **Initialize Drive** or **Erase Drive**, and press Enter.
2. If you select **Drive Erase**, highlight the **Erase Mode** field and press Enter.
3. Select **Simple**, **Normal**, or **Thorough** from the pop-up menu and press Enter.
4. Highlight **Go** and press Enter.
   The **Initialize Virtual Drive** message appears. (The message is similar to that of erasing a drive.)
5. Highlight **Confirm** and press the spacebar to confirm the operation.
6. Highlight **Yes** and press Enter.
   A message appears indicating that the initialization or erase operation has started.
7. Highlight **Yes** and press Enter to return to the previous window.
   This dialog displays a progress indicator that shows the percentage completion of the operation. It also displays a `Stop` command, as shown in the following figure.

Figure 138 Initialize Progress Indicator

8. To stop the initialization or erase process, highlight **Stop** and press Enter.

NOTE The progress indicator refreshes automatically.

5.7.1.13 Rebuilding a Drive

The manual Rebuild option is available only under certain conditions, as described here. If a hot spare drive is available, a rebuild starts automatically if a physical drive in a redundant array fails or is forced offline. If the Emergency Spare controller property is set to **Unconfigured Good** or **Unconfigured Good** and **Global Hotspare**, HII firmware automatically uses an Unconfigured Good drive to rebuild a failed or offline drive if no hot spares are available.

The manual Rebuild option is available only if a member drive of a virtual drive fails, there are no available hot spare drives, and the Emergency Spare controller property is set to **None**.

Follow these steps to start a manual Rebuild operation on an Unconfigured Good drive.

1. Open the pop-up drive operations menu, highlight **Rebuild**, and press Enter.
2. Highlight **Go** and press Enter.

   A progress indicator shows the percentage completion of the rebuild operation. This indicator refreshes automatically, and the **Rebuild Drive Success** message appears.

5.7.1.14 Securely Erasing a Drive

Perform these steps to securely erase the currently selected FDE-capable drive. This option is available only if the controller supports security and if security is configured.

ATTENTION All data on the drive is lost when you erase it. Before starting these operations, back up any data that you want to keep.

Perform these steps to securely erase an FDE-capable drive:

1. Open the pop-up drive operations menu, highlight **Secure Erase**, and press Enter.

2. Highlight **Go** and press Enter.
   The **Secure Erase** warning message appears.
3. Highlight **Confirm** and press the spacebar to confirm the operation.
4. Highlight **Yes** and press Enter.
   A message appears indicating that the secure erase operation has started.
5. Highlight **Yes** and press **Enter** to return to the previous dialog.
   This dialog now displays a progress bar and a `Stop` command.
6. To stop the secure erase process, highlight **Stop**, and press Enter.

NOTE A progress indicator shows the percentage completion of the operation. This indicator refreshes automatically.

5.7.1.15 Removing a Physical Drive

Perform these steps to remove a physical drive:

1. Open the pop-up drive operations menu, highlight **Prepare for Removal**, and press Enter.
2. Highlight **Go** and press Enter.
   A warning message appears.
3. Highlight **Confirm** and press the spacebar to confirm the operation.
4. Highlight **Yes** and press Enter.
   A message appears indicating that the action has been completed.
5. Highlight **Yes** and press **Enter** to return to the previous dialog.
   The drive is removed.

5.7.2 Viewing Advanced Drive Properties

The following dialog appears when you select **Advanced** on the **Drive Management** menu. The property information in this dialog is view-only, and cannot be modified.

Figure 139 Advanced Drive Properties 1

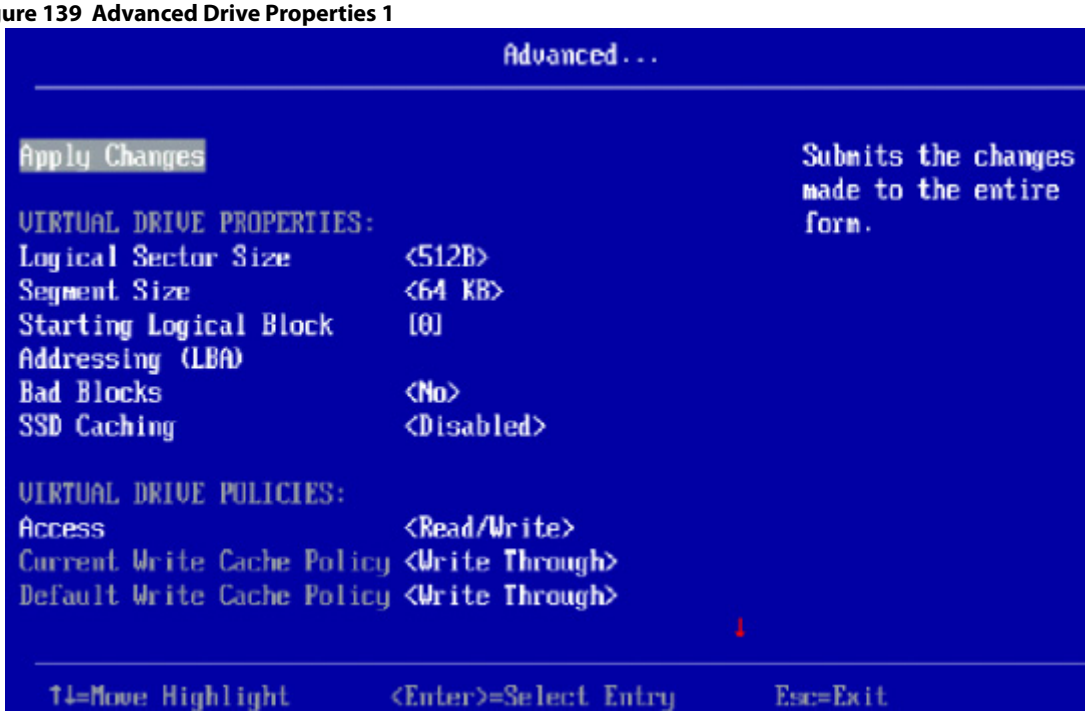

The small red arrow at the bottom of the dialog indicates that you can scroll down to view more physical drive properties.

NOTE The red arrow appears when there is too much information to display in one dialog. The amount of information that can be displayed in one dialog depends on the capabilities of the HII browser.

The following table describes all of the entries listed on the **Advanced Drive Properties** dialog.

Table 34 Advanced Drive Properties

Property	Description
Certified	Indicates whether the selected drive is vendor-certified. In some configurations you can only use certified drives to create configurations.
Logical Sector Size	The logical sector size of this drive. The possible options are **4 KB** or **512 B**.
Physical Sector Size	The physical sector size of this drive. The possible options are **4 KB** or **512 B**.
SMART Status	Indicates whether the Self-Monitoring Analysis and Reporting Technology (SMART) feature is enabled or disabled on the drive. The SMART feature monitors the internal performance of all motors, heads, and drive electronics to detect predictable drive failures.
Revision	The firmware revision level of the drive.
Connected Port	The port on which the drive is connected.
Media Errors	The number of physical errors detected on the disk media.
Predicted Fail Count	A property indicating the number of errors that have been detected on the disk media.
SAS Address	The World Wide Name (WWN) for the drive.
Emergency Spare	Indicates whether the drive is commissioned as an emergency spare.
Commissioned Hot Spare	Indicates if any hot spare drive (dedicated, global, or emergency) has actually been commissioned.
Cache Setting	Indicates if the drive cache is enabled or disabled.
Available Size (GB)	The available size of the drive, in GB.

Table 34 Advanced Drive Properties (Continued)

Property	Description
Used Space	The configured space of the drive, in GB.
Disk Protocol	Indicates whether the drive uses SAS or SATA protocol.
Negotiated Drive Transfer Speed	The negotiated link speed for data transfer to and from the drive.
Number of Connections	The number of connection on the drive. SAS drives have two ports.
FDE Capable	Indicates whether the drive is capable of encryption.
Secured	Indicates whether the drive is secured.

5.8 Managing Hardware Components

When you select **Hardware Components** on the **Main Menu**, the **Hardware Components** menu appears, as shown in the following figure.

Figure 140 Hardware Components Menu

The preceding figure lists the status of the temperature sensors, fans, power supplies, and other hardware components (such as batteries) installed in the system.

Select **Advanced** and press Enter to view more detailed information about the installed hardware components. The following dialog appears.

Figure 141 Advanced Hardware Components Menu

Select **Battery Management** or **Enclosure Management** to view more detailed information.

5.8.1 Managing Batteries

The following dialog appears when you select **Battery Management** on the **Advanced Hardware Components** menu.

Figure 142 Battery Management

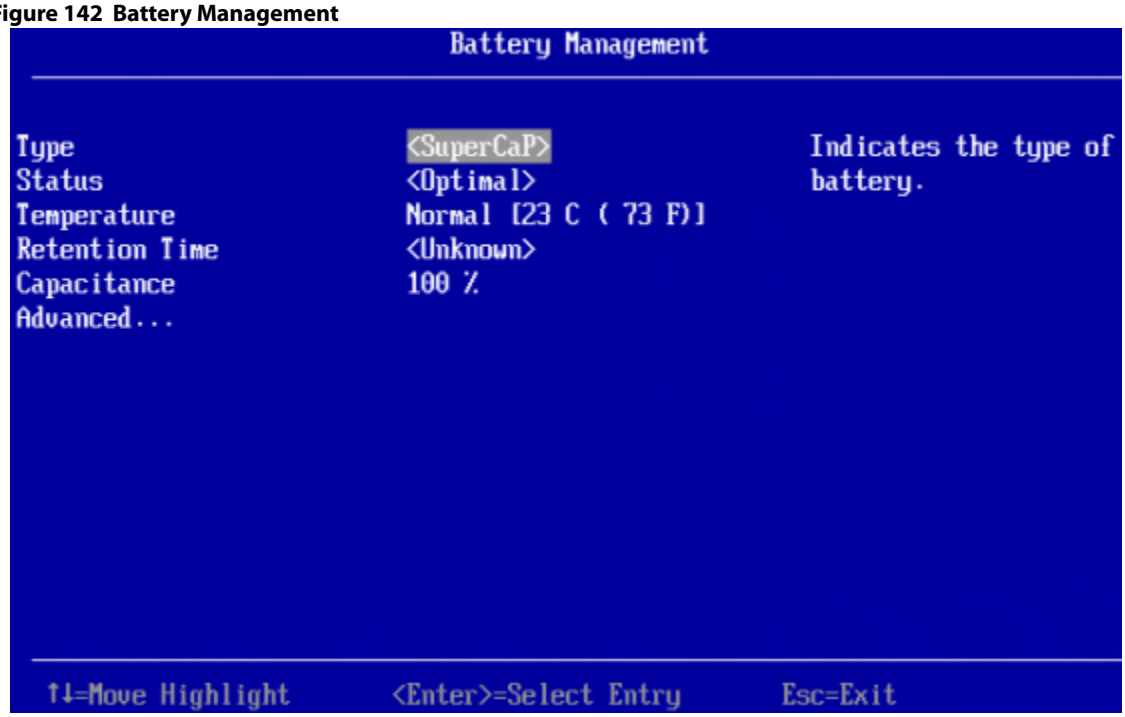

The following table describes the basic battery properties.

Table 35 Basic Battery Management Properties

Property	Description
Type	Type of the battery, such as Super Cap.
Status	Current status of the battery, such as **Optimal**. The battery status field has six states. If the battery operation is normal, the state is Optimal. ■ Optimal ■ Missing ■ Failed ■ Degraded ■ Degraded [Needs Attention] ■ Unknown
Temperature	Indicates the current temperature of the battery. Also indicates whether the current temperature of the battery is normal or high.
Retention Time	The number of hours the battery can support with the capacity it now has. The possible values are **48+ hours**, **Unknown**, or an exact number of hours between 1 and 48.
Capacitance	Available capacitance of the battery, stated as a percentage.

To view advanced battery properties, highlight **Advanced** and press Enter. The following dialog appears.

Figure 143 Advanced Battery Management

The small red arrow at the bottom of the dialog indicates that you can scroll down to view more Advanced Battery Management properties.

NOTE The red arrow appears when there is too much information to display in one dialog. The amount of information that can be displayed in one dialog depends on the capabilities of the HII browser.

The following table describes the advanced battery properties and the other options on this dialog. Properties marked with an asterisk are user-selectable. All other properties are view only.

Table 36 Advanced Battery Management Properties

Property	Description
Start Manual Learn Cycle*	Highlight this field and press Enter to start a manual battery learn cycle.
Set Automatic Learn Cycle Properties*	Highlight this field and press Enter to set the properties for an automatic battery learn cycle.
Manufacturer	Manufacturer of the battery.
Serial Number	Serial number of the battery.
Date of Manufacture	Manufacturing date of the battery.
Firmware Version	Firmware version of the battery.
Status	Status of the battery. If the status is Learning, Degraded, or Failed, a reason is listed for the status.
Voltage	Voltage level of the battery, in mV. Also indicates if the current battery voltage is normal or low.
Current	Current of the battery, in mA.
Design Capacity	Theoretical capacity of the battery.
Full Capacity	Full charge capacity of the battery.
Remaining Capacity	Remaining capacity of the battery.
Auto-learn Mode	Indicates whether auto-learn mode is enabled or disabled. A learn cycle is a battery calibration operation that the controller performed periodically to determine the battery condition. This operation cannot be disabled.
Next Learn Cycle Time	Date and hour of the next scheduled learn cycle.

5.8.1.1 Setting Automatic Learn Cycle Properties

The **Set Automatic Learn Cycle Properties** dialog appears when you select **Set Automatic Learn Cycle Properties** on the **Advanced Battery Management** dialog.

The small red arrow at the bottom of the dialog indicates that you can scroll down to view more options.

NOTE The red arrow appears when there is too much information to display in one dialog. The amount of information that can be displayed in one dialog depends on the capabilities of the HII browser.

To generate an event as a reminder to start a learn cycle manually, highlight the field next to **Generate an event...**, and press the spacebar.

To enable or disable automatic learn cycle mode, highlight the field next to **Learn Cycle**, press Enter, and make a selection from the pop-up menu.

The **Day**, **Time**, **No. of Days**, and **No. of Hours** fields are also user-selectable through popup menus. The **Next Learn Cycle Time** field shows the time of the next learn cycle.

Use the **Apply**, **OK**, and **Cancel** fields at the bottom of the selections (not visible in this figure) to apply, confirm or cancel any changes to the learn cycle options.

5.8.2 Managing Enclosures

To manage enclosures and view enclosure properties, select **Enclosure Management** from the **Advanced Hardware Components** menu.

The **Enclosure Management** dialog shows the Vendor ID, Enclosure ID, Enclosure Model, Enclosure Location, Product Revision Level, Number of slots for the selected enclosure.

Figure 144 Enclosure Management

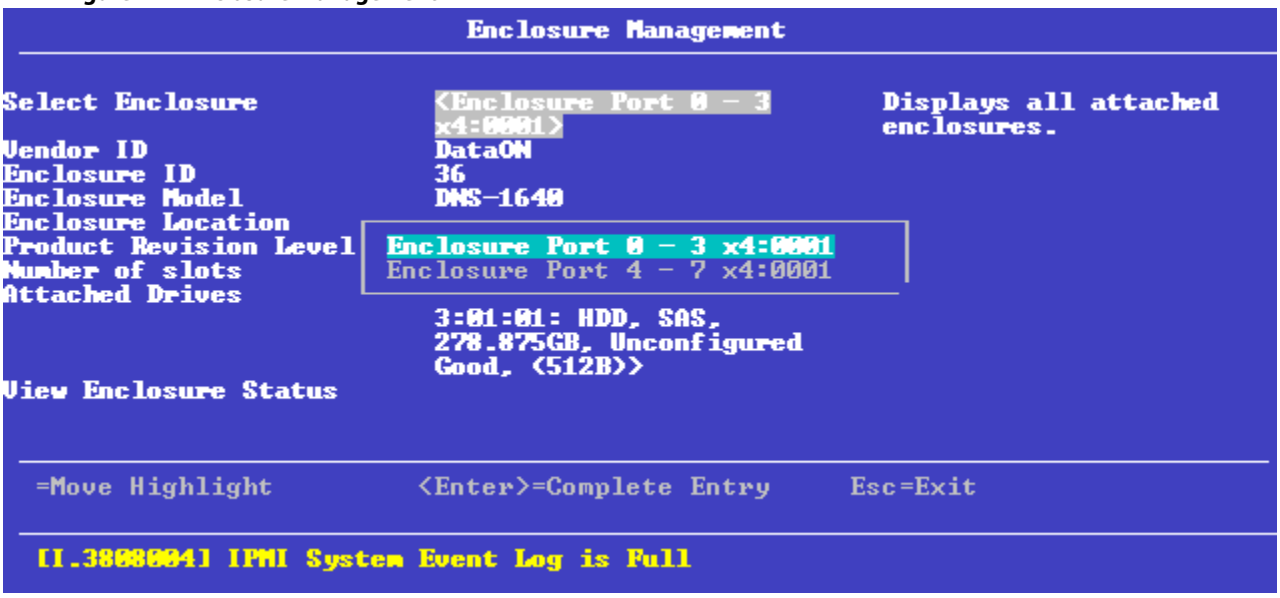

To select a different enclosure, highlight the **Select Enclosure** field, press Enter, and select the enclosure from the pop-up menu.

To view a pop-up menu of drives connected to the enclosure, highlight the **Attached Drives** field and press Enter.

To view more information about the enclosure status, highlight **View Enclosure Status** and press Enter. The following dialog appears.

Figure 145 View Enclosure Status

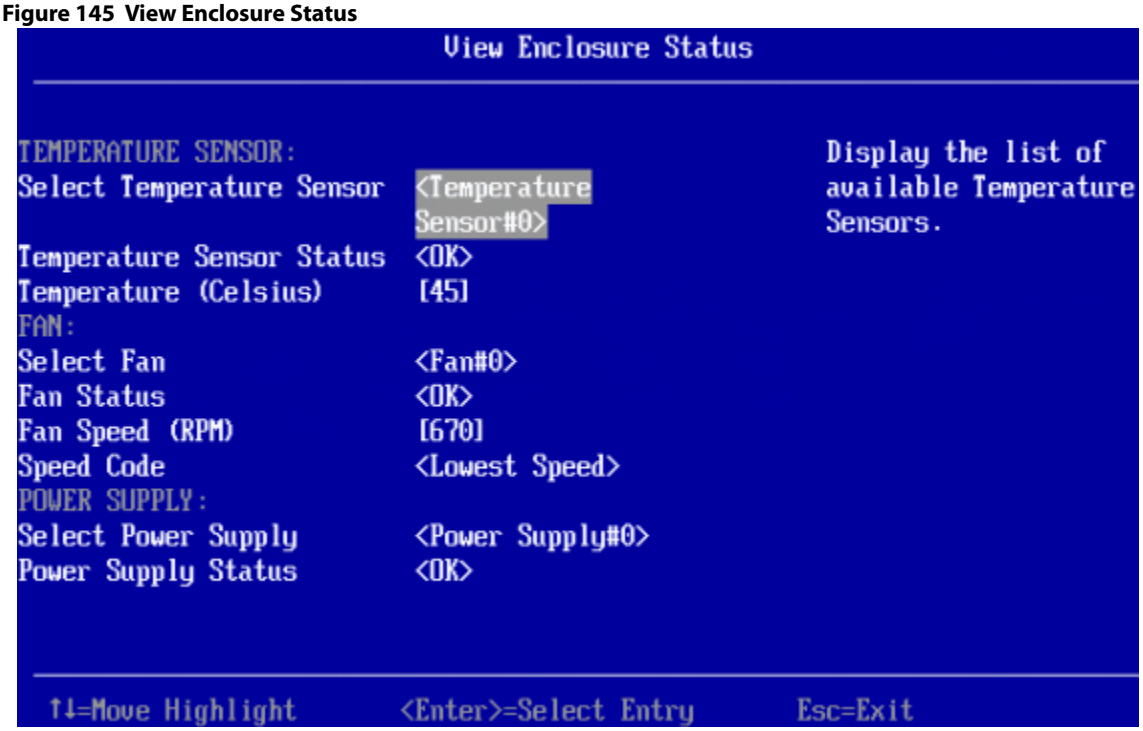

The **View Enclosure Status** dialog shows information about the temperature sensors, fans, and power supplies installed in the selected enclosure. To view a selectable pop-up menu of all of the installed sensors, fans, or power supplies, highlight the appropriate **Select** field, and press Enter.

Chapter 6: StorCLI

6.1 Overview

The Storage Command Line Interface (StorCLI) tool is the command line management software designed for the MegaRAID product line. The StorCLI tool is a command line interface that is designed to be easy to use, consistent, and easy to script. This chapter provides information on how to install and use the StorCLI tool and explains the various features of the StorCLI tool.

NOTE The legacy commands are deprecated from this guide.

6.2 Support for MegaCLI Commands

The MegaCLI commands can be executed on the StorCLI tool. A single binary is output for the StorCLI commands and its equivalent MegaCLI commands. See MegaCLI Commands to StorCLI Command Conversion for the information for conversion from MegaCLI commands to StorCLI commands.

6.3 Devices Supported by the StorCLI Tool

The StorCLI tool is designed to work with the MegaRAID product line. The StorCLI tool supports the following MegaRAID products.

- The 936*x* product line.
- MegaRAID SAS 9360-4i
- MegaRAID SAS 9360-8i
- MegaRAID SAS 9380-4i4e
- MegaRAID SAS 9380-8e
- MegaRAID SAS 9361 -8i
- MegaRAID SAS 9361-4i

6.4 Installation

The MegaRAID controllers can be used with the following operating systems for Intel® and AMD® 32-bit and 64-bit x86-based motherboards:

- Microsoft® Windows Server® 2008 R2
- Microsoft Windows® 7 (32 bit and 64 bit)
- Microsoft Windows 8.1
- Microsoft Windows 8.1 Update
- Red Hat® Enterprise Linux® 5.8 (32 bit and 64 bit)
- Red Hat Enterprise Linux 5.9
- Red Hat Enterprise Linux 6.1
- Red Hat Enterprise Linux 6.2 (32 bit and 64 bit)

- Red Hat Enterprise Linux 6.6
- Red Hat Enterprise Linux 6.7
- Red Hat Enterprise Linux 7.0 (32 bit and 64 bit)
- Red Hat Enterprise Linux 7.1 (32 bit and 64 bit)
- Unbreakable Enterprise Kernel Release 3 Update 2 for Oracle® Linux 6.4 (64 bit and later)
- Unbreakable Enterprise Kernel Release 3 Update 3
- Unbreakable Enterprise Kernel Release 3 Update 4
- Unbreakable Enterprise Kernel Release 3 Update 5
- Oracle Virtual Machine 3.3
- Oracle Linux 6.4
- Oracle Linux 7.0
- SuSE® Linux Enterprise Server 11 SP2 (32 bit and 64 bit) and SuSE Linux Enterprise Server 11 SP4 (32 bit and 64 bit)
- SuSE Linux Enterprise Server 10 SP4 (32 bit and 64 bit)
- SuSE Linux Enterprise Server 12
- SLES 11 SP3
- Fedora® Core Linux 15
- Fedora 18
- Fedora 20
- VMware® ESX™ 4.0
- VMware ESX 4.1 U2
- VMware ESXi 4.1 U2
- VMware ESXi 5.0 U1
- VMware ESXi 5.1 U3
- VMware ESXi 5.5 U2
- VMware 5.0 Update 2
- VMware 5.1 Update 1
- VMware OP
- VMware vSphere® 5.5 U1
- VMware vSphere 2015/ESXi 6.0
- Solaris®
- Solaris SPARC
- Solaris 11 Update 1 x86
- FreeBSD®
- FreeBSD 9.3
- EFI
- Citrix® XenServer® 6.1
- Ubuntu® 14.04
- Ubuntu 14.10
- Ubuntu 15.05
- Unreal Development Kit 2010
- CentOS™ 7.0

6.4.1 Installing the StorCLI Tool on Microsoft Windows Operating Systems

The Windows StorCLI binary is provided in a binary format, and no separate installation is required.

1. Copy the binary file from the CD or from the company website.

2. Place the binary file in the directory from which you want to run the Storage Command Line Interface, and run the tool.

NOTE The StorCLI tool must be run with the administrator privileges.

Because Windows PowerShell is not fully supported by the StorCLI tool, use either one of the following techniques to run commands in the StorCLI tool in Windows PowerShell:

— Enclose commands in double quotes. As an example,

```
storcli "/cx show"
```

— Launch the Command Prompt from within Windows PowerShell to run the StorCLI commands.

6.4.2 Installing the StorCLI Tool on Linux Operating Systems

To install the StorCLI tool on Linux operating systems, perform the following steps:

1. Unzip the StorCLI tool package.
2. To install the StorCLI RPM feature, run the `rpm -ivh <StorCLI-x.xx-x.noarch.rpm>` command.
3. To upgrade the StorCLI RPM feature, run the `rpm -Uvh <StorCLI-x.xx-x.noarch.rpm>` command.

6.4.3 Installing the StorCLI Tool on Ubuntu Operating Systems

To install the StorCLI tool on the Ubuntu operating systems, perform the following steps:

NOTE Run all the commands using the super user (sudo) login.

1. Run the `sudo dpkg -i storcli_1.0_all.deb` command to install the Debian® package.
2. Run the `dpkg -l | grep -i storcli` command to verify that the Debian package was installed successfully.
3. To uninstall the Debian package, run the `sudo dpkg -r storcli` command.

6.4.4 Installing the StorCLI Tool on VMware Operating Systems

To install the StorCLI tool on VMware operating systems, run the following from the command line:

```
esxcli software vib install -v=<path-to-vib-package>
```

Example:

```
esxcli software vib install
-v=/vmfs/volumes/datastore1/StorCliMN/vmware-esx-StorCli-1.01.04.vib
```

6.4.5 Installing the StorCLI Tool on FreeBSD Operating Systems

The FreeBSD StorCLI binary is provided in a binary format, and no separate installation is required.

1. Copy the binary file from the CD or from the company website.
2. Place the binary file in the directory from which you want to run the Storage Command Line Interface, and run the tool.

6.4.6 Installing the StorCLI Tool on Microsoft EFI

The EFI StorCLI binary is provided in a binary format, and no separate installation is required.

1. Copy the binary file from the CD or from the company website.
2. Place the binary file in the directory from which you want to run the Storage Command Line Interface tool, and run the tool.

6.4.7 Installing the StorCLI Tool on Solaris Operating Systems

To install the StorCLI tool on Solaris operating systems, run the following command:

```
pkgadd -d Storcli.pkg
```

6.5 StorCLI Tool Command Syntax

This chapter describes the StorCLI command syntax and the valid values for each parameter in the general command syntax.

NOTE To get the output in JSON format, add `J` at the end of the command syntax. For example:

```
storcli /cx show <property1>|<property2> J
```

NOTE JSON format output is not supported in the EFI operating system. The EFI platform ignores the `J` when it is added at the end of the command syntax.

NOTE Background operations are blocked in the EFI and HII environments and these operations are resumed in the operating system environments.

The StorCLI tool syntax uses the following general format:

```
<[object identifier]> <verb> <[adverb | attributes | properties]> <[key=value]>
```

The StorCLI tool supports the object identifiers listed in the following table.

Table 37 Object Identifiers in the StorCli Command Syntax

Object Identifier	Description
No object identifier specified	If no object identifier exists, the command is a system command.
/c*x*	This object identifier is for controller *x*.
/c*x*/v*x*	This object identifier is for a virtual drive *x* on controller *x*.
/c*x*/vall	This object identifier is for all virtual drives on controller *x*.
/c*x*/e*x*	This object identifier is for an enclosure x on controller *x*.
/c*x*/eall	This object identifier is for all enclosures on controller *x*.
/c*x*/f*x*	This object identifier is for a foreign configuration x on controller *x*.
/c*x*/fall	This object identifier is for all foreign configurations on controller *x*.
/c*x*/e*x*/s*x*	This object identifier is for the drive is slot *x* on enclosure *x* on controller *x*.
/c*x*/s*x*	This object identifier represents the drives that are directly attached to controller *x*.

Table 37 Object Identifiers in the StorCli Command Syntax (Continued)

Object Identifier	Description
/cx/ex/sall	This object identifier is for all the drives on enclosure *x* on controller *x*.
/cx/dx	This object identifier is for the drive group *x* on enclosure *x* on controller *x*.
/cx/dall	This object identifier is for the all drive groups on enclosure *x* on controller *x*.
/cx/px	This object identifier is for a phy operation *x* on controller *x*.
/cx/pall	This object identifier is for all phy operations on controller *x*.
/cx/bbu	This object identifier is for a BBU *x* on controller *x*.
/cx/cv	This object identifier is for a cache vault *x* on controller *x*.

NOTE If enclosures are not used to connect physical drives to the controller, you do not specify the enclosure ID in the command.

The StorCLI tool supports the following verbs.

Table 38 Verbs in the StorCli Command Syntax

Verbs	Description
add	This verb adds virtual drives, JBODs, and so on to the object identifier.
del	This verb deletes a drive, value, or property of the object identifier.
set	This verb sets a value of the object identifier.
show	This verb shows the value and properties of the object identifier.
pause	This verb pauses an ongoing operation.
resume	This verb resumes paused operation.
compare	This verb compares an input value with a system value.
download	This verb downloads and flashes a file to the target.
start	This verb starts an operation.
flush	This verb flushes a controller cache or a drive cache.
stop	This verb stops an operation that is in progress. A stopped process cannot be resumed.
import	This verb imports the foreign configuration into the drive.
expand	This verb expands the size of the virtual drive.
insert	This verb replaces the configured drive that is identified as missing, and starts an automatic rebuild.
flasherase	This verb erases the flash memory on the controller.
transform	This verb downgrades the firmware memory on the controller.
restart	This verb restarts the controller without a system reboot.
apply	This verb applies the activation Key to a WarpDrive® card.

- <[adverb | attributes | properties]>
  Specifies what the verb modifies or displays.
- <[key=value]>
  Specifies a value, if a value is required by the command.

6.6 Working with the Storage Command Line Interface Tool

This section describes the commands supported by the Storage Command Line Interface Tool.

NOTE The Storage Command Line Interface Tool is not case sensitive.

ATTENTION The order in which you specify the command options should be the same as in this document; otherwise, the command will fail.

NOTE The Storage Command Line Interface Tool does not support the Snapshot feature.

6.6.1 System Commands

6.6.1.1 System Show Commands

The Storage Command Line Interface Tool supports the following system show commands:

```
storcli show
storcli show all
storcli show ctrlcount
storcli show help
storcli -v
```

The detailed description for each command follows.

storcli show

This command shows a summary of controller and controller-associated information for the system. The summary includes the number of controllers, the host name, the operating system information, and the overview of existing configuration.

storcli show all

This command shows the list of controllers and controller-associated information, information about the drives that need attention, and advanced software options.

storcli show ctrlcount

This command shows the number of controllers detected in the server.

storcli show help

This command shows help for all commands at the server level.

storcli -v

This command shows the version of the Storage Command Line Interface Tool.

6.6.2 Controller Commands

Controller commands provide information and perform actions related to the specified controller, such as the `/c0` controller. The Storage Command Line Interface Tool supports the controller commands described in this section.

6.6.2.1 Show and Set Controller Properties Commands

Table 39 Controller Commands Quick Reference Table

Commands	Value Range	Description
`show <properties>`	See Table 40	Shows specific controller properties.
`set <properties>`	See Table 40	Sets controller properties.
`show`	`all`: Shows all properties of the virtual drive. `freespace`: Shows the freespace in the controller. See Controller Show Commands.	Shows physical drive information.

This section provides command information to show and set controller properties.

NOTE You cannot set multiple properties with a single command.

storcli /*cx* show <*property*>

This command shows the current value of the specified property on the specified controller.

General example output:

```
Status Code = 0
Status = Success
Description = None
Controller: 0
Property_name = Property_value
```

You can show the following properties using the `storcli /cx show <property1>|<property2>` command.

```
storcli /cx show abortcconerror
storcli /cx show activityforlocate
storcli /cx show alarm
storcli /cx show backplane
storcli /cx show batterywarning
storcli /cx show bgirate
storcli /cx show bootwithpinnedcache
storcli /cx show cachebypass
storcli /cx show cacheflushint
storcli /cx show ccrate
storcli /cx show coercion
storcli /cx show consistencycheck|cc
storcli /cx show copyback
storcli /cx show directpdmapping
storcli /cx show dimmerswitch|ds
storcli /cx show eccbucketleakrate
storcli /cx show eccbucketsize
```

```
storcli /cx show eghs
storcli /cx show jbod
storcli /cx show loadbalancemode
storcli /c0 show largeiosupport
storcli /cx show maintainpdfailhistory
storcli /cx show migraterate
storcli /cx show ncq
storcli /cx show patrolread|pr
storcli /cx show perfmode
storcli /cx show pi
storcli /cx show prcorrectunconfiguredareas
storcli /cx show prrate
storcli /cx show rebuildrate
storcli /cx show rehostinfo
storcli /cx show restorehotspare
storcli /cx show safeid
storcli /cx show smartpollinterval
storcli /cx show spinupdelay
storcli /cx show spinupdrivecount
storcli /cx show time
storcli /cx show usefdeonlyencrypt
storcli /cx show badblocks
storcli /cx show wbsupport
storcli /cx show DPM
storcli /cx show SGPIOforce
storcli /cx show failpdonsmarterror
Storcli /cx show flushwriteverify
```

storcli /*cx* set <*property*> = <*value*>

General example output:

```
Status Code = 0
Status = Success
Description = None


Controller 0, new Property_name = Property_value
```

The following commands are examples of the properties that can be set using the `storcli /cx set <property>=<value>` command:

```
storcli /cx set  abortcconerror=<on|off>
storcli /cx set  termlog[=on|off|offthisboot]
storcli /cx set  activityforlocate=<on|off>
storcli /cx set  alarm=<on|off|silence>
storcli /cx set  batterywarning=<on|off>
storcli /cx set  bgirate=<value>
storcli /cx set  bootwithpinnedcache=<on|off>
storcli /cx set  cachebypass=<on|off>
storcli /cx set  cacheflushinterval=<value>
storcli /cx set  ccrate=<value>
storcli /cx set  coercion=<value>
storcli /cx set  consistencycheck|cc=[off|seq|conc][delay=value]
[starttime=yyyy/mm/dd hh] [excludevd=x-y,z]
storcli /cx set  copyback=<on|off> type=<smartssd|smarthdd|all>
storcli /cx set  directpdmapping=<on|off>
storcli /cx set  eccbucketleakrate=<value>
storcli /cx set  eccbucketsize=<value>
storcli /cx set  eghs [state=<on|off>][smarter=<on|off>][eug=<on|off>]
storcli /cx set  backplane [mode=<0-3>][expose=<on|off>]
storcli /cx set  dimmerswitch|ds=<on|off type=1|2|4>
storcli /cx set  foreignautoimport=<on|off>
storcli /cx set  jbod=<on|off>
storcli /cx set  loadbalancemode=<value>
storcli /cx set  maintainpdfailhistory=<on|off>
storcli /cx set  migraterate=<value>
storcli /cx set  ncq=<on|off>
storcli /cx set  patrolread|pr {=on mode=<auto|manual>}|{off}
storcli /cxvset  perfmode=<value>
storcli /cx set  pi [state=<on|off>][import=<on|off>]
storcli /cx set  prcorrectunconfiguredareas=<on|off>
storcli /cx set  prrate=<value>
storcli /cx set  rebuildrate=<value>
storcli /cx set  restorehotspare=<on|off>
storcli /cx set  smartpollinterval=<value>
storcli /cx set  spinupdelay=<value>
storcli /cx set  spinupdrivecount=<value>
```

```
storcli /cx set  stoponerror=<on|off>
storcli /cx set  usefdeonlyencrypt=<on|off>
storcli /cx set time=yyyymmdd hh:mm:ss|systemtime
storcli /cx set  usefdeonlyencrypt=<on|off>
storcli /cx set DPM=<on|off>
storcli /cx set  supportssdpatrolread=<on|off>
storcli /cx set  SGPIOforce=<on|off>
storcli /cx set immediateio=<on|off>
storcli /cx set driveactivityled=<on|off>
storcli /cx set sesmonitoring=[on|off]
storcli /cx set failpdonsmarterror=<on|off>
storcli /c0 set flushwriteverify=<on|off>
storcli /c0 set largeiosupport=on|off
```

The following table lists and describes the properties for the `show` and `set` commands.

Table 40 Properties for Show and Set Commands

Property Name	Set Command Range	Description
`abortcconerror`	`on|off`	Aborts consistency check when it detects an inconsistency.
`activityforlocate`	`on|off`	Enables/disables drive activity, drive activity locates function for systems without SGPIO/SES capabilities.
`alarm`	`on|off|silence` `silence`: Silences the alarm.	Enables/disables alarm on critical errors.
`batterywarning`	`on|off`	Enables/disables battery warnings.
`bgirate`	0 to 100	Sets background initialization rate in percentage.
`cachebypass`	`on|off`	Enables/disables the cache bypass performance improvement feature.
`cacheflushint`	0 to 255, default value `4`	Sets cache flush interval in seconds.
`ccrate`	0 to 100	Sets consistency check rate in percentage.
`coercion`	`0`: No coercion `1`: 128 MB `2`: 1 GB	Sets drive capacity in coercion mode.
`consistencycheck`	See Consistency Check.	See Consistency Check.

Table 40 Properties for Show and Set Commands (Continued)

Property Name	Set Command Range	Description
copyback	on|off type = smartssd|smarthdd|all smartssd: Copy back enabled for SSD drives. smarthdd: Copy back enabled for HDD drives. all: Copy back enabled for both ssd drives and HDD drives. Example: storcli /cx set copyback=on type=all	Enables/disables copy back for drive types.
directpdmapping	on|off	Enables/disables direct physical drive mapping. When enclosures are used, this feature is disabled; otherwise it should be enabled.
eccbucketleakrate	0 to 65535	Sets the leak rate of the single-bit bucket in minutes (one entry removed per leak-rate).
eccbucketsize	0 to 255	Sets the size of ECC single-bit-error bucket (logs event when full).
eghs state	on|off	Enables/disables the commissioning of otherwise incompatible global hot spare drives as Emergency Hot Spare (EHSP) drives.
eghs smarter	on|off	Enables/disables the commissioning of Emergency Hot Spare (EHSP) drives for Predictive Failure (PFA) events.
eghs eug	on|off	Enables/disables the commissioning of Unconfigured Good drives as Emergency Hot Spare (EHSP) drives.
backplane mode	0: Use autodetect logic of backplanes, such as SGPIO and I2C SEP using GPIO pins. 1: Disable autodetect SGPIO. 2: Disable I2C SEP autodetect. 3: Disable both the autodetects.	Configures enclosure detection on a non-SES/expander backplane.
backplane expose	on|off	Enables/disables device drivers to expose enclosure devices; for example, expanders, SEPs.
dimmerswitch|ds	See Dimmer Switch Commands.	See Dimmer Switch Commands.
foreignautoimport	on|off	Imports a foreign configuration automatically, at boot.
jbod	on|off	Enables/disables JBOD mode; by default, drives become system drives. Not supported by all controllers.Enables/disables JBOD mode; by default, drives become system drives.Not supported by all controllers. **NOTE** If you try to disable the JBOD mode, and if any of the JBOD has an operating system/file system, then the StorCLI tool displays a warning message indicating that the JBOD has an operating system or a file system on it and prompts you to use the force option to proceed with the disable operation.

Table 40 Properties for Show and Set Commands (Continued)

Property Name	Set Command Range	Description
loadbalancemode	on|off	Enables/disables automatic load balancing between SAS phys or ports in a wide port configuration.
largeiosupport	on|off	Sets the current settings on the controller for large I/O support.
maintainpdfailhistory	on|off	Maintains the physical drive fail history.
migraterate	0 to 100	Sets data migration rate in percentage.
patrolread|pr	See Patrol Read.	See Patrol Read.
perfmode	0: Tuned to provide best IOPS, currently applicable to non-FastPath 1: Tuned to provide least latency, currently applicable to non-FastPath	Performance tuning setting for the controller.
pi	on|off	Enables/disables data protection on the controller.
pi import	on|off	Enables/disables import data protection drives on the controller.
prcorrectunconfiguredareas	on|off	Correct media errors during PR by writing 0s to unconfigured areas of the disk.
prrate	0 to 100	Sets the patrol read rate of the virtual drives in percentage.
rebuildrate	0 to 100	Sets the rebuild rate of the drive in percentage.
reconrate	0 to 100	Sets the reconstruction rate for a drive, as a percentage.
restorehotspare	on|off	Becomes a hot spare on insertion of a failed drive.
smartpollinterval	0 to 65535	Set the time for polling of SMART errors, in seconds.
spinupdrivecount	0 to 255	Sets the number of drives that are spun up at a time.
spinupdelay	0 to 255	Sets the spin-up delay between a group of drives or a set of drives, in seconds.
stoponerror	on|off	Stops the MegaRAID BIOS during POST, if any errors are encountered.
time	Valid time in *yymmdd hh:mm:ss* format or systemtime	Sets the controller time to your input value or the system time (local time in 24-hour format).
usefdeonlyencrypt	on|off	Enables/disables FDE drive-based encryption.
DPM	on|off	Enables/disables drive performance monitoring
supportssdpatrolread	on|off	Enables/disables patrol read for SSD drives.
SGPIOforce	on|off	Forces the SGPIO status per port only for four drives; affects HPC controllers.
immediateio	on|off	Enables or disables Immediate I/O transactions.
driveactivityled	on|off	Activate or deactivate the Drive Activity LED.
sesmonitoring	on|off	Enables or disables SES monitoring.
failpdonsmarterror	on|off	Enables or disables the *Fail PD on SMARTer* property.
flushwriteverify	on|off	Enables or disables the Write Verify feature. This feature verifies if the data was written correctly to the cache before flushing the controller cache.

6.6.2.2 Controller Show Commands

The Storage Command Line Tool supports the following show commands:

```
storcli /cx show
storcli /cx show all [logfile[=filename]]
storcli /cx show freespace
```

The detailed description for each command follows.

storcli /*cx* show

This command shows the summary of the controller information. The summary includes basic controller information, foreign configurations, drive groups, virtual drives, physical drives, enclosures, and BBU information.

Input example:

```
storcli /c1 show
```

storcli /*cx* show all [logfile[=*filename*]]

The `show all` command shows all of the controller information, which includes basic controller information, bus information, controller status, advanced software options, controller policies, controller defaults, controller capabilities, scheduled tasks, miscellaneous properties, foreign configurations, drive groups, virtual drives, physical drives, enclosures, and BBU information.

If you use the `logfile` option in the command syntax, the logs are written to the specified file. If you do not specify a file name, then the logs are written to the `storsas.log` file. If you do not use the `logfile` option in the command syntax, the entire log output is printed to the console.

Input example:

```
storcli /c0 show all [logfile[=log.txt]]
```

NOTE The PCI information displayed as a part of `storcli /cx show` and `storcli /cx show all` commands is not applicable for the FreeBSD operating system. Hence, the PCI information fields are displayed as `N/A`.

storcli /*cx* show freespace

This command shows the usable free space in the controller.

Input example:

```
storcli /c0 show freespace
```

6.6.2.3 Controller Background Tasks Operation Commands

6.6.2.3.1 Rebuild Rate

```
storcli /cx set rebuildrate=<value>
storcli /cx show rebuildrate
```

The detailed description for each command follows.

storcli /*cx* set rebuildrate=<*value*>

This command sets the rebuild task rate of the specified controller. The input value is in percentage.

Input example:

```
storcli /c0 set rebuildrate=30
```

NOTE A high rebuild rate slows down I/O transaction processing.

storcli /*cx* show rebuildrate

This command shows the current rebuild task rate of the specified controller in percentage.

Input example:

```
storcli /c0 show rebuildrate
```

6.6.2.3.2 Patrol Read

The Storage Command Line Interface Tool supports the following patrol read commands:

```
storcli /cx resume patrolread
storcli /cx set patrolread ={{on mode=<auto|manual>}|{off}}
storcli /cx set patrolread [starttime=<yyyy/mm/dd hh>] [maxconcurrentpd=<value>]
[includessds=<on|off>] [uncfgareas=<on|off>]
storcli /cx set patrolread  delay=<value>
storcli /cx show patrolread
storcli /cx start patrolread
storcli /cx stop patrolread
storcli /cx pause patrolread
```

NOTE A patrol read operation is scheduled for all the physical drives of the controller.

The detailed description for each command follows.

storcli /*cx* resume patrolread

This command resumes a suspended patrol read operation.

Input example:

```
storcli /c0 resume patrolread
```

storcli /*cx* set patrolread {=on mode=<auto|manual>}|{off}

This command turns the patrol read scheduling on and sets the mode of the patrol read to automatic or manual.

Input example:

```
storcli /c0 set patrolread=on mode=manual
```

storcli /*cx* set patrolread [starttime=<*yyyy/mm/dd hh*>] [maxconcurrentpd=<*value*>] [includessds=<on|off>] [uncfgareas=on|off]

This command schedules a patrol read operation. You can use the following options for patrol read command operations.

Table 41 Set Patrol Read Input Options

Option	Value Range	Description
`starttime`	A valid date and hour in 24 hours format	Sets the start time in *yyyy/mm/dd hh* format.
`maxconcurrentpd`	Valid number of physical drives present	Sets the number of physical drives that can be patrol read at a single time.
`includessds`	—	Include SSDs in the patrol read operation.
`uncfgareas`	—	Include the areas not configured in the patrol read process.

NOTE Controller time is taken as a reference for scheduling a patrol read operation.

Input example:

```
storcli /c0 set patrolread=on starttime=2012/02/21 00
```

storcli /c*x* set patrolread [delay=<*value*>]

This command delays the scheduled patrol read in hours.

Input example:

```
storcli /c0 set patrolread delay=30
```

storcli /c*x* show patrolRead

This command shows the current state of the patrol read operation along with other details such as the **PR Mode**, **PR Execution Delay**, **PR iterations completed**, and **PR on SSD**. This command also shows the start time and the date when the patrol read operation started.

The values shown for the current state of the patrol read operation are **Ready**, **Active**, **Paused**, **Aborted**, **Stopped**, or **Unknown**.

If the state of the patrol read is active, a numeric value is shown along with the state which depicts the number of physical drives that have completed the patrol read operation. As an example, `Active 1` means that the one physical drive has completed the patrol read operation.

Input example:

```
storcli /c0 show patrolread
```

storcli /c*x* start patrolread

This command starts the patrol read operation. This command starts a patrol read immediately.

Input example:

```
storcli /c0 start patrolread
```

storcli /c*x* stop patrolread

This command stops a running patrol read operation.

Input example:

```
storcli /c0 stop patrolread
```

NOTE You cannot resume a stopped patrol read.

storcli /c*x* pause patrolread

This command pauses a running patrol read operation.

Input example:

```
storcli /c0 pause patrolread
```

NOTE You can run this command only when a patrol read operation is running on the controller.

6.6.2.3.3 Consistency Check

The Storage Command Line Interface Tool supports the following commands to schedule, perform, and view the status of a consistency check (CC) operation:

```
storcli /cx set consistencycheck|cc=[off|seq|conc][delay=value]
starttime=yyyy/mm/dd hh [excludevd=x-y,z]
```

```
storcli /cx show cc
```

```
storcli /cx show ccrate
```

The detailed description for each command follows.

storcli /*cx* set consistencycheck|cc=[off|seq|conc][delay=*value*] starttime=*yyyy/mm/dd hh* [excludevd=*x-y,z*]

This command schedules a consistency check (CC) operation. You can use the following options with the consistency check command.

Table 42 Set CC Input Options

Option	Value Range	Description
cc	seq: Sequential mode. conc: Concurrent mode. off: Turns off the consistency check.	Sets CC to either sequential mode, or concurrent mode, or turns off the CC. **NOTE** The concurrent mode slows I/O processing.
delay	-1 and any integer value.	Delay a scheduled consistency check. The value is in hours. A value of 0 makes the CC runs continuously with no delay (in a loop). **NOTE** Only scheduled consistency checks can be delayed.
starttime	A valid date and hour in 24-hours format.	Start time of a consistency check is *yyyy/mm/dd hh* format.
excludevd	The range should be less than the number of virtual drives.	Excludes virtual drives from the consistency checks. To exclude particular virtual drives, you can provide list of virtual drive names (V*x*,V*y* … format) or the range of virtual drives that you want to exclude from a consistency check (V*x*-V*y* format). If this option is not specified in the command, no virtual drives are excluded.

Input example:

```
storcli /c0 set CC=on starttime=2012/02/21 00 excludevd v0-v3
```

storcli /*cx* show cc

This command shows the consistency check schedule properties for a controller.

Input example:

```
storcli /c0 show cc
```

storcli /*cx* show ccrate

This command checks the status of a consistency check operation. The CC rate appears in percentage.

Input example:

```
storcli /c0 show ccrate
```

NOTE A high CC rate slows I/O processing.

6.6.2.4 Premium Feature Key Commands

The Storage Command Line Interface Tool supports the following commands for premium feature keys:

```
storcli /cx set advancedsoftwareoptions(aso) key=<value> [preview]
storcli /cx aso [transfertovault][rehostcomplete][deactivatetrialkey]
storcli /cx show safeid
```

The detailed description for the command follows.

storcli /*cx* set advancedsoftwareoptions(aso) key=<*value*> [preview]

This command activates advanced software options (ASO) for a controller. You can use the following options with the advanced software options command.

Table 43 Set Advanced Software Options Input Options

Option	Value Range	Description
key	40 alpha numeric characters.	Key to activate ASO on the controller. **NOTE** After they are activated, ASOs cannot be removed from the controller.
deactivatetrialkey	—	Deactivates the trial key applied on the specified controller.
rehostcomplete	—	Enables rehosting on the specified controller.
transfertovault	—	Transfers the ASO key to the vault and disables the ASO.

Input example:

```
storcli /c0 set Aso key=LSI0000
```

storcli /*cx* show safeid

This command shows the Safe ID of the specified controller.

Input example:

```
storcli /c0 show safeid
```

6.6.2.5 Controller Security Commands

The Storage Command Line Interface Tool supports the following controller security commands:

```
storcli /cx compare securitykey=ssssss
storcli /cx delete securitykey
storcli /cx set securitykey keyid=kkkk
storcli /cx set securitykey=sssss [passphrase=sssss][keyid=sssss]
storcli /cx set securitykey=sssss  oldsecuritykey=ssss [passphrase=sssss]
[keyid=sssss]
```

The detailed description for each command follows.

storcli /*cx* show securitykey keyid

This command shows the security key on the controller.

Input example:

```
storcli /c0 show securityKey keyid
```

storcli /c*x* compare securitykey=*ssssss*

This command compares and verifies the security key of the controller.

storcli /c*x* delete securitykey

This command deletes the security key of the controller.

Input example:

```
storcli /c0 delete securitykey
```

storcli /c*x* set securitykey keyId=*kkkk*

This command sets the key ID for the controller. The key ID is unique for every controller.

storcli /c*x* set securitykey=sssss [passphrase=sssss][keyid=*sssss*]

This command sets the security key for the controller. You can use the following options with the set security key command.

Table 44 Set Security Key Input Options

Option	Value Range	Description
passphrase	Should have a combination of numbers, upper case letters, lower case letters and special characters. Minimum of 8 characters and maximum of 32 characters.	String that is linked to the controller and is used in the next bootup to encrypt the lock key. If the passphrase is not set, the controller generates it by default.
keyid	—	Unique ID set for different controllers to help you specify a passphrase to a specific controller.

Input example:

```
storcli /c0 set securitykey=Lsi@12345 passphrase=Lsi@123456 keyid=1
```

storcli /c*x* set securitykey=*sssss* oldsecuritykey=*ssss* [passphrase=*sssss*][keyid=*sssss*]

This command changes the security key for the controller.

Input example:

```
storcli /c0 set securitykey=Lsi@12345 oldsecuritykey=pass123 passphrase=Lsi@123456
keyid=1
```

6.6.2.6 Flashing Controller Firmware Command

NOTE The Flashing Controller Firmware command is not supported in Embedded MegaRAID.

The following command is used to flash the controller firmware.

storcli /c*x* download file=*filepath* [fwtype=<*value*>] [nosigchk] [noverchk] [resetnow]

This command flashes the firmware with the ROM file to the specified adapter from the given file location (*filepath* is the absolute file path). See Online Firmware Upgrade Support for limitations.

You can use the following options in the table to flash the firmware:

Table 45 Flashing Controller Firmware Input Options

Option	Value Range	Description
`nosigchk`	—	The application flashes the firmware even if the check word on the file does not match the required check word for the controller. **NOTE** You can damage the controller if a corrupted image is flashed using this option.
`noverchk`	—	The application flashes the controller firmware without checking the version of the firmware image.
`fwtype`	`0`: Application `1`: TMMC `2`: GC-Enhanced	The firmware type to be downloaded. The application downloads the firmware for the controller. The TMMC downloads the firmware for the TMMC battery only. Default is `0` (application).
`resetnow`	—	Invokes online firmware update on the controller; you do not need to reboot the controller to make the update effective. **NOTE** The `resetnow` option is not supported in the UEFI mode.

6.6.2.7 Controller Cache Command

The following command flushes the controller cache.

storcli /c*x* flush|flushcache

This command flushes the controller cache.

Input example:

```
storcli /c0 flushcache
```

6.6.2.8 Controller Configuration Commands

The following command works with the controller configuration.

storcli /c*x* set config file=*file_name*

This command saves the controller configuration and its properties to the specified file.

NOTE You cannot load a saved configuration over an existing configuration when there are existing virtual drives. To load a saved configuration, you must first clear the existing configuration on the target controller.

Input example:

```
storcli /c0 set config file= log.txt
```

storcli /c*x* get config file=*file_name*

This command obtains the controller configuration and its properties from the specified file.

Input example:

```
storcli /c0 get config file= log.txt
```

6.6.3 Drive Commands

This section describes the drive commands, which provide information and perform actions related to physical drives. The following table describes frequently used virtual drive commands.

Table 46 Physical Drives Commands Quick Reference Table

Commands	Value Range	Description
`set`	`missing`: Sets the drive status as missing. `good`: Sets the drive status to unconfigured good. `offline`: Sets the drive status to offline. `online`: Sets the drive status to online.	Sets physical drive properties.
`show`	`all`: shows all properties of the physical drive. See Drive Show Commands.	Shows virtual drive information.

6.6.3.1 Drive Show Commands

The Storage Command Line Interface Tool supports the following drive show commands:

```
storcli /cx[/ex]/sx show
storcli /cx[/eall]/sall show
storcli /cx[/ex]/sx|sall show all
storcli /cx/[ex]/sx show smart
```

NOTE If enclosures are used to connect physical drives to the controller, specify the enclosure ID in the command. If no enclosures are used, you must specify the controller ID and slot ID.

The detailed description for each command follows.

storcli /*cx*[/*ex*]/*sx* show

This command shows the summary of the physical drive for a specified slot in the controller.

Input example:

```
storcli /c0/e0/s4 show
```

storcli /*cx*[/eall]/sall show

This command shows the summary information for all the enclosures and physical drives connected to the controller.

Input example:

```
storcli /c0/eall/sall show
```

storcli /*cx*[/*ex*]/*sx*|sall show all

This command shows all information of a physical drive for the specified slot in the controller. If you use the `all` option, the command shows information for all slots on the controller. *x* stands for a number, a list of numbers, a range of numbers, or all numbers.

This command also shows the NCQ (Native Command Queuing) status (**Enabled**, **Disabled**, or **N/A**) which is applicable only to SATA drives. If the controller to which the SATA drive is connected supports NCQ and NCQ is enabled on the SATA drive, the status is shown as **Enabled**; otherwise it is shown as **Disabled**. If NCQ is not a supported drive operation on the controller, the status is shown as **N/A**.

Input examples:

```
storcli /c0/e3/s0-3 show all
storcli /c0/e35/sall show all
```

NOTE The `storcli /cx/sx show all` command shows tape drives information.

storcli /*cx*/[*ex*]/*sx* show smart

This command displays the SMART information of a SATA drive.

Input example:

```
storcli /c0/e5/s1 show smart
```

6.6.3.2 Missing Drives Commands

The Storage Command Line Interface Tool supports the following commands to mark and replace missing physical drives:

```
storcli /cx[/ex]/sx insert dg=A array=B row=C
storcli /cx[/ex]/sx set missing
storcli /cx[/ex]/sx set offline
storcli /cx/dall
```

The detailed description for each command follows.

storcli /*cx*[/*ex*]/*sx* insert dg=A array=*B* row=*C*

This command replaces the configured drive that is identified as missing, and then starts an automatic rebuild.

Input example:

```
storcli /c0/e25/s3 insert dg=0 array=2 row=1
```

storcli /*cx*[/*ex*]/*sx* set missing

This command marks a drive as missing.

Input example:

```
storcli /c0/s4 set missing
```

storcli /cx/dall

This command is used to find the missing drives.

storcli /*cx*[/*ex*]/*sx* set offline

This command marks the drive in an array as offline.

NOTE To set a drive that is part of an array as *missing*, first set it as *offline*. After the drive is set to *offline*, you can then set the drive to *missing*.

6.6.3.3 Set Drive State Commands

The Storage Command Line Interface Tool supports the following commands to set the status of physical drives:

```
storcli /cx[/ex]/sx set jbod
storcli /cx[/ex]/sx set good [force]
storcli /cx[/ex]/sx set offline
storcli /cx[/ex]/sx set online
storcli /cx[/ex]/sx set missing
storcli /cx[/ex]/sx set bootdrive=<on|off>
```

The detailed description for each command follows.

storcli /*cx*[/*ex*]/*sx* set jbod

This command sets the drive state to JBOD.

Input example:

```
storcli /c1/e56/s3 set jbod
```

storcli /*cx*[/*ex*]/*sx* set good [force]

This command changes the drive state to unconfigured good.

Input example:

```
storcli /c1/e56/s3 set good
```

NOTE If the drive has an operating system or a file system on it, the StorCLI tool displays an error message and fails the conversion. If you want to proceed with the conversion, use the `force` option as shown in the following command.

Input example:

```
storcli /c1/e56/s3 set good [force]
```

storcli /*cx*[/*ex*]/*sx* set offline

This command changes the drive state to offline.

Input example:

```
storcli /c1/e56/s3 set offline
```

storcli /*cx*[/*ex*]/*sx* set online

This command changes the drive state to online.

Input example:

```
storcli /c1/e56/s3 set online
```

storcli /*cx*[/*ex*]/*sx* set missing

This command marks a drive as missing.

Input example:

```
storcli /c1/e56/s3 set missing
```

storcli /*cx*[/*ex*]/*sx* set bootmode=<on|off>

This command sets or unsets a physical drive as a boot drive.

Input example:

```
storcli /c1/e56/s3 set bootmode=on
```

6.6.3.4 Drive Initialization Commands

When you initialize drives, all the data from the drives is cleared. The Storage Command Line Interface Tool supports the following commands to initialize drives:

```
storcli /cx[/ex]/sx show initialization
storcli /cx[/ex]/sx start initialization
storcli /cx[/ex]/sx stop initialization
```

The detailed description for each command follows.

storcli /*cx*[/*ex*]/*sx* show initialization

This command shows the current progress of the initialization progress in percentage.

The estimated time (in minutes) left to complete the operation is also shown.

Input example:

```
storcli /c0/e31/s4 show initialization
```

storcli /*cx*[/*ex*]/*sx* start initialization

This command starts the initialization process on a drive.

Input example:

```
storcli /c0/e31/s4 start initialization
```

storcli /*cx*[/*ex*]/*sx* stop initialization

This command stops an initialization process running on the specified drive. A stopped initialization process cannot be resumed.

Input example:

```
storcli /c0/e56/s1 stop initialization
```

6.6.3.5 Drive Firmware Download Commands

The Storage Command Line Interface Tool supports the following command to download drive firmware:

storcli /*cx*[/*ex*]/*sx* download src=filepath [satabridge] [mode= 5|7]

This command flashes the drive firmware with the specified file.

The `satabridge` option lets you download the SATA bridge firmware in online mode.

The `mode` options specify the SCSI write buffer mode. The description follows:

- `5` – The drive firmware file is downloaded in chunks of 32KB.
- `7` – The entire drive firmware file is downloaded at once.

NOTE The default mode is `7`.

Input example:

```
storcli /c0/e56/s1 download src=c:\file1.bin
```

Input example:

```
storcli /c0/e56/s1 download src=c:\file1.bin mode=5
```

6.6.3.6 Locate Drives Commands

The Storage Command Line Interface Tool supports the following commands to locate a drive and activate the physical disk activity LED:

```
storcli /cx[/ex]/sx start locate
```

```
storcli /cx[/ex]/sx stop locate
```

The detailed description for each command follows.

storcli /*cx*[/*ex*]/*sx* start locate

This command locates a drive and activates the drive's LED.

Input example:

```
storcli /c0/e56/s1 start locate
```

storcli /*cx*[/*ex*]/*sx* stop locate

This command stops a locate operation and deactivates the drive's LED.

Input example:

```
storcli /c0/e56/s1 stop locate
```

6.6.3.7 Prepare to Remove Drives Commands

The Storage Command Line Interface Tool supports the following commands to prepare the physical drive for removal:

```
storcli /cx[/ex]/sx spindown
storcli /cx[/ex]/sx spinup
```

The detailed description for each command follows.

storcli /*cx*[/*ex*]/*sx* spindown

This command spins down an unconfigured drive and prepares it for removal. The drive state is unaffiliated and it is marked offline.

Input example:

```
storcli /cx/e34/s4 spindown
```

storcli /*cx*[/*ex*]/*sx* spinup

This command spins up a spun-down drive and the drive state is unconfigured good.

Input example:

```
storcli /cx/e34/s4 spinup
```

NOTE The `spinup` command works on a physical drive only if the user had previously issued a `spindown` command on the same physical drive.

6.6.3.8 Drive Security Command

The Storage Command Line Interface Tool supports the following drive security commands:

```
storcli /cx[/ex]/sx show securitykey keyid
```

storcli /*cx*[/*ex*]/*sx* show securitykey keyid

This command shows the security key for secured physical drives.

Input example:

```
storcli /c0/[e252]/s1 show SecurityKey keyid
```

storcli /*cx*/[*ex*]/*sx* set security = on

This command enables security on a JBOD.

Input example:

```
storcli /c0/[e252]/s1 set security = on
```

6.6.3.9 Drive Secure Erase Commands

The Storage Command Line Interface Tool supports the following drive erase commands:

```
storcli /cx[/ex]/sx secureerase [force]
```

```
storcli /cx[/ex]/sx show erase

storcli /cx[/ex]/sx start erase [simple|normal|thorough] [patternA=<value1>]
[patternB=<value2>]

storcli /cx[/ex]/sx stop erase
```

The detailed description for each command follows.

storcli /*cx*[/*ex*]/*sx* secureerase [force]

This command erases the drive's security configuration and securely erases data on a drive. You can use the `force` option as a confirmation to erase the data on the drive and the security information.

Input example:

```
storcli /c0/e25/s1 secureerase
```

NOTE This command deletes data on the drive and the security configuration and this data is no longer accessible. This command is used for SED drives only.

storcli /*cx*[/*ex*]/*sx* show erase

This command provides the status of erase operation on non-SEDs.

Input example:

```
storcli /c0/e25/s1 show erase
```

storcli /*cx*[/*ex*]/*sx* start erase [simple|normal|thorough|standard] [patternA=<*val1*>] [patternB=<*val2*>]

This command securely erases non-SED drives. The drive is written with erase patterns to make sure that the data is securely erased. You can use the following options with the start erase command:

Table 47 Drive Erase Command Options

Options	Value Range	Description
`erase`	`simple`: Single pass, single pattern write `normal`: Three pass, three pattern write `thorough`: Nine pass, repeats the normal write 3 times	Secure erase type.
`patternA`	8-bit value	Erase pattern A to overwrite the data.
`patternB`	8-bit value	Erase pattern B to overwrite the data.

Input example:

```
storcli /c0/e25/s1 start erase thorough patternA=10010011 patternB=11110000
```

6.6.3.10 Rebuild Drives Commands

The following commands rebuild drives in the Storage Command Line Interface Tool:

```
storcli /cx[/ex]/sx pause rebuild

storcli /cx[/ex]/sx resume rebuild

storcli /cx[/ex]/sx show rebuild

storcli /cx[/ex]/sx start rebuild

storcli /cx[/ex]/sx stop rebuild
```

NOTE If enclosures are used to connect physical drives to the controller, specify the enclosure ID in the command.

The detailed description for each command follows.

storcli /c*x*[/e*x*]/s*x* pause rebuild

This command pauses an ongoing rebuild process. You can run this command only for a drive that is currently rebuilt.

Input example:

```
storcli /c0/s4 pause rebuild
```

storcli /c*x*[/e*x*]/s*x* resume rebuild

This command resumes a paused rebuild process. You can run this command only when a paused rebuild process for the drive exists.

Input example:

```
storcli /c0/s4 resume rebuild
```

storcli /c*x*[/e*x*]/s*x* show rebuild

This command shows the progress of the rebuild process in percentage.

The estimated time (in minutes) left to complete the operation is also shown.

Input example:

```
storcli /c0/s5 show rebuild
```

storcli /c*x*[/e*x*]/s*x* start rebuild

This command starts a rebuild operation for a drive.

Input example:

```
storcli /c0/s4 start rebuild
```

storcli /c*x*[/e*x*]/s*x* stop rebuild

This command stops a rebuild operation. You can run this command only for a drive that is currently rebuilt.

Input example:

```
storcli /c0/s4 stop rebuild
```

6.6.3.11 Drive Copyback Commands

The Storage Command Line Interface Tool supports the following commands for drive copyback:

```
storcli /cx[/ex]/sx pause copyback
storcli /cx[/ex]/sx resume copyback
storcli /cx[/ex]/sx show copyback
storcli /cx[/ex]/sx start copyback target=eid:sid
storcli /cx[/ex]/sx stop copyback
```

The detailed description for each command follows.

NOTE In the copyback commands, `cx[/ex]/sx` indicates the source drive and `eid:sid` indicates the target drive.

NOTE When a copyback operation is enabled, the alarm continues to beep even after a rebuild is complete; the alarm stops beeping only when the copyback operation is completed.

storcli /c*x*[/e*x*]/s*x* pause copyback

This command pauses a copyback operation. You can run this command only when there is a copyback operation running.

Input example:

```
storcli /c0/e25/s4 pause copyback
```

storcli /c*x*[/e*x*]/s*x* resume copyback

This command resumes a paused copyback operation. You can run this command only when there is a paused copyback process for the drive.

Input example:

```
storcli /c0/e25/s4 resume copyback
```

storcli /c*x*[/e*x*]/s*x* show copyback

This command shows the progress of the copyback operation in percentage.

The estimated time (in minutes) left to complete the operation is also shown.

Input example:

```
storcli /c0/e25/s4 show copyback
```

storcli /c*x*[/e*x*]/s*x* start copyback target=eid:sid

This command starts a copyback operation for a drive.

Input example:

```
storcli /c0/e25/s4 start copyback target=25:8
```

storcli /c*x*[/e*x*]/s*x* stop copyback

This command stops a copyback operation. You can run this command only on drives that have the copyback operation running.

Input example:

```
storcli /c0/e25/s4 stop copyback
```

NOTE A stopped rebuild process cannot be resumed.

6.6.3.12 Hot Spare Drive Commands

The following commands create and delete hot spare drives:

```
storcli /cx[/ex]/sx add hotsparedrive
{dgs=<n|0,1,2...>}[enclaffinity][nonrevertible]
```

```
storcli /cx/[ex]/sx delete hotsparedrive
```

NOTE If enclosures are used to connect the physical drives to the controller, specify the enclosure ID in the command.

The detailed description for each command follows.

storcli /c*x*[/e*x*]/s*x* add hotsparedrive [{dgs=<*n*|0,1,2...>}] [enclaffinity][nonrevertible]

This command creates a hot spare drive. You can use the following options to create a hot spare drive.

Table 48 Add Hot Spare Drive Input Options

Option	Value Range	Description
dgs	Valid drive group number	Specifies the drive group to which the hot spare drive is dedicated.
enclaffinity	Valid enclosure number	Specifies the enclosure with which the hot spare is associated. If this option is specified, affinity is set; if it is not specified, there is no affinity. **NOTE** Affinity cannot be removed after it is set for a hot spare drive.
nonrevertible	—	Sets the drive as a nonrevertible hot spare.

Input example:

```
storcli /c0/e3/s4,5 add hotsparedrive
```

This command sets the drives /c0/e3/s4,5 as Global Hot spare.

Input example:

```
storcli /c0/e3/s6,8 add hotsparedrive dgs=0,1
```

This command sets /c0/e3/s6,8 as Dedicated Hot spare for disk groups 0,1.

storcli /*cx*/[*ex*]/*sx* delete hotsparedrive

This command deletes a hot spare drive.

Input example:

```
storcli /c0/e3/s4,5 delete hotsparedrive
```

6.6.3.13 Drive Performance Monitoring Commands

The Storage Command Line Interface Tool supports the following commands for drive performance monitoring:

```
Storcli /cx show pdfailevents [lastoneday] [fromSeqNum=xx] [file=filename]
```

```
Storcli /cx set pdfaileventoptions detectiontype=val correctiveaction=val
errorrthreshold=val
```

The detailed description for each command follows.

Storcli / *cx* show pdfailevents *[lastoneday] [fromSeqNum=xx][file=filename]*

This command shows all of the drive predictive failure events.

Input Example 1:

```
storcli /c0 show pdfailevents
```

This command shows all of the drive predictive failure events from the oldest sequence number.

Input Example 2:

```
storcli /c0 show pdfailevents lastoneday
```

This command shows all of the drive predictive failure events that occurred in the last 24 hours.

Input Example 3:

```
storcli /c0 show pdfailevents fromSeqNum
```

This command shows all of the drive predictive failure events generated from the specified sequence number.

NOTE While running these commands, if you provide a file name, the events are written to the specified file as values separated by commas.

Storcli / *cx* set pdfaileventoptions detectiontype=*val* correctiveaction=*val* errorrthreshold=*val*

This command provides the current settings of the `pdfaileventoptions` set on the controller and the various options to change these settings.

Input Example 1:

```
storcli /c0 set pdfaileventoptions detectiontype=x
```

This command sets the detection type for the drive. The valid range is 0 to 3.

NOTE For the changes to take effect, a reboot is required.

Input Example 2:

```
storcli /c0 set pdfaileventoptions correctiveaction=x
```

This command sets the corrective actions to be taken when the media error is detected. The valid value is 0 or 1.

Input Example 3:

```
storcli /c0 set pdfaileventoptions errorrthreshold=x
```

This command sets the error threshold for the controller. The valid range is 0 to 3

6.6.4 Virtual Drive Commands

The Storage Command Line Interface Tool supports the following virtual drive commands. The following table describes frequently used virtual drive commands.

Table 49 Virtual Drives Commands Quick Reference Table

Commands	Value Range	Description
`add`	See the following **Add RAID Configuration Input Options** tables.	Creates virtual drives.
`delete`	`cc` or `cachecade`: Deletes CacheCade virtual drives. `force`: Deletes the virtual drive where operating system is present.	Deletes a virtual drive.
`set`	See the following **Add RAID Configuration Input Options** tables and**Change Virtual Properties Commands** section.	Sets virtual drive properties.
`show`	`all`: Shows all properties of the virtual drive. `cc`: Shows properties of CacheCade virtual drives. See the **Virtual Drive Show Command** section.	Shows virtual drive information.

6.6.4.1 Add Virtual Drives Commands

The Storage Command Line Interface Tool supports the following commands to add virtual drives:

```
storcli /cx add vd raid[0|1|5|6|00|10|50|60][Size=<VD1_Sz>,<VD2_Sz>,..|all]
[name=<VDNAME1>,..] drives=e:s|e:s-x,y;e:s-x,y,z [PDperArray=x][SED]
[pdcache=on|off|default][pi] [DimmerSwitch(ds)=default|automatic(auto)|
none|maximum(max)|MaximumWithoutCaching(maxnocache)]

[wt|wb|awb] [nora|ra] [direct|cached][cachevd] [Strip=<8|16|32|64|128|256|1024>]
[AfterVd=X][EmulationType=0|1|2] [Spares = [e:]s|[e:]s-x|[e:]s-x,y]
[force][ExclusiveAccess]
```

NOTE The supported strip size can vary from a minimum of 64 KB to 1 MB for MegaRAID controllers and only 64 KB for Integrated MegaRAID controllers.

```
storcli /cx add vd each raid0 [name=<VDNAME1>,..] [drives=e:s|e:s-x|e:s-x,y] [SED]
[pdcache=on|off|default][pi] [DimmerSwitch(ds)=default|automatic(auto)|
none|maximum(max)|MaximumWithoutCaching(maxnocache)] [wt|wb|awb] [nora|ra]
[direct|cached][EmulationType=0|1|2]
[Strip=<8|16|32|64|128|256|1024>][ExclusiveAccess]
```

NOTE The supported strip size can vary from a minimum of 64 KB to 1 MB for MegaRAID controllers and only 64 KB for Integrated MegaRAID controllers.

```
storcli /cx add VD cachecade|cc raid[0,1,10] drives = [e:]s|[e:]s-x|[e:]s-x,y
[WT|WB|AWB] [assignvds = 0,1,2]
```

This command creates a RAID configuration. You can use the following options to create the RAID volume:

NOTE * indicates default values.

The detailed description for each command follows.

storcli /*cx* add vd raid[0|1|5|6|00|10|50|60][Size=<*VD1_Sz*>,<*VD2_Sz*>,..|*all] [name=<*VDNAME1*>,..] drives=*e:s*|*e:s-x*|*e:s-x,y;e:s-x,y,z* [PDperArray=x][SED] [pdcache=on|off|*default][pi] [DimmerSwitch(ds)=default|automatic(auto)| *none|maximum(max)|MaximumWithoutCaching(maxnocache)][cachevd][ExclusiveAccess|SharedAccess*] [wt|*wb |awb] [nora|*ra] [*direct|cached] [EmulationType=0][Strip=<8|16|32|64|128|256|1024>] [AfterVd=*X*] [Spares = [*e:*]*s*|[*e:*]*s-x*|[*e:*]*s-x,y*] [force]**

NOTE The supported strip size can vary from a minimum of 64 KB to 1 MB for MegaRAID controllers and only 64 KB for Integrated MegaRAID controllers.

Table 50 Add RAID Configuration Input Options

Option	Value Range	Description
raid	[0|1|5|6|00|10|50|60].	Sets the RAID type of the configuration.
size	Maximum size based on the physical drives and RAID level.	Sets the size of each virtual drive. The default value is for the capacity of all referenced disks.
name	15 characters of length.	Specifies the drive name for each virtual drive.
drives	Valid enclosure number and valid slot numbers for the enclosure.	In *e*:*s*|*e*:*s*-*x*|*e*:*s*-*x*,*y*: ■ *e* specifies the enclosure ID. ■ *s* represents the slot in the enclosure. ■ *e*:*s*-*x* is the range convention used to represent slots *s* to *x* in the enclosure *e* (250 characters max.). **NOTE** Make sure that the same block size (in a physical drive) is used in each [*e*:*s*] pair. As an example, if you use 4096 bytes in the *e0*:*s0* pair, use 4096 bytes in the *e1*:*s1* pair too. Mixing of block sizes between the [*e*:*s*] pairs is not supported.
pdperarray	1–16.	Specifies the number of physical drives per array. The default value is automatically chosen.
sed	—	Creates security-enabled drives.

Table 50 Add RAID Configuration Input Options (Continued)

Option	Value Range	Description
pdcache	on|off|default.	Enables or disables PD cache.
pi	—	Enables protection information.
dimmerswitch	default: Logical device uses controller default power-saving policy. automatic (auto): Logical device power savings are managed by firmware. none: No power-saving policy. maximum (max): Logical device uses maximum power savings. MaximumWithoutCaching (maxnocache): Logical device does not cache write to maximize power savings.	Specifies the power-saving policy. Sets to default automatically.
direct|cached	cached: Cached I/O. direct: Direct I/O.	Sets the logical drive cache policy. Direct I/O is the default.
EmulationType	0: Default emulation, which means if there are any 512e drives in the configured ID, then the physical bytes per sector is shown as 512e(4k). If there are no 512e drives then the physical bytes per sector will be 512n. 1: Disable, which means even though there are no 512e drives in the configured ID, the physical bytes per sector will be shown 512n. 2=Force, which means even though there are no 512e drives in the configured ID, the physical bytes per sector will be shown as 512e (4k).	
wt|wb|awb	wt: Write through.wb: Write back.awb: Always Write Back.	Enables write through. Write back is the default.
nora|ra	ra: Read ahead.nora: No read ahead.	Disables read ahead. Enabled is the default.
cachevd	—	Enables SSD caching on the created virtual drive.
strip	8, 16, 32, 64, 128, 256, 512, 1024. **NOTE** The supported strip size can vary from a minimum of 64 KB to 1 MB for MegaRAID controllers and only 64 KB for Integrated MegaRAID controllers.	Sets the strip size for the RAID configuration.
aftervd	Valid virtual drive number.	Creates the VD in the adjacent free slot next to the specified VD.
spares	Number of spare physical drives present.	Specifies the physical drives that are to be assigned to a disk group for spares.
force	—	Forces a security-capable physical drive to be added to a drive group without security.

Input example:

```
storcli /c0 add vd raid10 size=2gb,3gb,4gb names=tmp1,tmp2,tmp3 drives=252:2-3,5,7
pdperarray=2
```

storcli /cx add vd cc|cachecade raid[0,1,10] drives=[*e*:]*s*|[*e*:]*s*-*x*|[*e*:]*s*-*x*,*y* [[wt|*wb|awb]] [assignvds=0,1,2]

This command creates CacheCade virtual drives and associates existing virtual drives to CacheCade virtual drives. You can use the following options to create the CacheCade virtual drive.

Table 51 Add RAID Configuration Input Options

Option	Value Range	Description
cachecade	—	Creates a CacheCade virtual drive.
raid	0, 1, 10	Sets the RAID type of the CacheCade virtual drive.
drives	Valid enclosure number and valid slot number	See the drivesrow in the previous table for format.
wt|*wb|awb	wt: Enables write through. wb: Enables write back. awb Enables always write back.	Enables or disables write cache.
assignvds	Valid virtual drive number (0 to 63)	Specifies the list of virtual drives associated with the new CacheCade virtual drives.

Input example:

```
storcli /c0 add vd raid10 size=2gb,3gb,4gb names=tmp1,tmp2,tmp3 drives=252:2-3, 7
```

6.6.4.2 Delete Virtual Drives Commands

The Storage Command Line Interface Tool supports the following virtual drive delete commands:

```
storcli /cx/vx|vall del
storcli /cx/vx|vall del cachecade
storcli /cx/vx|vall del force
storcli /cx/vx del [cachecade] [discardcache] [force]
```

NOTE If the virtual drive has user data, you must use the force option to delete the virtual drive.
A virtual drive with a valid master boot record (MBR) and a partition table is considered to contain user data.

If you delete a virtual drive with a valid MBR without erasing the data and then create a new virtual drive using the same set of physical drives and the same RAID level as the deleted virtual drive, the old unerased MBR still exists at block0 of the new virtual drive, which makes it a virtual drive with valid user data. Therefore, you must provide the force option to delete this newly created virtual drive.

The detailed description for each command follows.

storcli /cx/vx|vall del

This command deletes a particular virtual drive or, when the vall option is used, all the virtual drives on the controller are deleted.

Input example:

```
storcli /c0/v2 del
```

ATTENTION This command deletes virtual drives. Data located on these drives will no longer be accessible.

storcli /cx/vx|vall del cachecade

This command deletes a specific CacheCade virtual drive on a controller, or all the CacheCade configuration for a controller.

Input example:

```
storcli /c0/vall del cachecade
```

ATTENTION This command deletes virtual drives. Data located on these drives will no longer be accessible.

storcli /*cx*/*vx*|vall del force

This command deletes a virtual drive only after the cache flush is completed. With the `force` option, the command deletes a virtual drive without waiting for the cache flush to complete.

Input example:

```
storcli /c0/v2 del force
```

ATTENTION This command deletes the virtual drive where the operating system is present. Data located on these drives and the operating system of the drive will no longer be accessible

storcli /*cx*/*vx* del [cachecade] [discardcache] [force]

This command with the `discardCache` option deletes the virtual drive without flushing the cached data.

Input example:

```
storcli /c0/v2 delete discardcache
```

6.6.4.3 Virtual Drive Show Commands

The Storage Command Line Interface Tool supports the following virtual drive show commands:

```
storcli /cx/vx show
storcli /cx/vx show all [logfile[=filename]]
storcli /cx/vx show hoqrebuild
```

The detailed description for each command follows.

storcli /*cx*/*vx* show

This command shows the summary of the virtual drive information.

Input example:

```
storcli /c0/v0 show
```

storcli /*cx*/*vx* show all [logfile[=*filename*]]

The `show all` command shows all of the virtual drive information, which includes the virtual drive information, physical drives used for the virtual drives, and virtual drive properties.

If you use the `logfile` option in the command syntax, the logs are written to the specified file. If you do not specify a file name, then the logs are written to the `storsas.log` file. If you do not use the `logfile` option in the command syntax, the entire log output is printed to the console.

Input example:

```
storcli /c0/v0 show all [logfile[=log.txt]]
```

storcli /*cx*/*vx* show hoqrebuild

This command shows the current status of the head of queue rebuild for the virtual drive.

Input example:

```
storcli /c0/v0 show hoqrebuild
```

6.6.4.4 Preserved Cache Commands

If a virtual drive becomes offline or is deleted because of missing physical disks, the controller preserves the dirty cache from the virtual disk. The Storage Command Line Interface Tool supports the following commands for preserved cache:

```
storcli /cx/vx delete preservedCache [force]
storcli /cx show preservedCache
```

The detailed description for each command follows.

storcli /*cx*/*vx* delete preservedcache

This command deletes the preserved cache for a particular virtual drive on the controller in missing state. Use the `force` option to delete the preserved cache of a virtual drive in offline state.

Input example:

```
storcli /c0/v1 delete preservedcache
```

storcli /*cx* show preservedCache

This command shows the virtual drive that has preserved cache and whether the virtual drive is offline or missing.

Input example:

```
storcli /c0 show preservedCache
```

6.6.4.5 Change Virtual Properties Commands

The Storage Command Line Interface Tool supports the following commands to change virtual drive properties:

```
storcli /cx/vx set accesspolicy=<rw|ro|blocked|rmvblkd>
storcli /cx/vx set iopolicy=<cached|direct>
storcli /cx/vx set name=<namestring>
storcli /cx/vx set pdcache=<on|off|default>
storcli /cx/vx set rdcache=<ra|nora>
storcli /cx/vx|vall set ssdcaching=<on|off>
storcli /cx/vx|vall set HostAccess=ExclusiveAccess|SharedAccess
storcli /cx/vx set wrcache=<wt|wb|awb>
storcli /cx/vx set emulationType=0|1|2
storcli /cx/vx set ds=Default|Auto|None|Max|MaxNoCache
storcli /cx/vx set autobgi=On|Off
storcli /cx/vx set pi=Off
storcli /cx/vx set bootdrive=<On|Off>
storcli /cx/vx set hidden=On|Off
storcli /cx/vx set hoqrebuild=On|Off
storcli /cx/vx set cbsize=0|1|2 cbmode=0|1|2|3|4|7
```

The detailed description for each command follows.

storcli /*cx*/*vx* set accesspolicy=<rw|ro|blocked|rmvblkd>

This command sets the access policy on a virtual drive to read write, read only, or blocked or rmvblkd (remove blocked).

Input example:

```
storcli /c0/v0 set accesspolicy=rw
```

storcli /c*x*/v*x* set iopolicy=<cached|direct>

This command sets the I/O policy on a virtual drive to cached I/O or direct I/O.

Input example:

```
storcli /c0/v0 set iopolicy=cached
```

storcli /c*x*/v*x* set name=<*namestring*>

This command names a virtual drive. The name is restricted to 15 characters.

Input example:

```
storcli /c1/v0 set name=testdrive123
```

storcli /c*x*/v*x* set pdcache=<on|off|default>

This command sets the current disk cache policy on a virtual drive to on, off, or default setting.

Input example:

```
storcli /c0/v0 set pdcache=on
```

storcli /c*x*/v*x* set rdcache=<ra|nora>

This command sets the read cache policy on a virtual drive to read ahead or no read ahead.

Input example:

```
storcli /c0/v0 set rdcache=nora
```

storcli /c*x*/v*x*|vall set ssdcaching=<on|off>

This command assigns CacheCade virtual drives. If `ssdcaching=off`, the CacheCade virtual drive is removed.

Input example:

```
storcli /c0/v0 set ssdcaching=on
```

storcli /cx/vx|vall set HostAccess=ExclusiveAccess|SharedAccess

This command sets the host access policy for the virtual drive. when the host access policy is exclusive access, a server has exclusive access to the virtual drive. The virtual drive cannot be shared between servers. If the host policy is shared access, the virtual drive can be shared between servers.

Input example:

```
storcli /c0/v0 set HostAccess=ExclusiveAccess
```

storcli/cx/vx set wrcache=<wt|wb|awb>

This command sets the write cache policy on a virtual drive to write back, write through, or always write back.

Input example:

```
storcli /c0/v0 set wrcache=wt
```

storcli /c*x*/v*x* set hidden=on|off

This command hides or unhides a virtual drive. If `hidden=on`, the virtual drive is hidden.

Input example:

```
storcli /c0/v0 set hidden=on
```

storcli /c*x*/v*x* set hoqrebuild=on|off

This command enables or disables the head of the queue drive rebuild on a virtual drive.

Input example:

```
storcli /c0/v0 set hoqrebuild=on
```

storcli /*cx*/*vx* set cbsize=0|1|2 cbmode=0|1|2|3|4|7

This command sets the Cache bypass size and the Cache bypass mode on a virtual drive.

The `cbsize` option follows:

- `0` – 64k Cache bypass.
- `1` – 128k Cache bypass.
- `2` – 256k Cache bypass.

The `cbmode` option follows:

- `0` – Enable the intelligent mode Cache bypass.
- `1` – Enable the standard mode Cache bypass.
- `2` – Enable the custom mode Cache bypass 1.
- `3` – Enable the custom mode Cache bypass 2.
- `4` – Enable the custom mode Cache bypass 3.
- `7` – Disable Cache bypass.

NOTE When `cbmode` is set to 7, the user given `cbsize` value is ignored

Input example:

```
storcli /c0/v0 set cbsize=1 cbmode=2
```

6.6.4.6 Virtual Drive Initialization Commands

The Storage Command Line Interface Tool supports the following commands to initialize virtual drives:

```
storcli /cx/vx show init
storcli /cx/vx start init [full][Force]
storcli /cx/vx stop init
```

NOTE If the virtual drive has user data, you must use the `force` option to initialize the virtual drive.
A virtual drive with a valid MBR and partition table is considered to contain user data.

The detailed description for each command follows.

storcli /*cx*/*vx* show init

This command shows the initialization progress of a virtual drive in percentage.

The estimated time (in minutes) left to complete the operation is also shown.

Input example:

```
storcli /c0/v2 show init
```

storcli /*cx*/*vx* start init [full]

This command starts the initialization of a virtual drive. The default initialization type is fast initialization. If the `full`option is specified, full initialization of the virtual drive starts.

Input example:

```
storcli /cx/vx start init [full]
```

storcli /*cx*/*vx* stop init

This command stops the initialization of a virtual drive. A stopped initialization cannot be resumed.

Input example:

```
storcli /c0/v0 stop init
```

6.6.4.7 Virtual Drive Erase Commands

The Storage Command Line Interface Tool supports the following commands to erase virtual drives:

```
storcli /cx/vx erase
storcli /cx/vx show erase
```

The detailed description for each command follows.

storcli /*cx*/*vx* erase

This command erases the data on the virtual drive.

Input example:

```
storcli /c0/v0 erase
```

storcli /*cx*/*vx* show erase

This command shows the status of the erase operation on the virtual drive.

Input example:

```
storcli /c0/v0 show erase
```

6.6.4.8 Virtual Drive Migration Commands

NOTE The virtual drive migration commands are not supported in Embedded MegaRAID.

The Storage Command Line Interface Tool supports the following commands for virtual drive migration (reconstruction):

```
storcli /cx/vx show migrate
storcli /cx/vx start migrate <type=raidx> [option=<add|remove>
drives=[e:]s|[e:]s-x|[e:]s-x,y] [Force]
```

The detailed description for each command follows.

storcli /*cx*/*vx* show migrate

This command shows the progress of the virtual drive migrate operation in percentage.

The estimated time (in minutes) left to complete the operation is also shown.

Input example:

```
storcli /c0/v0 show migrate
```

storcli /*cx*/*vx* start migrate <type=*raidlevel*> [option=<add | remove> drives=<e1:s1,e2:s2 ...>]

This command starts the reconstruction on a virtual drive to the specified RAID level by adding or removing drives from the existing virtual drive. You can use the following options with the start migrate command.

Table 52 Virtual Drive Migration Command Options

Options	Value Range	Description
`type =RAID level`	`RAID [0|1|5|6]`	The RAID level to which the virtual drive must be migrated.
`[option=<add | remove> drives=<e1:s1,e2:s2, …>]`	`add`: Adds drives to the virtual drive and starts reconstruction. `remove`: Removes drives from the virtual drive and starts reconstruction. `drives`: The enclosure number and the slot number of the drives to be added to the virtual drive. **NOTE** Make sure that the same block size (in a physical drive) is used in each [*e*:*s*] pair. As an example, if you use 4096 bytes in the *e0*:*s0* pair, use 4096 bytes in the *e1*:*s1* pair too. Mixing of block sizes between the [*e*:*s*] pairs is not supported.	Adds or removes drives from the virtual drive.

Virtual drive migration can be done between the following RAID levels.

Table 53 Virtual Drive Migration Table

Initial RAID level	Migrated RAID level
RAID 0	RAID 1
RAID 0	RAID 5
RAID 0	RAID 6
RAID 1	RAID 0
RAID 1	RAID 5
RAID 1	RAID 6
RAID 5	RAID 0
RAID 5	RAID 6
RAID 6	RAID 0
RAID 6	RAID 5

Input example: In the following example, `252` is the enclosure number and `0`, `1`, and `2` are the slot numbers.

```
storcli/c0/v0 start migrate type=raid0 option=add drives=252:0,252:1,252:2
```

6.6.4.9 Virtual Drive Consistency Check Commands

The Storage Command Line Interface Tool supports the following commands for virtual drive consistency checks:

```
storcli /cx/vx pause cc
storcli /cx/vx resume cc
storcli /cx/vx show cc
storcli /cx/vx start cc [force]
storcli /cx/vx stop cc
```

NOTE If enclosures are used to connect the physical drives to the controller, specify the IDs in the command.

The detailed description for each command follows.

storcli /*cx*/*vx* pause cc

This command pauses an ongoing consistency check process. You can resume the consistency check at a later time. You can run this command only on a virtual drive that has a consistency check operation running.

Input example:

```
storcli /c0/v4 pause cc
```

storcli /*cx*/*vx* resume cc

This command resumes a suspended consistency check operation. You can run this command on a virtual drive that has a paused consistency check operation.

Input example:

```
storcli /c0/v4 resume cc
```

storcli /*cx*/*vx* show cc

This command shows the progress of the consistency check operation in percentage.

The estimated time (in minutes) left to complete the operation is also shown.

Input example:

```
storcli /c0/v5 show cc
```

storcli /*cx*/*vx* start cc force

This command starts a consistency check operation for a virtual drive. Typically, a consistency check operation is run on an initialized virtual drive. Use the `force` option to run a consistency check on an uninitialized drive.

Input example:

```
storcli /c0/v4 start cc
```

storcli /*cx*/*vx* stop cc

This command stops a consistency check operation. You can run this command only for a virtual drive that has a consistency check operation running.

Input example:

```
storcli /c0/v4 stop cc
```

NOTE You cannot resume a stopped consistency check process.

6.6.4.10 Background Initialization Commands

The Storage Command Line Interface Tool supports the following commands for background initialization:

```
storcli /cx/vx resume bgi
storcli /cx/vx set autobgi=<on|off>
storcli /cx/vx show autobgi
storcli /cx/vx show bgi
storcli /cx/vx stop bgi
storcli /cx/vx suspend bgi
```

The detailed description for each command follows.

storcli /*cx*/*vx* resume bgi

This command resumes a suspended background initialization operation.

Input example:

```
storcli /c0/v0 resume bgi
```

storcli /*cx*/*vx* set autobgi=<on|off>

This command sets the auto background initialization setting for a virtual drive to on or off.

Input example:

```
storcli /c0/v0 set autobgi=on
```

storcli /*cx*/*vx* show autobgi

This command shows the background initialization setting for a virtual drive.

The estimated time (in minutes) left to complete the operation is also shown.

Input example:

```
storcli /c0/v0 show autobgi
```

storcli /*cx*/*vx* show bgi

This command shows the background initialization progress on the specified virtual drive in percentage.

The estimated time (in minutes) left to complete the operation is also shown.

Input example:

```
storcli /c0/v0 show bgi
```

storcli /*cx*/*vx* stop bgi

This command stops a background initialization operation. You can run this command only for a virtual drive that is currently initialized.

Input example:

```
storcli /c0/v4 stop bgi
```

storcli /*cx*/*vx* pause bgi

This command suspends a background initialization operation. You can run this command only for a virtual drive that is currently initialized.

Input example:

```
storcli /c0/v4 pause bgi
```

6.6.4.11 Virtual Drive Expansion Commands

The Storage Command Line Interface Tool supports the following commands for virtual drive expansion:

```
storcli /cx/vx expand size=<value> [expandarray]
storcli /cx/vx|vall show expansion
```

The detailed description for each command follows.

storcli /*cx*/*vx* expand size=<*value*> [expandarray]

This command expands the virtual drive within the existing array or if you replace the drives with drives larger than the size of the existing array. Even though the value provided by you may be in MB, the value of the expanded size is displayed based on the nearest possible unit. Depending on the input (value) provided by you, `storcli` recognizes the size from the input provided by you and rounds up the size to the nearest percentage of free space remaining on the drive group; hence, the actual expanded size may differ from the size requested by you. If the `expandarray` option is specified, the existing array is expanded. If this option is not specified, the virtual drive is expanded.

storcli /*cx*/*vx* show expansion

This command shows the expansion information on the virtual drive with and without array expansion.

Input example:

```
storcli /c0/v0 show expansion
```

6.6.4.12 Display the Bad Block Table

The Storage Command Line Interface Tool supports the following command to check for bad block entries of virtual drives on the selected controller:

```
storcli /cx/vx show bbmt
```

Input example:

```
storcli /c0/v0 show bbmt
```

6.6.4.13 Clear the LDBBM Table Entires

The Storage Command Line Interface Tool supports the following command to clear the LDBBM table entries:

```
storcli /cx/vx delete bbmt
```

Input example:

```
storcli /c0/v0 delete bbmt
```

6.6.5 Foreign Configurations Commands

The Storage Command Line Interface Tool supports the following commands to view, import, and delete foreign configurations:

```
storcli /cx/fall|fall del|delete [ securitykey=ssssssssss ]
storcli /cx/fall|fall import [preview][ securitykey=ssssssssss ]
storcli /cx/fall|fall show [all] [ securitykey=ssssssssss ]
```

NOTE Provide the security key when importing a locked foreign configuration created in a different machine that is encrypted with a security key.

The detailed description for each command follows.

storcli /*cx*/*fall*|fall del| delete [securitykey=*ssssssssss*]

This command deletes the foreign configuration of a controller. Input the security key if the controller is secured.

Input example:

```
storcli /c0/fall delete
```

storcli /*cx*/*fall*|fall import [preview] [securitykey=*ssssssssss*]

This command imports the foreign configurations of a controller. The `preview` option shows a summary of the foreign configuration before importing it.

Input example:

```
storcli /c0/fall import
```

storcli /*cx*/*fall*|fall show [all][securitykey=*ssssssssss*]

This command shows the summary of the entire foreign configuration for a particular controller. The `all` option shows all the information of the entire foreign configuration.

NOTE The EID:Slot column is populated for the foreign PDs that are locked.

Input example:

```
storcli /c0/fall show preview

storcli /c0/fall import preview

storcli /c0/fall show all
```

6.6.6 BIOS-Related Commands

The Storage Command Line Interface Tool supports the following BIOS commands:

```
storcli /cx set bios [state=<on|off>] [Mode=<SOE|PE|IE|SME>] [abs=<on|off>]
[DeviceExposure=<value>]
```

The detailed description for the command follows.

storcli /*cx* set bios [state=<on|off>] [Mode=<SOE|PE|IE|SME>] [abs=<on|off>] [DeviceExposure=<value>]

This command enables or disables the MegaRAID controller's BIOS, sets the BIOS boot mode, and enables the BIOS to select the best logical drive as the boot drive.The mode options abbreviations follow:

- SOE: Stop on Errors.
- PE: Pause on Errors.
- IE: Ignore Errors.
- SME: Safe mode on Errors.

NOTE The legacy BIOS can load a limited number of the PCI device's BIOS. Disable the MegaRAID BIOS to avoid issues during POST.

Input example:

```
storcli /c0 set bios[state=on][Mode=SOE][abs=on][deviceexposure=20]
```

6.6.6.1 OPROM BIOS Commands

The Storage Command Line Interface Tool supports the following OPROM BIOS commands:

```
storcli /cx/ex/sx set bootdrive=on|off

storcli /cx/vx set bootdrive=on|off

storcli /cx show bootdrive
```

The detailed description for each command follows.

storcli /*cx*/*ex*/*sx* set bootdrive=on|off

This command sets the specified physical drive as the boot drive. During the next reboot, the BIOS looks for a boot sector in the specified physical drive.

Input example:

```
storcli /c0/e32/s4 set bootdrive=on
```

storcli /*cx*/*vx* set bootdrive=on|off

This command sets the specified virtual drive as the boot drive. During the next reboot, the BIOS looks for a boot sector in the specified virtual drive.

Input example:

```
storcli /c0/v0 set bootdrive=on
```

storcli/*cx*/*vx* show bootdrive

This command shows the boot drive for the controller. The boot drive can be a physical drive or a virtual drive.

Input example:

```
storcli /c0/v0 show bootdrive
```

6.6.7 Drive Group Commands

This section describes the drive group commands.

6.6.7.1 Drive Group Show Commands

The Storage Command Line Interface Tool supports the following drive group commands:

```
storcli /cx/dall show
storcli /cx/dall show all
storcli /cx/dall show cachecade
storcli /cx/dx show
storcli /cx/dx show all
storcli /cx/dx set security=on
storcli /cx/dx split mirror
storcli /cx/dall show mirror
storcli /cx/dall add mirror src=<val>[force]
storcli /cx/dx set hidden=<on|off>
```

storcli /*cx*/dall show

This command shows the topology information of all the drive group.

Input example:

```
storcli /c0/dall show
```

storcli /*cx*/dall show all

This command shows all available configurations in the controller which includes topology information, virtual drive information, physical drive information, free space, and free slot information.

Input example:

```
storcli /c0/dall show all
```

storcli /*cx*/dall show cachecade

This command shows all CacheCade virtual drive information.

Input example:

```
storcli /c0/dall show cachecade
```

storcli /*cx*/dx show

This command shows the topology information of the drive group.

Input example:

```
storcli /c0/dx show
```

storcli /*cx*/dx show all

This command shows the physical drive and the virtual drive information for the drive group.

Input example:

```
storcli /c0/dx show all
```

storcli /*cx*/dx set security=on

This command enables security on the specified drive group.

Input example:

```
storcli /c0/dx set security=on all
```

storcli /*cx*/dx split mirror

This command enables you to perform a break mirror operation on a drive group. The break mirror operation enables a RAID 1 configured drive group to be broken into two volumes. You can use one of the volumes in another system and replicate it without making a copy of the virtual drive.

Input example:

```
storcli /c0/dx split mirror
```

storcli /*cx*/dall show mirror

This command shows information about the mirror associated with the drive group.

Input example:

```
storcli /c0/dall show mirror
```

storcli /*cx*/dall add mirror src=<val>[force]

This command joins the virtual drive with its mirror. The possible values to be used are 0, 1, or 2.

Input example:

```
storcli /c0/dall add mirror src=<1>[force]
```

storcli /*cx*/d*x* set hidden=<on|off>

This command hides or unhides a drive group.

Input example:

```
storcli /c0/d0 set hidden=on
```

6.6.7.2 Set Transport Ready Drive Group Command

This command sets the Transport Ready policy for the drive group.

NOTE To know more about the Transport Ready feature, see the Transport Ready Pre-requisites section.

Input Example:

```
storcli /cx/dx set transport=on  [EDHSP=on/off  [SDHSP=on/off]]
```

storcli /c0/d0 set transport=on

If the user selects the Transport Ready operation and the corresponding DG has Dedicated Hot Spare (Exclusive and shared) associated with it, then an option EDHSP and SDHSP with on/off can be used. The following table provides the details of the different EDHSP and SDHSP options:

Table 54 EDHDP and SDHSP options Table

Command	Transport option
Storcli /c0/d0 set transport=on	Include all Exclusive Dedicated Hot Spares and Shared Dedicated Hot Spares
Storcli /c0/d0 set transport=on EDHSP=on	Include only Exclusive Dedicated Hot Spare
Storcli /c0/d0 set transport=on EDHSP=off	Exclude all Exclusive Dedicated Hot Spares and Shared Dedicated Hot Spares
Storcli /c0/d0 set transport=on SDHSP=on	Include all Exclusive Dedicated Hot Spares and Shared Dedicated Hot Spares. **NOTE** Use this command only when no DHSP is present
Storcli /c0/d0 set transport=on SDHSP=off	Exclude all Exclusive Dedicated Hot Spares and Shared Dedicated Hot Spares
Storcli /c0/d0 set transport=on EDHSP=on SDHSP=on	Include all Exclusive Dedicated Hot Spares and Shared Dedicated Hot Spares
Storcli /c0/d0 set transport=on EDHSP=on SDHSP=off	Include only Exclusive Dedicated Hot Spare
Storcli /c0/d0 set transport=on EDHSP=off SDHSP=on	Include all Exclusive Dedicated Hot Spares and Shared Dedicated Hot Spares. **NOTE** Use this command only when no DHSP is present
Storcli /c0/d0 set transport=on EDHSP=off SDHSP=off	Exclude all Exclusive Dedicated Hot Spares and Shared Dedicated Hot Spares

storcli /cx/dx set transport=off

This command clears the Transport Ready policy for the drive group.

Input example:

```
storcli /c0/d0 set transport=off
```

6.6.8 Dimmer Switch Commands

6.6.8.1 Change Virtual Drive Power Settings Commands

The Storage Command Line Interface Tool supports the following command to change the Dimmer Switch setting.

storcli /*cx*/*vx* set ds=<default | auto | none | max | maxnocache>

This command changes the power-saving properties on an unconfigured drive and a hot spare drive. See `dimmerswitch` in the following table for values.

Input example:

```
storcli /cx/vx set ds=default
```

You can use the following combinations for the Dimmer Switch commands:

```
storcli /cx set ds=off type=1|2|4
```

```
storcli /cx set ds=on type=1|2 [properties]
storcli /cx set ds=on type=4 defaultldtype=<value> [properties]
storcli /cx set ds=on [properties]
```

The following table describes the power-saving options.

Table 55 Dimmer Switch Input Options

Option	Value Range	Description
`dimmerswitch` or `ds`	`on|off`	Turns the Dimmer Switch option on.
`type`	`1`: Unconfigured `2`: Hot spare `4`: All of the drives (unconfigured drives and hot spare drives).	Specifies the type of drives that the Dimmer Switch feature is applicable. By default, it is activated for unconfigured drives and hot spare drives.
`defaultldtype`	`auto`: Logical device power savings are managed by the firmware. `none`: No power saving policy. `max`: Logical device uses maximum power savings. `maxnocache`: Logical device does not cache write to maximise power savings.	Specifies the default logical drive type that is created by the Dimmer Switch option; set to none automatically.
`properties`	`disableldps`: Interval in hours or time in *hh:mm* format `spinupdrivecount`: Valid enclosure number (0 to 255) `SpinUpEncDelay`: Valid time in seconds	Sets the interval or time in which the power-saving policy for the logical drive is turned off. Specifies the number of drives in the enclosure that are spun up. Specifies the delay of spin-up groups within an enclosure in seconds.

storcli/*cx* show DimmerSwitch(ds)

This command shows the current Dimmer Switch setting for the controller.

Input example:

```
storcli/c0 show ds
```

6.6.9 BBU Commands

The Storage Command Line Interface Tool supports the following battery backup unit (BBU) commands:

NOTE To increase the life of a battery, the battery is not fully charged. Band Gap charging keeps the maximum battery charge within a band comfortably above the data retention time requirement instead of keeping the battery charged to the maximum level. However, when a learn cycle is required, the battery is fully charged because a learn cycle starts only once the battery is fully charged.

```
storcli /cx/bbu show
storcli /cx/bbu show all
storcli /cx/bbu set autolearnmode=<value>
storcli /cx/bbu set bbuMode=<value>
storcli /cx/bbu set learndelayinterval=<value>
```

```
storcli /cx/bbu set powermode=sleep
storcli /cx/bbu set writeaceess=sealed
storcli /cx/bbu set learnStartTime=[DDD HH|off]
storcli /cx/bbu show modes
storcli /cx/bbu show properties
storcli /cx/bbu show status
storcli /cx/bbu start learn
storcli /cx/bbu start retentiontest
```

The detailed description for each command follows.

storcli /*cx*/bbu show

This command shows the summary information for the BBU of a controller.

Input example:

```
storcli /c0/bbu show
```

storcli /*cx*/bbu show all

This command shows all the information of the BBU.

Input example:

```
storcli /c0/bbu show all
```

storcli /*cx*/bbu set autolearnmode=<value>

This command starts the automatic learn cycle on the battery. The possible values are **0** - Enabled, **1**- Disabled, and **2** - WarnViaEvent.

Input example:

```
storcli /c0/bbu set autolearnmode=0
```

storcli /*cx*/bbu set bbuMode=<*value*>

This command sets the BBU mode for the BBU. The following table shows the various BBU modes.

Table 56 BBU Mode

Mode	Description
0	48 hours of retention[a] at 60 °C, 1-year Service Life.
1	12 hours of retention at 45 °C, 5-year Service Life, transparent learn.[b]
2	12 hours of retention at 55 °C, 3-year Service Life, transparent learn.
3	24 hours of retention at 45 °C, 3-year Service Life, transparent learn.
4	48 hours of retention at 45 °C, 3-year Service Life.
5	48 hours of retention at 55 °C, 1-year Service Life.
6	Same as the description for BBU mode 5. The BBU mode 6 enables you to receive events when the battery capacity reaches suboptimal and critical thresholds.

a. Indicates how long the battery can hold data in the controller's memory in case of accidental system shutdown.

b. The controller's performance is not affected during the battery's learn cycle.

Input example:

```
storcli /c0/bbu set bbuMode=2
```

NOTE BBU modes are supported on any iBBU08/09 bbu/controller combo and later-generation controllers.

storcli /*cx*/bbu set learndelayinterval=<*value*>

This command sets the learn delay interval for the BBU in hours. The value must be between 0 to 168 hours (7 days).

Input example:

```
storcli /c0/bbu set learnDelayInterval=30
```

storcli /*cx*/bbu set powermode=sleep

This command places the battery in low-power storage mode. The battery automatically exits this state after 5 seconds.

Input example:

```
storcli /c0/bbu set powermode=sleep
```

storcli /*cx*/bbu set writeaccess=sealed

This command seals the gas gauge EEPROM write access.

NOTE Use the `set writeaccess=sealed` command at manufacturing time.

Input example:

```
storcli /c0/bbu set writeaccess=sealed
```

storcli /*cx*/bbu set writeaccess=sealed

This command seals the gas gauge EEPROM write access.

NOTE Use the `set writeaccess=sealed` command at manufacturing time.

Input example:

```
storcli /c0/bbu set writeaccess=sealed
```

storcli /*cx*/bbu set learnStartTime=[DDD HH| off]

This command sets the learn start time for the BBU in hours on the day specified. **DDD** refers to the day of the week (SUN,MON,....SAT), **HH** refers to the hours (0-23 hours), and **off**sets the learn start to off.

Input example:

```
storcli /c0/bbu set learnStartTime=MON 12
```

storcli /*cx*/bbu show properties

This command shows the BBU Learn properties for a controller.

Input example:

```
storcli /c0/bbu show properties
```

storcli /*cx*/bbu show status

This command shows the battery information, firmware status, and the gas gauge status.

Input example:

```
storcli /c0/bbu show status
```

storcli /*cx*/bbu start learn

This command starts the BBU learning cycle. The battery learn cycle is immediately started and no other parameters are required for this command.

Input example:

```
storcli /c0/bbu start learn
```

storcli /*cx*/bbu start retentiontest

This command starts the battery retention test. This command requires you to reboot your system.

Input example:

```
storcli /c0/bbu start retentiontest
```

6.6.10 CacheVault Commands

The Storage Command Line Interface Tool supports the following CacheVault® command:

```
storcli /cx/cv show all
```

storcli /*cx*/cv show all

This command shows all the information of a CacheVault that is connected to a controller.

NOTE This command only works when a CacheVault is connected to the controller; otherwise, an error message appears.

Input example:

```
storcli /c0/cv show all
```

6.6.11 Enclosure Commands

The Storage Command Line Interface Tool supports the following enclosure commands:

```
storcli /cx/ex download src=filepath[forceActivate]
storcli /cx/ex show all
storcli /cx/ex show status
```

The detailed description for each command follows.

storcli /*cx*/*ex* download src=filepath [forceactivate]

This command flashes the firmware with the file specified at the command line. The enclosure performs an error check after the operation. The following option can be used with the enclosure firmware download command.

Table 57 Enclosure Firmware Download Command Options

Option	Value Range	Description
`forceactivate`	—	Issues a command descriptor block (CDB) with write command with no data with command mode 0x0F (flash download already in progress). **NOTE** This option is used primarily to activate Scotch Valley Enclosures.

NOTE The firmware file that is used to flash the enclosure can be of any format. The StorCLI utility assumes that you provide a valid firmware image.

Input example:

```
storcli /c0/e0 download src=c:\file2.bin
```

storcli /*cx*/*ex* show all

This command shows all enclosure information, which includes general enclosure information, enclosure inquiry data, a count of enclosure elements, and information about the enclosure elements.

Input example:

```
storcli /c0/e0 show all
```

storcli /*cx*/*ex* show status

This command shows the enclosure status and the status of all the enclosure elements.

Input example:

```
storcli /c0/e0 show status
```

6.6.12 PHY Commands

The Storage Command Line Interface Tool supports the following phy commands:

```
storcli /cx/px|pall set linkspeed=0(auto)|1.5|3|6|12
storcli /cx/px|pall show
storcli /cx/px|pall show all
storcli /cx/ex show phyerrorcounters
storcli /cx[/ex]/sx show phyerrorcounters
storcli /cx[/ex]/sx reset phyerrorcounters
```

The detailed description for each command follows.

storcli /*cx*/*px*|pall set linkspeed=0(auto)|1.5|3|6|12

This command sets the PHY link speed. You can set the speed to 1.5 Gb/s, 3 Gb/s, 6 Gb/s, or 12 Gb/s. The linkspeed is set to `auto` when you specify `linkspeed = 0`.

Input example:

```
storcli /c0/p0 set linkspeed=1.5
```

storcli /*cx*/*px*|pall show

This command shows the basic PHY layer information.

Input example:

```
storcli /c1/p0 show
```

storcli /*cx*/*px*|pall show all

This command shows all the PHY layer information.

Input example:

```
storcli /c1/p0 show all
```

storcli /*cx*/*ex* show phyerrorcounters

This command shows the enclosure/expander phy error counters.

Input example:

```
storcli /c1/e0 show phyerrorcounters
```

storcli /*cx*[/*ex*]/*sx* show phyerrorcounters

This command shows the drive phy error counters.

Input example:

```
storcli /c1/e0/s0 show phyerrorcounters
```

storcli /*cx*[/*ex*]/*sx* reset phyerrorcounters

This command resets the drive phy error counters.

Input example:

```
storcli /c1/e0/s0 reset phyerrorcounters
```

6.6.13 Logging Commands

The Storage Command Line Interface Tool supports the following commands to generate and maintain log files:

```
storcli /cx clear events
storcli /cx delete termlog
storcli /cx show events file=<absolute path>
storcli /cx show eventloginfo
storcli /cx show termlog type=config|contents [logfile[=filename]]
storcli /cx show dequeue log file =<filepath>
Storcli /cx show alilog [logfile[=filename]]
```

The detailed description for each command follows.

storcli /*cx* delete events

This command deletes all records in the event log.

Input example:

```
storcli /c0 delete events
```

storcli /*cx* delete termlog

This command clears the TTY (firmware log for issue troubleshooting) logs.

Input example:

```
storcli /c0 delete termlog
```

storcli /*cx* show events file=<*absolute path*>

This command prints the system log to a text file and saves the file in the specified location.

Input example:

```
storcli /c0 show events file=C:\Users\brohan\test\eventreports
```

storcli /*cx* show eventloginfo

This command shows the history of log files generated.

Input example:

```
storcli /c0 show eventloginfo type=config
```

storcli /*cx* show termlog type=config|contents [logfile[=*filename*]]

This command shows the firmware logs. The `config` option shows the term log configuration (settings of TTY BBU buffering), the `contents` option shows the term log. The `contents` option is the default.

If you use the `logfile` option in the command syntax, the logs are written to the specified file. If you do not specify a file name, then the logs are written to the `storsas.log` file. If you do not use the `logfile` option in the command syntax, the entire log output is printed to the console.

Input example:

```
storcli /c0 show termlog=contents [logfile[=log.txt]]
```

storcli /*cx* show dequeue log =<filepath>

This command shows the debug log from the firmware.

Input example:

```
storcli /c0 show dequeue log=<c:\test\log.txt>
```

storcli /*cx*show alilog [logfile[=*filename*]]

This command writes the system logs to the specified file.

Input example:

```
storcli /c0 show alilog [logfile[=log.txt]]
```

6.6.14 Automated Physical Drive Caching Commands

The Storage Command Line Interface Tool supports the following automated physical drive caching commands:

```
storcli /cx set autopdcache=<off|r0>[immediate]
storcli /cx show autopdcache
```

The detailed description for each command follows.

storcli /*cx* set autopdcache=<off|r0>[immediate]

This command lets you set the controller's automated physical drive cache policy to RAID 0. When set to RAID-0, all un-configured physical drives are configured as a single RAID 0 drive, until the maximum virtual drive limit is reached. The `immediate` option lets this command execute the conversion (to RAID 0) operation only on all the existing physical drives. Any newly physical drives connected in the future do not get converted to RAID 0. If you omit the `immediate` option in this command, conversion to RAID 0 takes place on newly connected physical drives too. Automatic conversion to RAID 0 can be turned off by setting the autopdcache policy to `off`.

Input example:

```
storcli /c0 set autopdcache=r0 immediate
```

storcli /*cx* show autopdcache

This command lets you view the automatic physical drive caching property.

Input example:

```
storcli /c0 show autopdcache
```

6.7 Frequently Used Tasks

6.7.1 Showing the Version of the Storage Command Line Interface Tool

The following command shows the version of the command line tool:

```
Storcli -v
```

6.7.2 Showing the StorCLI Tool Help

The following command shows the StorCLI tool help:

```
Storcli -h
```

Help appears for all the StorCLI tool commands.

6.7.3 Showing System Summary Information

The following command shows the summary of all the controller information:

```
Storcli -show [all]
```

6.7.4 Showing Free Space in a Controller

The following command shows the free space available in the controller:

```
Storcli /cx show freespace
```

6.7.5 Adding Virtual Drives

The following command creates a virtual drive:

```
Storcli /cx add vd type=raid[0|1|5|6|10|50|60][Size=<VD1_Sz>,<VD2_Sz>,..|*all]
[name=<VDNAME1>,..] drives=e:s|e:s-x|e:s-x,y [PDperArray=x|auto*]
[SED] [pdcache=on|off|*default][pi] [DimmerSwitch(ds)=default|automatic(auto)|
*none|maximum(max)|MaximumWithoutCaching(maxnocache)] [wt|*wb|awb] [nora|*ra]
[*direct|cached]
[strip=<8|16|32|64|128|256|512|1024] [AfterVd=x] [Spares=[e:]s|[e:]s-x|[e:]s-x,y]
```

NOTE The supported strip size can vary from a minimum of 64 KB to 1 MB for MegaRAID controllers and only 64 KB for Integrated MegaRAID controllers.

```
[Cbsize = 0|1|2 Cbmode = 0|1|2]
[force]
```

The following inputs can be used when adding virtual drives:

- The controller in which the virtual drives are created.
- The RAID type of the virtual drives.
  The supported RAID types are 0, 1, 5, 6, 10, 50, 60.

- The size of each virtual drive.
- The drives that are used to create the virtual drives.

  drives = `e:s|e:s-x|e:s-x,y`

  Where:
  - *e* specifies the enclosure ID.
  - *s* represents the slot in the enclosure.
  - `e:s-ex` is the range conventions used to represents slots `s` to `x` in the enclosure `e`.
- The physical drives per array.

  The physical drives per array can be set to a particular value.
- The `SED` option creates security-enabled drives.
- The `PDcache` option can be set to `on` or `off`.
- The `pi` option enables protection information.
- The Dimmer Switch is the power save policy. It can be set to default or `automatic *`, `none`, `maximum(max)`, or `MaximumWithoutCaching(maxnocache)`.
- The `wt` option disables write back.
- The `nora` option disables read ahead.
- The `cached` option enables the cached memory.
- The `strip` option sets the strip size.

  It can take the values `8`, `16`, `32`, `64`, `128`, `256`, `512`, `1024`.

  NOTE The supported strip size can vary from a minimum of 64 KB to 1 MB for MegaRAID controllers and only 64 KB for Integrated MegaRAID controllers.
- The `AfterVdX` option creates the virtual drives in the adjacent free slot next to the specified virtual drives.

  NOTE The * indicates default values used in the creation of the virtual drives. If values are not specified, the default values are taken.

Example: `/cxadd vd type=r1 drives=0:10-15 WB Direct strip=64`

This command creates a RAID volume of RAID 1 type from drives in slots 10 to slot 15 in enclosure 0. The strip size is 64kb.

6.7.6 Setting the Cache Policy in a Virtual Drive

The following command sets the write cache policy of the virtual drive:

`Storcli /cx/v(x|all) set wrcache=wt|wb|awb`

The command sets the write cache to write back, write through, or always write back.

6.7.7 Showing Virtual Drive Information

The following command shows the virtual drive information for all the virtual drives in the controller:

`storcli /cx show [all]`

6.7.8 Deleting Virtual Drives

The following command deletes virtual drives:

```
storcli /cx/v(x|all) del [cc|cachecade]
```

The following inputs are required when deleting a virtual drive:

- The controller on which the virtual drive or virtual drives is present.
- The virtual drives that must be deleted; or you can delete all the virtual drives on the controller using the `vall` option.
- The `cc` or `cachecade` option to confirm that the deleted drive is a CacheCade drive.

6.7.9 Flashing Controller Firmware

The following command is used to flash the controller firmware.

```
storcli /cx download file=filepath [fwtype=<value>] [nosigchk]
[noverchk][resetnow]
```

For more information, see Flashing Controller Firmware Command. For limitations, see Online Firmware Upgrade Support.

Chapter 7: MegaRAID Storage Manager Overview and Installation

This chapter provides a brief overview of the MegaRAID Storage Manager software and explains how to install it on the supported operating systems.

7.1 Overview

The MegaRAID Storage Manager software enables you to configure, monitor, and maintain storage configurations on Avago SAS controllers. The MegaRAID Storage Manager graphical user interface (GUI) makes it easy for you to create and manage storage configurations.

7.1.1 Creating Storage Configurations

The MegaRAID Storage Manager software enables you to easily configure the controllers, drives, and virtual drives on your workstation or on the server. The Configuration wizard greatly simplifies the process of creating drive groups and virtual drives. The wizard allows you to easily create new storage configurations and modify the configurations.

You can create configurations using the following modes:

- **Simple configuration** specifies a limited number of settings and has the system select drives for you. This option is the easiest way to create a virtual drive.
- **Advanced configuration** lets you choose additional settings and customize the creation of virtual drives. This option provides greater flexibility when creating virtual drives for your specific requirements because you can select the drives and the virtual drive settings when you create a virtual drive. In addition, you can use the advanced configuration procedure to create spanned drive groups.

In addition, the Modify Drive Group wizard enables you to increase the capacity of a virtual drive and to change the RAID level of a drive group.

NOTE The Modify Drive Group wizard was previously known as the Reconstruction wizard.

7.1.2 Monitoring Storage Devices

The MegaRAID Storage Manager software displays the status of controllers, virtual drives, and drives on the workstation or on the server that you are monitoring. The system errors and events are recorded in an event log file and are displayed on the dialog. Special device icons appear on the window to notify you of drive failures and other events that require immediate attention.

7.1.3 Maintaining Storage Configurations

You can use the MegaRAID Storage Manager software to perform system maintenance tasks, such as running patrol read operations, updating firmware, and running consistency checks on drive groups that support redundancy.

7.2 Hardware and Software Requirements

The hardware requirements for the MegaRAID Storage Manager software are as follows:

- PC-compatible computer with an IA-32 (32-bit) Intel® Architecture processor or an EM64T (64-bit) processor; also compatible with SPARC V9 architecture-based systems.
- Minimum 256 MB of system memory (512 MB recommended).
- A hard drive with at least 400 MB available free space; Solaris® 10 x86 and Solaris 10 SPARC, Solaris 11 x86 and Solaris 11 SPARC requires a minimum of 640 MB.

The supported operating systems for the MegaRAID Storage Manager software are as follows:

- Microsoft® Windows Server® 2003, Microsoft Windows Server 2008, Microsoft Windows Server 2008 R2, Microsoft Windows® XP, Microsoft Windows Vista®, Microsoft Windows 7, Microsoft Windows 8, Microsoft Windows 8.1, Microsoft Windows 8.1 Update, Windows 10, and Microsoft Windows Server 2012.

  NOTE Support of the SNMP Agent is deprecated during the default installation of the MegaRAID Storage Manager software on Microsoft Windows 8.1, Microsoft Windows 8.1 Update, Microsoft Windows 2012, and later versions. However, in a custom installation, if you select SNMP as one of the utilities, the SNMP Agent is installed.

- Oracle® Enterprise Linux® 5 U6 and U7, Oracle Enterprise Linux 6 and U1, and Oracle Enterprise Linux 5.10, 6.5, and 7.0.
- Red Hat® Linux (RHEL) 3.0, 4.0, 5.0, 5.8, 5.9, 5.10, 5.11, 6.0, 6.5, 6.6, 6.7, 7.0, and 7.1. The MegaRAID Storage Manager software supports 64-bit environment from RHEL 6 onwards.
- Solaris® 10 x86, Solaris SPARC, Solaris 11 x86, Solaris 11 SPARC, Solaris 11 Update 1 x86, and Solaris 11 Update 1 SPARC.
- SUSE® Linux/SLES 9, 10, 11, 11 SP2, 11 SP3, 11 SP4, and 12 with the latest updates and service packs.
- VMware® ESX 4.0 and 4.1.
- VMware ESXi 4.0, 4.1, 5.0, 5.0 Update 2, 5.1, 5.1 Update 1, 5.1 Update 3, 5.5, 5.5 Update 2, and 6.0.
- UEK R3 Update 3 for Oracle Linux 6.4 (64 bit and later) and UEK R3 Update 4.

Refer to your server documentation and to the operating system documentation for more information on hardware and operating system requirements.

NOTE The MegaRAID Storage Manager software also is supported in the Network Address Translation (NAT) environment. If the server is installed in a remote machine and you want to connect to that server over a NAT environment, through a remote client, you can connect to the remote server by providing the NAT IP address.

NOTE The MegaRAID Storage Manager software uses the local IP address in the same subnet as the SMTP server to deliver email notifications to the SMTP server.

You can use the MegaRAID Storage Manager software to remotely monitor the systems running the VMware ESXi (3.5 and above) operating system.

NOTE Storelib libraries need the capability to be installed with more than one version. All the Storelib libraries have been moved to a private location. Perform a clean uninstallation and then install only the MegaRAID Software Manager software to avoid any conflicts.

7.3 Installing MegaRAID Storage Manager

This section explains how to install (or reinstall) the MegaRAID Storage Manager software on your workstation or on your server for the supported operating systems: Microsoft Windows, Red Hat Linux, SuSE Linux, Solaris 10 x86, and Solaris SPARC.

7.3.1 Prerequisite for MegaRAID Storage Manager Installation

The MegaRAID Storage Manager software installation script also installs the Avago SNMP agent, Red Hat Package Manager (RPM). The Avago SNMP agent application depends upon the standard SNMP-Util package.

Make sure that the SNMP-Util package is present in the system before you install the MegaRAID Storage Manager software.

The SNMP-Util package includes the `net-snmp-libs` and the `net-snmp-utils` RPMs and additional dependent RPMs. Make sure that these RPMs are installed from the operating system media before you install the MegaRAID Storage Manager software.

7.3.2 Installing the MegaRAID Storage Manager Software on Microsoft Windows

Perform the following steps to install the MegaRAID Storage Manager software on a system running the Microsoft Windows Server 2003, Microsoft Windows Server 2008, Microsoft Server 2008 R2, Microsoft Windows XP, Microsoft Windows Vista, Microsoft Windows 7, Microsoft Windows 8, Microsoft Windows 8.1, Microsoft Windows 8.1 Update, Windows 10, or Microsoft Windows Server 2012 operating system:

1. Insert the MegaRAID Storage Manager software installation CD in the CD-ROM drive.

   If necessary, find and double-click the `setup.exe` file to start the installation program.
2. In the **Welcome** screen that appears, click **Next**.

   If the MegaRAID Storage Manager software is already installed on this system, then an upgraded installation occurs.
3. Read and accept the user license and click **Next**.

   The **Customer Information** window appears, as shown in the following figure.

Figure 146 Customer Information Window

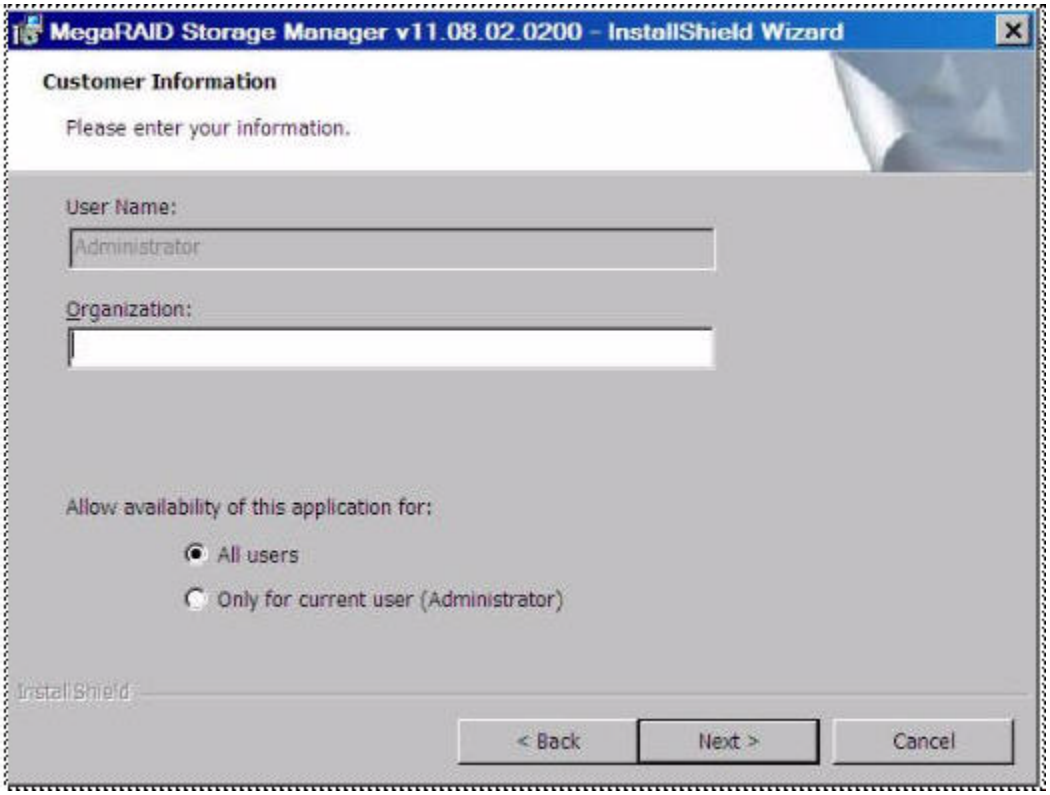

4. Enter your user name and organization name.

   In the bottom part of the screen, select an installation option:

   — If you select the **All users** radio button, any user with administrative privileges can use this version of the MegaRAID Storage Manager software to view or change storage configurations.

   — If you select the **Only for current user** radio button, the MegaRAID Storage Manager software shortcuts and associated icons are available only to the user with this user name.

5. Click **Next** to continue.
6. Accept the default destination folder, or click **Change** to select a different destination folder, as shown in the following figure.

Figure 147 Destination Folder Window

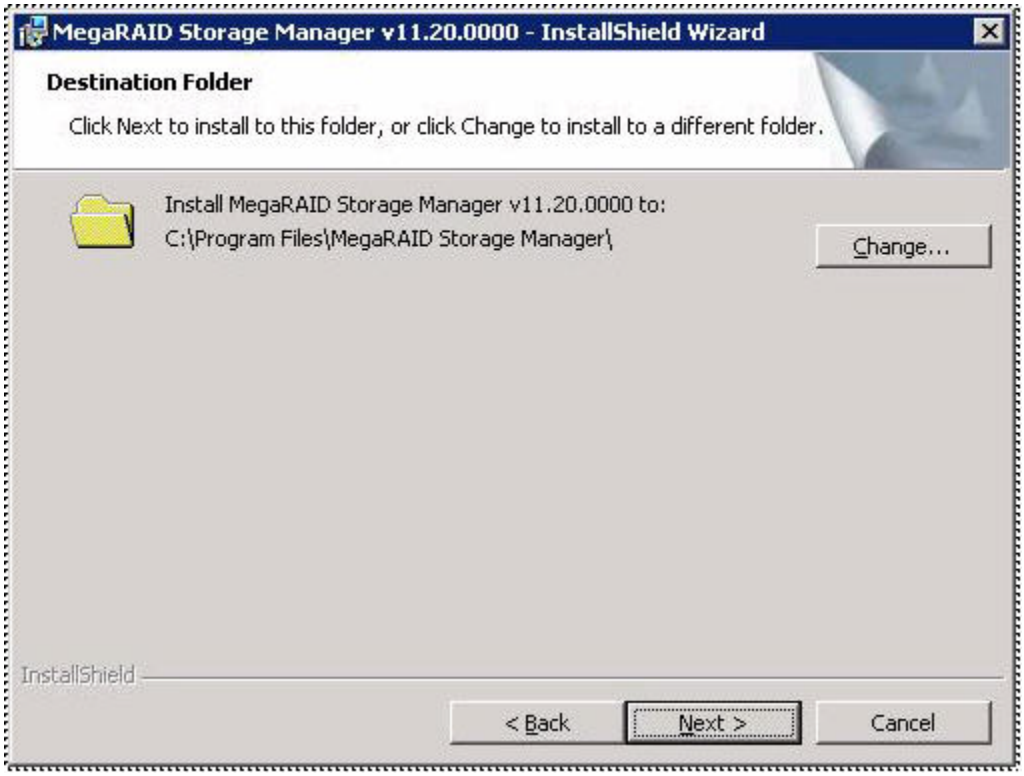

7. Click **Next** to continue.

   The **Setup Type** window appears, as shown in the following figure.

Figure 148 Setup Type Window

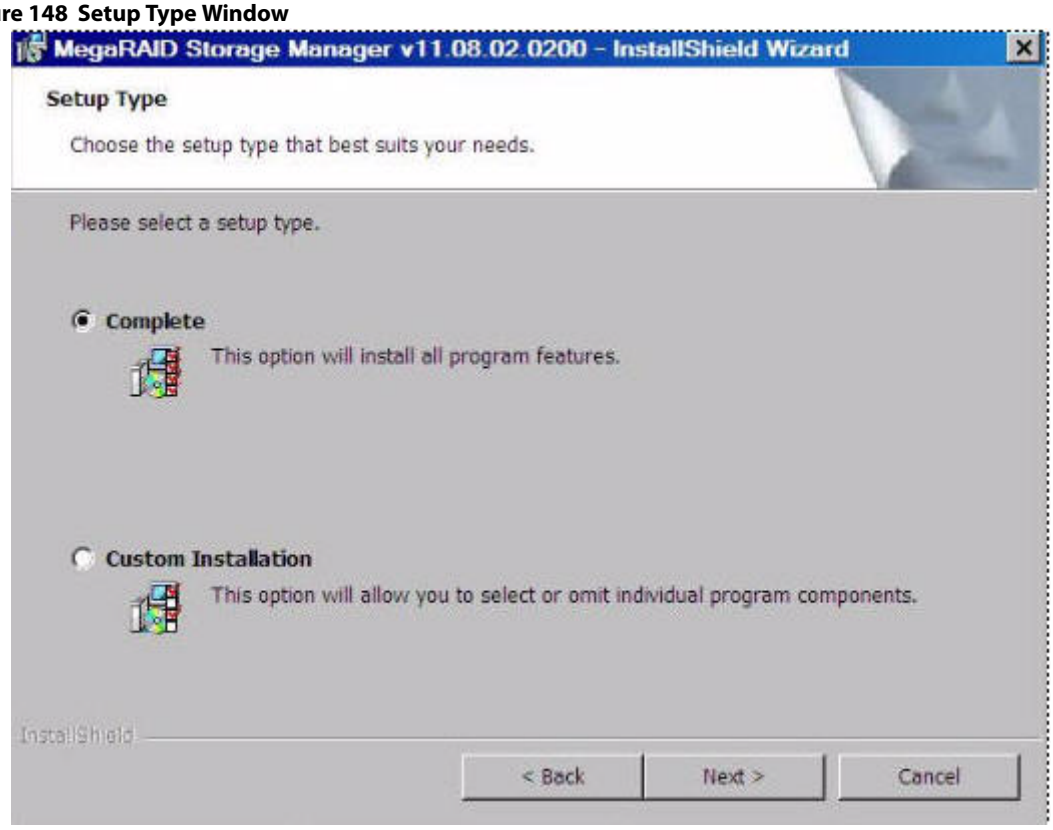

8. Select one of the setup options.

   The options are fully explained in the window text.

   — Select the **Complete** radio button if you are installing the MegaRAID Storage Manager software on a server.

   — Select the **Custom Installation** radio button if you want to select individual program components.

9. Click **Next** to continue.

   If you select **Custom Installation** as your setup option, the second **Setup Type** dialog appears, as shown in the Custom Setup Window figure. If you select **Complete** as your setup option, the **LDAP Login Information** dialog appears.

Figure 149 LDAP Logon Information

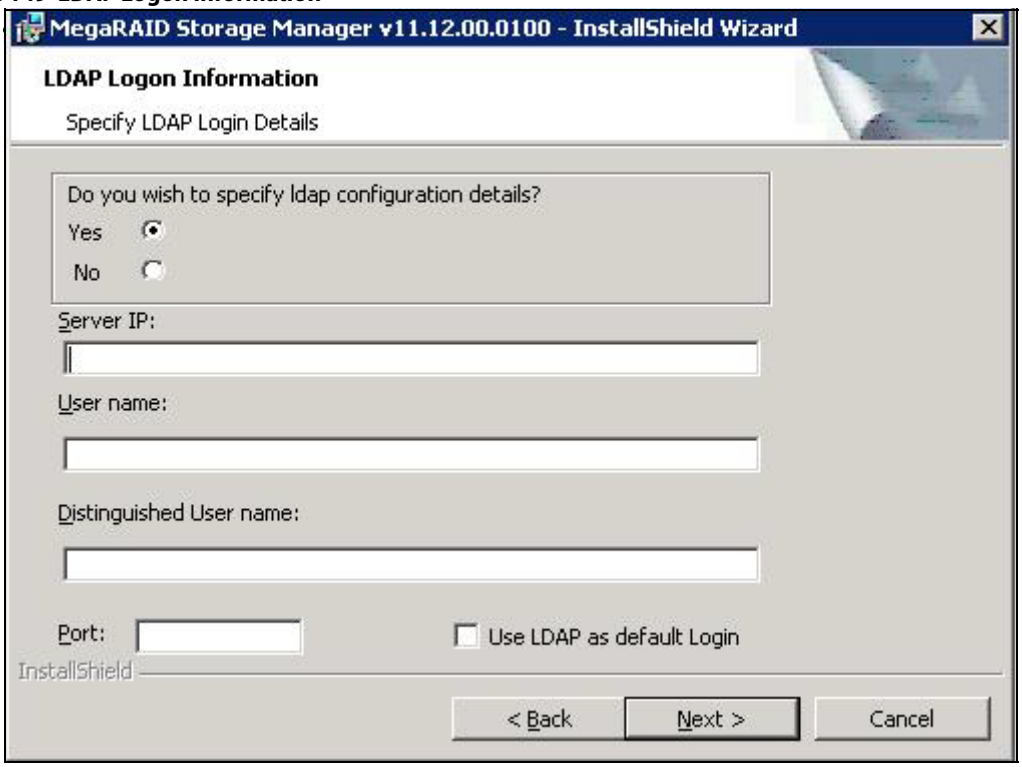

10. To specify LDAP configuration details, select **Yes**, and perform the following substeps, or if you do not want to specify LDAP configuration details, click **No** and click **Next**.
    a. Enter the LDAP server's IP address in the **Server IP** field.
    b. Enter the LDAP server's user name in the **User name** field.
       An example of a user name can be `username@testldap.com`.
    c. Enter the name of the Domain Controller in the **Distinguished User name** field.
       As an example, the Domain Controller name can be `dc= TESTLDAP, dc=com`.
    d. Enter the LDAP server's port number in the **Port** field.
    e. Select the **Use LDAP as default Login** check box to always connect to the LDAP server.

    All the values entered in this dialog are saved in the `ldap.properties` file.
11. Click **Next**.
    The **SelfSigned key details** dialog appears.
12. In the drop-down list, click either **2048** or **1024**.
13. Click **Next**.
    The **Alert notifications of user choice** dialog appears.

14. Select one of the options. The options are explained in the window text.
    - **Select of Last reboot** – If you select this option, the MegaRAID Storage Manager software retrieves events from the last reboot.
    - **Select of Last reboot** – If you select this option, the MegaRAID Storage Manager software retrieves events from the last log clear.
    - **Select of Last Shutdown** – If you select this option, the MegaRAID Storage Manager software retrieves events from the last clean shutdown.

    NOTE These options work only if the MegaRAID Storage Manager software alert notification history files (`SASAdapterInfo_<adapter_index>`) are not found.
15. Click **Next**.
16. In the dialog that appears, click **Install** to begin the installation.
17. Select one of the setup options.

    See Setup Options for specific information.

Figure 150 Custom Setup Window

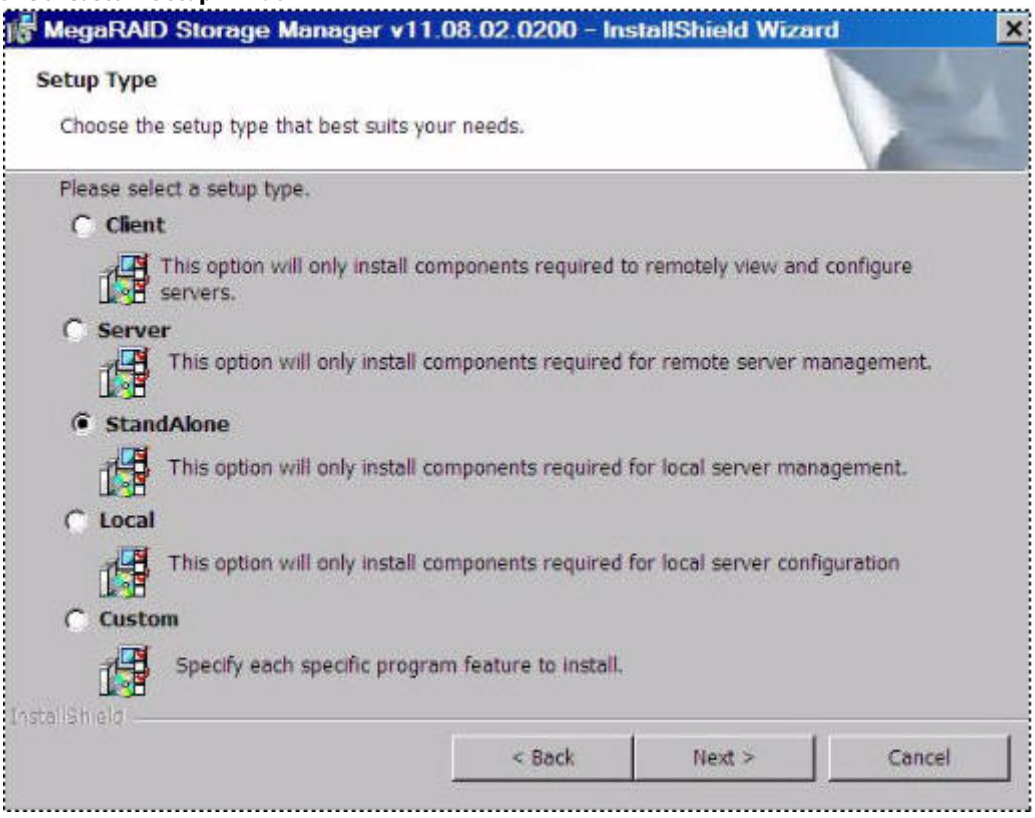

18. Click **Next** to proceed.
19. Click **Install** to install the program.
20. When the final **Configuration Wizard** window appears, click **Finish**.

    If you select **Client** installation for a computer that is used to monitor servers, and if no available servers exist with a registered framework on the local subnet (that is, servers with a complete installation of the MegaRAID Storage Manager software), the server window appears. The **MegaRAID Storage Manager – Host View** window does not list any servers. You can use the **MegaRAID Storage Manager – Host View** window to manage systems remotely.

7.3.2.1 Setup Options

The MegaRAID Storage Manager software enables you to select from one of the following setup options when you install it:

- Select the **Client** radio button if you are installing the MegaRAID Storage Manager software on a computer that will be used to view and configure servers over a network.

  To begin installation, click **Install** on the next window that appears.

  In the Client mode of installation, the MegaRAID Storage Manager software installs only client-related components, such as the MegaRAID Storage Manager GUI.

  Use this mode when you want to manage and monitor servers remotely. When you install the MegaRAID Storage Manager software in Client mode on a laptop or a desktop, you can log in to a specific server by providing the IP address.

- Select the **Server** radio button to install only those components required for remote server management.

  To begin installation, click **Install** on the next window that appears.

- Select the **StandAlone** radio button if you will use the MegaRAID Storage Manager software to create and manage storage configurations on a stand-alone workstation.

  NOTE If you select *Client* or *Standalone* as your setup option, the LDAP Logon Information dialog appears.

  To begin installation, click **Install** on the next window that appears.

- Select the **Local** radio button if you want to view only the workstation that has the MegaRAID Storage Manager software installed.

  You will not be able to discover other remote servers and other remote servers will also not be able to connect to your workstation. In a local mode installation, you will be using the loopback address instead of the IP address.

- Select the **Custom** radio button if you want to specify individual program features to install.

  If you select **Custom**, a window listing the installation features appears. Select the features you want on this window.

7.3.3 Uninstalling the MegaRAID Storage Manager Software on Microsoft Windows

You can uninstall the MegaRAID Storage Manager software from a system running on Microsoft Windows operating system through the Control Panel, the command prompt, or the MegaRAID Storage Manager uninstallation utility.

7.3.3.1 Uninstalling the MegaRAID Storage Manager Software through the Control Panel

To uninstall the MegaRAID Storage Manager software through the Control Panel, follow these steps:

1. Select **Add/Remove Programs** from the **Control Panel**.
2. Select MegaRAID Storage Manager from the list of the **Add/Remove Programs** window.
3. Click **Remove**.

7.3.3.2 Uninstalling the MegaRAID Storage Manager Software Using the Command Prompt

To uninstall the MegaRAID Storage Manager software using the command prompt, follow these steps:

1. Go to the command prompt.
2. Go to the folder `MSM_INSTALLATION_FOLDER`.

3. Run either of the two commands in the Command Prompt:
   - `Uninstaller.exe` (for interactive mode of uninstallation).
   - `Uninstaller.exe -silent` (for Silent uninstallation).

7.3.3.3 Uninstalling the MegaRAID Storage Manager Software Using the MegaRAID Storage Manager Uninstallation Utility

To uninstall the MegaRAID Storage Manager software through the MegaRAID Storage Manager uninstallation utility, follow these steps:

1. Go to **Start > MegaRAID Storage Manager**.
2. Click **MegaRAID Storage Manager Uninstall**.
3. Follow the prompts to complete the uninstallation procedure.

7.3.4 Installing and Supporting the MegaRAID Storage Manager Software on Solaris and SPARC Operating Systems

This section discusses the installation of the MegaRAID Storage Manager software on the Solaris 10 (U5, U6, U7, U8, U9, and U10), Solaris 11 (x86 and x64) and Solaris SPARC operating systems.

7.3.4.1 Installing the MegaRAID Storage Manager Software for Solaris 10 x86

This section documents the installation of the MegaRAID Storage Manager software on the Solaris 10 U5, U6, U7, U8 x86 and x64 operating systems.

Follow these steps to install the MegaRAID Storage Manager software on a system running the Solaris 10 x86 operating system:

1. Copy the `MegaRaidStorageManager-SOLX86-…..tar.gz` file to a temporary folder.
2. Untar the `MegaRaidStorageManager-SOLX86-…..tar.gz` file using the following command:

   ```
   tar -zxvf MegaRaidStorageManager-SOLX86-…..tar.gz
   ```

   This step creates a new disk directory.
3. Go to the new disk directory, and find and read the `readme.txt` file.
4. Enter the Bash shell.
5. Execute the command `./install.sh` present in the disk directory.
6. When prompted by the installation scripts, select Y to complete the installation.

7.3.4.2 Installing the MegaRAID Storage Manager Software for Solaris 10 SPARC

Perform the following steps to install the MegaRAID storage Manager software for Solaris 10 SPARC.

1. Copy the `MegaRaidStorageManager-SOLSPARC-8.10-…….tar.gz` file to a temporary folder.
2. Untar the `MegaRaidStorageManager-SOLSPARC-8.10-…….tar.gz` file using the following command:

   ```
   tar -zxvf MegaRaidStorageManager-SOLSPARC-8.10-…….tar.gz
   ```

   This step creates a new disk directory. Go to the new disk directory, and find and read the `readme.txt` file.
3. Enter the Bash shell.
4. Execute the command `./install.sh` present in the disk directory.
5. When prompted by the installation scripts, type `Y` to complete the installation.

NOTE The MegaRAID CacheCade Pro 2.0 software is not applicable in SPARC.

7.3.4.3 Installing the MegaRAID Storage Manager Software for Solaris 11 x86

Follow these steps to install the MegaRAID Storage Manager software on a system running the Solaris 11 x86 operating system.

1. Copy the `MegaRaidStorageManager-SOL11X86-.....tar.gz` file to a temporary folder.
2. Untar the `MegaRaidStorageManager-SOL11X86-.....tar.gz` file using the following command.

   `tar -zxvf MegaRaidStorageManager-SOL11X86-.....tar.gz`

   This step creates a new disk directory.
3. Go to the new disk directory, and read the `readme.txt` file.
4. Enter the Bash shell.
5. Execute the `./install.sh` command present in the disk directory.
6. When prompted by the installation scripts, type `Y` to complete the installation.

7.3.4.4 Installing the MegaRAID Storage Manager Software for Solaris 11 SPARC

Follow these steps to install the MegaRAID Storage Manager software on a system running Solaris 11 SPARC:

1. Copy the `MegaRaidStorageManager-SOL11SPARC-.....tar.gz` file to a temporary folder.
2. Untar the `MegaRaidStorageManager-SOL11SPARC-.....tar.gz` file using the following command:

   `tar -zxvf MegaRaidStorageManager-SOL11SPARC-.....tar.gz`

   This step creates a new disk directory.
3. Go to the new disk directory 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, type `Y` to complete the installation.

NOTE The Avago MegaRAID CacheCade Pro 2.0 software is not applicable in SPARC.

7.3.5 Uninstalling the MegaRAID Storage Manager Software on Solaris 10 (U5, U6, U7, U8, U9, and U10), Solaris 11 (x86 and x64), and Solaris SPARC

Follow these steps to uninstall the MegaRAID Storage Manager software on a system running Solaris operating systems:

1. Run the `Uninstaller.sh` file located in `/opt/MegaRaidStorageManager` directory.
2. When prompted by the uninstallation scripts, select `Y` to complete the installation.

To shut down the MegaRAID Storage Manager Framework service, run the `svcadm disable -t MSMFramework` command.

To start the Framework service, run the `svcadm enable MSMFramework` command.

When the service is in maintenance state, run the `svcadm clear MSMFramework` command.

To check the status of the MegaRAID Storage Manager services, run the `svcs-a|grep -i msm` command.

7.3.6 Prerequisites for Installing the MegaRAID Storage Manager Software on RHEL6.x x64 and RHEL7.x x64

Before installing the MegaRAID Storage Manager software on RHEL6.*x* x64 and RHEL7.*x* x64 operating systems, install the following RPMs. Without these RPMs, the MegaRAID Storage Manager software might not install correctly or might not work as expected.

- `libstdc++-4.4.4-13.el6.i686.rpm`
- `compat-libstdc++-33-3.2.3-69.i686.rpm`
- `libXau-1.0.5-1.el6.i686.rpm`
- `libXcb-1.5-1.el6.i686.rpm`
- `libX11-1.3-2.el6.i686.rpm`
- `libXext-1.1-3.el6.i686.rpm`
- `libXi-1.3-3.el6.i686.rpm`
- `libXtst-1.0.99.2-3.el6.i686.rpm`

The RHEL6.*x* x64 and RHEL7.*x* x64 operating systems installation is required for the MegaRAID Storage Manager software to work. The previous list of RPMs come as part of RHEL6.*x* x64 and RHEL7.*x* x64 Operating System DVDs. These RPMs might need additional dependent RPMs as well, and you must install all the dependent RPMs on the target system.

NOTE The RPM versions listed previous might change in future RHEL6.*x* x64 and RHEL7.*x* x64 releases. Install the corresponding RPMs from the operating system installation media.

NOTE The MegaRAID Storage Manager software currently provides an additional binary to run it in a native 64-bit Linux environment.

7.3.7 Installing the MegaRAID Storage Manager Software on RHEL or SLES/SuSE Linux

Follow these steps if you need to install the MegaRAID Storage Manager software on a system running the Red Hat Linux 3.0, 4.0, 5.0, 6.0, 6.6, 6.7, or 7.0 operating system or the SLES 9, 10, 11, 11 SP2, 11 SP3, 11 SP4, or 12 operating system:

NOTE To install the MegaRAID Storage Manager software on a SLES 64-bit platform, you need to create certain symbolic links that are mentioned in Running a CIM Plug-in on Red Hat Enterprise Linux 5.

1. Copy the `MSM_linux_installer-11.02.00-00.tar.gz` file to a temporary folder.
2. Untar the `MSM_linux_installer-11.02.00-00.tar.gz` file using the following command:

   ```
   tar -zxvf MSM_linux_installer-11.02.00-00-...tar.gz
   ```

   A new disk directory is created.
3. Go to the new `disk` directory.
4. In the `disk` directory, find and read the `readme.txt` file.
5. To start the installation, enter the following command:

```
csh install.csh -a
```

The preceding command works only if the `csh` shell is installed; otherwise, use the following command:

```
install.csh
```

If you select **Client** installation for a computer that is used to monitor servers, and if no available servers exist with a registered framework on the local subnet (that is, servers with a complete installation of the MegaRAID Storage Manager software), the **MegaRAID Storage Manager – Host Name** window appears. The **MegaRAID Storage Manager – Host Name** window does not list any servers. You can use this window to manage systems remotely.

To install the software using an interactive mode, execute the command `./install.csh` from the installation disk.

To install the product in a non-interactive or silent mode, use the command `./install.csh [-options] [-ru popup]` from the installation disk. The installation options are as follows:

- **Complete**
- **Client Component Only**
- **StandAlone**
- **Local**
- **Server**

The `-ru popup` command removes the pop-up from the installation list.

You also can run a non-interactive installation using the `RunRPM.sh` command.

The installer offers the following setup options:

- **Complete**

  This option installs all the features of the product.
- **Client Components Only**

  The StoreLib feature of the product is not installed in this type of installation. As a result, the resident system can only administer and configure all of the servers in the subnet, but it cannot serve as a server.
- **StandAlone**

  Only the networking feature will not be installed in this case. But the system can discover other servers in the subnet and can be discovered by the other servers in the subnet.
- **Local**

  This option lets you view only the workstation that has the MegaRAID Storage Manager software installed. You will not be able to discover other remote servers and other remote servers will also not be able to connect to your workstation. In a local mode installation, you will be using the loopback address instead of the IP address.
- **Server**

  This option installs components required for remote server management.

This installation helps you select any of the setup types, but if you run `RunRPM.sh`, it installs the complete feature.

NOTE To install and run the MegaRAID Storage Manager software on RHEL 5, you must disable SELinux.

7.3.8 Linux Error Messages

The following messages can appear while you are installing the MegaRAID Storage Manager software on a Linux operating system:

- `More than one copy of MegaRAID Storage Manager software has been installed.`

  This message indicates that the user has installed more than one copy of the MegaRAID Storage Manager software. (This step can be done by using the `rpm-force` command to install the `rpm` file directly, which is not recommended, instead of using the `install.sh` file.) In such cases, the user must uninstall all of the `rpm` files manually before installing the MegaRAID Storage Manager software with the procedure listed previously.

- `The version is already installed.`

  This message indicates that the version of the MegaRAID Storage Manager software you are trying to install is already installed on the system.

- `The installed version is newer.`

  This message indicates that a version of the MegaRAID Storage Manager software is already installed on the system, and it is a newer version than the version you are trying to install.

- `Exiting installation.`

  This is the message that appears when the installation is complete.

- `RPM installation failed.`

  This message indicates that the installation failed for some reason. Additional message text explains the cause of the failure.

7.3.9 Kernel Upgrade

If you want to upgrade the kernel in the Linux operating system, you must restart the MegaRAID Storage Manager Framework and Services in the same order by entering the following command.

```
/etc/init.d/vivaldiframeworkd restart
```

7.3.10 Uninstalling the MegaRAID Storage Manager Software on RHEL or SLES or SuSE Linux

To uninstall the MegaRAID Storage Manager software on a system running Linux, follow these steps:

1. Go to `/usr/local/MegaRAID Storage Manager`.
2. Run `./uninstaller.sh`.

This procedure uninstalls the MegaRAID Storage Manager software.

7.3.11 MegaRAID Storage Manager Software Customization

You can customize your Logo and Splash window by editing the `msm.properties` file present in the *<installation-directory\MegaRAID Storage Manager>* folder.

The `msm.properties` file has no values for the following keys:

- `CHANNELLOGO=`
- `CHANNELSPLASHSCREEN=`

No default values are assigned for these keys; therefore, the MegaRAID Storage Manager software uses the default Avago Logo and splash screen.

To customize the Logo and splash screen, enter the Logo and Splash screen file name against these entries.

To enter the file names follow these steps:

1. Open the `msm.properties` file in the *<installation-directory\MegaRAID Storage Manager>* folder.
2. Enter the value for the logo file against the `CHANNELLOGO` key.
3. Enter the value for the splash screen file against the `CHANNELSPLASHSCREEN` key.
4. Save the file.
5. Place these two images in the *<installation-directory\MegaRAID Storage Manager>* folder.
6. Start the application.

Following are some of important points that you need to keep in mind:

- File names for both entries should not have any spaces.

  For example, the valid file name would be: `logo_test_1.png`, `LogoTest1.png`, or `TEST_SPLASH_FILE.jpg`.
- The logo image dimensions should not exceed 160 pixels x 85 pixels (width x height).
- The splash screen image dimensions should not exceed 390 pixels x 260 pixels (width x height).

After making the changes mentioned previously, when you log into the MegaRAID Storage Managers software, you will be able to view the changes with the new splash screen and logo in the MegaRAID Storage Manager software.

7.3.12 Updating the Strength of Public and Private RSA keys

The size of the RSA public key for the MegaRAID Software Manager software is upgraded from 1024 to 2048. The public key communicates with the client (GUI) and the server (Framework Service and Pop-up Process). Because of this change, you must upgrade the renewal process across the client (GUI) and the server (Framework Service and Pop-up Process) with the new key. The same subnet might have a mixed-mode of installation, which might have the latest software version with the renewed key size of 2048 and older versions with the key size of 1024. It also might contain different versions of clients and servers. This situation poses a compatibility issue in the mixed-mode installation because the old public key (size 1024) cannot perform a handshake with latest public key (2048) as an artifact of cryptography. To address this compatibility issue, with a known limitation that both the keys cannot be loaded concurrently, you must edit the `vivaldikey.properties` file for vivaldikey selection to manage both the old and the new installation.

The following figure shows how the handshake works after editing the vivaldikey.

Figure 151 Vivaldikey Handshake

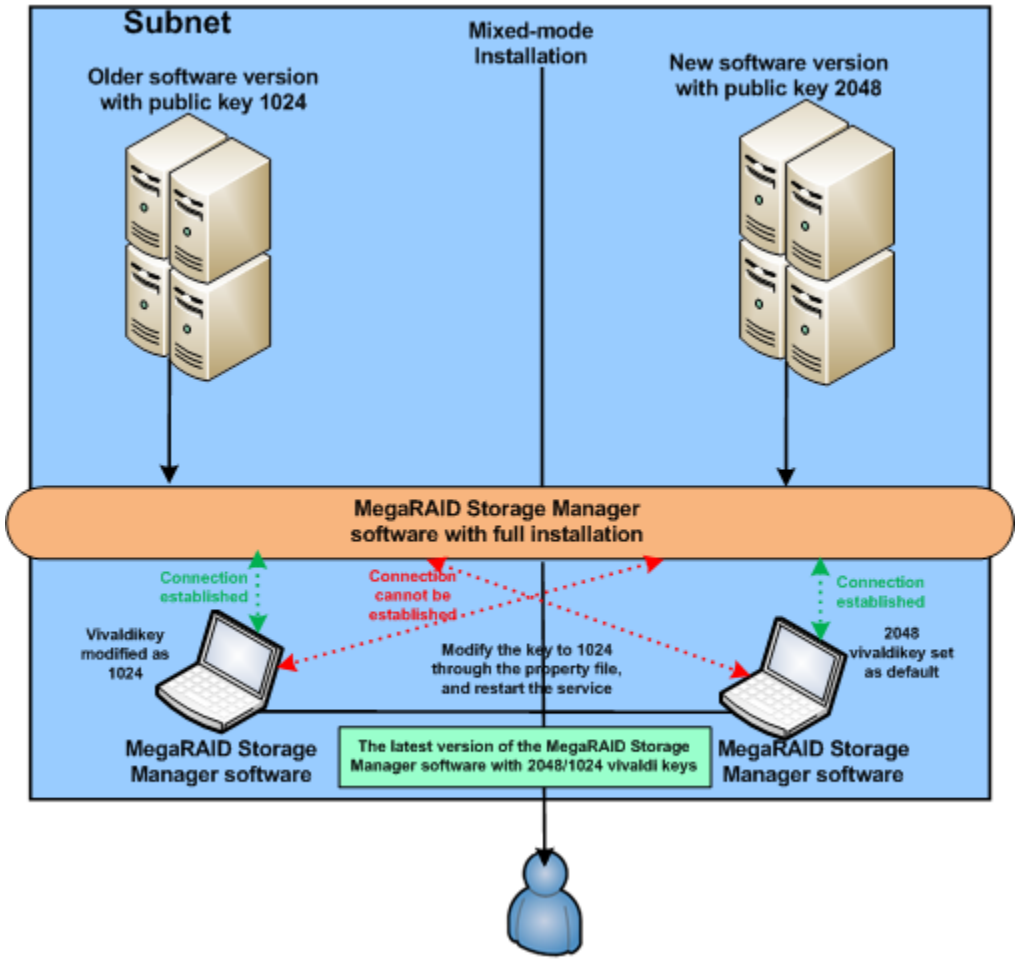

7.3.12.1 Limitations

The following are the limitations:

- The latest installation version is capable of managing both the old installation and the new installation sequentially, which is followed by a property switch and the MSM Framework Service restart.
- The old installation version cannot manage the newly installed servers.

Table 58 Expected Behavior

Client	Server	Expected Behavior
The latest version of the MegaRAID Storage Manager software (vivaldikey 2048).	The latest version of the MegaRAID Storage Manager software (vivaldikey 2048).	The server and the client are connected.
The old version of the MegaRAID Storage Manager software (vivaldikey 1024).	The old version of the MegaRAID Storage Manager software (vivaldikey 1024).	The server and the client are connected.
The latest version of the MegaRAID Storage Manager software (vivaldikey 2048).	The old version of the MegaRAID Storage Manager software (vivaldikey 1024).	The connection fails.
The old version of the MegaRAID Storage Manager software (vivaldikey 1024).	The latest version of the MegaRAID Storage Manager software (vivaldikey 2048).	The connection fails.

7.3.12.2 Updating the Property File and Vivaldikeys

Follow these steps to update the property file and the vivaldikeys.

1. Stop the Pop-up Process.
2. Open the `vivaldikey.properties` file from all of the following locations:
   - `C:\Program Files (x86)\MegaRAID Storage Manager`
   - `C:\Program Files (x86)\MegaRAID Storage Manager\Framework`
   - `C:\Program Files (x86)\MegaRAID Storage Manager\MegaPopup`
3. Edit the `VIVALDI_KEY_FILE` field to `vivaldikeys_2048`.
4. Save the `vivaldikey.properties` file.
5. Start the MegaRAID Storage Manager Framework service.
6. Start the Pop-up Process.

NOTE The latest installation version helps you to manage the old and the new installations. However, the old installation versions of the servers cannot manage the latest installed servers.

7.3.13 Stopping the Pop-Up Notification Process

The pop-up notification is started automatically when you login to the operating system. To stop the pop-up notification, you must follow steps based on your operating system.

7.3.13.1 Windows Operating System

To stop the pop-up notification process on the Windows operating system, follow these steps:

1. Go to the command prompt.
2. Navigate to the *<MSM_INSTALLATION_FOLDER>*`\MegaPopup` folder.
3. Run the command, `popup -stop`.
   After running the preceding command, the pop-up process stops.

7.3.13.2 Linux, Solaris x86, and Solaris SPARC Operating Systems

To stop the pop-up notification process on Linux, Solaris x86, or Solaris SPARC operating systems, follow these steps:

1. Go to the command prompt.
2. Go to the *<MSM_INSTALLATION_FOLDER>*`\MegaPopup` folder.
3. Run the script, `shutdownpopup -sh` in the console.
   After running the preceding command, the pop-up process stops.

7.3.14 Restarting the Pop-Up Notification Process

When you restart the MegaRAID Storage Manager Framework Service in Windows, Linux, Solaris x86, or Solaris SPARC operating systems, and if you want to see the pop-up notifications, you need to start the Pop-up Process.

- For the Windows operating system, you must first stop the pop-up process (see Windows Operating System) and then restart the same.
  After stopping the pop-up process, run the `Popup.exe` command in the same console. The pop-up process is started again.

- For the Linux operating system, you must first stop the pop-up process (see Linux, Solaris x86, and Solaris SPARC Operating Systems) and then restart the same.

  After stopping the pop-up process, run the `./popup&` command from the same console. The pop-up process is started again.
- For the Solaris x86 or Solaris SPARC operating system, you must first stop the pop-up process (see Linux, Solaris x86, and Solaris SPARC Operating Systems) and then restart the same.

  After stopping the pop-up process, run the `./popup` command from the same console. The pop-up process is started again.

7.4 Installing and Supporting the MegaRAID Storage Manager Software on VMware

This section documents the installation of the MegaRAID Storage Manager software on VMware ESX (also known as Classic) and on the VMware ESXi operating system.

7.4.1 Prerequisites for Installing the MegaRAID Storage Manager for VMware

For the VMware 3.5 operating system, it is necessary to install the `libstdc++34-3.4.0-1.i386.rpm` file before installing the MegaRAID Storage Manager software. You can download the `rpm` file from: `http://rpm.pbone.net/index.php3/stat/4/idpl/1203252/com/libstdc++34-3.4.0-1.i386.rpm.html`.

For the VMware 4.1 operating system, it is necessary to create a soft link as follows before installing the MegaRAID Storage Manager software. Run the following command to create the necessary soft link required for the MegaRAID Storage Manager software to work.

```
sudo ln -sf /lib/libgcc_s.so.1/usr/lib/vmware/lib/libgcc_s.so.1
```

For VMware ESXi 5.0 to work with the MegaRAID Storage Manager software, the SMI-S Provider must be installed.

7.4.2 Installing the MegaRAID Storage Manager Software on VMware ESX (VMware Classic)

The VMware operating system does not support any graphics components. To install the MegaRAID Storage Manager software on the VMware operating system, run the script `./vmware_install.sh` from the installation disk.

NOTE Make sure that on a 32-bit or on a 64-bit VMware operating system, you install the 32-bit MegaRAID Storage Manager software.

The installer lets you accept the license agreement, operating system, and Storelib package as follows:

- End user license agreement
- Operating system (VMware 4.*x* operating system)
- Select the Storelib (Inbox Storelib or Storelib from the MegaRAID Storage Manager package)

NOTE VMware Classic is not supported on VMware 5.*x* and higher versions.

7.4.3 Uninstalling the MegaRAID Storage Manager Software for VMware

To uninstall the Server Component of the MegaRAID Storage Manager software on VMware, either use the `Uninstall` command in the Program menu, or run the script `/usr/local/MegaRAID Storage Manager/uninstaller.sh`.

You must keep in mind the following points:

- A MegaRAID Storage Manager upgrade is supported in this release.

  Future releases can update this release.
- To shut down the MegaRAID Storage Manager Framework service, run the following command:

```
/etc/init.d/vivaldiframeworkd stop
```

The Linux RPM of the MegaRAID Storage Manager software works under the console with minimal changes. Hardware RAID is currently supported in ESX 4.x.

NOTE There is a known limitation that virtual drives that are created or deleted will not be reflected to the kernel. The workaround is to reboot the server or to run the `esxcfg-rescan <vmhba#>` command from a COS shell.

7.4.4 MegaRAID Storage Manager Support on the VMware ESXi Operating System

This section outlines the product requirements needed to support the VMware ESXi operating system. Classic VMware includes a service console that is derived from the Linux 2.4 kernel, but with reduced functionality.

The MegaRAID Storage Manager server part cannot be installed directly in the VMware ESXi operating system. Management is performed through the MegaRAID Storage Manager software installed on a Linux/Windows machine in the same subnet.

NOTE For VMware ESXi 5.0 to work with the MegaRAID Storage Manager software, the SMI-S Provider must be installed.

Remote management of VMware ESXi is supported only in a complete installation of the MegaRAID Storage Manager software on the following operating systems:

- Microsoft Windows Server
- RHEL
- SUSE Linux

Network communication is a key element for a proper communication between the ESXi CIM provider and the Avago management software. Make sure that the network settings are correct by making the following changes:

- Provide a proper host name and an IP address while performing the initial configurations for the ESXi host.
- For networks that do not have DNS configured, the “hosts” file in the machine on which the MegaRAID Storage Manager software is installed must be edited as follows:
  a. Add an entry to map the VMware host’s IP address with the host name.

     This is for the discovery process to happen correctly. In the absence of this entry, the VMware host would be discovered as 0.0.0.0.
  b. Add an entry to map the actual IP address of the localhost with its hostname (an entry for the loopback address would be present by default in the hosts file and it should not be removed).

     This is to ensure that the Asynchronous Event Notifications (AENs) are delivered correctly.

     For example, if 135.24.228.136 is the IP address of your VMware host and 135.24.228.137 is the IP address of your Linux host, the following entries must be added in the hosts file:

```
135.24.228.136 dhcp-135-24-228-136.lsi.com dhcp-135-24-228-136 #VMware
135.24.228.137 dhcp-135-24-228-137.lsi.com dhcp-135-24-228-137 #Linux
```

7.4.5 Limitations of Installation and Configuration

The following are the limitations of this installation and configuration.

- No status information exists for the controller.
- Events are collected as long as the MegaRAID Storage Manager software runs on the client.
- The MegaRAID Storage Manager software on VMware responds slower as compared to the response of the MegaRAID Storage Manager software on the Windows, Linux, or Solaris operating systems.

  Events are collected from the time a client logs in to an ESXi machine for the first time, and it continues to be collected as long as the Framework is running.

7.4.5.1 Differences in the MegaRAID Storage Manager Software for VMware ESXi

The following are some of the differences in the MegaRAID Storage Manager utility when you manage a VMware server.

- The following limitations apply to the system information exposed through the application:
  - Only the IP address and the host name appear.
  - No support exists for the controller health information.
  - The OS Name and Controller Health information is displayed as `Unknown` in the **Host View** Screen.

Figure 152 Host View Screen

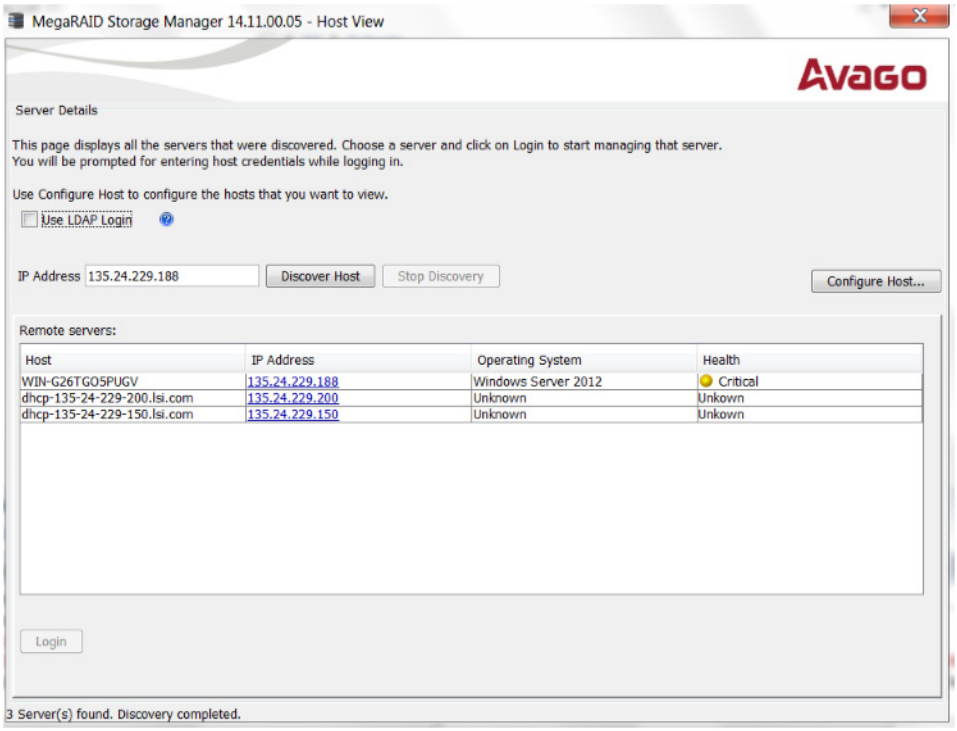

- Authentication support:
  - The MegaRAID Storage Manager software allows CIMOM server authentication with the user ID and the password for VMware.
  - Access to VMware ESXi hosts is controlled based on the user privileges. Only root users can have full access, while the non-root users can have only view only access.
  - Multiple root users can simultaneously login using 'Full Access' mode to access the VMware ESXi server.

- Event logging:

  Event logging support is available for the VMware ESXi operating system, but it works differently than the normal MegaRAID Storage Manager framework mode. The event logging feature for the MegaRAID Storage Manager Client connected to a VMware ESXi system behaves as follows:

  — The System logs are logged in the remote server instead of logging in the ESXi server. For differentiating between the events received from the remote server and the ESXi server, the MegaRAID Storage Manager software appends the ESxi server's IP address on the events received from the ESxi server.
  — The "View Log" option allows you to view the logs saved in a text file on the Event Logger dialog.
  — Refreshing of the MegaRAID Storage Manager GUI after any updates on the firmware is slower for a client connected to VMware ESXi hosts, compared to one that is connected to a Windows/Linux/Solaris host.
- VMware ESXi is supported only on a full installation of the MegaRAID Storage Manager software; standalone, client-only, server-only, and local modes do not support VMware ESXi management.
- VMware ESXi is supported on following operating systems:
  — Microsoft Windows Server
  — RHEL
  — SUSE Linux

7.5 Installing and Configuring a CIM Provider

This section describes the installation and configuration of the MegaRAID Common Information Model (CIM) provider. The Common Information Model offers common definitions of management information for networks, applications, and services, and allows you to exchange management information across systems throughout a network.

On a VMware ESXi system, management is possible only through a CIM provider, and it is performed through the MegaRAID Storage Manager software installed on a remote machine running a Linux or Windows operating system.

The VMware ESXi system comes with the Small Footprint CIM Broker (SFCB) CIM Object Manager (or CIMOM). A CIMOM manages communication between providers, which interact with the hardware, and a CIM client, where the administrator manages the system.

SFCB supports Common Manageability Programming Interface (CMPI)-style providers. CMPI defines a common standard used to interface manageability instrumentation (providers, instrumentation) to management brokers (CIM Object Manager). CMPI standardizes manageability instrumentation, which lets you write and build instrumentation once and run it in different CIM environments (on one platform).

7.5.1 Installing a CIM SAS Storage Provider on the Linux Operating System

The following procedure documents how to install and uninstall the Avago 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 running the `rpm -qa | grep LsiSASProvider` command.

- To install a CIM SAS Storage Provider on a Linux system, install the SAS Provider using the Red Hat Package Manager (RPM) by entering the following command:

  ```
  rpm -ivh
  ```

  The RPM installs all of the necessary files and the Managed Object Format (MOF), and it registers the libraries. The SAS Provider is now ready to use.

NOTE After you install the Avago CIM SAS Provider, the MOF file `LSI_SASRaid.mof` is available under the `/etc/lsi_cimprov/sas/pegasus/common` directory.

- To uninstall a CIM SAS Storage Provider on a Linux system, remove the Avago 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 the `PATH` variable of `/etc/profile`, and hence, are defined in all environments of the system.

7.5.2 Running the CIM SAS Storage Provider on Pegasus

To run the CIM SAS Storage Provider on Pegasus version 2.5.*x*, perform the following steps:

1. After you install the Avago SAS Pegasus provider, verify that the `libLsiSASProvider.so` file and the `libLsiSASProvider.so.1` file are in the `/usr/lib/Pegasus/providers` directory.

   If these files are not present, copy the `libLsiSASProvider.so.1` file from `/opt/tog-pegasus/providers/lib` to `/usr/lib/Pegasus/providers`, and create a `libLsiSASProvider.so` symbolic link to `/usr/lib/Pegasus/providers/libLsiSASProvider.so.1` at `/usr/bin/Pegasus/providers`.
2. Restart the Pegasus CIM Server and Avago Server by performing the following steps:
   - To start the tog-pegasus server, run the following command:

     `# /etc/init.d/tog-pegasus restart`
   - To start LSISAS Server, run the following command:

     `# /etc/init.d/LsiSASd restart`

7.5.3 Installing a CIM SAS Storage Provider on Windows

The following procedure describes how to install and uninstall the Avago CIM SAS Storage Provider on a system running on a Windows operating system.

Perform the following steps to install a CIM SAS Storage Provider on a Windows system:

1. Go to DISK1.
2. Run `setup.exe`.

   The installer installs all of the necessary files and the MOF, and registers the COM DLL. The CIM SAS Provider is now ready to use.

Perform the following steps to uninstall a CIM SAS Storage Provider on a Windows operating system.

1. Select **Control Panel > Add/Remove Program**.
2. Remove the Avago 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.

7.6 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.

NOTE The MegaRAID Storage Manager application uses the local IP address in the same subnet as the SMTP server to deliver email notifications to the SMTP server.

This section describes the installation and configuration of the Avago MegaRAID SNMP agent on Linux, Solaris, and Windows operating systems.

NOTE The complete installation of the MegaRAID Storage Manager software installs the SNMP agent. However, you can install the SNMP agent (installer) on a system separately, without the MegaRAID Storage Manager software being installed.

7.6.1 Prerequisite for the Avago SNMP Agent RPM Installation

The Avago SNMP agent application depends upon the standard SNMP Utilities package. Make sure that the SNMP-Util package is present in the system before you install the Avago SNMP agent RPM.

The SNMP-Util package includes the `net-snmp-libs` and the `net-snmp-utils` RPMs and additional dependent RPMs.

Make sure that these RPMs are installed from the operating system media before you install the Avago SNMP agent RPM.

7.6.2 Installing an SNMP Agent on Windows

This section explains how to install and configure the SAS SNMP Agent for the Windows operating system.

7.6.2.1 Installing an SNMP Agent

Perform the following steps to install an SNMP Agent:

1. Run `setup.exe` from DISK1.
2. Use the SNMP Manager to retrieve the SAS data (it is assumed that you have compiled the `LSI-AdapterSAS.mib` file already).

   The `LSI-AdapterSAS.mib` file is available under the `%ProgramFiles%\LSI Corporation\SNMPAgent\SAS` directory.
3. Use a trap utility to get the traps.

NOTE Before you install the Agent, make sure that the SNMP Service is already installed in the system.

7.6.2.2 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 the **Control Panel**.

2. Select **Add/Remove Windows Components** in the left side of the **Add/Remove Programs** window.
3. Select **Management and Monitoring Tools**.
4. Click **Next**, and follow any prompts to complete the installation procedure.

7.6.2.3 Configuring SNMP Service on the Server Side

Perform the following steps to configure SNMP Service on the server side.

1. Select **Administrative Tools** from the **Control Panel**.
2. Select **Services** in the **Administrative Tools** window.
3. Select **SNMP Service** in the **Services** window.
4. Open the **SNMP Service**.
5. Click the **Security** tab, and make sure that **Accept SNMP Packets from any host** is selected.
6. Click the **Traps** tab, and select the list of host IP addresses to which you want the traps to be sent with the community name.

7.6.2.4 Installing the SNMP Service for the Windows 2008 Operating System

Before you install the Avago Agent, make sure that SNMP Service is already installed in the system.

If you do not have SNMP Service installed on your system, perform the following steps to install SNMP Service for Windows 2008 system.

1. Select **Program and Features** from the **Control Panel**.
2. Click **Turn windows feature on/off** to select the windows components to install.
3. Select **Features** from the menu.
4. Click **Add Features**.
5. Select **SNMP Services**.
6. Click **Next**.
7. Click **Install,** and the SNMP installation starts.

   You will be prompted for the Windows 2008 CD during the installation.
8. Insert the CD, and click **Ok**.

   The installation resumes.

   After the installation is complete, the system displays a message saying that the installation is successful.

7.6.2.5 Configuring the SNMP Service on the Server Side for the Windows 2008 Operating System

To configure SNMP service on the server side for Windows 2008 operating system, perform the following steps:

1. Select **Administrative Tools** from the **Control Panel**.
2. Select **Services** from the **Administrative Tools** window.
3. Select **SNMP Service** from the **Services** window.
4. Open **SNMP Service**, and go to its properties.
5. Go to the **Security** tab, and make sure that **Accept SNMP Packets from any host** is selected.
6. Click the **Traps** tab, and select the list of host IP addresses to which you want the traps to be sent with the community name.

7.6.3 Prerequisite for Installing the SNMP Agent on a Linux Server

To install the SNMP application, you need the `libstdc++.so.6.` library. This library is present in the `/usr/lib` directory.

You can install the SNMP application (`net-snmp`) from the Linux software component RPM that provides these libraries. These RPMs are available in the Linux OS DVD.

7.6.4 Installing and Configuring an SNMP Agent on Linux

This section explains how to install and configure the SAS SNMP Agent for the SuSE Linux and Red Hat Linux operating systems.

Perform the following steps to install and configure the SAS SNMP Agent for the SuSE Linux and Red Hat Linux operating systems:

NOTE This procedure requires that you have the Net-SNMP agent installed on the Linux machine. The RPM has not been created to support -U version. The RPM -U will probably fail with this RPM.

1. Install the Avago SAS SNMP Agent using the `rpm -ivh <sas rpm>` command.

   NOTE Before installation, check whether any pass command exists that starts with 1.3.6.1.4.1.3582 OID in `snmpd.conf`. If so, delete all of the old pass commands that start with 1.3.6.1.4.1.3582 OID. (This situation could occur if an earlier version of the Avago SNMP Agent was installed in the system.)

   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.

   The `snmpd.conf` file structure should be the same as the file structure l`si_mrdsnmpd.conf`. For reference, a sample configuration file (`lsi_mrdsnmpd.conf`) is in the `/etc/lsi_mrdsnmp` directory.

2. To run an SNMP query from a remote machine, add the IP address of that machine in the `snmpd.conf` file, as in this example:

   ```
   com2sec     snmpclient     172.28.136.112     public
   ```

   Here, the IP address of the remote machine is 172.28.136.112.

3. To receive an SNMP trap to a particular machine, add the IP address of that machine in the `com2sec` section of the `snmpd.conf` file.

   For example, to get a trap in 10.0.0.144, add the following to the `snmpd.conf` file.

   ```
   #           sec.name       source          community
   com2sec     snmpclient     10.0.0.144      public
   ```

4. To send SNMPv1 traps to a custom port, add the following configuration information to the `snmpd.conf` file:

   ```
   Trapsink HOST [community [port] ]
   ```

   Specify the custom port number; otherwise, the default SNMP trap port, 162, is used to send traps.

5. To run or stop the `snmpd` daemon, enter the following command:

   ```
   /etc/init.d/snmpd start
   /etc/init.d/snmpd stop
   ```

6. To start or stop the SAS SNMP Agent daemon before issuing a SNMP query, enter the following command:

```
/etc/init.d/lsi_mrdsnmpd start
/etc/init.d/lsi_mrdsnmpd stop
```

   You can check the status of the SAS SNMP Agent daemon by checked by entering the following command:

```
/etc/init.d/lsi_mrdsnmpd status
```

7. Issue an SNMP query in this format:

```
snmpwalk -v1 -c public localhost .1.3.6.1.4.1.3582
```

8. 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 the `disableAuthorization yes` line in the `snmptrapd.conf` file and then run the `sudo snmptrapd -P -F "%02.2h:%02.2j TRAP%w.%q from %A %v\n` command.

NOTE It is assumed that the `snmpd.conf` file is located in the `/etc/snmp` directory for the Red Hat operating system and the `/etc` directory for the SLES operating system. You can change the file location from the `/etc/init.d/lsi_mrdsnmpd` file.

You can install SNMP without the trap functionality. To do so, set the TRAPIND environment variable to `N` before running RPM.

Before you install a new version, you must uninstall all previous versions.

For the SLES 10 operating system, perform the following steps to run SNMP:

1. Copy the `/etc/snmp/snmpd.conf` file to the `/etc/snmpd.conf` file.
2. Modify the `/etc/init.d/snmpd` file, and change the `SNMPDCONF=/etc/snmp/snmpd.conf` entry to `SNMPDCONF=/etc/snmpd.conf`.
3. Run `LSI SNMP rpm`.

7.6.5 Installing and Configuring an SNMP Agent on Solaris

This section explains how to install and configure the SAS SNMP Agent for the Solaris operating system.

7.6.5.1 Prerequisites

This package requires that you have the Solaris System Management Agent installed on the Solaris machine.

NOTE While installing the SAS SNMP Agent on Solaris 11, the **net-snmp** package must be installed on the machine.

7.6.5.2 Installing the SNMP Agent on Solaris

To install SNMP for the Solaris operating system, perform the following steps:

1. Unzip the Avago SAS SNMP Agent package.

2. Run the install script by using the following command:

```
# ./install.sh
```

The installation exits if any existing versions of the `storelib` and `sassnmp` utilities are installed on the Solaris machine. Uninstall the existing version by using the following commands:

`# pkgrm sassnmp` (to uninstall the Avago SAS SNMP Agent)

`# pkgrm storelib` (to uninstall the `storelib` library)

7.6.5.3 Avago SAS SNMP MIB Location

After you install the Avago SAS SNMP Agent package, the `LSI-AdapterSAS.mib` MIB file is installed under the `/etc/lsi_mrdsnmp/sas` directory.

7.6.5.4 Starting, Stopping, and Checking the Status of the Avago SAS SNMP Agent

The following commands are used to start, stop, restart, and check the status of the Solaris System Management Agent (`net snmpd`) daemon on Solaris 10 x86 and Solaris 10 SPARC:

- 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`

The following commands are used to start, stop, restart, and check the status of the Solaris System Management Agent (`net snmpd`) daemon on Solaris 11 x86:

- Start: `# svcadm enable svc:/application/management/net-snmp`
- Stop: `# svcadm disable svc:/application/management/net-snmp`
- Restart: `# svcadm restart svc:/application/management/net-snmp`
- Status: `# svcs svc:/application/management/net-snmp`

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 on Solaris 10 x86, Solaris 10 SPARC, and Solaris 11 x86:

- 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`

7.6.5.5 Configuring the snmpd.conf File

By default, you can run the SNMP queries (`walk`, `get`) from any remote machine without any changes to the `snmpd.conf` file. To quickly add a new community and client access, perform the following steps:

1. Stop the SMA service by running the following command:

```
# svcadm disable svc:/application/management/sma:default
```

2. Add `read-only` and `read-write` community names.
   a. Add a `read-only` community name and client/hostname/ipaddress under `SECTION: Access Control Setup` in the `/etc/sma/snmp/snmpd.conf` file, as shown in the following excerpt.

```
###############################################
```

```
# SECTION: Access Control Setup
# This section defines who is allowed to talk to
# your running SNMP Agent.
# rocommunity: a SNMPv1/SNMPv2c read-only access
# community name
# arguments: community
# [default|hostname|network/bits] [oid]
# rocommunity snmpclient 172.28.157.149
################################################
```

NOTE In Solaris 11 x86, add a `read-only` community name and client/hostname/ipaddress under `SECTION: Access Control Setup` in the `/etc/net-snmp/snmp/snmpd.conf` file as shown in the preceding excerpt.

b. Add a `readwrite` 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.
# rwcommunity: a SNMPv1/SNMPv2c read-write access
# community name
# arguments: community
# [default|hostname|network/bits] [oid]
# rwcommunity snmpclient 172.28.157.149
################################################
```

NOTE In Solaris 11 x86, add a `read-only` community name and client/hostname/ipaddress under `SECTION: Access Control Setup` in the `/etc/net-snmp/snmp/snmpd.conf` file as shown in the preceding excerpt.

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.

NOTE In Solaris 11 x86, you must start the `net-snmpd` daemon service, by executing the following command: `# svcadm enable svc:/application/management/net-snmp`

7.6.5.6 Configuring SNMP Traps

To receive SNMP traps, perform the following steps:

1. Stop the Avago SAS SNMP Agent by using the following command:

```
#/etc/init.d/lsi_mrdsnmpd stop
```

2. Edit the `/etc/lsi_mrdsnmp/sas/sas_TrapDestination.conf` file, and add the IP address as shown in the following excerpt.

```
################################################
# Agent Service needs the IP addresses to sent trap
# The trap destination may be specified in this file
# or using snmpd.conf file. Following indicators can
# be set on "TrapDestInd" to instruct the agent to
# pick the IPs as the destination.
# 1 - IPs only from snmpd.conf
# 2 - IPs from this file only
# 3 - IPs from both the files
################################################
TrapDestInd 2
############# Trap Destination IP #################
# Add port no after IP address with no
# space after
# colon to send the SNMP trap
# message to custom port.
# Alternatively, you can also use
# trapsink command
# in snmpd.conf to send the SNMP trap
# message to
# custom port, else default SNMP trap
# port 162 shall be used.
127.0.0.1 public
145.147.201.88:1234 testComm
################################################
```

NOTE Solaris also supports Custom community support.

3. If `'TrapDestInd'` above is set to 1, the IP addresses shall be taken from the `/etc/sma/snmp/snmpd.conf` file in the following format: `'com2sec snmpclient 172.28.157.149 public'` the `'Trapsink'` and `'TrapCommunity'` tokens are supported for sending customized SNMP traps.

NOTE In Solaris 11 x86, the file will be taken from the `/etc/net-snmp/snmp/snmpd.conf` file.

4. Start the Avago SAS SNMP Agent by entering the following command:

```
#/etc/init.d/lsi_mrdsnmpd start
```

7.6.5.7 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
```

7.7 MegaRAID Storage Manager Remotely Connecting to VMware ESX

When the MegaRAID Storage Manager software is used to connect to a VMware ESX machine from a remote machine (Windows or Linux), for long running operations (such as volume creation, deletion) to complete in a shorter time, perform the following steps:

1. Login to the VMware ESX machine.
2. Open `/etc/sfcb/sfcb.cfg`.
3. Increase the `keepaliveTimeout` value from 1 to 100 or to a higher value.
4. Restart `sfcbd` (`/etc/init.d/sfcbd-watchdog restart`).
5. Restart the MegaRAID Storage Manager Framework on the MegaRAID Storage Manager client machine.
   - For Windows – Restart the framework service.
   - For Linux – Restart the vivaldi framework service.
6. Relaunch the **MegaRAID Storage Manager** window.

7.8 Prerequisites to Running MegaRAID Storage Manager Remote Administration

The MegaRAID Storage Manager software requires ports 3071 and 5571 to be open to function. Follow these steps to prepare to run the MegaRAID Storage Manager Remote Administration.

1. Configure the system with a valid IP address.

   Make sure the IP address does not conflict with another in the sub network.

   Ports, such as 3071 and 5571, are open and available for the MegaRAID Storage Manager framework communication.
2. Disable all security manager and firewall.
3. Configure the multicasting.

   Make sure Class D multicast IP addresses are registered (at least 229.111.112.12 should be registered for the MegaRAID Storage Manager software to work); if not, create a static route using the following command:

```
Route add 229.111.112.12 dev eth1
```

4. Install the MegaRAID Storage Manager software.

   If the MegaRAID Storage Manager software is already installed, restart the MegaRAID Storage Manager Framework.

7.9 CLI Packaging Details

The following table describes the strategies followed while packaging the CLI binaries with the MegaRAID Storage Manager software.

Table 59 CLI Packaging Strategies

Release Details	Packaging Details
MegaRAID Storage Manager Software Major Release (*N*[a]).	*N-1* GA or the latest point release of the CLI is bundled.
MegaRAID Storage Manager Software Point Release.	*N* GA or the latest point release of the CLI is bundled.

a. Where *N* is the release stream.

Chapter 8: MegaRAID Storage Manager Window and Menus

This chapter explains how to start the MegaRAID Storage Manager software and describes the MegaRAID Storage Manager window and menus.

8.1 Starting the MegaRAID Storage Manager Software

You must have administrative privileges to use the MegaRAID Storage Manager software in either full-access or in view-only mode. Follow these steps to start the MegaRAID Storage Manager software on various platforms.

- To start the MegaRAID Storage Manager software on a Microsoft Windows operating system, select **Start > Programs > MegaRAID Storage Manager > StartupUI**, or double-click the MegaRAID Storage Manager shortcut on the desktop.

  NOTE If a warning appears stating that Windows firewall has blocked some features of the program, click Unblock to allow the MegaRAID Storage Manager software to start. (The Windows firewall sometimes blocks the operation of programs that use Java® Technology.)

- To start the MegaRAID Storage Manager software on a Red Hat Linux operating system, select **Applications > System Tools > MegaRAID Storage Manager StartupUI**.
- To start the MegaRAID Storage Manager software on an SUSE Linux or SLES operating system, select **Start > System > More Programs > MegaRAID Storage Manager**.
- To start the MegaRAID Storage Manager software on a Solaris x86 and Solaris SPARC operating system, select **Launch > Applications > Utilities > MegaRAID Storage Manager StartupUI**.

8.2 Discovery and Login

You can start the MegaRAID Storage Manager software from a remote Windows/ Linux machine that has the MegaRAID Storage Manager software installed in complete mode.When the program starts, the **Host View** dialog appears, as shown in the following figure. The remote servers are displayed, along with their IP addresses, operating system, and health status.

NOTE If you do a local mode installation, as shown in Section Installing MegaRAID Storage Manager Software on Microsoft WIndows, the following figure will not be displayed. It will directly prompt you to the login dialog as shown in the Server Login.

Figure 153 Host View

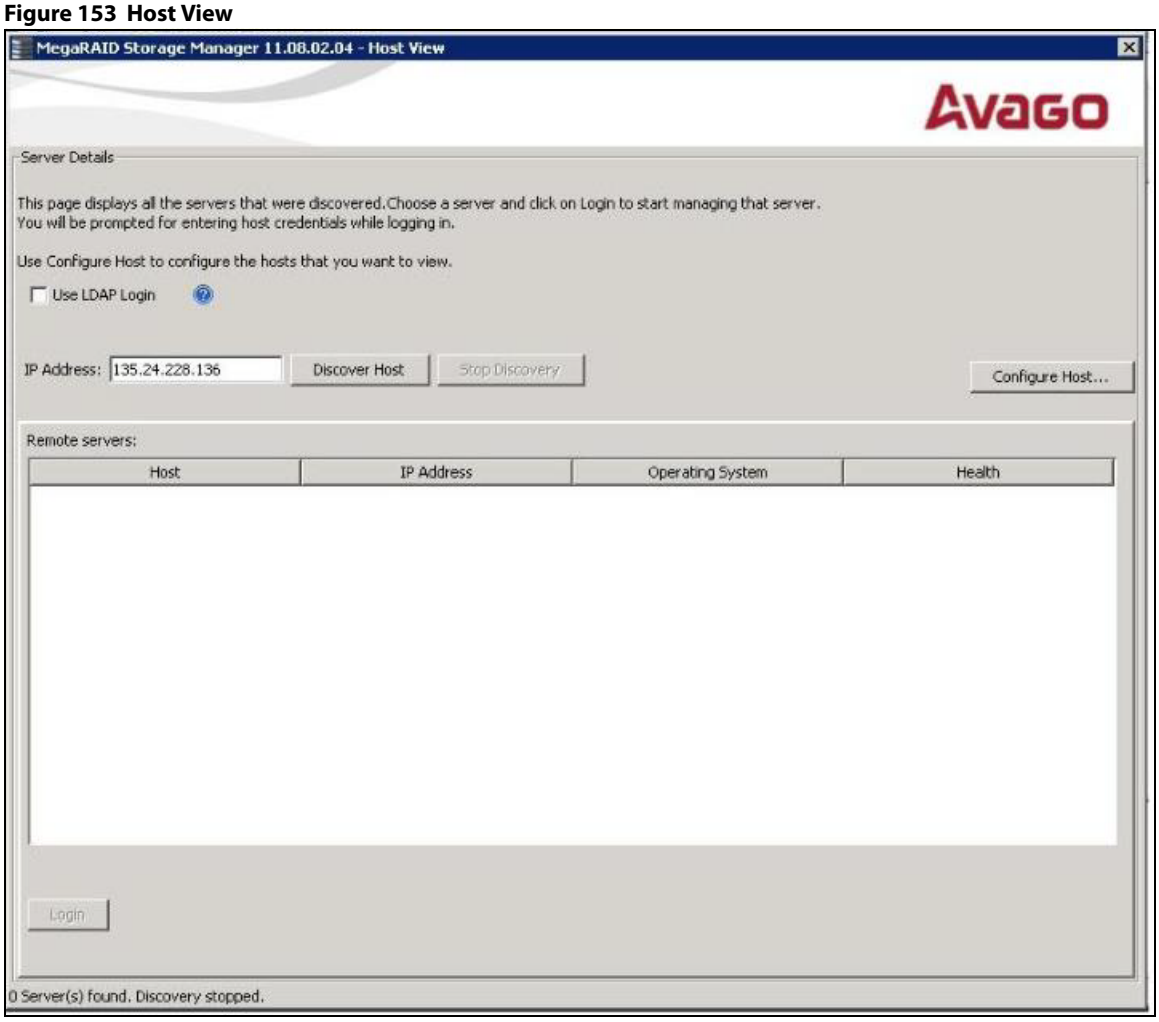

If High Availability DAS is supported on the controller, instead of the above **Host View** dialog, the **Host View – High Availability DAS** dialog appears, as shown in Host View - High Availability DAS.

The **Host View** dialog shows an icon for each server on which the MegaRAID Storage Manager software is installed. The servers are color-coded with the following definitions:

- Green: The server is operating properly.
- Yellow: The server is running in a partially degraded state (possibly because a drive in a virtual drive has failed).
- Orange: The server is running in a degraded state.
- Red: The server storage configuration has failed.

NOTE Do not enter the VMware ESXi server's IP address in the **IP Address** field in the previous figure. Instead enter a valid MegaRAID Storage Manager server's IP address and select the **Display all the systems in the Network of the local server** option in the following figure.

1. Click **Configure Host** to configure the hosts.

   The **Configure Host** dialog appears, as shown in the following figure.

Figure 154 Configure Host

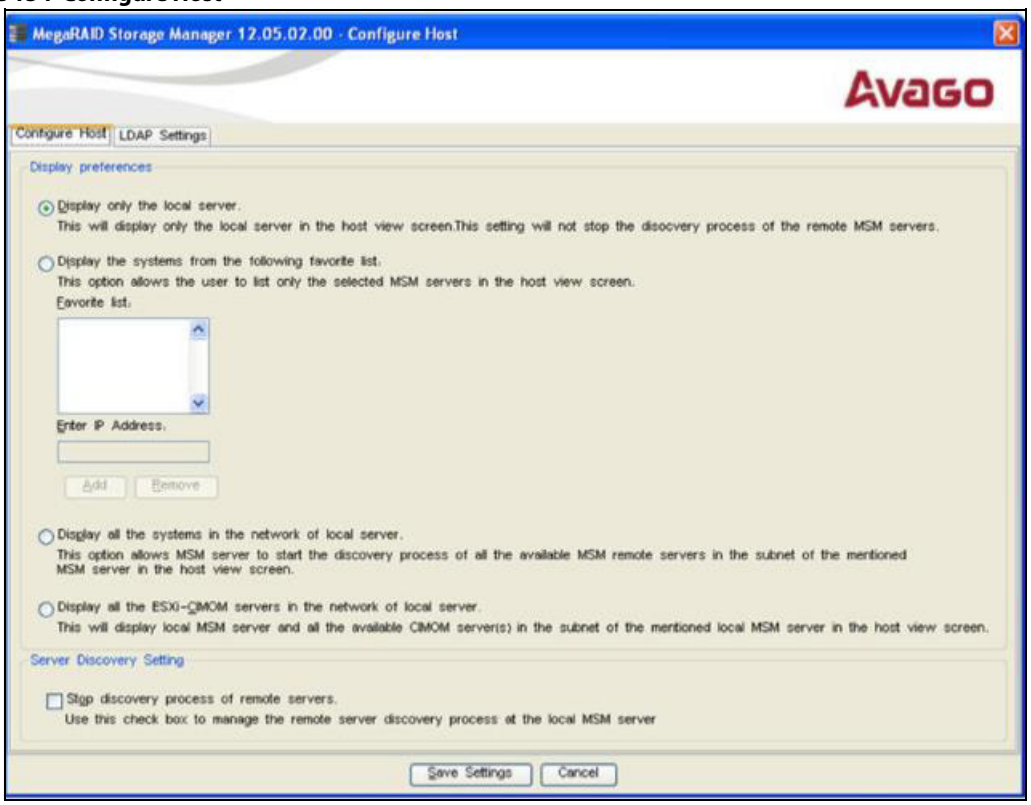

The following options are available to configure the host.

- **Display only the local server** – Select this option to display only the Local server or the Server of the IP address entered in the Host View screen.
- **Display the systems from the following favorite list** – Allows you to enter IP addresses of the MegaRAID Storage Manager servers and discovers only those servers. You can enter an IP address in the **Enter IP Address** field and click **Add**. The server corresponding to the IP address appears in the **Favorite list**.
- **Display all the systems in the Network of the local server** – Discovers all the MegaRAID Storage Manager servers available in the network.
- **Display all the ESXi-CIMOM servers in the network of local server** - Discovers the local MegaRAID Storage Manager server and all the available ESXi servers in the network.

NOTE If the controller supports High Availability DAS, and you want to view the cluster information in a single pane, select either of the options: **Display the systems from the following favorite list** or **Display all the systems in the Network of the local server**.

2. Click **Save Settings** to save your setting, or on **Cancel** to quit without saving.

   If you click **Save Settings**, a confirmation dialog appears asking you to confirm your settings. Click **OK** in the confirmation dialog to start the discovery process.

3. Select the **Stop discovery process of remote servers** check box and click on **Save Settings**, to abort the discovery process which has already begun. This option is enabled only when there is an active discovery process.

NOTE The VMware server does not show the system health and the operating system details. It shows only the host name and the IP address of the server. The Operating system and Controller Health are displayed as `Unknown` in **Host View** screen. When connecting to a

VMware server on a different subnet, one or more frameworks have to be running in the subnet to connect to the CIMOM.

Figure 155 Host View Window

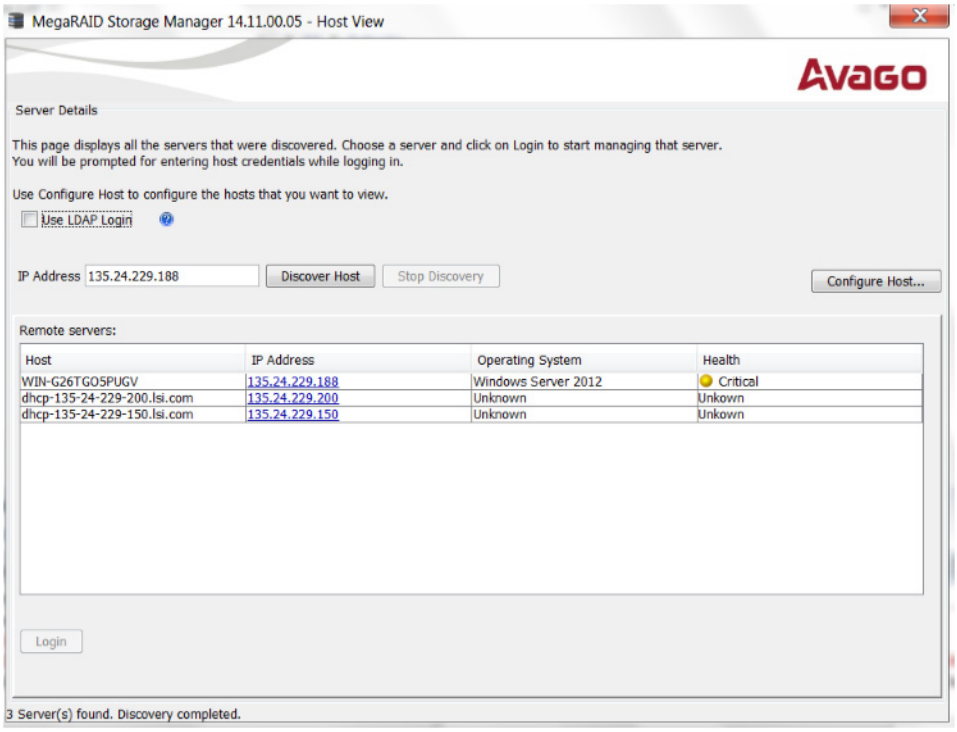

The servers appear in the list of found hosts in the **Host View** dialog.

4. Double-click the icon of the server that you want to access.

   The **Server Login** window appears, as shown in the following figure.

Figure 156 Server Login

5. Enter the root account name and password of the host in the **User Name** and **Password** fields respectively.

NOTE In the **User Name** field, you can also enter the domain name along with the user name; for example, `Avago\abc`, where `Avago` is the domain name and `abc` is your user name.

The question mark icon opens a dialog box that explains what you need for full access to the server and for view-only access to the server. You are allowed three attempts to Log in.

NOTE When connected to VMware operating system, the **Server Login** window shows only one label for access, Full Access. Multiple users can have full access to the VMware server.

6. Select an access mode from the drop-down menu for **Login Mode**, and click **Login**.
   - Select **Full Access** if you need to both view and change the current configuration.
   - Select **View Only** if you need to only view and monitor the current configuration.

NOTE If the computer is networked, this login is for the computer itself, not the network login.

Enter the root or administrator user name and password to use Full Access mode.

NOTE In Linux, users belonging to the root group can log in. You do not have to be the user root.

If your user name and password are correct for the Login mode you have chosen, the MegaRAID Storage Manager main menu appears.

8.3 High Availability DAS Support

If High Availability DAS is supported on the controller when you launch the MegaRAID Storage Manager application, the following dialog appears.

Figure 157 Host View - High Availability DAS

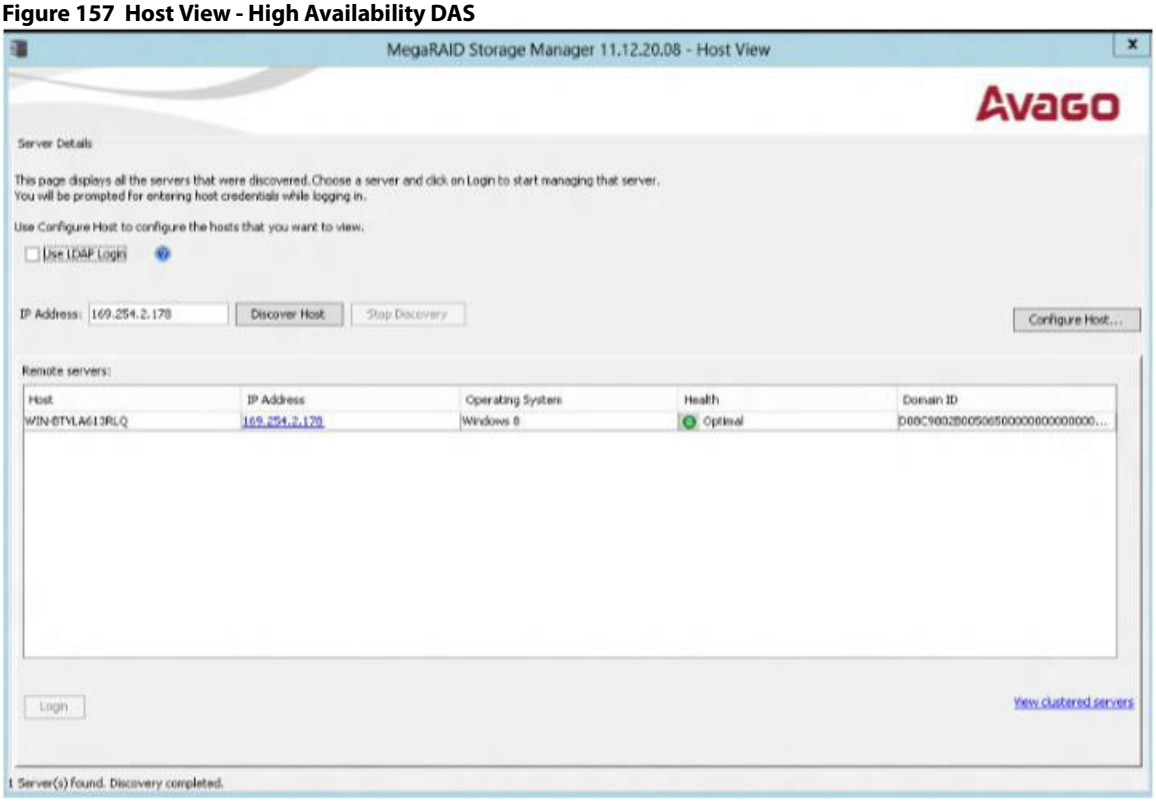

1. Click **View clustered servers** to view all the High Availability cluster servers available.
   The **View Clustered Servers** dialog appears, as shown in the following figure.

Figure 158 View Clustered Servers

2. Click on a server link to log into that server.
   The **Server Login** window appears.
3. Enter the login details in the **Server Login** window.

8.4 LDAP Support

The MegaRAID Storage Manager application supports the discovery of remote MegaRAID Storage Managers servers using LDAP. To enable LDAP support, the MegaRAID Storage Manager servers must be registered with the LDAP server.

NOTE LDAP supports only Windows Active Directory LDAP Server Implementation.

NOTE ESXi servers are not discovered during LDAP discovery.

To register the MegaRAID Storage Manager servers with the LDAP server, define a new attribute, `ou`, on the machine on which the LDAP server is configured, and give this attribute the value MSM. This registration enables the discovery of only the MegaRAID Storage Manager servers that have been registered with the LDAP server.

To use LDAP support, follow these steps:

1. Double-click the MegaRAID Storage Manager software shortcut icon on your desktop.

   The **Select Server** dialog appears.

2. Select the **Use LDAP Login** check box, and click **Discover Host**.

   All the MegaRAID Storage Manager servers registered with the LDAP server are displayed in the **Remote servers** box.

   NOTE If the **Use LDAP Login** check box is selected, the **IP Address** field is disabled.

3. Click on a server link to connect to the LDAP server.

   NOTE Based on the privileges allotted to you, the MegaRAID Storage Manager servers are launched with full access rights or read-only rights.

If you have selected the **Do not prompt for credentials when connecting to LDAP** check box (in the LDAP Settings tab in the **Configure Host** dialog), you are directly connected to the LDAP server; otherwise, the **LDAP Login** dialog appears.

Figure 159 LDAP Login

Follow these steps to enter the LDAP login details:

1. Enter the IP address of the LDAP server in the **LDAP Server IP Address** field.
2. Enter the LDAP server's user name and password in the **User Name** and **Password** fields, respectively.
   An example of a user name can be `username@testldap.com`.
3. Enter the name of the Domain Controller in the **Distinguished Name** field.
   As an example, the Domain Controller name can be `dc= TESTLDAP, dc=com`.

   NOTE The **LDAP Server IP Address**, **User Name**, **Password**, and **Distinguished Name** fields are already populated if their corresponding values have been stored in the LDAP Settings tab in the **Configure Host** dialog.

4. Perform one of these actions:
   — If you want to use the default port number, select the **Use Default Port** check box.
   The default port number, 389, appears in the **Port** field.
   — If you do not want to use the default port number, uncheck the **Use Default Port** check box, and enter a port number in the **Port** field.
5. Select the **Remember my Login Details** check box if you want to save all the values entered in this dialog in the LDAP Settings tab in the **Configure Host** dialog.
6. Click **Login** to log in to the LDAP server.

8.5 Configuring LDAP Support Settings

To configure settings for LDAP support, follow these steps:

1. Navigate to the **Configure Host** dialog, and click the **LDAP Settings** tab.
   The following fields appear.

Figure 160 Configure Host LDAP

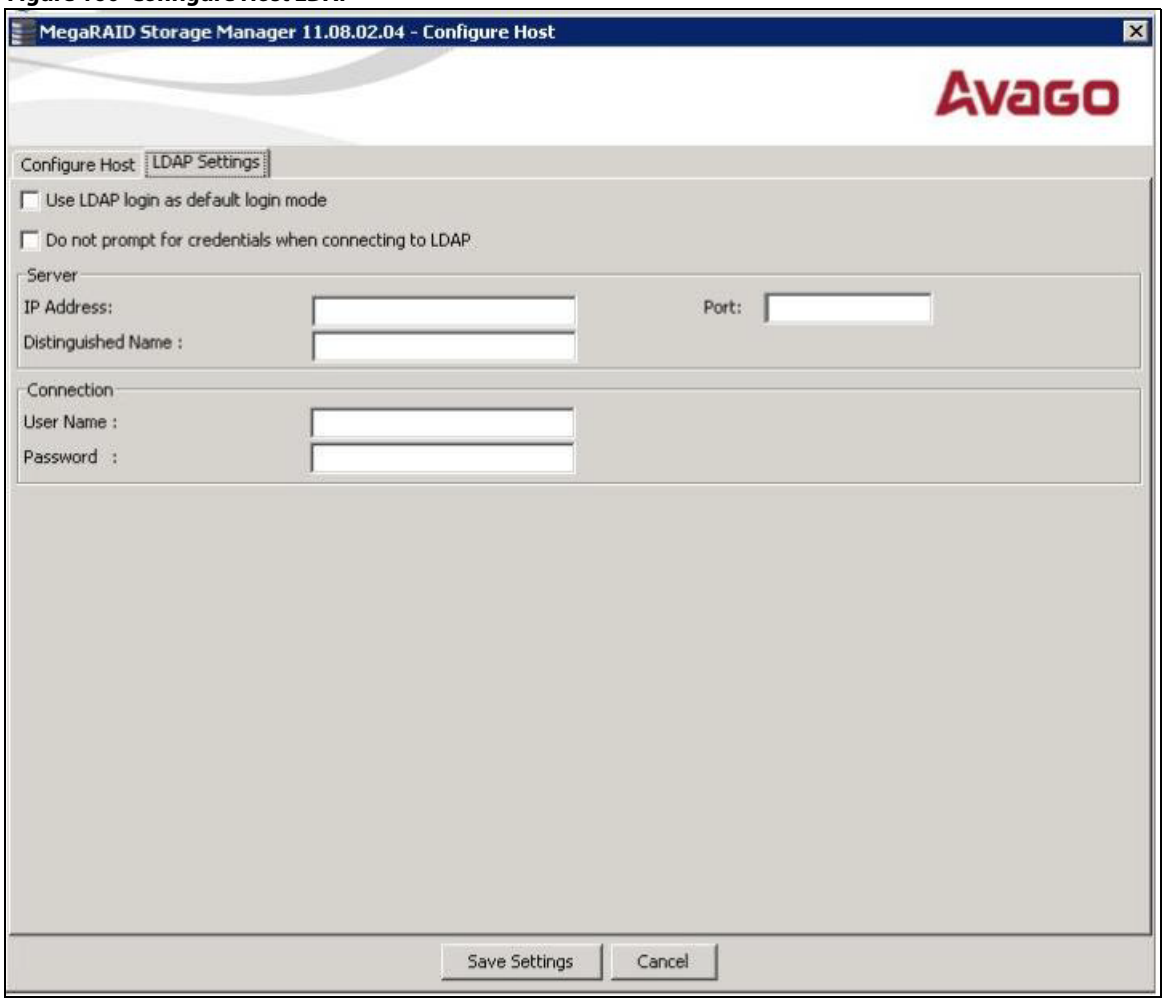

2. Select the **Use LDAP login as default login mode** check box to always connect to the LDAP server.
3. Select the **Do not prompt for credentials when connecting to LDAP** check box if you do not want the **LDAP Login** dialog to appear when connecting to the LDAP server.
4. Enter the IP address of the LDAP server in the **IP Address** field.
5. Enter the port number in the **Port** field.
6. Enter the name of the Domain Controller in the **Distinguished Name** field.
7. Enter the user name and password for logging into the LDAP server in the **User Name** and **Password** fields, respectively.
8. Click **Save Settings** to save all the values entered in the fields in the `msm.properties` file.

8.6 MegaRAID Storage Manager Main Menu

This section describes the MegaRAID Storage Manager main menu window.

- Dashboard View, Physical View, and Logical View
- Properties and Graphical View Tabs
- Event Log Panel

8.6.1 Dashboard View, Physical View, and Logical View

The MegaRAID Storage Manager software displays the *Dashboard* view, the *Physical* view, or the *Logical* view. Depending on which tab is selected, you can view information about the system and the attached devices.

Dashboard View

The *Dashboard* view shows an overview of the system and covers the following features:

- Status of the controller cards that are connected to the system. When multiple controllers are connected, they are sorted based on the bus device function. The controllers are indexed with numbers 0, 1, 2, and so on.
- Properties of the virtual drives, physical drives, enclosures, and expanders.
- Total capacity, configured capacity, and unconfigured capacity.
- Background operations in progress.
- The MegaRAID Storage Manager software features and their status (enabled or disabled).
- Actions you can perform, such as creating a virtual drive and updating the firmware.
- Links to online help.

Figure 161 MegaRAID Storage Manager Dashboard View

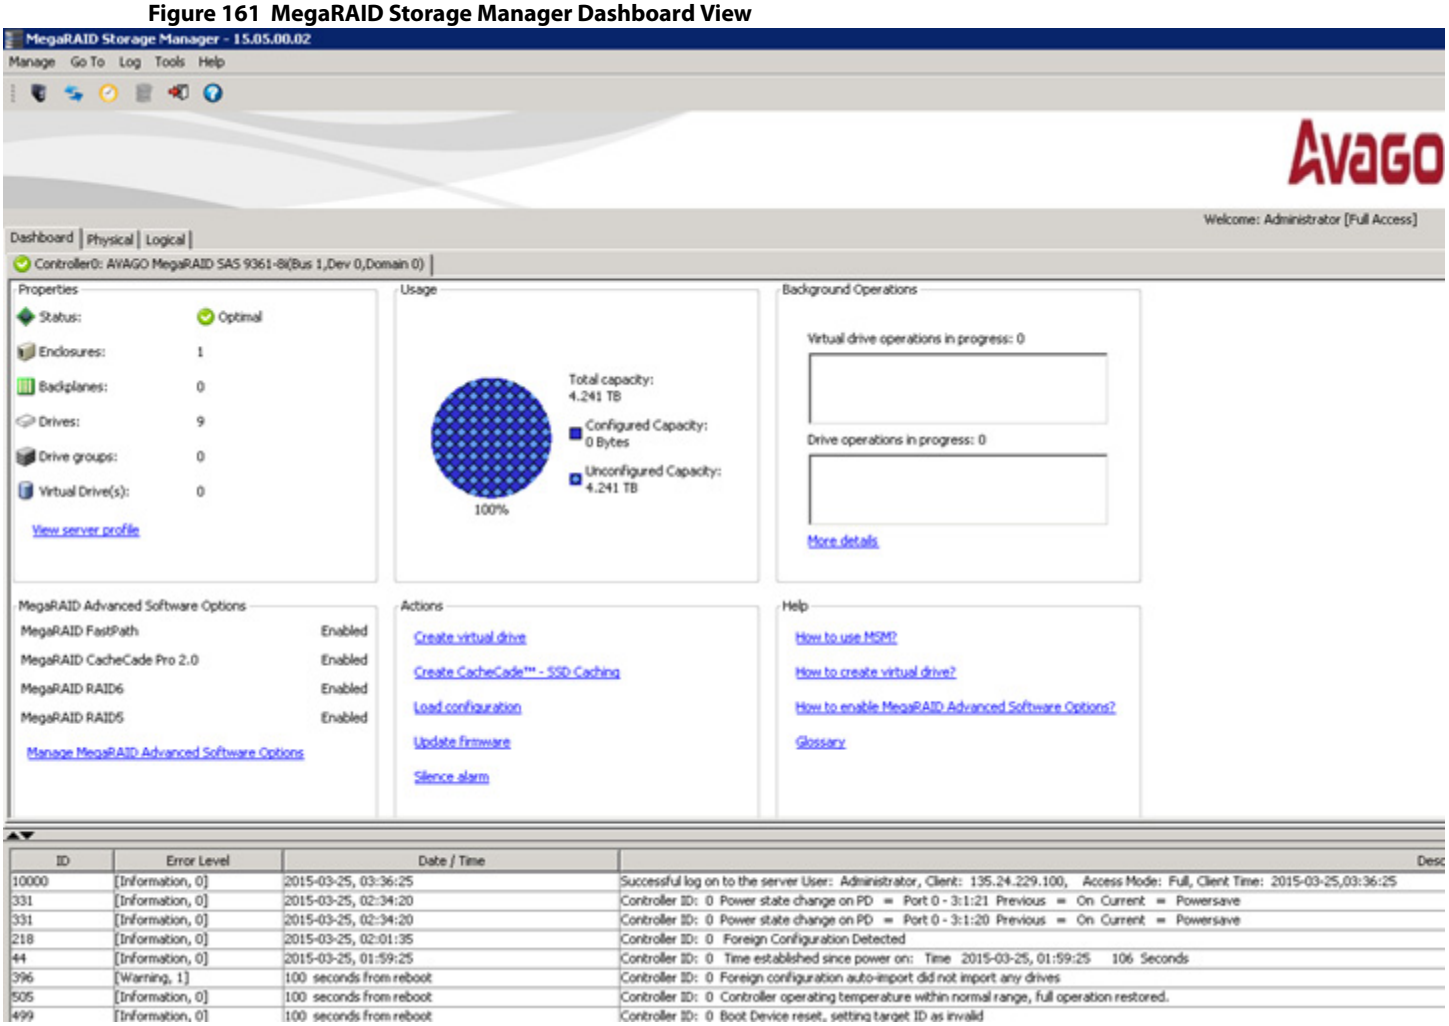

NOTE Some *Dashboard* view screens in this document do not show the controller indexing feature.

NOTE If the controller supports High Availability DAS, the **HA Peer Controller Status** field appears in the above dialog and displays one of the following values: **Active** (both the servers in the cluster are running), **Inactive** (only one server in the cluster is running), or **Incompatible** (there is incompatibility between the servers).

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. Also, when multiple controllers are connected, they are sorted based on the bus device function. The controllers are indexed with numbers 0, 1, 2, and so on. Drives and other devices are attached to the ports. The properties for each item appear in the right panel of the screen.

Figure 162 MegaRAID Storage Manager Physical View

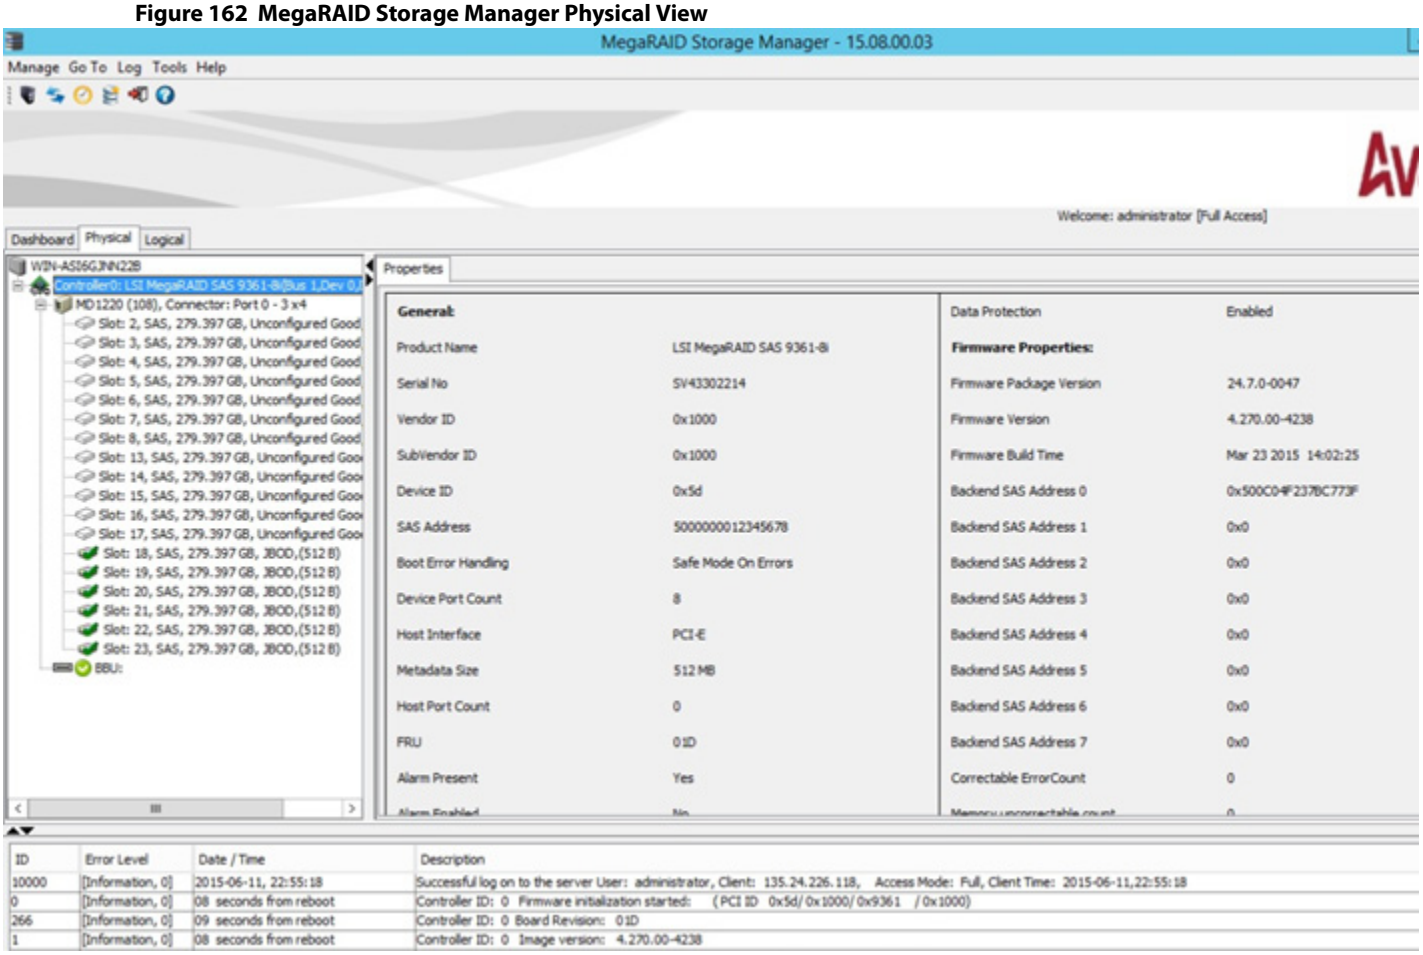

NOTE Some *physical* view screens in this document do not show the controller indexing and port enumeration features.

If the controller supports High Availability DAS, an additional parent mode, **Server Domain**, appears on the device tree in the **Physical** tab, as shown in the following figure.

Figure 163 Physical View for High Availability DAS

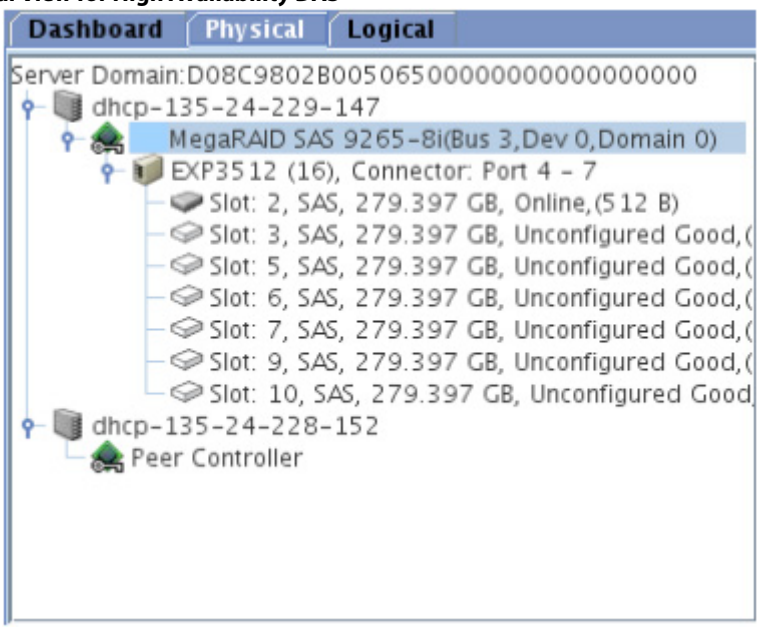

The **Server Domain** is the domain ID of the cluster and shows the two servers that belong to it as child nodes. Information that pertains to the logged-in server in the cluster (such as controller name, enclosures, physical drives) is shown in the **Physical** tab. For the peer server, no details are shown; the Physical tab just detects that a peer server exists and a controller is attached to it. Right-click **Server Domain** to view the properties of the cluster. A view-only properties dialog appears with two fields; **Domain ID** and **No. of Servers Tagged**.

Logical View

The *Logical* view shows the hierarchy of controllers, virtual drives, and the drives and drive groups that make up the virtual drives. When multiple controllers are connected, they are sorted based on the bus device function. The controllers are indexed with 0, 1, 2, and so on. The properties for these components appear in the right panel.

The following figure shows the *Logical* view.

Figure 164 MegaRAID Storage Manager Logical View

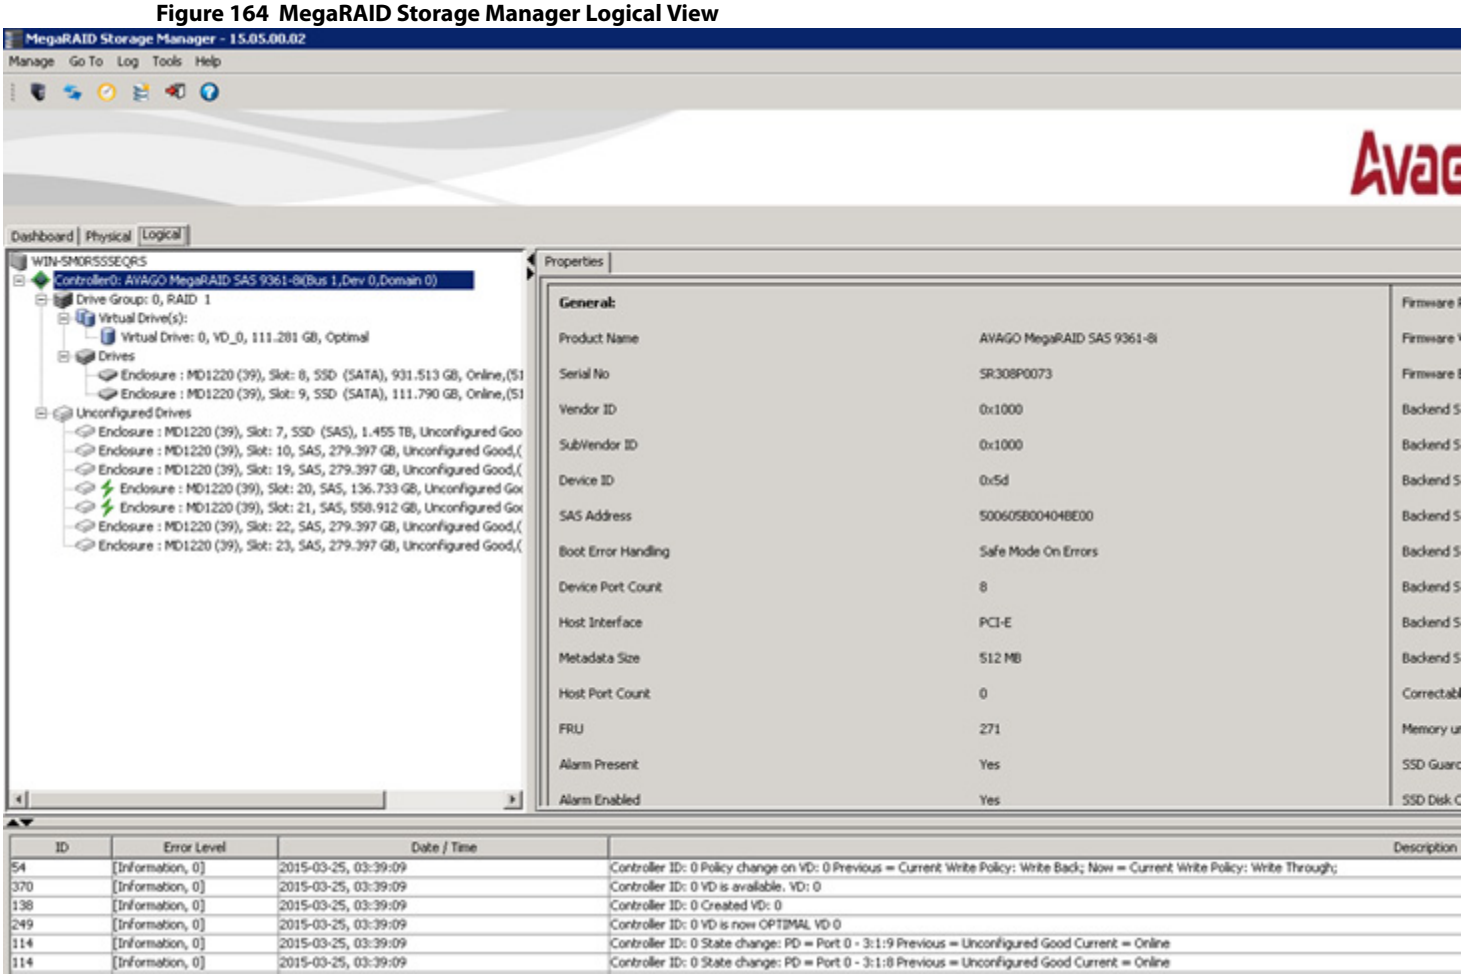

NOTE Some *Logical* view screens in this document do not show the controller indexing feature.

If the controller supports High Availability DAS, an additional parent mode, **Server Domain**, appears on the device tree in the *Logical* tab, as shown in the following figure.

Figure 165 Logical View for High Availability DAS

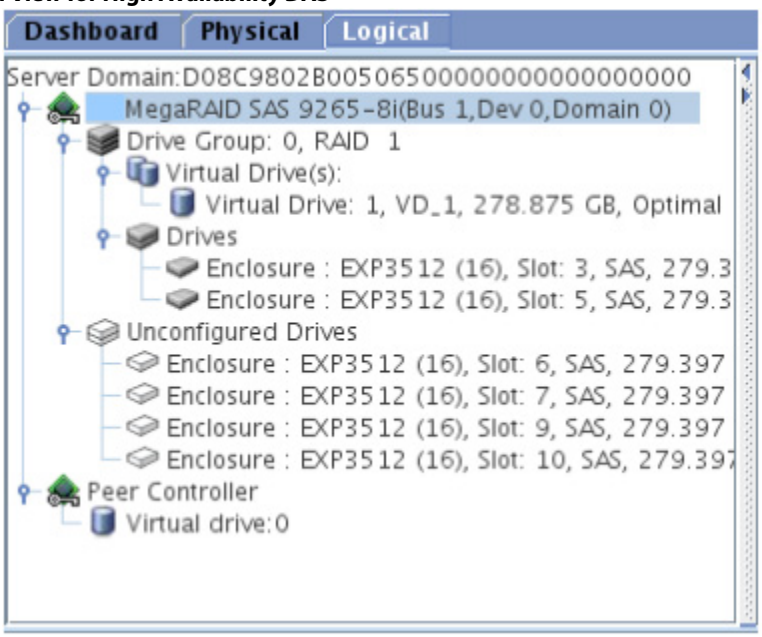

The **Server Domain** is the domain ID of the cluster and shows the two servers that belong to it as child nodes. Information that pertains to the logged-in server in the cluster (such as controller name, drive groups, virtual drives) is shown. For the peer server, the **Logical** tab detects that a peer server exists and a controller is attached to it and shows only the virtual drives created by the peer server. Right-click **Server Domain** to view the properties of the cluster. A view-only properties dialog shows with two fields; **Domain ID** and **No. of Servers Tagged**.

8.6.2 Physical Drive Temperature

The temperature for the physical drive appears in the following figure. You can scroll down to view the **Temperature** property.

Figure 166 Physical Drive Temperature

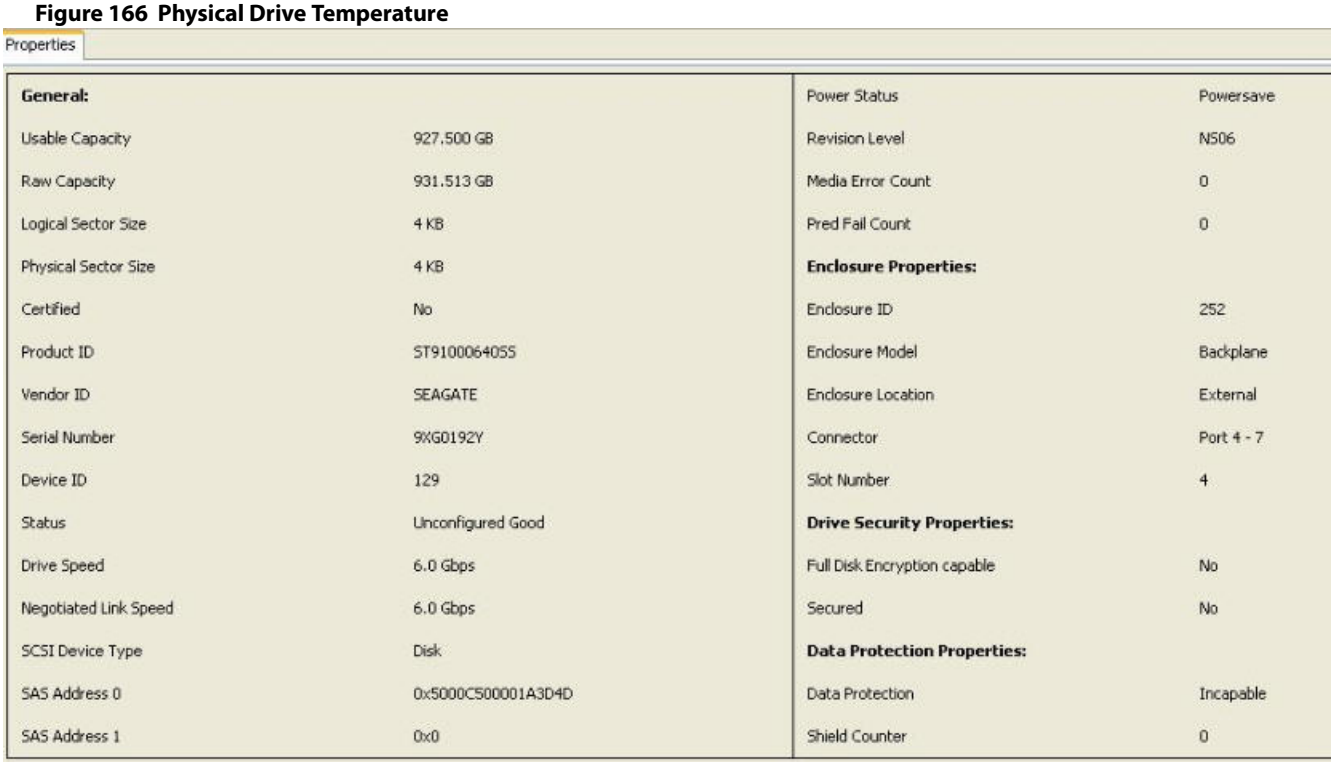

8.6.3 Shield State

This section describes the Shield state in the MegaRAID Storage Manager software.

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

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

The three possible Shield states are **Unconfigured – Shielded**, **Configured – Shielded**, and **Hotspare – Shielded**.

8.6.4 Shield State Physical View

Follow these steps to view the Shield state under the **Physical** view tab.

1. Click the **Physical** tab in the device tree.
   The red dot icon ( ) indicates a Shield state.
   The *Physical* view shield state is shown in the following figure.

Figure 167 Physical View Shield State

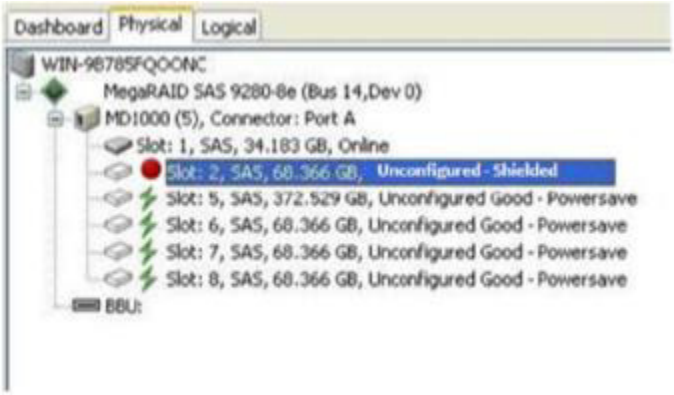

8.6.5 Logical View Shield State

Follow these steps to view the Shield state under the **Logical** tab.

1. Click the **Logical** tab in the device tree.
   The red dot icon ( ) indicates a Shield state.
   The *Logical* view Shield state is shown in the following figure.

Figure 168 Logical View Shield State

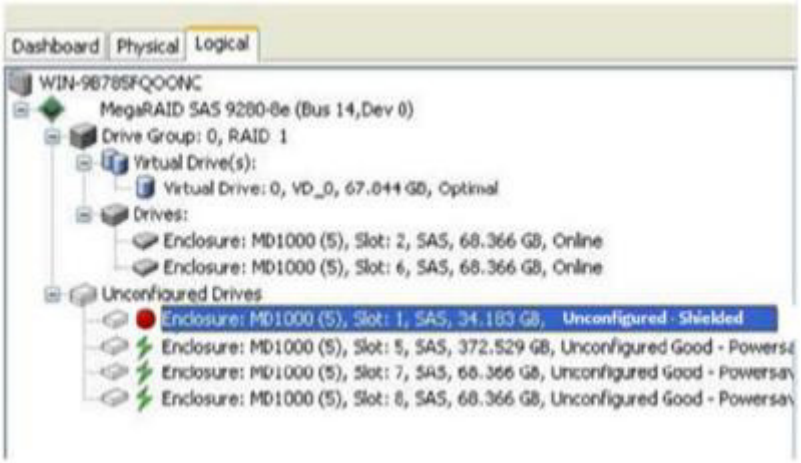

8.6.6 Viewing the Physical Drive Properties

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

1. Click the **Physical** tab or **Logical** tab in the device tree.
   The red dot icon ( ) indicates a Shield state.
2. Click the physical drive in Shield state on *Physical* view or *Logical* view of the device tree to view the properties.
   The device properties are displayed as shown in the following figure.

Figure 169 Physical Drive Properties of a Drive in Shield State

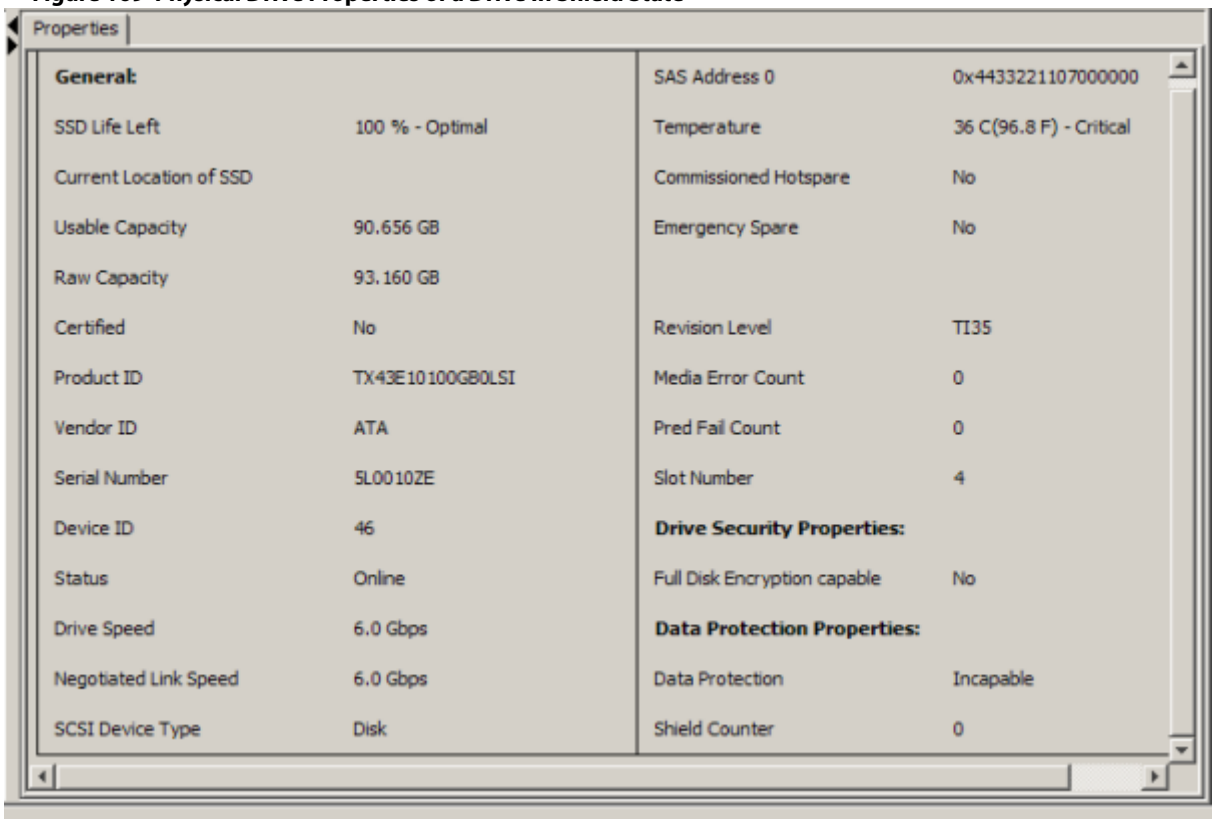

NOTE The Status of the drive must be of the Shield type.

8.6.7 Viewing the Server Profile of a Drive in Shield State

Perform these steps to view the server properties of the drive in Shield state.

1. Click the **Dashboard** tab in the device tree.
2. Click the **View Server Profile** link in the dashboard view.
   The server profile information is displayed, as shown in the following figure.

Figure 170 View of a Drive in Shield State

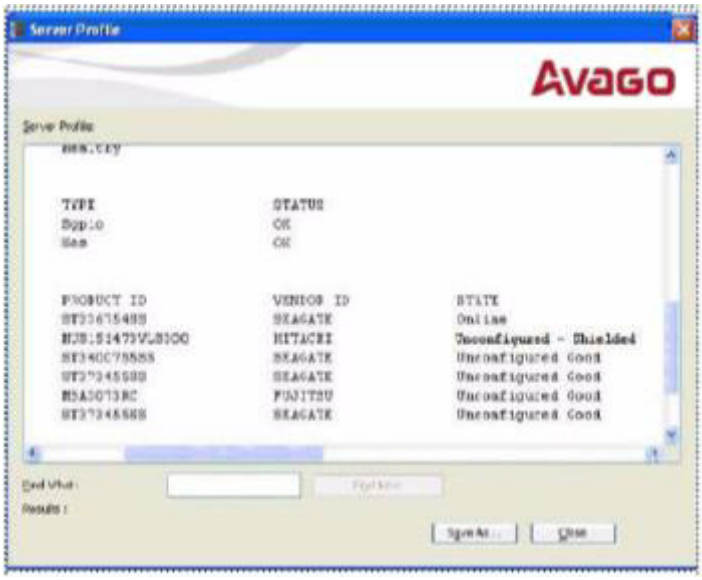

8.6.8 Displaying the Virtual Drive Properties

The MegaRAID Storage Manager application displays the following additional virtual drive statistics under controller properties.

- Parity size
- Mirror date size
- Metadata size

8.6.8.1 Parity Size

Parity size is used to store parity information on RAID 5, RAID 6, RAID 50, and RAID 60 virtual drives.

Follow these steps to view the Parity Size.

1. In the *Logical* view, click the **Virtual Drive** node.
2. For RAID 5, RAID 6, RAID 50, and RAID 60, the **Parity Size** is displayed, as shown in the following figure.

Figure 171 Parity Size

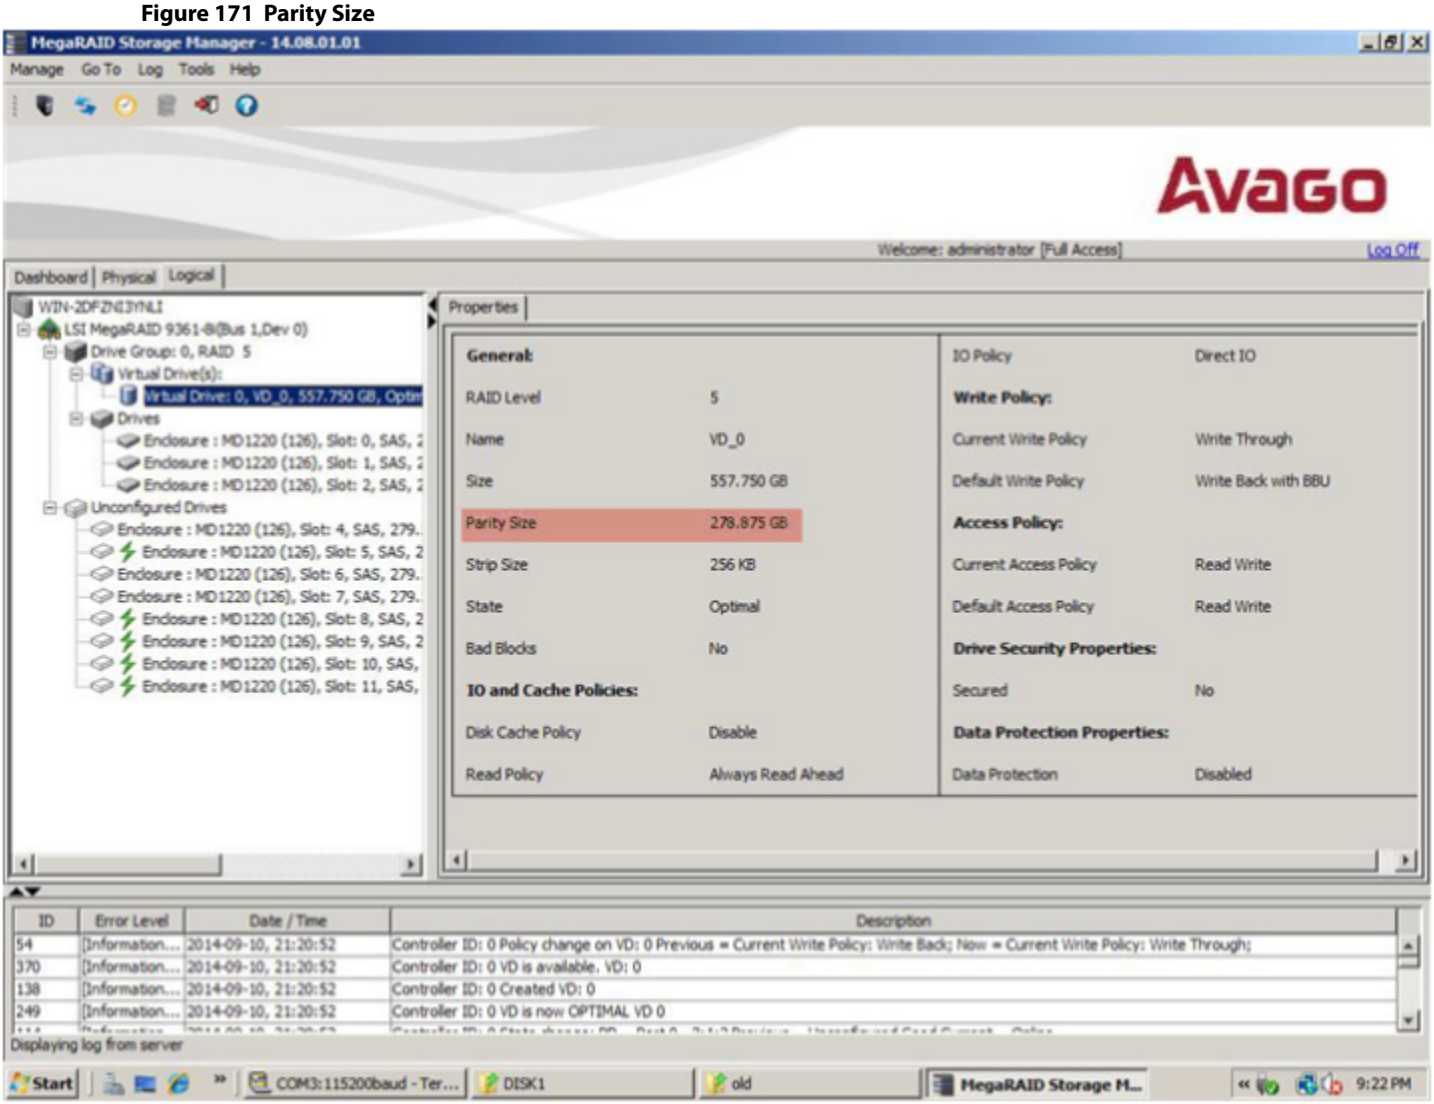

8.6.8.2 Mirror Data Size

Mirror Data Size is used to determine the size used for storing redundant information on RAID 1 and RAID 10 virtual drives.

Follow these steps to view the Mirror Data Size.

1. In the *Logical* view, click on the Virtual Drive node.
   The Mirror data size is displayed for RAID 1 and RAID 10 volumes, as shown in the following figure.

Figure 172 Mirror Data Size

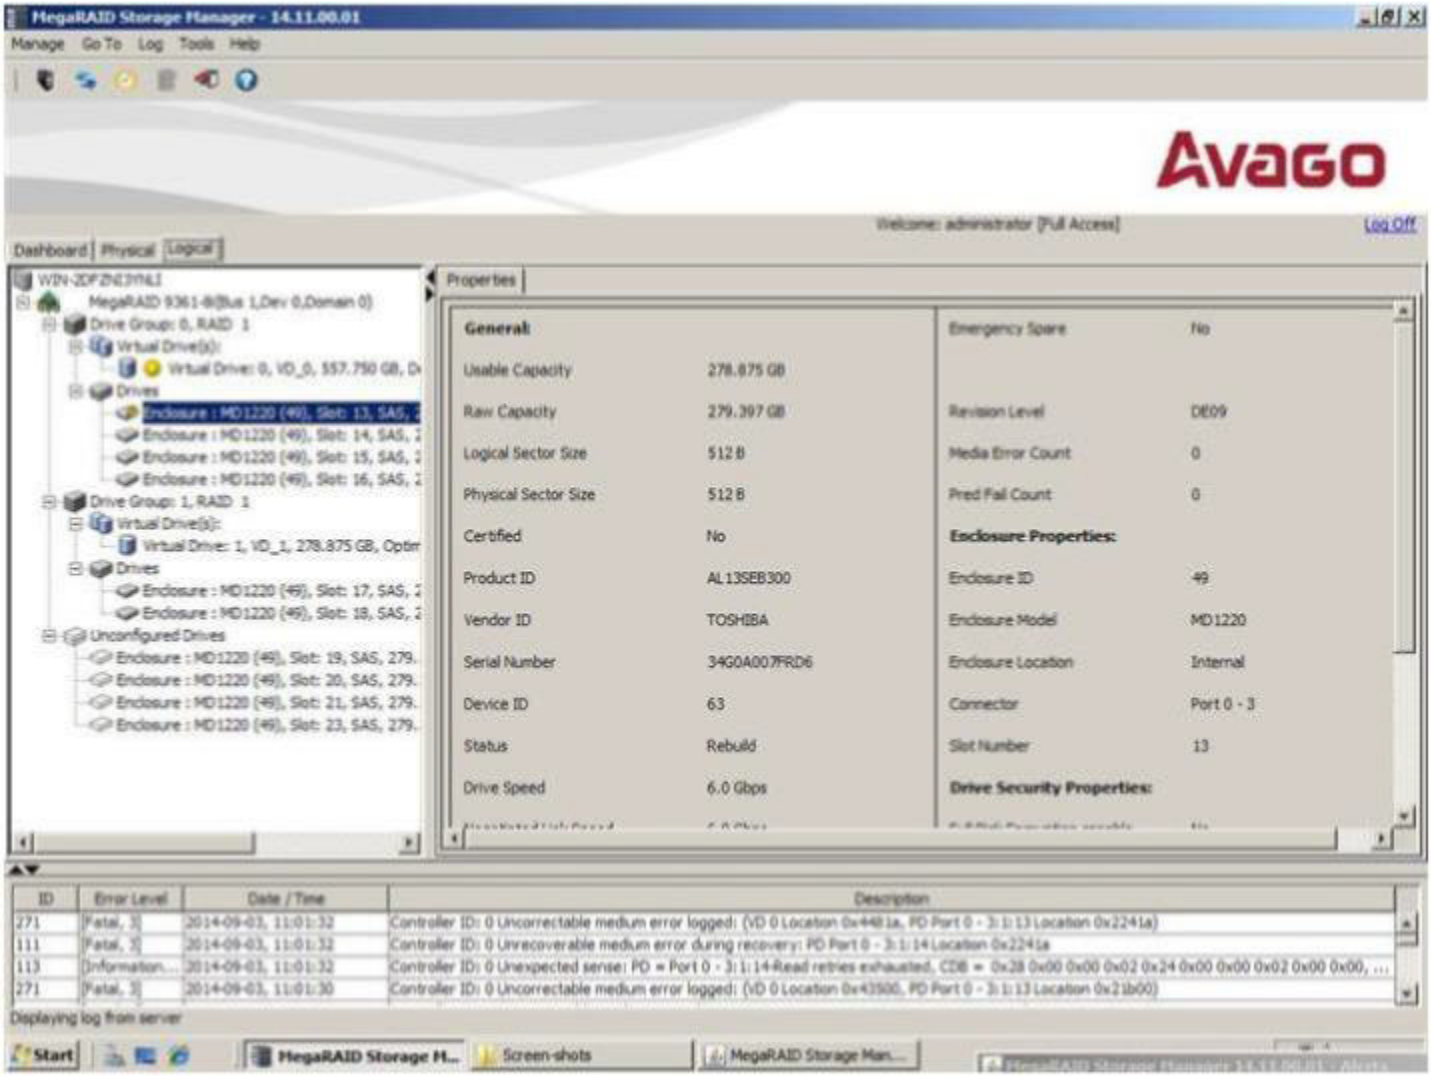

NOTE The parity size and mirror data size are not displayed for RAID 0 and RAID 00 volumes.

8.6.8.3 Metadata Size

The metadata size field displays the total space used for metadata.

Follow these steps to view the Metadata Size.

1. In the *Logical* view or the *Physical* view, click the controller node.
   The total space used for metadata is displayed in this field, as shown in the following figure.

Figure 173 Metadata Size

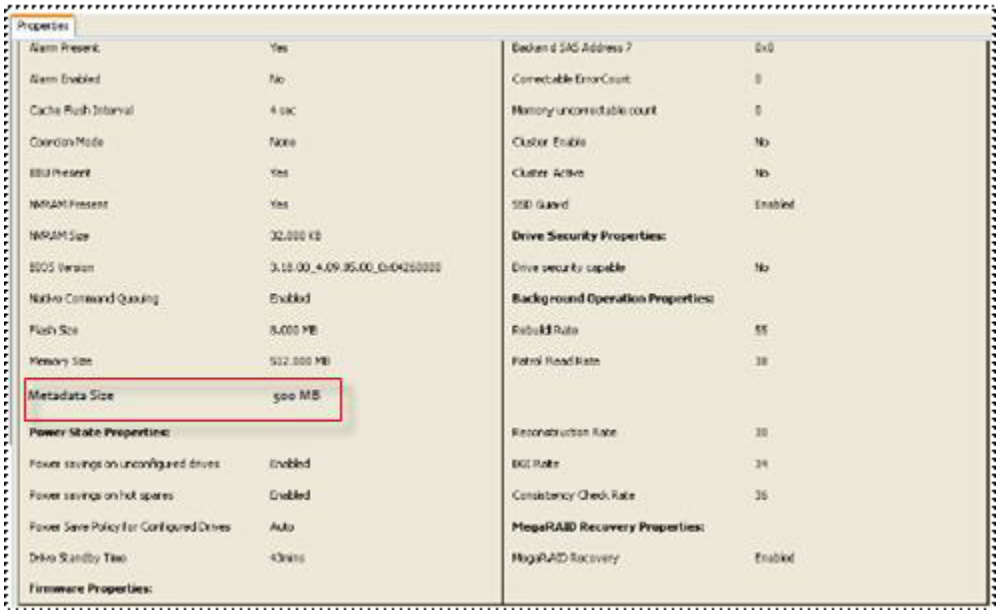

NOTE The size units displayed are the following: if the size is less than 1 MB (1024 KB), the size is displayed in KB. If the size is greater than or equal to 1 MB but less than 1 GB (1024 MB), the size is displayed in MB. If the size is greater than or equal to 1 GB, but less than 1 TB (1024 GB), the size is displayed in GB.

8.6.9 Emergency Spare

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

8.6.9.1 Emergency Spare for Physical Drives

The `Emergency Spare` property determines whether a particular drive is capable of becoming an emergency spare. This property is displayed under the controller properties only if the Global spare for emergency and the Unconfigured Good for emergency controller properties are enabled.

Follow these steps to view the `Emergency Spare` property.

1. Go to either the *Logical* view or the *Physical* view.
2. Click the drive for which you want to view the spare properties.

   The emergency spare is displayed under general properties. This property denotes whether a particular drive is commissioned as an emergency spare or not an emergency spare.

NOTE This property is displayed only for online physical drives.

Figure 174 Emergency Spare – Physical Drive Properties

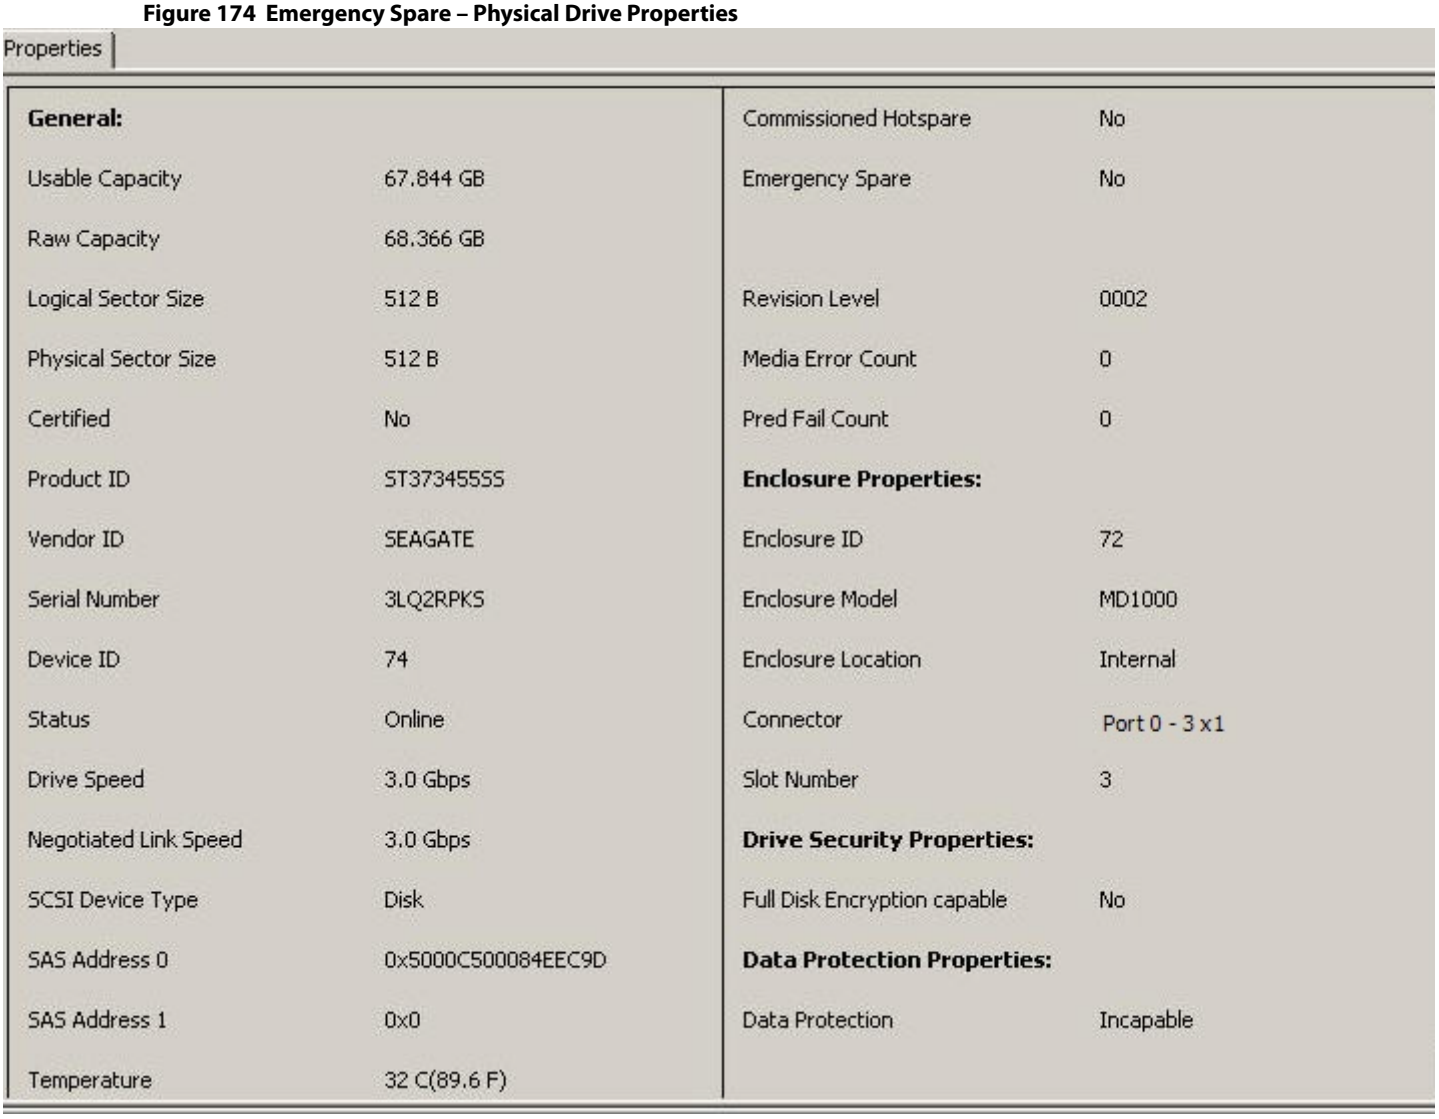

8.6.9.2 Emergency Spare Property for Controllers

The `Emergency Spare` properties under the controller properties are configured based on enabling or disabling the following properties:

- Emergency Spare
- Emergency for SMARTer

To view the `Emergency Spare` property for controllers, click the controller node in the device tree.

The `Emergency Spare` properties are displayed, as shown in the following figure.

Figure 175 Emergency Spare Properties for Controllers

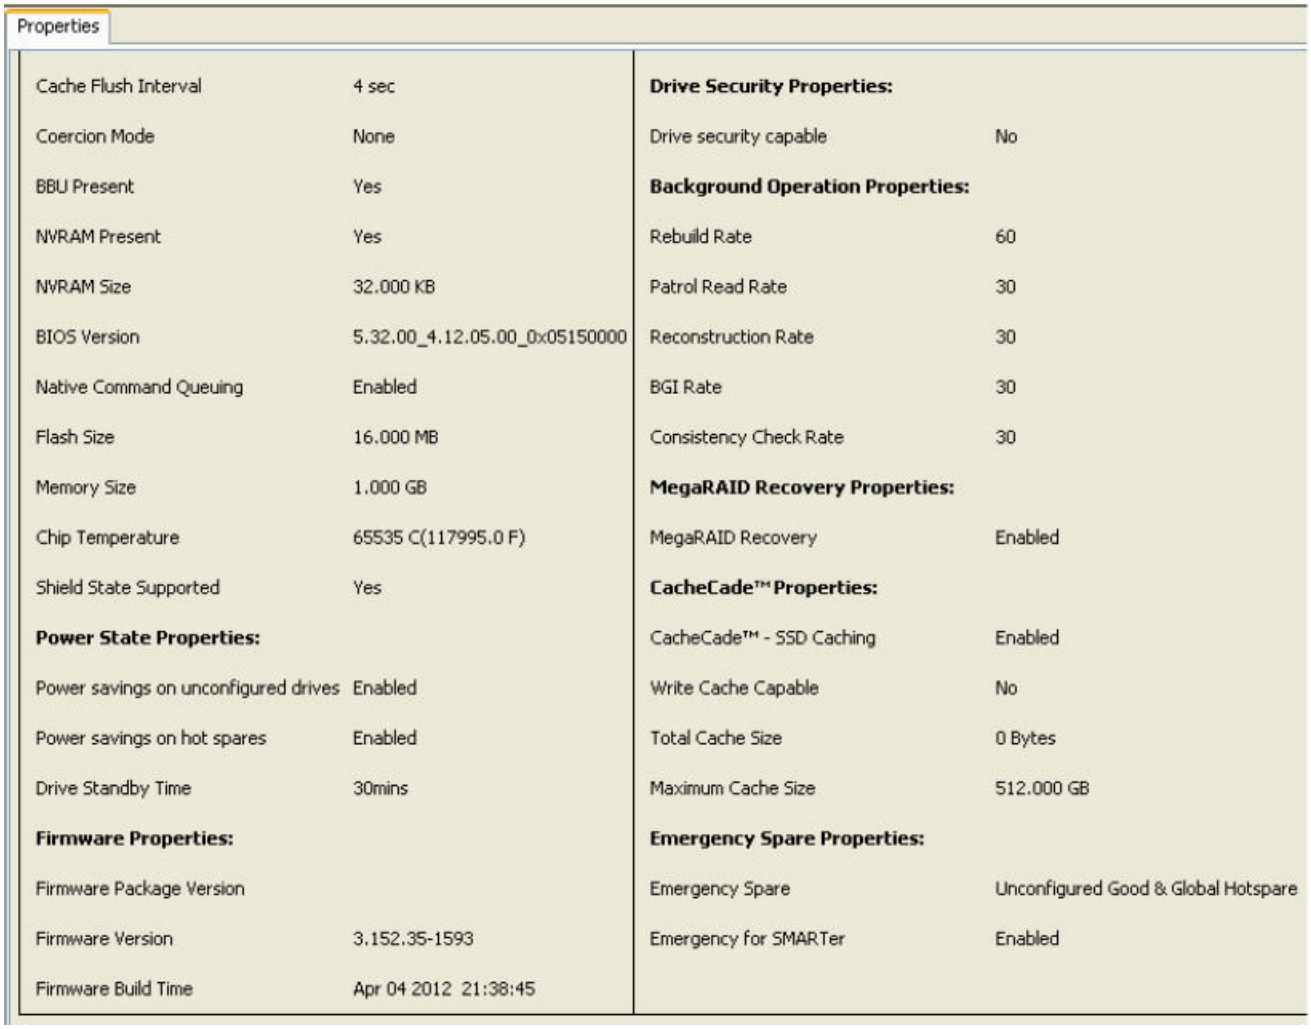

8.6.9.3 Commissioned Hotspare

The commissioned hotspare is used to determine whether the online drive has a Commissioned Hotspare.

To check if the drive is commissioned with a hotspare, click the online physical drive node in the device tree.

The `Commissioned Hotspare` property is displayed, as shown in the following figure. This property is displayed only for online physical drives.

Figure 176 Commissioned Hotspare

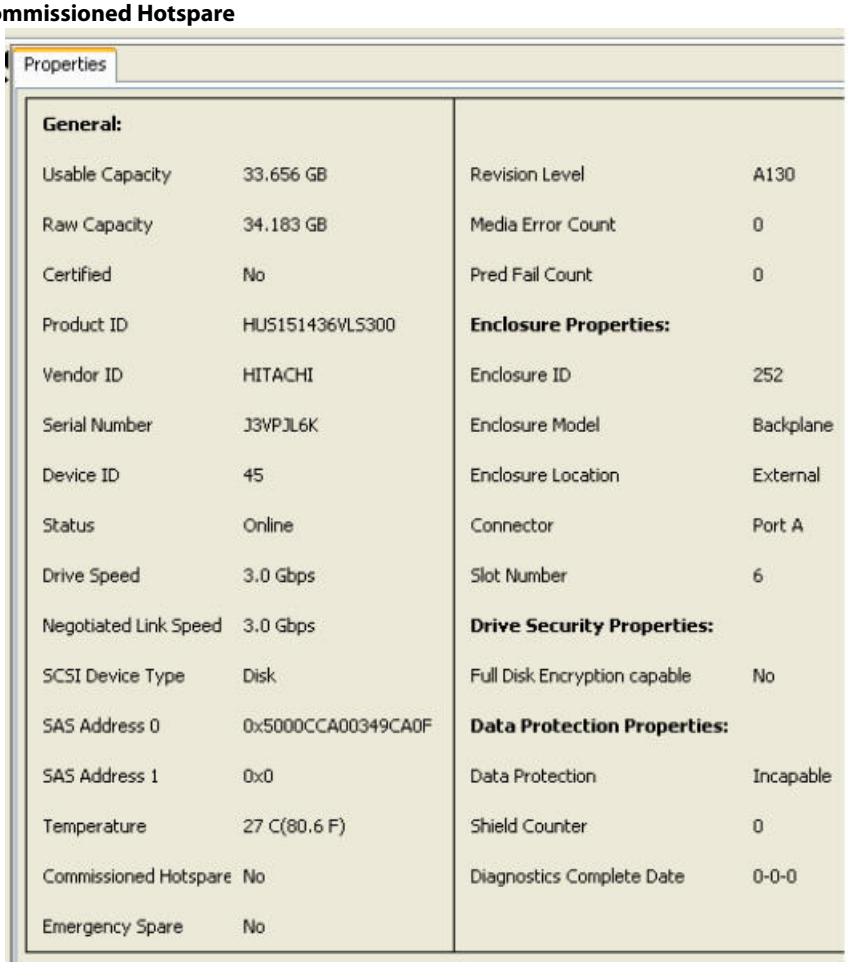

8.6.10 SSD Disk Cache Policy

The MegaRAID firmware provides support to change the write-cache policy for SSD media of individual physical drives.

The MegaRAID firmware does not allow any user application to modify the write-cache policies of any SSD media. The host applications can modify this property through a new logical device (LD) addition or a LD property change. When SSDs are configured in a mixed disk group with HDDs, the Physical Device Write-Cache Policy setting of all the participating drives are changed to match the SSD cache policy setting.

Follow these steps to view the SSD cache property.

1. Click the controller node in the device tree.

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

Figure 177 Controller Properties – SSD Disk Cache Policy

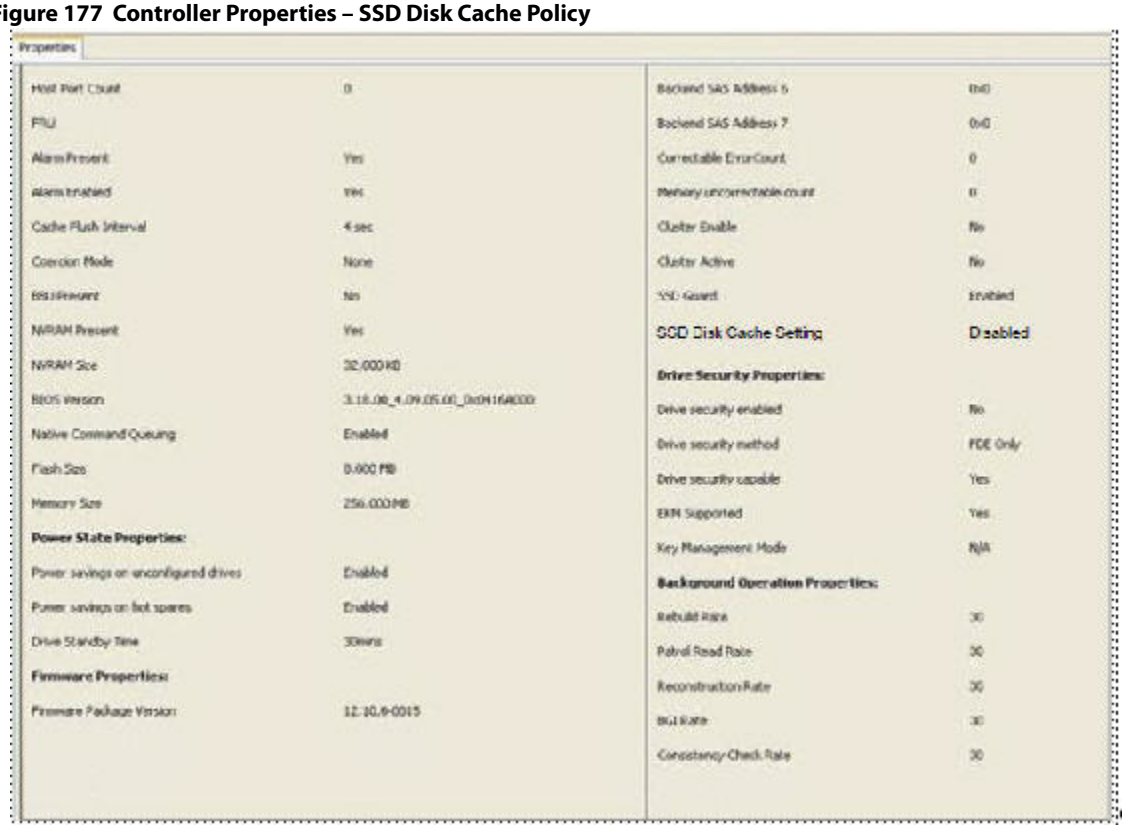

8.6.10.1 Virtual Drive Settings

If the SSD cache property is enabled in the controller properties screen as shown in Controller Properties – SSD Disk Cache Policy, then you cannot select the disk cache policy for the virtual drives having only SSD drives or a mix of SSD drives and HDD drives during virtual drive creation. The value of the disk cache policy is unchanged and the drop-down menu is disabled.

Follow these steps to view the virtual drive settings.

1. Right-click the controller node in the device tree.
2. Select the **Create Virtual Drive** menu option.
3. Select **Advanced Configuration**, and click **Next**.
4. Create **Drive Group**, and click **Next**.

   The **Create Virtual Drive – Virtual drive settings** dialog appears, as shown in the following figure.

Figure 178 Virtual Drive Settings

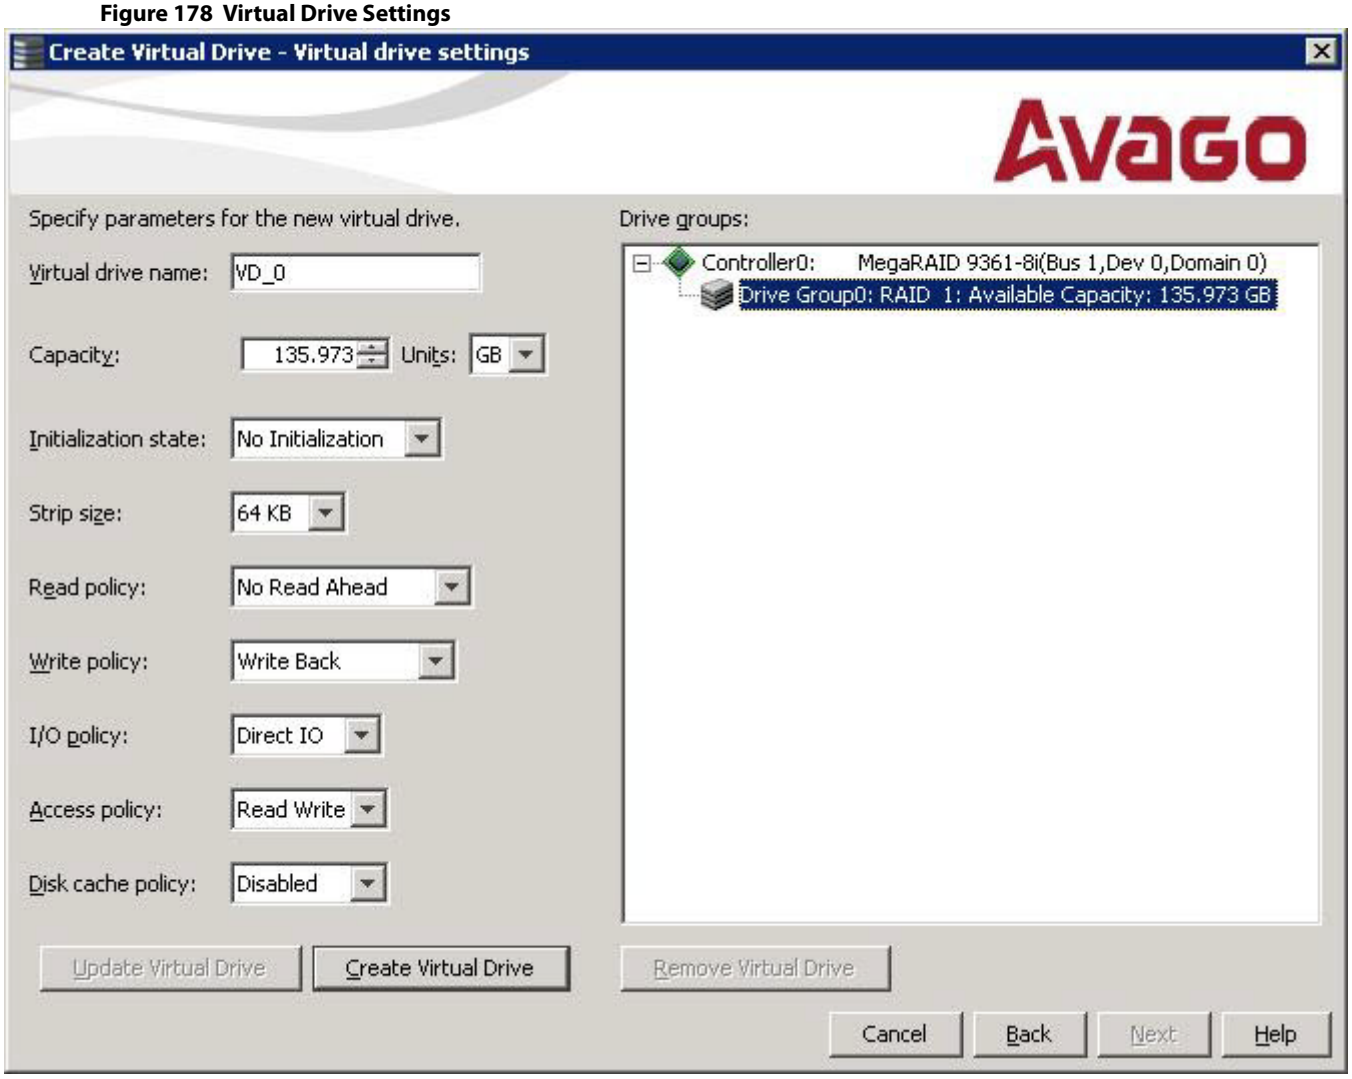

The **Disk Cache Policy** drop-down list is disabled.

8.6.10.2 Set the Virtual Drive Properties

Follow these steps to set the virtual drive properties.

1. Right-click the virtual drive node in the *Logical* view of the device tree.
2. Select **Set Virtual Drive Properties**.
   The **Set Virtual Drive Properties** dialog appears, as shown in the following figure.

Figure 179 Virtual Drive Properties

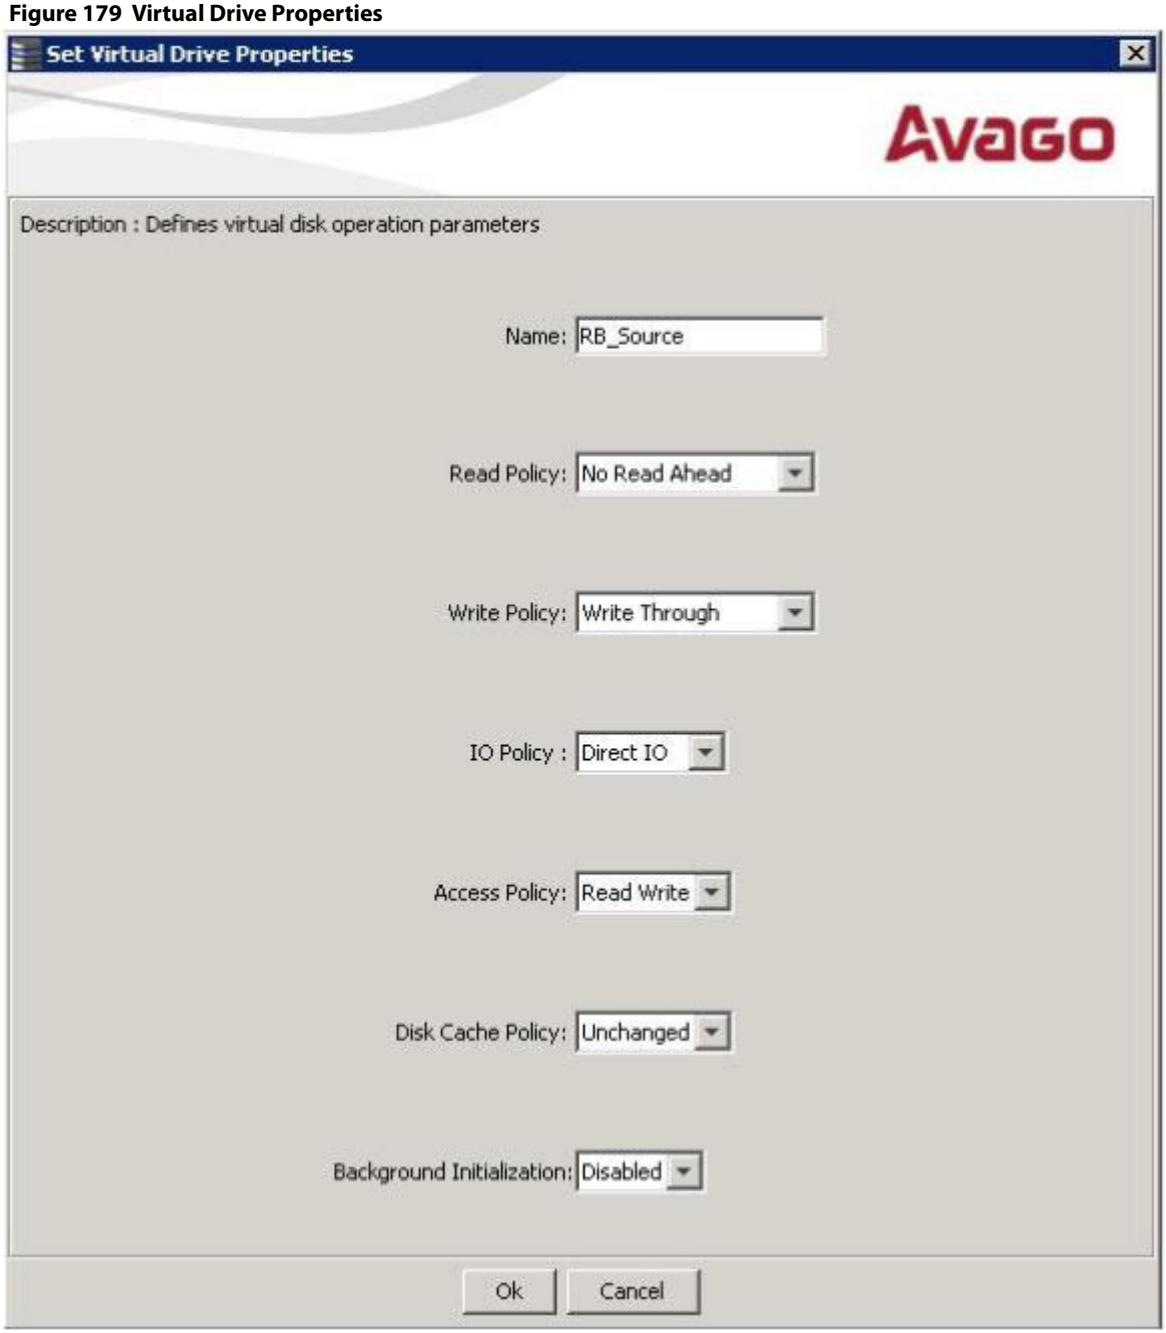

NOTE You cannot select the Disk cache policy for the virtual drives having only SSD drives or a mix of SSD and HDD during VD creation. The value of the Disk Cache Policy is Unchanged and can be set for only HDD drives.

8.6.11 Non-SED Secure Erase Support

This section describes the firmware changes required to securely erase data on non-SEDs (normal HDDs).

SEDs securely erase their internal encryption keys, effectively destroying all of the data present on the drive. For non–SED drives, the erase operation consists of a series of write operations to a drive that overwrites every

user-accessible sector of the drive with specified patterns. It can be repeated in multiple passes using different data patterns for enhanced security. The sanitization technique is more secure than a simple format operation and is commonly called a *clearing* operation, similar to the existing physical drive clear command.

Follow these steps to set physical drive properties.

1. In the *Physical* view, right click the **Physical Drive** node.
2. Select the **Drive Erase** option (Alt+E).

   The **Mode Selection – Drive Erase** dialog appears.

Figure 180 Mode Selection – Drive Erase Window

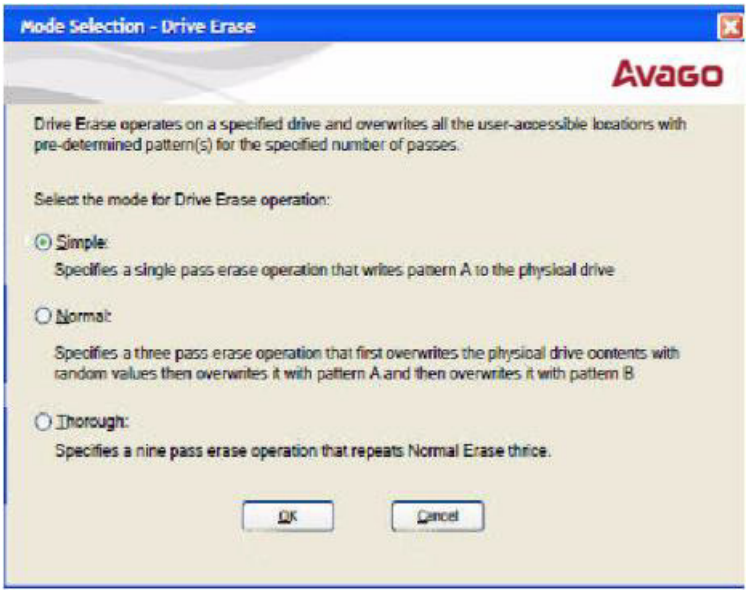

3. You can select the various modes available under the **Select the mode for Drive Erase operation**.
   - **Simple** – (Alt+S). When you select this option and click **OK**, the **Drive Erase** message box appears.

Figure 181 Drive Erase Message

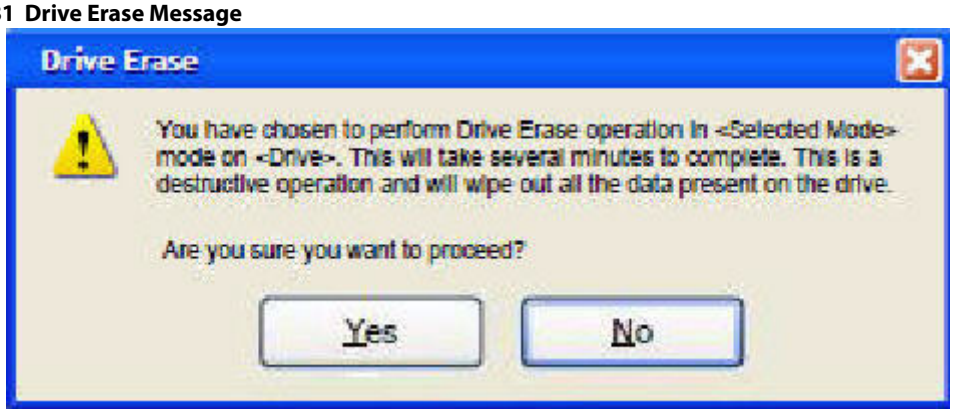

   - **Normal** – (Alt+N). Select this option and click **OK**. The **Drive Erase** message, as shown in the previous figure, appears.
   - **Thorough** – (Alt+T). Select this option and click **OK**. The **Drive Erase** message, as shown in the previous figure, appears.

8.6.11.1 Group Show Progress for Drive Erase

Physical drive erase operation is a time-consuming operation and is performed as a background task. It posts events to notify users of the progress.

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

1. Click the **Show Progress** toolbar icon in the MegaRAID Storage Manager.
   You can also select **Show Progress** from the dashboard or select **Show Progress** from the **Manage** menu.
2. Click the **More info** link under the Background Operations portlet.
   The progress bar appears.

Figure 182 Group Show Progress

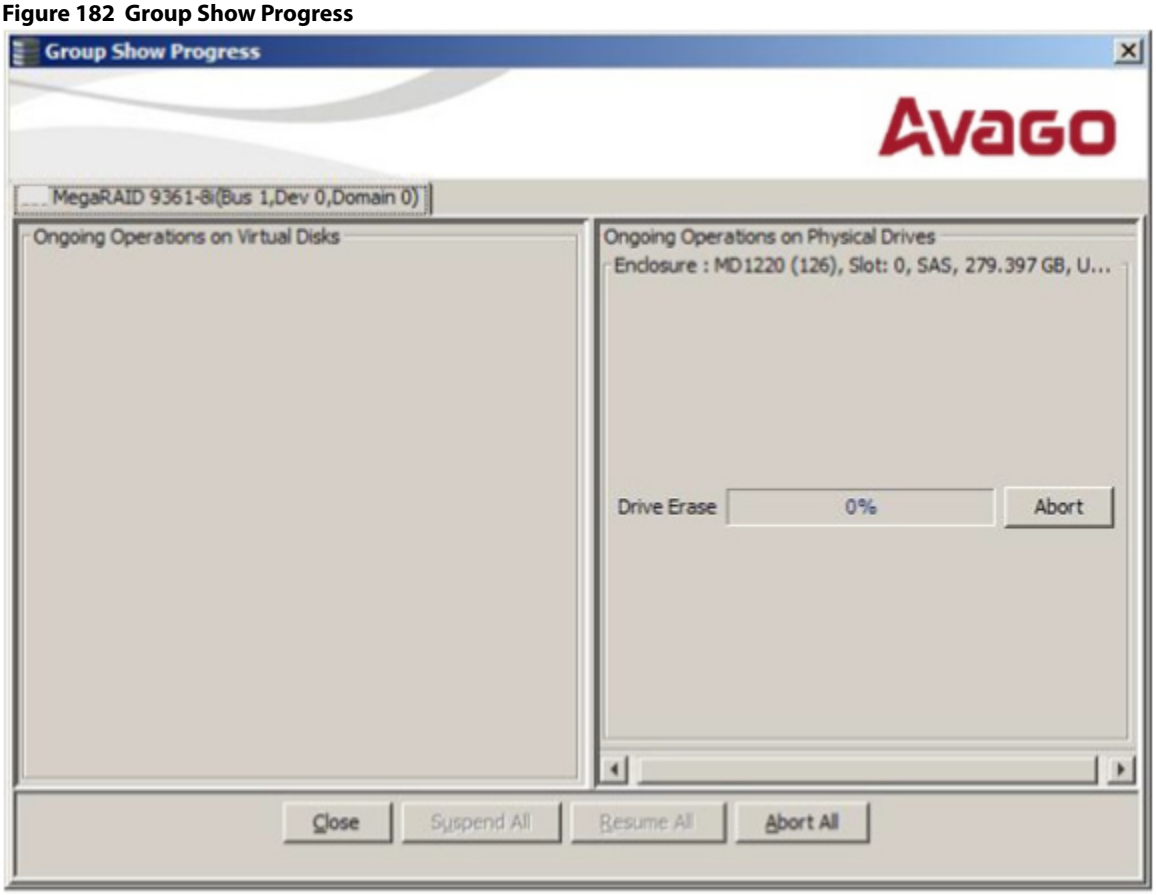

When you click the **Abort All** button, all Drive Erase operations stop, and the progress bar is not displayed.

8.6.11.2 Virtual Drive Erase

Virtual drive erase operates on a specified virtual drive and overwrites all user-accessible locations. It supports non-zero patterns and multiple passes. Virtual drive erase optionally deletes the virtual drive and erases the data within the virtual drive's LBA range. Virtual drive erase is a background operation, and it posts events to notify users of their progress.

Follow these steps to open the Virtual Drive Erase menu.

1. In the **Logical** view, right -click the **Virtual Drive** node.
2. Click on the **Virtual Drive** node, select top level navigation, and click **Go to**.
3. Select **Virtual Drive** and select **Events & Response**.

   The **Logical View – Virtual Drive Erase** menu appears.
4. Select **Virtual Drive Erase**.

   The **Virtual Drive Erase Menu** opens, as shown in the following figure.

Figure 183 Mode Selection – Virtual Drive Erase Dialog

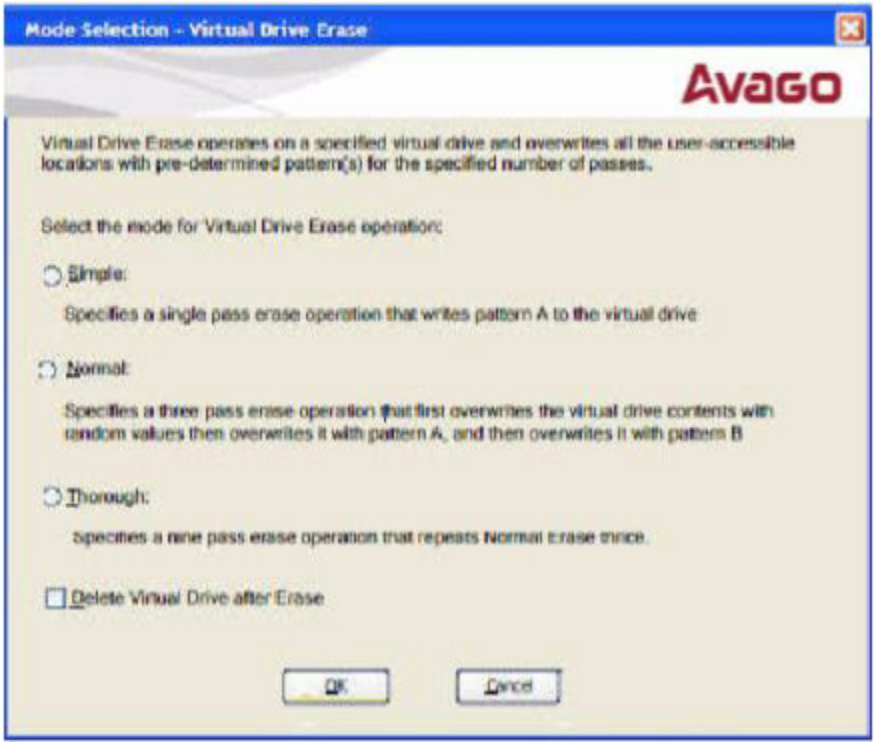

The menu has the following options.

— **Simple** (Alt+S)

After you select this option and click **OK**, and if **Delete Virtual Drive after Erase** is selected, the Warning Message for Virtual Drive Erase figure appears; otherwise, the Warning Message for Virtual Drive Erase without Virtual Drive Delete figure appears.

— **Normal** (Alt+N)

After you select this option and click **OK**, and if **Delete Virtual Drive after Erase** is selected, the Warning Message for Virtual Drive Erase figure appears; otherwise, the Warning Message for Virtual Drive Erase without Virtual Drive Delete figure appears.

— **Thorough** (Alt+T)

After you select this option and click **OK**, and if **Delete Virtual Drive after Erase** is selected, the Warning Message for Virtual Drive Erase figure appears; otherwise, the Warning Message for Virtual Drive Erase without Virtual Drive Delete figure appears.

— **Delete Virtual Drive after Erase** (Alt+D)

When you select this option, the virtual drive is erased and the Warning Message for Virtual Drive Erase figure appears; otherwise, the Warning Message for Virtual Drive Erase without Virtual Drive Delete figure appears.

— **OK** (Alt+O)

Click **OK** and if **Delete Virtual Drive after Erase** is checked, the Warning Message for Virtual Drive Erase figure appears; otherwise, the Warning Message for Virtual Drive Erase without Virtual Drive Delete figure appears.

— **Cancel** (Alt+C)

When you select this option, the dialog closes, and the MegaRAID Storage Manager navigates back to *Physical* view.

Figure 184 Warning Message for Virtual Drive Erase

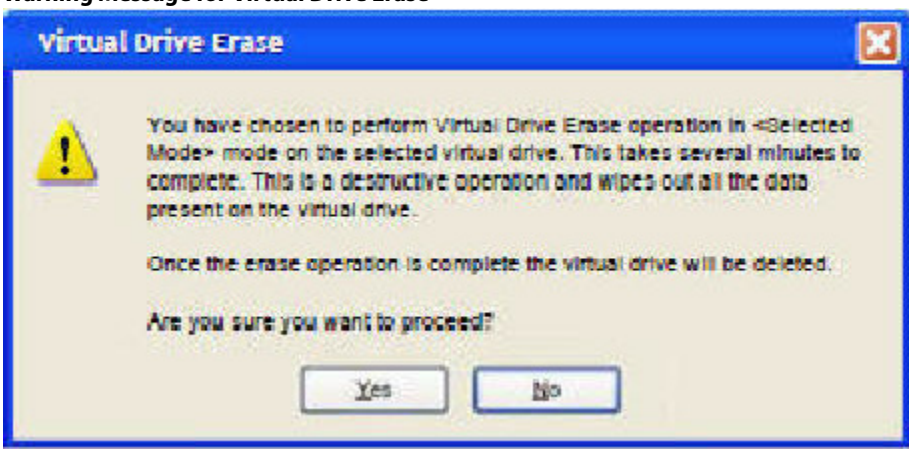

— Click **Yes** to erase the virtual drive.
— Click **No** to cancel the erase and close the dialog.

Figure 185 Warning Message for Virtual Drive Erase without Virtual Drive Delete

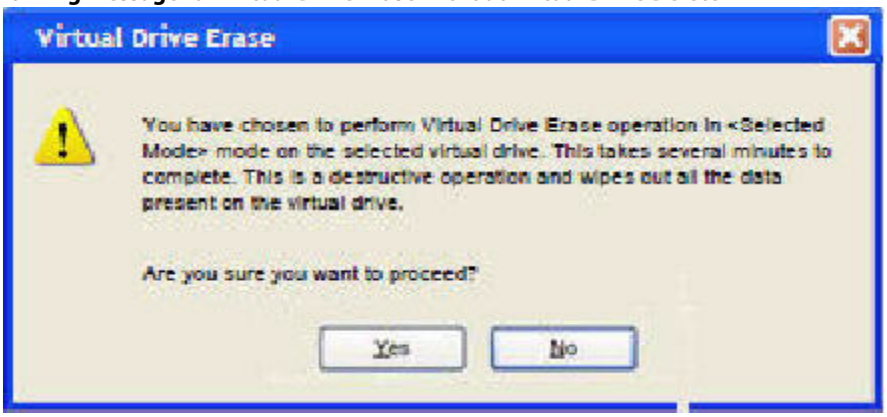

— Click **Yes** to erase the virtual drive.
— Click **No** to cancel the erase and close the dialog.

8.6.11.3 Group Show Progress for Virtual Drive Erase

The virtual drive erase operation is a time-consuming operation and is performed as a background task. It posts events to notify users of the progress.

To view the progress of Group Show Progress – Virtual Drive, click the **Show Progress** toolbar icon.

You can also either select **Show Progress** from the Manage menu, or select the **More info** Link under Background Operations portlet on the dashboard.

The Virtual Drive Erase progress bar appears, as shown in the following figure.

Figure 186 Group Show Progress – Virtual Drive

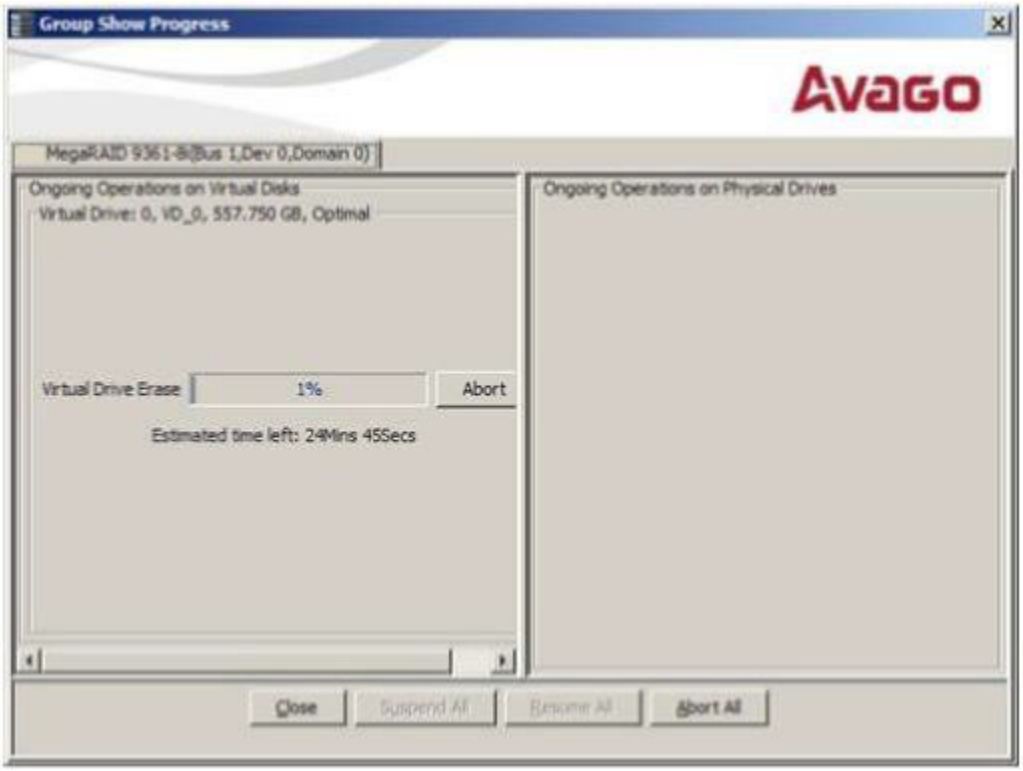

8.6.12 Rebuild Write Cache

MegaRAID firmware supports drive cache properties during a rebuild operation. The MegaRAID solution temporarily enables drive cache for the physical drive that is being rebuilt for the duration of the rebuild operation. Users can enable or disable this feature using the Mega CLI feature.

The MegaRAID software automatically changes the setting for a drive that is being rebuilt. If the PD_CACHE for the rebuilt drive is already set, the firmware does not need to do anything extra.

The firmware identifies and sets the cache policy of the drives whenever a rebuild operation starts and the catch policy is reflected in the event logs. The firmware also makes sure to flush the cache just before committing the drive to the disk group.

8.6.13 Background Suspend/Resume Support

MegaRAID provides a background Suspend or Resume Support feature that enhances the functionality where in the background operations running on a physical drive or a virtual drive can be suspended for some time, and resumed later using the Resume option.

The background operations, including consistency-check, rebuild, replace, and background initialization are supported by an abort operation. If any operation is stopped before completion, it is considered to be aborted. An aborted operation cannot be resumed from the place where it was stopped.

A suspended operation can be resumed later by using the **Resume** option, and the suspended operation resumes from the point where the operation was suspended last.

To perform a suspend and resume operation, go to the **Group Show Progress** dialog, and perform the tasks mentioned below. You also can select **Show Progress** from the **Manage** menu, or select the **More info** link under the **Background Operations** portlet on the dashboard.

The **Group Show Progress** dialog appears, as shown in the following figure. If Patrol Read is running, the **Group Show Progress Patrol Read** dialog appears.

Figure 187 Group Show Progress

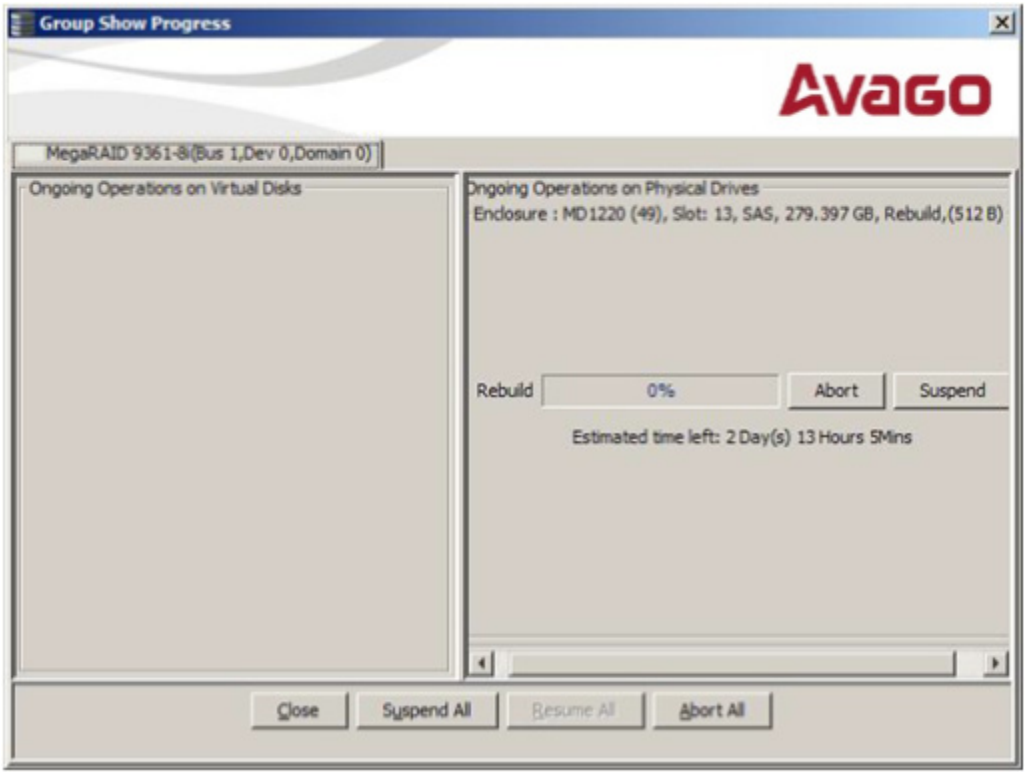

- **Suspend** (Alt+S) – Click the **Suspend** button to suspend the background operation taking place at that particular point of time. When the operations gets suspended, the **Resume** button appears instead of the **Suspend** button.
- **Resume** (Alt+E) – Click the **Resume** button to resume the operation from the point where it was suspended last.
- **Abort** (Alt+B) – Click the **Abort** button to abort the ongoing active operation.
- **Resume All** (Alt+R) – Click the **Resume All** button to resume all the suspended operations from the point they were suspended. This button is disabled if no operations are suspended.
- **Suspend All** (Alt+S) – Click the **Suspend All** button to suspend all the active operations. The **Suspend All** button is enabled only if one or more operations are in active state.
- **Abort All** (Alt+A) – Click the **Abort All** button to abort all the active operations.
- **Close** (Alt+C) – Click the **Close** button to close the dialog.

NOTE **Suspend**, **Resume**, **Suspend All**, and **Resume All** will be applicable only for background initialization, rebuild, replace, and consistency check operations.

Figure 188 Group Show Progress Patrol Read

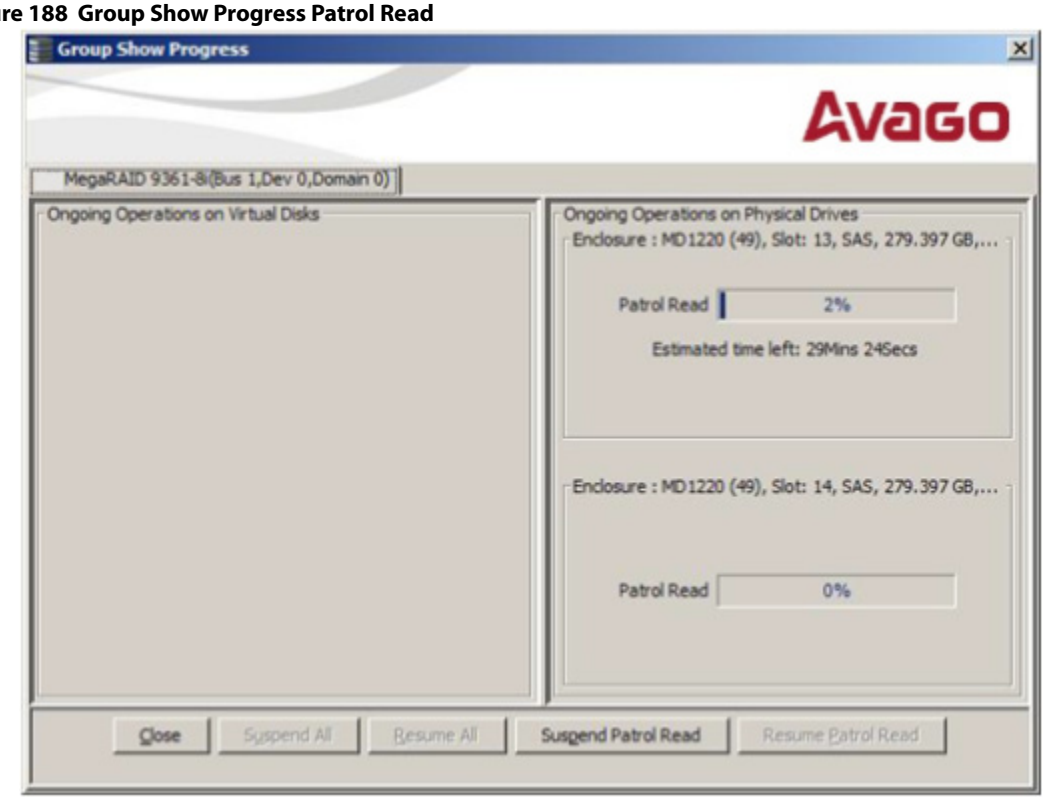

- **Suspend Patrol Read** – Click to suspend the patrol read operation.
- **Resume Patrol Read**- Click to resume the patrol read operation from the point where it was suspended last.

8.6.14 Enclosure Properties

To view the Enclosure properties, in the *Physical* view, click the **Enclosure** node.

The Enclosure properties are displayed, as shown in the following figure.

Figure 189 Enclosure Properties

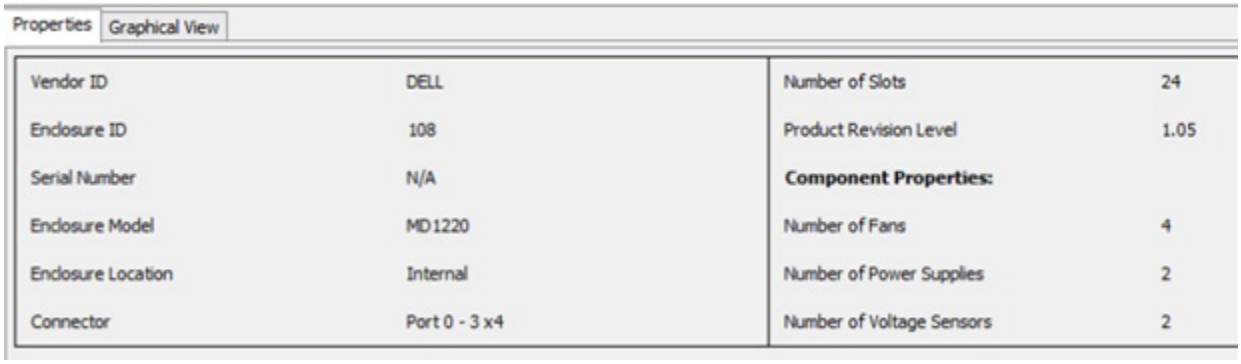

8.6.15 Expander Properties

To view the Expander properties, in the *Physical* view, click the **Expander** node.

The Expander properties are displayed, as shown in the following figure.

Figure 190 Expander Properties

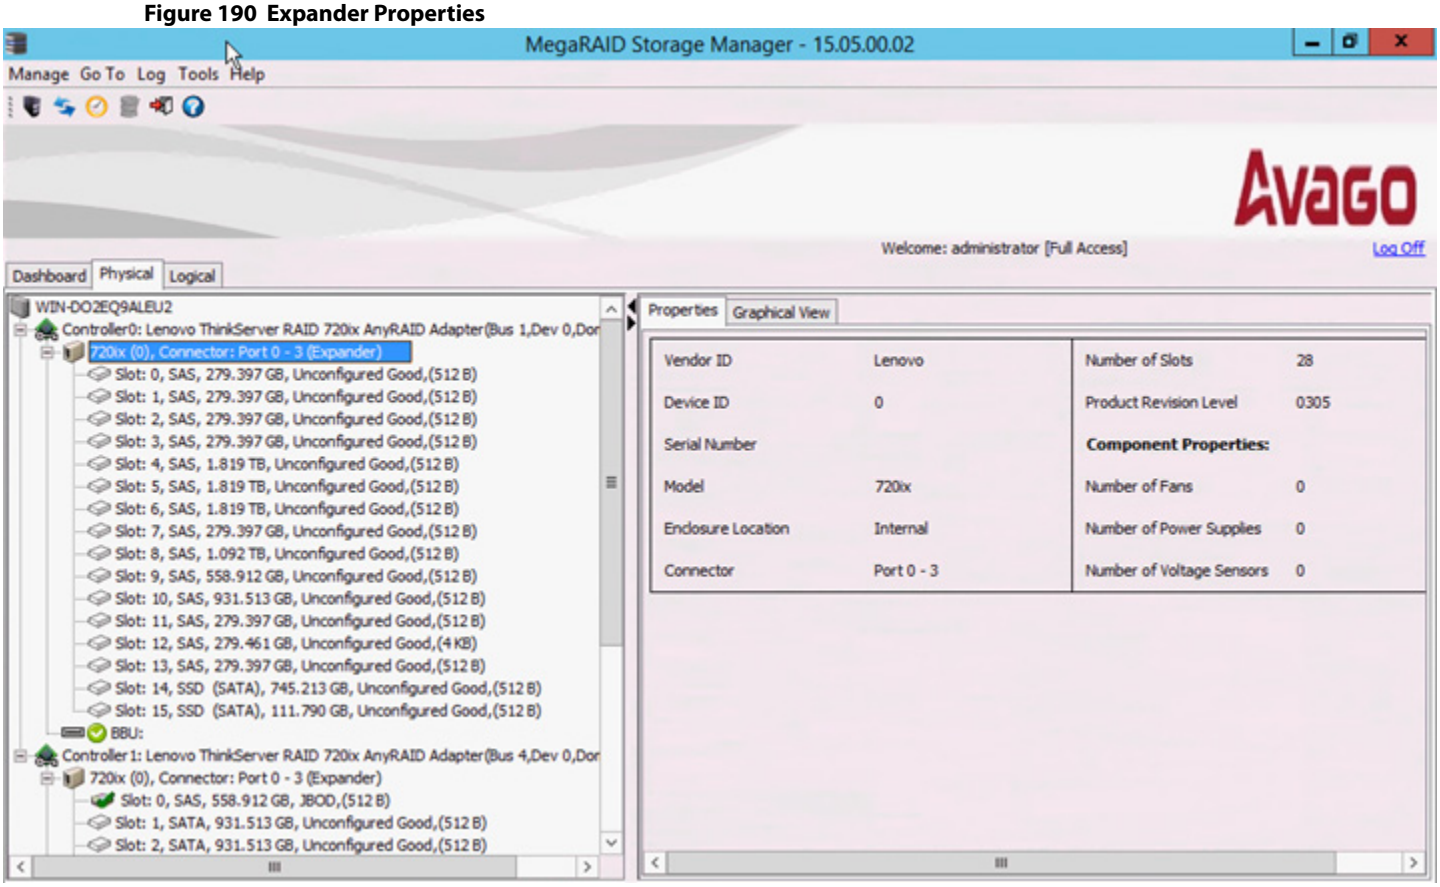

8.7 GUI Elements in the MegaRAID Storage Manager Window and Menus

This section describes the graphical user interface (GUI) elements used in the MegaRAID Storage Manager software.

8.7.1 Device Icons

The following icons in the left panel represent the controllers, drives, and other devices.

	Status
	System
	Controller
	Backplane
	Enclosure
	Port

	Drive group
	Virtual drive
	Online drive
	Power save mode
	Dedicated hotspare
	Global hotspare
	Battery backup unit (BBU)
	Tape drive
	CD-ROM
	Foreign drive
	Unconfigured drive
	Locked SED
	Unlocked SED

NOTE The MegaRAID Storage Manager software shows the icons for tape drive devices; however, no tape-related operations are supported by the utility. If these operations are required, use a separate backup application.

A red circle to the right of an icon indicates that the device has failed. For example, this icon indicates that a drive has failed: .

A yellow circle to the right of an icon indicates that a device is running in a partially degraded state. For example, this icon indicates that a virtual drive is running in a degraded state because a controller has failed.

An orange circle to the right of an icon indicates that a device is running in a degraded state.

8.7.2 Properties and Graphical View Tabs

The right panel of the MegaRAID Storage Manager window has one tab or two tabs, depending on which type of device you select in the left panel.

- The **Properties** tab displays information about the selected device.

  For example, if you select a controller icon in the left panel, the **Properties** tab lists information about the controller, such as the controller name, NVRAM size, and device port count.
- 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 191 Properties Tab and Graphical View Tab

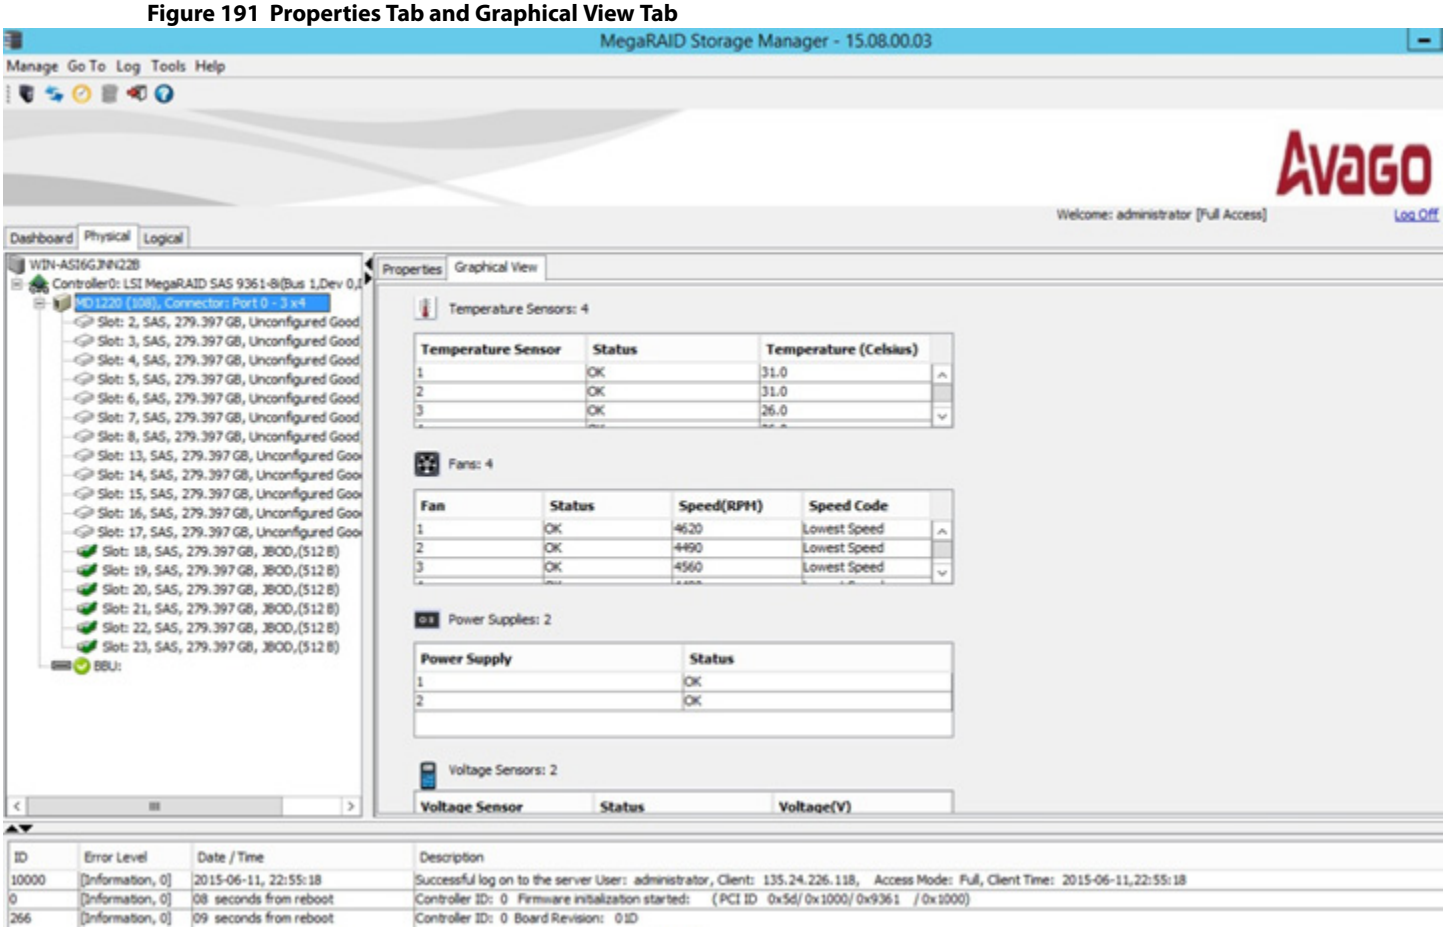

8.7.3 Event Log Panel

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 **Monitoring Controllers and Their Attached Devices**. For more information about the event log entries, see Events and Messages.

8.7.4 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 the **Configuration** and **Maintaining and Managing Storage Configurations** sections.

Manage Menu

The Manage menu has a `Refresh` option for updating the display in the **MegaRAID Storage Manager** window (refresh is seldom required; the display usually updates automatically) and an `Exit` option to end your session on MegaRAID Storage Manager. The `Server` option shows all the servers that were discovered by a scan. In addition, you can perform a check consistency, initialize multiple virtual groups, and show the progress of group operations on virtual drives.

Go To Menu

The **Go To** menu is available when you select a controller, drive group, physical drive, virtual drive, or battery backup unit in the main menu screen. The menu options vary depending on the type of device selected in the left panel of the MegaRAID Storage Manager main menu. The options also vary depending on the current state of the selected device. For example, if you select an offline drive, the **Make Drive Online** option appears in the **Physical Drive** menu.

Configuration options are also available. This is where you access the Configuration Wizard that you use to configure drive groups and virtual drives To access the Wizard, select the controller in the left panel, and then select **Go To > Controller > Create Virtual Drive**.

Log Menu

The **Log** menu includes options for saving and clearing the message log. For more information about the Log menu, see Events and Messages.

Tools Menu

On the **Tools** menu, you can select to access the **Configure Alerts** dialog, where you can set the alert delivery rules, event severity levels, exceptions, and e-mail settings. For more information, see Configuring Alert Notifications.

Help Menu

On the **Help** menu, you can select to view the MegaRAID Storage Manager online help file. You can select to view version information for the MegaRAID Storage Manager software.

NOTE When you use the MegaRAID Storage Manager online help, you might see a warning message that Internet Explorer® has restricted the file from showing active content. If this warning appears, click on the active content warning bar, and enable the active content.

NOTE If you are using the Linux operating system, you must install Firefox® browser or Mozilla® browser for the MegaRAID Storage Manager online help to display.

NOTE When connected to the VMware server, only the IP address and the host name information appear. The other information, such as the operating system name, version, and architecture do not appear.

Chapter 9: Configuration

This chapter explains how to use MegaRAID Storage Manager software to create and modify storage configurations on Avago SAS controllers.

The Avago 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 Introduction to RAID.

NOTE You cannot create or modify a storage configuration unless you are logged on to a server with administrator privileges.

NOTE The MegaRAID Storage Manager software supports 64 virtual drive creation. It does not support 240 virtual drive creation. If you have created 240 virtual drives with some other application and then launch the MegaRAID Storage Manager software, it displays only 64 virtual drives. For more information see the 240 Virtual Drive Feature Limitations appendix.

9.1 Creating a New Configuration

You can use the MegaRAID Storage Manager software to create new storage configurations on systems with SAS controllers. You can create the following types of configurations:

- **Simple configuration** specifies a limited number of settings and has the system select drives for you.
  This option is the easiest way to create a virtual drive.
- **Advanced configuration** lets you choose additional settings and customize virtual drive creation.
  This option provides greater flexibility when creating virtual drives for your specific requirements.

This section describes the virtual drive parameters and explains how to create simple and advanced storage configurations.

9.1.1 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** (Default)
    The new configuration is not initialized, and the existing data on the drives is not overwritten.
  - **Fast Initialization**
    The firmware quickly writes 0s to the first and last 8-MB regions of the new virtual drive and then completes the initialization in the background. This allows you to start writing data to the virtual drive immediately.
  - **Full Initialization**
    A complete initialization is done on the new configuration. You cannot write data to the new virtual drive until the initialization is complete. This process can take a long time if the drives are large.

NOTE BGI is supported only for RAID 5 and RAID 6 and not for any other RAID levels. New RAID 5 virtual drives require at least five drives for a background initialization to start. New RAID 6 virtual drives require at least seven drives for a background initialization to start. If there are fewer drives, the background initialization does not start.

- **Strip size**

  Strip sizes of 8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB, and 1024 KB are supported. The default is 64 KB. For more information, see the *striping* entry in the **Appendix J**.

  NOTE The supported strip size can vary from a minimum of 64 KB to 1 MB for MegaRAID controllers and only 64 KB for Integrated MegaRAID controllers. The LSISAS2108 controller supports strip size from 8 KB to 1 MB.

- **Read policy**

  Specify the read policy for this virtual drive:

  — **Always read ahead**

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

  — **No read ahead** (Default)

  Disables the read ahead capability.

- **Write policy**

  Specify the write policy for this virtual drive:

  — **Write Through**

  In this mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This option eliminates the risk of losing cached data in case of a power failure.

  — **Always Write Back**

  In this mode, the controller sends a data transfer completion signal to the host when the controller cache receives all the data in a transaction. If you select the Always Write Back policy and the battery is absent, the firmware is forced to use the Write Back policy.

  — **Write Back** (Default)

  In this mode, the controller sends a data transfer completion signal to the host when the controller cache receives all of the data in a transaction. If you select the Write Back policy and the battery is absent, the firmware disables the Write Back policy and defaults to the Write Through policy. This option provides a good balance between data protection and performance.

  NOTE The write policy depends on the status of the BBU. If the BBU is not present, is low, is failed, or is being charged, the current write policy switches to Write Through, which provides better data protection.

- **I/O policy**

  The I/O policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.

  — **Cached IO**

  In this mode, all reads are buffered in cache memory.

— **Direct IO** (Default)

In this mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory.

Cached IO provides faster processing, and **Direct IO** makes sure that the cache and the host contain the same data.

- **Access policy**

  Select the type of data access that is allowed for this virtual drive.

  — **Read/Write** (Default)

  Allow read/write access. This setting is the default value.

  — **Read Only**

  Allow read-only access.

  — **Blocked**

  Do not allow access.

- **Disk cache policy**

  Select a cache setting for this drive:

  — **Enabled**

  Enable the disk cache.

  — **Disabled** (Default)

  Disable the disk cache.

  — **Unchanged**

  Leave the current disk cache policy unchanged.

9.1.2 Optimum Controller Settings for CacheCade

Write Policy: Write Back/Write Through/Always Write Back

9.1.3 Optimum Controller Settings for Fast Path

Write Policy: Write Through

IO Policy: Direct IO

Read Policy: No Read Ahead

Stripe Size: 64 KB

9.1.4 Creating a Virtual Drive Using Simple Configuration

Simple configuration is the quickest and easiest way to create a new storage configuration. When you select simple configuration mode, the system creates the best configuration possible using the available drives.

NOTE You cannot create spanned drives using the simple configuration procedure. To create spanned drives, use the advanced configuration procedure described in Creating a Virtual Drive Using Advanced Configuration.

NOTE Physical drives which are in **Prepare for Removal** state are not considered for virtual drive creation.

Follow these steps to create a new storage configuration in simple configuration mode.

1. Perform either of the following steps:
   - Right-click the controller node in the device tree in the left frame of the **MegaRAID Storage Manager** window, and select **Create Virtual Drive**.
   - Select the controller node, and select **Go To > Controller > Create Virtual Drive** in the menu bar, as shown in the following figure.

Figure 192 Create Virtual Drive Menu Option

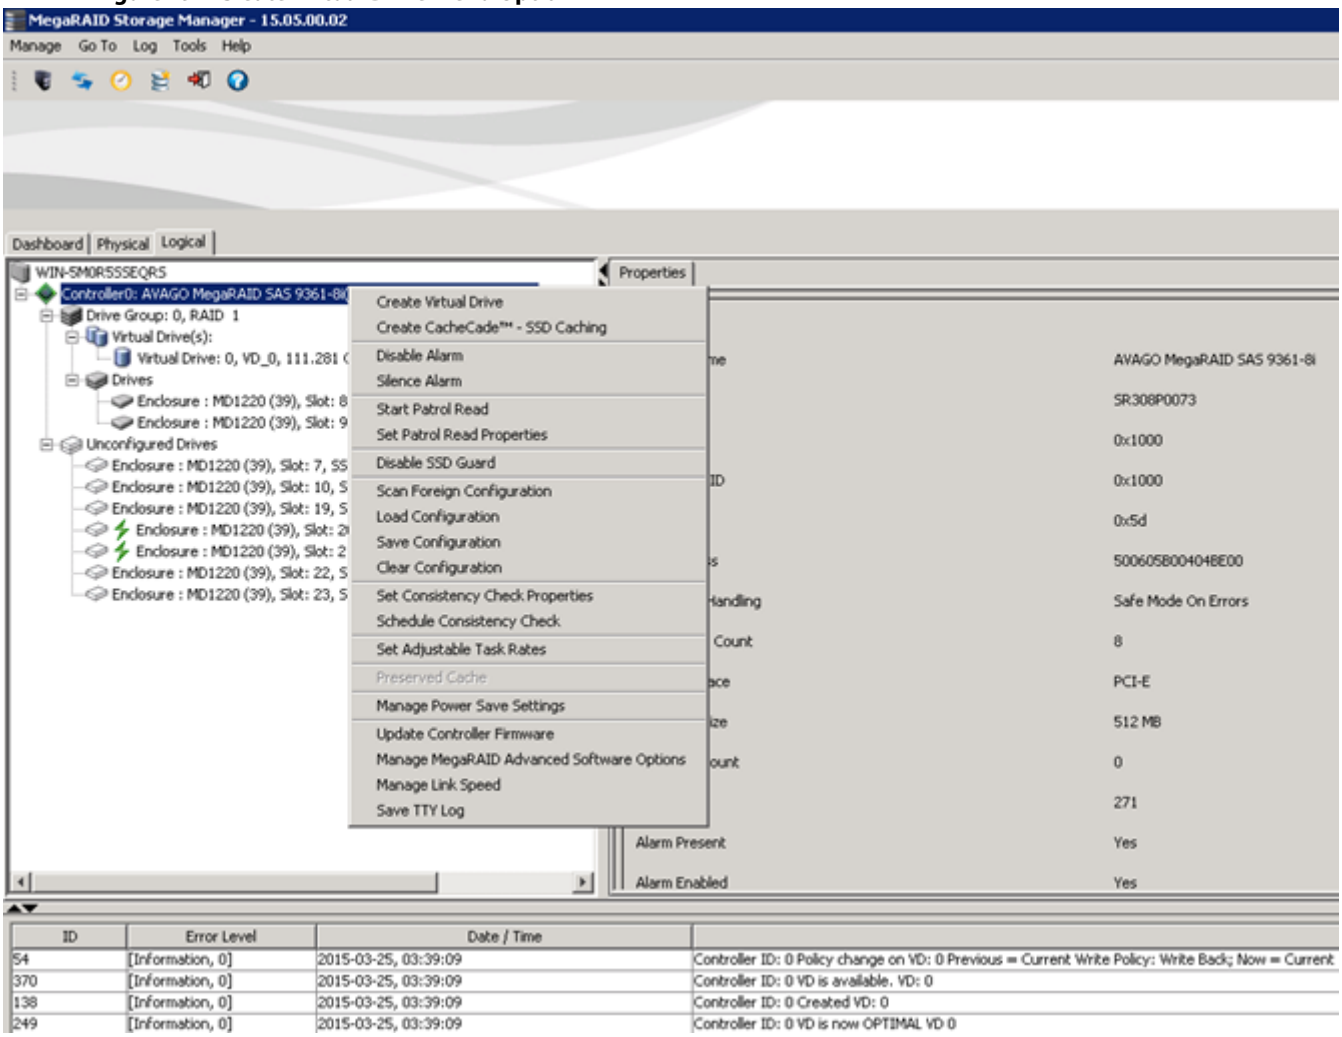

The dialog for the configuration mode (simple or advanced) appears, as shown in the following figure.

Figure 193 Create Virtual Drive – Choose mode

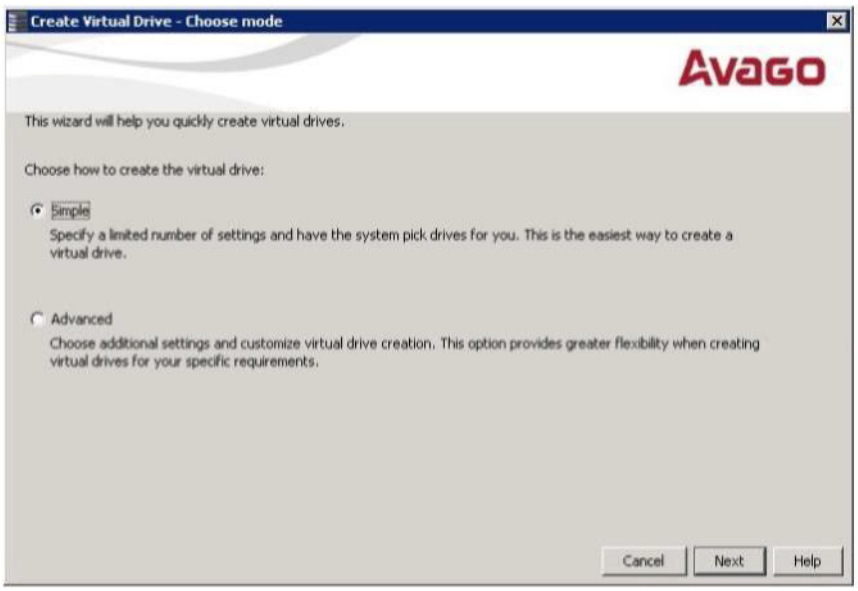

2. Select the **Simple** radio button, and click **Next**.

   The **Create Virtual Drive – Allocate capacity** dialog appears, as shown in the following figure. If unconfigured drives are available, you have the option to use those unconfigured drives. If unconfigured drives are available, the **Create Drive Group Settings** window appears, and you can go to step 4.

Figure 194 Using the Free Capacity of an Existing Drive Group

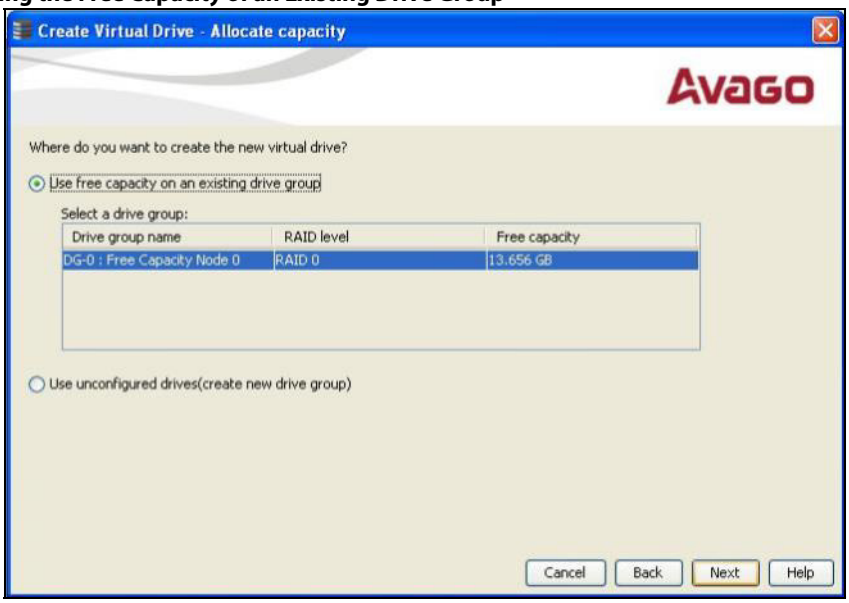

3. Perform either of the two options:
   — If a drive group exists, select the **Use free capacity on an existing drive group** radio button and click **Next**. Continue with step 4. The **Create Virtual Drive** window appears, as shown in the following figure. If different

types of drives are attached to the controller, such as HDD, SDD, SAS, and SATA, an option appears to allow drive type mixing.

— If unconfigured drives are available, select the radio button to use the unconfigured drives, and click **Next**. Continue with step 10. The **Summary** window appears as shown in the Create Virtual Drive – Summary Window figure.

Figure 195 Create Virtual Drive – Drive group and Virtual drive settings Dialog

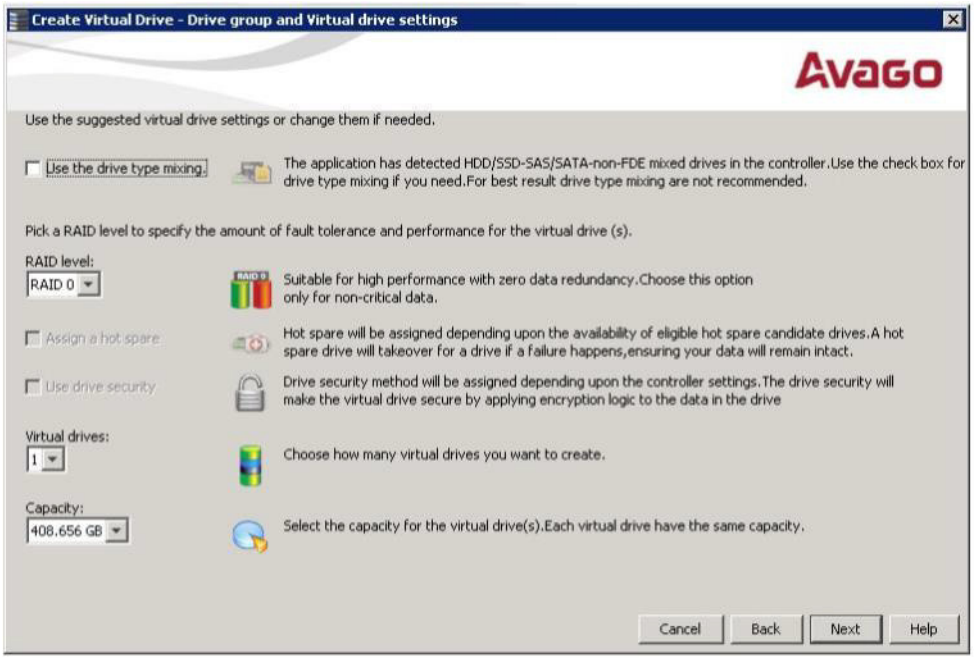

4. If you want to allow different types of drives in a configuration, select the **Use the drive type mixing** check box.

   NOTE For best results, do not use drive type mixing.

5. Select the RAID level desired for the virtual drive.

   When you use simple configuration, the RAID controller supports RAID levels 1, 5, and 6. In addition, it supports independent drives (configured as RAID 0). The window text gives a brief description of the RAID level that you select. The RAID levels that you can choose depend on the number of drives available.

6. Select the **Assign a hot spare** check box if you want to assign a dedicated hot spare to the new virtual drive.

   If an unconfigured good drive is available, that drive is assigned as a hot pare. Hot spares are drives that are available to replace failed drives automatically in a redundant virtual drive (RAID 1, RAID 5, or RAID 6).

7. Select the **Use drive security** check box if you want to set a drive security method.

   The Avago 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. See Avago MegaRAID SafeStore Encryption Services, for more information about the SafeStore feature.

8. Use the drop-down list in the **Virtual drives** field to choose how many virtual drives you want to create.

9. Select the capacity of the virtual drives.

   Each virtual drive has the same capacity.

10. Click **Next**.

    The **Create Virtual Drive – Summary** window appears, as shown in the following figure. This window shows the selections you made for simple configuration.

Figure 196 Create Virtual Drive – Summary Window

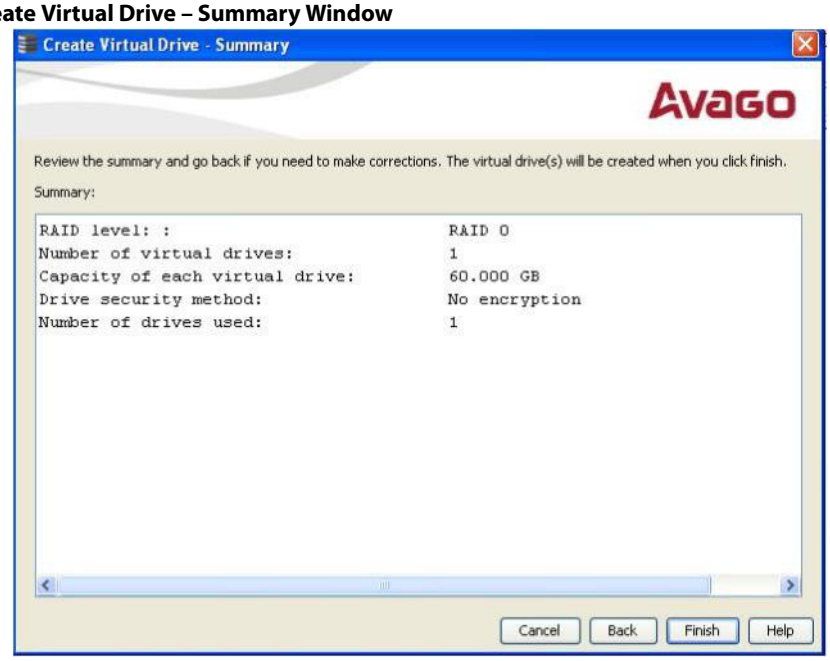

NOTE If High Availability DAS is supported on the controller and you are creating a virtual drive using simple configuration, by default, the virtual drive is shared with the other servers in that cluster.

11. Either click **Back** to return to the previous window to change any selections, or click **Finish** to accept and complete the configuration.

    The new virtual drive is created after you click **Finish**. After the configuration is completed, a dialog box notifies you that the virtual drives were created successfully.

NOTE If you create a large configuration using drives that are in Power-Save mode, it could take several minutes to spin up the drives. A progress bar appears as the drives spin up. If any of the selected unconfigured drives fail to spin up, a dialog box that identifies these drives appears.

9.1.5 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.

NOTE Physical drives which are in **Prepare for Removal** state are not listed for virtual drive creation.

Follow these steps to create a new storage configuration in the advanced configuration mode. This example shows the configuration of a spanned drive group.

1. Perform either of the following steps to bring up the **Configuration** wizard:
   - Right-click the controller node in the device tree in the left frame of the **MegaRAID Storage Manager** window, and select **Create Virtual Drive**.
   - Select the controller node, and select **Go To > Controller > Create Virtual Drive** in the menu bar.

   The dialog for the choosing the configuration mode (simple or advanced) appears, as shown in the following figure.

Figure 197 Create Virtual Drive - Choose mode Dialog

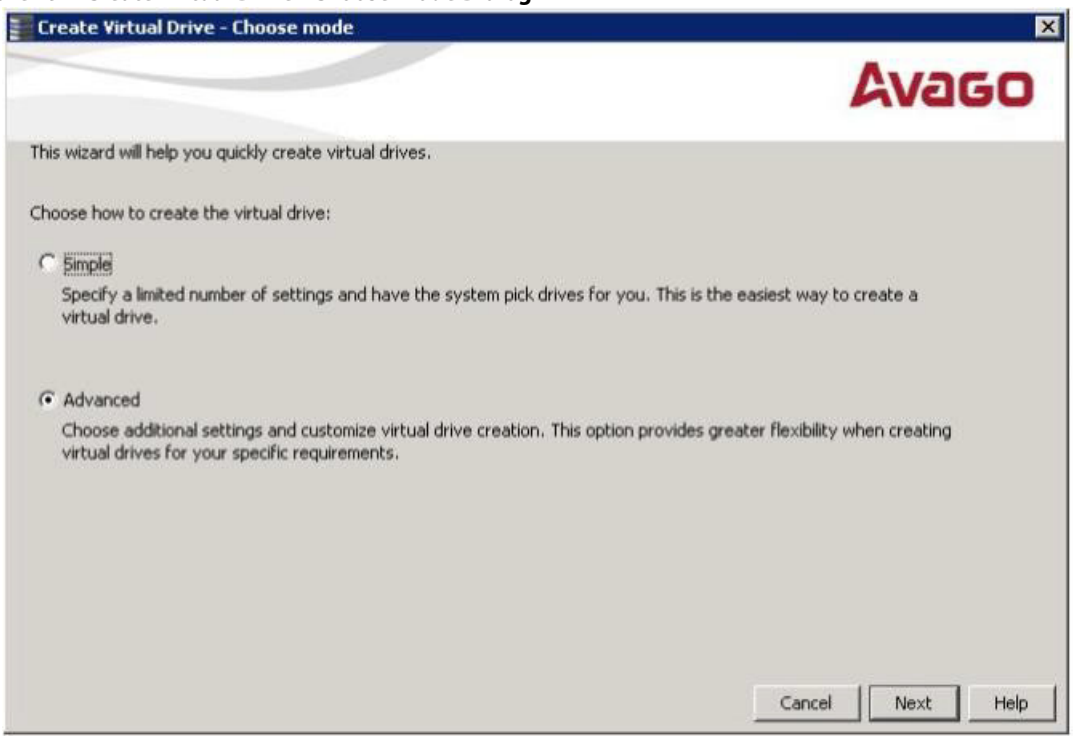

2. Select the **Advanced** radio button, and click **Next**.

   The **Create Drive Group Settings** window appears, as shown in the following figure.

Figure 198 Create Drive Group - Drive Group Settings Window

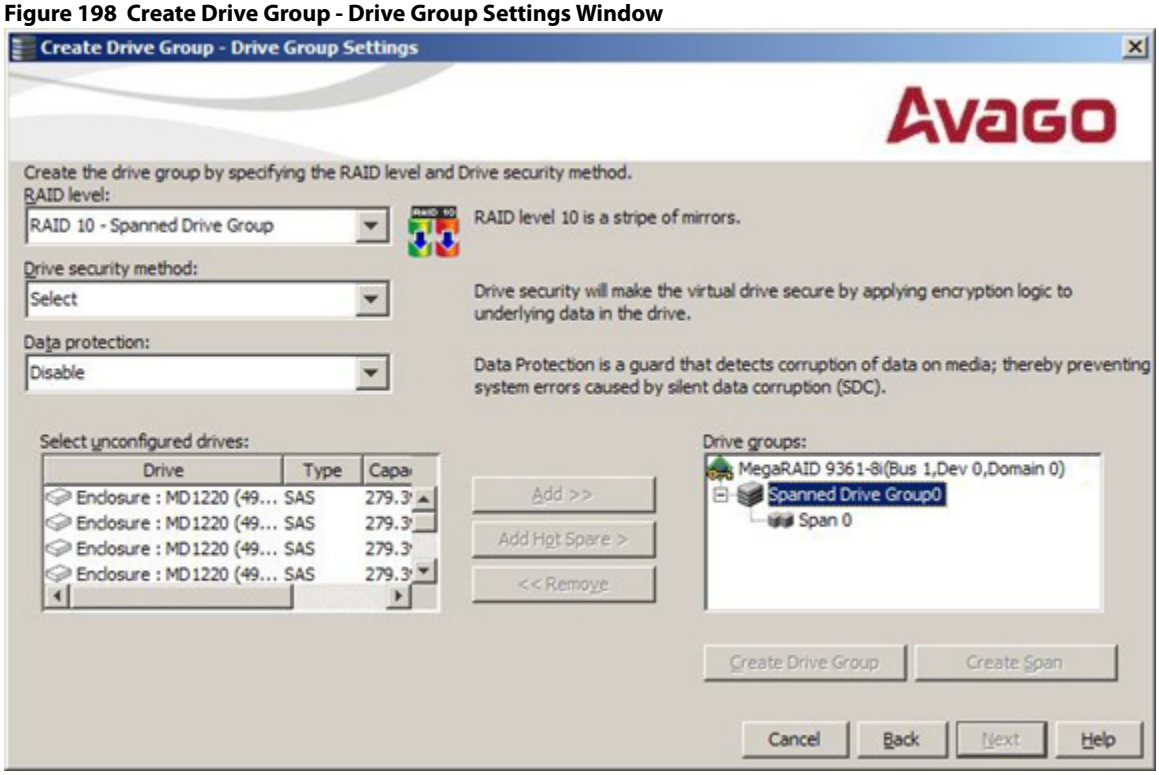

3. Select the following items on the **Create Drive Group - Drive Group Settings** window:
   a. Select the RAID level desired for the drive group from the drop-down menu. To make a spanned drive, select **RAID 10**, **RAID 50**, or **RAID 60** in the **RAID level** field.

      Drive Group 0 and **Span 0** appear in the **Drive groups** field when you select RAID 10, 50, or 60.

      The RAID controller supports RAID levels 1, 5, 6, 10, 50, and 60. In addition, it supports independent drives (configured as RAID 0 and RAID 00). The dialog text gives a brief description of the RAID level that you select. You can choose the RAID levels depending on the number of available drives.

   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 Avago MegaRAID SafeStore Encryption Services, for more information about drive security and encryption.

   c. Select *unconfigured* drives from the list of drives, and click **Add>** to add them to the drive group.

      The selected drives appear under **Span 0** below **Drive Group 0**, as shown in the following figure.

Figure 199 Span 0 of Drive Group 0

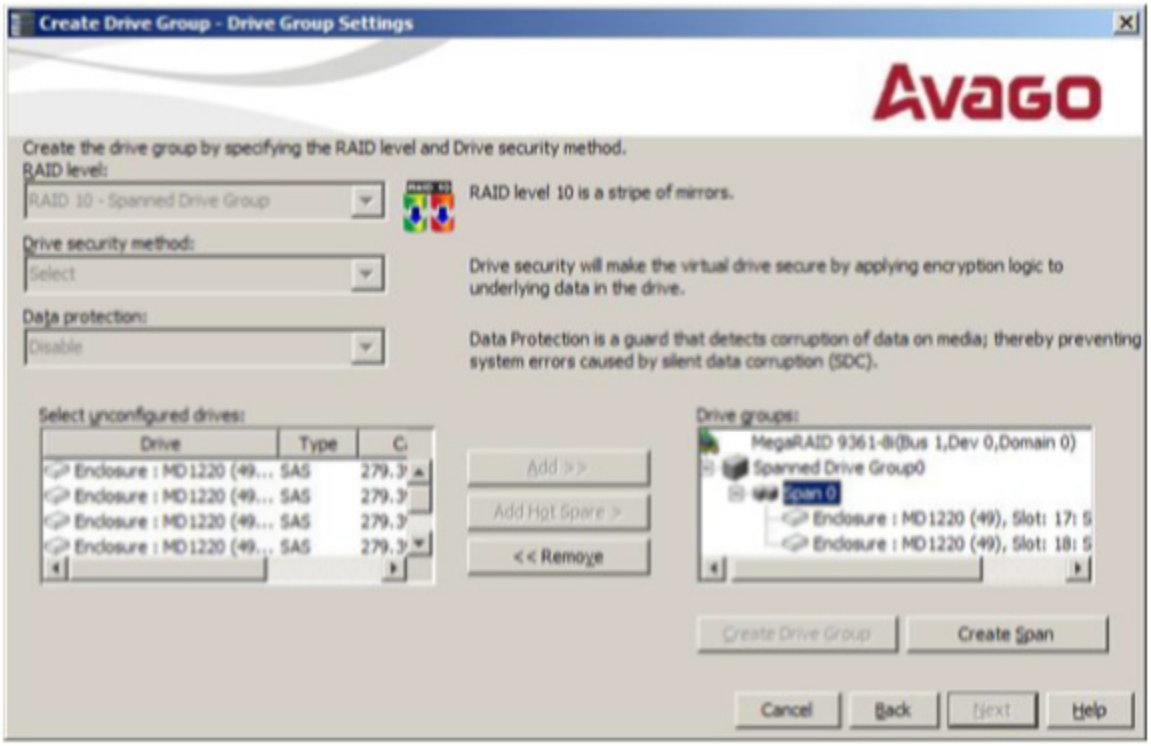

d. Click **Create Span** to create a second span in the drive group.

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 the following figure.

Figure 200 Span 0 and Span 1 of Drive Group 0

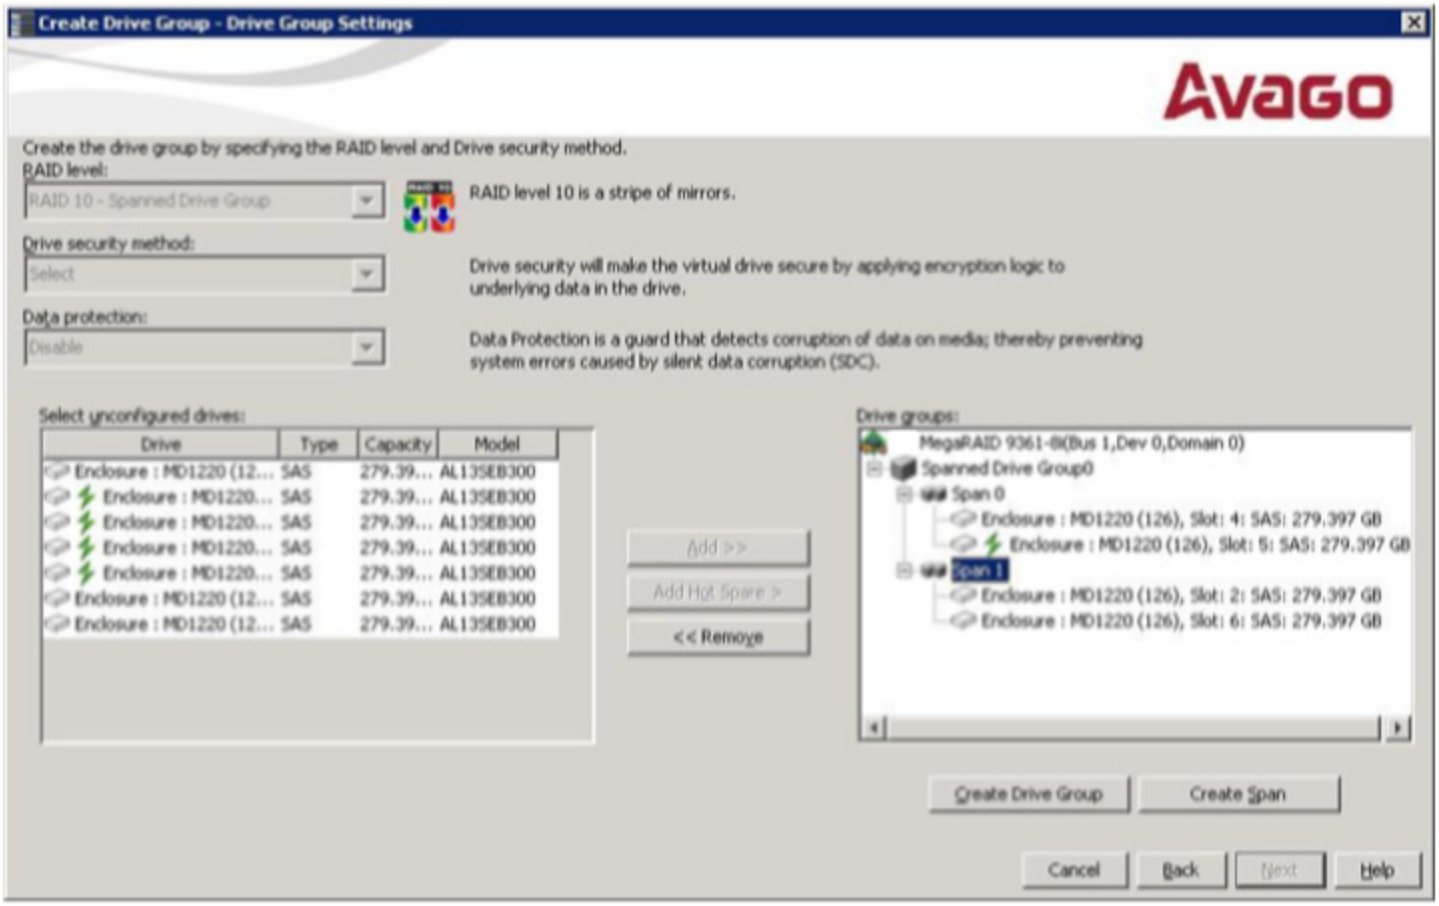

f. Click **Create Drive Group** to make a drive group with the spans.

g. Click **Next** to complete this step.

The **Create Virtual Drive - Virtual drive settings** window appears, as shown in the following figure. The drive group and the default virtual drive settings appear. The options to update the virtual drive or remove the virtual drive are grayed out until you create the virtual drive.

NOTE The parameters in the **Create Virtual Drive - Virtual drive settings** window display in Disabled mode (grayed out) for SAS-Integrated RAID (IR) controllers because these parameters do not apply to SAS-IR controllers.

Figure 201 Create Virtual Drive - Virtual Drive Settings Window

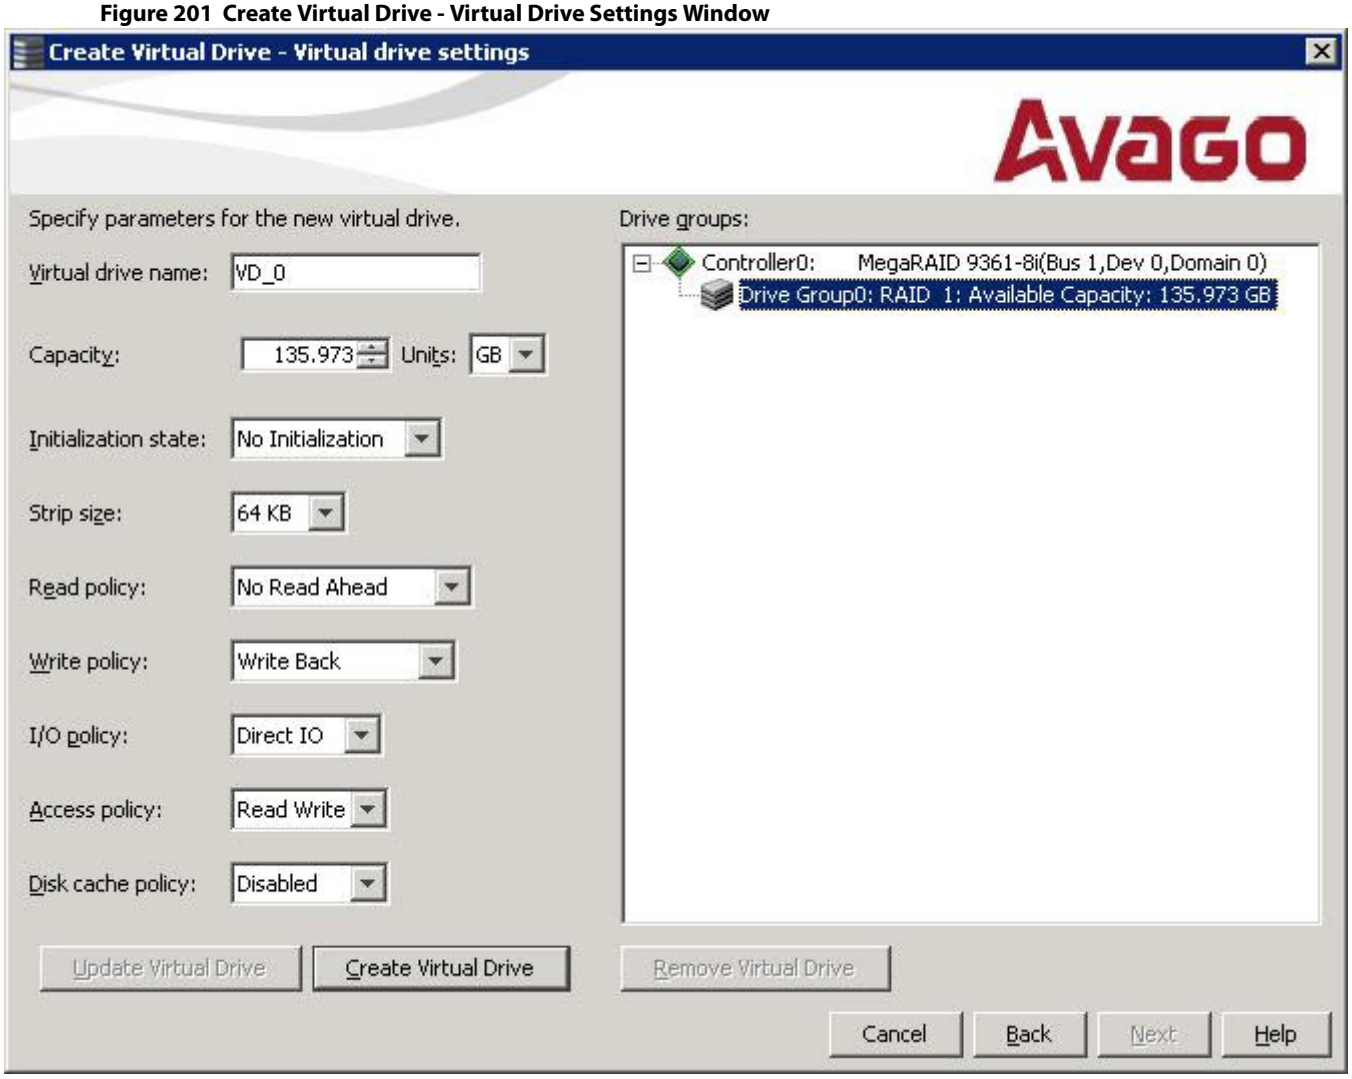

NOTE If you select **Write Back** as the write policy, and no battery exists, the battery is low or failed, or the battery is running through a re-learn cycle, the write policy switches to **Write Through**. This setting eliminates the risk of data loss in case of a power failure. A message window notifies you of this change.

NOTE If the controller supports High Availability DAS, the **Provide Shared Access** option appears in the above dialog. Select this option if you want the virtual drive to be shared between the two servers in a cluster.

4. Change any virtual drive settings, if desired.

   See Selecting Virtual Drive Settings, for more information about the virtual drive settings.

5. Click **Create Virtual Drive**.

   The new virtual drive appears under the drive group. The options **Update Virtual Drive** and **Remove Virtual Drive** are available. **Update Virtual Drive** allows you to change the virtual drive settings, and **Remove Virtual Drive** allows you to delete the virtual drive.

6. Click **Next**.

   The **Create Virtual Drive - Summary** window appears, as shown in the following figure. This window shows the selections you made for advanced configuration.

Figure 202 Create Virtual Drive - Summary Window

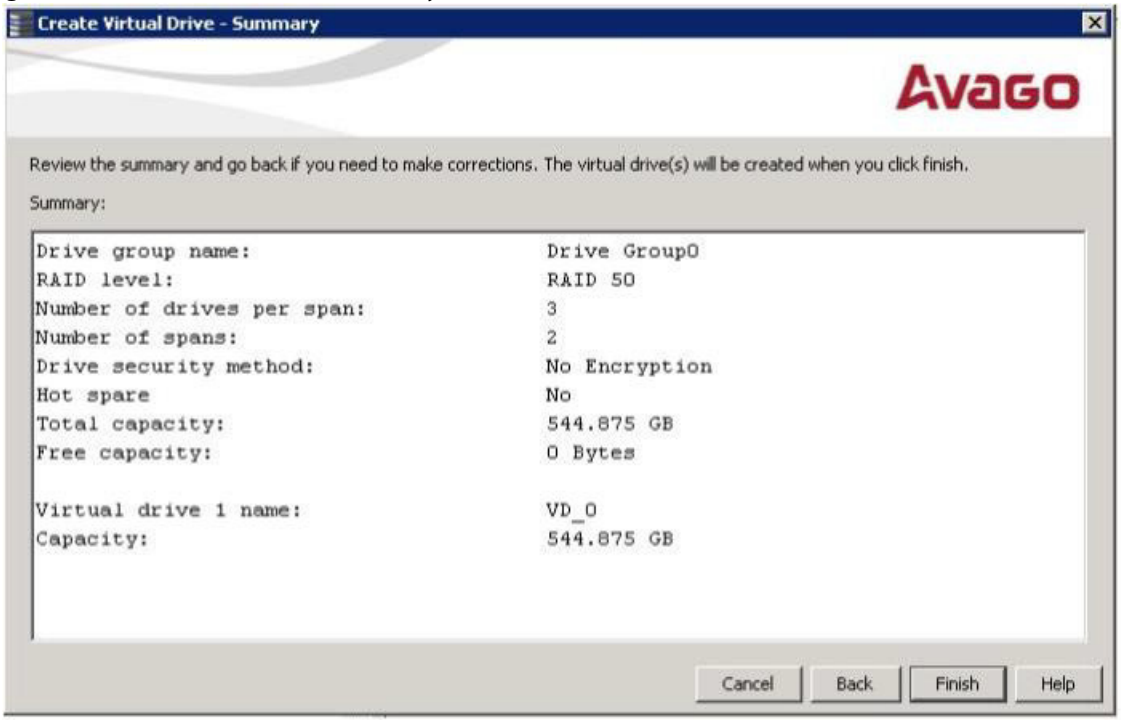

7. Click **Back** to return to the previous window to change any selections, or click **Finish** to accept and complete the configuration.

   After you click **Finish**, the new storage configuration is created and initialized according to the selected options.

   NOTE If you create a large configuration using drives that are in Power-Save mode, it can take several minutes to spin up the drives. A progress bar appears as the drives spin up. If any of the selected unconfigured drives fail to spin up, a dialog appears that identifies the drives.

   After the configuration is completed, a dialog notifies you that the virtual drives were created successfully.

8. Click **OK**.

   The **Enable SSD Caching on New Virtual Drives** dialog appears.

   The newly created virtual drive is enabled for SSD caching by default.

9. Click **OK** to confirm SSD caching on the virtual drive. Click **No** if you want to disable SSD caching on the virtual drive.

   The **All** check box is selected by default. To disable SSD caching on the virtual drives, deselect the **All** check box.

   If more drive capacity exists, the dialog asks whether you want to create more virtual drives. If no more drive capacity exists, you are prompted to close the configuration session.

10. Select either **Yes** or **No** to indicate whether you want to create additional virtual drives.

    If you select **Yes**, the system takes you to the Create Virtual Drive – Drive group and Virtual drive settings Dialog. If you select **No**, the utility asks whether you want to close the wizard.

11. If you selected **No** in the previous step, select either **Yes** or **No** to indicate whether you want to close the wizard.

    If you select **Yes**, the **Configuration** wizard closes. If you select **No**, the dialog closes, and you remain on the same page.

9.2 Converting JBOD Drives to Unconfigured Good

You can convert JBOD drives to Unconfigured Good using the **Create Virtual Drive** option or **Make Unconfigured Good** drive option with a single configuration.

NOTE MegaRAID SAS 9240-4i and MegaRAID SAS 9240-8i controllers support JBOD.

Perform the following steps to configure JBOD to Unconfigured Good drives:

1. Perform one of these actions:
   - Right-click the controller node in the device tree in the left frame of the **MegaRAID Storage Manager** window, and select **Create Virtual Drive**.
   - Select the controller node, and select **Go To > Controller > Create Virtual Drive**.

   The **Create Virtual Drive - JBOD to Unconfigured Good Conversion** wizard appears, as shown in the following figure.

Figure 203 Create Virtual Drive - JBOD to Unconfigured Good Conversion Dialog

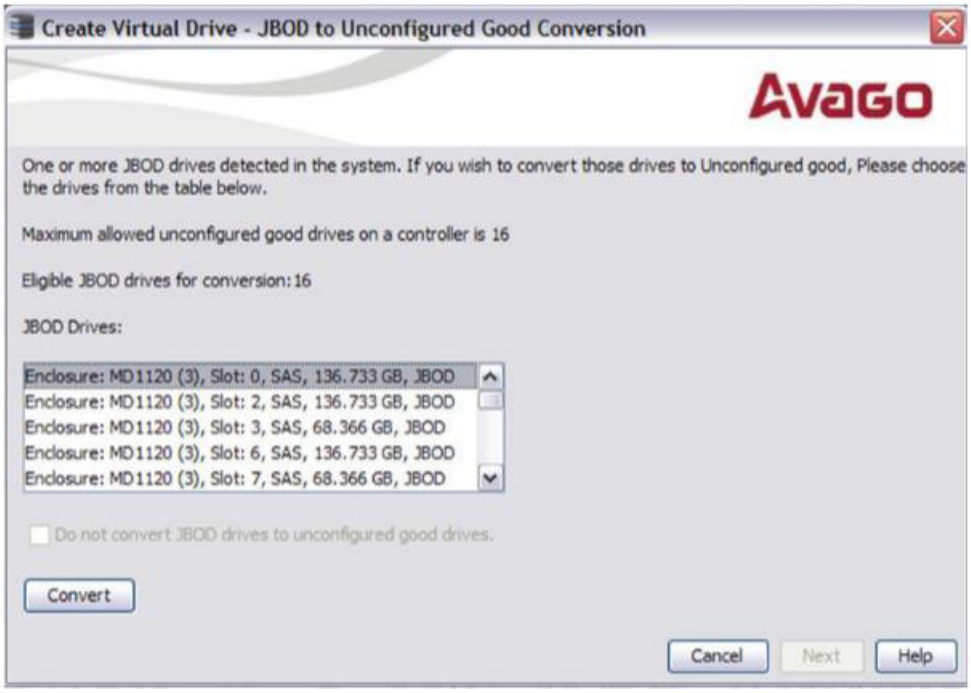

The **JBOD Drives** field displays the available JBOD drives available in the system.

NOTE If you do not want to make any JBOD as unconfigured good drives, select the **Do not convert JBOD drives to unconfigured good drives** check box, and the MegaRAID Storage Manager application skips changing any selected JBOD to unconfigured good drive.

2. Select the drives that you want configured as Unconfigured Good and then click **Convert**.

   A warning message appears stating that the JBOD drives will be removed and prompting for your confirmation.
3. Select **Confirm** and click **Yes** to proceed with the conversion.

NOTE If one or more JBOD drives have an OS or a file system installed on them, another warning message appears prior to conversion listing

those JBOD drives that have an OS or a file system installed on them. The message states that any attempt to convert the listed JBOD drives to unconfigured good drives would remove the existing data on the drives. Click **Yes** if you want to proceed with the conversion. Otherwise, click **No** to return to the previous screen and unselect those JBOD drives that have the OS installed on them.

4. Click **Next**.

   The **Create Virtual Drive - Drive group and Virtual drive settings** dialog appears.

9.2.1 Converting JBOD to Unconfigured Good from the MegaRAID Storage Manager Main Menu

You can also convert JBOD to Unconfigured Good by performing these steps:

1. Select **Controller >Make UnConfigured Good** from the main **MegaRAID Storage Manager** main menu.

   The **Make Configured Good** dialog appears, as shown in the following figure.

Figure 204 Make Configured Good Dialog

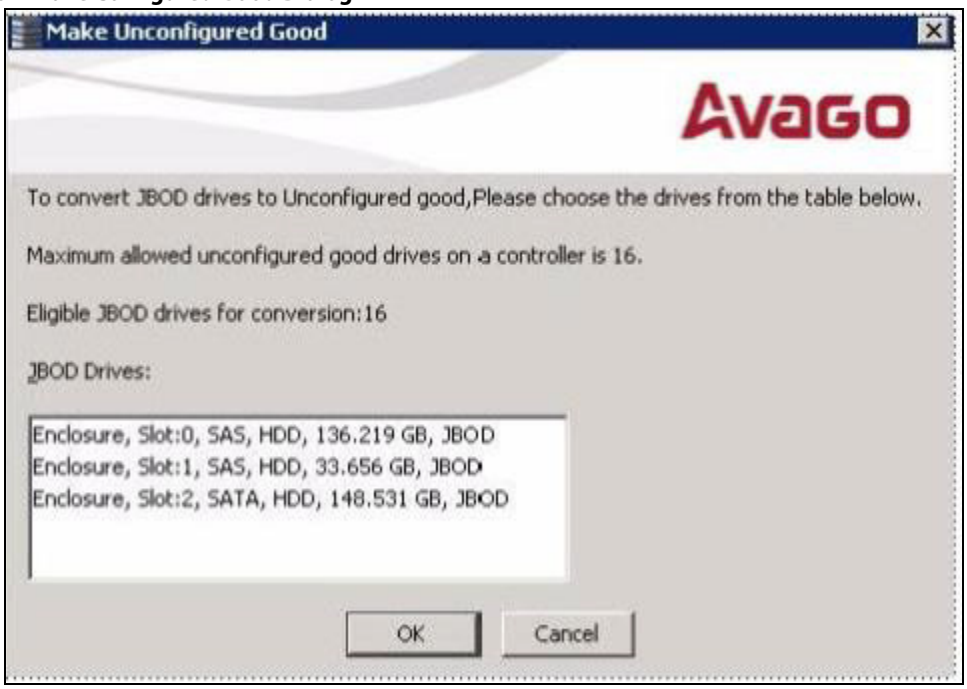

2. Select the JBOD drives to be configured as unconfigured good.
3. Click **OK**.

   A warning message appears stating that the JBOD drives will be removed and prompting for your confirmation.
4. Select **Confirm** and click **Yes** to proceed with the conversion.

   The selected JBOD drives are configured as unconfigured good.

NOTE If one or more JBOD drives have OS or file system installed on them, prior to conversion, another warning message appears listing those JBOD drives that have the OS or file system installed on them. The message states that any attempt to convert the listed JBOD drives to unconfigured good drives would remove the existing data on the drives. Click **Yes** if you want to proceed with the conversion. Else, click

No to return to the previous screen and unselect those JBOD drives that have the OS installed on them.

9.2.2 Removing a JBOD Drive

Follow these steps to remove a JBOD drive from the physical view:

1. Select the Physical View tab in the left panel of the MegaRAID Storage Manager window.
2. Right-click the physical Drive node, which is configured as a JBOD.
3. Select **Remove JBOD**.

   A warning message appears stating that the JBOD drive will be removed and prompting for your confirmation.
4. Select **Confirm** and click **Yes** to remove the JBOD drive.

   The selected JBOD drive is removed.

NOTE If the JBOD drive has an OS or a file system installed on it, another warning message appears prior to removal stating that any attempt to remove the JBOD drive would remove the existing data on the drive. Select **Confirm**, and click **Yes** if you want to proceed with the removal. Otherwise, click **No** to return to the previous screen.

9.3 Enabling Security on JBOD

The MegaRAID Storage Manager software supports Enable Security feature for security capable JBOD drives.

Follow these steps to enable security on security capable JBOD drives:

1. In the left panel of the **MegaRAID Storage Manager** window, select the Security capable JBOD drive for which you want to enable security.
2. In the **MegaRAID Storage Manager** window, select **Go To** > **Physical Drive** > **Enable Security** as shown in the following figure.

Figure 205 MegaRAID Storage Manager Window

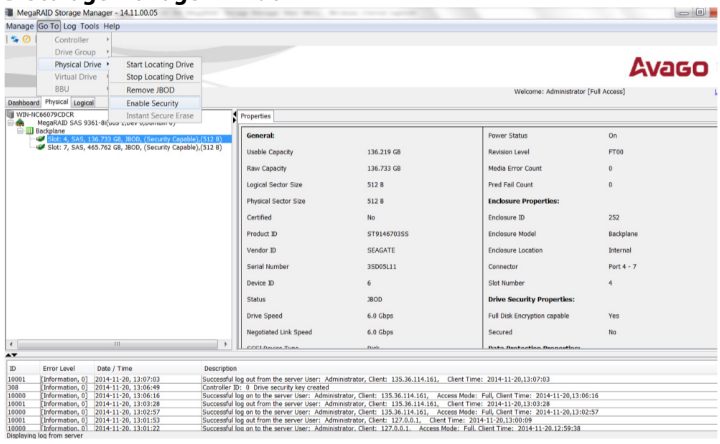

3. You can also make a JBOD secure as follows.
   a. In the **MegaRAID Storage Manager** window, right click on **Controller** >**Make JBOD Secure**.

      The **Make JBOD Secure** dialog appears, as shown in the following figure.

Figure 206 Make JBOD Secure Window

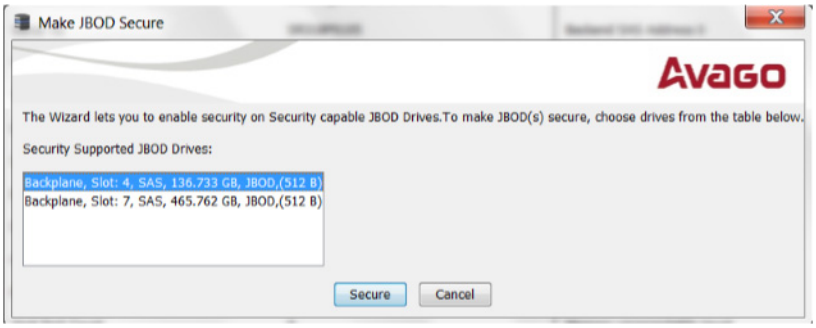

b. Select the JBOD drive for which you need to enable security.

c. Click **Secure**.

9.4 Creating 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 create a dedicated or global hot spare drive, follow these steps:

1. Select the **Physical** tab in the left panel of the MegaRAID Storage Manager main menu, and click the icon of an unused drive.

   For each drive, the window displays the port number, enclosure number, slot number, drive state, drive capacity, and drive manufacturer.
2. Either select **Go To > Physical Drive > Assign Global Hot Spare**, or select **Go To > Physical Drive > Assign Dedicated Hot Spare**.
3. Perform one of these actions:
   - If you selected **Assign Dedicated Hotspare**, select a drive group from the list that appears. The hot spare is dedicated to the drive group that you select.

     NOTE If the controller supports High Availability DAS, dedicated hot spares can be assigned to only one drive group. If you try to assign dedicated hot spares to more than one drive group, an error message appears.
   - If you selected **Assign Global Hotspare**, skip this step, and go to the next step. The hot spare is available to any virtual drive on a specific controller.
4. Click **Go** to create the hot spare.

   The drive state for the drive changes to dedicated or global hot spare, depending on your selection.

9.5 Changing Adjustable Task Rates

If you want to change the Rebuild rate and other task rates for a controller, you must first log onto the server in Full Access mode.

NOTE Leave the adjustable task rates at their default settings to achieve the best system performance. If you raise the task rates above the defaults, foreground tasks will run more slowly and it might seem that the

system is not responding. If you lower the task rates below the defaults, rebuilds and other background tasks might run very slowly and might not complete within a reasonable time. If you decide to change the values, record the original default values here so you can restore them later, if necessary: **Rebuild rate:** _____, **Background Initialization (BGI) rate:** _____, **Check consistency rate:** _____.

To change the adjustable task rates, perform the following steps:

1. Select a controller icon in the **Physical** tab or the **Logical** tab in the left panel of the **MegaRAID Storage Manager** window.
2. Select **Go To > Controller > Set Adjustable Task Rates** from the menu bar.

   The **Set Adjustable Task Rates** window appears, as shown in the following figure.

Figure 207 Set Adjustable Task Rates Menu

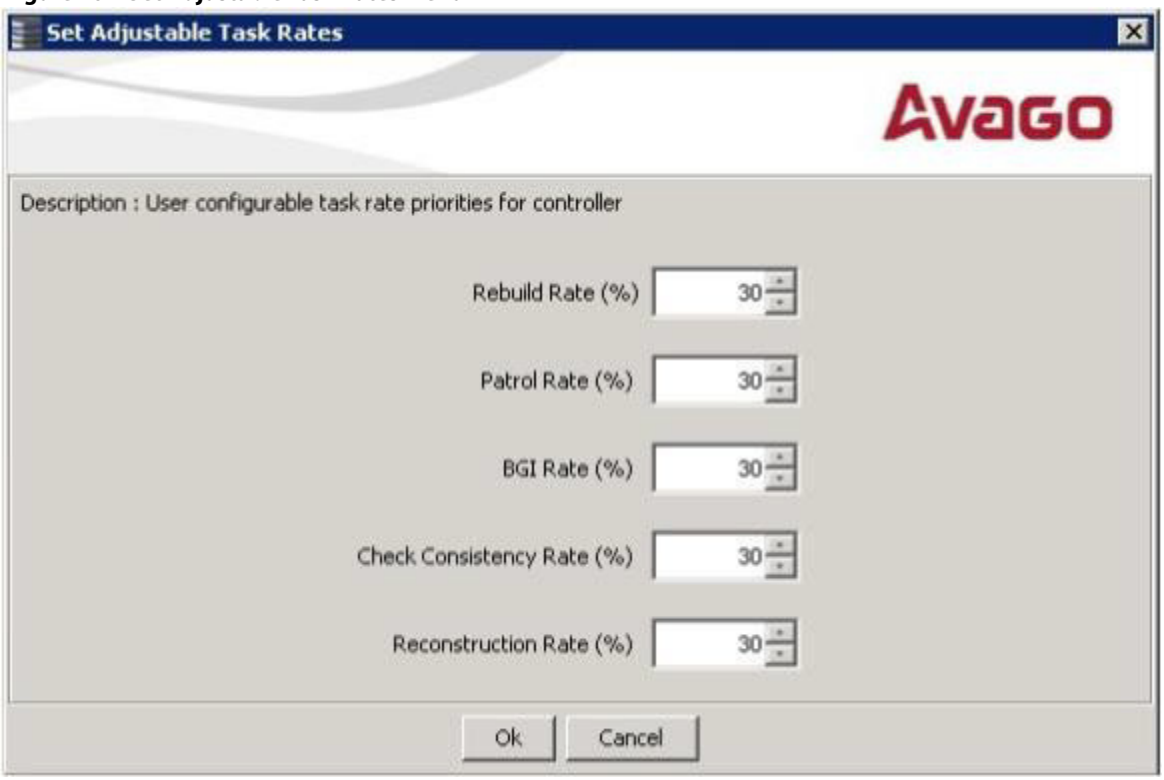

3. Enter changes, as needed, to the following task rates:
   - **Rebuild Rate**. Enter a number from 0 to 100 to control the rate at which a rebuild will be performed on a drive when one is necessary. The higher the number, the faster the rebuild will occur (and the system I/O rate may be slower as a result).
   - **Patrol Rate**. Enter a number from 0 to 100 to control the rate at which patrol reads will be performed. Patrol read monitors drives to find and resolve potential problems that might cause drive failure. The higher the number, the faster the patrol read will occur (and the system I/O rate may be slower as a result).
   - **Background Initialization (BGI) Rate**. Enter a number from 0 to 100 to control the rate at which virtual drives are initialized "in the background." Background initialization establishes mirroring or parity for a RAID virtual drive while allowing full host access to the virtual drive. The higher the number, the faster the initialization will occur (and the system I/O rate may be slower as a result).
   - **Check Consistency Rate**. Enter a number from 0 to 100 to control the rate at which a consistency check is done. A consistency check scans the consistency data on a fault tolerant virtual drive to determine if the data

has become corrupted. The higher the number, the faster the consistency check is performed (and the system I/O rate may be slower as a result).

— **Reconstruction Rate**. Enter a number from 0 to 100 to control the rate at which reconstruction of a virtual drive occurs. The higher the number, the faster the reconstruction occurs (and the system I/O rate may be slower as a result).

4. Click **Ok** to accept the new task rates.
5. When the warning message appears, click **OK** to confirm that you want to change the task rates.

9.6 Changing Power Settings

The RAID controller includes Dimmer Switch® technology that conserves energy by placing certain unused drives into Power-Save mode. In Power-Save mode, the drives use less energy, and the fan and the enclosure require less energy to cool and house the drives, respectively. Also, this technology helps avoid application timeouts caused by spin-up delays and drive wear caused by excessive spin-up/down cycles.

You can use the **Power Settings** field in the MegaRAID Storage Manager software to choose whether to allow unconfigured drives or Commissioned Hotspares to enter Power-Save mode.

NOTE The Dimmer Switch technology is enabled by default.

When they are in the Power-Save mode, unconfigured drives and drives configured as Commissioned Hotspares (dedicated or global) can be spun down. When spun down, the drives stay in Power-Save mode except for periodic maintenance, which includes the following:

- Periodic background media scans (Patrol Read) to find and correct media defects to avoid losing data redundancy (hot spare drives only)
- Use of a Commissioned Hotspare to rebuild a degraded drive group (Commissioned Hotspare drives only)
- Update of disk data format (DDF) and other metadata when you make changes to RAID configurations (Commissioned Hotspare drives and unconfigured drives)

NOTE If your controller does not support this option, the **Power Settings** field does not appear.

Follow these steps to change the power-save setting.

1. Select a controller icon in the **Physical** tab or the **Logical** tab in the left panel of the **MegaRAID Storage Manager** window.
2. Select **Go To > Controller > Manage Power Settings** from the menu bar.

   The **Manage Power Save Settings** dialog appears.

Figure 208 Manage Power Save Settings

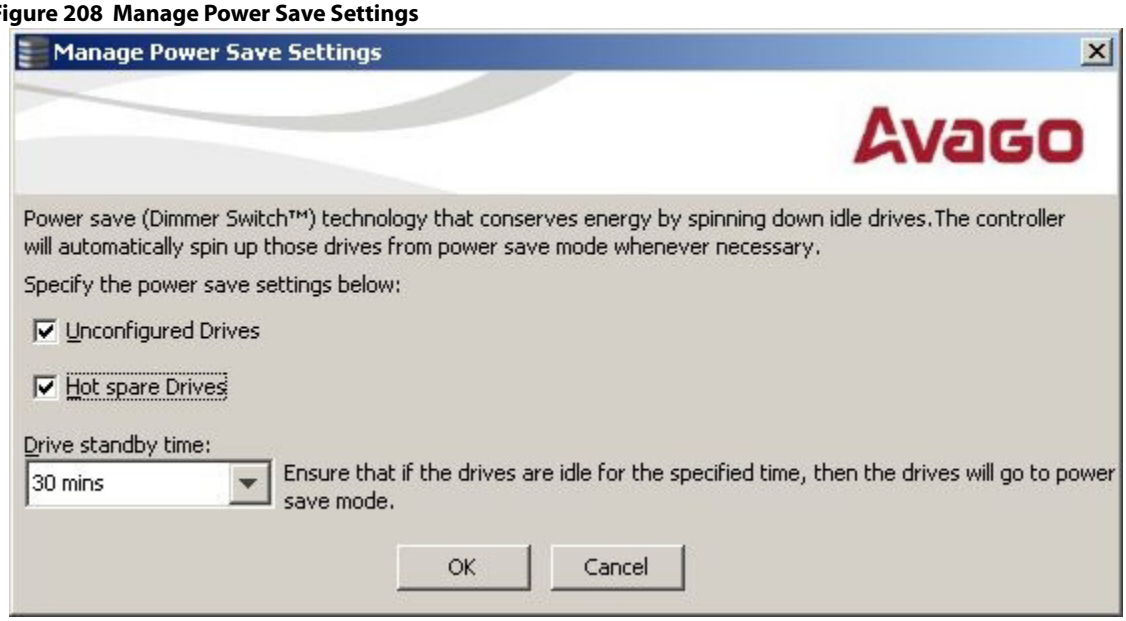

3. Select the **Unconfigured Drives** check box to let the controller enable the unconfigured drives to enter the Power-Save mode.
4. Select the **Hot spare Drives** check box to let the controller enable the Hot spare drives to enter the Power-Save mode.
5. Select the drive standby time (Alt+D) using the drop-down list from the **Drive standby time** field.

   NOTE The **Drive Standby time** drop-down list is enabled only if any of the check boxes above it are checked. The drive standby time can be 30 minutes, 1 hour, 1.30 hours, or 2 hours through 24 hours.

6. Click **OK**.

   The Power-Save settings are saved. After you click **OK**, a confirmation dialog appears prompting you to save your changes.

   If you do not specify the Power-Save settings in the **Manage Power Save Settings** dialog, a confirmation dialog appears. The confirmation dialog mentions that the system does not have power savings for any of the drives, and asks if you would like to proceed.

9.7 Recovering and Clearing Punctured Block Entries

You can recover and clear the punctured block area of a virtual drive.

ATTENTION This operation removes any data stored on the physical drives. Back up the good data on the drives before making any changes to the configuration.

When a Patrol Read or a Rebuild operation encounters a media error on the source drive, it punctures a block on the target drive to prevent the use of the data with the invalid parity. Any subsequent read operation to the punctured block completes but with an error. Consequently, the puncturing of a block prevents any invalid parity generation later while using this block.

To recover or clear the punctured block area of a virtual drive, run a Slow (or Full) Initialization to zero out and regenerate new parity causing all bad block entries to be removed from the bad block table.

To run a Slow (or Full) Initialization, see Selecting Virtual Drive Settings.

9.8 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.

ATTENTION Do not enable drive caching on a mirrored drive group (RAID 1 or RAID 1E). If you do, data can be corrupted or lost in the event of a sudden power loss. A warning appears if you try to enable drive caching for a mirrored drive group.

NOTE For virtual drives with SAS drives only, set the drive write cache policy set to **Disabled**, by default. For virtual drives with SATA drives only, set the drive write cache policy to **Enabled**, by default.

To change the virtual drive properties, perform the following steps:

1. Select a virtual drive icon in the **Physical** tab or the **Logical** tab in the left panel of the **MegaRAID Storage Manager** window.
2. Select **Go To > Virtual Drive > Set Virtual Drive Properties** from the menu bar.

   The **Set Virtual Drive Properties** dialog appears, as shown in the following figure.

Figure 209 Set Virtual Drive Properties Dialog

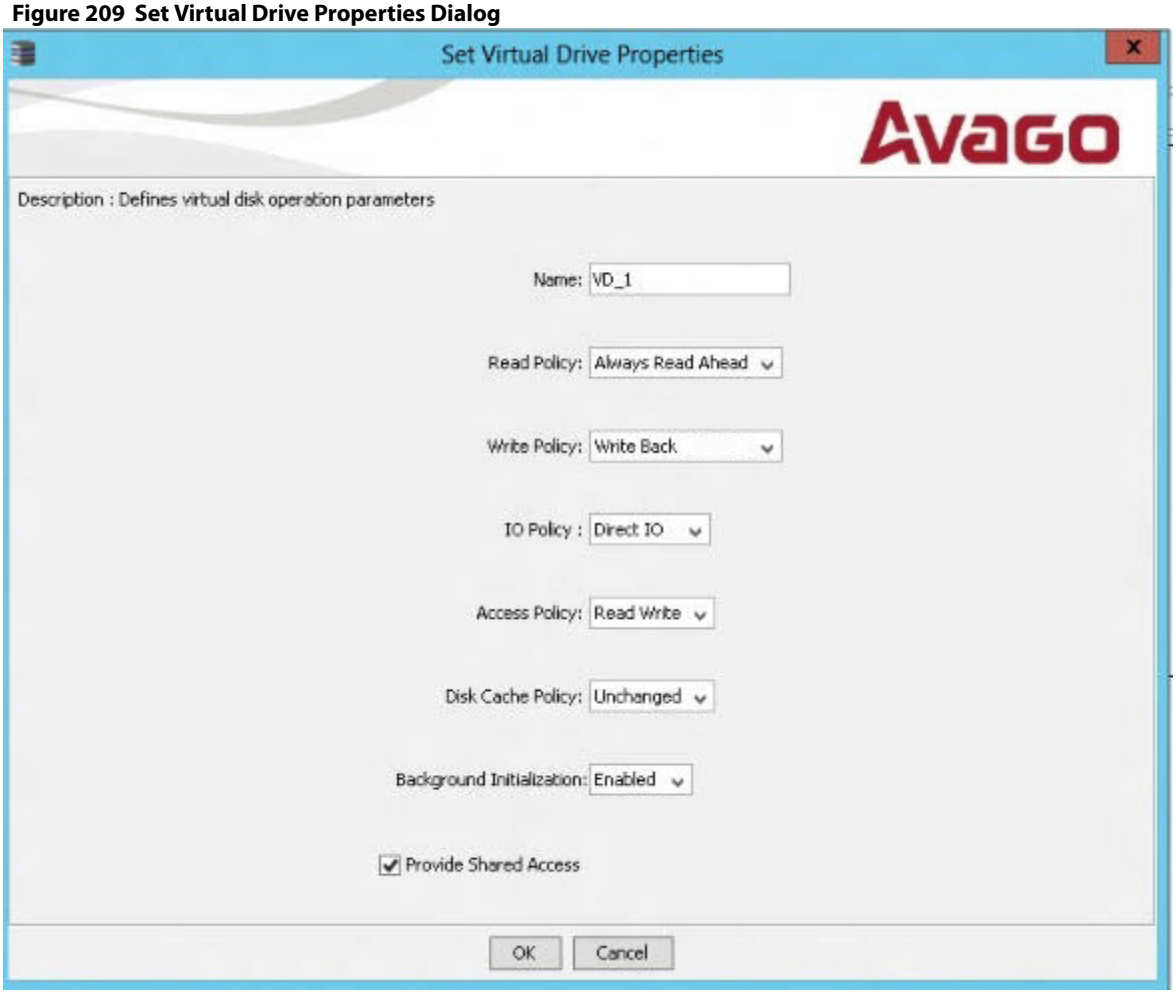

NOTE If the controller supports High Availability DAS, the **Provide Shared Access** check box appears in the above dialog. Select this option if you want the virtual drive to be shared between the two servers in a cluster.

3. Change the virtual drive properties as needed.
   For information about these properties, see Selecting Virtual Drive Settings.
4. Click **OK** to accept the changes.
   The virtual drive settings are updated.

9.9 Changing a Virtual Drive Configuration

You can use the **Modify Drive Group** wizard in the MegaRAID Storage Manager software to change the configuration of a virtual drive by adding drives to the virtual drive, removing drives from it, or changing its RAID level.

ATTENTION Be sure to back up the data on the virtual drive before you change its configuration.

NOTE You cannot change the configuration of a RAID 10, RAID 50, or RAID 60 virtual drive. You cannot change a RAID 0, RAID 1, RAID 5, or RAID 6 configuration if two or more virtual drives are defined on a single drive group. (The Logical tab shows which drive groups and drives are used by each virtual drive.) Physical drives which are in **Prepare for Removal** state are not listed for virtual drive migration

9.9.1 Accessing the Modify Drive Group Wizard

NOTE The **Modify Drive Group** wizard was previously known as the **Reconstruction** wizard.

Perform the following steps to access the **Modify Drive Group** wizard options:

1. Click the **Logical** tab in the left panel of the MegaRAID Storage Manager main menu window.
2. Select a drive group in the left panel of the window.
3. Select **Go To** > **Drive Group** > **Modify Drive Group** on the menu bar, or right-click the virtual drive icon to access the **Modify Drive Group** wizard.

   The following warning appears about rebooting virtual drives containing boot partitions that are undergoing RAID level migration or capacity expansion operations. Back up your data before you proceed.

Figure 210 Reboot Warning Message

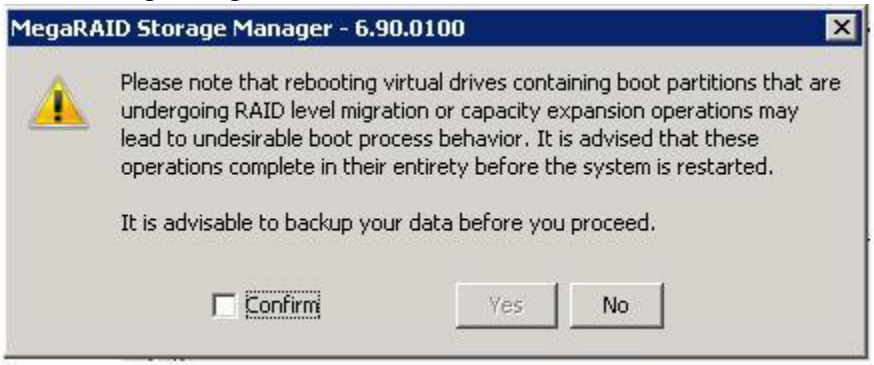

4. Select the **Confirm** check box, and click **Yes**.

   A warning to back up your data appears, as shown in the following figure.

Figure 211 Warning to Back Up Data Message

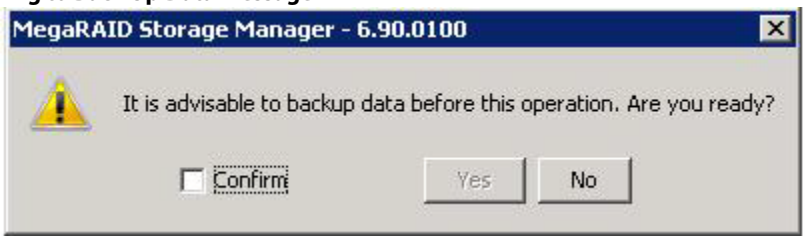

5. Select the **Confirm** check box, and click **Yes**.

   The **Modify Drive Group** wizard window appears, as shown in the following figure.

Figure 212 Modify Drive Group Wizard Window

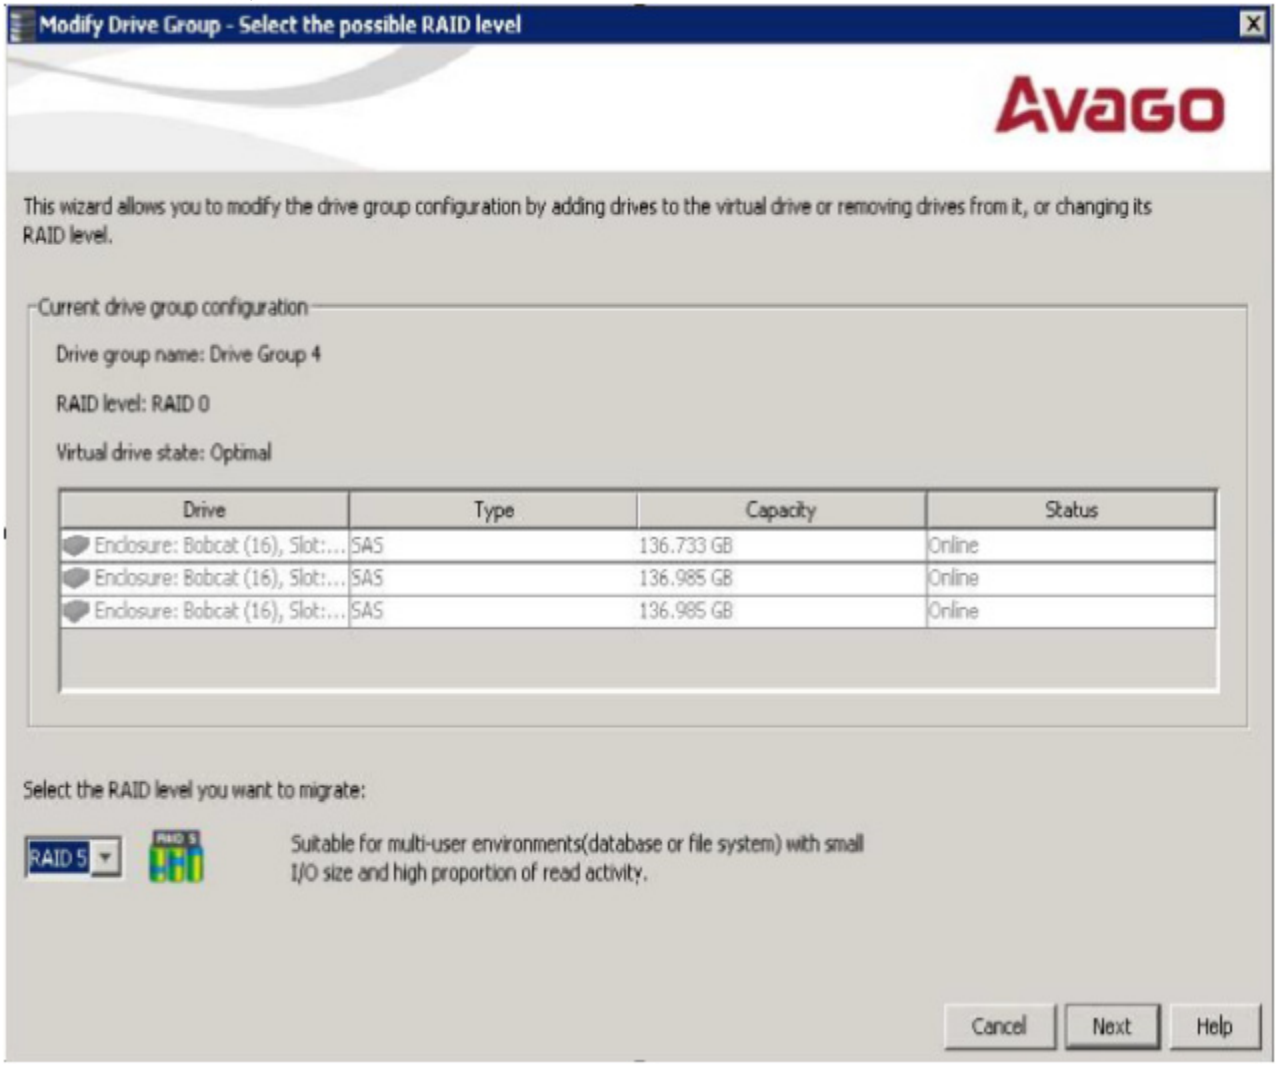

The following sections explain the **Modify Drive Group** wizard options.

9.9.2 Adding a Drive or Drives to a Configuration

ATTENTION Be sure to back up the data on the virtual drive before you add a drive to it.

Follow these steps to add a drive or drives to a configuration with the **Modify Drive Group** wizard.

1. Click the **Logical** tab in the left panel of the MegaRAID Storage Manager window.
2. Select a drive group in the left panel of the window.
3. Either select **Go To > Drive Group > Modify Drive Group** on the menu bar, or right-click the virtual drive icon to access the **Modify Drive Group** wizard.

   The **Modify Drive Group** wizard window appears.

Figure 213 Modify Drive Group Wizard Window

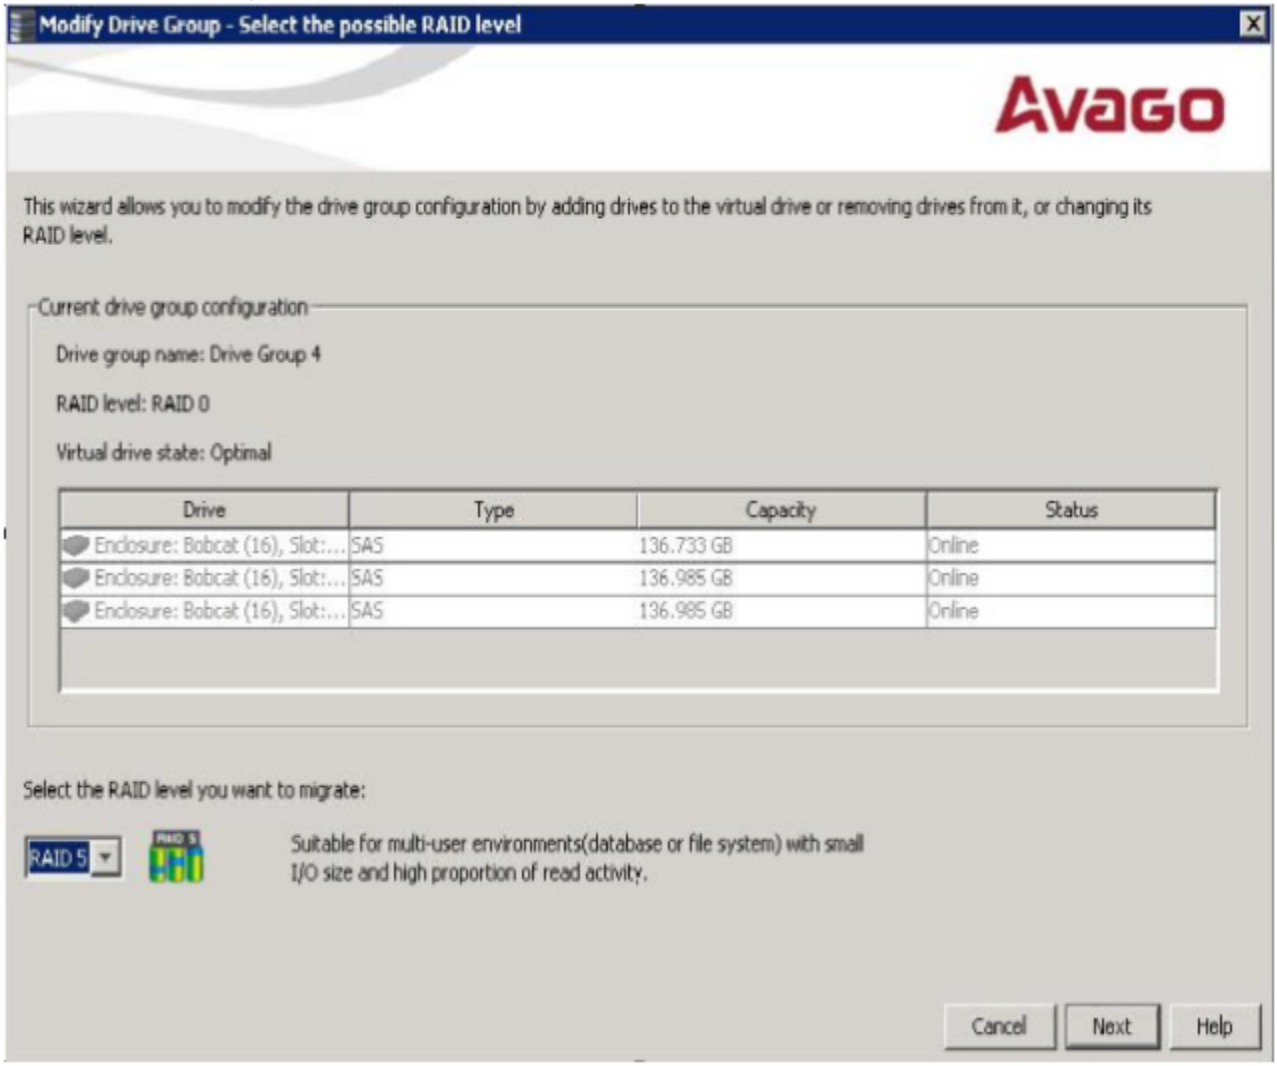

4. Select the RAID level to which you want to change ("migrate") the drive group, and click **Next**.

   The following window appears. It lists the drives you can add, and it states whether you have to add a minimum number of drives to change the RAID level from the current level to the new RAID level.

Figure 214 Modify Drive Group – Add Drives to the Current Configuration Window

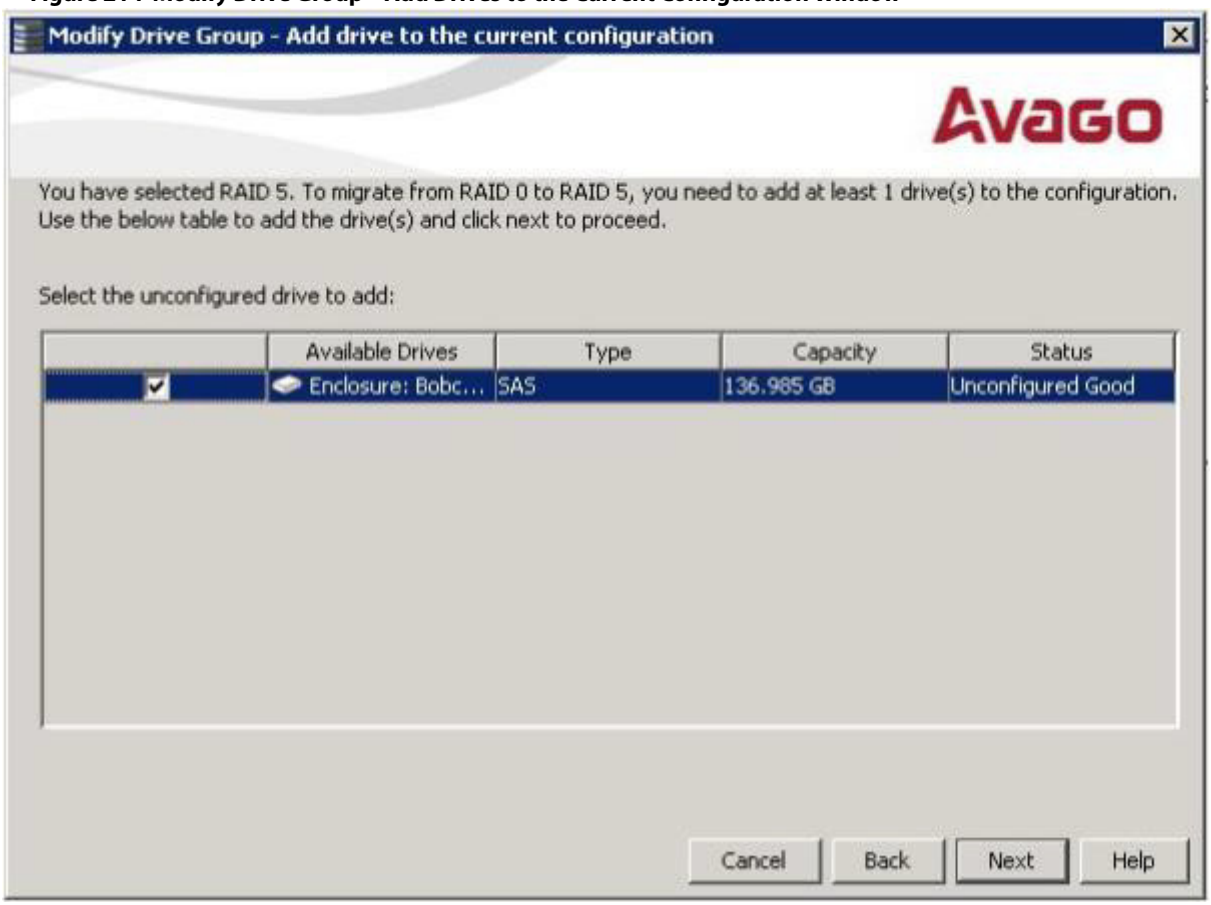

5. Click the check box next to any unconfigured drives that you want to add, and then click **Next**.

NOTE The drives you add must have the same capacity as or greater capacity than the drives already in the drive group, or you cannot change the RAID level.

The **Modify Drive Group - Summary** window appears. This window shows the current settings and what the settings will be after the drives are added.

Figure 215 Modify Drive Group - Summary Window

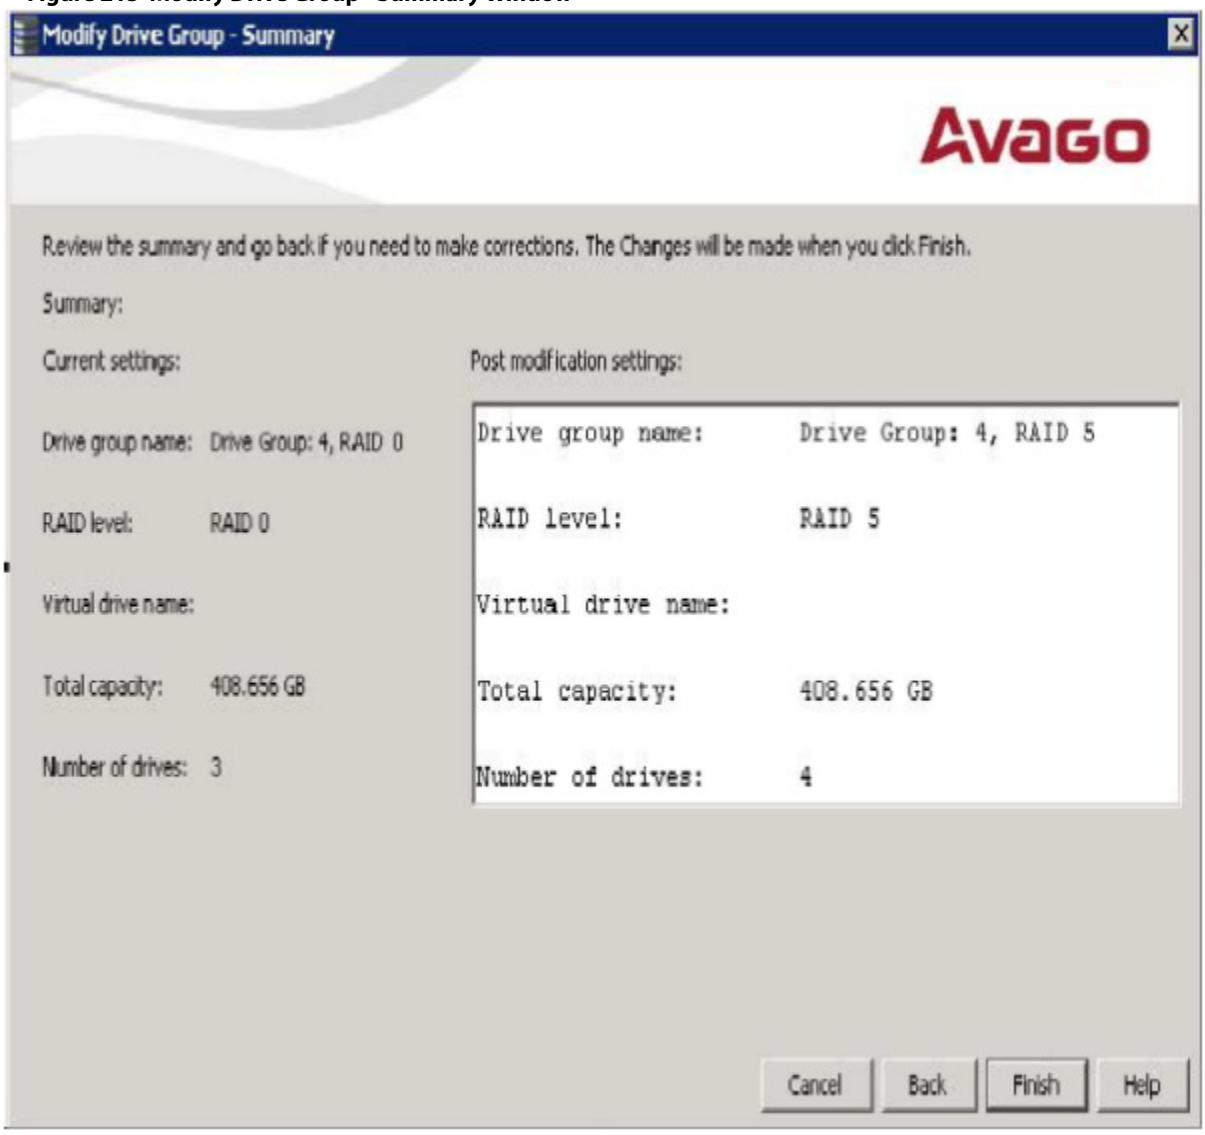

6. Review the configuration information.

   You can click **Back** if you need to change any selections.
7. Click **Finish** to accept the changes.

   A confirmation message appears. The message states that this operation cannot be aborted and asks whether you want to continue.
8. Click **Yes** to accept and complete the addition of the drives to the drive group.

9.9.3 Removing a Drive from a Configuration

ATTENTION Be sure to back up the data on the virtual drive before you remove a drive from it.

Follow these steps to remove a drive from a RAID 1, RAID 5, or RAID 6 configuration.

NOTE This option is not available for RAID 0 configurations.

1. Click the **Logical** tab in the left panel of the MegaRAID Storage Manager window.

2. Click a drive icon in the left panel of the window.
3. Either select **Go To > Physical Drive > Make Drive Offline** on the menu bar, or right-click the drive, and select **Make Drive Offline** from the menu.

   A confirmation message appears. The message states that this operation cannot be aborted and asks whether you want to continue.
4. Click **Yes** to accept and complete the removal of the drive from the drive group.

9.9.4 Replacing a Drive

ATTENTION Be sure to back up the data on the virtual drive before you replace a drive.

Follow these steps to add a replacement drive and copy the data from the drive that was removed to the replacement drive.

1. Click the **Logical** tab in the left panel of the MegaRAID Storage Manager window.
2. Select a drive in the left panel of the window.
3. Either select **Go To > Physical Drive > Replace Physical Drive** on the menu bar, or right-click the virtual drive icon to access the **Modify Drive Group** wizard.

   The dialog with the replacement drive appears, as shown in the following figure.

Figure 216 Drive Replacement Window

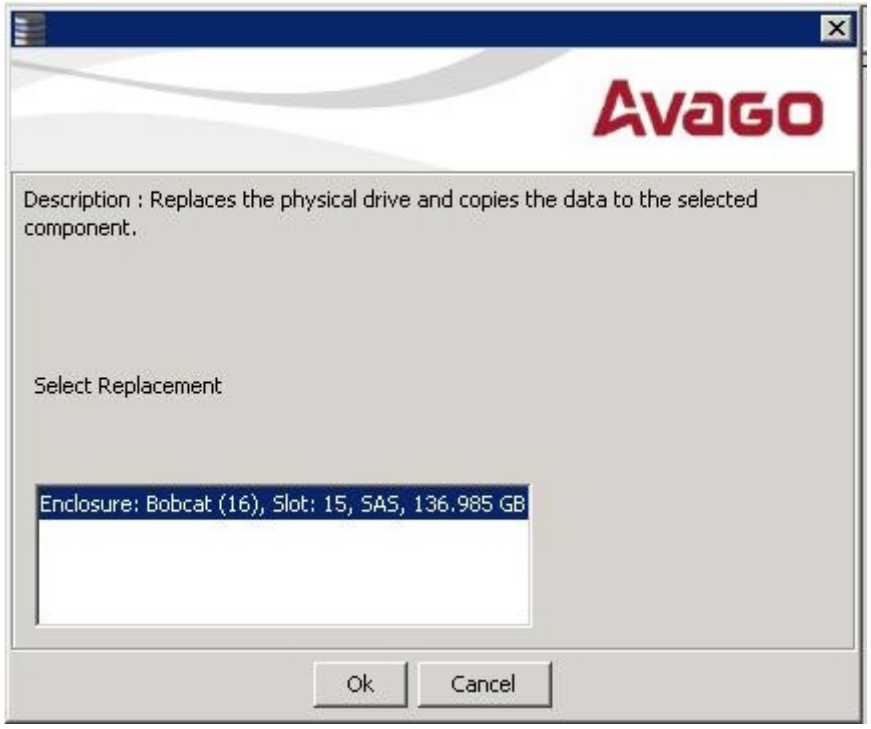

4. Select a replacement drive.

   A confirmation message appears.
5. Click **Yes**.

   This step replaces a drive and copies the data to the selected component.

9.9.5 Migrating the RAID Level of a Virtual Drive

As the amount of data and the number of drives in your system increase, you can use RAID-level migration to change a virtual drive from one RAID level to another. You do not have to power down or reboot the system when you make this change.

When you migrate a virtual drive to another RAID level, you can keep the same number of drives, or you can add drives. In some cases, you have to add a certain number of drives to migrate the virtual drive from one RAID level to another. The window indicates the minimum number of drives you are required to add.

ATTENTION Be sure to back up the data on the virtual drive before you change the RAID level.

Follow these steps to change the RAID level of the virtual drive with the **Modify Drive Group** wizard:

1. Click the **Logical** tab in the left panel of the MegaRAID Storage Manager window.
2. Select a drive group in the left panel of the window.
3. Either select **Go To > Drive Group > Modify Drive Group** on the menu bar, or right-click the virtual drive icon to access the **Modify Drive Group** wizard.

   The **Modify Drive Group** wizard appears.

Figure 217 Modify Drive Group Wizard Dialog

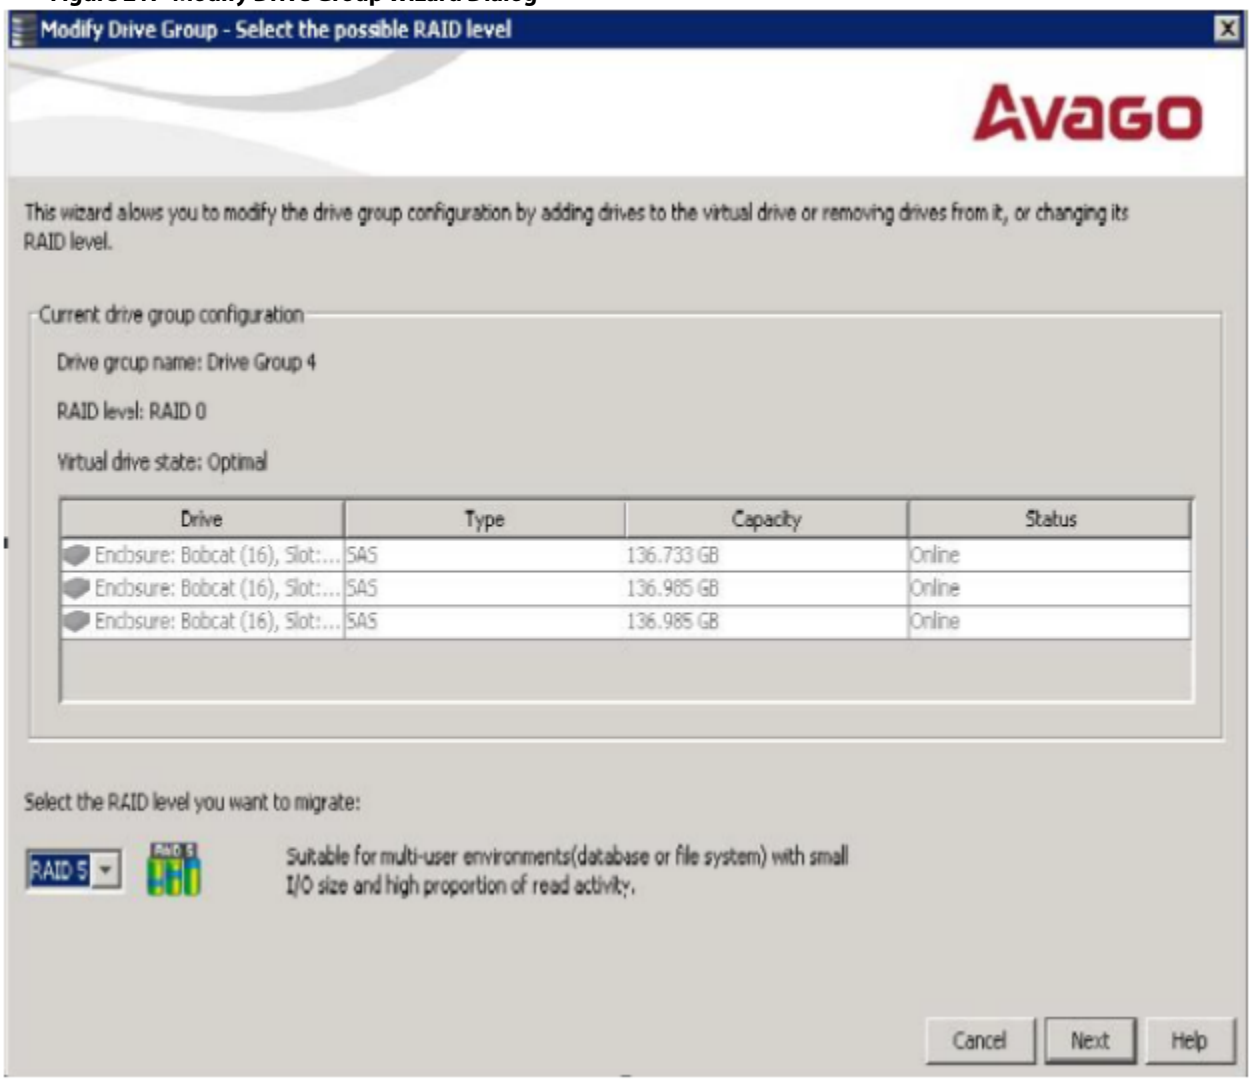

4. On the **Modify Drive Group Wizard** dialog, select the RAID level to which you want to change ("migrate") the drive group to, and click **Next**.

   The following dialog appears. The dialog states the number of drives that you have to add to change the RAID level from the current level to a new RAID level that requires more drives.

Figure 218 Modify Drive Group - Add drive to the current configuration Screen

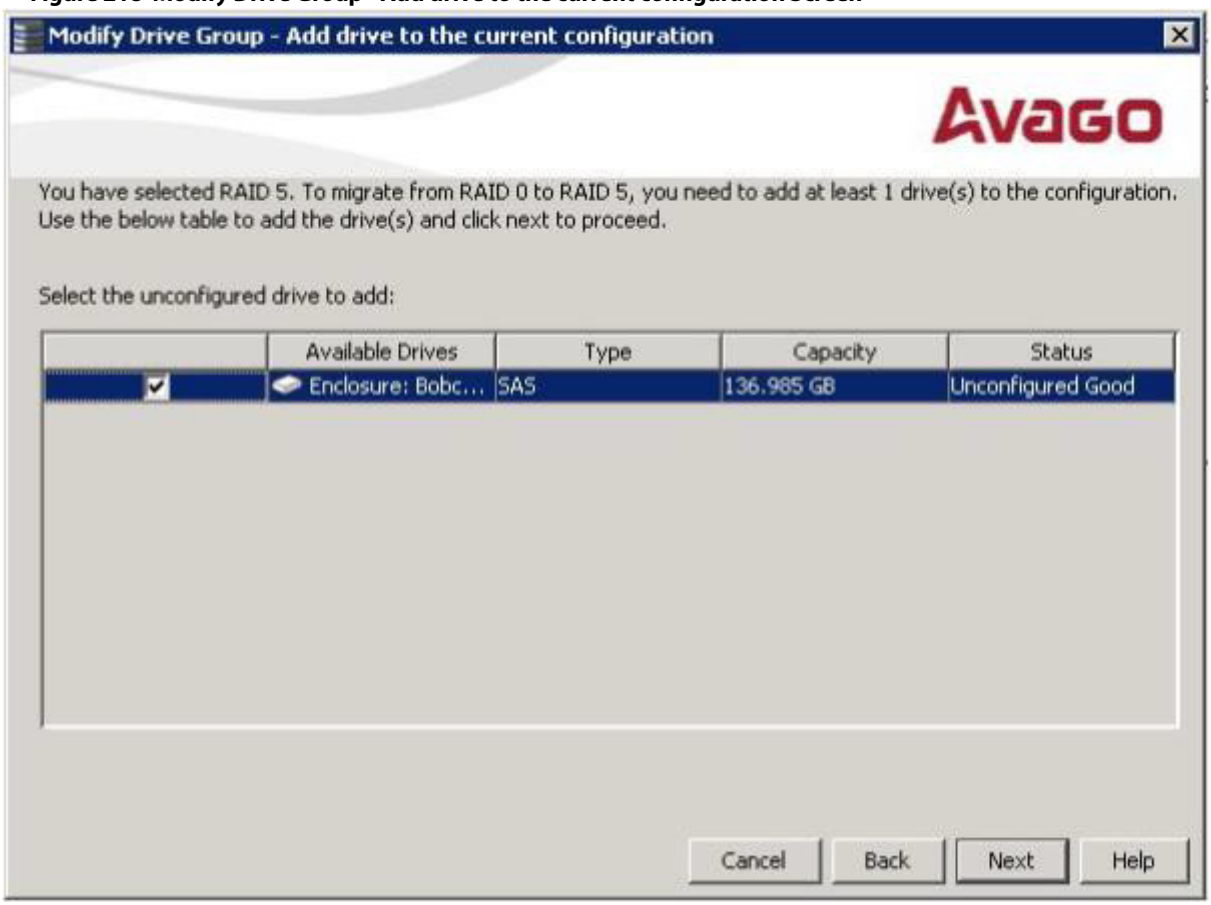

5. Select the unconfigured drive or drives to add, and click **Next**.

NOTE The drives you add must have the same capacity as or greater capacity than the drives already in the drive group, or you cannot change the RAID level.

The **Modify Drive Group – Summary** window appears. This window shows the current settings and what the settings will be after the drives are added.

Figure 219 Modify Drive Group - Summary Screen

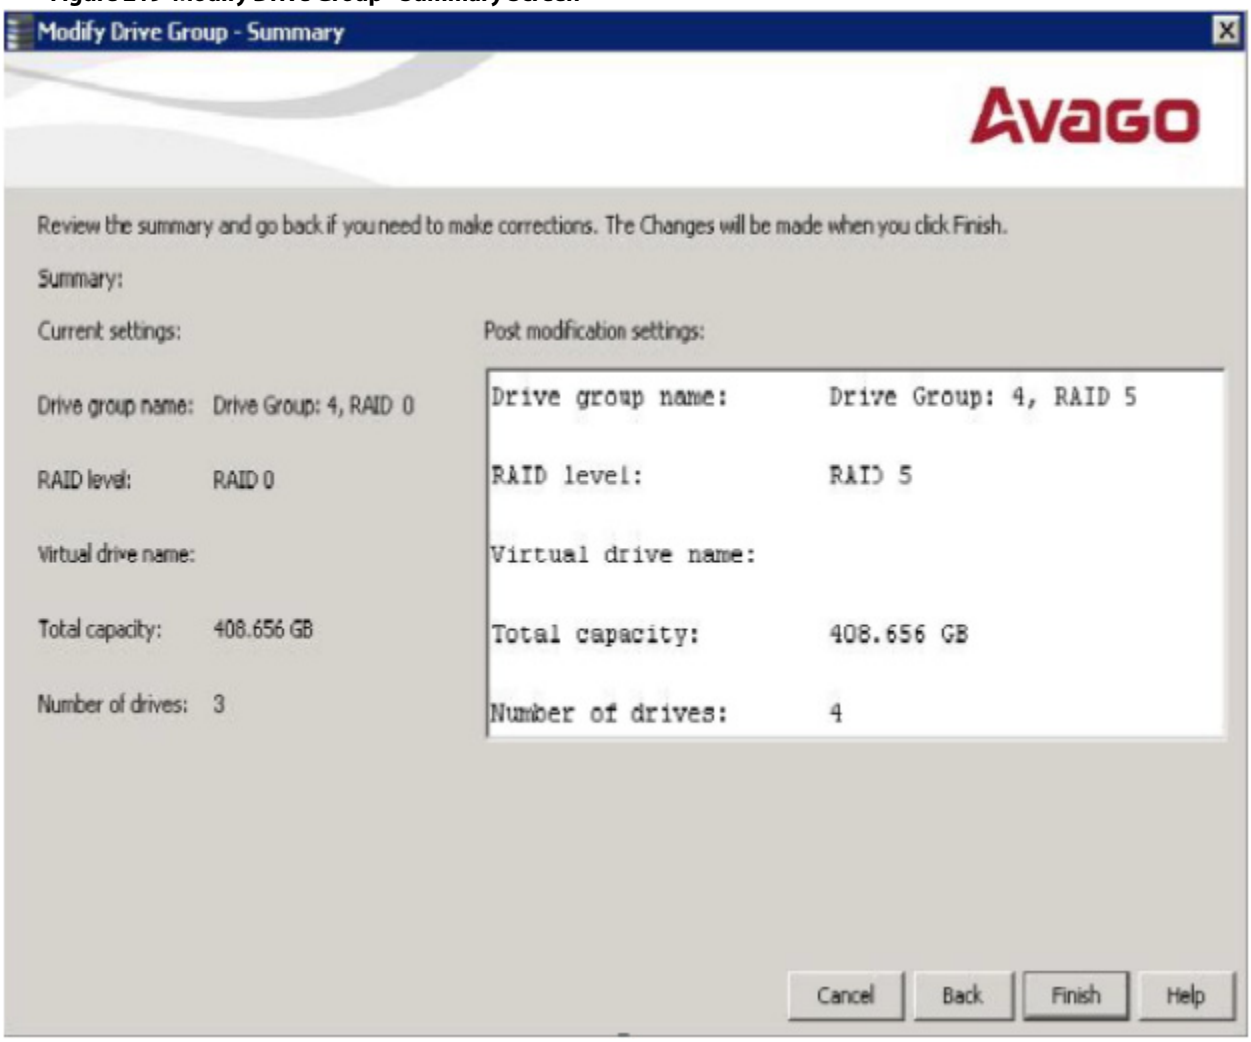

6. Review the configuration information.

   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.

9.10 Deleting a Virtual Drive

ATTENTION Make sure to back up the data that is on the virtual drive before you delete it. Make sure that the operating system is not installed on this virtual drive.

You can delete virtual drives to rearrange the storage space. To delete a virtual drive, follow these steps.

1. Back up all user data that is on the virtual drive you want to delete.

2. On the **MegaRAID Storage Manager** window, select the **Logical** tab, and click the icon of the virtual drive you want to delete.
3. Select **Go To > Virtual Drive > Delete Virtual Drive**.
4. When the warning messages appear, click **Yes** to confirm that you want to delete the virtual drive.

NOTE You are asked twice if you want to delete a virtual disk to avoid deleting the virtual disk by mistake.

9.11 Performing a Join Mirror Operation

You can perform a join mirror operation on a drive group to continue using the modified virtual drive or to reuse the original virtual drive.

Follow these steps to perform a join mirror operation:

1. Go to the Logical tab in the MegaRAID Storage Manager window.
2. Right click on the drive group on which you want to perform the join mirror operation and select **Join Mirror**.

   The **Join Mirror - Choose Option** dialog appears.

Figure 220 Join Mirror - Choose Option Dialog

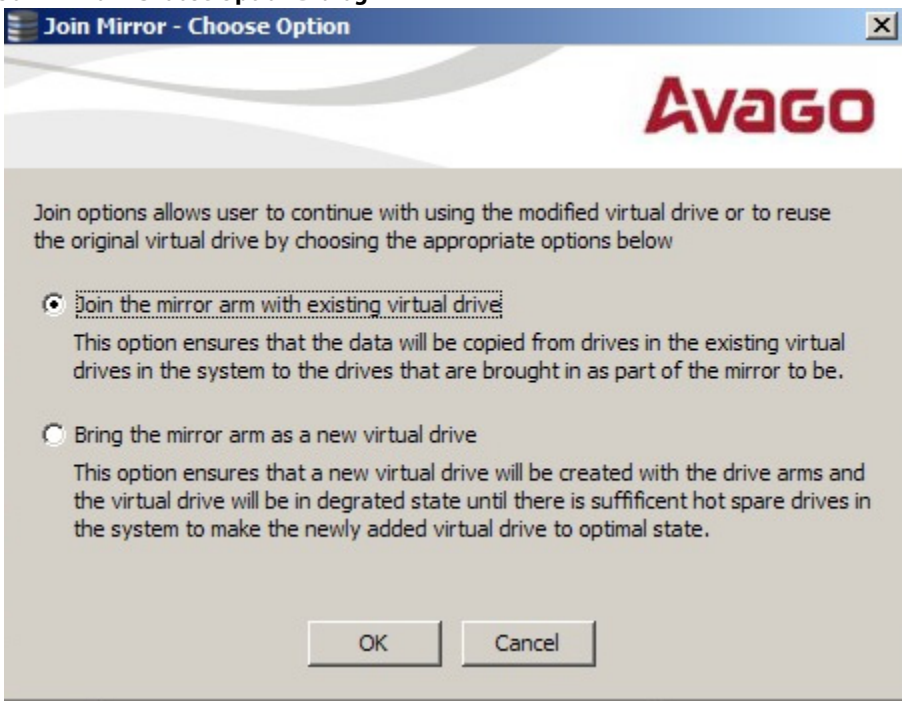

3. Select one of the two options and click **OK**.

   If you select **Join the mirror arm with existing virtual drive**, the following dialog appears.

Figure 221 Confirmation Message Dialog

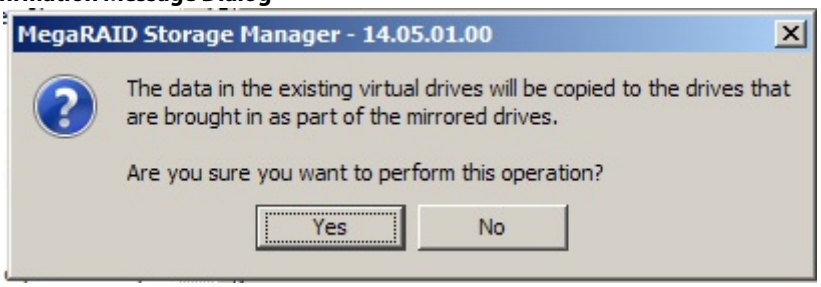

If you select **Bring the mirror arm as a new virtual drive**, the following dialog appears.

Figure 222 Confirmation Message Dialog

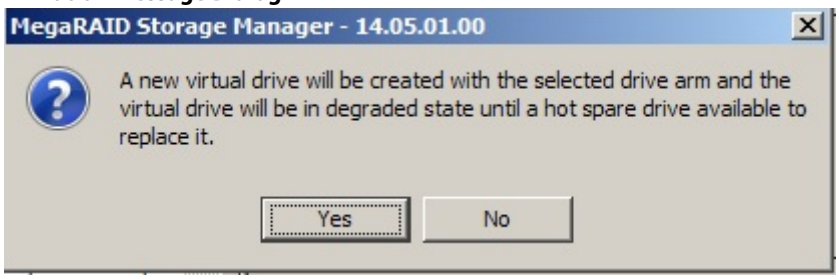

Click **Yes** to proceed with the operation.

9.11.0.1 Configuration

4. This chapter explains how to use MegaRAID Storage Manager software to create and modify storage configurations on Avago SAS controllers.
5. The Avago 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 Introduction to RAID.

NOTE You cannot create or modify a storage configuration unless you are logged on to a server with administrator privileges.

NOTE The MegaRAID Storage Manager software supports 64 virtual drive creation. It does not support 240 virtual drive creation. If you have created 240 virtual drives with some other application and then launch the MegaRAID Storage Manager software, it displays only 64 virtual drives. For more information see the 240 Virtual Drive Feature Limitations appendix.

9.11.0.1.1 Creating a New Configuration

6. You can use the MegaRAID Storage Manager software to create new storage configurations on systems with 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.
7. This section describes the virtual drive parameters and explains how to create simple and advanced storage configurations.

9.11.0.1.1.1 Selecting Virtual Drive Settings

8. 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** (Default)

The new configuration is not initialized, and the existing data on the drives is not overwritten.

— **Fast Initialization**

The firmware quickly writes 0s to the first and last 8-MB regions of the new virtual drive and then completes the initialization in the background. This allows you to start writing data to the virtual drive immediately.

— **Full Initialization**

A complete initialization is done on the new configuration. You cannot write data to the new virtual drive until the initialization is complete. This process can take a long time if the drives are large.

NOTE BGI is supported only for RAID 5 and RAID 6 and not for any other RAID levels. New RAID 5 virtual drives require at least five drives for a background initialization to start. New RAID 6 virtual drives require at least seven drives for a background initialization to start. If there are fewer drives, the background initialization does not start.

— **Strip size**

Strip sizes of 8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB, and 1024 KB are supported. The default is 64 KB. For more information, see the *striping* entry in ***UNRESOLVED***.

NOTE The supported strip size can vary from a minimum of 64 KB to 1 MB for MegaRAID controllers and only 64 KB for Integrated MegaRAID controllers. The LSISAS2108 controller supports strip size from 8 KB to 1 MB.

— **Read policy**

Specify the read policy for this virtual drive. By default, No Read Ahead cache policy is applied. The possible options follow:

— **Always Read Ahead**

Permits the controller to read sequentially ahead of the requested data and allows the controller to store the additional data in the cache memory. Here, the controller anticipates that the data is required frequently. Even though Always Read Ahead policy speeds up the reads for sequential data, but little improvement is seen when accessing the random data.

— **No Read Ahead**

Disables the Always Read Ahead capability of the controller.

— **Write policy**

Specify the write policy for this virtual drive:

— **Write Through**

In this mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This option eliminates the risk of losing cached data in case of a power failure.

— **Always Write Back**

In this mode, the controller sends a data transfer completion signal to the host when the controller cache receives all the data in a transaction. If you select the Always Write Back policy and the battery is absent, the firmware is forced to use the Write Back policy.

— **Write Back** (Default)

In this mode, the controller sends a data transfer completion signal to the host when the controller cache receives all of the data in a transaction. If you select the Write Back policy and the battery is absent, the firmware disables the Write Back policy and defaults to the Write Through policy. This option provides a good balance between data protection and performance.

NOTE The write policy depends on the status of the BBU. If the BBU is not present, is low, is failed, or is being charged, the current write policy switches to Write Through, which provides better data protection.

— **I/O policy**

The I/O policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.

— **Cached IO**

Cached I/O path is the path for data transfer between the hosts and disks that involves the cache.

— **Direct IO**

Direct IO is a direct path for the data transfer between the host and disks that does not involve the cache.

Cached IO provides faster processing, and **Direct IO** makes sure that the cache and the host contain the same data. The firmware might switch from the cached I/O path to the direct I/O path automatically based on the cache policy settings and the status of the battery backup unit.

— **Access policy**

Select the type of data access that is allowed for this virtual drive.

— **Read/Write** (Default)

Allow read/write access. This setting is the default value.

— **Read Only**

Allow read-only access.

— **Blocked**

Do not allow access.

— **Disk cache policy**

Select a cache setting for this drive:

— **Enabled**

Enable the disk cache.

— **Disabled** (Default)

Disable the disk cache.

— **Unchanged**

Leave the current disk cache policy unchanged.

9.11.0.1.1.2 Optimum Controller Settings for CacheCade

9. Write Policy: Write Back/Write Through/Always Write Back

9.11.0.1.1.3 Optimum Controller Settings for Fast Path

10. Write Policy: Write Through
11. IO Policy: Direct IO

12. Read Policy: No Read Ahead
13. Stripe Size: 64 KB

9.11.0.1.1.4 Creating a Virtual Drive Using Simple Configuration

14. Simple configuration is the quickest and easiest way to create a new storage configuration. When you select simple configuration mode, the system creates the best configuration possible using the available drives.

    NOTE You cannot create spanned drives using the simple configuration procedure. To create spanned drives, use the advanced configuration procedure described in Creating a Virtual Drive Using Advanced Configuration.

    NOTE When a physical drive is in the **Prepare for Removal** state, you cannot create a virtual drive using that physical drive. To create a virtual drive when the physical drive is in the Prepare for Removal state, you must manually undo the operation by navigating to the **Undo Removal** option.

15. Follow these steps to create a new storage configuration in simple configuration mode.

1. Perform either of the following steps:
   — Right-click the controller node in the device tree in the left frame of the **MegaRAID Storage Manager** window, and select **Create Virtual Drive**.
   — Select the controller node, and select **Go To > Controller > Create Virtual Drive** in the menu bar, as shown in the following figure.

Figure 223 Create Virtual Drive Menu Option

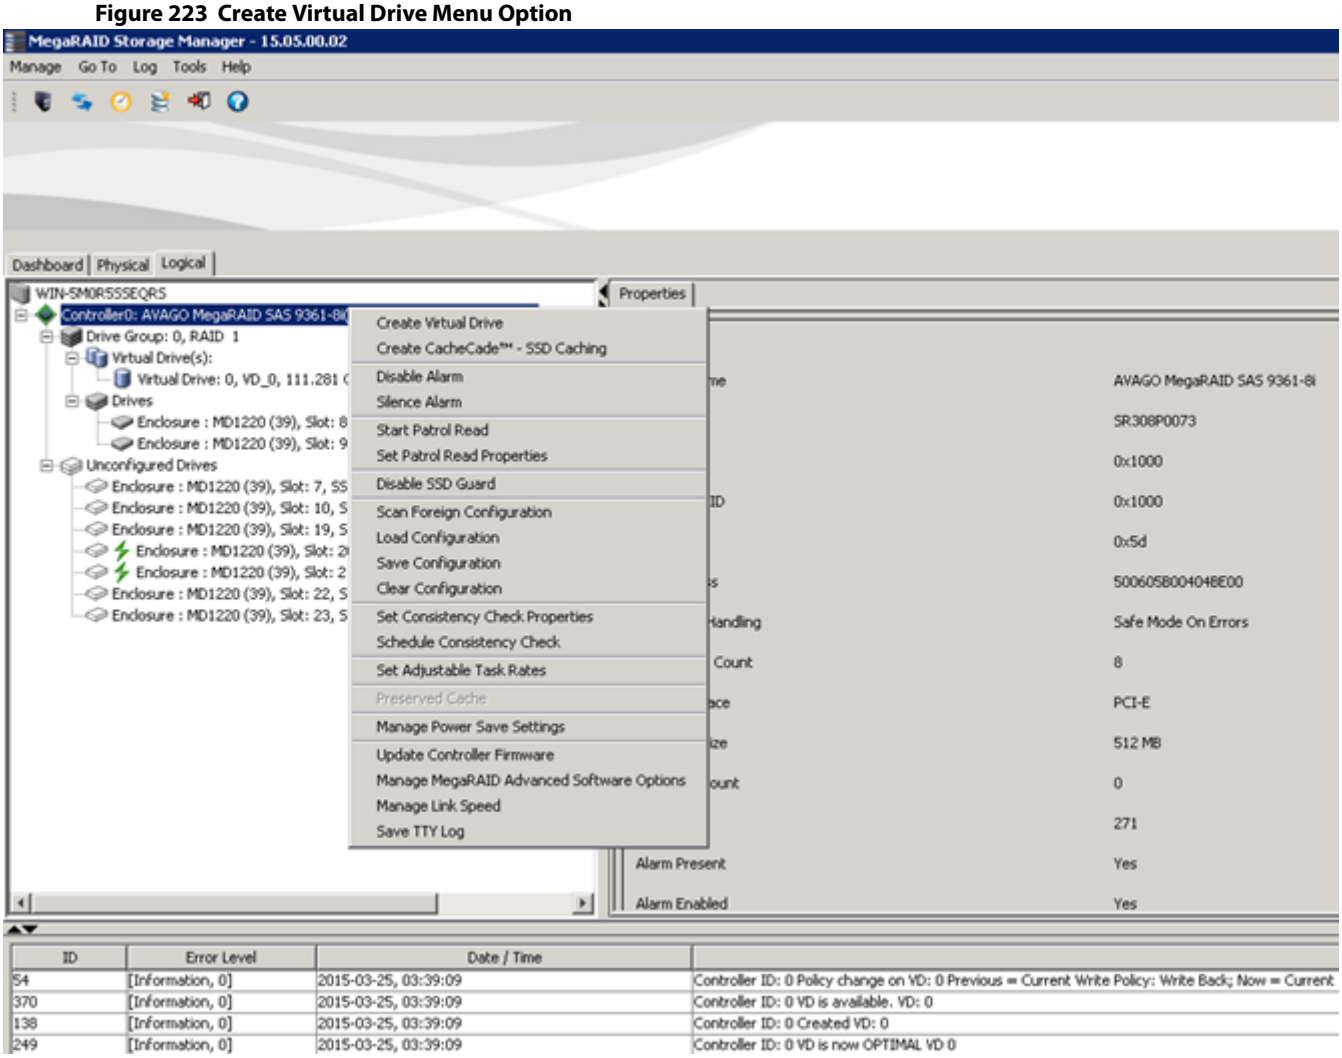

The dialog for the configuration mode (simple or advanced) appears, as shown in the following figure.

Figure 224 Create Virtual Drive – Choose mode

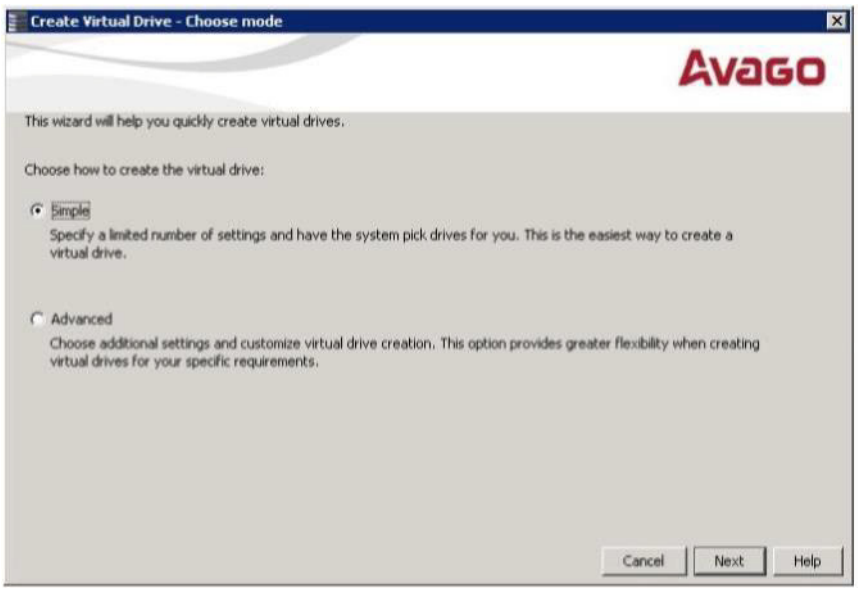

2. Select the **Simple** radio button, and click **Next**.

   The **Create Virtual Drive – Allocate capacity** dialog appears, as shown in the following figure. If unconfigured drives are available, you have the option to use those unconfigured drives. If unconfigured drives are available, the **Create Drive Group Settings** window appears, and you can go to step 4.

Figure 225 Using the Free Capacity of an Existing Drive Group

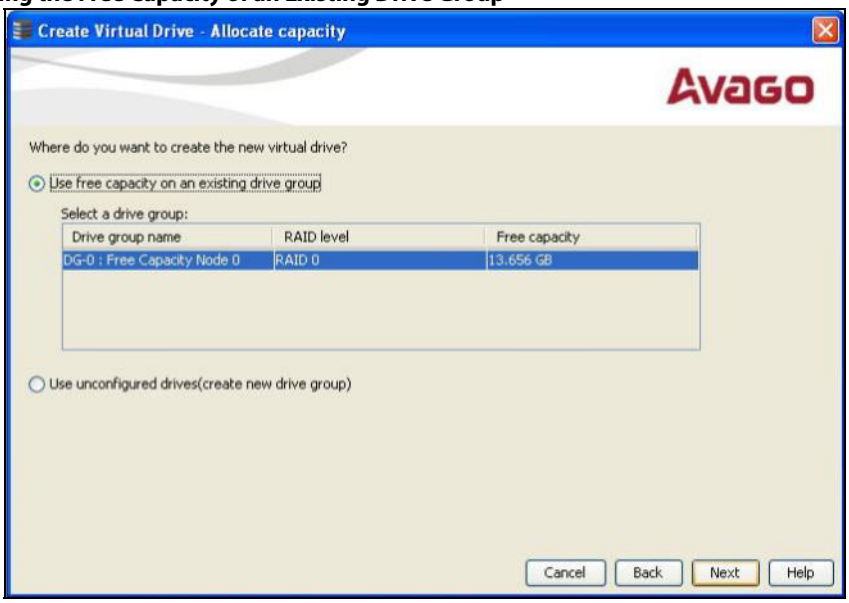

3. Perform either of the two options:
   — If a drive group exists, select the **Use free capacity on an existing drive group** radio button and click **Next**. Continue with step 4. The **Create Virtual Drive** window appears, as shown in the following figure. If different

types of drives are attached to the controller, such as HDD, SDD, SAS, and SATA, an option appears to allow drive type mixing.

— If unconfigured drives are available, select the radio button to use the unconfigured drives, and click **Next**. Continue with step 10. The **Summary** window appears as shown in the Create Virtual Drive – Summary Window figure.

Figure 226 Create Virtual Drive – Drive group and Virtual drive settings Dialog

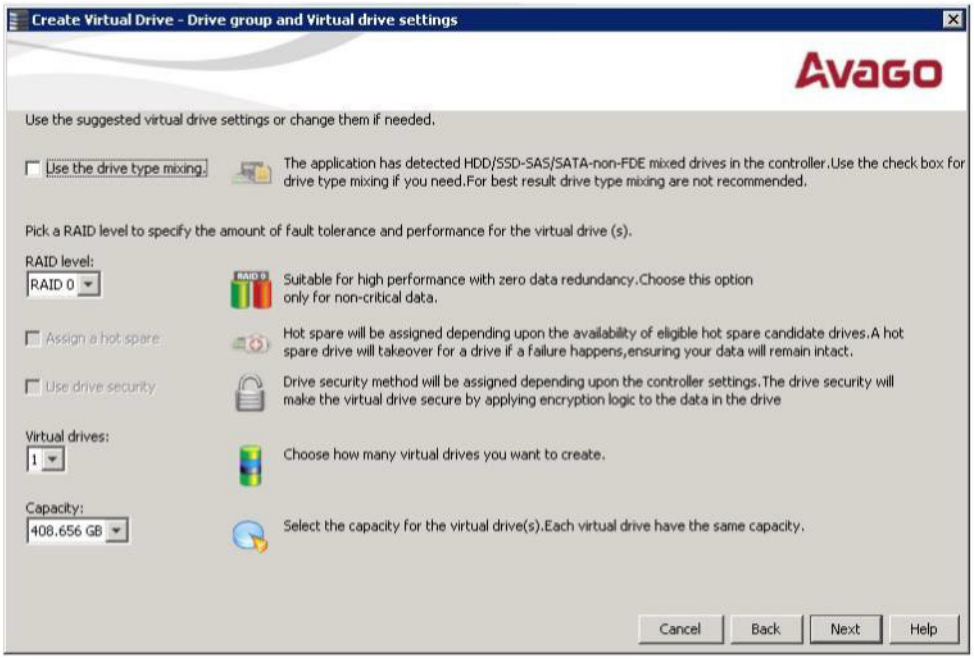

4. If you want to allow different types of drives in a configuration, select the **Use the drive type mixing** check box.

   NOTE For best results, do not use drive type mixing.

5. Select the RAID level desired for the virtual drive.

   When you use simple configuration, the RAID controller supports RAID levels 1, 5, and 6. In addition, it supports independent drives (configured as RAID 0). The window text gives a brief description of the RAID level that you select. The RAID levels that you can choose depend on the number of drives available.

6. Select the **Assign a hot spare** check box if you want to assign a dedicated hot spare to the new virtual drive.

   If an unconfigured good drive is available, that drive is assigned as a hot pare. Hot spares are drives that are available to replace failed drives automatically in a redundant virtual drive (RAID 1, RAID 5, or RAID 6).

7. Select the **Use drive security** check box if you want to set a drive security method.

   The Avago 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. See Avago MegaRAID SafeStore Encryption Services, for more information about the SafeStore feature.

8. Use the drop-down list in the **Virtual drives** field to choose how many virtual drives you want to create.

9. Select the capacity of the virtual drives.

   Each virtual drive has the same capacity.

10. Click **Next**.

    The **Create Virtual Drive – Summary** window appears, as shown in the following figure. This window shows the selections you made for simple configuration.

Figure 227 Create Virtual Drive – Summary Window

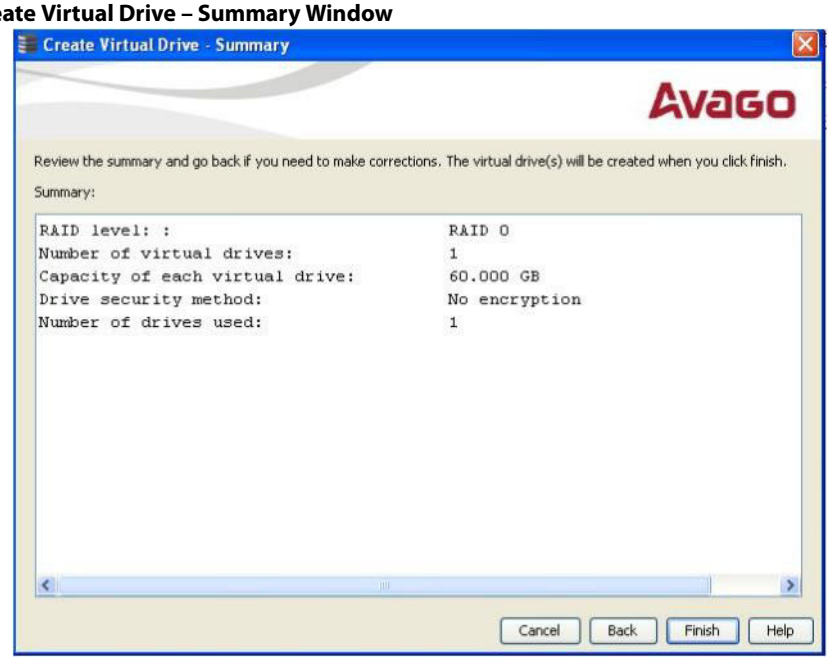

NOTE If High Availability DAS is supported on the controller and you are creating a virtual drive using simple configuration, by default, the virtual drive is shared with the other servers in that cluster.

11. Either click **Back** to return to the previous window to change any selections, or click **Finish** to accept and complete the configuration.

    The new virtual drive is created after you click **Finish**. After the configuration is completed, a dialog box notifies you that the virtual drives were created successfully.

NOTE If you create a large configuration using drives that are in Power-Save mode, it could take several minutes to spin up the drives. A progress bar appears as the drives spin up. If any of the selected unconfigured drives fail to spin up, a dialog box that identifies these drives appears.

9.11.0.1.1.5 Creating a Virtual Drive Using Advanced Configuration

12. 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.

NOTE When a physical drive is in the **Prepare for Removal** state, you cannot create a virtual drive using that physical drive.

13. Follow these steps to create a new storage configuration in the advanced configuration mode. This example shows the configuration of a spanned drive group.

1. Perform either of the following steps to bring up the **Configuration** wizard:
   - Right-click the controller node in the device tree in the left frame of the **MegaRAID Storage Manager** window, and select **Create Virtual Drive**.
   - Select the controller node, and select **Go To > Controller > Create Virtual Drive** in the menu bar.

   The dialog for the choosing the configuration mode (simple or advanced) appears, as shown in the following figure.

Figure 228 Create Virtual Drive - Choose mode Dialog

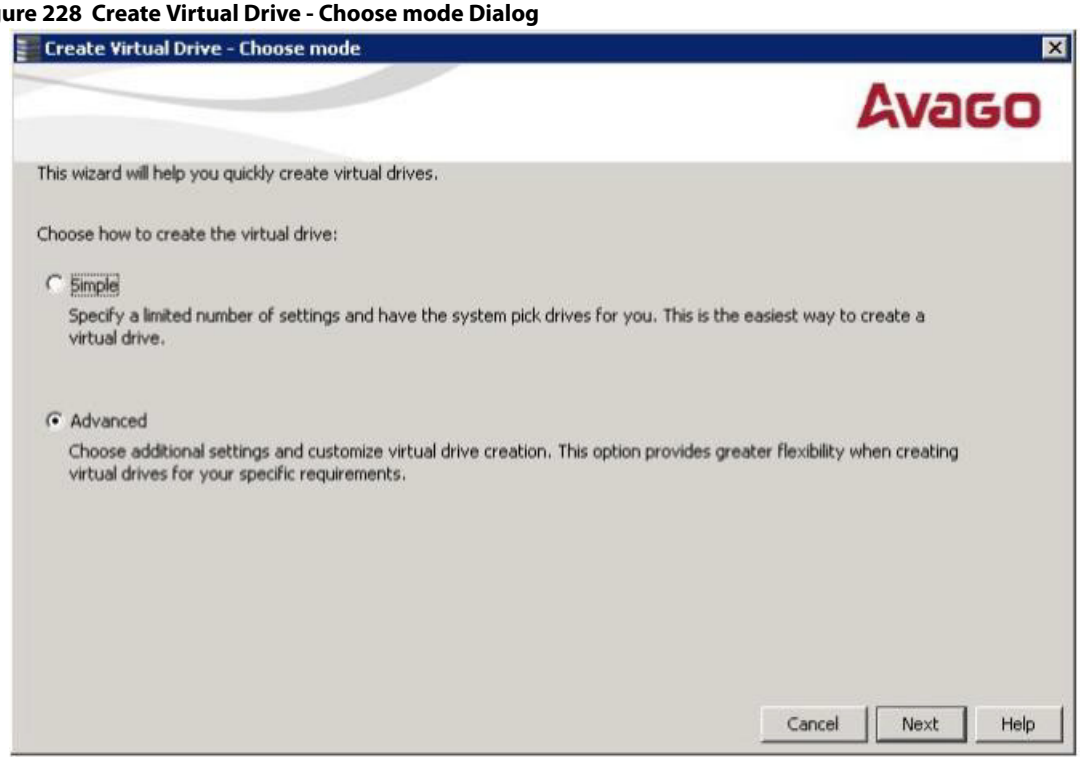

2. Select the **Advanced** radio button, and click **Next**.

   The **Create Drive Group Settings** window appears, as shown in the following figure.

Figure 229 Create Drive Group - Drive Group Settings Window

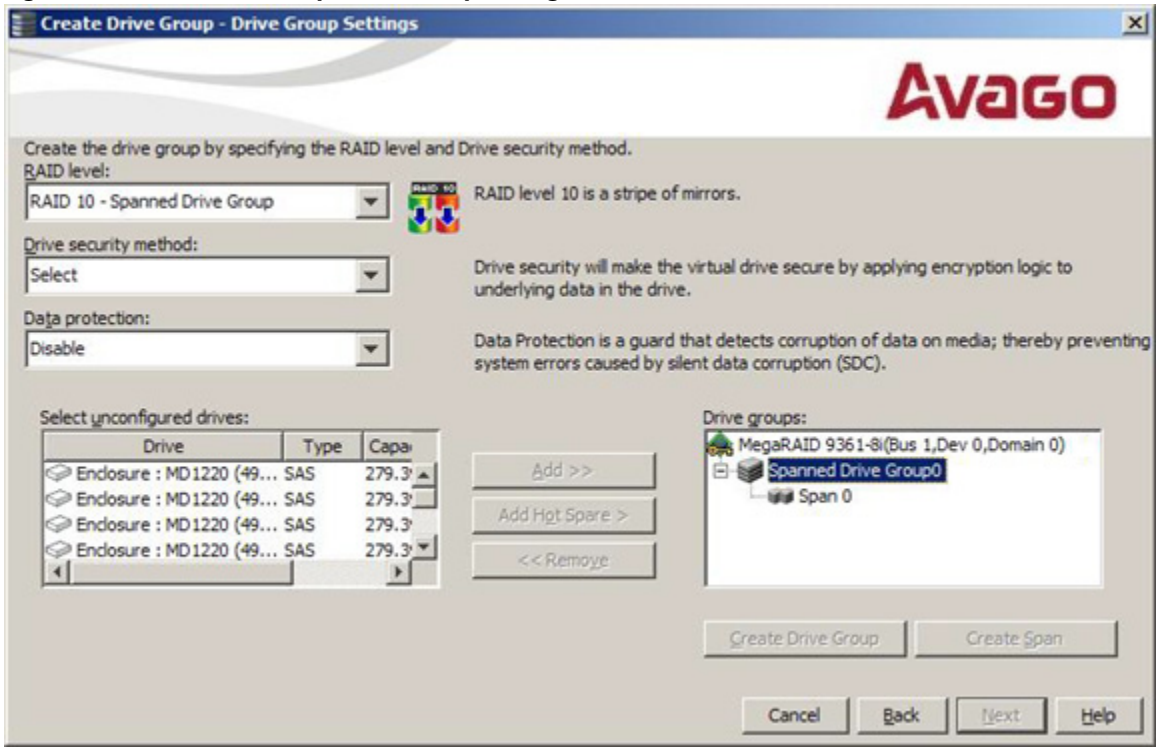

3. Select the following items on the **Create Drive Group - Drive Group Settings** window:
   a. Select the RAID level desired for the drive group from the drop-down menu. To make a spanned drive, select **RAID 10**, **RAID 50**, or **RAID 60** in the **RAID level** field.

      Drive Group 0 and **Span 0** appear in the **Drive groups** field when you select RAID 10, 50, or 60.

      The RAID controller supports RAID levels 1, 5, 6, 10, 50, and 60. In addition, it supports independent drives (configured as RAID 0 and RAID 00). The dialog text gives a brief description of the RAID level that you select. You can choose the RAID levels depending on the number of available drives.
   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 Avago MegaRAID SafeStore Encryption Services for more information about drive security and encryption.
   c. Select *unconfigured* drives from the list of drives, and click **Add>** to add them to the drive group.

      The selected drives appear under **Span 0** below **Drive Group 0**, as shown in the following figure.

Figure 230 Span 0 of Drive Group 0

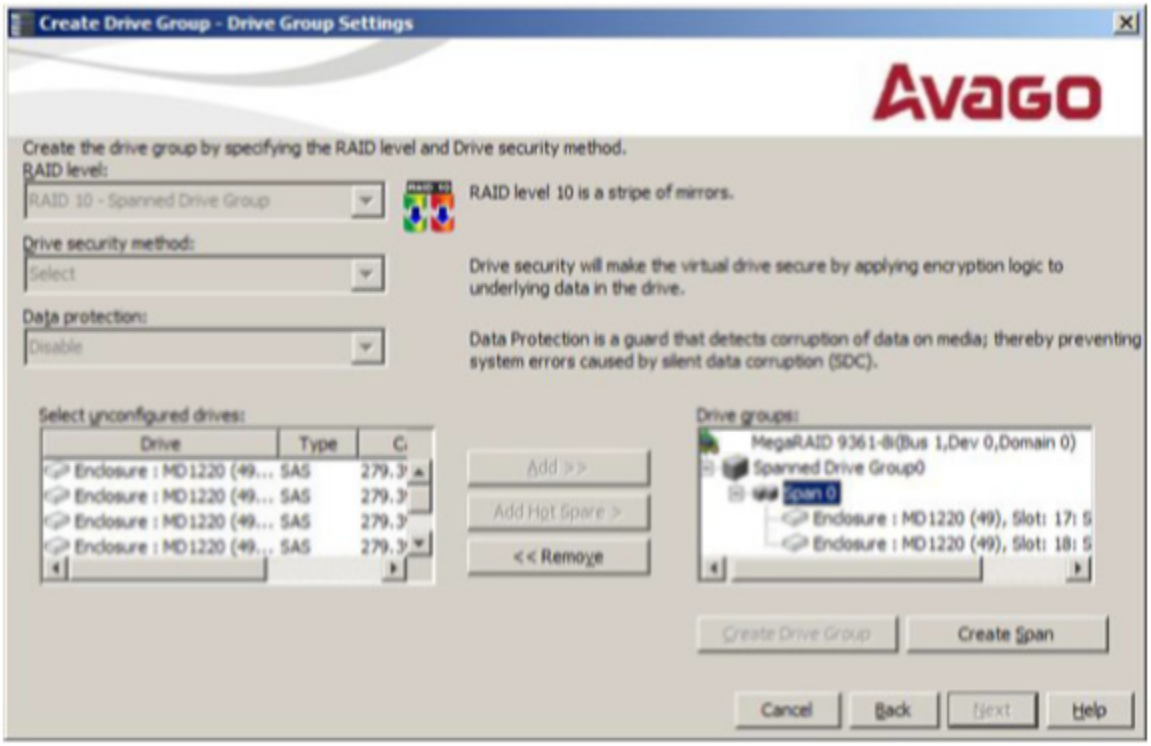

   d. Click **Create Span** to create a second span in the drive group.
   e. Select *unconfigured* drives from the list of drives, and click **Add>** to add them to the second drive group.

      The selected drives appear under **Span 1** below **Drive Group 0**, as shown in the following figure.

Figure 231 Span 0 and Span 1 of Drive Group 0

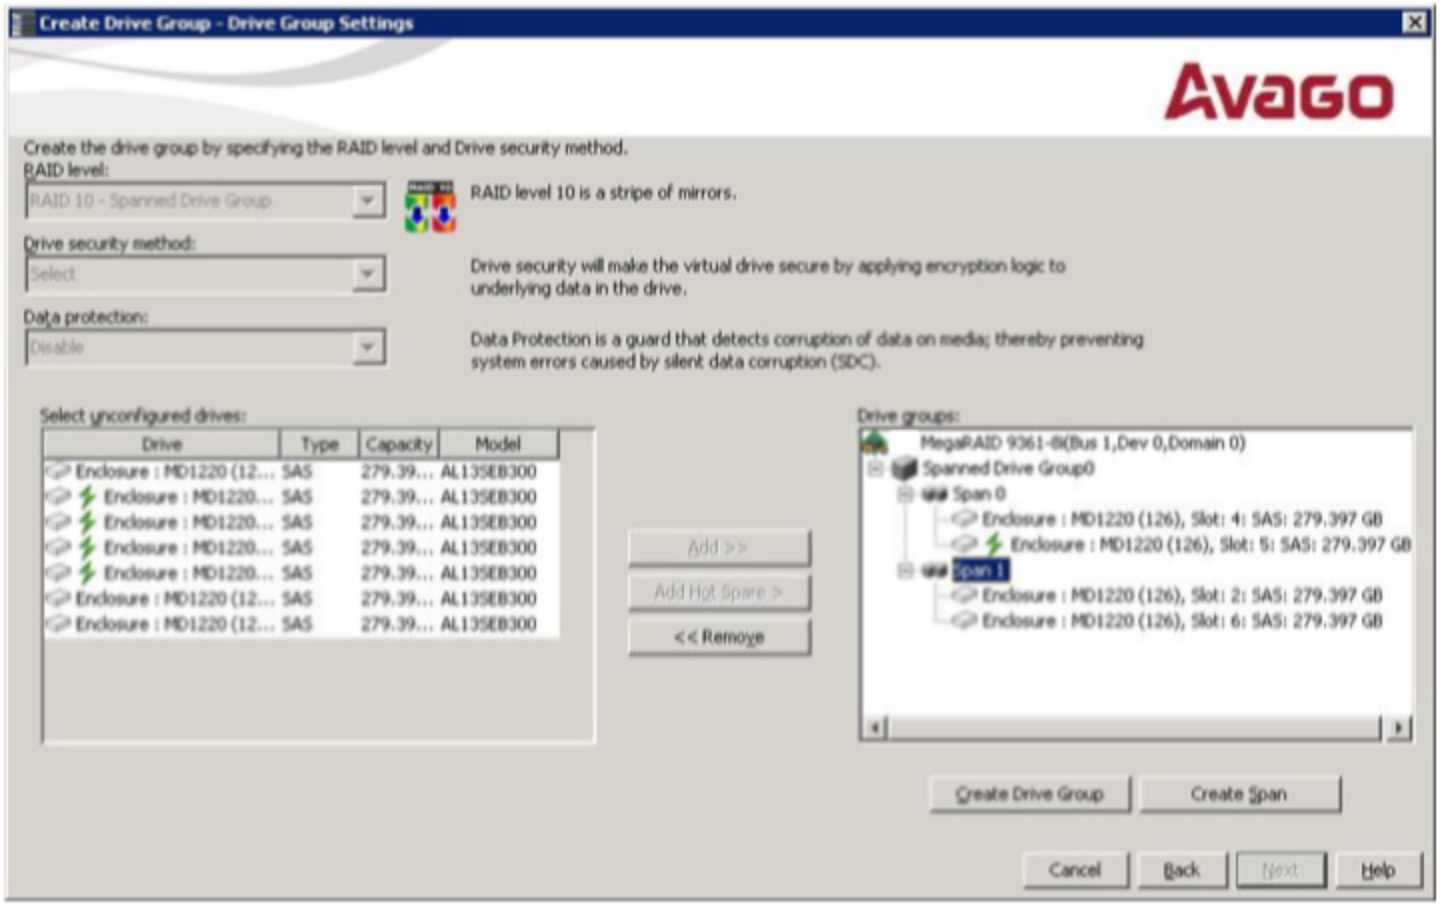

f. Click **Create Drive Group** to make a drive group with the spans.

g. Click **Next** to complete this step.

The **Create Virtual Drive - Virtual drive settings** window appears, as shown in the following figure. The drive group and the default virtual drive settings appear. The options to update the virtual drive or remove the virtual drive are grayed out until you create the virtual drive.

NOTE The parameters in the **Create Virtual Drive - Virtual drive settings** window display in Disabled mode (grayed out) for SAS-Integrated RAID (IR) controllers because these parameters do not apply to SAS-IR controllers.

Figure 232 Create Virtual Drive - Virtual drive settings Window

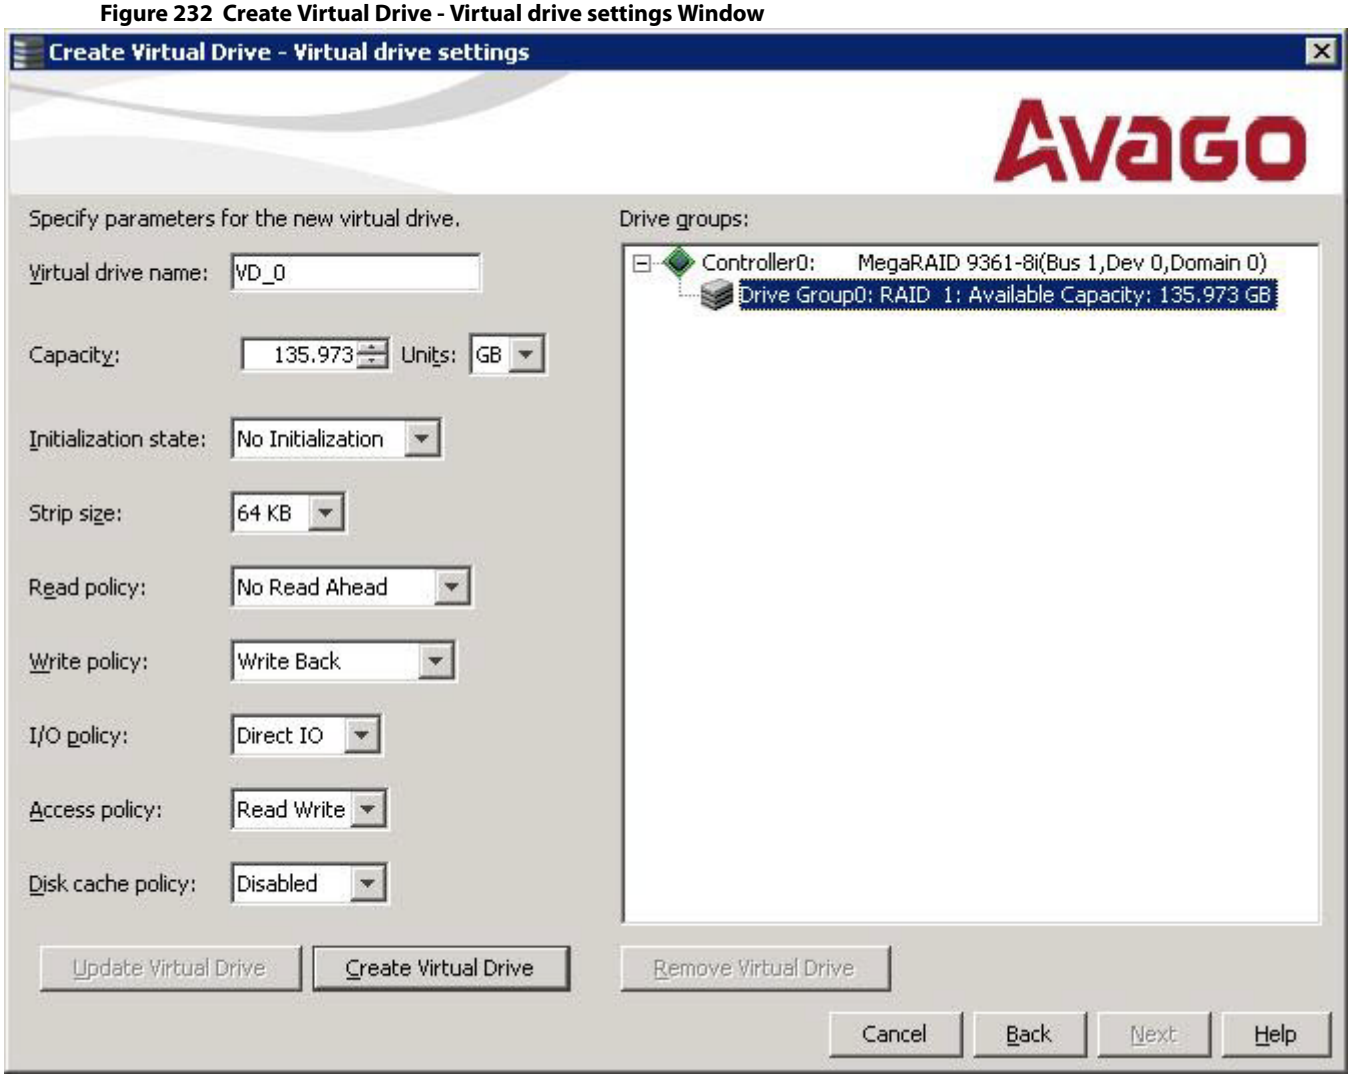

NOTE If you select **Write Back** as the write policy, and no battery exists, the battery is low or failed, or the battery is running through a re-learn cycle, the write policy switches to **Write Through**. This setting eliminates the risk of data loss in case of a power failure. A message window notifies you of this change.

NOTE If the controller supports High Availability DAS, the **Provide Shared Access** option appears in the above dialog. Select this option if you want the virtual drive to be shared between the two servers in a cluster.

4. Change any virtual drive settings, if desired.

   See Selecting Virtual Drive Settings for more information about the virtual drive settings.

5. Click **Create Virtual Drive**.

   The new virtual drive appears under the drive group. The options **Update Virtual Drive** and **Remove Virtual Drive** are available. **Update Virtual Drive** allows you to change the virtual drive settings, and **Remove Virtual Drive** allows you to delete the virtual drive.

6. Click **Next**.

   The **Create Virtual Drive - Summary** window appears, as shown in the following figure. This window shows the selections you made for advanced configuration.

Figure 233 Create Virtual Drive - Summary Window

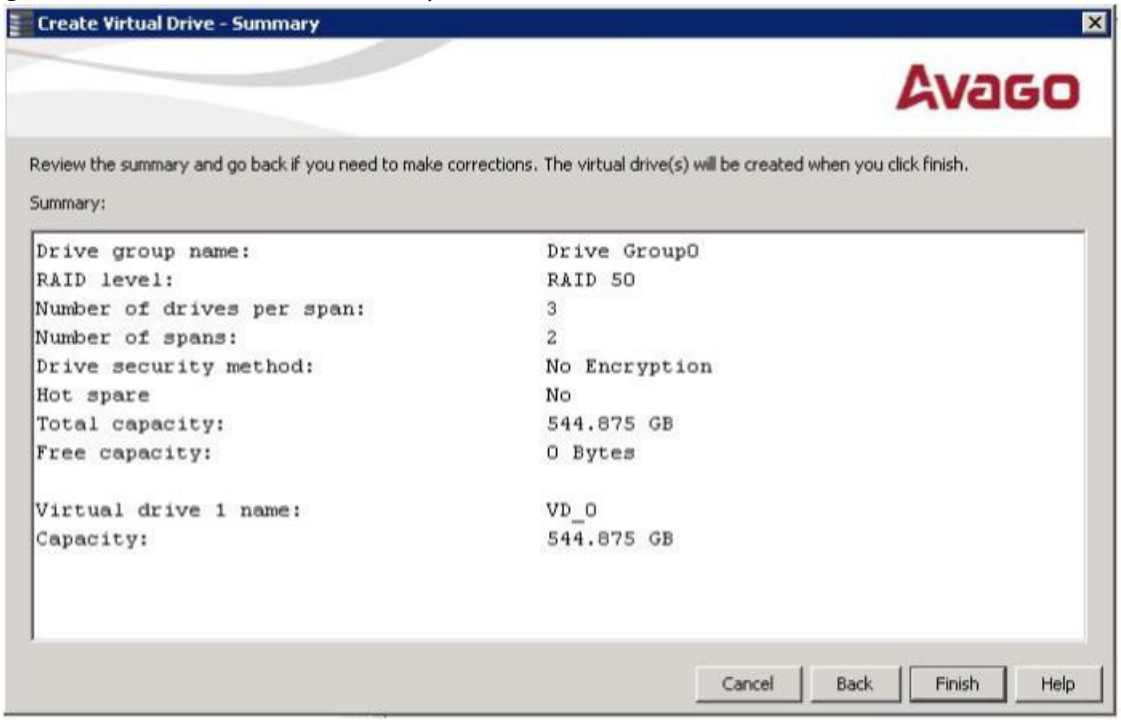

7. Click **Back** to return to the previous window to change any selections, or click **Finish** to accept and complete the configuration.

   After you click **Finish**, the new storage configuration is created and initialized according to the selected options.

   NOTE If you create a large configuration using drives that are in Power-Save mode, it can take several minutes to spin up the drives. A progress bar appears as the drives spin up. If any of the selected unconfigured drives fail to spin up, a dialog appears that identifies the drives.

   After the configuration is completed, a dialog notifies you that the virtual drives were created successfully.

8. Click **OK**.

   The **Enable SSD Caching on New Virtual Drives** dialog appears.

   The newly created virtual drive is enabled for SSD caching by default.

9. Click **OK** to confirm SSD caching on the virtual drive. Click **No** if you want to disable SSD caching on the virtual drive.

   The **All** check box is selected by default. To disable SSD caching on the virtual drives, deselect the **All** check box.

   If more drive capacity exists, the dialog asks whether you want to create more virtual drives. If no more drive capacity exists, you are prompted to close the configuration session.

10. Select either **Yes** or **No** to indicate whether you want to create additional virtual drives.

    If you select **Yes**, the system takes you to the Create Virtual Drive – Drive group and Virtual drive settings Dialog. If you select **No**, the utility asks whether you want to close the wizard.

11. If you selected **No** in the previous step, select either **Yes** or **No** to indicate whether you want to close the wizard.

    If you select **Yes**, the **Configuration** wizard closes. If you select **No**, the dialog closes, and you remain on the same page.

9.11.0.1.2 Converting JBOD Drives to Unconfigured Good

12. You can convert JBOD drives to Unconfigured Good using the **Create Virtual Drive** option or **Make Unconfigured Good** drive option with a single configuration.

    NOTE MegaRAID SAS 9240-4i and MegaRAID SAS 9240-8i controllers support JBOD.

13. Perform the following steps to configure JBOD to Unconfigured Good drives:

1. Perform one of these actions:
   - Right-click the controller node in the device tree in the left frame of the **MegaRAID Storage Manager** window, and select **Create Virtual Drive**.
   - Select the controller node, and select **Go To > Controller > Create Virtual Drive**.

   The **Create Virtual Drive - JBOD to Unconfigured Good Conversion** wizard appears, as shown in the following figure.

Figure 234 Create Virtual Drive - JBOD to Unconfigured Good Conversion Dialog

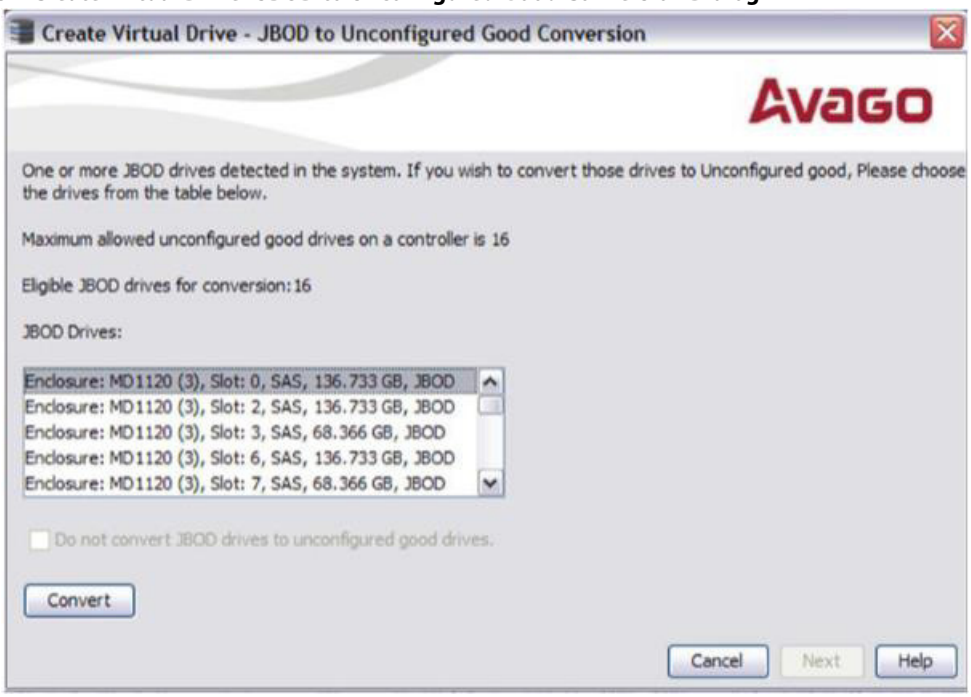

The **JBOD Drives** field displays the available JBOD drives available in the system.

NOTE If you do not want to make any JBOD as unconfigured good drives, select the **Do not convert JBOD drives to unconfigured good drives** check box, and the MegaRAID Storage Manager application skips changing any selected JBOD to unconfigured good drive.

2. Select the drives that you want configured as Unconfigured Good and then click **Convert**.

   A warning message appears stating that the JBOD drives will be removed and prompting for your confirmation.

3. Select **Confirm** and click **Yes** to proceed with the conversion.

NOTE If one or more JBOD drives have an OS or a file system installed on them, another warning message appears prior to conversion listing those JBOD drives that have an OS or a file system installed on them. The message states that any attempt to convert the listed JBOD drives

to unconfigured good drives would remove the existing data on the drives. Click **Yes** if you want to proceed with the conversion. Otherwise, click **No** to return to the previous screen and unselect those JBOD drives that have the OS installed on them.

4. Click **Next**.
   The **Create Virtual Drive - Drive group and Virtual drive settings** dialog appears.

9.11.0.1.2.1 Converting JBOD to Unconfigured Good from the MegaRAID Storage Manager Main Menu

5. You can also convert JBOD to unconfigured good by performing these steps:
1. Select **Controller >Make Unconfigured Good** from the main **MegaRAID Storage Manager** main menu.
   The **Make Configured Good** dialog appears, as shown in the following figure.

Figure 235 Make Configured Good Dialog

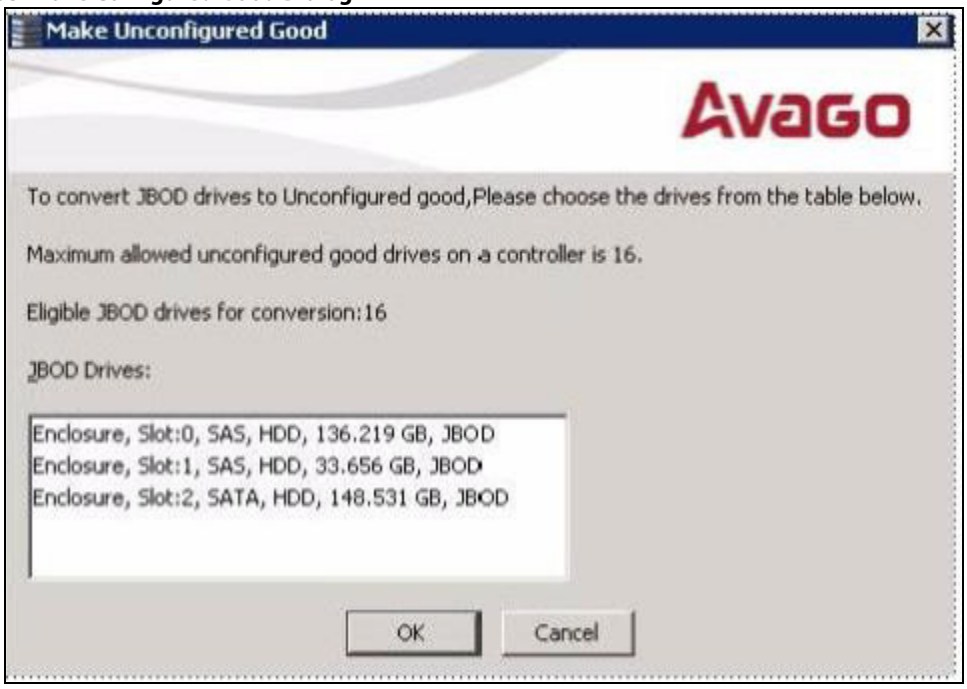

2. Select the JBOD drives to be configured as unconfigured good.
3. Click **OK**.
   A warning message appears stating that the JBOD drives will be removed and prompting for your confirmation.
4. Select **Confirm** and click **Yes** to proceed with the conversion.
   The selected JBOD drives are configured as unconfigured good.

NOTE If one or more JBOD drives have OS or file system installed on them, prior to conversion, another warning message appears listing those JBOD drives that have the OS or file system installed on them. The message states that any attempt to convert the listed JBOD drives to unconfigured good drives would remove the existing data on the drives. Click **Yes** if you want to proceed with the conversion. Else, click **No** to return to the previous screen and unselect those JBOD drives that have the OS installed on them.

9.11.0.1.2.2 Removing a JBOD Drive

5. Follow these steps to remove a JBOD drive from the physical view:

1. Select the Physical View tab in the left panel of the MegaRAID Storage Manager window.
2. Right-click the physical Drive node, which is configured as a JBOD.
3. Select **Remove JBOD**.

   A warning message appears stating that the JBOD drive will be removed and prompting for your confirmation.
4. Select **Confirm** and click **Yes** to remove the JBOD drive.

   The selected JBOD drive is removed.

NOTE If the JBOD drive has an OS or a file system installed on it, another warning message appears prior to removal stating that any attempt to remove the JBOD drive would remove the existing data on the drive. Select **Confirm**, and click **Yes** if you want to proceed with the removal. Otherwise, click **No** to return to the previous screen.

9.11.0.1.3 Enabling Security on JBOD

5. The MegaRAID Storage Manager software supports Enable Security feature for security capable JBOD drives.
6. Follow these steps to enable security on security capable JBOD drives:

1. In the left panel of the **MegaRAID Storage Manager** window, select the Security capable JBOD drive for which you want to enable security.
2. In the **MegaRAID Storage Manager** window, select **Go To** > **Physical Drive** > **Enable Security** as shown in the following figure.

Figure 236 MegaRAID Storage Manager Window

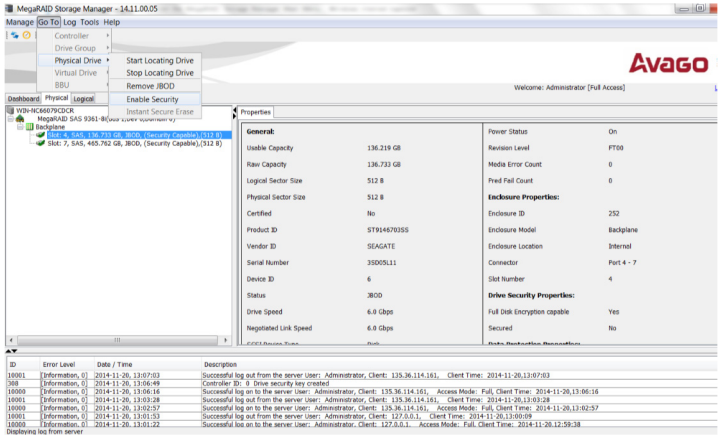

3. You can also make a JBOD secure as follows.

   a. In the **MegaRAID Storage Manager** window, right click on **Controller** >**Make JBOD Secure**.

      The **Make JBOD Secure** dialog appears, as shown in the following figure.

Figure 237 Make JBOD Secure Window

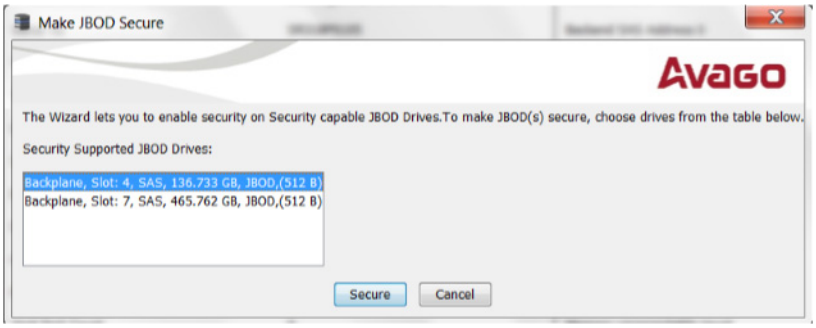

b. Select the JBOD drive for which you need to enable security.

c. Click **Secure**.

9.11.0.1.4 Creating Hot Spare Drives

4. 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.
5. To create a dedicated or global hot spare drive, follow these steps:

1. Select the **Physical** tab in the left panel of the MegaRAID Storage Manager main menu, and click the icon of an unused drive.

   For each drive, the window displays the port number, enclosure number, slot number, drive state, drive capacity, and drive manufacturer.
2. Either select **Go To > Physical Drive > Assign Global Hot Spare**, or select **Go To > Physical Drive > Assign Dedicated Hot Spare**.
3. Perform one of these actions:
   — If you selected **Assign Dedicated Hotspare**, select a drive group from the list that appears. The hot spare is dedicated to the drive group that you select.

   NOTE If the controller supports High Availability DAS, dedicated hot spares can be assigned to only one drive group. If you try to assign dedicated hot spares to more than one drive group, an error message appears.

   — If you selected **Assign Global Hotspare**, skip this step, and go to the next step. The hot spare is available to any virtual drive on a specific controller.
4. Click **Go** to create the hot spare.

   The drive state for the drive changes to dedicated or global hot spare, depending on your selection.

9.11.0.1.5 Changing Adjustable Task Rates

5. If you want to change the Rebuild rate and other task rates for a controller, you must first log onto the server in Full Access mode.

   NOTE Leave the adjustable task rates at their default settings to achieve the best system performance. If you raise the task rates above the defaults, foreground tasks will run more slowly and it might seem that the system is not responding. If you lower the task rates below the defaults, rebuilds and other background tasks might run very slowly and might not complete within a reasonable time. If you decide to change the values, record the original default values here so you can restore them later, if necessary: **Rebuild rate:** _____,

Background Initialization (BGI) rate: _____, Check consistency rate: _____.

6. To change the adjustable task rates, perform the following steps:

1. Select a controller icon in the **Physical** tab or the **Logical** tab in the left panel of the **MegaRAID Storage Manager** window.
2. Select **Go To > Controller > Set Adjustable Task Rates** from the menu bar.

   The **Set Adjustable Task Rates** window appears, as shown in the following figure.

Figure 238 Set Adjustable Task Rates Window

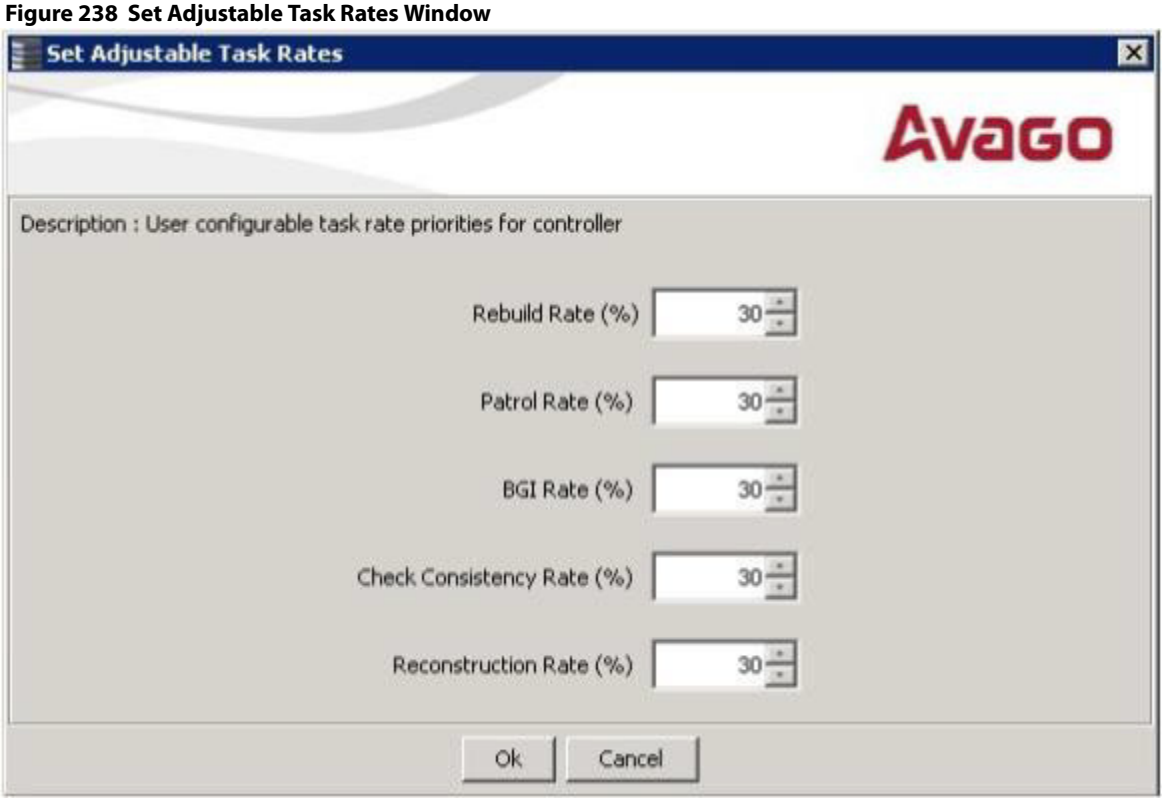

3. Enter changes, as needed, to the following task rates:
   - **Rebuild Rate**. Enter a number from 0 to 100 to control the rate at which a rebuild will be performed on a drive when one is necessary. The higher the number, the faster the rebuild will occur (and the system I/O rate may be slower as a result).
   - **Patrol Rate**. Enter a number from 0 to 100 to control the rate at which patrol reads will be performed. Patrol read monitors drives to find and resolve potential problems that might cause drive failure. The higher the number, the faster the patrol read will occur (and the system I/O rate may be slower as a result).
   - **Background Initialization (BGI) Rate**. Enter a number from 0 to 100 to control the rate at which virtual drives are initialized "in the background." Background initialization establishes mirroring or parity for a RAID virtual drive while allowing full host access to the virtual drive. The higher the number, the faster the initialization will occur (and the system I/O rate may be slower as a result).
   - **Check Consistency Rate**. Enter a number from 0 to 100 to control the rate at which a consistency check is done. A consistency check scans the consistency data on a fault tolerant virtual drive to determine if the data

has become corrupted. The higher the number, the faster the consistency check is performed (and the system I/O rate may be slower as a result).

— **Reconstruction Rate**. Enter a number from 0 to 100 to control the rate at which reconstruction of a virtual drive occurs. The higher the number, the faster the reconstruction occurs (and the system I/O rate may be slower as a result).

4. Click **Ok** to accept the new task rates.
5. When the warning message appears, click **OK** to confirm that you want to change the task rates.

9.11.0.1.6 Changing Power Settings

6. The RAID controller includes Dimmer Switch® technology that conserves energy by placing certain unused drives into Power-Save mode. In Power-Save mode, the drives use less energy, and the fan and the enclosure require less energy to cool and house the drives, respectively. Also, this technology helps avoid application timeouts caused by spin-up delays and drive wear caused by excessive spin-up/down cycles.
7. You can use the **Power Settings** field in the MegaRAID Storage Manager software to choose whether to allow unconfigured drives or Commissioned Hotspares to enter Power-Save mode.

NOTE The Dimmer Switch technology is enabled by default.

8. When they are in the Power-Save mode, unconfigured drives and drives configured as Commissioned Hotspares (dedicated or global) can be spun down. When spun down, the drives stay in Power-Save mode except for periodic maintenance, which includes the following:
   — Periodic background media scans (Patrol Read) to find and correct media defects to avoid losing data redundancy (hot spare drives only)
   — Use of a Commissioned Hotspare to rebuild a degraded drive group (Commissioned Hotspare drives only)
   — Update of disk data format (DDF) and other metadata when you make changes to RAID configurations (Commissioned Hotspare drives and unconfigured drives)

NOTE If your controller does not support this option, the **Power Settings** field does not appear.

9. Follow these steps to change the power-save setting.

1. Select a controller icon in the **Physical** tab or the **Logical** tab in the left panel of the **MegaRAID Storage Manager** window.
2. Select **Go To > Controller > Manage Power Settings** from the menu bar.

   The **Manage Power Save Settings** dialog appears.

Figure 239 Manage Power Save Settings

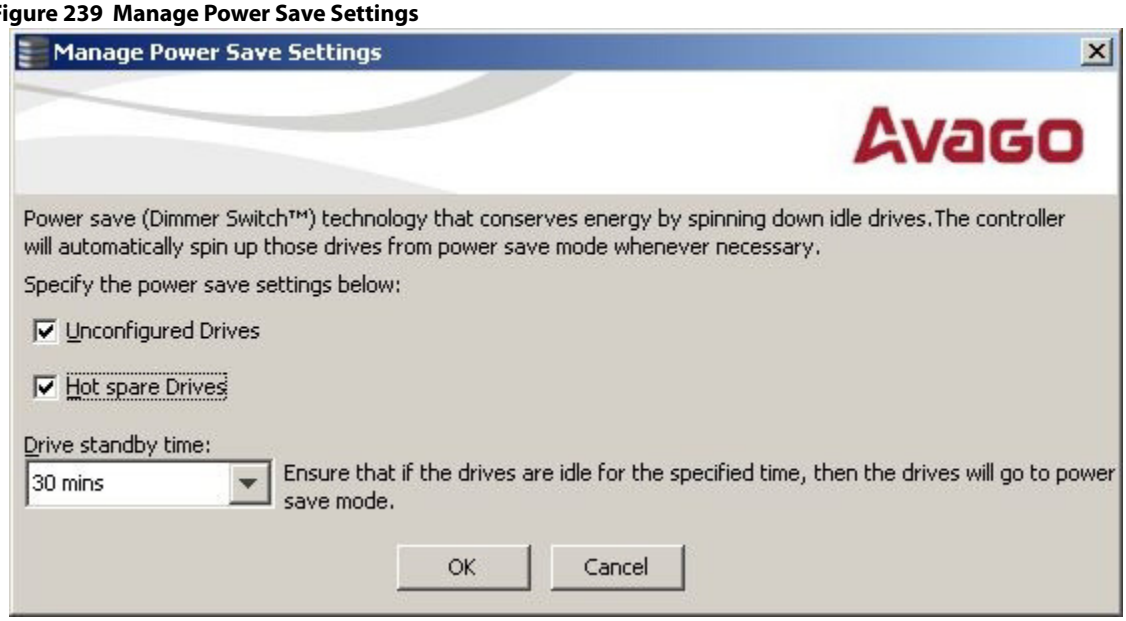

3. Select the **Unconfigured Drives** check box to let the controller enable the unconfigured drives to enter the Power-Save mode.
4. Select the **Hot spare Drives** check box to let the controller enable the Hot spare drives to enter the Power-Save mode.
5. Select the drive standby time (Alt+D) using the drop-down list from the **Drive standby time** field.

   NOTE The **Drive Standby time** drop-down list is enabled only if any of the check boxes above it are checked. The drive standby time can be 30 minutes, 1 hour, 1.30 hours, or 2 hours through 24 hours.

6. Click **OK**.

   The Power-Save settings are saved. After you click **OK**, a confirmation dialog appears prompting you to save your changes.

   If you do not specify the Power-Save settings in the **Manage Power Save Settings** dialog, a confirmation dialog appears. The confirmation dialog mentions that the system does not have power savings for any of the drives, and asks if you would like to proceed.

9.11.0.1.7 Recovering and Clearing Punctured Block Entries

7. You can recover and clear the punctured block area of a virtual drive.

   ATTENTION This operation removes any data stored on the physical drives. Back up the good data on the drives before making any changes to the configuration.

8. When a Patrol Read or a Rebuild operation encounters a media error on the source drive, it punctures a block on the target drive to prevent the use of the data with the invalid parity. Any subsequent read operation to the punctured block completes but with an error. Consequently, the puncturing of a block prevents any invalid parity generation later while using this block.
9. To recover or clear the punctured block area of a virtual drive, run a Slow (or Full) Initialization to zero out and regenerate new parity causing all bad block entries to be removed from the bad block table.
10. To run a Slow (or Full) Initialization, see Selecting Virtual Drive Settings.

9.11.0.1.8 Changing Virtual Drive Properties

11. You can change the read policy, write policy, and other virtual drive properties at any time after a virtual drive is created.

ATTENTION Do not enable drive caching on a mirrored drive group (RAID 1 or RAID 1E). If you do, data can be corrupted or lost in the event of a sudden power loss. A warning appears if you try to enable drive caching for a mirrored drive group.

NOTE For virtual drives with SAS drives only, set the drive write cache policy set to **Disabled**, by default. For virtual drives with SATA drives only, set the drive write cache policy to **Enabled**, by default.

12. To change the virtual drive properties, perform the following steps:

1. Select a virtual drive icon in the **Physical** tab or the **Logical** tab in the left panel of the **MegaRAID Storage Manager** window.
2. Select **Go To > Virtual Drive > Set Virtual Drive Properties** from the menu bar.

   The **Set Virtual Drive Properties** dialog appears, as shown in the following figure.

Figure 240 Set Virtual Drive Properties Dialog

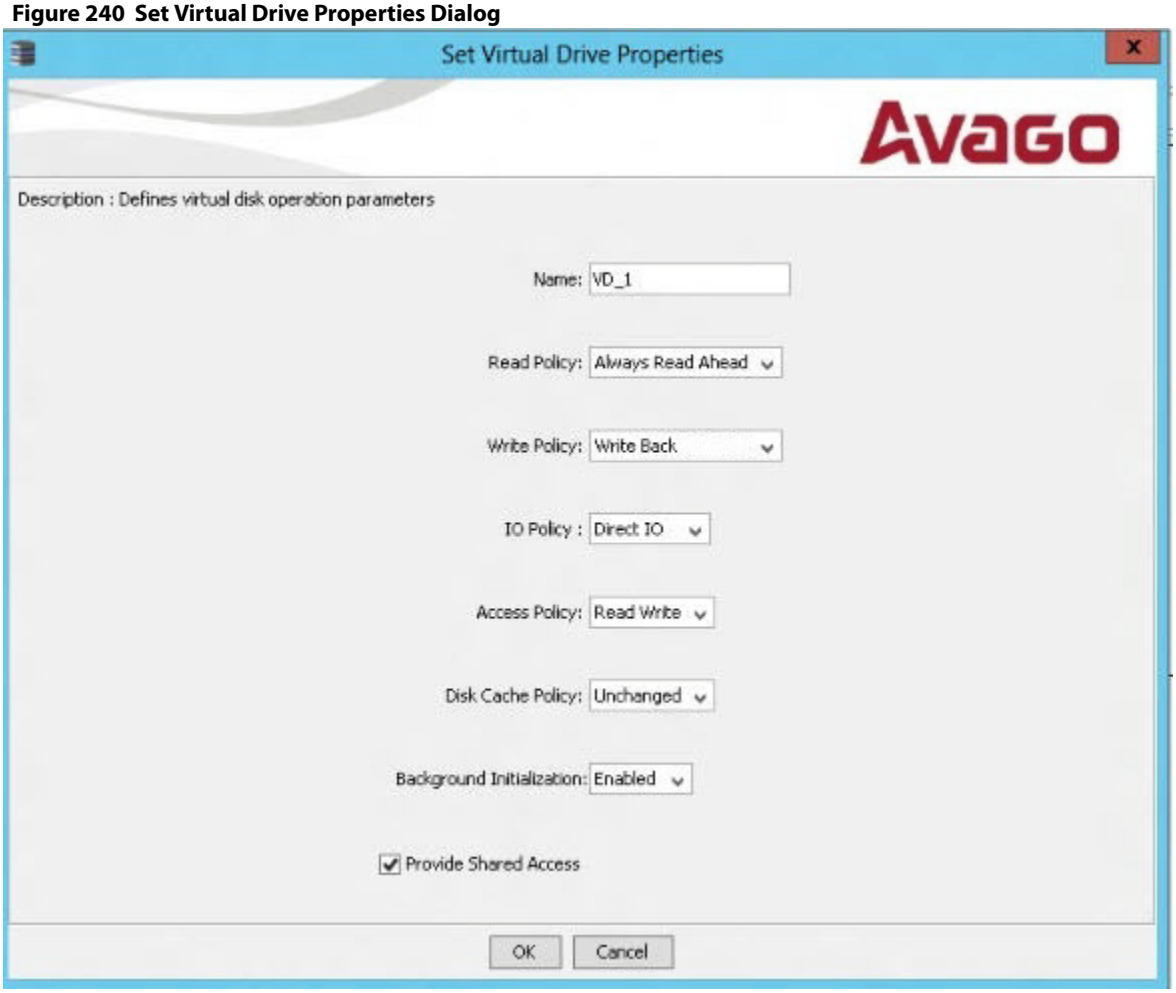

NOTE If the controller supports High Availability DAS, the **Provide Shared Access** check box appears in the above dialog. Select this option if you

want the virtual drive to be shared between the two servers in a cluster.

3. Change the virtual drive properties as needed.

   For information about these properties, see Selecting Virtual Drive Settings.

4. Click **OK** to accept the changes.

   The virtual drive settings are updated.

9.11.0.1.9 Changing a Virtual Drive Configuration

5. You can use the **Modify Drive Group** wizard in the MegaRAID Storage Manager software to change the configuration of a virtual drive by adding drives to the virtual drive, removing drives from it, or changing its RAID level.

   ATTENTION Be sure to back up the data on the virtual drive before you change its configuration.

   NOTE You cannot change the configuration of a RAID 10, RAID 50, or RAID 60 virtual drive. You cannot change a RAID 0, RAID 1, RAID 5, or RAID 6 configuration if two or more virtual drives are defined on a single drive group. (The Logical tab shows which drive groups and drives are used by each virtual drive.)

9.11.0.1.9.1 Accessing the Modify Drive Group Wizard

   NOTE The **Modify Drive Group** wizard was previously known as the **Reconstruction** wizard.

6. Perform the following steps to access the **Modify Drive Group** wizard options:
1. Click the **Logical** tab in the left panel of the MegaRAID Storage Manager main menu window.
2. Select a drive group in the left panel of the window.
3. Select **Go To > Drive Group > Modify Drive Group** on the menu bar, or right-click the virtual drive icon to access the **Modify Drive Group** wizard.

   The following warning appears about rebooting virtual drives containing boot partitions that are undergoing RAID level migration or capacity expansion operations. Back up your data before you proceed.

Figure 241 Reboot Warning Message

4. Select the **Confirm** check box, and click **Yes**.

   A warning to back up your data appears, as shown in the following figure.

Figure 242 Warning to Back Up Data Message

5. Select the **Confirm** check box, and click **Yes**.
   The **Modify Drive Group** wizard window appears, as shown in the following figure.

Figure 243 Modify Drive Group Wizard Window

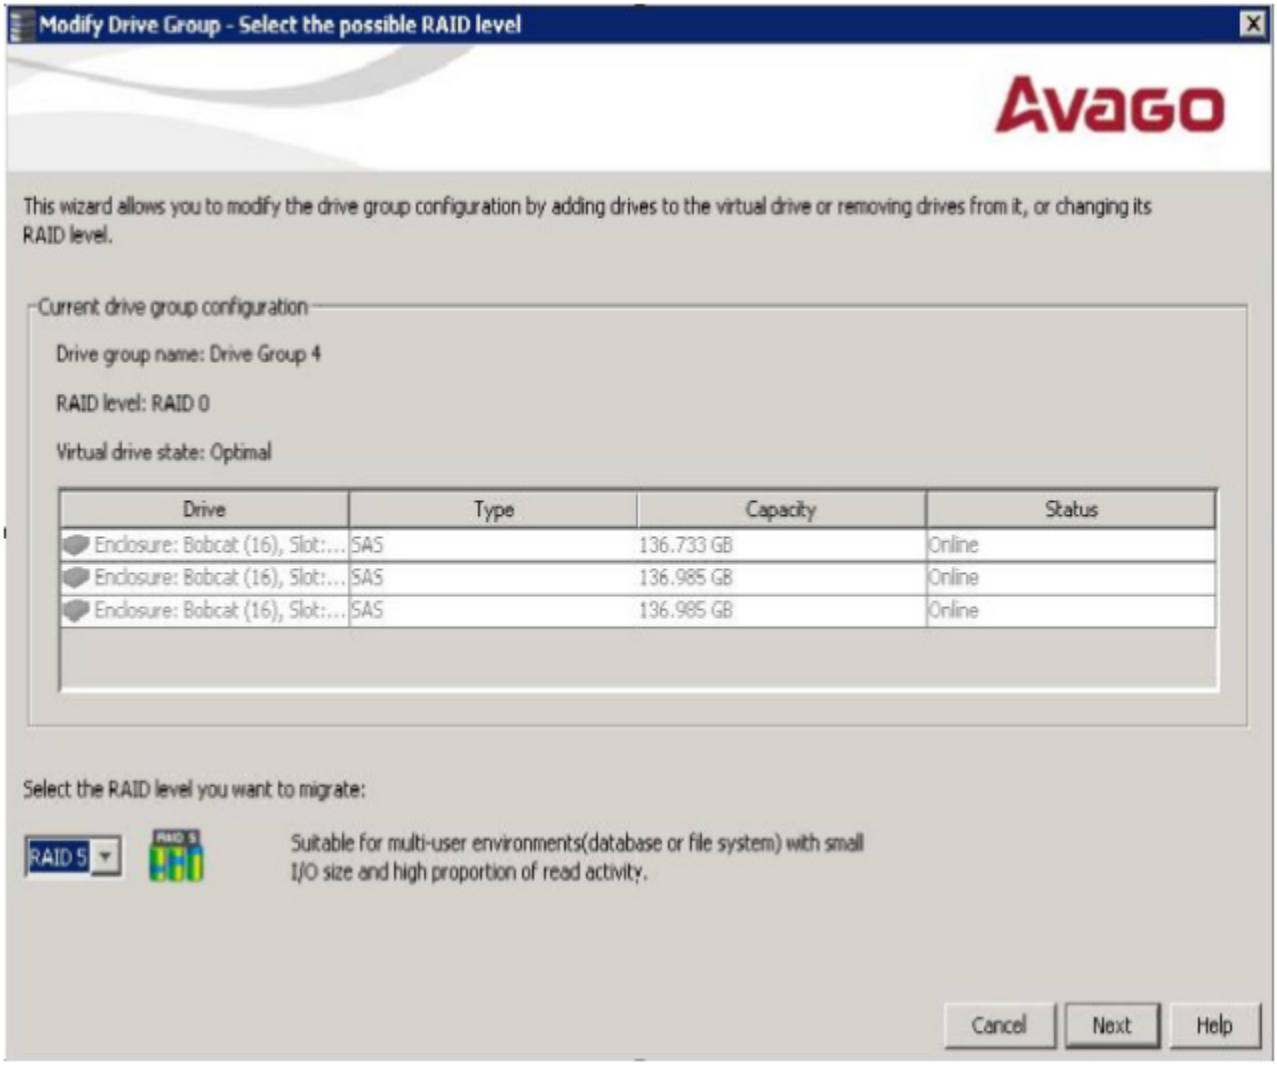

6. The following sections explain the **Modify Drive Group** wizard options.

9.11.0.1.9.2 Adding a Drive or Drives to a Configuration

ATTENTION Be sure to back up the data on the virtual drive before you add a drive to it.

7. Follow these steps to add a drive or drives to a configuration with the **Modify Drive Group** wizard.

1. Click the **Logical** tab in the left panel of the MegaRAID Storage Manager window.
2. Select a drive group in the left panel of the window.
3. Either select **Go To > Drive Group > Modify Drive Group** on the menu bar, or right-click the virtual drive icon to access the **Modify Drive Group** wizard.

   The **Modify Drive Group** wizard window appears.

Figure 244 Modify Drive Group - Select the possible RAID level Window

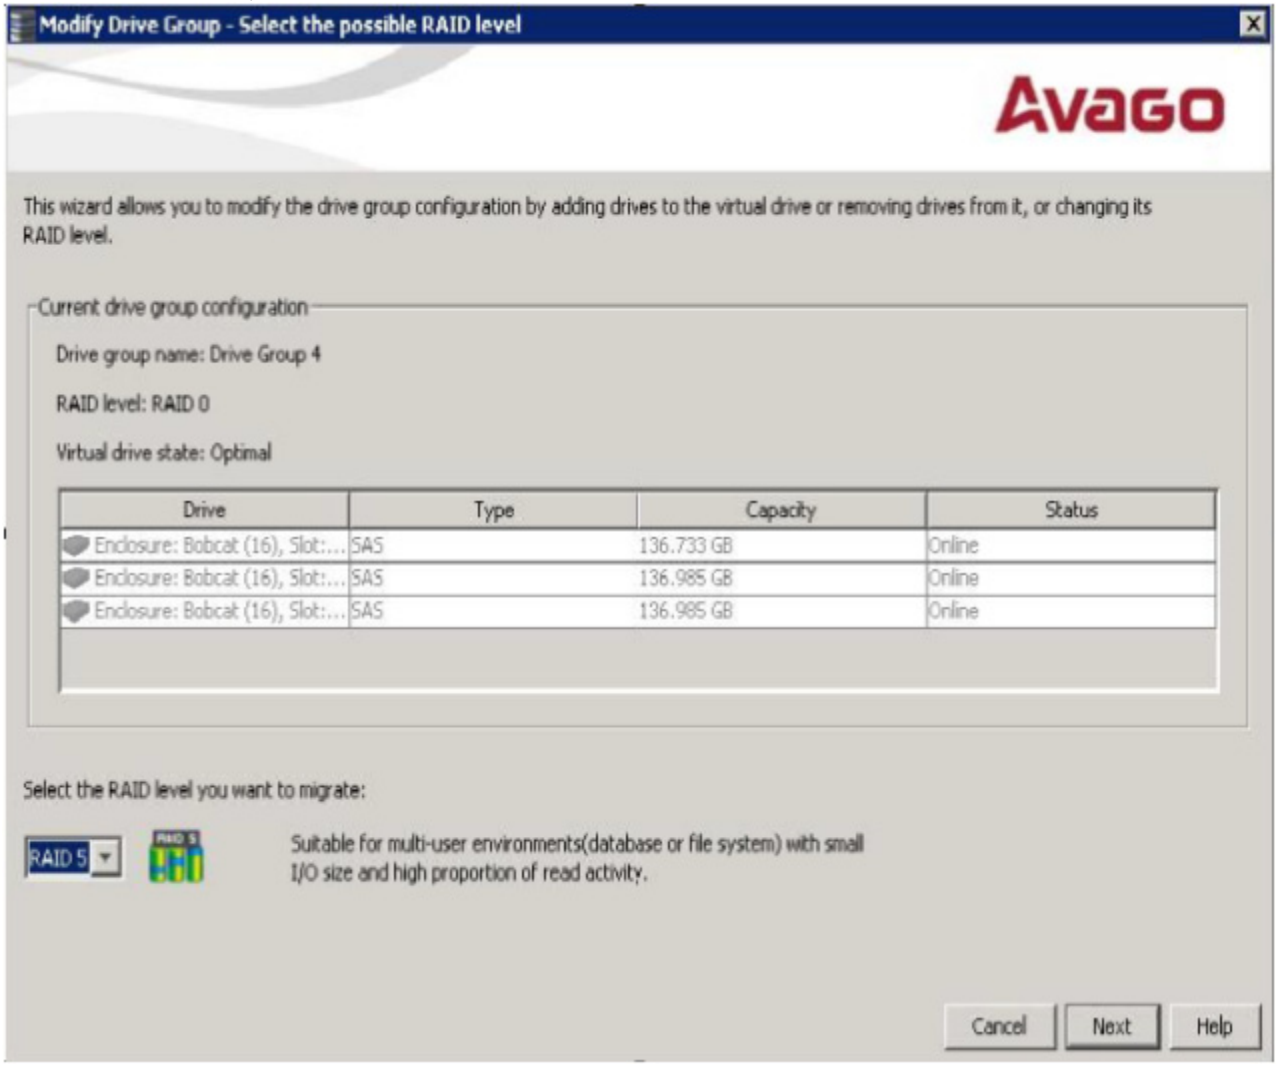

4. Select the RAID level to which you want to change ("migrate") the drive group, and click **Next**.

   The following window appears. It lists the drives you can add, and it states whether you have to add a minimum number of drives to change the RAID level from the current level to the new RAID level.

Figure 245 Modify Drive Group – Add drive to the current configuration Window

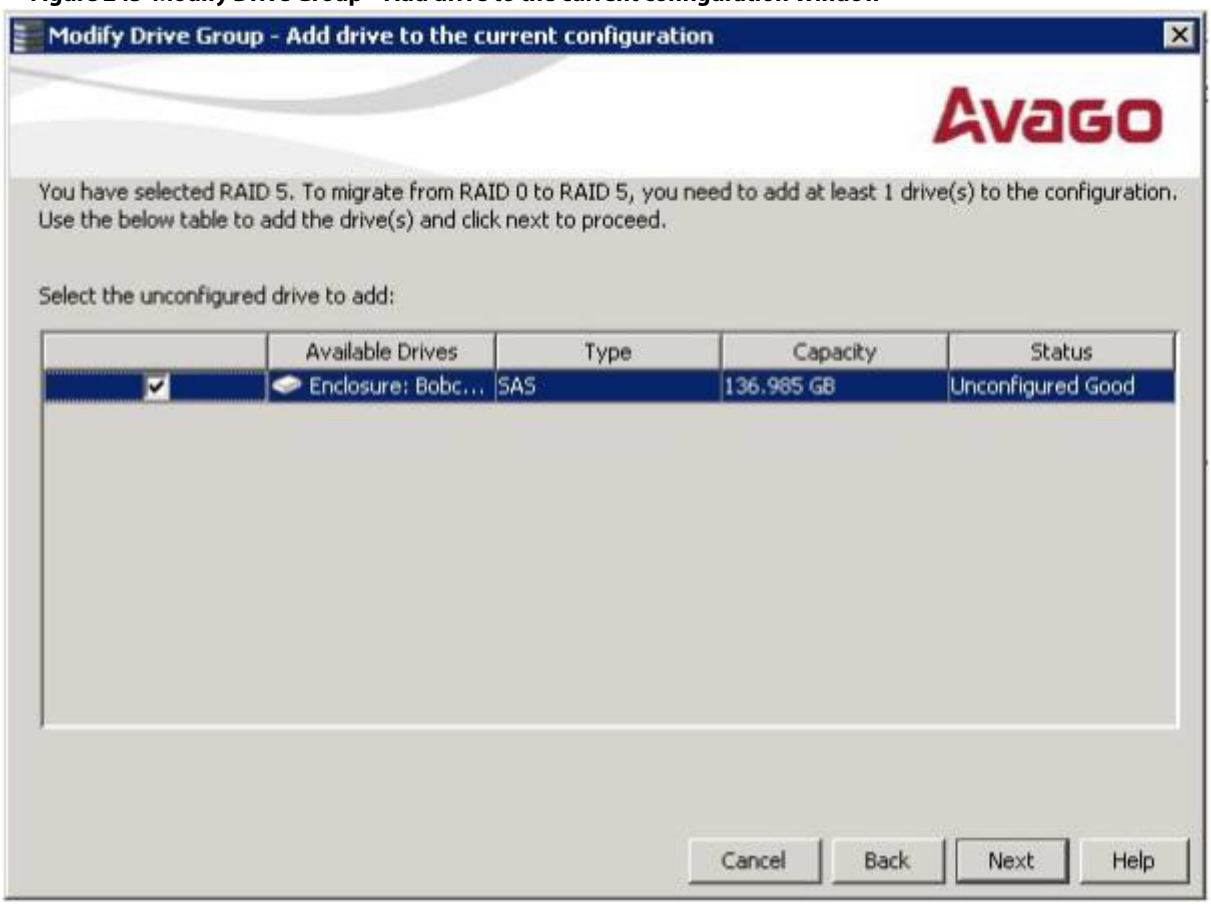

5. Click the check box next to any unconfigured drives that you want to add, and then click **Next**.

NOTE The drives you add must have the same capacity as or greater capacity than the drives already in the drive group, or you cannot change the RAID level.

The **Modify Drive Group - Summary** window appears. This window shows the current settings and what the settings will be after the drives are added.

Figure 246 Modify Drive Group - Summary Window

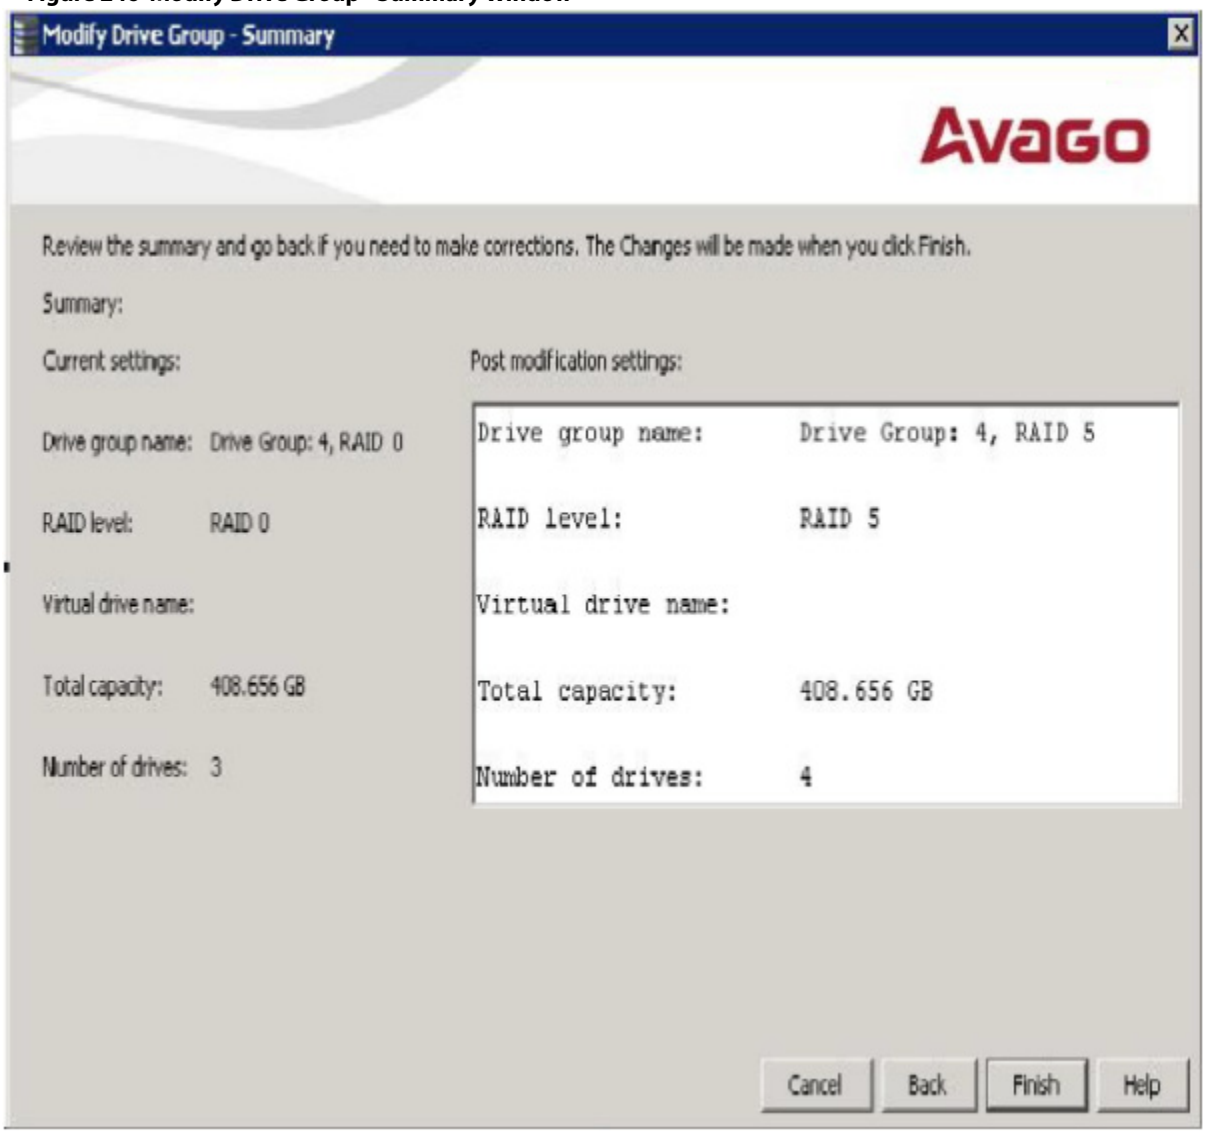

6. Review the configuration information.

   You can click **Back** if you need to change any selections.

7. Click **Finish** to accept the changes.

   A confirmation message appears. The message states that this operation cannot be aborted and asks whether you want to continue.

8. Click **Yes** to accept and complete the addition of the drives to the drive group.

9.11.0.1.9.3 Removing a Drive from a Configuration

ATTENTION Be sure to back up the data on the virtual drive before you remove a drive from it.

9. Follow these steps to remove a drive from a RAID 1, RAID 5, or RAID 6 configuration.

NOTE This option is not available for RAID 0 configurations.

1. Click the **Logical** tab in the left panel of the MegaRAID Storage Manager window.

2. Click a drive icon in the left panel of the window.
3. Either select **Go To > Physical Drive > Make Drive Offline** on the menu bar, or right-click the drive, and select **Make Drive Offline** from the menu.

   A confirmation message appears. The message states that this operation cannot be aborted and asks whether you want to continue.
4. Click **Yes** to accept and complete the removal of the drive from the drive group.

9.11.0.1.9.4 Replacing a Drive

ATTENTION Be sure to back up the data on the virtual drive before you replace a drive.

5. Follow these steps to add a replacement drive and copy the data from the drive that was removed to the replacement drive.

1. Click the **Logical** tab in the left panel of the MegaRAID Storage Manager window.
2. Select a drive in the left panel of the window.
3. Either select **Go To > Physical Drive > Replace Physical Drive** on the menu bar, or right-click the virtual drive icon to access the **Modify Drive Group** wizard.

   The dialog with the replacement drive appears, as shown in the following figure.

Figure 247 Drive Replacement Window

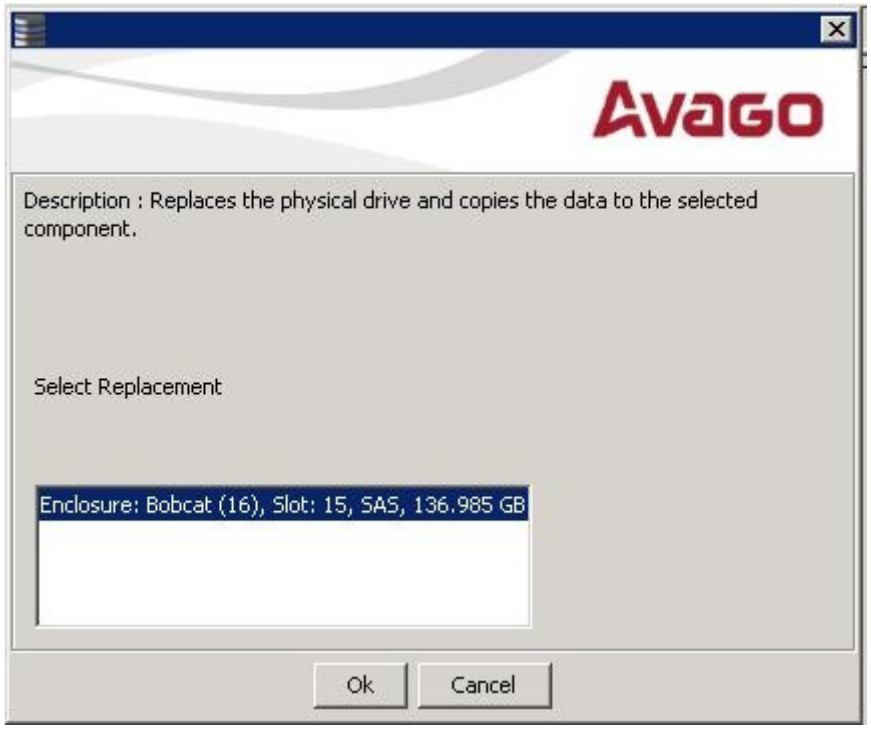

4. Select a replacement drive.

   A confirmation message appears.
5. Click **Yes**.

   This step replaces a drive and copies the data to the selected component.

9.11.0.1.9.5 Migrating the RAID Level of a Virtual Drive

6. 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.
7. When you migrate a virtual drive to another RAID level, you can keep the same number of drives, or you can add drives. In some cases, you have to add a certain number of drives to migrate the virtual drive from one RAID level to another. The window indicates the minimum number of drives you are required to add.

   ATTENTION Be sure to back up the data on the virtual drive before you change the RAID level.

8. Follow these steps to change the RAID level of the virtual drive with the **Modify Drive Group** wizard:

1. Click the **Logical** tab in the left panel of the MegaRAID Storage Manager window.
2. Select a drive group in the left panel of the window.
3. Either select **Go To > Drive Group > Modify Drive Group** on the menu bar, or right-click the virtual drive icon to access the **Modify Drive Group** wizard.

   The **Modify Drive Group** wizard appears.

Figure 248 Modify Drive Group - Select the possible RAID level Dialog

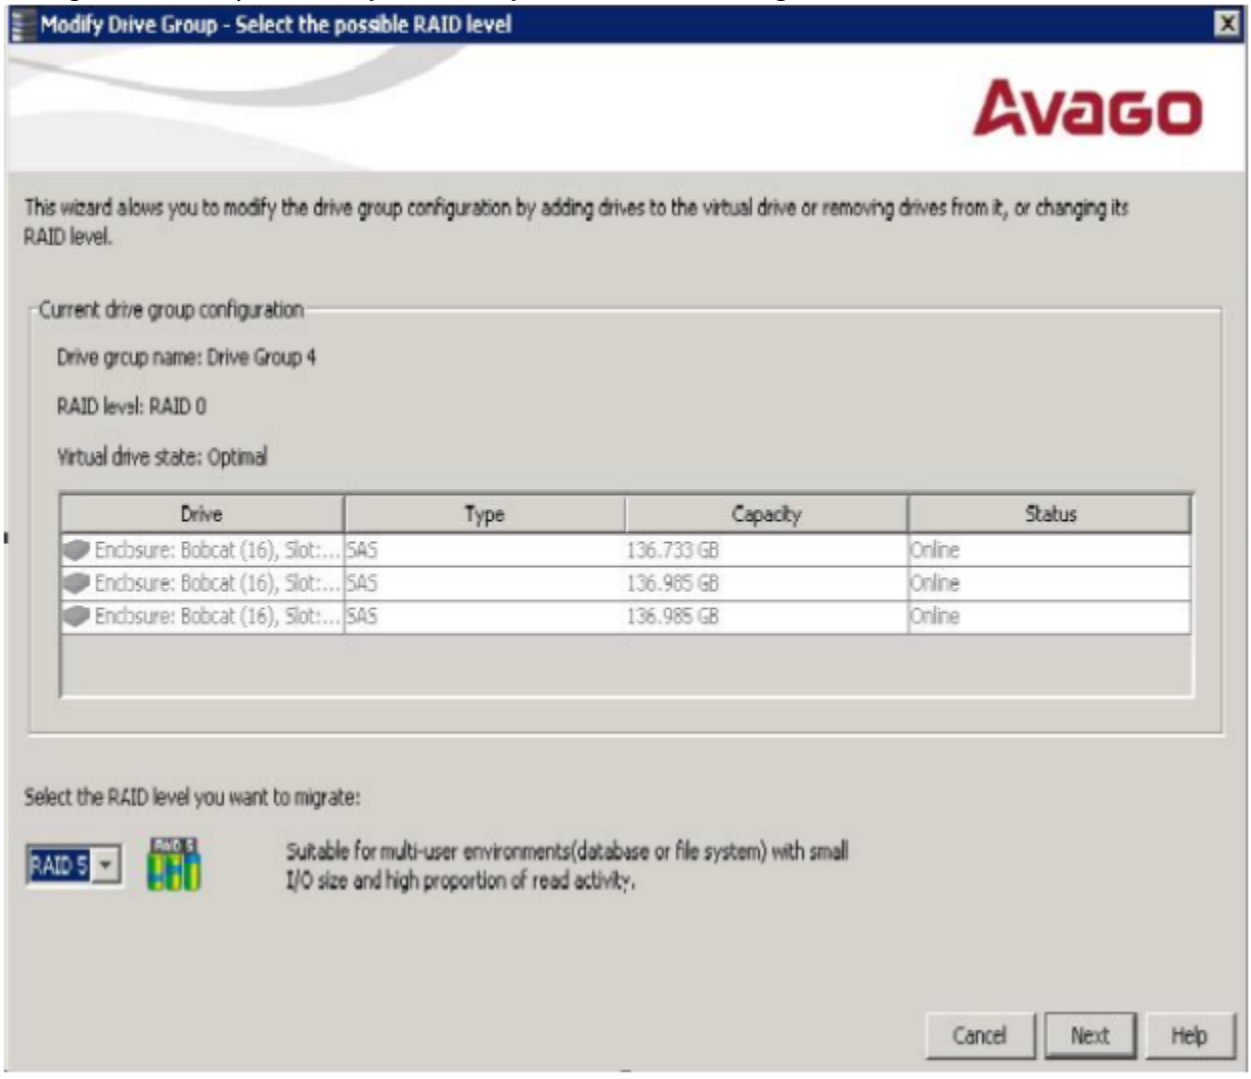

4. On the **Modify Drive Group - Select the possible RAID level** dialog, select the RAID level to which you want to change ("migrate") the drive group to, and click **Next**.

   The following dialog appears. The dialog states the number of drives that you have to add to change the RAID level from the current level to a new RAID level that requires more drives.

Figure 249 Modify Drive Group - Add drive to the current configuration Dialog

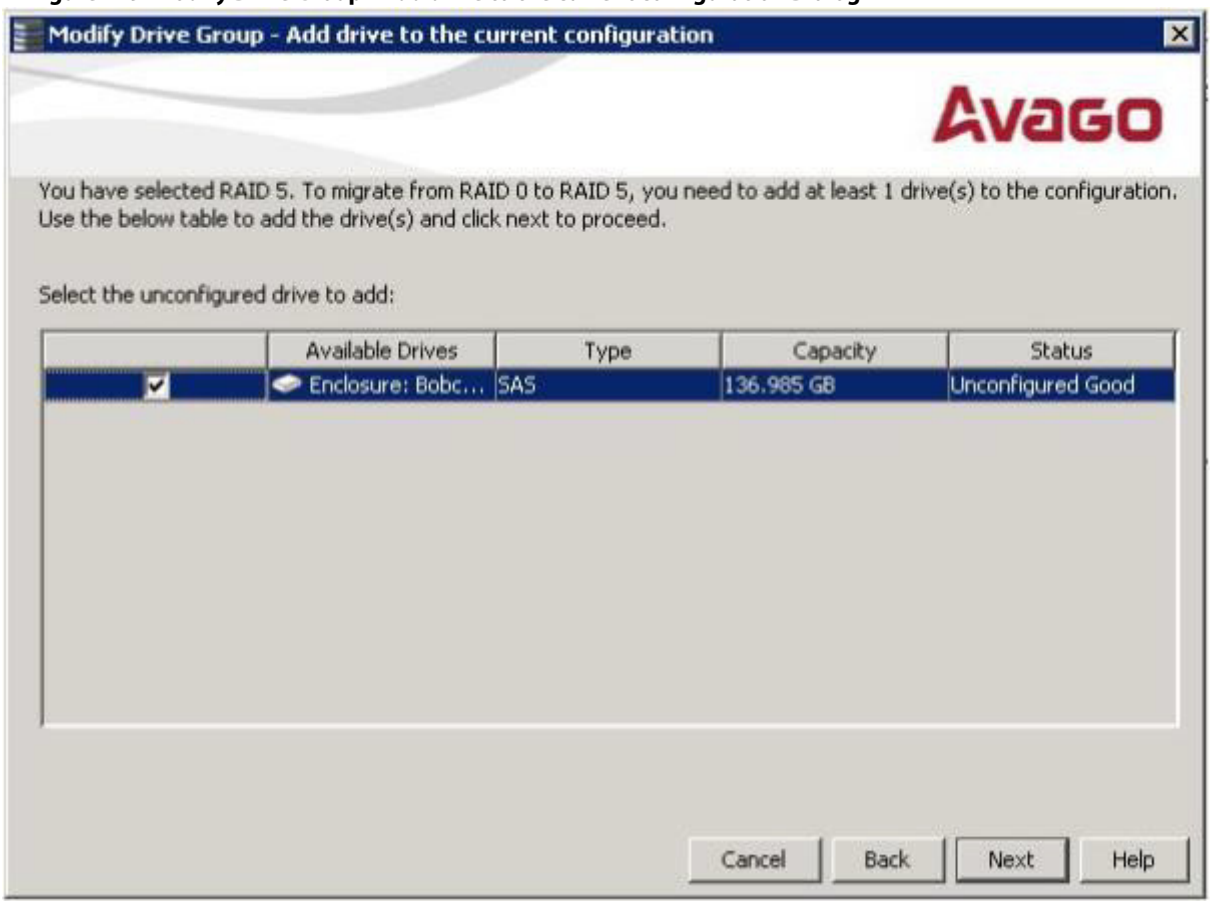

5. Select the unconfigured drive or drives to add, and click **Next**.

   NOTE The drives you add must have the same capacity as or greater capacity than the drives already in the drive group, or you cannot change the RAID level.

   The **Modify Drive Group – Summary** window appears. This window shows the current settings and what the settings will be after the drives are added.

Figure 250 Modify Drive Group - Summary Window

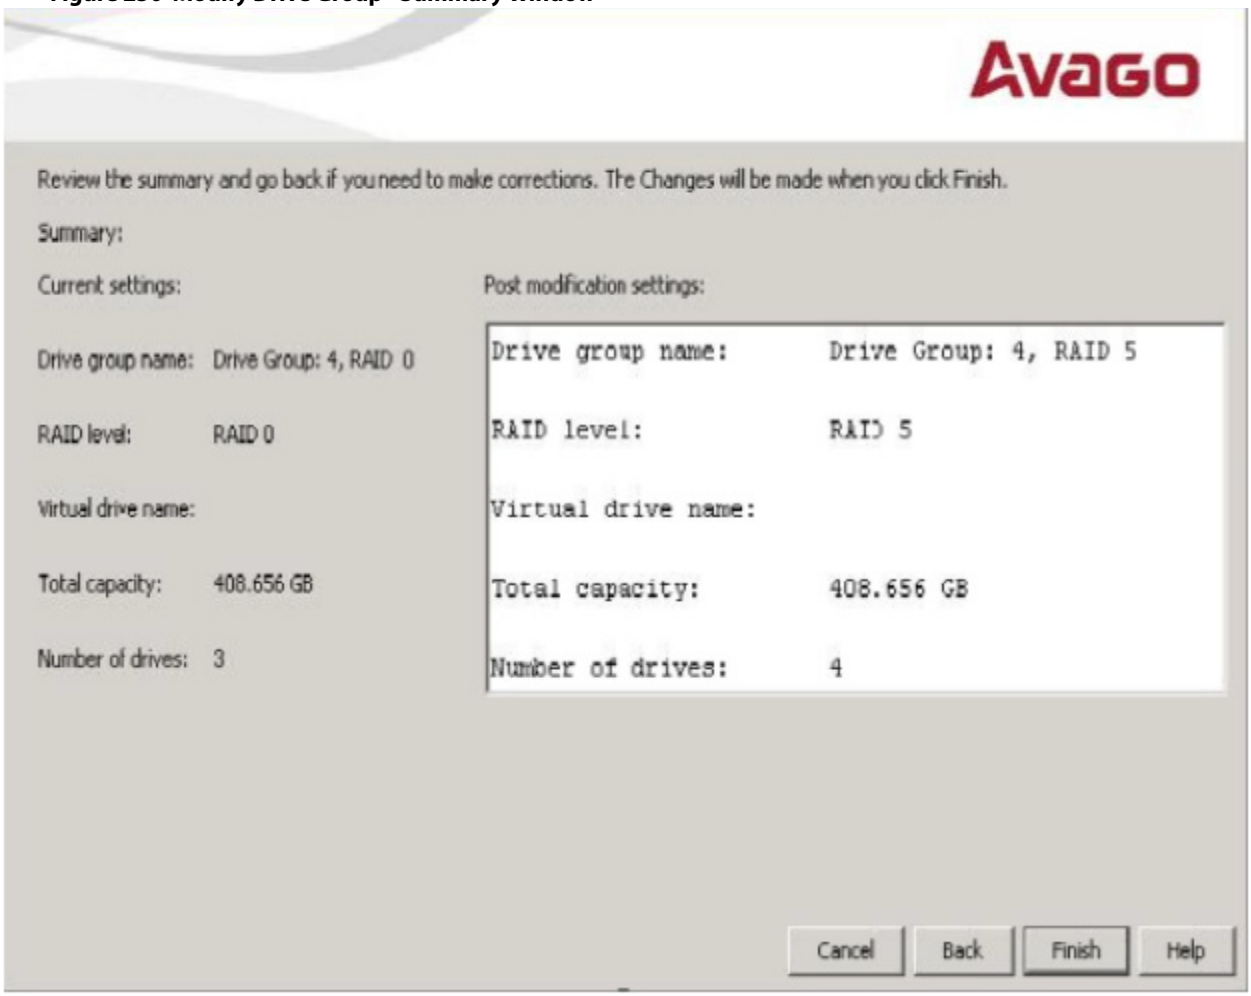

6. Review the configuration information.

   You can click **Back** if you need to change any selections.
7. Click **Finish** to accept the changes.

   A confirmation message appears. The message states that this operation cannot be aborted and asks whether you want to continue.
8. Click **Yes** to accept and complete the 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.

9.11.0.1.10 Deleting a Virtual Drive

ATTENTION Make sure to back up the data that is on the virtual drive before you delete it. Make sure that the operating system is not installed on this virtual drive.

9. You can delete virtual drives to rearrange the storage space. To delete a virtual drive, follow these steps.
1. Back up all user data that is on the virtual drive you want to delete.
2. On the **MegaRAID Storage Manager** window, select the **Logical** tab, and click the icon of the virtual drive you want to delete.

3. Select **Go To > Virtual Drive > Delete Virtual Drive**.
4. When the warning messages appear, click **Yes** to confirm that you want to delete the virtual drive.

NOTE You are asked twice if you want to delete a virtual disk to avoid deleting the virtual disk by mistake.

9.11.0.1.11 Performing a Join Mirror Operation

5. You can perform a join mirror operation on a drive group to continue using the modified virtual drive or to reuse the original virtual drive.
6. Follow these steps to perform a join mirror operation:
1. Go to the Logical tab in the MegaRAID Storage Manager window.
2. Right click on the drive group on which you want to perform the join mirror operation and select **Join Mirror**.

   The **Join Mirror - Choose Option** dialog appears.

Figure 251 Join Mirror - Choose Option Dialog

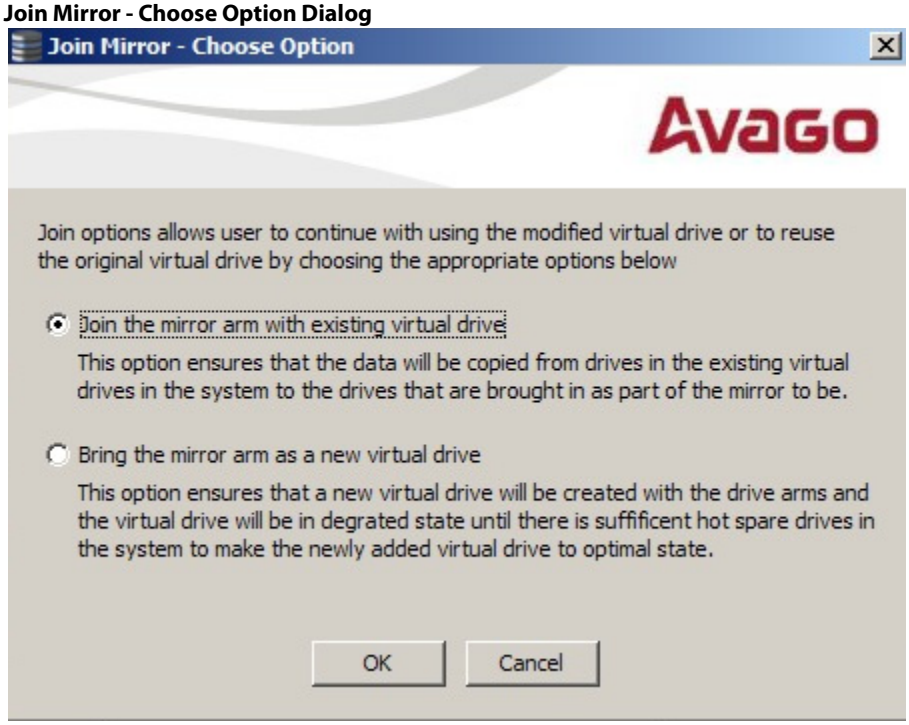

3. Select one of the two options and click **OK**.

   If you select **Join the mirror arm with existing virtual drive**, the following dialog appears.

Figure 252 Confirmation Message Dialog

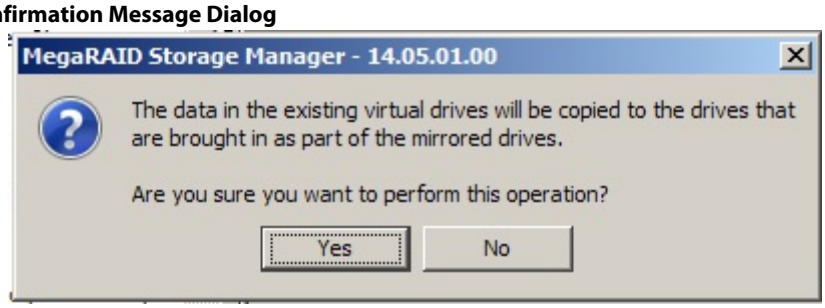

If you select **Bring the mirror arm as a new virtual drive**, the following dialog appears.

Figure 253 Confirmation Message Dialog

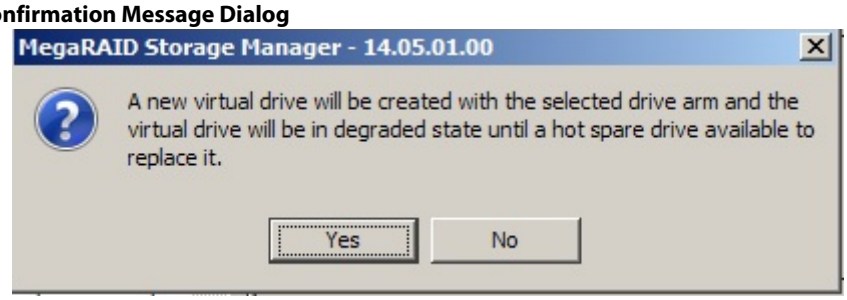

4. Click **Yes** to proceed with the operation.

9.12 Hiding and Unhiding a Virtual Drive or a Drive Group

You can hide or unhide either a virtual drive or a drive group on the controller. After the hide/unhide operation is performed successfully and the access policy is defaulted to Hidden, you will not have any other access policy options such as ReadWrite, ReadOnly, and so on available to you.

9.12.1 Hiding a Virtual Drive

You can hide a virtual drive on the controller.

Perform the following steps to hide a virtual drive:

1. Go to the **Logical tab** in the MegaRAID Storage Manager window.
2. Select a virtual drive that you want to hide.
3. Right-click the selected virtual drive select **Hide**.
   A message box appears, asking you to confirm the operation.
4. Select the **Confirm** checkbox and click **OK** to hide the virtual drive.

9.12.2 Unhiding a Virtual Drive

You can unhide a virtual drive on the controller.

Perform the following steps to unhide a virtual drive:

1. Go to the **Logical tab** in the MegaRAID Storage Manager window.
2. Select a virtual drive that you want to unhide.
3. Right-click the selected virtual drive select **Unhide**.
   A message box appears, asking you to confirm the operation.
4. Select the **Confirm** checkbox and click **OK** to unhide the virtual drive.

9.12.3 Hiding a Drive Group

You can hide a drive group on the controller. If you hide a drive group, all of the virtual drives that are a part of this drive group become hidden.

Perform the following steps to hide a drive group:

1. Go to the **Logical** tab in the MegaRAID Storage Manager window.
2. Select a drive group that you want to hide.
3. Right-click the selected drive group and select **Hide All Virtual Drives**.

   A message box appears, asking you to confirm the operation.
4. Select the **Confirm** checkbox and click **OK** to hide the drive group.

9.12.4 Unhiding a Drive Group

You can unhide a drive group on the controller. If you unhide a drive group, all of the virtual drives that are a part of this drive group become unhidden.

Perform the following steps to unhide a drive group:

1. Go to the **Logical** tab in the MegaRAID Storage Manager window.
2. Select a drive group that you want to unhide.
3. Right-click the selected drive group and select **Unhide All Virtual Drives**.

   A message box appears, asking you to confirm the operation.
4. Select the **Confirm** checkbox and click **OK** to Unhide the drive group.

9.13 Performing a Transport Ready Operation

The transport ready operation allows you to move a drive group and all the shared or exclusive dedicated hot spare drives associated with that drive group from one server to another server non-disruptively, in a real time.

NOTE To know more about the Transport Ready feature, see Transport Ready Pre-requisites section.

Follow these steps to perform a transport ready operation:

1. Go to the **Logical** tab in the MegaRAID Storage Manager window.
2. Right-click the drive group on which you want to perform the transport ready operation and select **Transport Ready**.

Figure 254 Transport Ready

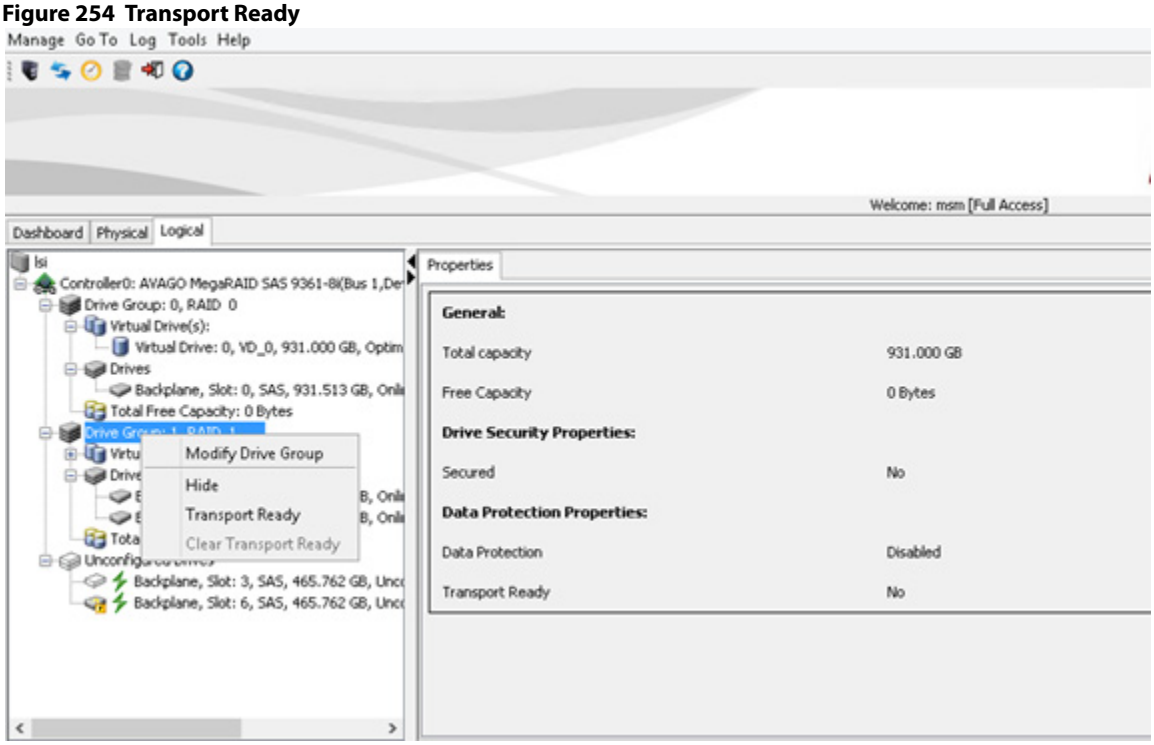

A warning message is displayed requesting you to confirm your action.

3. Select the **Confirm** check box and click **Yes**.

   If you have any shared dedicated host spare drives associated with the drive group on which you are performing the Transport Ready operation, a warning message is displayed requesting you to confirm your action.

Figure 255 Confirm Transport Ready Dialog

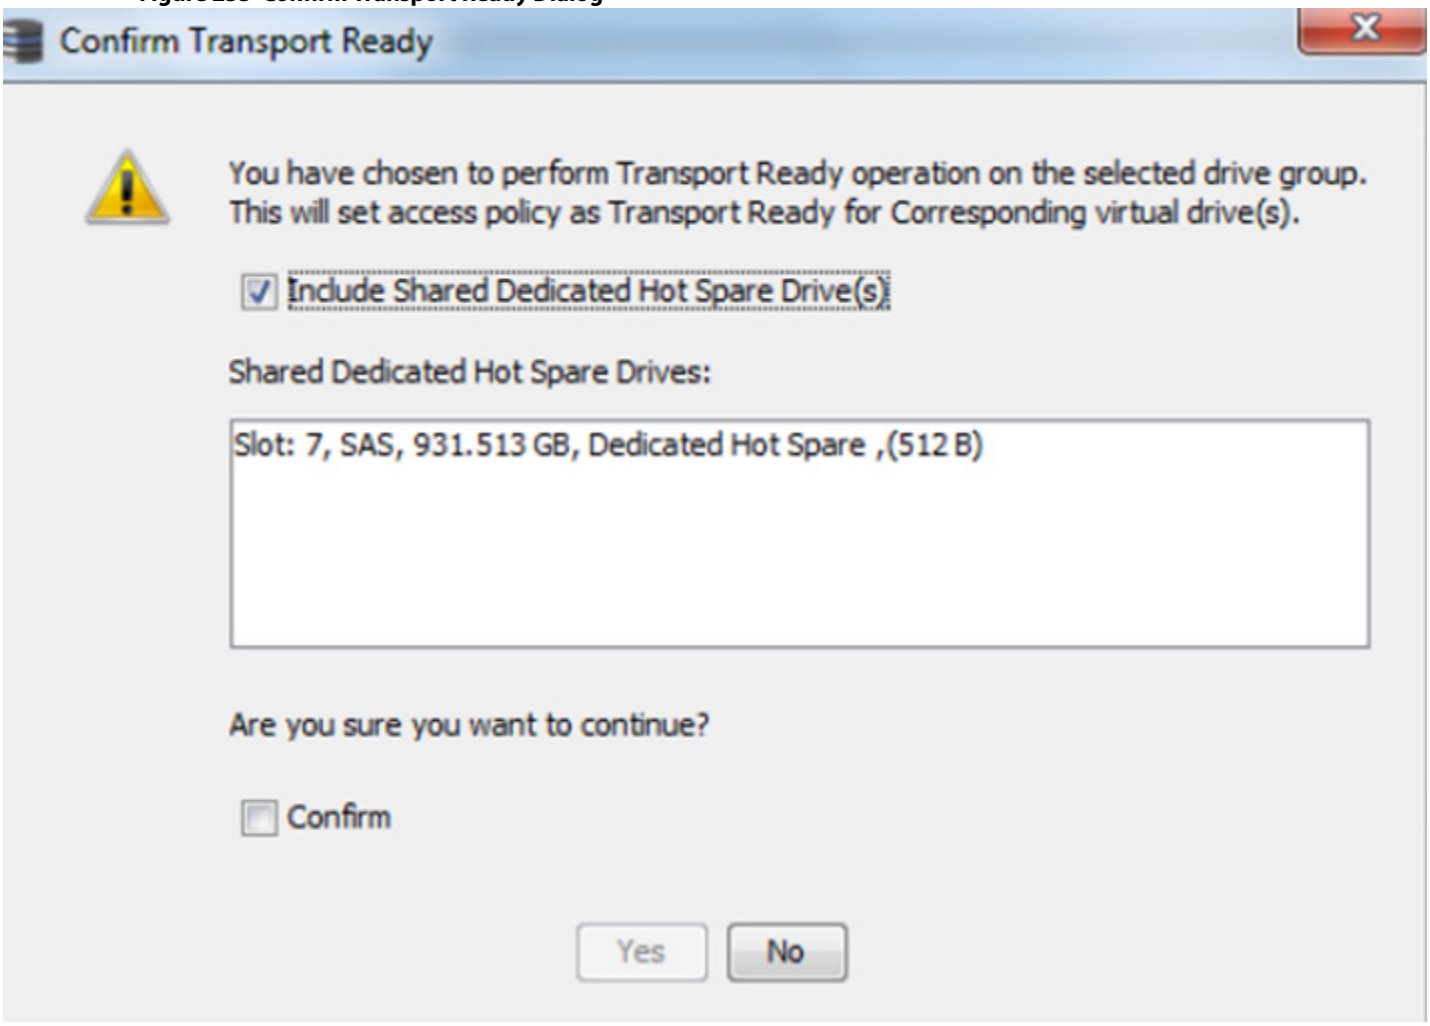

a. Select the **Include Shared Dedicated Hot Spare Drive(s)** check box.
b. Select the **Confirm** check box.
c. Click **Yes**.

If you have any `exclusive dedicated host spare drives` associated with the drive group on which you are performing the Transport Ready operation, a warning message is displayed requesting you to confirm your action.

Figure 256 Confirm Transport Ready Dialog

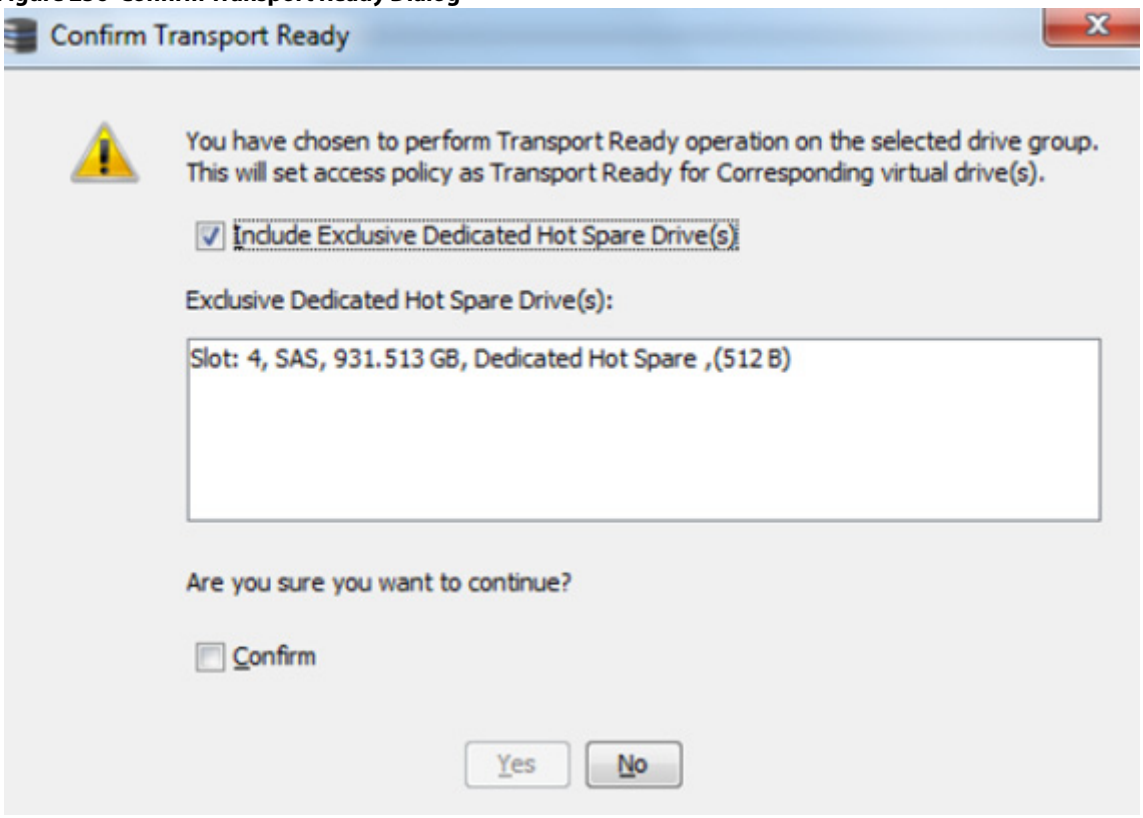

a. Select the **Include Exclusive Dedicated Hot Spare Drive(s)** check box.
b. Select the **Confirm** check box.
c. Click **Yes**.

If you have both exclusive and shared dedicated host spare drives associated with the drive group on which you are performing the Transport Ready operation, a warning message is displayed requesting you to confirm your action.

Figure 257 Confirm Transport Ready Dialog

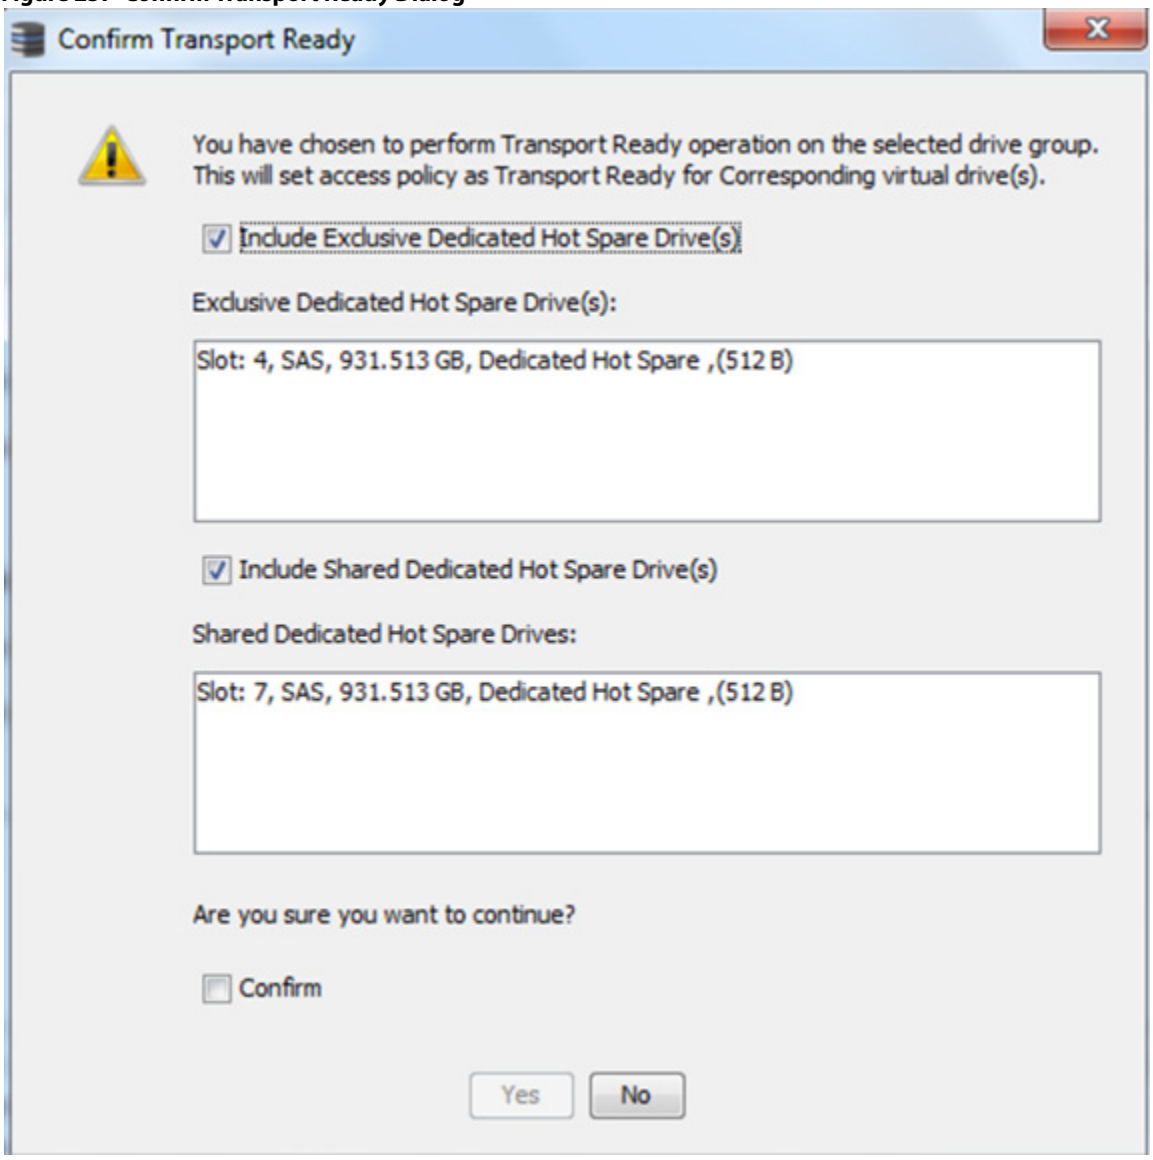

The drive group on which you have performed the Transport Ready operation is appended, and a Transport Ready property is displayed on the drive group. Once the Transport Ready property is displayed on the drive group, you can remove the drive group.

Figure 258 Transport Ready Property Display

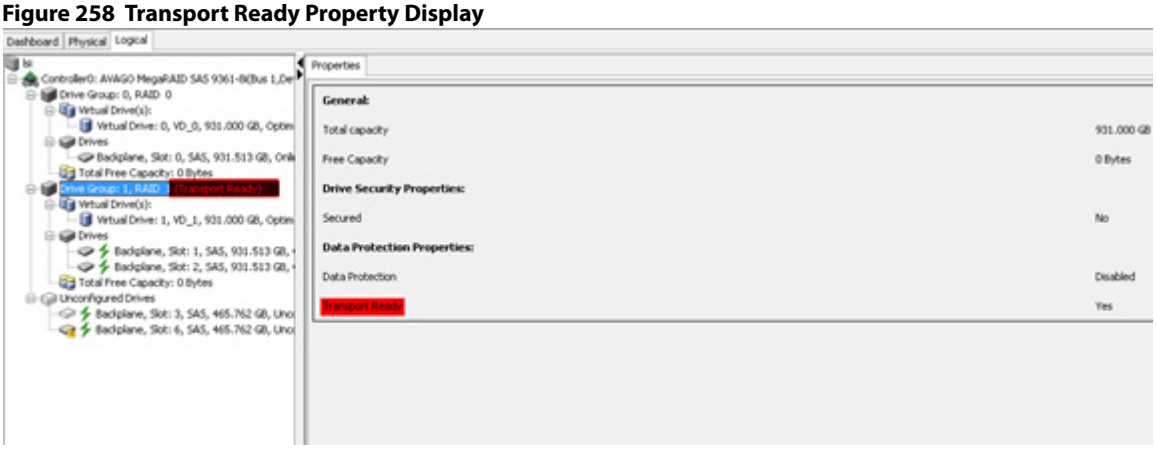

After the Transport Ready operation is performed successfully, the access policy of the virtual drive associated with the drive group on which the Transport Ready operation has been performed is defaulted to Transport Ready. You will not have any other access policy options (such as Read Write, Read Only, and so on) available, as shown in the following figure (in the following figure, some of the highlighted areas are for illustration purposes only).

Figure 259 Transport Ready - VD Access Policy

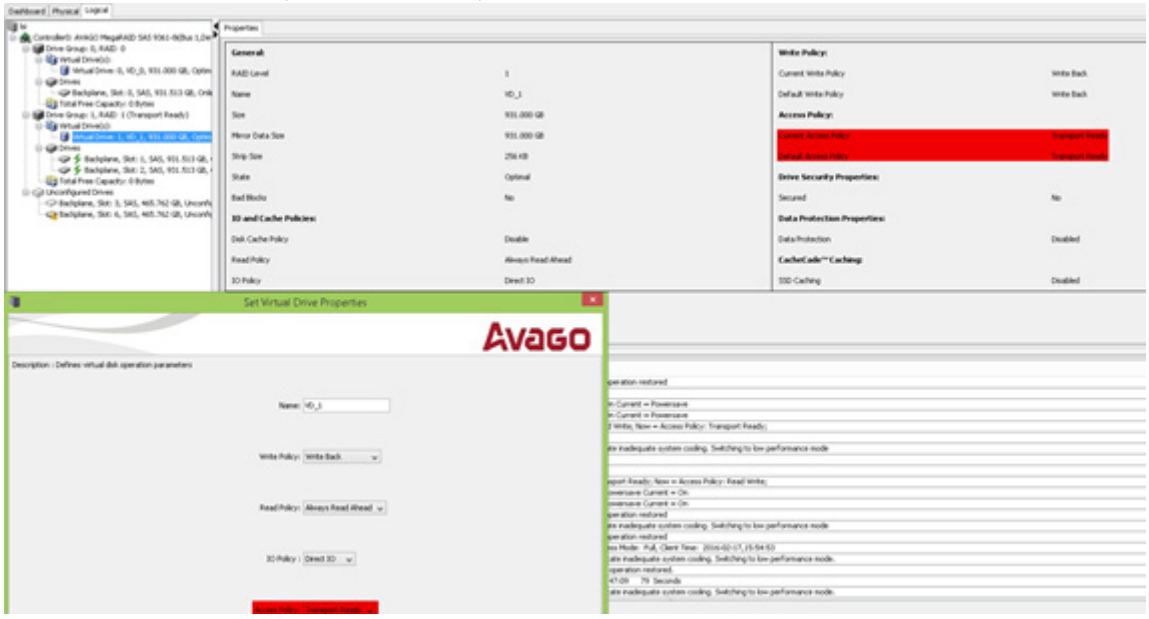

NOTE The firmware preserves the last virtual drive state for a Transport Ready VD. Inserting or removing a physical drive does not change the state of the VD.

For example, assume that you have created a RAID 5 virtual drive using 3 physical drives; the VD state is optimal; and you have performed a Transport Ready operation on that drive group. After the Transport Ready operation is completed, you remove one physical drive from the RAID 5 virtual drive. In this scenario, virtual drive will still be in `Optimal` state, not in the Degraded state.

The MegaRAID Storage Manager allows you to assign dedicated hot spares per span for spanned drive groups.

9.13.1 Clearing a Transport Ready Operation

You have performed the Transport Ready operation on a drive group and if you want to reverse it (undo your previous action), you can do so using the Clear Transport Ready option.

Follow these steps to clear a Transport Ready operation:

1. Select the drive group on which you have performed the Transport Ready operation.
2. Right-click and select **Clear Transport Ready**.

Figure 260 Clearing Transport Ready

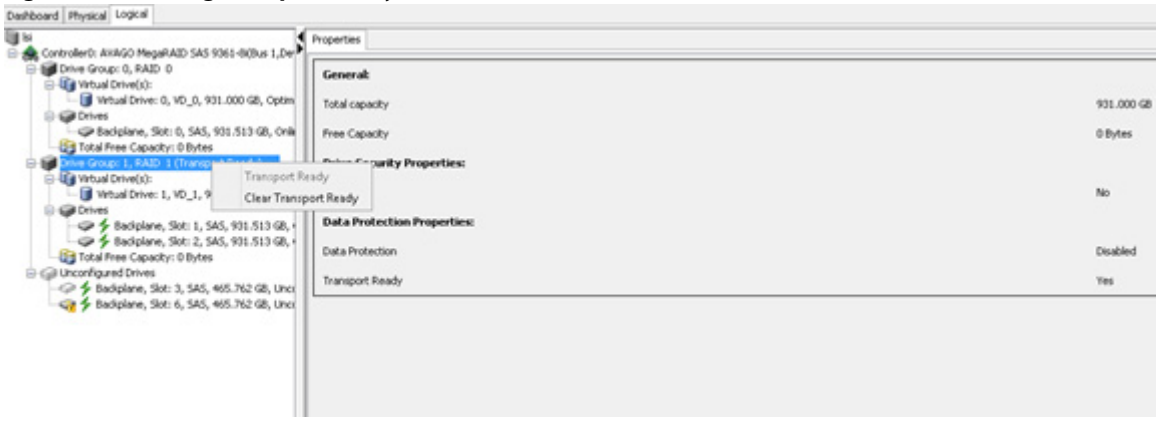

A warning message is displayed requesting you to confirm your action.

3. Select the **Confirm** check box and click **Yes**.

Chapter 10: Monitoring Controllers and Their Attached Devices

This chapter explains how to use the MegaRAID Storage Manager software to monitor the status of drives, virtual drives, and other storage devices.

The MegaRAID Storage Manager software enables you to monitor the activity of all the controllers present in the system and the devices attached to them.

The MegaRAID Storage Manager software does a background check every one hour to verify if the controller and the system time are in synch. If the time difference between the controller and the system is more than 90 seconds, the MegaRAID Storage Manager software synchronizes the time so that the controller time and the system time are in sync.

When you perform an operation on devices (such as the creation of a new virtual drive) or when devices automatically go from an optimal state to a different state (such as a created virtual drive goes to a degraded state or a Battery Backup Unit goes bad), the MegaRAID Storage Manager software gets those events from the controller and gives a notification to you, using different alert delivery methods.

10.1 Alert Delivery Methods

Based on the severity level (Information, Warning, Critical and Fatal), the default alert delivery methods change. By default, each severity level has one or more alert delivery methods configured for it, as shown in the following table. To modify these alert delivery methods, see Configuring Alert Notifications. The different alert delivery methods are as follows:

- Vivaldi Log/MegaRAID Storage Manager Log
- System Log
- Pop-up Notification
- Email Notification

Table 60 Severity Level and Default Alert Delivery Methods

Severity Level	Default Alert Delivery Method	Meaning
Information	Vivaldi log/MegaRAID Storage Manager log and System log	Informational message. No user action is necessary.
Warning	Vivaldi log/MegaRAID Storage Manager log and System log	Some component might be close to a failure point.
Critical	Vivaldi log/MegaRAID Storage Manager log, System log, and Popup Notification	A component has failed, but the system has not lost data.
Fatal	Vivaldi log/MegaRAID Storage Manager log, System log, Popup Notification, and Email Notification	A component has failed, and data loss has occurred or will occur.

10.1.1 Vivaldi Log/MegaRAID Storage Manager Log

By default, all the severity events appear in the Vivaldi log/MegaRAID Storage Manager log and are displayed at the bottom of the **MegaRAID Storage Manager** main menu window. Each message that appears in this log has a severity level that indicates the importance of the event (severity), a date and timestamp (when it occurred), and a brief description, as show in the following figure.

Figure 261 Vivaldi Log

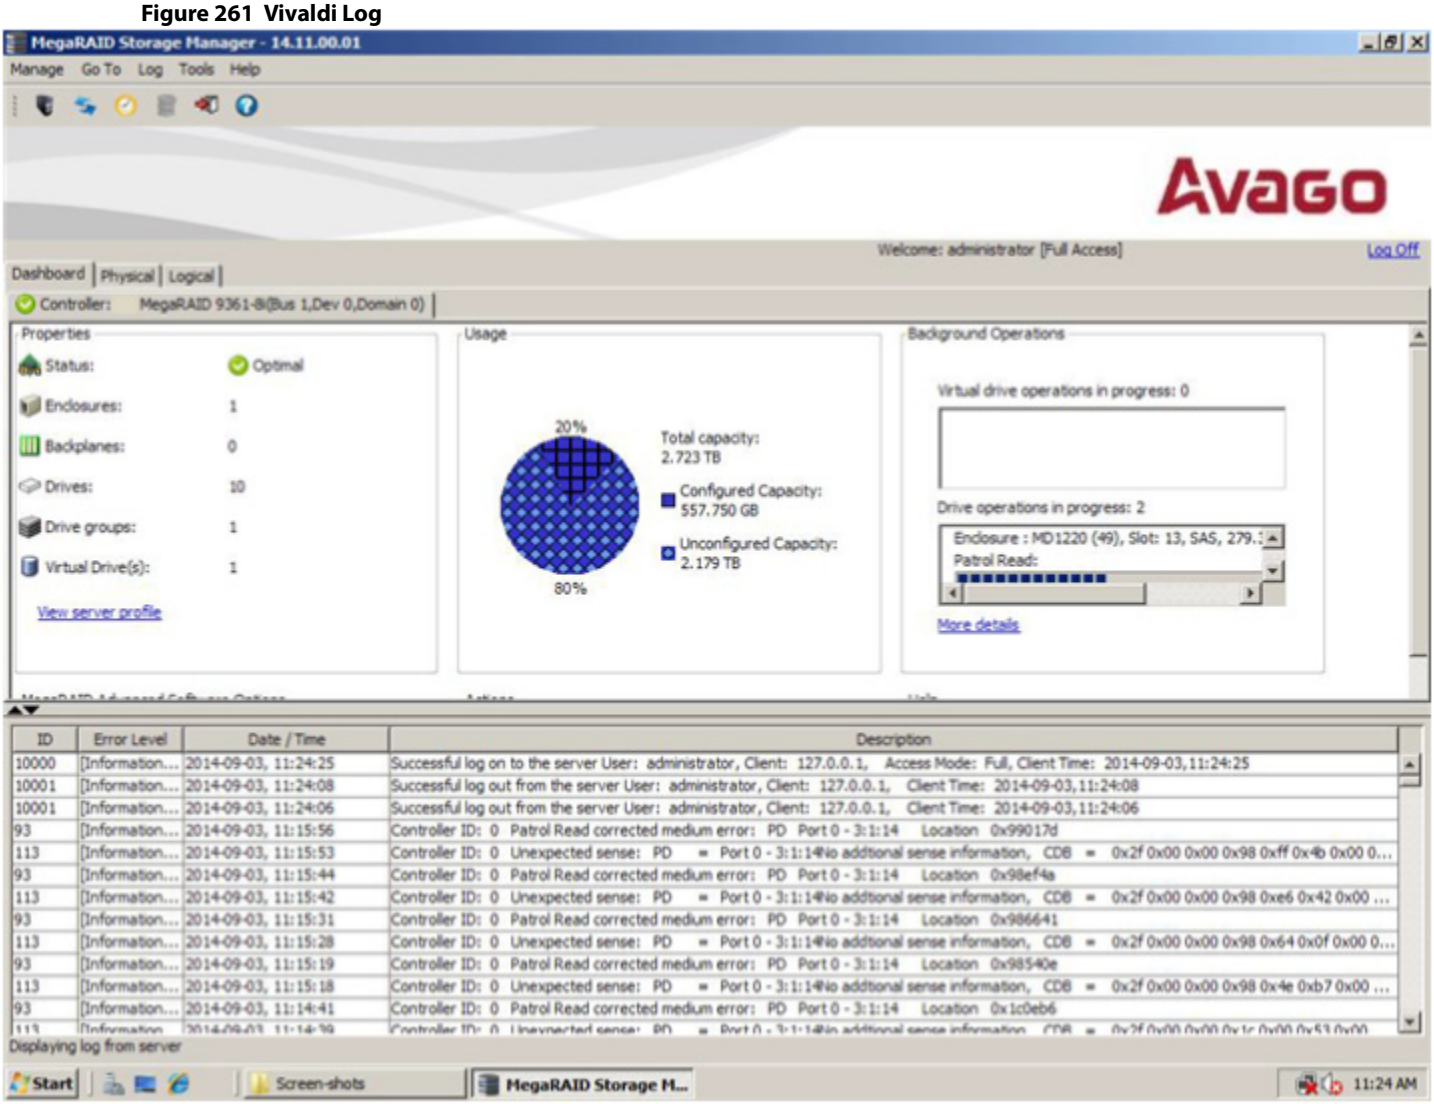

The following events appear in the log when the MegaRAID Storage Manager application is connected to the server.

- Successful log on to the server.
- Successful log out from the server.
- Server log cleared.
- Full access denied on the server.

You can double click on an event to display the same information in a separate window. The status bar at the bottom of the screen indicates whether the log is a MegaRAID Storage Manager server log or a locally stored log file.

When a Vivaldi log/MegaRAID Storage Manager log appears, the **Log** menu has the following options:

- **Save Log**

  Saves the current log to a `.log` file.
- **Save Log Text**

  Saves the current log in `.txt` format.
- **Load**

  Enables you to load a local `.log` file in the bottom of the **MegaRAID Storage Manager** main menu window. If you select the **Load** menu, you will not be able to view the current log.

- **Rollback to Current Log**

  This menu appears if we have loaded the logs from a local `.log` file. When you select this menu, you can view the current log.

- **Clear Log**

  Clears the current log information, if you have full access (versus view-only access). You have the option to save the log first.

10.1.2 System Log

By default, all the severity events are logged in the local system log (`syslog`). Based on the operating system you are using, the system log is logged in the following `syslog` locations:

- In Windows, the system log is logged in **Event Viewer > Application**.
- In Linux, the system log is logged in `/var/log/messages`.
- In Solaris, the system log is logged in `/var/adm/messages`.

10.1.2.1 Setting Up the Custom Facility Level in the System Log File for the Solaris x86 Operating System

In the Solaris operating system, the MegaRAID Storage Manager software logs the system messages in the `/var/adm/messages` directory with the facility level, `LOG_USER` by default. You can edit the `config-current.xml` file to specify a custom facility level to log the system messages. Follow these steps to edit the `config-current.xml` file:

1. Run the `./popup stop` command from the from *<MSM_HOME>*`\MegaPopup` directory to stop the Pop-Up Process.
2. Run the `svcadm disable -t MSMFramework` command to stop the MegaRAID Storage Manager Framework service.
3. Edit the `config-current.xml` file in the *<MSM_HOME>*`\MegaMonitor` directory to set the custom facility level.

   An example follows:

```
Default: Log level - LOG_USER----------------------------------
<systemlog>    <facility-level>8</facility-level></systemlog>
Edit: Log level - LOG_UUCP-----------------------------
<systemlog>    <facility-level>64</facility-level></systemlog>
```

4. Run the `svcadm enable MSMFramework` command to start the MegaRAID Storage Manager Framework service.
5. Run the `./popup start` command from the from *<MSM_HOME>*`\MegaPopup` directory to start the Pop-Up Process.

Available Custom Facility Levels

Choose any one of the following facility level while customizing the system log. The default facility level is `LOG_USER 8`.

- `LOG_USER 8`
- `LOG_MAIL 16`
- `LOG_DAEMON 24`
- `LOG_AUTH 32`
- `LOG_SYSLOG 40`
- `LOG_LPR 48`

- LOG_NEWS 56
- LOG_UUCP 64
- LOG_LOCAL0 128
- LOG_LOCAL1 136
- LOG_LOCAL2 144
- LOG_LOCAL3 152
- LOG_LOCAL4 160
- LOG_LOCAL5 168
- LOG_LOCAL6 176
- LOG_LOCAL7 184

10.1.3 Pop-up Notification

By default, fatal and critical events are displaying in a pop-up notification. A pop-up Notification is started automatically when you login to the operating system. Through this feature, you can view multiple events in a single pop-up window as shown in the following figure.

If the MegaRAID Storage Manager Framework connects to a VMware ESXi server, an additional read only field **Event From** appears in the following dialog (next to the **Controller ID** field) showing the IP address of the VMware ESXi server.

Figure 262 Pop-up Notification

10.1.4 Email Notification

By default, fatal events are displayed as email notifications. Based on your configuration, the email notifications are delivered to you as shown in the following figure.

In the email notification, besides the event's description, the email also contains system information and the controller's image details. Using this additional information, you can find out the system and the controller on which the fatal error occurred.

If the MegaRAID Storage Manager Framework connects to a VMware ESXi server, an additional read only field **Event From** appears in the following dialog showing the IP address of the VMware ESXi server.

Figure 263 Email Notification

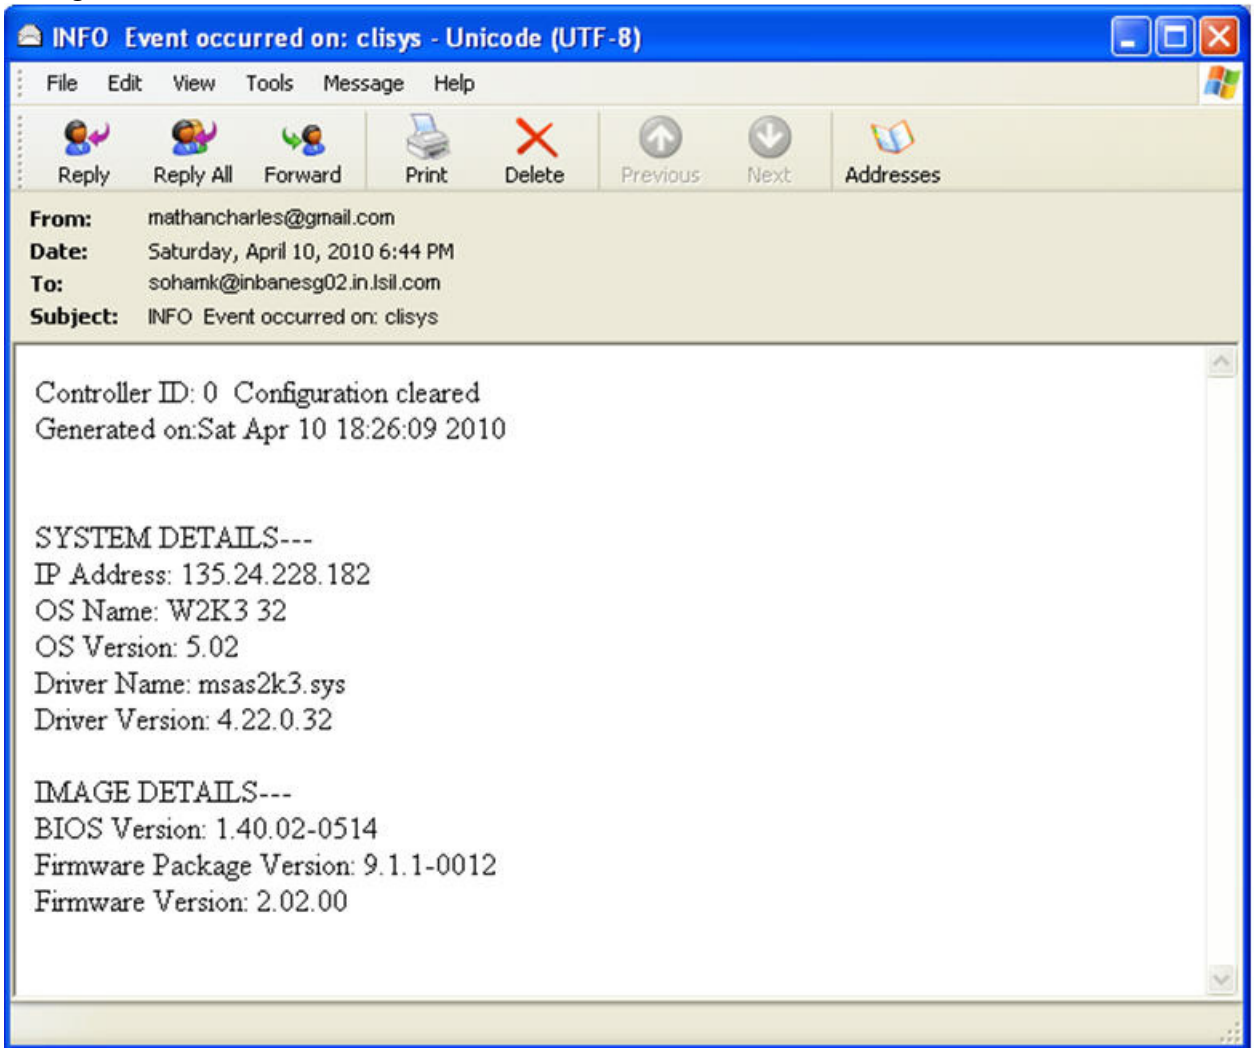

10.2 Configuring Alert Notifications

The Alert Notification Configuration feature allows you to control and configure the alerts that the MegaRAID Storage Manager software sends when various system events occur.

Select **Tools > Configure Alerts** on the main menu screen.

NOTE The **Configure Alerts** option differs based on your configuration. If the MegaRAID Storage Manager Framework connects to a Linux, Solaris, or a Windows server, the **Tools** menu shows the **Configure Alerts** option. If Monitor Plugin is configured on the server, the Tools menu shows the **Monitor Configure Alerts** option. If the MegaRAID Storage Manager Framework connects with a VMware ESXi server, the Tools menu shows the **CIMOM Configure Alerts** option.

The **Configure Alerts** window appears, as shown in the following figure. The window contains three tabs: **Alert Settings**, **Mail Server**, and **Email**.

Figure 264 Configure Alerts

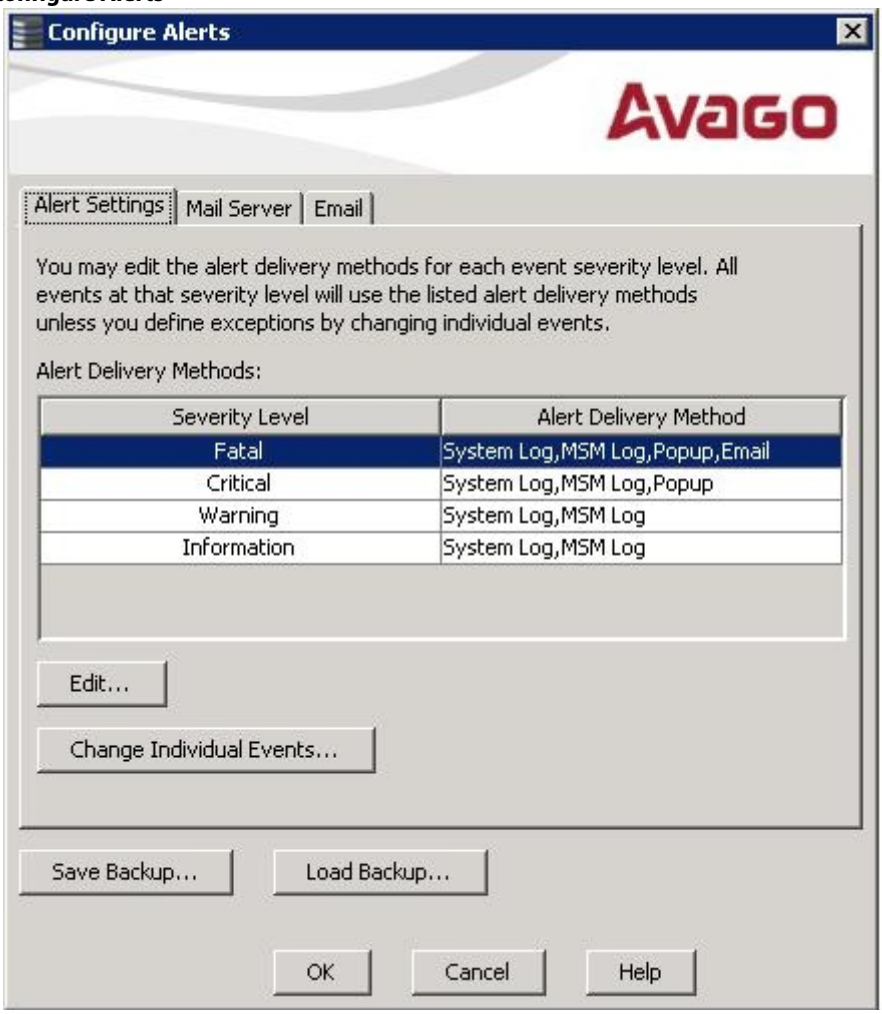

You can select the **Alert Settings** tab to perform the following actions:

- Edit the alert delivery method for different severity levels.
- Change the method of delivery for each individual event.
- Change the severity level of each individual event.
- Save an `.xml` backup file of the entire alert configuration.
- Load all the values from a previously saved backup into the dialog to edit or save these values as the current alert notification configuration.

NOTE When you load a saved backup file, all unsaved changes made in the current session will be lost.

You can select the **Mail Server** tab to perform the following actions:

- Enter or edit the sender email address.
- Enter the SMTP server name or the IP address.
- Enter the SMTP server authentication related information (user name and password).

NOTE These fields are optional and are filled only when the SMTP server requires authentication.

- Save an `.xml` backup file of the entire alert configuration.
- Load all of the values from a previously saved backup into the dialog to edit or save these values as the current alert notification configuration.

ATTENTION When you load a saved backup file, all unsaved changes made in the current session will be lost.

You can select the **Email** tab to perform the following actions:

- Add new email addresses for recipients of alert notifications.
- Send test messages to the recipient email addresses.
- Remove email addresses of recipients of alert notifications.
- Save an `.xml` backup file of the entire alert configuration.
- Load all of the values from a previously saved backup into the dialog to edit or save these values as the current alert notification configuration.

ATTENTION When you load a saved backup file, all unsaved changes made in the current session will be lost.

10.3 Editing Alert Delivery Methods

You can edit the default alert delivery methods, such as pop-up, email, system log, or the Vivaldi Log/MegaRAID Storage Manager log to a different severity level (Information, Warning, Critical and Fatal).

Perform the following steps to edit the alert delivery methods:

1. On the **Configure Alerts** window, click the **Alerts Setting** tab.
2. Under the **Alerts Delivery Methods** heading, select one of the severity levels.
3. Click **Edit**.

   The **Edit** dialog appears.

Figure 265 Edit Dialog

4. Select the desired alert delivery methods for alert notifications at the event severity level.
5. Click **OK** to set the delivery methods used for the severity level that you selected.

10.4 Changing Alert Delivery Methods for Individual Events

You can change the alert delivery options for an event without changing the severity level.

1. On the **Configure Alerts** window, click the **Alerts Setting** tab.
   The **Alerts Setting** portion of the window appears.
2. Click **Change Individual Events**.
   The **Change Individual Events** dialog appears, as shown in the following figure. The dialog shows the events by their ID number, description, and the severity level.

Figure 266 Change Individual Events

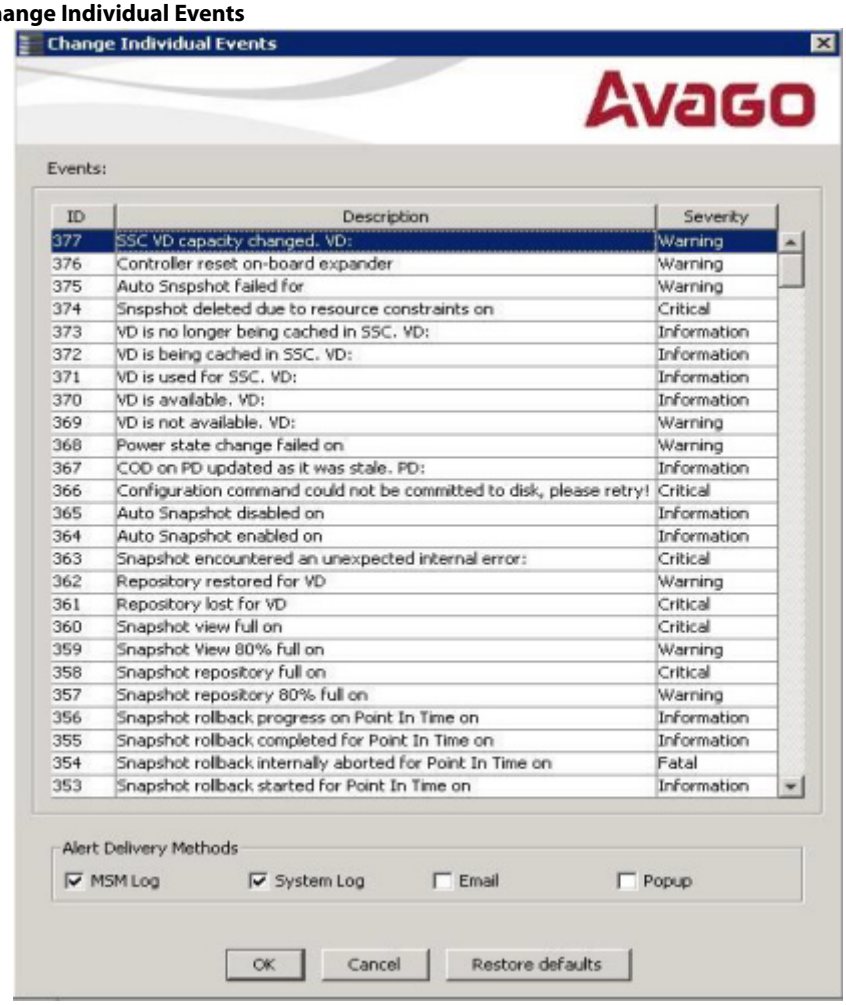

3. Click an event in the list to select it.
   The current alert delivery methods appear for the selected event in the **Alert Delivery Methods** frame.
4. Select the desired alert delivery methods for the event.
5. Click **OK** to return to the **Configure Alerts** window.
6. You may click **Cancel** to discard your current changes and to go back to the **Configure Alerts** window.
7. In the **Configure Alerts** window, click **OK**.

NOTE You can click **Restore Defaults** to revert back to the default alert delivery method and the default severity level of an individual event.

For more information, see Roll Back to the Default Individual Event Configuration.

10.5 Changing the Severity Level for Individual Events

To change the event severity level for a specific event, perform the following steps:

1. On the **Configure Alerts** window, click the **Alerts Setting** tab.
   The **Alerts Setting** portion of the window appears.
2. Click **Change Individual Events**.
   The **Change Individual Events** dialog appears. The dialog shows the events by their ID number, description, and severity level.
3. Click an event in the list to select it.
   The current severity appears in the **Severity** cell for the selected event.
4. Click the **Severity** cell for the event.
   The **Event Severity** drop-down menu appears for that event, as shown in the following figure.

Figure 267 Change Individual Events Severity Level Menu

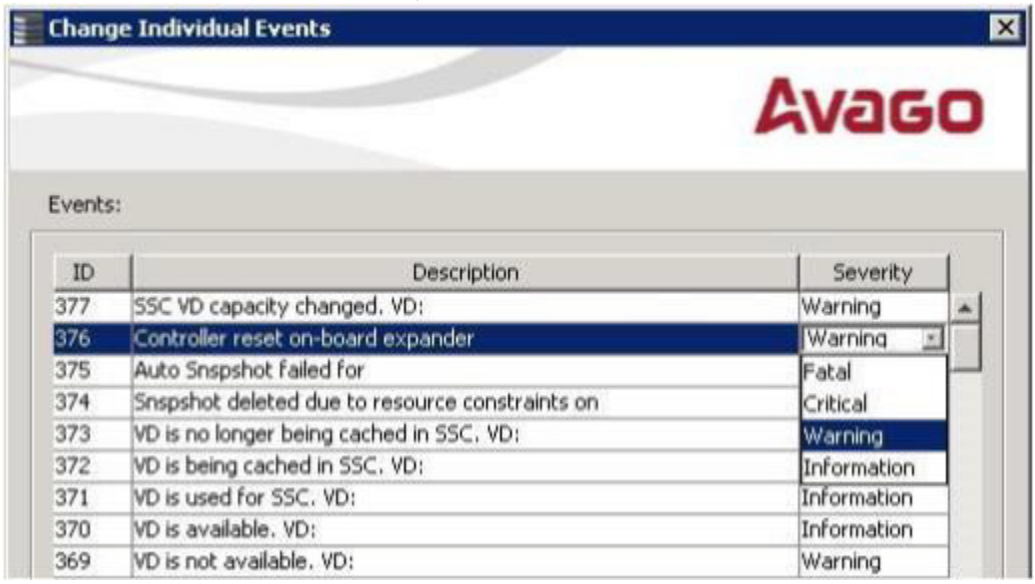

5. Select a different severity level for the event from the menu.
6. Click **OK** to return to the **Configure Alerts** window.
7. You may click **Cancel** to discard your current changes and to go back to the **Configure Alerts** window.
8. In the **Configure Alerts** window, click **OK** to save all the changes made to the events.

10.6 Roll Back to the Default Individual Event Configuration

To revert back to the default alert delivery method and the default severity level of an individual event, perform the following steps:

1. On the **Configure Alerts** window, click the **Alerts Setting** tab.
   The **Alerts Setting** portion of the window appears.
2. Click **Change Individual Events**.
   The **Change Individual Events** dialog appears, as shown in the Change Individual Events figure. The dialog shows the events by their ID number, description, and the severity level.
3. Click **Restore Defaults**.
   The **Change Individual Events** dialog appears with the default alert delivery method and the default severity level of all individual events.
4. Click **OK** to return to the **Configure Alerts** window.
5. In the **Configure Alerts** window, click **OK** to save all the changes made to the events.

10.7 Entering or Editing the Sender Email Address and SMTP Server

You can use the **Configure Alerts** window to enter or edit the sender email address and the SMTP server.

1. On the **Configure Alerts** window, click the **Mail Server** tab.
   The Mail Server options appear, as shown in the following figure.

Figure 268 Mail Server Options

2. Enter a sender's email address in the **Sender email address** field, or edit the existing sender email address.
3. Enter your SMTP server name/IP Address in the **SMTP Server** field, or edit the existing details.

4. Clear the **Use Default** check box to enter the desired port number in the **Port** field.
5. Click **OK**.

The MegaRAID Storage Manager software does *not* support e-mail functionality using a secured SMTP server such as Gmail or Yahoo.

If an SMTP server uses its own self signed certificate, communication cannot be established to the MegaRAID Storage Manager server for security reasons. The MegaRAID Storage Manager software can communicate with all mail clients that either use a MegaRAID Storage Manager software certificate or do not use their own self signed certificate.

10.8 Authenticating the SMTP Server

The MegaRAID Storage Manager software supports a SMTP authentication mechanism called *Login*. This feature provides an extra level of security, while sending an email from the MegaRAID Storage Manager server.

To enter or modify the SMTP server authentication information, perform the following steps:

1. On the **Configure Alerts** window, click the **Mail Server** tab.

   The Mail Server options appear, as shown in the Mail Server Options figure.
2. If on your SMTP server, the authentication mechanism is enabled and if you want to enable this feature on the MegaRAID Storage Manager software, then you need to select the **This Server requires authentication** check box and enter the authentication details in the corresponding fields (**User name** and **Password**).

   If you do not want to enable this feature on the MegaRAID Storage Manager software or if you know that your SMTP server does not support the *Login* mechanism, then de-select the **This Server requires authentication** check box.

   NOTE The **This Server requires authentication** check box is selected by default.
3. Enter a user name in the **User name** field.

   This step is optional if **This Server requires authentication** check box is selected.
4. Enter the password in the **Password** field.

   This step is optional if **This Server requires authentication** check box is selected.
5. Click **OK**.

10.9 Adding Email Addresses of Recipients of Alert Notifications

The **Email** tab in the **Configure Alerts** window shows the email addresses of the recipients of the alert notifications. The MegaRAID Storage Manager software sends alert notifications to those email addresses. Use the **Configure Alerts** window to add or remove email addresses of recipients and to send test messages to recipients that you add.

To add email addresses of recipients of the alert notifications, perform the following steps:

1. Click the **Email** tab in the **Configure Alerts** window.

Figure 269 Adding Email Settings

2. Enter the email address you want to add in the **New recipient email address** field.
3. Click **Add**.
   The new email address appears in the **Recipient email addresses** field.

10.10 Testing Email Addresses of Recipients of Alert Notifications

Use the **Email** tab in the **Configure Alerts** window to send test messages to the email addresses that you added for the recipients of alert notifications.

1. Click the **Email** tab on the **Configure Alerts** window.
   The **Email** section of the window appears, as shown in the Adding Email Settings figure.
2. Click an email address in the **Recipient email addresses** field.
3. Click **Test**.
4. Confirm whether the test message was sent to the email address.
   A pop-up message indicates if the test message sent to the email address was successful. If the MegaRAID Storage Manager software cannot send an email message to the email address, an error message appears.

10.11 Removing Email Addresses of Recipients of Alert Notifications

Use the **Email** tab in the **Configure Alerts** window to remove email addresses of the recipients of alert notifications.

1. Click the **Email** tab on the **Configure Alerts** window.
   The **Email** section of the window appears, as shown in the Adding Email Settings figure.
2. Click an email address in the **Recipient email addresses** field.
   The **Remove** button, which was grayed out, is now active.

3. Click **Remove**.

   The email address is deleted from the list.

10.12 Saving Backup Configurations

You can save an `.xml` backup file of the entire alert configuration. This includes all the settings on the three tabs (**Alert Settings**, **Mail Server**, and **Email**).

1. On the **Configure Alerts** window, click the **Alert Setting** tab, the **Mail Server** tab, or the **Email** tab.
2. Click **Save Backup**.

   The drive directory appears.
3. Enter a filename with an `.xml` extension for the backup configuration (in the format `filename.xml`).
4. Click **Save**.

   The drive directory disappears.
5. Click **OK**.

   The backup configuration is saved, and the **Configure Alerts** window closes.

10.13 Loading Backup Configurations

You can load all of the values from a previously saved backup into the **Configure Alerts** window (all tabs) to edit or save these values as the current alert notification configuration.

NOTE If you choose to load a backup configuration and the **Configure Alerts** window currently contains changes that have not yet been saved as the current alert notification configuration, the changes will be lost. You are prompted to confirm your choice.

1. On the **Configure Alerts** window, click the **Alert Setting** tab, the **Mail Server** tab, or the **Email** tab.
2. Click **Load Backup**.

   You are prompted to confirm your choice. The drive directory appears from which you can select a backup configuration to load.
3. Select the backup configuration file (it should be in `.xml` format).
4. Click **Open**.

   The drive directory disappears.
5. Click **OK**.

   The backup configuration is saved, and the **Configure Alerts** window closes.

10.14 Monitoring Server Events

The MegaRAID Storage Manager software enables you to monitor the activity of MegaRAID Storage Manager users in the network.

When a user logs on/logs off from the application, the event message appears in the log displayed at the bottom of the MegaRAID Storage Manager screen (the Vivaldi log/MegaRAID Storage Manager Log). These event message have a

severity level, a date and timestamp (User log on / log off time), and a brief description that contains a user name, client IP address, an access mode (full/view only) and a client system time.

10.15 Monitoring Controllers

When the MegaRAID Storage Manager software is running, you can see the status of all the controllers in the left panel.

If a controller is operating normally, the controller icon looks like this: If a controller has failed, a small red circle appears next to the icon.

To display the complete controller information, click on a controller icon in the left panel of the MegaRAID Storage Manager main menu. The controller properties appear in the right panel as shown in the following figure. Most of the information on this tab is self-explanatory.

Figure 270 Controller Properties

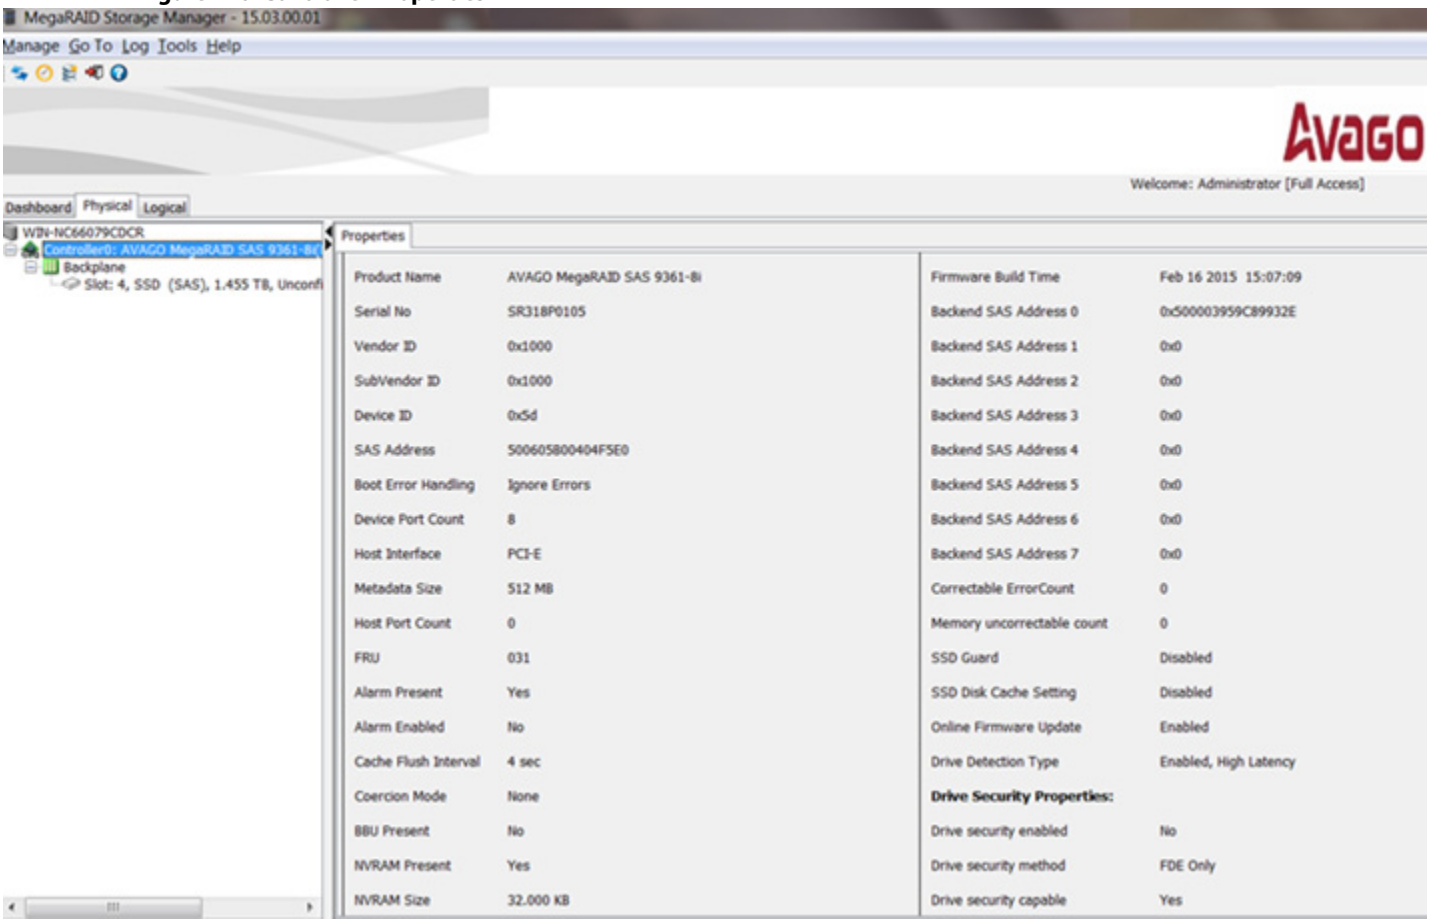

In the above dialog, the following properties appear under the **High Availability Properties** heading if the controller supports High Availability DAS:

- **Topology Type** - Indicates whether clustering is supported or not on the controller. Possible values for this field are **Server Storage Cluster**, or **None**.
- **Maximum Controller Nodes** - Indicates the total number of servers in a cluster.
- **Domain ID** - Shows the domain ID of the two servers in a cluster. The domain ID for both the servers is the same.
- **Peer Controller Status** - Indicates if both the servers in a cluster are running or not. The possible values are **Active**, **Inactive**, or **Incompatible**.

- **Incompatibility Details** - Indicates the reason for the incompatibility between the servers in a cluster. The possible values are **FW Level Mismatch**, **HW Incompatibility**, **Controller Property Mismatch**, **Premium Features Mismatch**, or **None**.

  NOTE If the controller does not support High Availability DAS, only the **Topology Type** property appears under the **High Availability Properties** heading, with the value **None**.

The Rebuild rate, Patrol read rate, Reconstruction rate, Consistency check rate, and BGI rate (background initialization) are all user selectable. For more information, see Changing Adjustable Task Rates.

The **BBU Present** field indicates whether a battery backup unit is installed.

The **Alarm Enabled** field indicates whether the controller has an alarm to alert the user with an audible tone when there is an error or a problem on the controller. Options are available for disabling or silencing the alarm by right clicking on a controller icon or by selecting **Go To > Controller** menu.

10.16 Monitoring Drives

When the MegaRAID Storage Manager software is running, you can see the status of all the drives in the left panel. If a drive is operating normally, the icon looks like this: If a drive has failed, a small red circle appears to the right of the icon.

To display the complete drive information, click on a drive icon in the left panel of the MegaRAID Storage Manager main menu. The drive properties appear in the right panel as shown in the following figure. The information on this tab 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.

Figure 271 Drive Properties

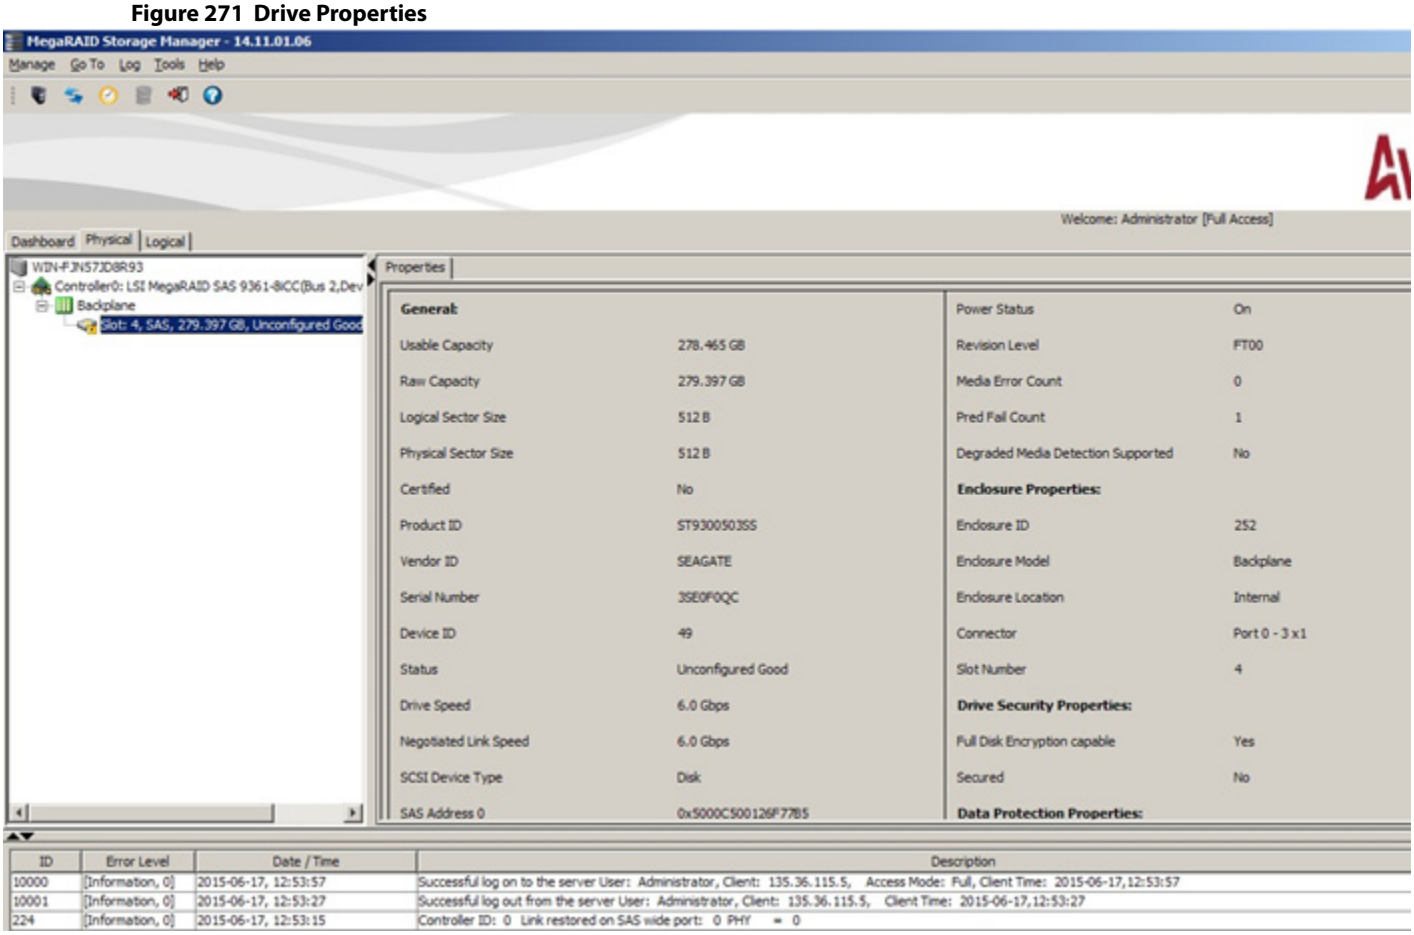

The **Power Status** property displays the status On when a drive is spun up and displays the status Powersave when a drive is spun down. Note that SSD drives and other drives that never spin down still show On.

If the drives are in a disk enclosure, you can identify which drive is represented by a disk icon on the left. To do this, follow these steps:

1. Click the drive icon in the left panel.
2. Select **Go To > Physical Drive > Start Locating Drive** tab in the right panel.
   The LED on the drive in the enclosure starts blinking to show its location.

   NOTE LEDs on drives that are global hot spares do not blink.
3. To stop the drive light on the enclosure from blinking, select **Go To > Physical Drive > Stop Locating Drive**.

10.17 Running a Patrol Read

A patrol read periodically verifies all sectors of the drives connected to a controller, including the system reserved area in the RAID configured drives. You can run a patrol read for all RAID levels and for all hot spare drives. A patrol read is initiated only when the controller is idle for a defined period and has no other background activities.

You can set the patrol read properties and start the patrol read operation, or you can start the patrol read without changing the properties.

1. Click a controller icon in the left panel.

2. Select **Go To > Controller > Set Patrol Read Properties**, or right-click on a controller and select **Set Patrol Read Properties** from the menu.

   The **Patrol Read - Set properties** window appears, as shown in the following figure.

Figure 272 Patrol Read - Set Properties

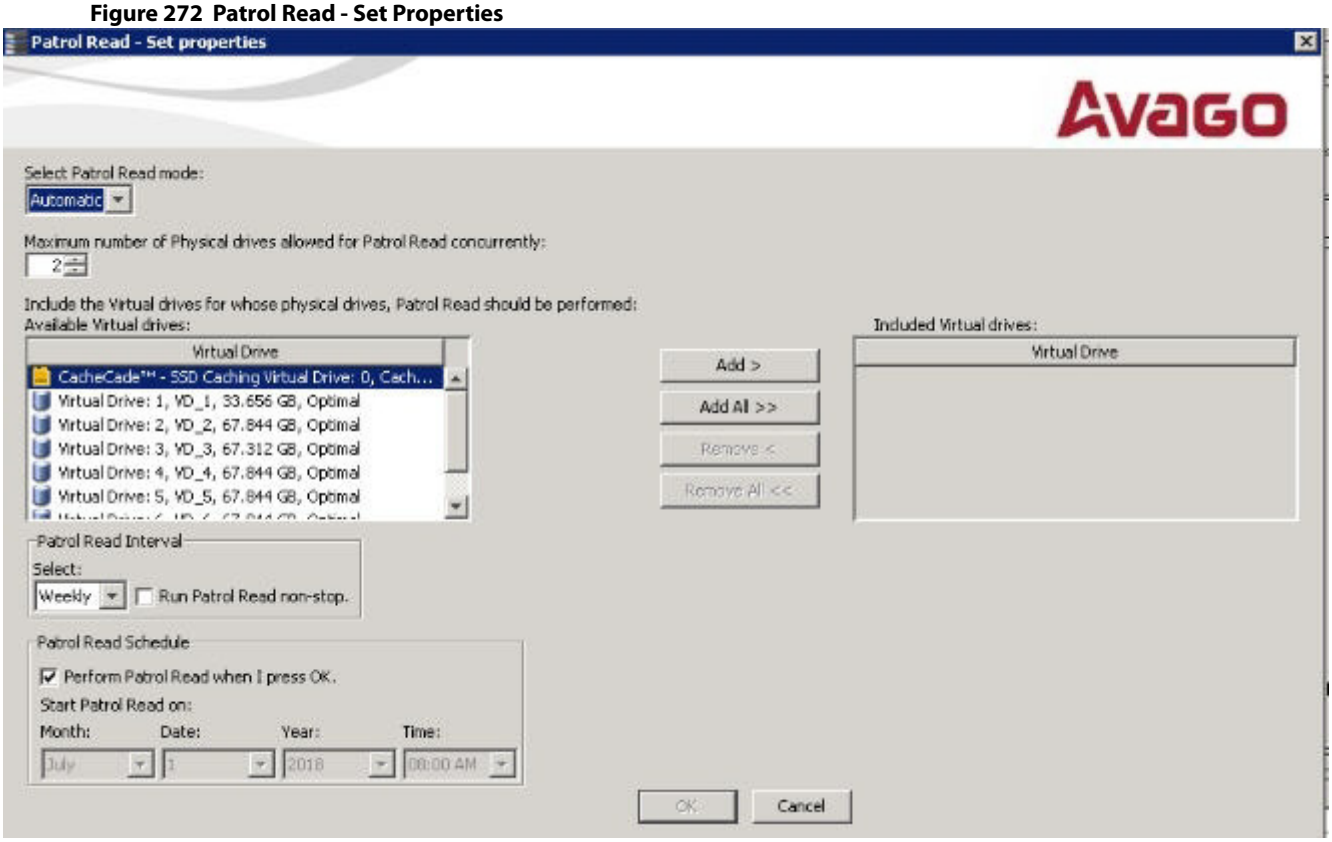

3. Select an operation mode for patrol read from the following options:
   - **Automatic**: Patrol read runs automatically at the time interval you specify on this window.
   - **Manual**: Patrol read runs only when you manually start it, by selecting Start Patrol Read from the controller options window.
   - **Disabled**: Patrol read does not run.
4. (Optional) Specify a maximum count of drives to include in the patrol read.

   The count must be a number from 1 to 255.
5. (Optional) Click virtual drives in the list under the heading **Virtual Drive** 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 runs.

   The default frequency is weekly (168 hours), which is suitable for most configurations. The other options are hourly, daily, and monthly.

NOTE Leave the patrol read frequency and other patrol read settings at the default values to achieve the best system performance. If you decide to change the values, record the original default values here so you can restore them later, if necessary: **Patrol Read Frequency**: _____, **Continuous Patrolling**: Enabled/Disabled, **Patrol Read Task Rate**: _____.

7. (Optional) Set Patrol Read to run at a specific time.

   The default setting for the patrol read is to start when you click **OK** on this window. To change the default setting so that the patrol read starts at a specific time, follow these steps (otherwise, skip this step and proceed to step 8):

   a. Deselect the **Perform Patrol Read when I press OK** check box.
   b. Select the month, year, day, and time to start the 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 only in the event log.

9. Click **Go** to enable these Patrol Read options.

To start a patrol read without changing the patrol read properties, follow these steps:

1. Click a controller icon in the left panel of the MegaRAID Storage Manager main menu screen.
2. Select **Go To > Controller > Start Patrol Read** in the menu bar, or right-click a controller and select **Start Patrol Read** from the menu.
3. When prompted, click **Yes** to confirm that you want to start a patrol read.

10.17.1 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. Leave the patrol read task rate at its default setting.

If you raise the task rate above the default, the foreground tasks run slowly, and it might appear that the system is not responding. If you lower the task rate less than the default, rebuilds and other background tasks might run very slowly and might not complete within a reasonable time.

10.18 Monitoring Virtual Drives

When the MegaRAID Storage Manager software is running, you can see the status of all virtual drives. If a virtual drive is operating normally, the icon looks like this: . Color-coded circles appear next to the icon to indicate the following:

- Green: The server is operating properly.
- Yellow: The server is running in a partially degraded state (for example, if a drive has failed); the data is still safe, but data could be lost if another drive fails.
- Orange: The server is running in a degraded state.
- Red: The server storage configuration has failed.

When the **Logical** tab is selected, the panel on the left shows which drives are used by each virtual drive. The same drive can be used by multiple virtual drives.

To display complete virtual drive information, click the **Logical** tab in the left panel, and click on a virtual drive icon in the left panel. The properties appear in the right panel as shown in the following figure. The RAID level, strip size, and access policy of the virtual drive are set when the virtual drive is configured.

Figure 273 Virtual Drive Properties

If High Availability DAS is supported on the controller, two additional virtual drive properties, **GUID** and **Host Access Policy** appear on the Properties page.

- **GUID** - Indicates a unique ID assigned to this virtual drive by the firmware.
- **Host Access Policy** - Indicates whether or not the virtual drive is shared between the servers in a cluster. The values for this property are **Shared**, **Exclusive**, and **Exclusive to Peer Controller**.

You can change the read policy, write policy, and other virtual drive properties. To change these properties, see Changing Virtual Drive Properties.

NOTE You can change the Read Policy, Write Policy, and other virtual drive properties by selecting the virtual drive icon and then selecting **Go To > Virtual Drive > Set Virtual Drive Properties** in the menu bar.

If the drives in the virtual drive are in a disk enclosure, you can identify them by making their LEDs blink. To identify the drives, follow these steps:

1. Click the virtual drive icon in the left panel.
2. Either select **Go To > Virtual Drive > Start Locating Virtual Drive**, or right-click a virtual drive and select **Start Locating Virtual Drive** from the menu.

   The LEDs on the drives in the virtual drive start blinking (except for the hot spare drives).
3. To stop the LEDs from blinking, select **Go To > Virtual Drive > Stop Locating Virtual Drive**, or right-click a virtual drive and select **Stop Locating Virtual Drive** from the menu.

10.19 Monitoring Enclosures

When the MegaRAID Storage Manager software is running, you can see the status of all enclosures connected to the server by selecting the **Physical** tab in the left panel. If an enclosure is operating normally, the icon looks like this: . If an enclosure is not functioning normally—for example, if a fan has failed—an orange, yellow, or red circle appears to the right of the icon.

Information about the enclosure appears in the right panel when you select the **Properties** tab on the main menu screen. A graphical display of enclosure information appears when you select the **Graphical View** tab.

The display in the center of the screen shows how many slots of the enclosure are populated by the drives and the lights on the drives show the drive status. The information on the right shows you the status of the temperature sensors, fans, and power supplies in the enclosure.

To view the enclosure properties, in the physical view click on the **Enclosure** node. The **Enclosure Properties** are displayed, as shown in the following figure.

Figure 274 Enclosure Properties

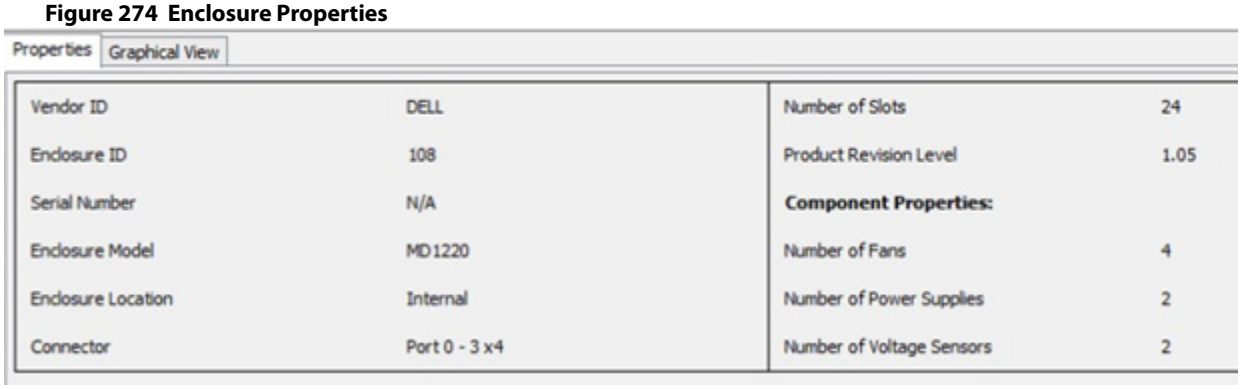

10.20 Monitoring Battery Backup Units

When the MegaRAID Storage Manager software is running, you can monitor the status of all of the BBUs connected to controllers in the server. If a BBU is operating normally, the icon looks like this: . If a BBU fails, a red dot appears next to the icon.

NOTE To increase the life of a battery, the battery is not fully charged. Band Gap charging keeps the maximum battery charge within a band comfortably above the data retention time requirement instead of keeping the battery charged to the maximum level. However, when a learn cycle is required, the battery is fully charged because a learn cycle starts only once the battery is fully charged.

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 pane, as shown in the following figure.

Figure 275 Battery Properties

Some fields like **Charge** appear only in the BBU property pages of batteries other than TMM-C battery. Similarly fields such as **Capacitance** appear only in the BBU property pages of TMM-C battery.

3. Click **Advanced Properties** to view additional BBU properties

   The **Advanced Properties** dialog appears.

Figure 276 Advanced Properties

Additional properties such as **Manufacturer**, **Serial Number**, **Full Capacity**, are displayed. You can also set battery learn cycles from the **Advanced Properties** dialog. For more details on battery learn cycles, see the following section.

10.21 Battery Learn Cycle

Learn cycle is a battery calibration operation that is performed by the controller periodically to determine the condition of the battery. You can start battery learn cycles manually or automatically. To choose automatic battery learn cycles, enable automatic learn cycles.

If you enable automatic learn cycles, you can delay the start of the learn cycles for up to 168 hours (7 days). If you select the **Generate an event to remind me when to start a learn cycle manually** check box in the **Set Automatic Learn Cycle Properties** dialog, the automatic learn cycle gets disabled and an event is generated to remind you when you need to start a learn cycle.

10.21.1 Setting Automatic Learn Cycle Properties

To set automatic learn cycle properties, perform the following steps:

NOTE For TMM-C battery you cannot set automatic learn cycles properties.

1. Click the **Physical** tab to open the Physical view.
2. Select the **BBU** icon in the left panel.
3. Select **Go To > BBU > Set Automatic Learn Cycle Properties**.

   The **Set Learn Cycle Properties** dialog appears, as shown in the following figure.

Figure 277 Set Learn Cycle Properties

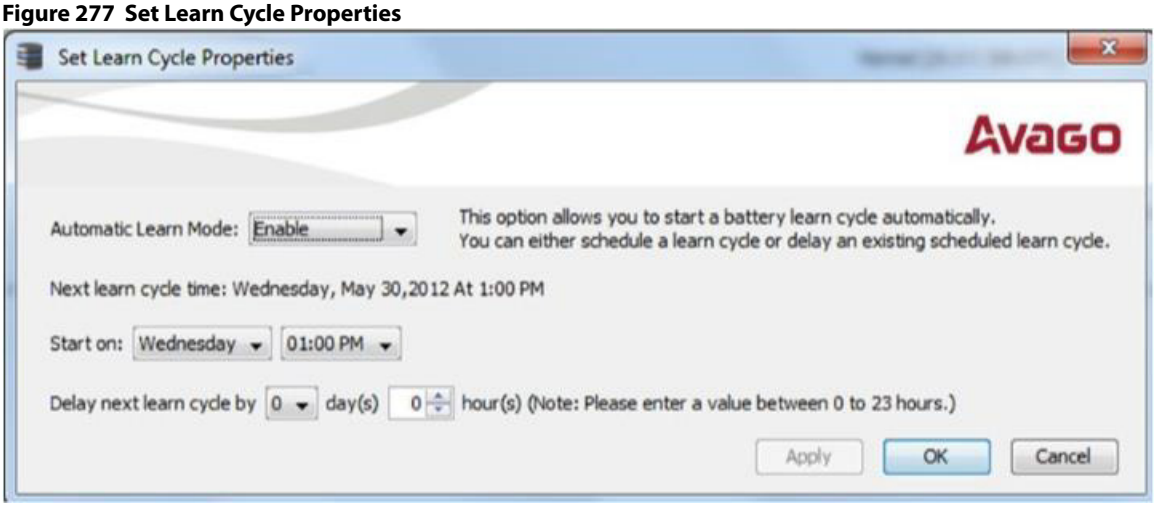

4. Select **Enable** from the **Automatic Learn Mode** drop-down list. The other two options are **Disable** and **Warn Via Event**.

   If you select **Disable**, the automatic battery learn cycle is disabled. The **Start on** and **Delay next learn cycle by** fields are also disabled.

   If you select **Warn Via Event**, an event is generated notifying you when to start a learn cycle manually.

   If a learn cycle is disabled or not scheduled, the value **None** appears in the **Next learn cycle time** field.

   If a learn cycle is already scheduled, the day of the week, date, and time of the next learn cycle appears in the **Next learn cycle time** field.

   NOTE After selecting **Disable**, if you select **Enable**, the controller firmware resets the battery module properties to initiate an immediate battery learn cycle. The **Next Learn cycle** field is updated only after the battery relearn is completed. Once the relearning cycle is completed, the value in the **Next Learn cycle** field displays the new date and the time of the next battery learn cycle.
5. In the **Start on** field, specify a day and time to start the automatic learn cycle.
6. You can delay the start of the next learn cycle up to 7 days (168 hours) by specifying the day and hours in the **Delay next learn cycle by** field.

   If changes are made to the **Set Learn Cycle Properties** dialog, click **Apply** to refresh the dialog with the updated settings, without closing the dialog.

   If you selected **Disable** in the **Automatic Learn Mode** drop-down list, and click **OK** or **Apply**, a warning dialog appears asking for your confirmation to disable the automatic learn cycle.

10.21.2 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. Perform one of these actions:
   — Select **Go To > BBU > Start Manual Learn Cycle**.
   — Right-click the **BBU** icon, and select **Start Manual Learn Cycle** from the pop-up menu.

10.22 Monitoring Rebuilds and Other Processes

The MegaRAID Storage Manager software lets you monitor the progress of rebuilds and other lengthy processes in the **Group Show Progress** window.

To monitor the progress of these operations, open the show progress window by selecting **Manage > Show Progress** on the menu bar.

The **Group Show Progress** dialog appears.

Figure 278 Group Show Progress Window

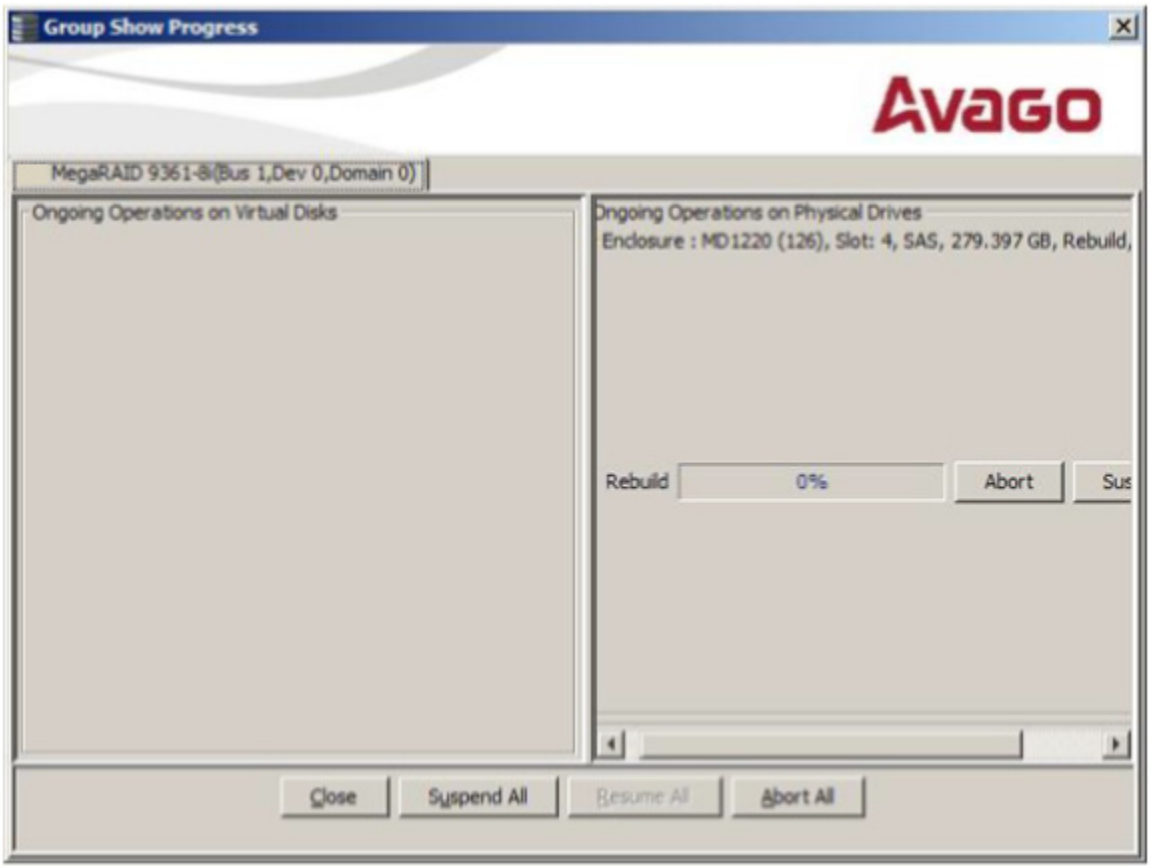

The **Group Show Progress** window displays a percent-complete indicator for drive rebuilds. Rebuilds might take a long time to complete. An up-arrow appears above the drive icon while it is being rebuilt.

Operations on virtual drives appear in the left panel of the window, and operations on drives appear in the right panel. The type of operations that appear in this window are as follows:

- Initialization of a virtual drive
- Rebuild
- Consistency check
- Non FDE Physical Drive Erase

- Virtual Drive Erase
- Patrol Read
- LD Reconstruction
- LD Disassociate
- PD Clear
- Replace
- Background Initialization (BGI)

A Modify Drive Group process cannot be aborted. To abort any other ongoing process, click the **Abort** button next to the status indicator. Click **Abort All** to abort all ongoing processes. Click **Close** to close the window.

10.23 Managing Link Speed

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

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

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

To change the link speed, perform the following steps:

1. Perform one of these actions:
   — Right-click a controller in the left frame of the MegaRAID Storage Manager main menu, and select **Manage Link Speed**.
   — 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 appears, as shown in the following figure.

Figure 279 Manage Link Speed Dialog

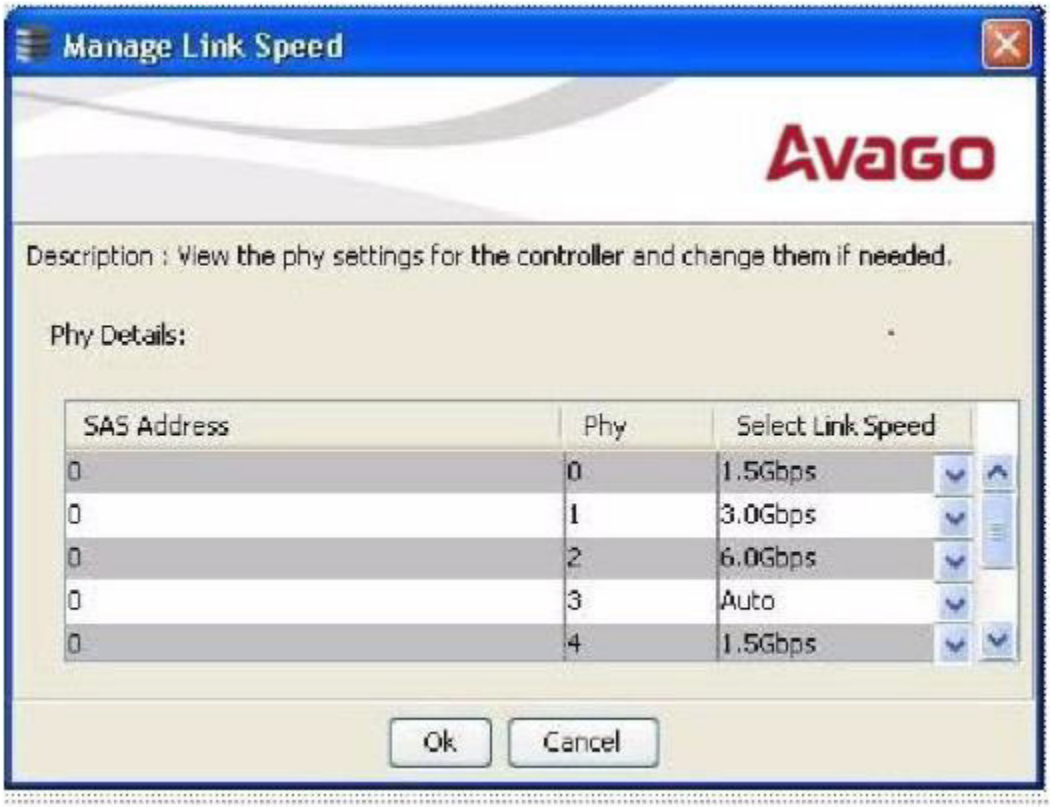

— The **SAS Address** column displays the SAS address that uniquely identifies a device in the SAS domain.
— The **Phy** column displays the system-supported phy link values. The phy link values are from 0 through 7.
— The **Select Link Speed** column displays the phy link speeds.

2. Select the desired link speed from the **Select Link Speed** field using the drop-down selector.
   The link speed values are Auto,1.5 Gbps, 3.0 Gbps, or 6.0 Gbps.

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

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 280 System Restart Required Message

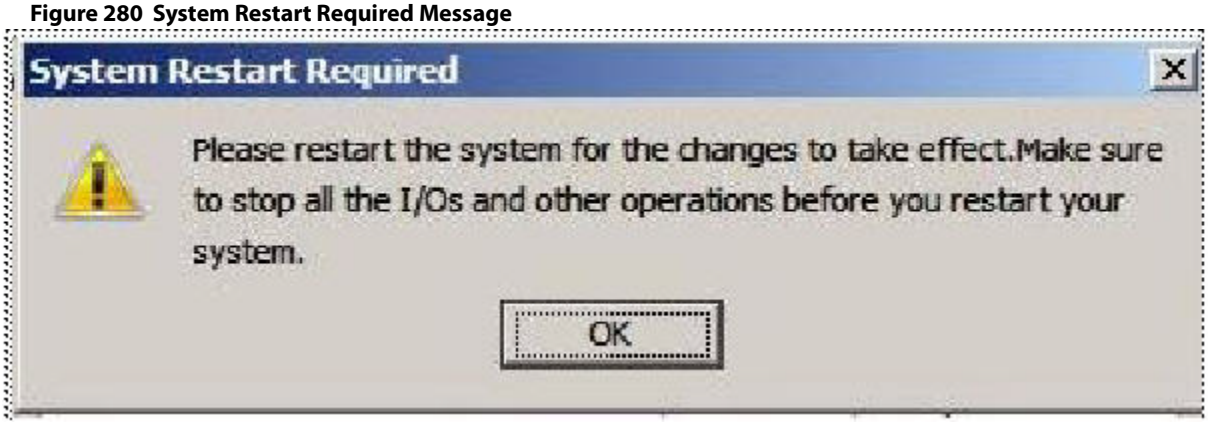

Chapter 11: Maintaining and Managing Storage Configurations

This chapter explains how to use the MegaRAID Storage Manager software to maintain and manage storage configurations. Log on to the server in Full Access mode to perform the maintenance and management tasks.

11.1 Initializing a Virtual Drive

When you create a new virtual drive with the **Configuration** Wizard, you can select the Fast Initialization or Full Initialization option to initialize the disk immediately. However, you can select No Initialization if you want to initialize the virtual drive later.

To initialize a virtual drive after completing the configuration process, perform these steps:

1. Select the **Logical** tab in the left panel of the **MegaRAID Storage Manager** window, and click the icon of the virtual drive that you want to initialize.
2. Select **Go To > Virtual Drive > Start Initialization**.

   The **Initialize** dialog appears.
3. Select the virtual drives to initialize.

   ATTENTION Initialization erases all data on the virtual drive. Make sure to back up any data you want to keep before you initialize a virtual drive. Make sure the operating system is not installed on the virtual drive you are initializing.
4. Select the **Fast Initialization** check box if you want to use this option.

   If you leave the box unselected, the MegaRAID Storage Manager software runs a Full Initialization on the virtual drive.
5. Click **Start** to begin the initialization.

   You can monitor the progress of the initialization. See Monitoring Rebuilds and Other Processes for more information.

11.1.1 Running a Group Initialization

Initialization prepares the storage medium for use. You can run initialization on multiple drives at one time. Follow these steps to run a group consistency check.

1. Select **Manage > Initialize**.

   The **Group Initialization** dialog appears.

Figure 281 Group Initialization Dialog

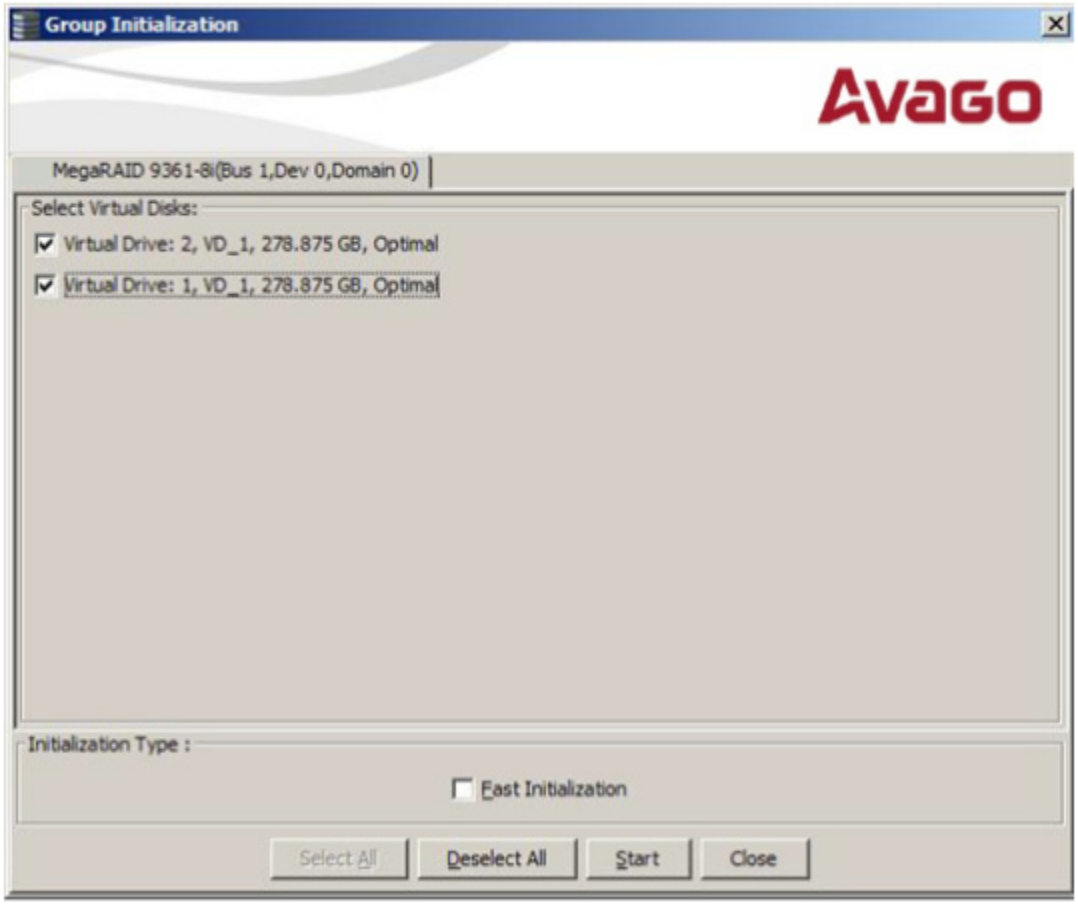

2. Either check the virtual drives on which to run the initialization, or click **Select All** to select all of the virtual drives.
3. Click **Start**.

   You can monitor the progress of the group initialization. See Monitoring Rebuilds and Other Processes for more information.

11.2 Running a Consistency Check

You should periodically run a consistency check on fault-tolerant virtual drives (RAID 1, 5, 6, 10, 50, or 60 configurations; RAID 0 does not provide data redundancy). A consistency check scans the virtual drive to determine whether the data has become corrupted and needs to be restored.

For example, in a system with parity, checking consistency means computing the data on one drive and comparing the results to the contents of the parity drive. You must run a consistency check if you suspect that the data on the virtual drive might be corrupted.

ATTENTION Make sure to back up the data before running a consistency check if you think the data might be corrupted.

To run a consistency check, first set the consistency check properties, and then schedule the consistency check. This section explains how to set the properties, schedule the check, and run the consistency check.

11.2.1 Setting the Consistency Check Settings

Follow these steps to set the properties for a consistency check:

1. Click the **Physical** tab or the **Logical** tab and select a controller.
2. Click **Go To > Controller > Set Consistency Check Properties**.
   The **Set Consistency Check Properties** dialog appears.

Figure 282 Set Consistency Check Properties Dialog

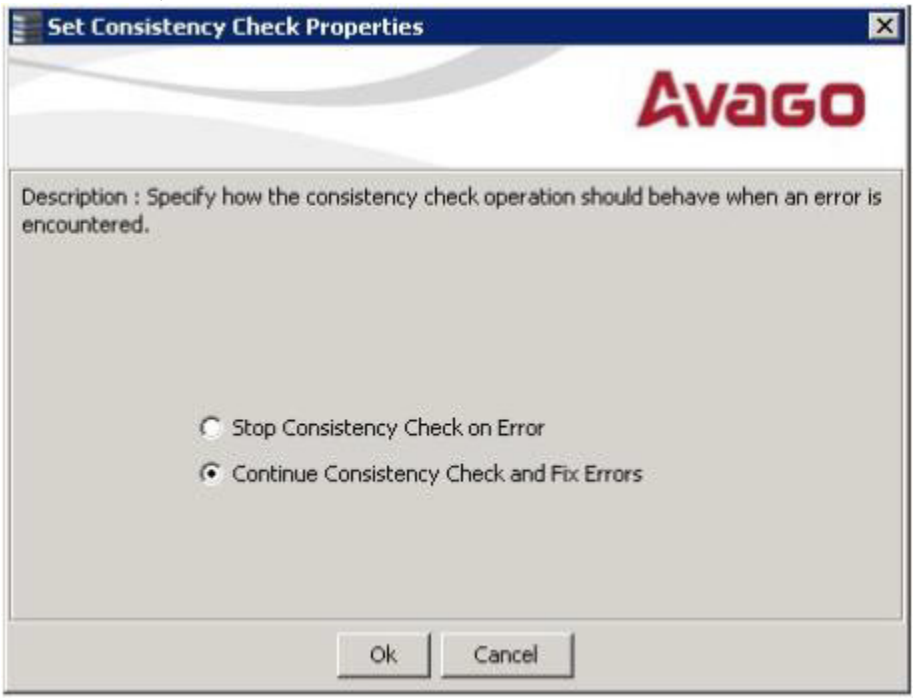

3. Choose one of the two options:
   - **Stop Consistency Check on Error**: The RAID controller stops the consistency check operation if the utility finds an error.
   - **Continue Consistency Check and Fix Errors**: The RAID controller continues the consistency check if the utility finds and error, and then fixes the errors.
4. Click **Ok**.

11.2.2 Scheduling a Consistency Check

Follow these steps to set the properties for a consistency check:

1. Click the **Physical** tab or the **Logical** tab, and select the controller.
2. Select **Go To > Controller > Schedule Consistency Check**.
   The **Schedule Consistency Check** dialog appears.

Figure 283 Schedule Consistency Check Dialog

3. Perform the following steps to schedule the consistency check:
   a. Select how often to run the consistency check from the drop-down list.
      You can click **Advanced** for more detailed date options.
   b. (Optional) Select the **Run consistency check continuously** check box.
   c. Select the month, day, and year on which to start the consistency check.
   d. Select the time of day to start the consistency check.
4. Click **Ok**.
   You can monitor the progress of the consistency check. See Monitoring Rebuilds and Other Processes for more information.

11.2.3 Running a Group Consistency Check

You can run a consistency check on multiple drives at one time. Follow these steps to run a group consistency check.

1. Select **Manage > Check Consistency**.
   The **Group Consistency Check** dialog appears.

Figure 284 Group Consistency Check Dialog

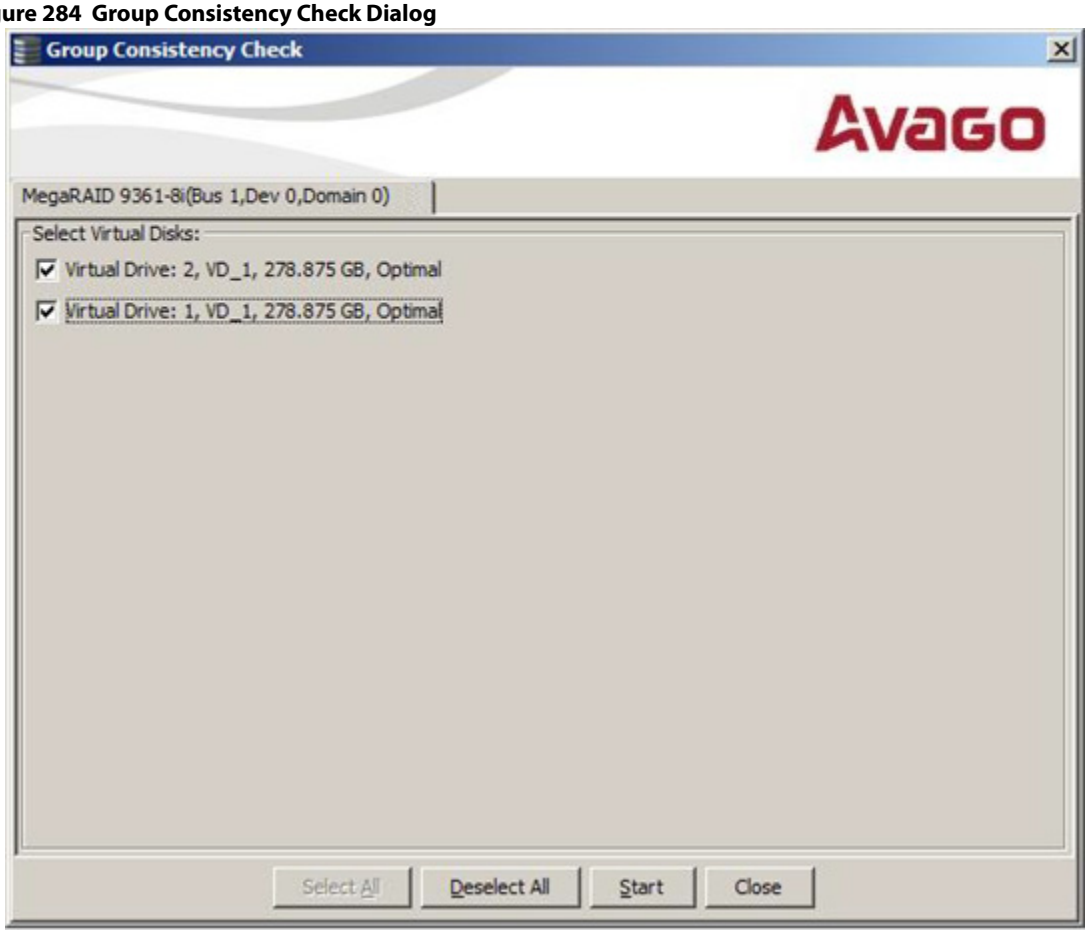

2. Either check the virtual drives on which to run the consistency check, or click **Select All** to select all of the virtual drives.
3. Click **Start**.
   You can monitor the progress of the group consistency check. See Monitoring Rebuilds and Other Processes for more information.

11.3 Scanning for New Drives

You can use the **Scan for Foreign Configuration** option to find drives with foreign configurations. A foreign configuration is a RAID configuration that already exists on a replacement set of physical disks that you install in a computer system. In addition, if one or more drives are removed from a configuration, by a cable pull or drive removal, for example, the configuration on those drives is considered a foreign configuration by the RAID controller. Drives that are foreign are listed on the physical drives list with a special symbol in the MegaRAID Storage Manager software.

The utility allows you to import the existing configuration to the RAID controller or clear the configuration so you can create a new configuration using these drives. You can preview the foreign configuration before you decide whether to import it.

The MegaRAID Storage Manager software usually detects newly installed drives and displays icons for them in the **MegaRAID Storage Manager** window. If for some reason the MegaRAID Storage Manager software does not detect a new drive (or drives), you can use the Scan for Foreign Configuration command to find it.

Follow these steps to scan for a foreign configuration:

1. Select a controller icon in the left panel of the **MegaRAID Storage Manager** window.
2. Select **Go To > Controller > Scan Foreign Configuration**.

   If the MegaRAID Storage Manager software detects any new drives, it displays a list of them on the window. If not, it notifies you that no foreign configuration is found.
3. Follow the instructions on the window to complete the drive detection.

11.4 Rebuilding a Drive

If a drive in a redundant virtual drive (RAID 1, 5, 6, 10, 50, or 60) fails, the MegaRAID Storage Manager software automatically rebuilds the data on a hot spare drive to prevent data loss. The rebuild is a fully automatic process, so it is not necessary to issue a `Rebuild` command. You can monitor the progress of drive rebuilds in the **Group Show Progress** window. To open this window, select **Manage > Show Progress**.

If a single drive in a RAID 1, RAID 5, RAID 10, or RAID 50 virtual drive fails, the system is protected from data loss. A RAID 6 virtual drive can survive two failed drives. A RAID 60 virtual drive can survive two failed drives in each span in the drive group. Data loss is prevented by using parity data in RAID 5, RAID 6, RAID 50, and RAID 60, and data redundancy in RAID 1 and RAID 10.

The failed drive must be replaced, and the data on the drive must be rebuilt on a new drive to restore the system to fault tolerance. You can choose to rebuild the data on the failed drive if the drive is still operational. If dedicated hot spares or global hot spare disks are available, the failed drive is rebuilt automatically without any user intervention.

A red circle to the right of the drive icon indicates that a drive has failed. A yellow circle appears to the right of the icon of the virtual drive that uses this drive which indicates that the virtual drive is in a degraded state; the data is still safe, but data could be lost if another drive fails.

Follow these steps to rebuild a drive:

1. Right-click the icon of the failed drive, and select **Rebuild**.
2. Click **Yes** when the warning message appears. If the drive is still good, a rebuild starts.

   You can monitor the progress of the rebuild in the **Group Show Progress** window by selecting **Manage > Show Progress**. If the drive cannot be rebuilt, an error message appears. Continue with the next step.
3. Shut down the system, disconnect the power cord, and open the computer case.
4. Replace the failed drive with a new drive of equal capacity.
5. Close the computer case, reconnect the power cord, and restart the computer.
6. Restart the MegaRAID Storage Manager software.

   When the new drive spins up, the drive icon changes back to normal status, and the rebuild process begins automatically. You can monitor the progress of the rebuild in the **Group Show Progress** window by selecting **Manage > Show Progress**.

If you want to force a drive into Fail status to trigger a rebuild, right-click the drive icon, and select **Make Drive Offline**. A red circle appears next to the drive icon. Right-click the icon, and select **Rebuild** from the pop-up menu.

11.4.1 New Drives Attached to a MegaRAID Controller

When you insert a new drive on a MegaRAID system and if the inserted drive does not contain valid DDF metadata, the drive displays as JBOD for MegaRAID entry-level controllers, such as the SAS 9240-4i/8i. If the drive does contain valid DDF metadata, its drive state is Unconfigured Good.

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

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

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

11.5 Making a Drive Offline or Missing

If a drive is currently part of a redundant configuration and you want to use it in another configuration, you can use the MegaRAID Storage Manager commands to remove the drive from the first configuration and change the drive state to Unconfigured Good.

ATTENTION After you perform this procedure, *all data on that drive is lost*.

To remove the drive from the configuration without harming the data on the virtual drive, follow these steps:

1. In the **MegaRAID Storage Manager** window, select **Go To > Physical Drive > Make Drive Offline**.
   The drive status changes to Offline.
2. Select **Go To > Physical Drive > Mark Drive as Missing**.
   The drive status changes to Unconfigured Good.

   ATTENTION 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.
   When a hot spare is available, the data on the virtual drive is rebuilt. You can now use the removed drive for another configuration.

   ATTENTION If the MegaRAID Storage Manager software detects that a drive in a virtual drive has failed, it makes the drive offline. If this situation occurs, you must remove the drive and replace it. You cannot make the drive usable for another configuration by using the **Mark physical disk as missing** command and the **Rescan** commands.

11.6 Removing a Drive

You may sometimes need to remove a non-failed drive that is connected to the controller. For example, you may need to replace the drive with a larger drive. Follow these steps to remove a drive safely:

1. Click the icon of the drive in the left panel, and click the **Operations** tab in the right panel.

2. Select **Prepare for Removal**, and click **Go**.
3. Wait until the drive spins down and remove it.

   If you change your mind, select **Undo Prepare for Removal**, and click **Go**.

11.7 Upgrading Firmware

The MegaRAID Storage Manager software enables you to easily upgrade the controller firmware. For limitations, see Online Firmware Upgrade Support.

To avoid data loss because of dirty cache on the controller, the utility forces the virtual disks into Write Through mode after a firmware upgrade. It is in this mode until the server reboots. In Write Through mode, the controller sends a data transfer completion signal to the host when the disk subsystem has received all of the data in a transaction. This way, in case of a power outage, the controller does not discard the dirty cache.

Follow these steps to upgrade the firmware:

1. In the left panel of the **MegaRAID Storage Manager** window, click the icon of the controller you want to upgrade.
2. In the **MegaRAID Storage Manager** window, select **Go To > Controller > Update Controller Firmware**.
3. Click **Browse** to locate the `.rom` update file, as shown in the following figure.

Figure 285 Update Controller Firmware Dialog

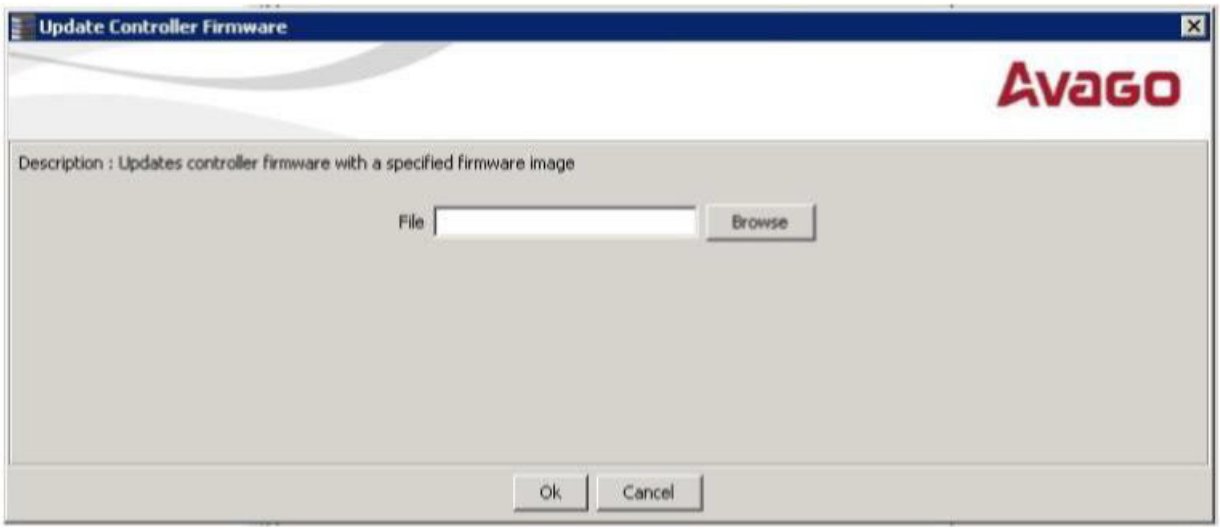

4. After you locate the file, click **Open**.

   The MegaRAID Storage Manager software displays the version of the existing firmware.
5. When you are prompted to indicate whether you want to upgrade the firmware, click **Yes**.

   A progress bar appears along with messages that indicate when an image opens and when an image downloads.
6. After an image has been downloaded and if Online Firmware Update is supported on the controller, a confirmation message box appears that asks for your confirmation.

NOTE If Online Firmware Update is not supported on the controller, the confirmation message box does not appear. Instead, after an image is downloaded, a message appears that indicates an image is being flashed. The controller is updated with the new firmware code

contained in the `.rom` file. Reboot the system after the new firmware is flashed. The new firmware does not take effect until reboot.

If you click **Yes** in the confirmation message box, the progress bar continues with a message that indicates that an image is being flashed.

After the progress bar disappears, either of the following two messages appear in a message box.

— New Firmware Version is flashed successfully. Online Firmware Update is not possible in this case. System reboot is required for the new firmware *<version number>* to take effect.

— `New Firmware Version is flashed successfully. Controller Reset will start now.`

If the first message appears, reboot your system.

If the second message appears, the MegaRAID Storage Manager main menu window reappears. A `Restart Started` event appears in the log (at the bottom of the MegaRAID Storage Manager main menu window) and a progress bar appears that states `Controller reset is in progress`.

After the controller reset process is completed, the controller is updated with the new firmware code contained in the `.rom` file.

NOTE While performing the Online Firmware Update method, there is a small window of time where the I/O transactions are held and the controller is automatically reset. This results in a timeout to your virtualized environments and causes I/O transaction errors. Choose the traditional firmware update method to avoid the controller reset.

11.7.1 Upgrading the CPLD Version

The MegaRAID Storage Manager software supports the Complex Programmable Logic Device (CPLD) version check feature.

To avoid updating an incorrect CPLD ROM file on the controller, the MegaRAID Storage Manager software compares the base part number of the CPLD USERCODE of the existing controller with the CPLD ROM file version. In the USERCODE of the existing controller, the first 20 bits are considered as the base part number and the rest 12 bits as revision level. If the base part number matches, then the MSM software checks for the revision level. If the version matches, the MSM software allows upgrading CPLD version of the controller. When an upgrade operation is initiated, MSM software prints both the versions and then upgrades the CPLD version.

Chapter 12: Using MegaRAID Advanced Software

This chapter describes the MegaRAID advanced software offered by the MegaRAID Storage Manager software for certain MegaRAID SAS 12Gb/s RAID controllers and explains how to use these features.

12.1 MegaRAID Advanced Software

The MegaRAID advanced software are features that the MegaRAID Storage Manager software supports on certain MegaRAID SAS 12Gb/s RAID controllers. The following MegaRAID SAS 12Gb/s RAID controllers support advanced software features that offer improved performance, data protection, and availability:

- MegaRAID SAS 9360-4i
- MegaRAID SAS 9360-8i
- MegaRAID SAS 9380-4i4e
- MegaRAID SAS 9380-8e
- MegaRAID SAS 9361 -8i
- MegaRAID SAS 9361-4i

NOTE Record your controller serial number in a safe location in case you need to contact Avago Technical Support.

ATTENTION Back up your data before you make a change in the system configuration. Failure to do so could result in data loss.

The MegaRAID advanced software includes the following features:

- MegaRAID FastPath
- MegaRAID CacheCade SSD Read Caching software
- MegaRAID CacheCade Pro 2.0 SSD Read/Write Caching software
- MegaRAID SafeStore

12.2 MegaRAID Software Licensing

The MegaRAID Software licensing authorizes you to enable the MegaRAID advanced software features present in the MegaRAID Storage Manager application. You have to obtain the activation key to enable, and use the advanced software features present in the controller.

12.3 Managing MegaRAID Advanced Software

The **MegaRAID Advanced Software** wizard allows you to use the advanced software features. Perform the following steps to enable the *activation key* to use the advanced controller features:

1. Select the **Physical** tab or the **Logical** tab in the left panel of the **MegaRAID Storage Manager** window, and click a controller icon.

2. Choose either of the following options:
   — Select **Go To > Controller > Manage MegaRAID Advanced Software Options**.
   — Click **Manage MegaRAID Advanced Software Options** from the dashboard under the feature portlet.

   The Manage MegaRAID Advanced Software Options wizard appears.
   — If none of the advanced software options present in the controller are in a boot mode, the second dialog appears, as shown in the following figure. You cannot activate any advanced software options from this window as this is a view-only window.
   — If even one of the advanced software options present in the controller is in a boot mode, the first dialog appears, as shown in the following figure.

Figure 286 Manage MegaRAID Advanced Software Options Dialog View-Only Mode

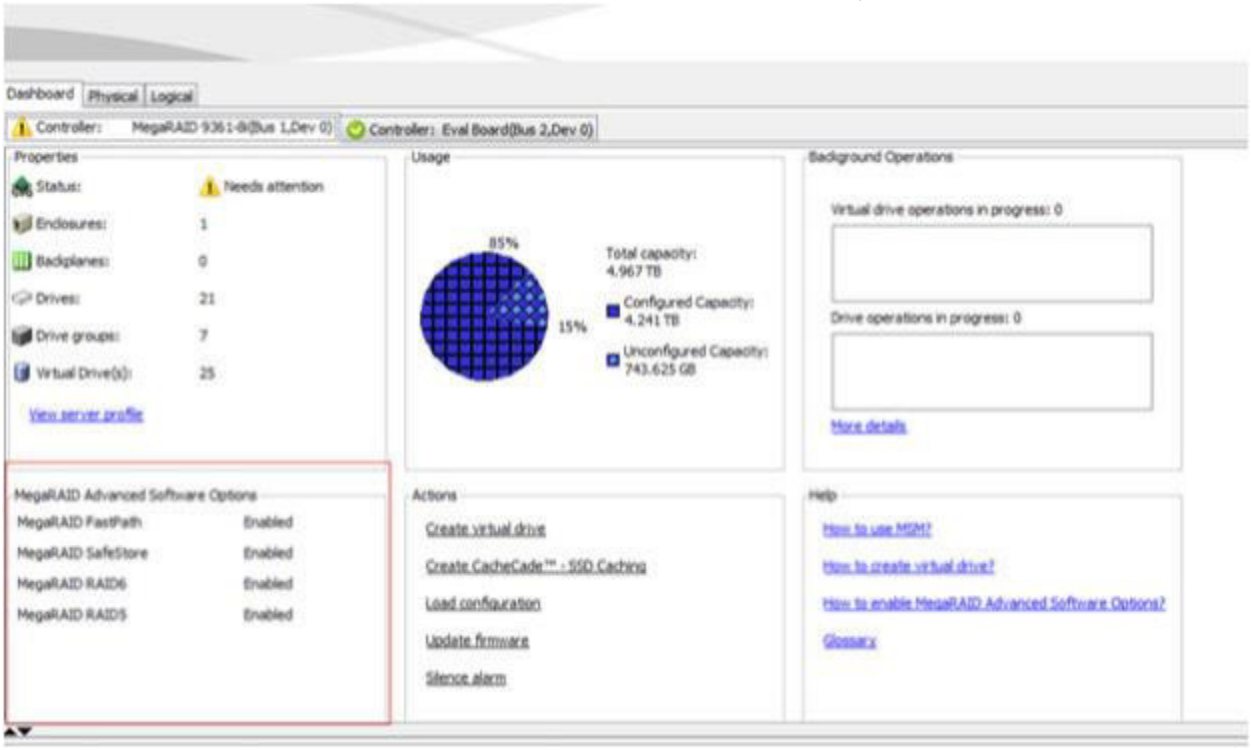

Figure 287 Manage MegaRAID Advanced Software Options Dialog

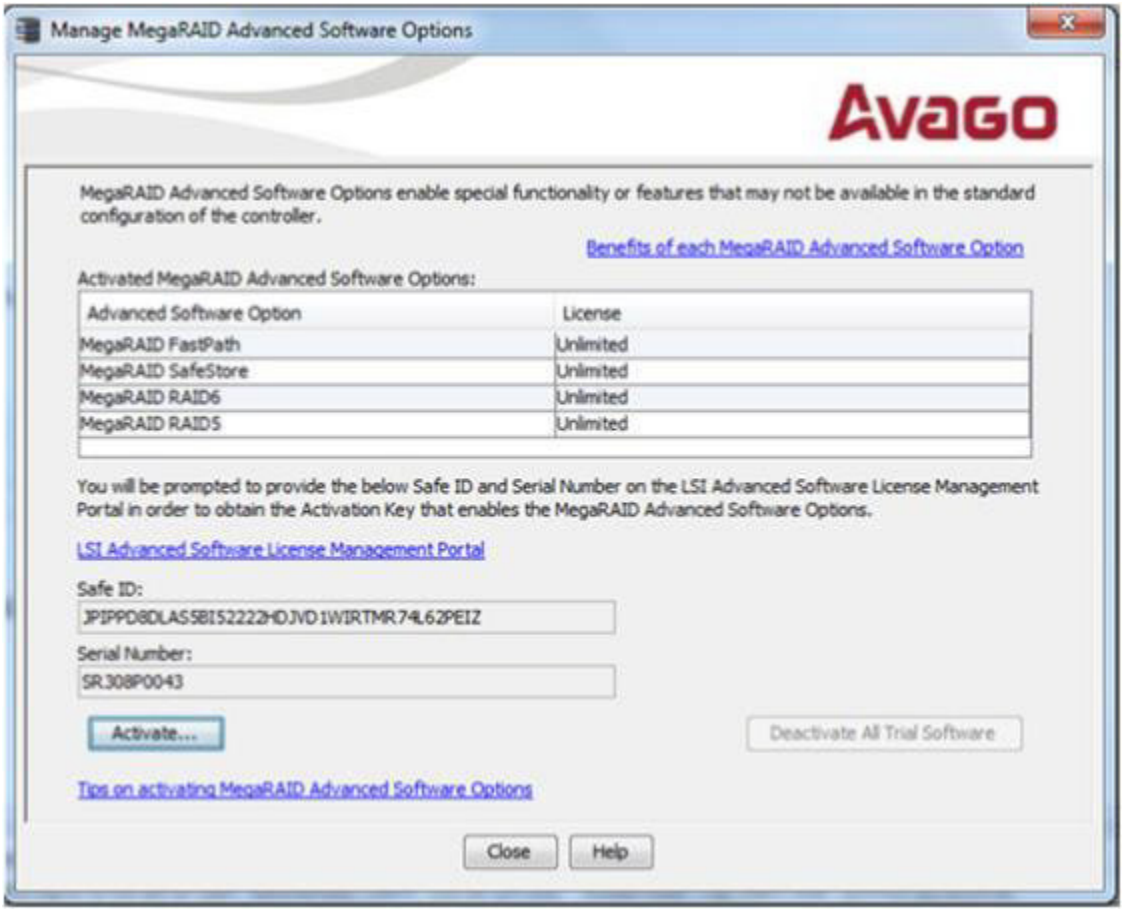

The **Activated MegaRAID Advanced Software Options** table consists of the **Advanced Software Option** and the **License** columns.

— The **Advanced Software Option** column displays the list of advanced software options present in the controller.

— The **License** column displays the license details for the list of advanced software options present in the **Advanced Software Option** column. The license details validates if the software is under a trial period, or if it can be used without any trial period (Unlimited).

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.avagotech.com/products/server-storage/raid-controllers/#tab-Adva3. If you click the **Tips on activating MegaRAID Advanced Software Options** link, you can access http://www.avagotech.com/products/server-storage/raid-controllers/advanced-software-licensing.

   Both the **Safe ID** field and the **Serial Number** field consists of a pre-defined value generated by the controller. Alternatively, you can copy the value and paste it in the text box for the applicable field.

4. Click **Activate**.

   The **Activate MegaRAID Advanced Software – Choose Method** wizard appears, as shown in Figure 288.

12.4 Activation Key

Figure 288 Activate MegaRAID Advanced Software Options - Choose Method Dialog

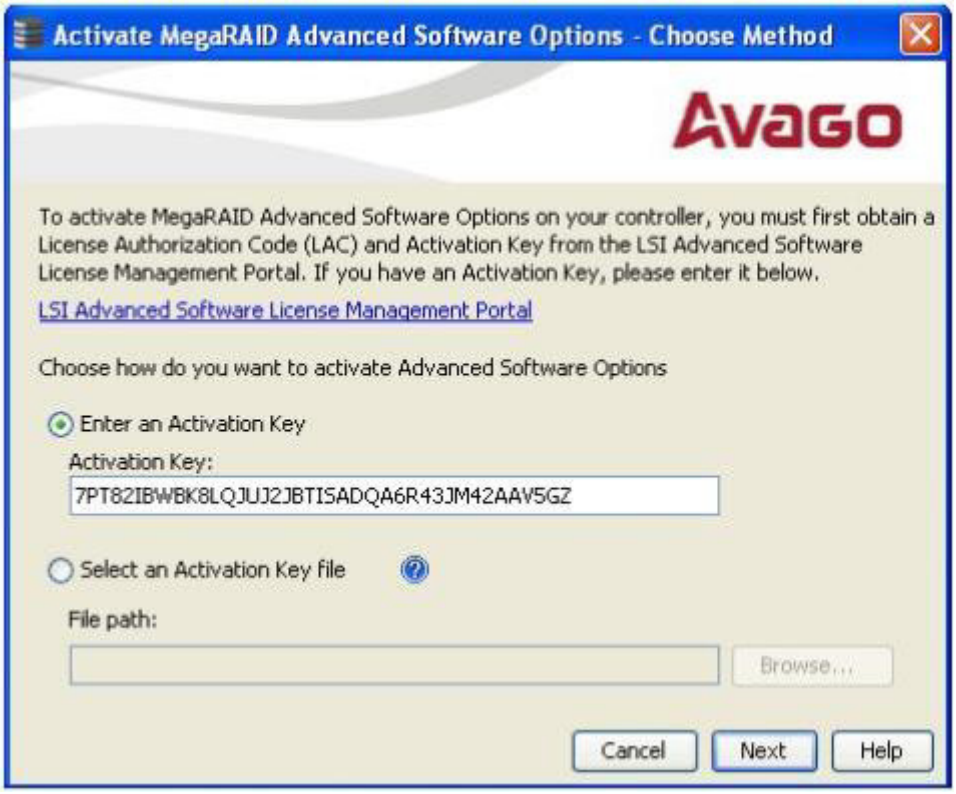

Perform the following steps to enter the activation key:

1. Click the **Avago 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** radio button, and enter the activation key in the text box provided below the **Activation Key** field.
   - Select the **Select an Activation Key file** radio button, and click **Browse** to get the path of the activation key file.
3. Click **Next**.

   After you click **Next**, one of the following two scenarios occurs:
   - The **Activate MegaRAID Advanced Software Options – Summary** dialog appears as shown in the Activate MegaRAID Advanced Software Options - Summary Dialog figure.
   - Depending on the relevant scenarios, the application responds by displaying corresponding messages as shown in Application Scenarios and Messages.

12.5 Advanced MegaRAID Software Status Summary

After you enter the activation key and click **Next**, the Activate **MegaRAID Advanced Software Option – Summary** wizard (as shown in the following figure) displays the list of the advanced softwares along with their *former status* and *new status* in the controller.

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

Figure 289 Activate MegaRAID Advanced Software Options - Summary Dialog

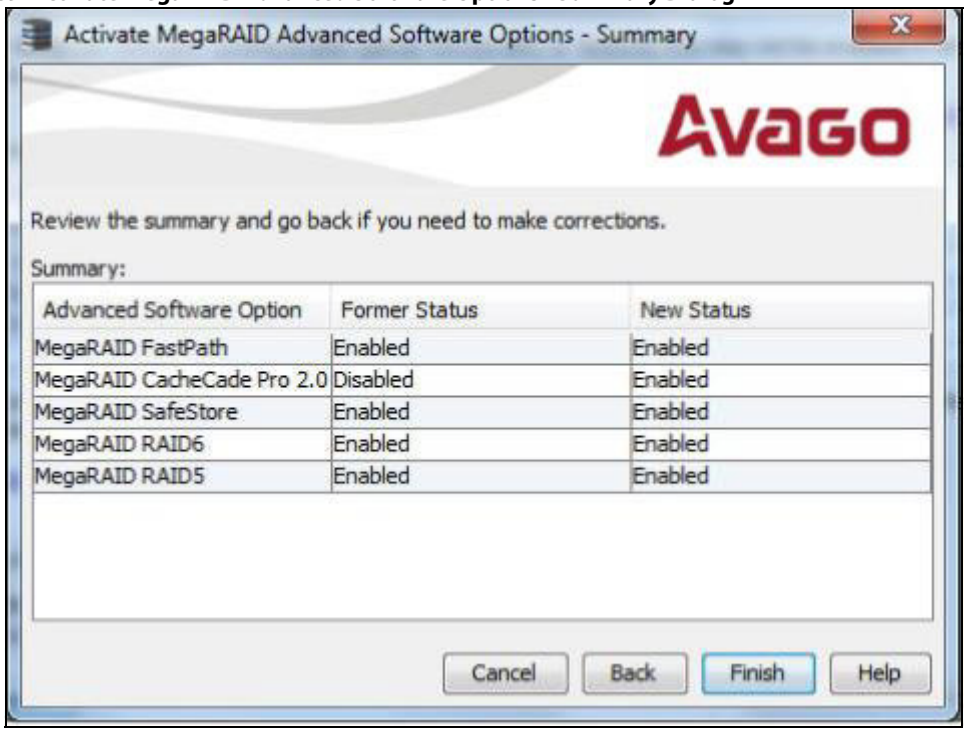

1. Click **Finish**.
   The status of the advanced software is enabled, and the advanced features are secured in the Key Vault.
2. Click **Cancel** to cancel this action.

12.6 Application Scenarios and Messages

Scenario # 1

If you enter an *invalid* activation key, the following message appears.

Figure 290 Invalid Activation Key Message

Scenario # 2

If you enter an *incorrect* activation key file, the following message appears.

Figure 291 Incorrect Activation Key File Message

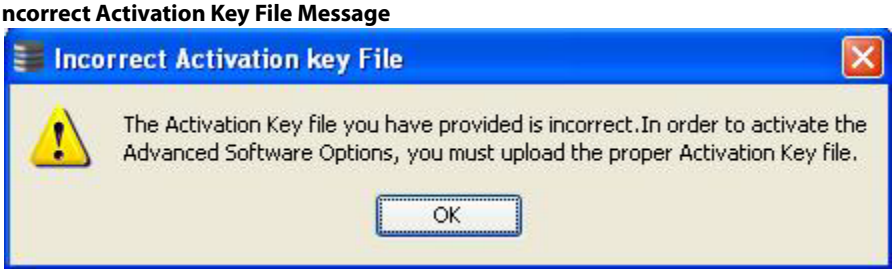

Scenario # 3

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

Figure 292 Activation Key Mismatch Message

NOTE Entering a space in the **Activation Key** field disables the **Next** button in the Activate MegaRAID Advanced Software Options - Choose Method Dialog figure.

If you click **Cancel** in the **Activate MegaRAID Advanced Software – Choose Method** dialog, as shown in the Activate MegaRAID Advanced Software Options - Choose Method Dialog figure, the following confirmation dialog box appears.

Figure 293 Activate MegaRAID Advanced Software - Confirmation Dialog

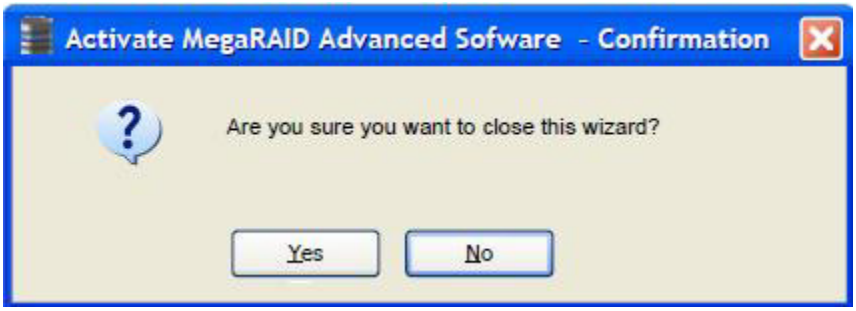

12.7 Activating an Unlimited Key over a Trial Key

When you activate an unlimited key over a trial key, a message, `The existing trial key will be deactivated and all the advanced software associated with it will be disabled`, appears (indicated in pink text in the following figure).

Figure 294 Activating an Unlimited Key over a Trial Key

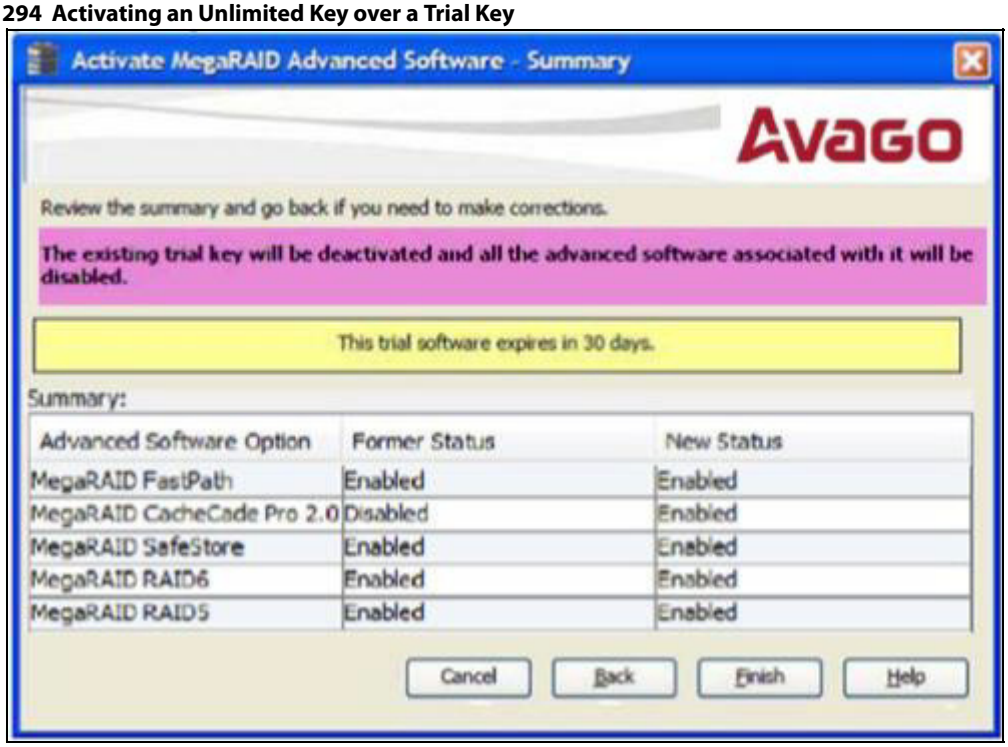

NOTE Except for the yellow shading, the other shadings of the text are provided for easy understanding in the relevant dialogs.

12.7.1 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 295 Activating a Trial Software

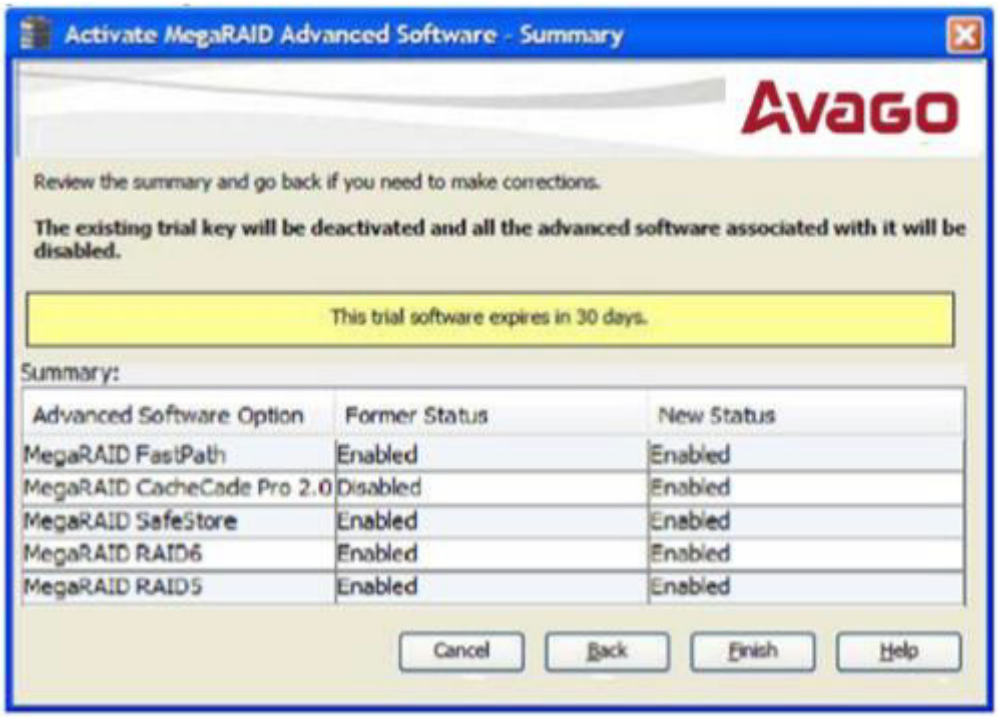

12.7.2 Activating an Unlimited Key

When you activate an unlimited key or a trial key, a message `Review the summary and go back if you need to make corrections` appears (indicated in green text in the following figure).

Figure 296 Activating an Unlimited Key

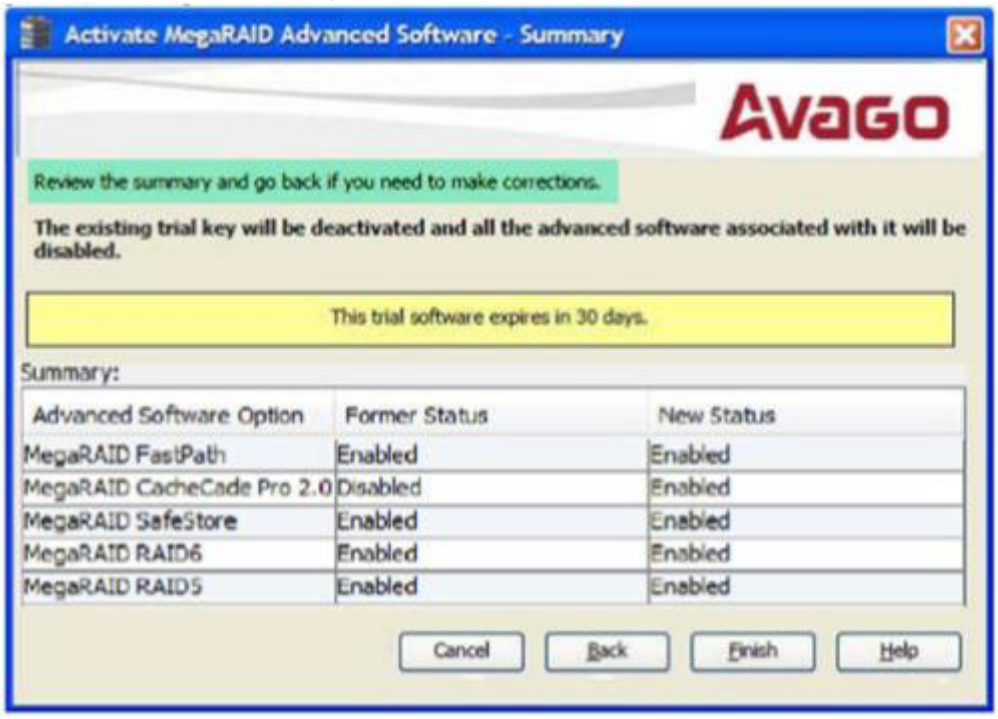

12.7.3 Reusing the Activation Key

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

Figure 297 Reusing the Activation Key

12.7.4 Securing Advanced MegaRAID Software

When you want to transfer the advanced software from the controller to the Key Vault, use the **Securing Advanced MegaRAID Software - Confirmation** wizard. This wizard is conditional, and appears only when the Key Vault and the unsecured keys exist.

1. Select any one of the following options to view the **Securing Advanced MegaRAID Software - Confirmation** wizard.
   - Select the **Physical** tab in the left panel of the MegaRAID Storage Manager window, and select a controller icon.
   - Select **Go To > Controller > Manage MegaRAID Advanced Software Options** wizard.

Figure 298 Secure MegaRAID Advanced Software - Confirmation Dialog

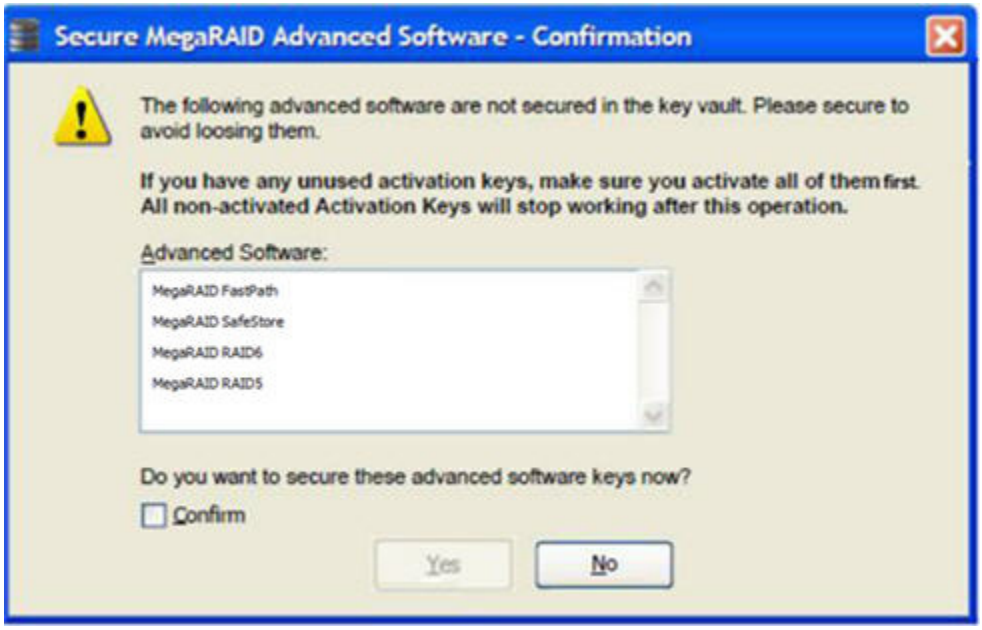

2. Select the **Confirm** check box, if you want to secure the advanced software.

   After you select the check box, the **Yes** button is enabled. This situation implies that the advanced software is secured in the keyvault.

   If the advanced software is not secured, the **Secure MegaRAID Advanced Software - Confirmation** dialog appears, as shown in the Activate MegaRAID Advanced Software - Confirmation Dialog figure.

12.8 Configuring Key Vault (Re-hosting Process)

Re-hosting is a process of transferring the advanced software features from one controller to another. To implement the re-hosting process, you must configure the **Configure Key Vault** button in the **Manage MegaRAID Software Options** wizard.

1. Choose any one of the following options to configure the Key Vault.
   - Click the **Configure Key Vault** button in the Manage **MegaRAID Advanced Software Options** wizard.
   - Select **Go To > Controller > Manage Premium Feature**.

   The **Configure Key Vault-Confirm Re-hosting Process** wizard appears, as shown in the following figure.

Figure 299 Configure Key Vault

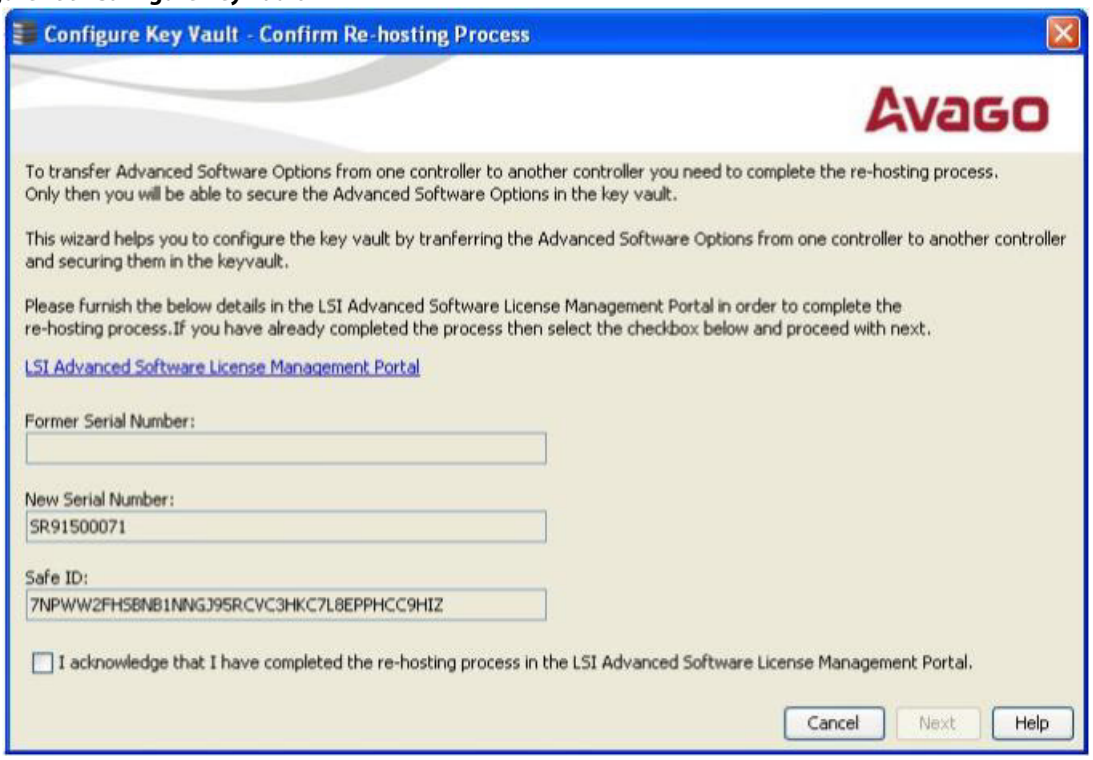

2. Select the **I acknowledge that I have completed the re-hosting process in the Avago Advanced Software License Management Portal** check box.
3. Click **Next**.

   The **Configure key Vault- Secure Advanced Software Options** wizard appears, as shown in the following figure.

Figure 300 Configure Key Vault - Secure Advanced Software Options Dialog

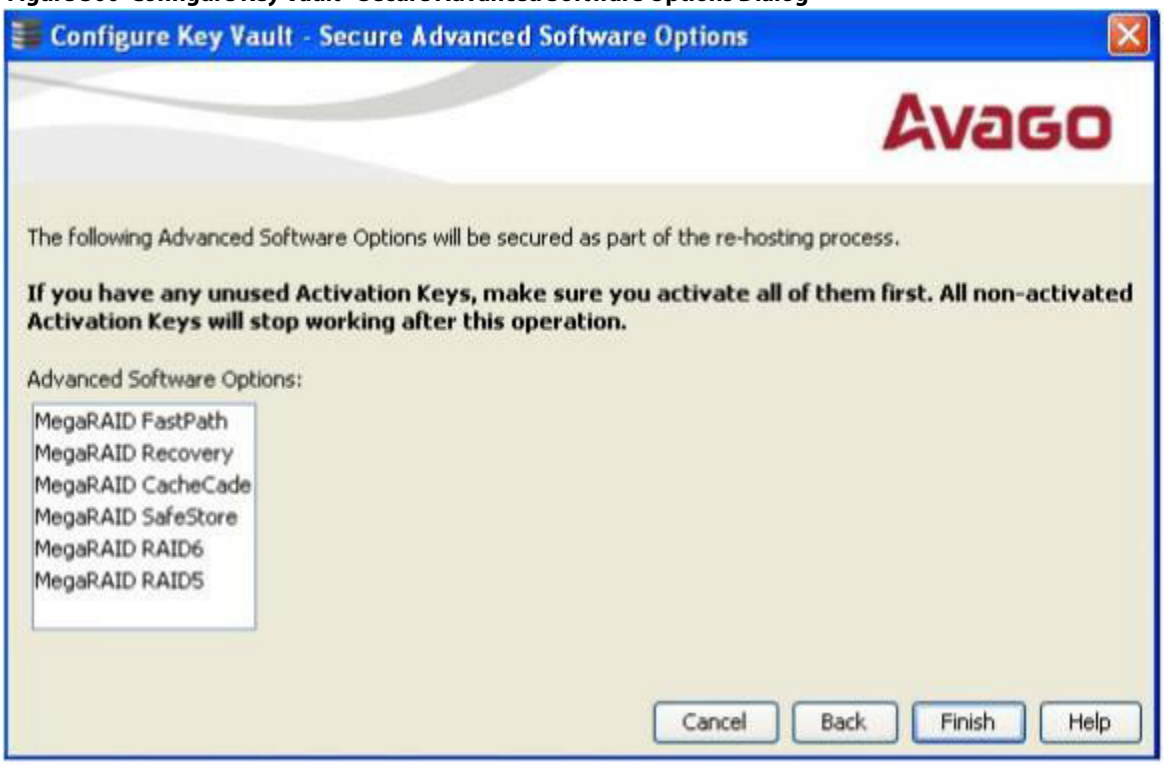

4. Click **Finish** and the advanced software options are secured in the key vault.

NOTE The **Next** button in the **Configure Key Vault** wizard is enabled only if you select the check box. This wizard is conditional and appears only if the re-hosting process is necessary, and when both the key vault and the unsecured keys are present at the same time.

12.9 Re-hosting Complete

If you want to transfer the advanced software options from one controller to another, use the re-hosting process. The re-hosting process makes sure that these options are secured in the Key Vault. You have to configure the Key Vault to complete the re-hosting process.

1. Choose any one of the following options to complete the re-hosting process.
   - Click the **Configure Key Vault** button from the **Manage MegaRAID Advanced Software Options** wizard.
   - Select **Go To > Controller > Manage MegaRAID Advanced Software Options** wizard.

   The **Re-Hosting Process - Complete** wizard appears, as shown in the following figure.

Figure 301 Re-Hosting Process - Complete Dialog

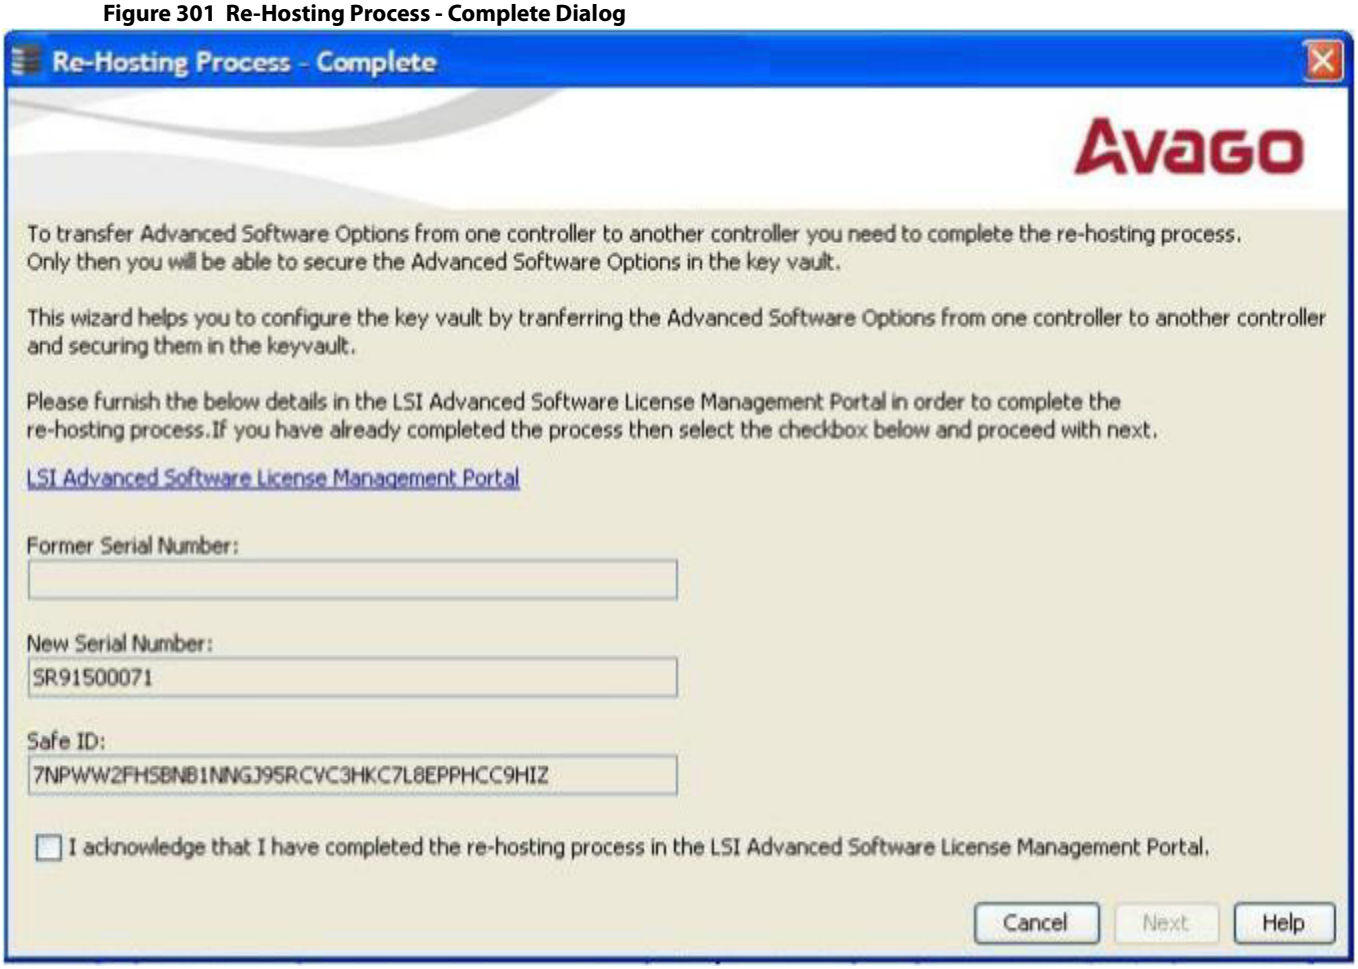

2. Select the **I acknowledge that I have completed the re-hosting process in the Avago Advanced Software License Management Portal** check box if you want to complete the re-hosting process.
   This setting makes sure that the advanced software features are transferred to the controller.
3. Click **Cancel** if you do not want to activate the re-hosting process.

12.10 Deactivate Trial Software

When you want to deactivate a trial software, use the **Deactivate All Trial Software** wizard.

Perform the following steps to enable the deactivate trial software button:

1. Click **Deactivate All Trial Software** in the **Manage MegaRAID Advanced Software Options** dialog.
   The **Deactivate All Trial Software - Confirmation** dialog appears, as shown in the following figure.

Figure 302 Deactivate All Trial Software - Confirmation Dialog

2. Select the **Confirm** check box, if you want to deactivate the software applications, that are used with a trial key.
3. Click **Yes**.
   The trial software is deactivated.

12.11 Using the MegaRAID CacheCade Advanced Software

The MegaRAID CacheCade software provides you with read caching capability.

Perform the following steps to use the CacheCade advanced software.

1. Click a RAID controller icon in the left frame.
2. Select **Go To > Controller > Create CacheCade - SSD Caching** on the menu bar.
   The wizard dialog appears.
3. Click unconfigured CacheCade - SSD Caching drives in the left frame to select the drives for the CacheCade drive group, as shown in the following figure.

Figure 303 Create CacheCade Drive group Dialog

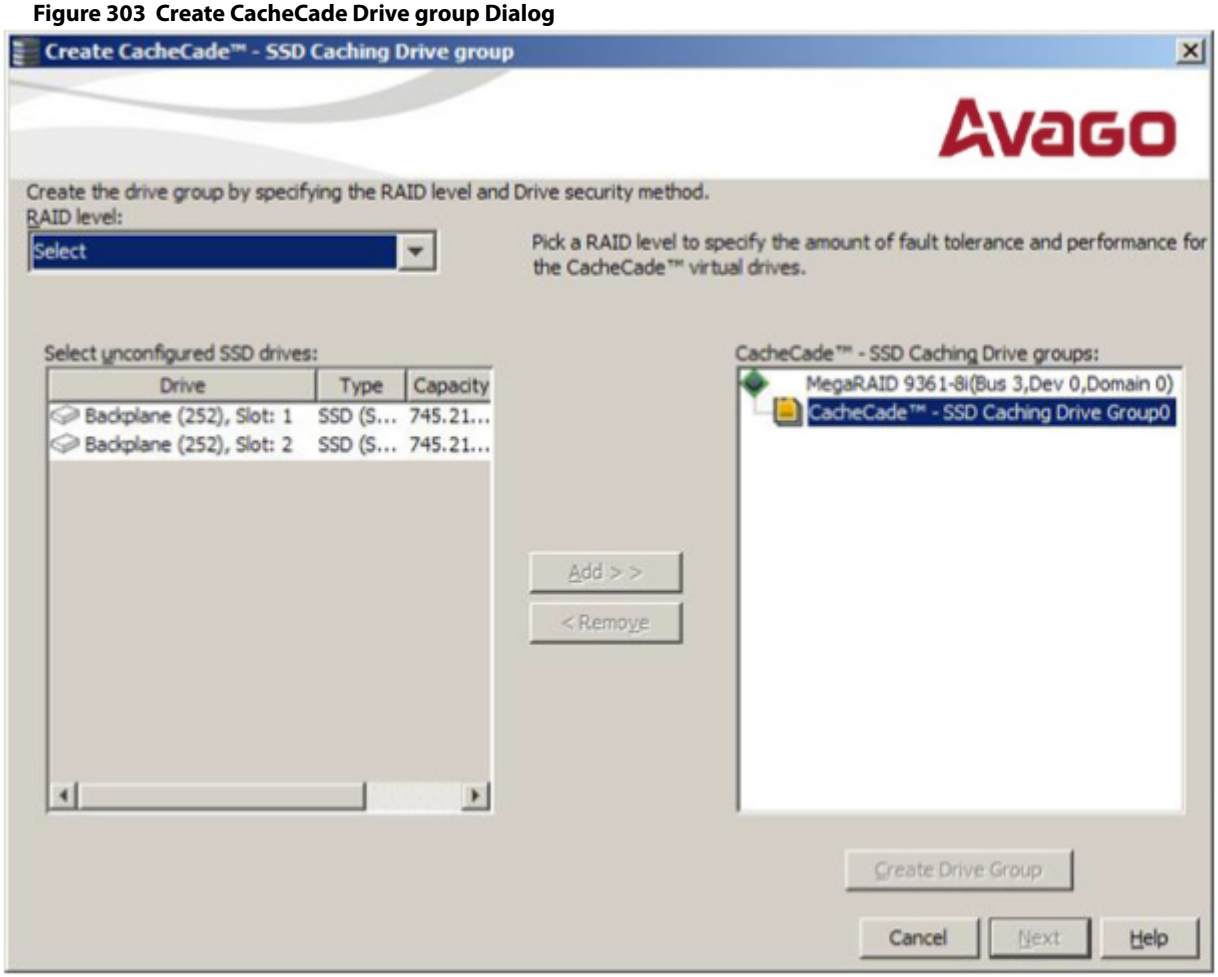

After you select the unconfigured drives, the **Add >>** button is available.

4. Click **Add >>** to move the selected drives to the drive group in the right frame, as shown in the following figure.

Figure 304 Create CacheCade Drive group Dialog

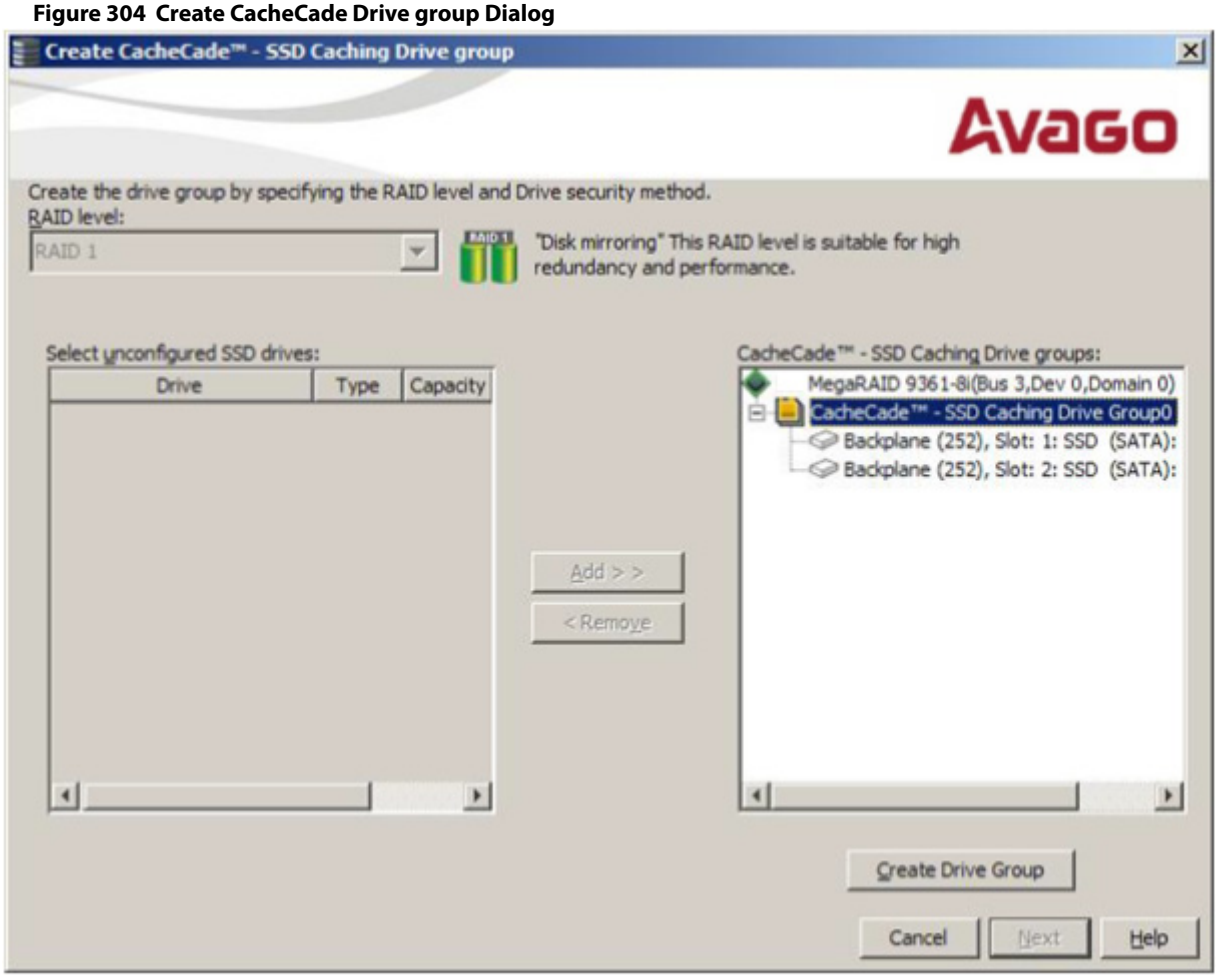

After you move the selected drives, the **Create Drive Group** button is available.

5. Click **Create Drive Group**.
6. Click **Next**.

   Use the next dialog that appears to select parameters for the cache disk.
7. Enter a name for the CacheCade - SSD Caching virtual drive in the **CacheCade - SSD Caching VD name** field, and click **Create Virtual Drive**.

   Depending on the number of drives, you might have the option to set the capacity of the CacheCade - SSD Caching drive.

   The CacheCade drive group icon appears in the menu dialog, as shown in the following figure.

Figure 305 Create CacheCade™ - Summary Dialog

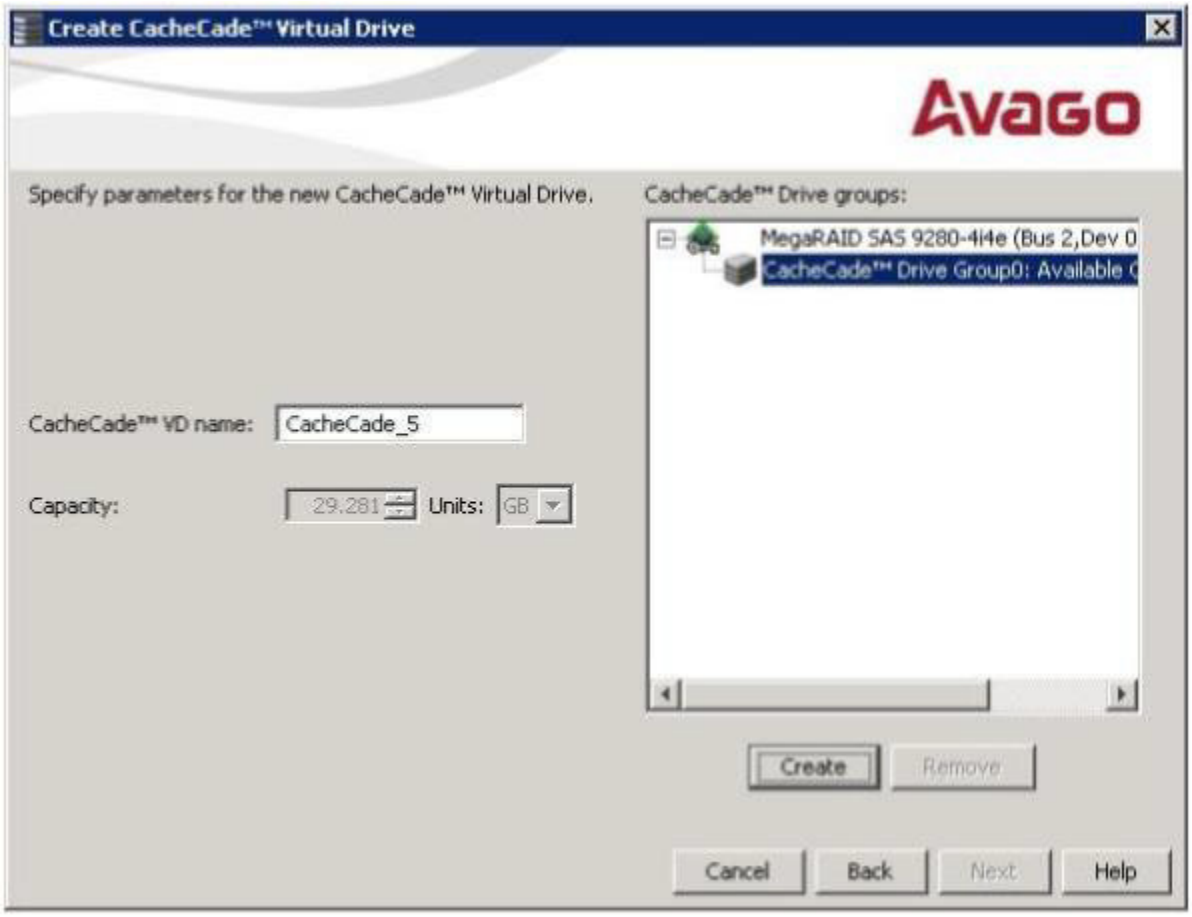

8. Click **Next**.

   The summary dialog appears, as shown in the following figure. This dialog displays the drive group name, the number of drives, the total capacity, the free capacity, the CacheCade virtual drive name, and the capacity being used.

Figure 306 CacheCade Virtual Drive Summary Dialog

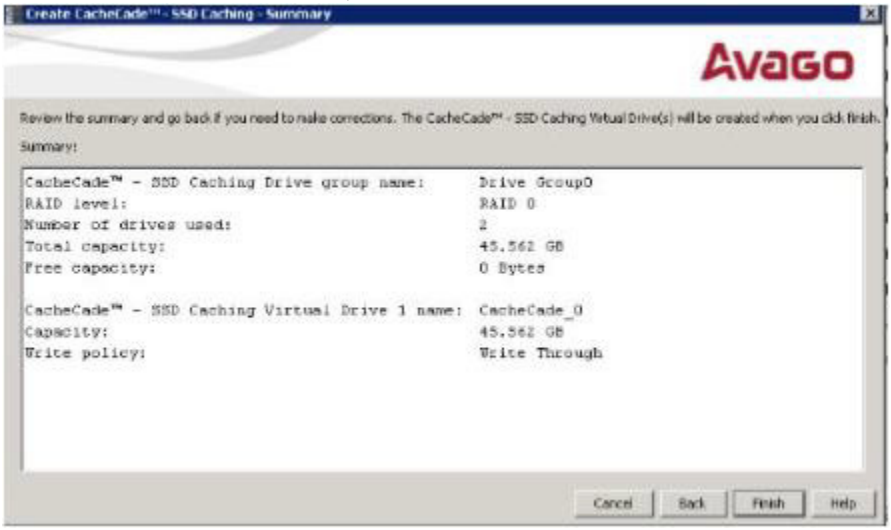

9. Click **Finish**.

   A confirmation message displays after the CacheCade virtual drive is successfully created.

   The CacheCade drive icon appears next to the RAID controller in the left frame, in the MegaRAID Storage Manager main window.

12.12 Using the MegaRAID CacheCade Pro 2.0 Software

The MegaRAID CacheCade Pro 2.0 software provides you with read and write caching capability.

NOTE The MegaRAID firmware has the provision to monitor I/O performance; changes have been made to accommodate the CacheCade Pro 2.0 software statistics. The CacheCade Pro 2.0 software metrics are captured for each logical drive that has CacheCade enabled. The CacheCade Pro 2.0 software gathers information about the cache windows allocated for a logical drive, the number of new windows allocated in this metrics collection period, the number of windows that are actively used, and the window hit rates.

Perform the following steps to use the CacheCade Pro 2.0 software:

1. Perform one of these actions:
   - Right-click on a controller in the device tree in the left frame of the **MegaRAID Storage Manager** window and select **Create CacheCade SSD Caching**.
   - Select a controller and select **Go To > Controller > Create CacheCade SSD Caching** in the menu bar.

   The **CacheCade SSD Caching** wizard appears, as shown in the following figure.

Figure 307 CacheCade SSD Caching Wizard - First Screen

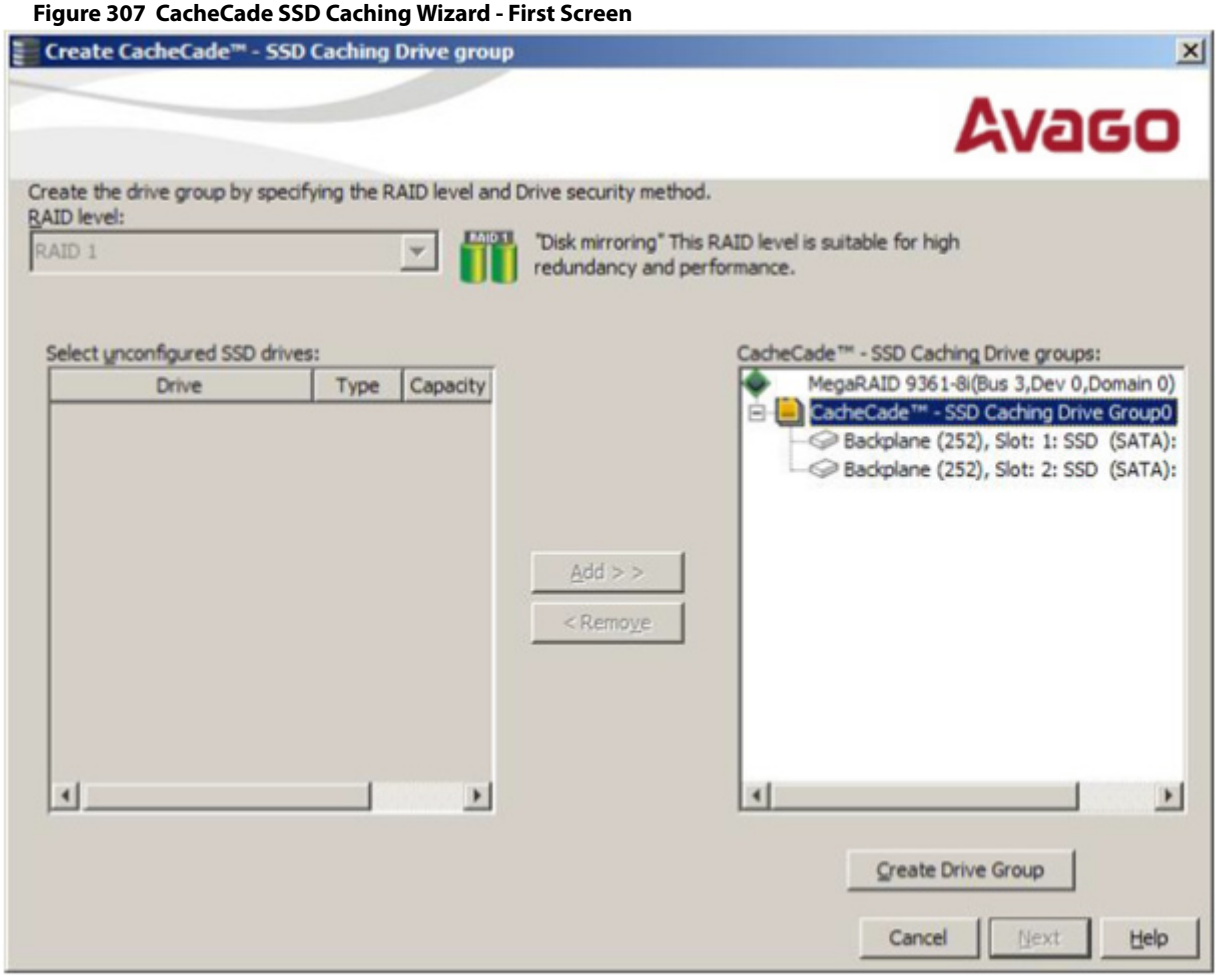

2. Select a RAID level for the CacheCade virtual drive in the **RAID level** field.
3. Select an unconfigured SSD drive, for the selected RAID level, from **Select unconfigured SSD Drives** in the left frame.

   After you select an unconfigured SSD Drive, the **Add** button is enabled.
4. Click **Add** to add the selected drive to the CacheCade - SSD Caching Drive groups in the right frame.

   After you click **Add**, the **Create Drive Group** button is enabled.
5. Click **Create Drive Group**.

   The newly created drive group appears in CacheCade SSD Caching Drive groups in the right frame.
6. Click **Next**.

   The next wizard screen appears.

Figure 308 Parameters for CacheCade SSD Caching Virtual Drive

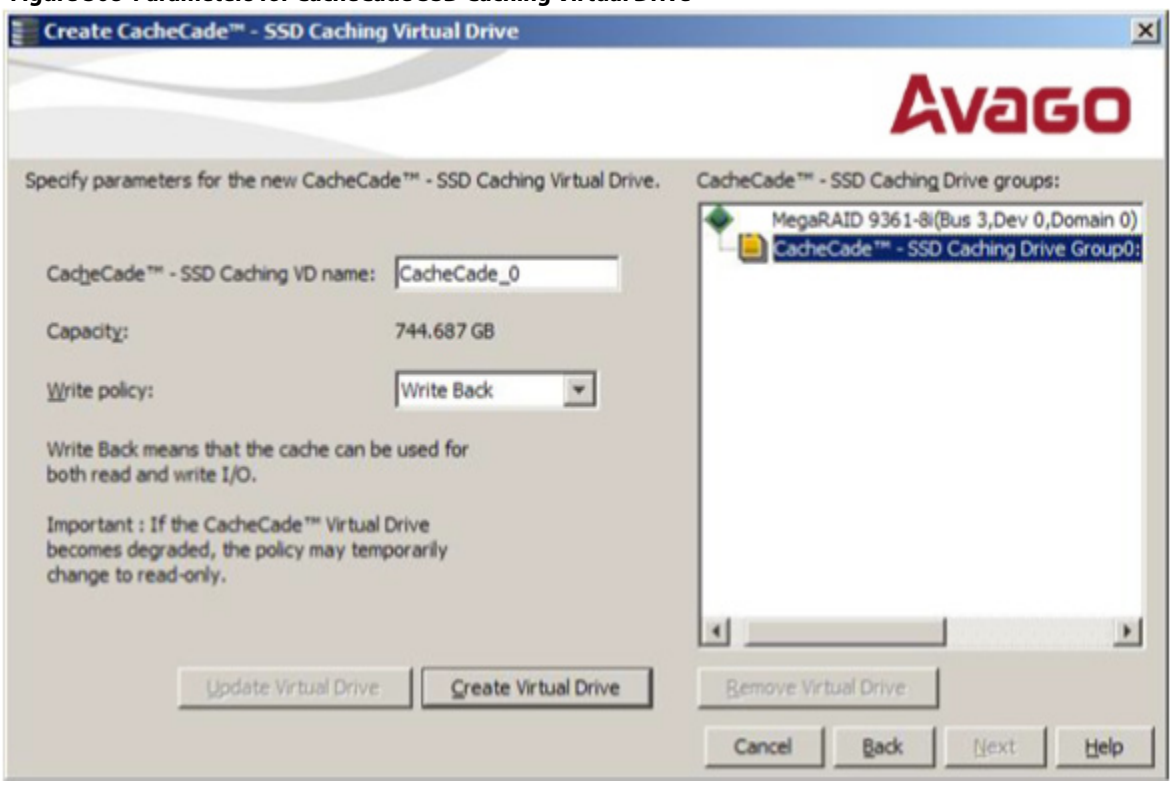

7. Enter a name for the CacheCade virtual drive in the **CacheCade - SSD Caching VD name** field.
8. Select a write policy from the **Write policy** drop-down list.

   A description of the selected write policy appears below.
9. Click **Create Virtual Drive**.

   The newly created virtual drive appears in the CacheCade SSD Caching Drive groups in the right frame. The **Remove Virtual Drive** button is enabled. You can select the newly created virtual drive and click **Remove Virtual Drive** to delete the virtual drive.
10. Click **Next**.

    The summary screen appears.

Figure 309 Create CacheCade - SSD Caching - Summary

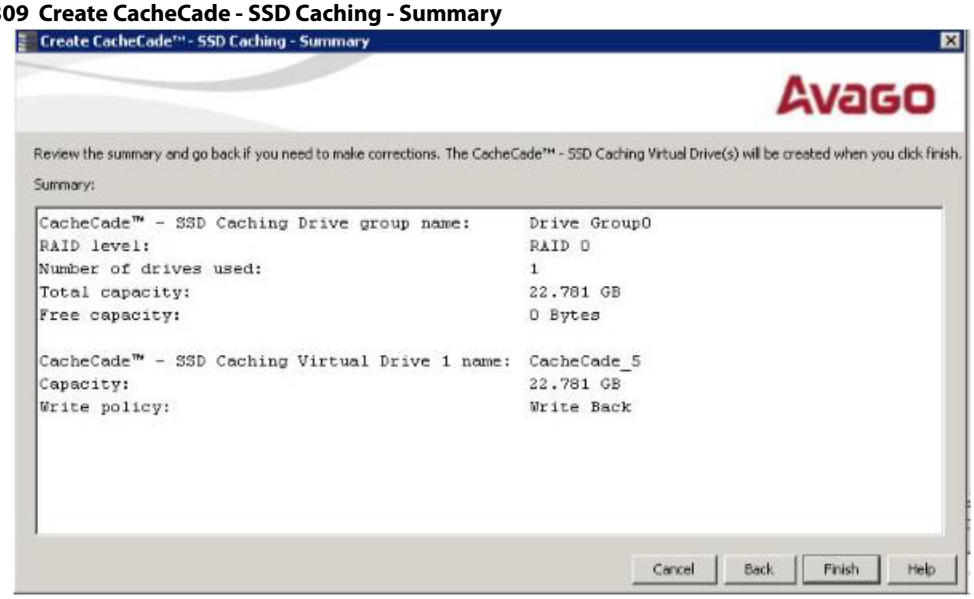

This screen displays the drive group name, the RAID level, the number of drives, the total capacity, the free capacity, the CacheCade virtual drive name, the capacity being used, and the write policy.

11. Click **Finish**.

    A confirmation message displays after the CacheCade virtual drive is successfully created. The CacheCade drive icon appears next to the RAID controller in the left frame in the **MegaRAID Storage Manager** window.

12.12.1 Modifying the CacheCade Virtual Drive Properties

You can modify the name and the write policy of a CacheCade virtual drive any time after a CacheCade virtual drive is created. Perform the following steps to change the virtual drive properties:

1. Perform one of these actions:
   - Right-click on a controller in the device tree in the left frame of the **MegaRAID Storage Manager** window, and select **Set Virtual Drive Properties**.
   - Select a controller, and select **Go To > Virtual Drive > Set Virtual Drive Properties**.

   The **Set Virtual Drive Properties** dialog appears, as shown in the following figure.

Figure 310 Set Virtual Drive Properties

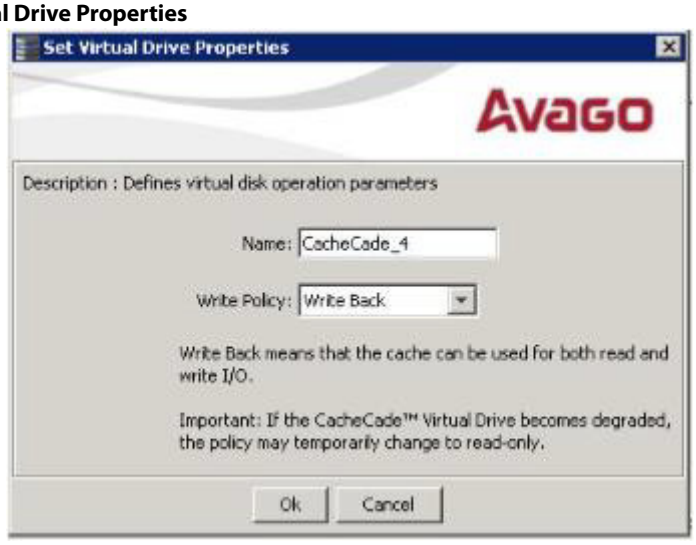

2. Edit the name of a CacheCade virtual drive in the **Name** field.
3. Select a write policy from the **Write Policy** drop-down list.
4. Click **OK**.

   A confirmation dialog appears with a warning note.
5. Select the **Confirm** check box, and click **OK**.

12.12.2 Enabling SSD Caching on a Virtual Drive

You can enable SSD caching on a virtual drive. When you enable SSD caching on a virtual drive, that virtual drive becomes associated with an existing or with a future CacheCade SSD Caching virtual drive. This option is only available when the virtual drive's caching is currently disabled.

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

1. Perform one of these actions:
   - Right-click on a virtual drive in the left frame of the **MegaRAID Storage Manager** window, and select **Enable SSD Caching**.
   - Select a virtual drive, and select **Go To > Virtual Drive > Enable SSD Caching**.

   The **Enable SSD Caching** dialog appears, as shown in the following figure.

Figure 311 Enable SSD Caching

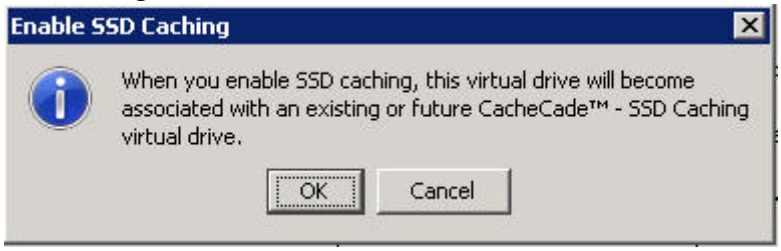

2. Click **OK** to enable caching for that virtual drive.

12.12.3 Disabling SSD Caching on a Virtual Drive

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

Perform the following steps to enable SSD Caching on a virtual drive:

1. Perform one of these actions:
   - Right-click on a virtual drive in the left frame of the **MegaRAID Storage Manager** window, and select **Disable SSD Caching**.
   - Select a virtual drive, and select **Go To > Virtual Drive > Disable SSD Caching**.

   The **Disable SSD Caching** dialog appears, as shown in the following figure.

Figure 312 Disable SSD Caching

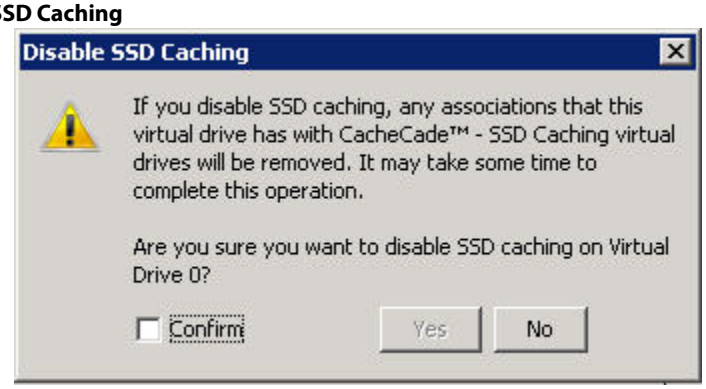

2. Select the **Confirm** check box, and click **Yes** to disable caching for that virtual drive.

12.12.4 Enabling or Disabling SSD Caching on Multiple Virtual Drives

You can enable or disable SSD caching on multiple virtual drives at one go.

Perform the follow steps to enable or disable SSD caching on multiple drives:

1. Perform one of these actions:
   - Right-click a controller in the left frame of the **MegaRAID Storage Manager** window, and select **Manage SSD Caching**.
   - Select a controller, and select **Go To > Controller > Manage SSD Caching**.

   The **Manage SSD Caching** dialog appears, as shown in the following figure.

Figure 313 Manage SSD Caching

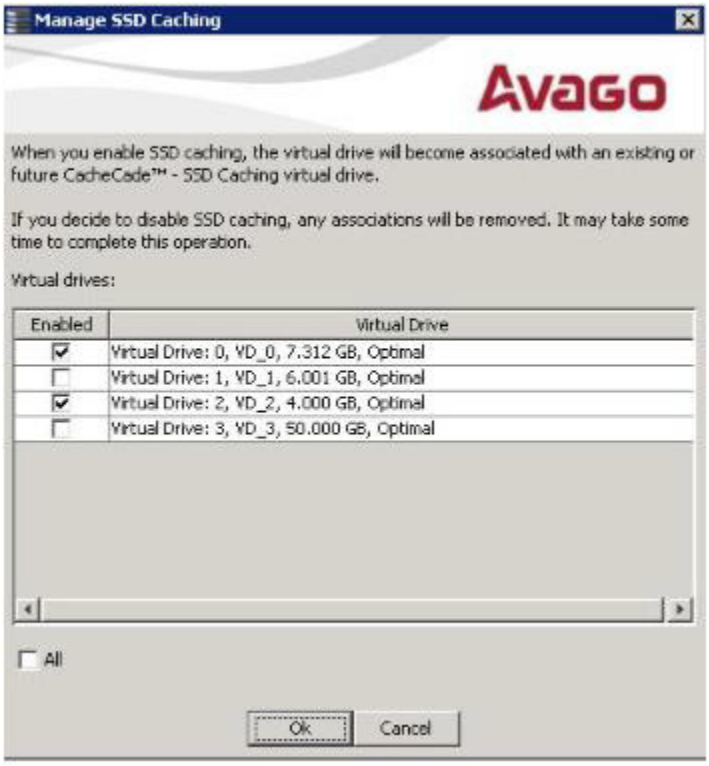

The virtual drives that have SSD caching enabled, have the check boxes next to them selected. The virtual drives that have SSD caching disabled, have deselected check boxes.

2. Select or deselect a check box to change the current setting of a virtual drive.
3. Click **Ok**.

   If you select the **All** check box, all the virtual drives are enabled. If you deselect the **All** check box, all the virtual drives are disabled.

   If you disable SSD caching on a virtual drive, the **Disable SSD Caching** dialog appears.
4. Select the **Confirm** check box, and click **OK** to enable/disable SSD caching on the selected virtual drives.

12.12.5 Modifying a CacheCade Drive Group

Perform the following steps to modify an existing CacheCade SSD caching drive group:

1. Delete the drive group.
2. Create a new CacheCade drive group.

12.12.6 Clearing Configuration on CacheCade Pro 2.0 Virtual Drives

You can clear all existing configurations on a selected controller that has CacheCade Pro 2.0 virtual drives.

1. Perform one of these actions:
   - Right-click on a controller in the left frame of the **MegaRAID Storage Manager** window, and select **Clear Configuration**.
   - Select a controller, and select **Go To > Controller > Clear Configuration**.

   The **Confirm Clear Configuration** dialog appears as shown, in the following figure.

Figure 314 Confirm Clear Configuration

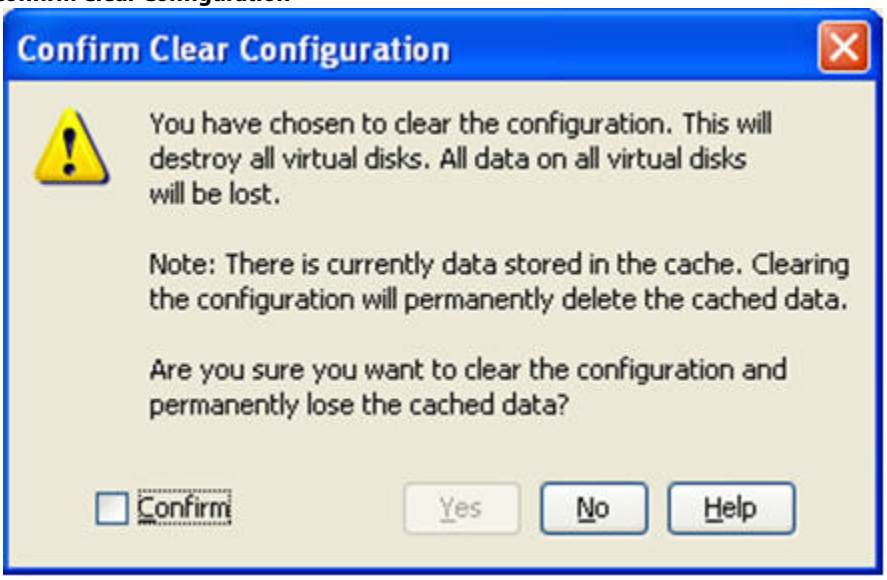

2. Select the **Confirm** check box, and click **Yes**.

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

3. Select the **Confirm** check box, and click **Yes**.

12.12.7 Removing Blocked Access

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

An icon appears in front of the affected virtual drive, next to the *Optimal* status.

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

If it is not possible to bring the CacheCade virtual drive to its optimal status, follow these steps to remove the blocked access from the virtual drive:

1. Right-click on the icon on the virtual drive with the blocked access, and select **Remove Blocked Access**.

   The **Confirm Remove Blocked Access** dialog appears, as shown in the following figure.

Figure 315 Confirm Remove Blocked Access

2. Select the **Confirm** check box, and click **Yes**.

12.12.8 Deleting a Virtual Drive with SSD Caching Enabled

You can delete a virtual drive that has SSD caching enabled on it.

Perform the following steps to delete the virtual drive:

1. Perform one of these actions:
   - Right-click on a CacheCade virtual drive, and select **Delete Virtual Drive**.
   - Select a CacheCade virtual drive and click **Go To > Virtual Drive > Delete Virtual Drive**.

   The **Confirm Delete Virtual Disk** dialog appears, as shown in the following figure.

Figure 316 Confirm Delete Virtual Disk

2. Select the **Confirm** check box, and click **Yes**.

ATTENTION If you select the **Force the delete to complete quickly** check box to delete the virtual drive, the data is not flushed before deleting the virtual drive. In this scenario, if you create this virtual drive after deleting it, there will be no data available.

12.13 Fast Path Advanced Software

MegaRAID Fast Path is a high-performance I/O accelerator for the CacheCade software drive groups connected to a MegaRAID controller card. The CacheCade software has a read performance advantage over HDDs and uses less power. This feature dramatically boosts storage subsystem bandwidth and overall transactional application performance when used with a 12Gb/s MegaRAID SATA+SAS controller.

Fast Path is a high-performance I/O accelerator for SSDs. Fast Path improves the I/O performance. If no SSDs are attached, Fast Path is not used.

12.13.1 Setting Fast Path Options

Perform the following steps to use the Fast Path advanced software:

1. Select the **Logical** tab on the **MegaRAID Storage Manager** window for the Logical view.
2. Select a virtual drive icon in the left frame.

3. Select **Go To > Virtual Drive > Set Virtual Drive Properties** on the menu bar.

   The **Set Virtual Drive Properties** dialog appears. It shows the default settings for the Fast Path advanced software:

   — Write Policy: **Write Thru**
   — IO Policy: **Direct IO**
   — Read Policy: **No Read Ahead**
   — Disk Cache Policy: **Disabled**
   — Strip Size: **64KB**
4. Click **OK**.

   A confirmation dialog displays.
5. Select the **Confirm** check box, and click **Yes** to confirm that you want to set the virtual drive properties.

12.14 Avago MegaRAID SafeStore Encryption Services

Avago SafeStore Encryption Services offer the ability to encrypt data on the drives and use the drive-based key management to provide data security. This solution provides data protection in the event of theft or loss of physical drives. If you remove a self-encrypting drive from its storage system or the server in which it resides, the data on that drive is encrypted, and becomes useless to anyone who attempts to access it without the appropriate security authorization.

12.14.1 Enabling Drive Security

This section describes how to enable, change, and disable the drive security, and how to import a foreign configuration using the SafeStore Encryption Services advanced software.

To enable drive security, the following details must be specified:

- **Security key identifier** - The controller, by default, assigns a security key identifier. However, you can change this security key identifier as per your requirement. If you have more than one security key identifier, the controller helps you to deteSection 317, Enable Drive Security – Security Key IdentifierSection 9.1.4rmine which security key identifier to enter.
- **Security key** - Provides you with an option to create secure virtual drives by specifying the security key. The security key provided by you is used to lock each self-encrypted drive attached to the controller.
- **Suggest Security Key** - Alternatively, you can click this option to have the system create a security key for you.
- **Password** - You can also specify a password to provide additional drive security.
- **Pause for password at boot time** and **Enforce strong password security** - If you select the **Pause for password at boot time**, you are prompted to provide the password each time you restart your server. If you select **Enforce strong password security**, the system enforces you to specify a strong password.
- **Show Key and Show Password** - You can either select or clear the**Show Key** and **Show Password** check boxes. By default, they are unchecked.

To enable drive security, perform the following steps:

1. Navigate to the **Physical** tab in the left panel of the **MegaRAID Storage Manager** window, and select a controller.
2. Select **Go To > Controller > Enable Drive Security**.

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

Figure 317 Enable Drive Security – Security Key Identifier

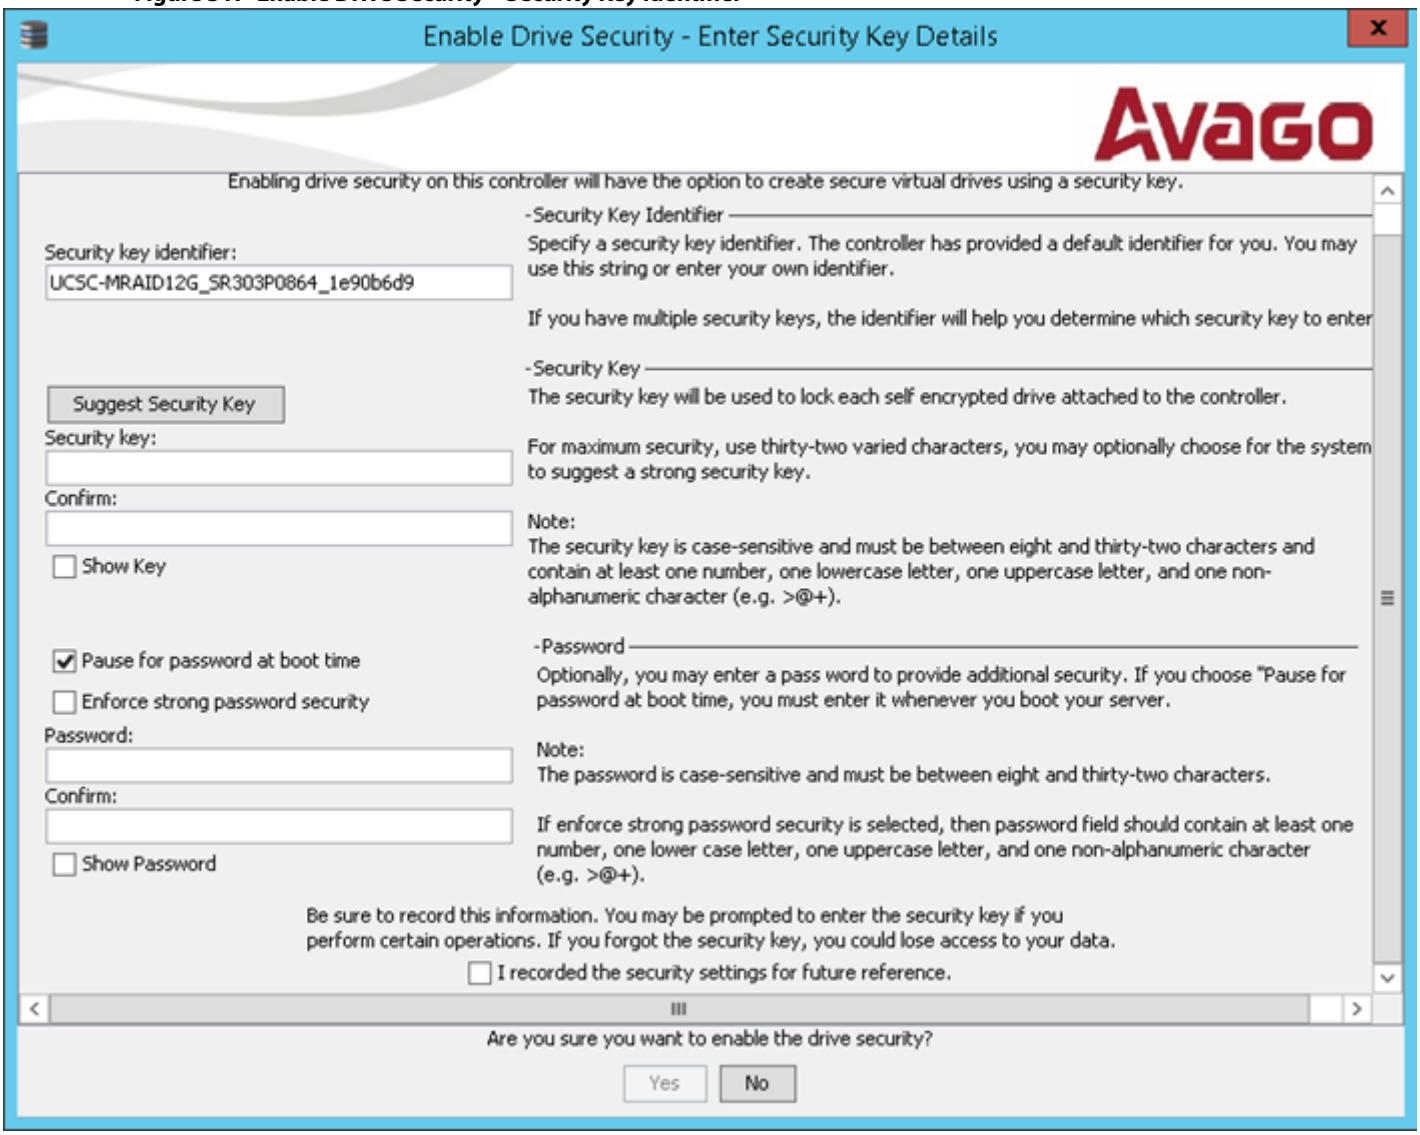

3. Either use the default security key identifier provided by the controller or specify a new security key identifier.

   NOTE If you create more than one security key, ensure that you change the security key identifier. Otherwise, you cannot differentiate between the security keys.

4. Either click **Suggest Security Key** to have the system create a security key for you or enter a new security key and confirm.
5. (Optional) Select the **Show Key** check box.

   If you choose this option, the security key that you specify or the security key that is created by the system if you have clicked on Suggest Security Key, will be visible to you. If you do not select this option, the security key will not be visible to you.

CAUTION **Ensure that you note down this security key somewhere for future reference. If you are unable to provide the security key when it is required by the system, you will lose access to your data.**

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 (for example, < > @ +). The space character is not permitted.

NOTE Non-U.S. keyboard users must be careful not to enter double-byte character set (DBCS) characters in the security key field. The firmware works with the ASCII character set only.

6. (Optional) Select the **Pause for password at boot time** check box.

   If you choose this option, you are prompted to provide the password each time you restart your server.

7. (Optional) Select the **Enforce strong password security** check box.

   If you choose this option, make sure the password is between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g. < > @ +). The space character is not permitted. The password is case-sensitive.

8. (Optional) Enter a password in the **Password** field and confirm the same password once again in the **Confirm** field.

9. (Optional) Select the **Show Password** check box.

   If you choose this option, the password that you specify will be visible to you. If you do not select this option, the password will not be visible to you.

   Warning messages appear if there is a mismatch between the characters entered in the **Password** field and the **Confirm** field, or if you have entered an invalid character.

   CAUTION **Ensure that you note down this password somewhere for future reference. If you are unable to provide the password when it is required by the system, you will lose access to your data.**

10. Select the **I recorded the security settings for future reference** check box, then click **Yes** to confirm that you want to enable drive security on this controller and have recorded the security settings for future reference.

    The **MegaRAID Storage Manager**enables drive security and returns to the main menu.

12.14.2 Enabling Drive Security

This section describes how to enable, change, and disable the drive security, and how to import a foreign configuration using the SafeStore Encryption Services advanced software.

To enable security on the drives, you need to perform the following actions to set drive security:

- Enter a security key identifier.
- A security key identifier appears whenever you have to enter a security key. If you have more than one security key, the identifier helps you determine which security key to enter.
- Enter a security key.
- After you create a security key, you have the option to create secure virtual drives using the key. You have to use the security key to perform certain operations.
- You can improve security by entering a password. To provide additional security, you can require the password whenever anyone boots the server.

Perform the following steps to enable drive security.

1. Select the **Physical** tab in the left panel of the **MegaRAID Storage Manager** window, and select a controller icon.

2. Select **Go To > Controller > Enable Drive Security**.
   The **Enable Drive Security** dialog appears, as shown in the following figure.

Figure 318 Enable Drive Security – Security Key Identifier

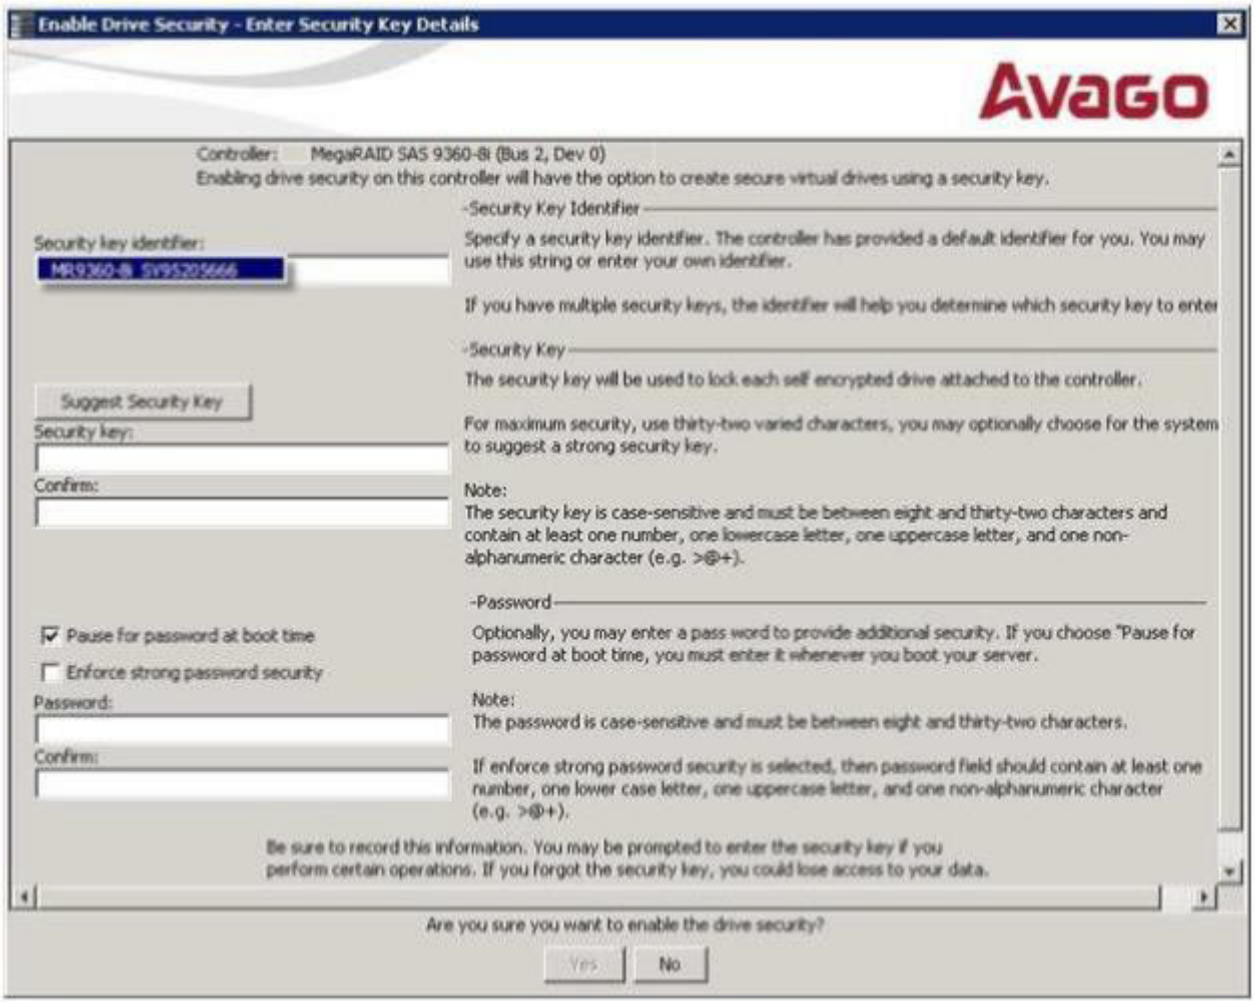

3. Either use the default security key identifier, or enter a new security key identifier.

   NOTE If you create more than one security key, make sure that you change the security key identifier. Otherwise, you cannot differentiate between the security keys.

4. Either click **Suggest Security Key** to have the system create a security key, or you can enter a new security key.
5. Enter the new security key again to confirm.

   CAUTION **If you are prompted for the security key and you forgot it or don't have access to it, you will lose access to your data.** Make 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 (for example, < > @ +). The space character is not permitted.

NOTE Non-U.S. keyboard users must be careful not to enter double-byte character set (DBCS) characters in the security key field. The firmware works with the ASCII character set only.

The following figure shows the security key entered and confirmed on this dialog.

Figure 319 Enable Drive Security - Security Key

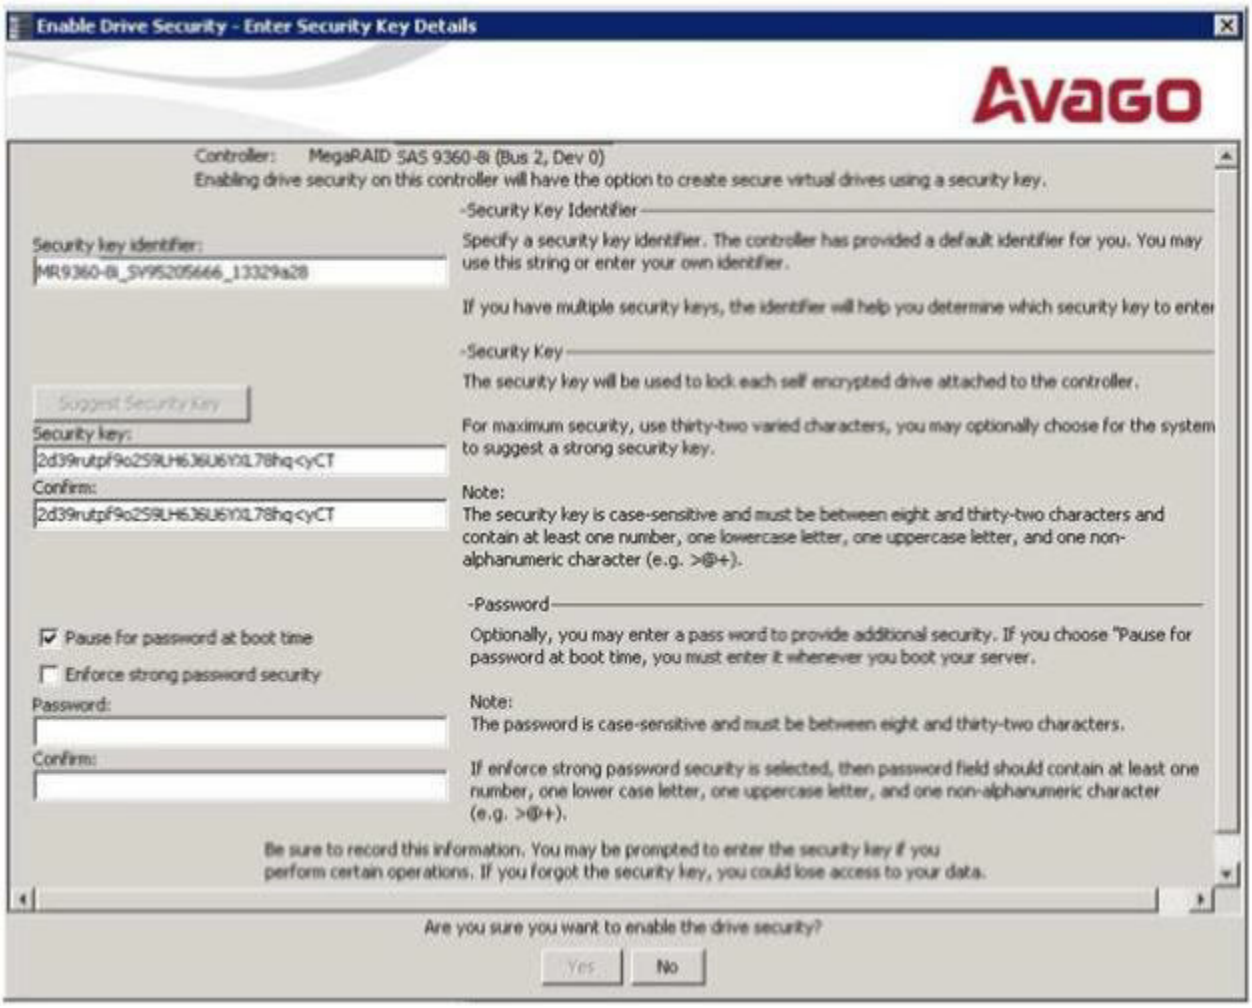

6. (Optional) Select the **Pause for password at boot time** check box.

   If you choose this option, you must enter the password whenever you boot the server.

7. (Optional) Select the **Enforce strong password security** check box.

   If you choose this option, make sure the password is between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g. < > @ +). The space character is not permitted. The password is case-sensitive.

8. (Optional) Enter a password in the **Password** field and then enter the same password in the **Confirm** field.

   Warning messages appear if a mismatch exists between the characters entered in the **Password** field and the **Confirm** field, or if there is an invalid character entered.

   NOTE Be sure to record the password. If you lose the password, you could lose access to your data.

   The following figure shows the password entered and confirmed on this dialog.

Figure 320 Enable Drive Security - Password

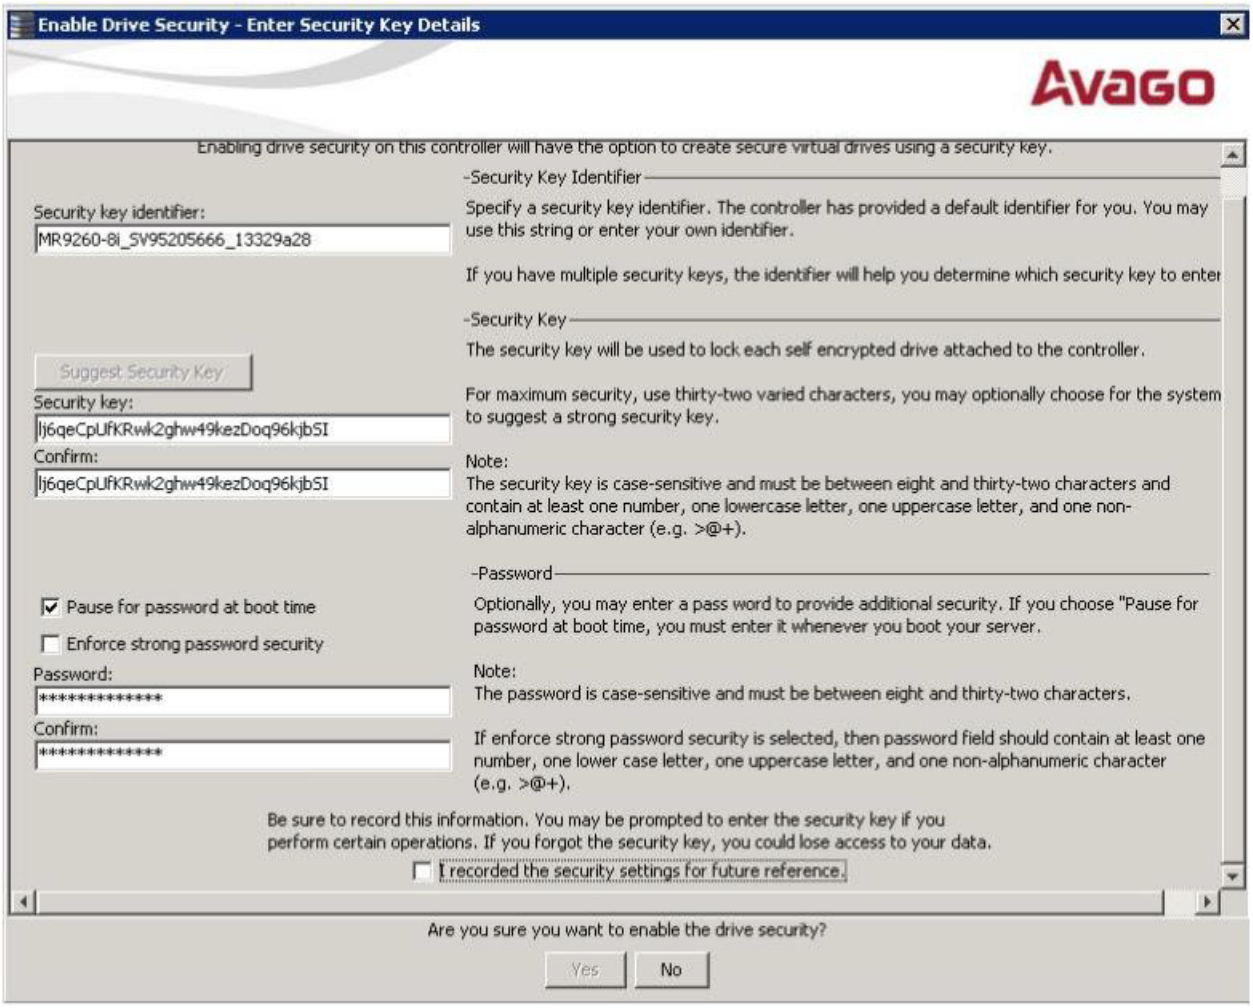

ATTENTION **If you forget the security key, you will lose access to your data.** Be sure to record your security key. You might need to enter the security key to perform certain operations.

9. Select the **I recorded the security settings for future reference** check box, and click **Yes** to confirm that you want to enable drive security on this controller and have recorded the security settings for future reference.

   The MegaRAID Storage Manager software enables drive security and returns you to the main menu.

12.14.3 Changing 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** window, and select a controller icon.
2. Select **Go To > Controller > Change Drive Security**.

   The **Change Security Settings – Introduction** dialog appears. This dialog lists the actions you can perform, which include editing the security key identifier, security key, and the password.
3. Either keep the existing security key identifier, or enter a new security key identifier.

NOTE If you change the security key, you need to change the security key identifier. Otherwise, you cannot differentiate between the security keys.

4. Either select the **Use the existing drive security key** radio button to use the existing drive security key, or enter a new security key and then enter the new security key again to confirm.

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

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

   NOTE Non-U.S. keyboard users must be careful not to enter DBCS characters in the Security Key field. The firmware works with the ASCII character set only.

5. If desired, click the option to use a password in addition to the security key.
6. 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 password can hold up to 32 characters. The key must be at least 8 characters.

   The next dialog that appears describes the changes you made and asks you whether you want to confirm these changes.
7. Click the check box to confirm that you have recorded the security settings for future reference, and click **Yes** to confirm that you want to change the drive security settings.

   The **Authenticate Drive Security Settings**dialog appears. Authentication is required for the changes that you requested to the drive security settings.
8. Enter the current security key to authenticate the changes.

   The MegaRAID Storage Manager software updates the existing configuration on the controller to use the new security settings and returns you to the main menu.

12.14.4 Disabling Drive Security

ATTENTION 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. If there are any secure drive groups on the controller, you cannot disable drive security. A warning dialog appears if you attempt to do so. To disable drive security, you must first delete the virtual drives on all of the secure drive groups.

Perform the following steps to disable drive security:

1. Select the **Physical View** tab in the left panel of the **MegaRAID Storage Manager** window, and select a controller icon.
2. Select **Go To** > **Controller** > **Disable Drive Security**.

   The **Confirm Disable Drive Security** dialog appears.

3. To disable drive security, click **Yes**.

   The MegaRAID Storage Manager software disables drive security and returns you to the main menu.

   ATTENTION If you disable drive security, you cannot create any new encrypted virtual drives and the data on all encrypted unconfigured drives will be erased. Disabling drive security does not affect the security or data of foreign drives.

12.14.5 Importing or Clearing a Foreign Configuration

A foreign configuration is a RAID configuration that already exists on a replacement set of drives that you install in a computer system. You can use the MegaRAID Storage Manager software to import the foreign configuration to the RAID controller or to clear the foreign configuration so you can create a new configuration using these drives.

To import a foreign configuration, you must perform the following tasks:

- Enable security to allow importation of locked foreign configurations. (You can import unsecured or unlocked configurations when security is disabled.)
- Run a scan for foreign configurations.
- If a locked foreign configuration is present and security is enabled, enter the security key, and unlock the configuration.
- Import the foreign configuration.

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

Verify whether any drives are left to import because the locked drives can use different security keys. If there are any drives left, repeat the import process for the remaining drives. After all the drives are imported, there is no configuration to import.

NOTE When you create a new configuration, the MegaRAID Storage Manager software shows only the unconfigured drives. Drives that have existing configurations, including foreign configurations, do not appear. To use drives with existing configurations, you must first clear the configuration on those drives.

Perform the following steps to import or clear a configuration:

1. Enable drive security to allow importation of locked foreign drives.
2. After you create a security key, right-click the controller icon, and select **Scan for Foreign Configuration**.

   If locked drives (security is enabled) exist, the **Unlock Foreign Drives** dialog appears.
3. Enter the security key to unlock the configuration.

   The **Foreign Configuration Detected** dialog appears, as shown in the following figure.

Figure 321 Foreign Configuration Detected Dialog

4. Choose one of the following options:
   - Click **Import** to import the foreign configuration from all of the foreign drives.
   - Click **Clear** to remove the configuration from all foreign drives.
   - Click **Advanced** to preview and import specific foreign configurations.
5. Click **OK**.

NOTE The operation cannot be reversed after it is started. Imported drives display as *Online* in the **MegaRAID Storage Manager** window.

6. Repeat the import process for any remaining drives.

   Because locked drives can use different security key, you must verify whether there are any remaining drives to be imported.

NOTE When you create a new configuration, the MegaRAID Storage Manager software shows only the unconfigured drives. Drives that have existing configurations, including foreign configurations, do not appear. To use drives with existing configurations, you must first clear the configuration on those drives.

12.14.5.1 Foreign Configurations in Cable Pull and Drive Removal Scenarios

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

The following scenarios can occur with cable pulls or drive removals. Use the **Foreign Configuration Preview** dialog to import or clear the foreign configuration in each case.

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

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

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

NOTE Start a consistency check immediately after the rebuild is complete to ensure data integrity for the virtual drives. See 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 Running a Consistency Check, for more information about checking data consistency.

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

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

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

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

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 main menu window. The messages are also logged in the Windows Application log (Event Viewer).

A.1 Error Levels

Each message that appears in the event log has a Severity level that indicates the severity of the event, as shown in the following table.

Table 61 Event Error Levels

Severity Level	Meaning
Information	Informational message. No user action is necessary.
Warning	Some component might be close to a failure point.
Critical	A component has failed, but the system has not lost data.
Fatal	A component has failed, and data loss has occurred or will occur.

A.2 Event Messages

The following table 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 62 Event Messages

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0000	Information	MegaRAID firmware initialization started (PCI ID %04x/%04x/%04x/%04x)	Logged at firmware initialization.
0x0001	Information	MegaRAID firmware version %s	Logged at firmware initialization to display firmware version.
0x0002	Fatal	Unable to recover cache data from TBBU	Currently not logged.
0x0003	Information	Cache data recovered from TBBU successfully	Currently not logged.
0x0004	Information	Configuration cleared	Logged when controller configuration is cleared.
0x0005	Warning	Cluster down; communication with peer lost	Currently not logged.
0x0006	Information	Virtual drive %s ownership changed from %02x to %02x	Currently not logged.
0x0007	Information	Alarm disabled by user	Logged when user disables alarm.

Table 62 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0008	Information	Alarm enabled by user	Logged when user enables alarm.
0x0009	Information	Background initialization rate changed to %d%%	Logged to display background initialization progress indication in percentage.
0x000a	Fatal	Controller cache discarded due to memory/battery problems	Logged on cache discard due to hardware problems.
0x000b	Fatal	Unable to recover cache data due to configuration mismatch	Currently not logged.
0x000c	Information	Cache data recovered successfully	Logged when cache data is successfully recovered after reboot.
0x000d	Fatal	Controller cache discarded due to firmware version incompatibility	Logged when cache data discarded because of firmware version mismatch.
0x000e	Information	Consistency Check rate changed to %d%%	Logged to display Consistency check progress indication percentage.
0x000f	Fatal	Fatal firmware error: %s	Logged in case of fatal errors and also while entering debug monitor.
0x0010	Information	Factory defaults restored	Logged while controller is reset to factory defaults.
0x0011	Information	Flash downloaded image corrupt	Logged to inform downloaded flash image is corrupt.
0x0012	Critical	Flash erase error	Logged in case of flash erase failure, generally after flash update.
0x0013	Critical	Flash timeout during erase	Logged to indicate flash erase operation timed out.
0x0014	Critical	Flash error	Generic unknown internal error during flash update flash.
0x0015	Information	Flashing image: %s	Logged to display flash image name string before getting updated to controller.
0x0016	Information	Flash of new firmware images complete	Logged to inform successful update of flash image(s).
0x0017	Critical	Flash programming error	Logged to notify, write failure during flash update, not being allowed usually due to internal controller settings.
0x0018	Critical	Flash timeout during programming	Logged to indicate flash write operation timed out.
0x0019	Critical	Flash chip type unknown	Logged during flash update tried with unsupported flash chip type.
0x001a	Critical	Flash command set unknown	Logged while unsupported flash command set detected, most likely because of unsupported flash chip.
0x001b	Critical	Flash verify failure	Logged when compare operation fails between written flash data and original data.
0x001c	Information	Flush rate changed to %d seconds	Logged to notify modified cache flush frequency in seconds.
0x001d	Information	Hibernate command received from host	Logged to inform about reception of hibernation command from host to controller, generally during host shutdown.
0x001e	Information	Event log cleared	Logged when controller log has been cleared.
0x001f	Information	Event log wrapped	Logged when controller log has been wrapped around, when the maximum logs are written.
0x0020	Fatal	Multi-bit ECC error: ECAR=%x, ELOG=%x, (%s)	Logged to notify ECC multi bit error in memory, ELOG: ecc info (source, type, syndrome), ECAR:ecc address.
0x0021	Warning	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s)	Logged to notify ECC single bit error in memory, ELOG: ecc info (source, type, syndrome), ECAR:ecc address.
0x0022	Fatal	Not enough controller memory	Logged to notify fatal controller condition, when you run out of memory to allocate.

Table 62 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0023	Information	Patrol Read complete	Logged when patrol read completes.
0x0024	Information	Patrol Read paused	Logged when patrol read is paused.
0x0025	Information	Patrol Read Rate changed to %d%%	Logged to indicate progress of patrol read in percentage.
0x0026	Information	Patrol Read resumed	Logged when patrol read is resumed.
0x0027	Information	Patrol Read started	Logged when patrol read is started.
0x0028	Information	Reconstruction rate changed to %d%%"	Logged to indicate progress of reconstruction in percentage.
0x0029	Information	Drive group modification rate changed to %d%%	Logged to indicate the change in Drive group modification frequency.
0x002a	Information	Shutdown command received from host	Logged when shutdown command is received from host to controller.
0x002b	Information	Test event: %s	General controller event, with a generic string.
0x002c	Information	Time established as %s; (%d seconds since power on)	Logged when controller time was set from host, also displaying time since power on in seconds.
0x002d	Information	User entered firmware debugger	Logged when user enters controller debug shell.
0x002e	Warning	Background Initialization aborted on %s	Logged to inform about user aborted background initialization on displayed LD number.
0x002f	Warning	Background Initialization corrected medium error (%s at %lx	logged to inform about corrected medium error on displayed LD number, LBALBA number, PD number and PDLBA number in that order.
0x0030	Information	Background Initialization completed on %s	Logged to inform Background Initialization completion on displayed LD.
0x0031	Fatal	Background Initialization completed with uncorrectable errors on %s	Logged to inform Background Initialization completion with error on displayed LD.
0x0032	Fatal	Background Initialization detected uncorrectable double medium errors (%s at %lx on %s)	Logged to inform Background Initialization completion with double medium error on displayed PD, PDLBA and LD in that order.
0x0033	Critical	Background Initialization failed on %s	Logged to inform Background Initialization failure on displayed LD.
0x0034	Progress	Background Initialization progress on %s is %s	Logged to inform Background Initialization progress in percentage of displayed LD.
0x0035	Information	Background Initialization started on %s	Logged to inform Background Initialization started for displayed LD.
0x0036	Information	Policy change on %s from %s to %s	Logged to inform the changed policy for displayed LD with old and new policies.
0x0038	Warning	Consistency Check aborted on %s	Logged to inform aborted Consistency check for displayed LD.
0x0039	Warning	Consistency Check corrected medium error (%s at %lx	Logged when Consistency check corrected medium error.
0x003a	Information	Consistency Check done on %s	Logged when Consistency check has completed successfully on the LD.
0x003b	Information	Consistency Check done with corrections on %s	Logged when Consistency check completed and inconsistency was found during check and was corrected.
0x003c	Fatal	Consistency Check detected uncorrectable double medium errors (%s at %lx on %s)	Logged when uncorrectable double medium error are detected while consistency check.

Table 62 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x003d	Critical	Consistency Check failed on %s	Logged when Consistency check failed as fatal error was found.
0x003e	Fatal	Consistency Check completed with uncorrectable data on %s	Logged when Uncorrectable error occurred during consistency check.
0x003f	Warning	Consistency Check found inconsistent parity on %s at strip %lx	Logged when consistency check finds inconsistency parity on a strip.
0x0040	Warning	Consistency Check inconsistency logging disabled on %s (too many inconsistencies)	Logged when consistency check finds too many inconsistent parity (greater than 10) and the inconsistency parity logging is disabled.
0x0041	Progress	Consistency Check progress on %s is %s	Logs Consistency Check progress, the progress is logged only if the progress is greater than 1% at an interval of every 15 seconds.
0x0042	Information	Consistency Check started on %s	Logged when consistency check has started
0x0043	Warning	Initialization aborted on %s	Logged when consistency check is aborted by you or for some other reason.
0x0044	Critical	Initialization failed on %s	Logged when initialization has failed.
0x0045	Progress	Initialization progress on %s is %s	Logs initialization progress, the progress is logged only if the progress is greater than 1% at an interval of every 15 seconds.
0x0046	Information	Fast initialization started on %s	Logged when quick initialization has started on a LD. The parameter to decide Quick init or Full init is passed by you.
0x0047	Information	Full initialization started on %s	Logged when full initialization has started.
0x0048	Information	Initialization complete on %s	Logged when initialization has completed successfully.
0x0049	Information	LD Properties updated to %s (from %s)	Logged when LD properties has been changed.
0x004a	Information	Reconstruction complete on %s	Logged when reconstruction has completed successfully.
0x004b	Fatal	Reconstruction of %s stopped due to unrecoverable errors	Logged when reconstruction has finished because of failure (unrecoverable errors).
0x004c	Fatal	Reconstruct detected uncorrectable double medium errors (%s at %lx on %s at %lx)	Logged while reconstructing if an unrecoverable double medium error is encountered.
0x004d	Progress	Reconstruction progress on %s is %s	Logs reconstruction progress, the progress is logged only if the progress is greater than 1% at an interval of every 15 seconds.
0x004e	Information	Reconstruction resumed on %s	Logged when reconstruction resumes after a power cycle.
0x004f	Fatal	Reconstruction resume of %s failed due to configuration mismatch	Logged when reconstruction resume failed due to configuration mismatch.
0x0050	Information	Reconstruction started on %s	Logged on start of reconstruction on a LD.
0x0051	Information	State change on %s from %s to %s	Logged when there is change in LD state. The event gives the new and old state. The state could be one of the following, LDS_OFFLINE, LDS_PARTIALLY_DEGRADED, LDS_DEGRADED, LDS_OPTIMAL.
0x0052	Information	Drive Clear aborted on %s	Logged when PD clear is aborted.
0x0053	Critical	Drive Clear failed on %s (Error %02x)	Logged when drive clear is failed and the even is logged along with error code.
0x0054	Progress	Drive Clear progress on %s is %s	Logs drive clear progress, the progress is logged only if the progress is greater than 1% at an interval of every 15 seconds.
0x0055	Information	Drive Clear started on %s	Logged when drive clear started on a PD.

Table 62 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0056	Information	Drive Clear completed on %s	Logged when PD clear task is completed successfully on a PD.
0x0057	Warning	Error on %s (Error %02x)	Logged if Read returns with Uncorrectable error or same errors on both the drives or write long returns with an error (ie. puncture operation could failed).
0x0058	Information	Format complete on %s	Logged when Format has completed.
0x0059	Information	Format started on %s	Logged when format unit is started on a PD.
0x005a	Critical	Hot Spare SMART polling failed on %s (Error %02x)	Currently not logged.
0x005b	Information	Drive inserted: %s	Logged when drive is inserted and slot/enclosure fields of PD are updated.
0x005c	Warning	Drive %s is not supported	Logged when the drive is not supported; reason could be the number of drive has exceeded the MAX supported drives or an unsupported drive is inserted like a SATA drive in SAS only enclosure or could be a unsupported drive type.
0x005d	Warning	Patrol Read corrected medium error on %s at %lx	Logged when Patrol read has successfully completed recovery read and recovered data.
0x005e	Progress	Patrol Read progress on %s is %s	Logs patrol read progress, the progress is logged only if the progress is greater than 1% at an interval of every 15 seconds.
0x005f	Fatal	Patrol Read found an uncorrectable medium error on %s at %lx	Logged when Patrol read is unable to recover data.
0x0060	Critical	Predictive failure: CDB: %s	Logged when a failure is found during smart (predictive failure) poll.
0x0061	Fatal	Patrol Read puncturing bad block on %s at %lx	Logged when patrol read punctures a block due to unrecoverable medium error.
0x0062	Information	Rebuild aborted by user on %s	Logged when the user aborts a rebuild operation.
0x0063	Information	Rebuild complete on %s	Logged when the rebuild operation on a logical drive on a physical drive (which may have multiple LDs) is completed.
0x0064	Information	Rebuild complete on %s	Logged when rebuild operation is completed for all logical drives on a given physical drive.
0x0065	Critical	Rebuild failed on %s due to source drive error	Logged if one of the source drives for the rebuild operation fails or is removed.
0x0066	Critical	Rebuild failed on %s due to target drive error	Logged if the target rebuild drive (on which rebuild operation is going on) fails or is removed from the controller.
0x0067	Progress	Rebuild progress on %s is %s	Logged to indicate the progress (in percentage) of the rebuild operation on a given physical drive.
0x0068	Information	Rebuild resumed on %s	Logged when the rebuild operation on a physical drive resumes.
0x0069	Information	Rebuild started on %s	Logged when the rebuild operation is started on a physical drive.
0x006a	Information	Rebuild automatically started on %s	Logged when the rebuild operation kicks in on a spare.
0x006b	Critical	Rebuild stopped on %s due to loss of cluster ownership	Logged when the rebuild operation is stopped due to loss of ownership.
0x006c	Fatal	Reassign write operation failed on %s at %lx	Logged when a check condition or medium error is encountered for a reassigned write.
0x006d	Fatal	Unrecoverable medium error during rebuild on %s at %lx	Logged when the rebuild I/O encounters an unrecoverable medium error.
0x006e	Information	Corrected medium error during recovery on %s at %lx	Logged when recovery completed successfully and fixed a medium error.

Table 62 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x006f	Fatal	Unrecoverable medium error during recovery on %s at %lx	Logged when the recovery for a failed I/O encounters a medium error.
0x0070	Information	Drive removed: %s	Logged when a drive is removed from the controller.
0x0071	Warning	Unexpected sense: %s, CDB%s, Sense: %s	Logged when an I/O fails due to unexpected reasons and sense data needs to be logged.
0x0072	Information	State change on %s from %s to %s	Logged when the state of a drive is changed by the firmware or by you.
0x0073	Information	State change by user on %s from %s to %s	Not logged by the firmware.
0x0074	Warning	Redundant path to %s broken	Not logged by the firmware.
0x0075	Information	Redundant path to %s restored	Not logged by the firmware
0x0076	Information	Dedicated Hot Spare Drive %s no longer useful due to deleted drive group	Not logged by the firmware.
0x0077	Critical	SAS topology error: Loop detected	Logged when device discovery fails for a SAS device as a loop was detected.
0x0078	Critical	SAS topology error: Unaddressable device	Logged when device discovery fails for a SAS device as an unaddressable device was found.
0x0079	Critical	SAS topology error: Multiple ports to the same SAS address	Logged when device discovery fails for a SAS device multiple ports with same SAS address were detected.
0x007a	Critical	SAS topology error: Expander error	Not logged by the firmware.
0x007b	Critical	SAS topology error: SMP timeout	Logged when device discovery fails for a SAS device due to SMP timeout.
0x007c	Critical	SAS topology error: Out of route entries	Logged when device discovery fails for a SAS device as expander route table is out of entries.
0x007d	Critical	SAS topology error: Index not found	Logged when device discovery fails for a SAS device as expander route table out of entries.
0x007e	Critical	SAS topology error: SMP function failed	Logged when device discovery fails for a SAS device due to SMP function failure.
0x007f	Critical	SAS topology error: SMP CRC error	Logged when device discovery fails for a SAS device due to SMP CRC error.
0x0080	Critical	SAS topology error: Multiple subtractive	Logged when device discovery fails for a SAS device as a subtractive-to-subtractive link was detected.
0x0081	Critical	SAS topology error: Table to table	Logged when device discovery fails for a SAS device as table-to-table link was detected.
0x0082	Critical	SAS topology error: Multiple paths	Not logged by the firmware.
0x0083	Fatal	Unable to access device %s	Logged when the inserted drive is bad and unusable.
0x0084	Information	Dedicated Hot Spare created on %s (%s)	Logged when a drive is configured as a dedicated spare.
0x0085	Information	Dedicated Hot Spare %s disabled	Logged when a drive is removes as a dedicated spare.
0x0086	Critical	Dedicated Hot Spare %s no longer useful for all drive groups	Logged when an array with a dedicated spare is resized. The hot spare (dedicated to this array and possibly others) will not be applicable to other arrays.
0x0087	Information	Global Hot Spare created on %s (%s)	Logged when a drive is configured as a global hot spare.
0x0088	Information	Global Hot Spare %s disabled	Logged when a drive configured as global host spare fails or is unconfigured by you.

Table 62 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0089	Critical	Global Hot Spare does not cover all drive groups	Logged when the global hotspare is too small (or doesn't meet the SAS/SATA restricitons) to cover certain arrays.
0x008a	Information	Created %s}	Logged as soon as the new logical drive created is added to the firmware configuration.
0x008b	Information	Deleted %s}	Logged when the firmware removes an LD from it's configuration upon a user request from the applications.
0x008c	Information	Marking LD %s inconsistent due to active writes at shutdown	Logged when we have active writes on one of the target disks of a Raid 5 LD at the time of shutdown.
0x008d	Information	Battery Present	Logged during firmware initialization when we check if there is a battery present and the check turns out true. This event is also logged when a battery is inserted or replaced with a new one and the battery present check returns true.
0x008e	Warning	Battery Not Present	Logged if the user has not disabled "Battery Not Present" warning at the boot time or if a battery has been removed.
0x008f	Information	New Battery Detected	Logged when we have a subsequent boot after a new battery has been inserted.
0x0090	Information	Battery has been replaced	Logged when a new battery has been replaced with an old battery.
0x0091	Critical	Battery temperature is high	Logged when we detect that the battery temperature is high during the periodic battery status check.
0x0092	Warning	Battery voltage low	Not logged by the firmware.
0x0093	Information	Battery started charging	Logged as part of monitoring the battery status when the battery is getting charged.
0x0094	Information	Battery is discharging	Logged as part of monitoring the battery status when the battery is getting discharged.
0x0095	Information	Battery temperature is normal	Logged as part of monitoring the battery status when the temperature of the battery is normal.
0x0096	Fatal	Battery has failed and cannot support data retention. Please replace the battery.	Logged when there is not enough capacity left in battery for expected data retention time. Battery has to be replaced.
0x0097	Information	Battery relearn started	logged when the battery relearn started, initiated either by the user or automatically.
0x0098	Information	Battery relearn in progress	Logged as part of monitoring the battery status when the battery relearn is in progress.
0x0099	Information	Battery relearn completed	Logged as part of monitoring the battery status when the battery relearn is complete.
0x009a	Critical	Battery relearn timed out	Not logged by the firmware.
0x009b	Information	Battery relearn pending: Battery is under charge	Logged as part of monitoring the battery status when the battery relearn is requested but yet to start.
0x009c	Information	Battery relearn postponed	Logged as part of monitoring the battery status when the battery relearn is requested but postponed as there is valid pinned cache present. This event can also be logged when learn delay interval has been explicitly set.
0x009d	Information	Battery relearn will start in 4 days	Logged as part of providing battery learn cycle information when auto learn is enabled.
0x009e	Information	Battery relearn will start in 2 day	Logged as part of providing battery learn cycle information when auto learn is enabled.
0x009f	Information	Battery relearn will start in 1 day	Logged as part of providing battery learn cycle information when auto learn is enabled.

Table 62 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x00a0	Information	Battery relearn will start in 5 hours	Logged as part of providing battery learn cycle information when auto learn is enabled.
0x00a1	Information	Battery removed	Logged as part of periodic monitoring of the battery status when a battery has been removed.
0x00a2	Information	Current capacity of the battery is below threshold	Logged as part of monitoring the battery status when the capacity of the battery is below threshold.
0x00a3	Information	Current capacity of the battery is above threshold	Logged as part of monitoring the battery status when the capacity of the battery is above threshold.
0x00a4	Information	Enclosure (SES) discovered on %s	Logged when an Enclosure (SES) is discovered for the first time.
0x00a5	Information	Enclosure (SAFTE) discovered on %s	Not logged by the firmware.
0x00a6	Critical	Enclosure %s communication lost	Logged when the communication with an enclosure has been lost.
0x00a7	Information	Enclosure %s communication restored	Logged when the communication with an enclosure has been restored
0x00a8	Critical	Enclosure %s fan %d failed	Logged when an enclosure fan has failed.
0x00a9	Information	Enclosure %s fan %d inserted	Logged when an enclosure fan has been inserted newly.
0x00aa	Critical	Enclosure %s fan %d removed	Logged when an enclosure fan has been removed.
0x00ab	Critical	Enclosure %s power supply %d failed	Not logged by the firmware.
0x00ac	Information	Enclosure %s power supply %d inserted	Logged when power supply has been inserted to an enclosure.
0x00ad	Critical	Enclosure %s power supply %d removed	Logged when power supply has been removed from an enclosure.
0x00ae	Critical	Enclosure %s SIM %d failed	Logged when the enclosure SIM has failed.
0x00af	Information	Enclosure %s SIM %d inserted	Logged when an enclosure SIM has been inserted.
0x00b0	Critical	Enclosure %s SIM %d removed	Logged when an enclosure initialization was completed but later the SIM was removed.
0x00b1	Warning	Enclosure %s temperature sensor %d below warning threshold	Logged when the enclosure services process has detected a temperature lower than a normal operating temperature or lower than the value indicated by the LOW WARNING THRESHOLD field in the Threshold In diagnostic page.
0x00b2	Critical	Enclosure %s temperature sensor %d below error threshold	Logged when the enclosure services process has detected a temperature lower than a safe operating temperature or lower than the value indicated by the LOW CRITICAL THRESHOLD field in the Threshold In diagnostic page.
0x00b3	Warning	Enclosure %s temperature sensor %d above warning threshold	Logged when the enclosure services process has detected a temperature higher than a normal operating temperature or higher than the value indicated by the HIGH WARNING THRESHOLD field in the Threshold In diagnostic page.
0x00b4	Critical	Enclosure %s temperature sensor %d above error threshold	Logged when the enclosure services process has detected a temperature higher than a safe operating temperature or higher than the value indicated by the HIGH CRITICAL THRESHOLD field in the Threshold In diagnostic page.
0x00b5	Critical	Enclosure %s shutdown	Logged when an unrecoverable condition is detected in the enclosure.
0x00b6	Warning	Enclosure %s not supported; too many enclosures connected to port	Logged when the maximum allowed enclosures per port is exceeded.
0x00b7	Critical	Enclosure %s firmware mismatch	Logged when two ESMs have different firmware versions.

Table 62 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x00b8	Warning	Enclosure %s sensor %d bad	Logged when the device is present on the phy, but the status does not indicate its presence.
0x00b9	Critical	Enclosure %s phy %d bad	Logged when the status indicates a device presence, but there is no corresponding SAS address is associated with the device.
0x00ba	Critical	Enclosure %s is unstable	Logged when the enclosure services process reports the sense errors.
0x00bb	Critical	Enclosure %s hardware error	Logged when a critical or an unrecoverable enclosure failure has been detected by the enclosure services process.
0x00bc	Critical	Enclosure %s not responding	Logged when there is no response from the enclosure.
0x00bd	Information	SAS/SATA mixing not supported in enclosure; Drive %s disabled	Logged when the SAS/SATA mixing in an enclosure is being violated.
0x00be	Information	Enclosure (SES) hotplug on %s was detected, but is not supported	Not reported to the user.
0x00bf	Information	Clustering enabled	Logged when the clustering is enabled in the controller properties.
0x00c0	Information	Clustering disabled	Logged when the clustering is disabled in the controller properties.
0x00c1	Information	Drive too small to be used for auto-rebuild on %s	Logged when the size of the drive is not sufficient for auto-rebuild.
0x00c2	Information	BBU enabled; changing WT virtual drives to WB	Logged when changing WT virtual drives to WB and the BBU status is good.
0x00c3	Warning	BBU disabled; changing WB virtual drives to WT	Logged when changing WB virtual drives to WT and the BBU status is bad.
0x00c4	Warning	Bad block table on drive %s is 80% full	Logged when the Bad block table on a drive is 80% full.
0x00c5	Fatal	Bad block table on drive %s is full; unable to log block %lx	Logged when the Bad block table on a drive is full and not able to add the bad block in the Bad block table.
0x00c6	Information	Consistency Check Aborted due to ownership loss on %s	Logged when the Consistency Check is aborted due to ownership is lost.
0x00c7	Information	Background Initialization (BGI) Aborted Due to Ownership Loss on %s	Logged when the Background Initialization (BGI) is aborted due to ownership loss.
0x00c8	Critical	Battery/charger problems detected; SOH Bad	Logged when the battery is not presented or removed and SOH is bad.
0x00c9	Warning	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); warning threshold exceeded	Logged when the Single-bit ECC errors exceeded the warning threshold.
0x00ca	Critical	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); critical threshold exceeded	Logged when the Single-bit ECC errors exceeded the critical threshold.
0x00cb	Critical	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); further reporting disabled	Logged when the Single-bit ECC errors exceeded all the thresholds and disable further logging.
0x00cc	Critical	Enclosure %s Power supply %d switched off	Logged when the enclosure services process has detected that the Enclosure Power supply is switched off and it was switched on earlier.

Table 62 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x00cd	Information	Enclosure %s Power supply %d switched on	Logged when the enclosure services process has detected that the Enclosure Power supply is switched on and it was switched off earlier.
0x00ce	Critical	Enclosure %s Power supply %d cable removed	Logged when the enclosure services process has detected that the Enclosure Power supply cable is removed and it was inserted earlier.
0x00cf	Information	Enclosure %s Power supply %d cable inserted	Logged when the enclosure services process has detected that the Enclosure Power supply cable is inserted and it was removed earlier.
0x00d0	Information	Enclosure %s Fan %d returned to normal	Logged when the enclosure services process has detected that the current status of a fan is good and it was failed earlier.
0x00d1	Information	BBU Retention test was initiated on previous boot	Logged when the Battery Retention test was initiated on previous boot.
0x00d2	Information	BBU Retention test passed	Logged when the Battery Retention test passed successfully.
0x00d3	Critical	BBU Retention test failed!	Logged when the Battery Retention test failed.
0x00d4	Information	NVRAM Retention test was initiated on previous boot	Logged when the NVRAM Retention test was initiated on previous boot.
0x00d5	Information	NVRAM Retention test passed	Logged when the NVRAM Retention test passed successfully.
0x00d6	Critical	NVRAM Retention test failed!	Logged when the NVRAM Retention test failed.
0x00d7	Information	%s test completed %d passes successfully	Logged when the controller diagnostics test passes successfully.
0x00d8	Critical	%s test FAILED on %d pass. Fail data: errorOffset=%x goodData=%x badData=%x	Logged when the controller diagnostics test fails.
0x00d9	Information	Self check diagnostics completed	Logged when Self check diagnostics is completed.
0x00da	Information	Foreign Configuration detected	Logged when Foreign Configuration is detected.
0x00db	Information	Foreign Configuration imported	Logged when Foreign Configuration is imported.
0x00dc	Information	Foreign Configuration cleared	Logged when Foreign Configuration is cleared.
0x00dd	Warning	NVRAM is corrupt; reinitializing	Logged when NVRAM is corrupt and re-initialized.
0x00de	Warning	NVRAM mismatch occurred	Logged when NVRAM mismatch occurs.
0x00df	Warning	SAS wide port %d lost link on PHY %d	Logged when SAS wide port lost link on a PHY.
0x00e0	Information	SAS wide port %d restored link on PHY %d	Logged when a SAS wide port restored link on a PHY.
0x00e1	Warning	SAS port %d, PHY %d has exceeded the allowed error rate	Logged when a SAS PHY on port has exceeded the allowed error rate.
0x00e2	Warning	Bad block reassigned on %s at %lx to %lx	Logged when a Bad block is reassigned on a drive from a error sector to a new sector.
0x00e3	Information	Controller Hot Plug® detected	Logged when a Controller Hot Plug is detected.
0x00e4	Warning	Enclosure %s temperature sensor %d differential detected	Logged when an Enclosure temperature sensor differential is detected.
0x00e5	Information	Drive test cannot start. No qualifying drives found	Logged when Disk test cannot start. No qualifying disks found.
0x00e6	Information	Time duration provided by host is not sufficient for self check	Logged when Time duration provided by the host is not sufficient for self check.

Table 62 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x00e7	Information	Marked Missing for %s on drive group %d row %d	Logged when a physical drive is Marked Missing on an array at a particular row.
0x00e8	Information	Replaced Missing as %s on drive group %d row %d	Logged when a physical drive is Replaced Missing on an array at a particular row.
0x00e9	Information	Enclosure %s Temperature %d returned to normal	Logged when an Enclosure temperature returns to normal.
0x00ea	Information	Enclosure %s Firmware download in progress	Logged when Enclosure a Firmware download is in progress.
0x00eb	Warning	Enclosure %s Firmware download failed	Logged when Enclosure a Firmware download failed.
0x00ec	Warning	%s is not a certified drive	Logged if the drive is not certified.
0x00ed	Information	Dirty cache data discarded by user	Logged when Dirty cache data is discarded by the user.
0x00ee	Information	Drives missing from configuration at boot	Logged when physical drives are missing from configuration at boot.
0x00ef	Information	Virtual drives (VDs) missing drives and will go offline at boot: %s	Logged when virtual drives missing drives and will go offline at boot.
0x00f0	Information	VDs missing at boot: %s	Logged when virtual drives missing at boot.
0x00f1	Information	Previous configuration completely missing at boot	Logged when Previous configuration completely missing at boot.
0x00f2	Information	Battery charge complete	Logged when Battery charge is completed.
0x00f3	Information	Enclosure %s fan %d speed changed	Logged when an Enclosure fan speed changed.
0x00f4	Information	Dedicated spare %s imported as global due to missing arrays	Logged when a Dedicated spare is imported as global due to missing arrays.
0x00f5	Information	%s rebuild not possible as SAS/SATA is not supported in an array	Logged when a rebuild is 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.	Logged when SEP has been rebooted as part of enclosure firmware download. It will be unavailable until reboot completes.
0x00f7	Information	Inserted PD: %s Info: %s	Logged when a physical drive is inserted.
0x00f8	Information	Removed PD: %s Info: %s	Logged when a physical drive is removed.
0x00f9	Information	VD %s is now OPTIMAL	Logged when a logical drive state changes to OPTIMAL.
0x00fa	Warning	VD %s is now PARTIALLY DEGRADED	Logged when a logical drive state changes to a partially degraded state.
0x00fb	Critical	VD %s is now DEGRADED	Logged when a logical drive state changes to degraded state.
0x00fc	Fatal	VD %s is now OFFLINE	Logged when a logical drive state changes to offline state.
0x00fd	Warning	Battery requires reconditioning; please initiate a LEARN cycle	Logged when a Battery requires reconditioning; please initiate a LEARN cycle.
0x00fe	Warning	VD %s disabled because RAID-5 is not supported by this RAID key	Logged when a virtual drive is 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	Logged when a virtual drive is disabled because RAID-6 is not supported by this controller.
0x0100	Warning	VD %s disabled because SAS drives are not supported by this RAID key	Logged when a virtual drive is disabled because SAS drives are not supported by this RAID key.
0x0101	Warning	PD missing: %s	Logged to provide information about the missing drive during boot.

Table 62 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0102	Warning	Puncturing of LBAs enabled	Currently not logged in the firmware.
0x0103	Warning	Puncturing of LBAs disabled	Currently not logged in the firmware.
0x0104	Critical	Enclosure %s EMM %d not installed	Logged when Enclosure SIM is not installed.
0x0105	Information	Package version %s	Prints the Package version number.
0x0106	Warning	Global affinity Hot Spare %s commissioned in a different enclosure	Logged when a hot spare that is a part of an enclosure is commissioned in a different enclosure.
0x0107	Warning	Foreign configuration table overflow	Logged when the number of GUIDs to import exceeds the total supported by the firmware.
0x0108	Warning	Partial foreign configuration imported, PDs not imported:%s	Logged when all the foreign configuration drives could not be imported.
0x0109	Information	Connector %s is active	Logged during initial boot when a SAS MUX connector is found for the controller.
0x010a	Information	Board Revision %s	Logged during boot.
0x010b	Warning	Command timeout on PD %s, CDB:%s	Logged when command to a PD Timesout.
0x010c	Warning	PD %s reset (Type %02x)	Logged when PD is reset.
0x010d	Warning	VD bad block table on %s is 80% full	Logged when number of Bad Blocks entries is at 80 % of what can be supported in the firmware.
0x010e	Fatal	VD bad block table on %s is full; unable to log block %lx (on %s at %lx)	Logged when number of Bad Blocks exceed what can be supported in the firmware.
0x010f	Fatal	Uncorrectable medium error logged for %s at %lx (on %s at %lx)	Logged when an uncorrectable medium error is detected.
0x0110	Information	VD medium error corrected on %s at %lx	Logged on the corrected medium error.
0x0111	Warning	Bad block table on PD %s is 100% full	Logged when Bad block table is 100 % Full. Any more media errors on this physical drive will not be logged in the bad block table.
0x0112	Warning	VD bad block table on PD %s is 100% full	Logged when Bad block table is 100 % Full. Any more media errors on this logical drive will not be logged in the bad block table.
0x0113	Fatal	Controller needs replacement, IOP is faulty	Currently not logged in the firmware.
0x0114	Information	Replace Drive started on PD %s from PD %s	Logged when Replace is started.
0x0115	Information	Replace Drive aborted on PD %s and src is PD %s	Logged when Replace is aborted.
0x0116	Information	Replace Drive complete on PD %s from PD %s	Logged when Replace is completed.
0x0117	Progress	Replace Drive progress on PD %s is %s	Logged to provide the progress of Replace.
0x0118	Information	Replace Drive resumed on PD %s from %s	Logged when Replace operation is resumed.
0x0119	Information	Replace Drive automatically started on PD %s from %s	Logged on automatic start of Replace.

Table 62 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x011a	Critical	Replace Drive failed on PD %s due to source %s error	Logged when the source physical drive of a Replace fails. The Replace stops and rebuild starts on the destination physical drive.
0x011b	Warning	Early Power off warning was unsuccessful	Currently not logged in the firmware.
0x011c	Information	BBU FRU is %s	Logged only for IBM.
0x011d	Information	%s FRU is %s	Logged if FRU data is present. Logged only for IBM.
0x011e	Information	Controller hardware revision ID %s	Currently not used in the firmware.
0x011f	Warning	Foreign import shall result in a backward incompatible upgrade of configuration metadata	Currently not used in the firmware.
0x0120	Information	Redundant path restored for PD %s	Logged when new path is added for the physical drives.
0x0121	Warning	Redundant path broken for PD %s	Logged when one path is removed.
0x0122	Information	Redundant enclosure EMM %s inserted for EMM %s	Logged when an enclosure is added.
0x0123	Information	Redundant enclosure EMM %s removed for EMM %s	Logged when an enclosure is removed
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	Logged when none of the disks can start PR.
0x0125	Information	Replace Drive aborted by user on PD %s and src is PD %s	Logged when Replace is aborted by the user.
0x0126	Critical	Replace Drive aborted on hot spare %s from %s, as hot spare needed for rebuild	Logged when Replace is aborted on a Hotspare.
0x0127	Warning	Replace Drive aborted on PD %s from PD %s, as rebuild required in the array	Logged when Replace is stopped for a higher priority rebuild operation on a drive.
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.	Logged when pinned cache lines are discarded for a LD.
0x0129	Information	Replace Drive cannot be started as PD %s is too small for src PD %s	Logged when destination PD is too small for Replace.
0x012a	Information	Replace Drive cannot be started on PD %s from PD %s, as SAS/SATA is not supported in an array	Logged when there is a SAS/SATA mixing violation for the destination PD.
0x012b	Information	Microcode update started on PD %s	Logged when PD Firmware download starts.
0x012c	Information	Microcode update completed on PD %s	Logged when PD Firmware download completes.
0x012d	Warning	Microcode update timeout on PD %s	Logged when PD Firmware download does not complete and times out.
0x012e	Warning	Microcode update failed on PD %s	Logged when PD Firmware download fails.
0x012f	Information	Controller properties changed	Logged when any of the controller properties has changed.

Table 62 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0130	Information	Patrol Read properties changed	Currently not logged in the firmware.
0x0131	Information	CC Schedule properties changed	Logged when consistency check scheduling property has changed.
0x0132	Information	Battery properties changed	Logged when any of the BBU properties has changed.
0x0133	Warning	Periodic Battery Relearn is pending. Please initiate manual learn cycle as Automatic learn is not enabled	Logged when BBU periodic relearn is pending.
0x0134	Information	Drive security key created	Logged when controller lock key is created.
0x0135	Information	Drive security key backed up	Logged when controller lock key is backed up.
0x0136	Information	Drive security key from escrow, verified	Logged when controller lock key is verified from escrow.
0x0137	Information	Drive security key changed	Logged when controller lock key is re-keyed.
0x0138	Warning	Drive security key, re-key operation failed	Logged when controller lock re-key operation failed.
0x0139	Warning	Drive security key is invalid	Logged when the controller lock is not valid.
0x013a	Information	Drive security key destroyed	Logged when the controller lock key is destroyed.
0x013b	Warning	Drive security key from escrow is invalid	Logged when the controller escrow key is not valid. This escrow key can not unlock any drive.
0x013c	Information	VD %s is now secured	Logged when secure LD is created.
0x013d	Warning	VD %s is partially secured	Logged when all the drives in the array are not secure.
0x013e	Information	PD %s security activated	Logged when PD security key is set.
0x013f	Information	PD %s security disabled	Logged when security key is removed from an FDE drive.
0x0140	Information	PD %s is reprovisioned	Logged when PD security is cleared.
0x0141	Information	PD %s security key changed	Logged when PD lock key is re-keyed.
0x0142	Fatal	Security subsystem problems detected for PD %s	Logged when PD security can not be set.
0x0143	Fatal	Controller cache pinned for missing or offline VD %s	Logged when LD cache is pinned.
0x0144	Fatal	Controller cache pinned for missing or offline VDs: %s	Logged when pinned cache is found during OCR.
0x0145	Information	Controller cache discarded by user for VDs: %s	Logged when LD pinned cache is discarded by the user.
0x0146	Information	Controller cache destaged for VD %s	Logged when LD pinned cache is recovered.
0x0147	Warning	Consistency Check started on an inconsistent VD %s	Logged when consistency check is started on an inconsistent LD.
0x0148	Warning	Drive security key failure, cannot access secured configuration	Logged when an invalid lock key is detected.
0x0149	Warning	Drive security password from user is invalid	Not logged.
0x014a	Warning	Detected error with the remote battery connector cable	Not logged.
0x014b	Information	Power state change on PD %s from %s to %s	Logged when PD power state (spun up, spun down, in-transition) changes.
0x014c	Information	Enclosure %s element (SES code 0x%x) status changed	Not logged.

Table 62 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x014d	Information	PD %s rebuild not possible as HDD/CacheCade software mix is not supported in a drive group	Logged when mixing violation occurs due to HDD/SSD mismatch.
0x014e	Information	Replace Drive cannot be started on PD %s from %s, as HDD/CacheCade software mix is not supported in a drive group	Logged when Replace could not be started on a PD because HDD/CacheCade software mix was not supported in a drive group.
0x014f	Information	VD bad block table on %s is cleared	Logged when a VD bad block table was cleared.
0x0150	Caution	SAS topology error: 0x%lx	Logged when a SAS topology error occurred.
0x0151	Information	VD cluster of medium errors corrected for %s at %lx (on %s at %lx)	Logged when medium errors were corrected for a PD for a LD.
0x0152	Information	Controller requests a host bus rescan	Logged when controller requested a host bus rescan.
0x0153	Information	Controller repurposed and factory defaults restored	Logged when controller repurposed and factory defaults were restored.
0x0154	Information	Drive security key binding updated	Logged when drive security key binding was updated.
0x0159	Critical	Controller encountered a fatal error and was reset	Logged when a controller encountered a fatal error and was reset.
0x015a	Information	Snapshots enabled on %s (Repository %s)	Logged when snapshot was enabled on a LD.
0x015b	Information	Snapshots disabled on %s (Repository %s) by the user	Logged when snapshot was disabled on a LD by the user.
0x015c	Critical	Snapshots disabled on %s (Repository %s), due to a fatal error	Logged when snapshot was disabled on a LD due to a fatal error.
0x015d	Information	Snapshot created on %s at %s	Logged when snapshot was created on a LD.
0x015e	Information	Snapshot deleted on %s at %s	Logged when snapshot was deleted on a LD.
0x015f	Information	View created at %s to a snapshot at %s for %s	Logged when view was created at a LD.
0x0160	Information	View at %s is deleted, to snapshot at %s for %s	Logged when View at a LD was deleted
0x0161	Information	Snapshot rollback started on %s from snapshot at %s	Logged when snapshot rollback was started on a LD.
0x0162	Fatal	Snapshot rollback on %s internally aborted for snapshot at %s	Logged when snapshot rollback was internally aborted.
0x0163	Information	Snapshot rollback on %s completed for snapshot at %s	Logged when snapshot rollback on a LD was completed.
0x0164	Information	Snapshot rollback progress for snapshot at %s, on %s is %s	Logged to report snapshot rollback progress on a LD.
0x0165	Warning	Snapshot space for %s in snapshot repository %s, is 80%% full	Logged when snapshot space for a LD in a snapshot repository was 80% full.
0x0166	Critical	Snapshot space for %s in snapshot repository %s, is full	Logged when snapshot space for a LD in a snapshot repository was full.
0x0167	Warning	View at %s to snapshot at %s, is 80%% full on snapshot repository %s	Logged when view at a LD to a snapshot was 80% full on a snapshot repository.
0x0168	Critical	View at %s to snapshot at %s, is full on snapshot repository %s	Logged when view at a LD to a snapshot was full on a snapshot repository.
0x0169	Critical	Snapshot repository lost for %s	Logged when snapshot repository was lost for a LD.
0x016a	Warning	Snapshot repository restored for %s	Logged when snapshot repository was restored for a LD.

Table 62 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x016b	Critical	Snapshot encountered an unexpected internal error: 0x%lx	Logged when snapshot encountered an unexpected internal error.
0x016c	Information	Auto Snapshot enabled on %s (snapshot repository %s)	Logged when auto snapshot was enabled.
0x016d	Information	Auto Snapshot disabled on %s (snapshot repository %s)	Logged when auto Snapshot was disabled.
0x016e	Critical	Configuration command could not be committed to disk, please retry	Logged when configuration command could not be committed to disk and was asked to retry.
0x016f	Information	COD on %s updated as it was stale	Logged when COD in DDF is updated due to various reasons.
0x0170	Warning	Power state change failed on %s (from %s to %s)	Logged when power state change failed on a PD.
0x0171	Warning	%s is not available	Logged when a LD was not available.
0x0172	Information	%s is available	Logged when a LD was available.
0x0173	Information	%s is used for CacheCade with capacity 0x%lx logical blocks	Logged when a LD was used for CacheCade with the indicated capacity in logical blocks.
0x0174	Information	%s is using CacheCade %s	Logged when a LD was using CacheCade.
0x0175	Information	%s is no longer using CacheCade %s	Logged when a LD was no longer using CacheCade.
0x0176	Critical	Snapshot deleted due to resource constraints for %s in snapshot repository %s	Logged when the snapshot is deleted due to resource constraints in snapshot repository.
0x0177	Warning	Auto Snapshot failed for %s in snapshot repository %s	Logged when the Auto Snapshot is failed for a VD in snapshot repository.
0x0178	Warning	Controller reset on-board expander	Logged when the chip reset issued to on-board expander.
0x0179	Warning	CacheCade (%s) capacity changed and is now 0x%lx logical blocks	Logged when the CacheCade capacity is changed along with the current capacity.
0x017a	Warning	Battery cannot initiate transparent learn cycles	Logged when the Battery cannot initiate transparent learn cycles.
0x017b	Information	Premium feature %s key was applied for - %s	Logged when the Premium feature key was applied.
0x017c	Information	Snapshot schedule properties changed on %s	Logged when the Snapshot schedule properties changed.
0x017d	Information	Snapshot scheduled action is due on %s	Logged when the Snapshot scheduled action is due.
0x017e	Information	Performance Metrics: collection command 0x%lx	Logged during the Performance Metrics collection.
0x017f	Information	Premium feature %s key was transferred - %s	Logged when the Premium feature key was transferred.
0x0180	Information	Premium feature serial number %s	Logged when displaying the Premium feature serial number.
0x0181	Warning	Premium feature serial number mismatched. Key-vault serial num - %s	Logged when Premium feature serial number mismatched.
0x0182	Warning	Battery cannot support data retention for more than %d hours. Please replace the battery	Logged during the Battery monitoring and it displays the remaining data retention time of the battery.
0x0183	Information	%s power policy changed to %s (from %s)	Logged when the power policy of an LD is changed.

Table 62 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0184	Warning	%s cannot transition to max power savings	Logged when LD cannot transition to max power savings.
0x0185	Information	Host driver is loaded and operational	This event is not reported to the user.
0x0186	Information	%s mirror broken	Logged when the mirror is broken for an LD.
0x0187	Information	%s mirror joined	Logged when joining the LD with its broken mirror.
0x0188	Warning	%s link %d failure in wide port	This event is not reported to the user.
0x0189	Information	%s link %d restored in wide port	This event is not reported to the user.
0x018a	Information	Memory module FRU is %s	This event is not reported to the user.
0x018b	Warning	Cache-vault power pack is sub-optimal. Please replace the pack	This event is not reported to the user.
0x018c	Warning	Foreign configuration auto-import did not import any drives	Logged when the Foreign configuration auto-import did not import any drives.
0x018d	Warning	Cache-vault microcode update required	Logged when the BMU is not in Normal mode and Cache-vault microcode update required.
0x018e	Warning	CacheCade (%s) capacity exceeds maximum allowed size, extra capacity is not used	Logged when CacheCade capacity exceeds maximum allowed size, extra capacity is not used.
0x018f	Warning	LD (%s) protection information lost	Logged when the protection information is lost for an LD.
0x0190	Information	Diagnostics passed for %s	Logged when the SHIELD™ Diagnostics passed for a PD.
0x0191	Critical	Diagnostics failed for %s	Logged when the SHIELD Diagnostics failed for a PD.
0x0192	Information	Server Power capability Diagnostic Test Started	Logged when the Server Power capability Diagnostic Test starts.
0x0193	Information	Drive Cache settings enabled during rebuild for %s	Logged when the Drive Cache settings enabled during rebuild for a PD.
0x0194	Information	Drive Cache settings restored after rebuild for %s	Logged when the Drive Cache settings restored after rebuild for a PD.
0x0195	Information	Drive %s commissioned as Emergency spare	Logged when the Drive commissioned as Emergency spare.
0x0196	Warning	Reminder: Potential non-optimal configuration due to drive %s commissioned as emergency spare	Logged when the PD being imported is an Emergency Spare.
0x0197	Information	Consistency Check suspended on %s	Logged when the Consistency Check is suspended on an LD.
0x0198	Information	Consistency Check resumed on %s	Logged when the Consistency Check is resumed on an LD.
0x0199	Information	Background Initialization suspended on %s	Logged when the Background Initialization is suspended on an LD.
0x019a	Information	Background Initialization resumed on %	Logged when the Background Initialization is resumed on an LD.
0x019b	Information	Reconstruction suspended on %s	Logged when the Reconstruction is suspended on an LD.
0x019c	Information	Rebuild suspended on %	Logged when the Rebuild is suspended on a PD.
0x019d	Information	Replace Drive suspended on %s	Logged when the Replace is suspended on a PD.
0x019e	Information	Reminder: Consistency Check suspended on %	Logged as a reminder when the Consistency Check is suspended on an LD.
0x019f	Information	Reminder: Background Initialization suspended on %s	Logged as a reminder when the Background Initialization is suspended on an LD.
0x01a0	Information	Reminder: Reconstruction suspended on %s	Logged as a reminder when the Reconstruction is suspended on an LD.

Table 62 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x01a1	Information	Reminder: Rebuild suspended on %s	Logged as a reminder when the Rebuild is suspended on a PD.
0x01a2	Information	Reminder: Replace Drive suspended on %s	Logged as a reminder when Replace is suspended on a PD.
0x01a3	Information	Reminder: Patrol Read suspended	Logged as a reminder when the Patrol Read is suspended.
0x01a4	Information	Erase aborted on %s	Logged when the Erase is aborted on a PD.
0x01a5	Critical	Erase failed on %s (Error %02x)	Logged when the Erase is failed on a PD along with the error.
0x01a6	Progress	Erase progress on %s is %s	Logged to display the Erase progress on a PD along with its current progress.
0x01a7	Information	Erase started on %s	Logged when Erase is started on a PD.
0x01a8	Information	Erase completed on %s	Logged when the Erase is completed on a PD.
0x01a9	Information	Erase aborted on %s	Logged when the Erase is aborted on an LD.
0x01aa	Critical	Erase failed on %s	Logged when the Erase is failed on an LD.
0x01ab	Progress	Erase progress on %s is %s	Logged to display the Erase progress on an LD along with its current progress.
0x01ac	Information	Erase started on %s	Logged when the Erase is started on an LD.
0x01ad	Information	Erase complete on %s	Logged when the Erase is complete on an LD.
0x01ae	Warning	Potential leakage during erase on %s	Logged to inform the Potential leakage during erase on an LD.
0x01af	Warning	Battery charging was suspended due to high battery temperature	Logged when the Battery charging was suspended due to high battery temperature.
0x01b0	Information	NVCache firmware update was successful	This event is not reported to the user.
0x01b1	Warning	NVCache firmware update failed	This event is not reported to the user.
0x01b2	Fatal	%s access blocked as cached data in CacheCade is unavailable	This event is not reported to the user.
0x01b3	Information	CacheCade disassociate started on %s	This event is not reported to the user.
0x01b4	Information	CacheCade disassociate completed on %s	This event is not reported to the user.
0x01b5	Critical	CacheCade disassociate failed on %s	This event is not reported to the user.
0x01b6	Progress	CacheCade disassociate progress on %s is %s	This event is not reported to the user.
0x01b7	Information	CacheCade disassociate aborted by user on %s	This event is not reported to the user.
0x01b8	Information	Link speed changed on SAS port %d and PHY %d	Logged when the Link speed changed on SAS port and PHY.
0x01b9	Warning	Advanced Software Options was deactivated for - %s	This event is not reported to the user.
0x01ba	Information	%s is now accessible	This event is not reported to the user.
0x01bb	Information	%s is using CacheCade	This event is not reported to the user.
0x01bc	Information	%s is no longer using CacheCade	This event is not reported to the user.
0x01bd	Warning	Patrol Read aborted on %s	Logged when the Patrol Read is aborted on a PD.

Table 62 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x01c2	Information	Periodic Battery Relearn was missed, and rescheduled to %s	Logged if Battery Relearn was missed at the scheduled time due to a system power off then the controller will reschedule automatically when you power on the system.
0x01c3	Information	Controller reset requested by host	Logged when the Controller Reset process started on the corresponding controller.
0x01c4	Information	Controller reset requested by host, completed	Logged when the Controller Reset process completed on the corresponding controller.
0x01c7	Warning	Controller booted in headless mode with errors	Logged when the Controller is booted to safe mode due to warning errors.
0x01c8	Critical	Controller booted to safe mode due to critical errors	Logged when the Controller is booted to safe mode due to critical errors.
0x01c9	Warning	Warning Error during boot - %s	Logged when a warning error occurs during booting the controller to safe mode.
0x01ca	Critical	Critical Error during boot - %s	Logged when a critical error occurs during booting the controller to safe mode
0x01cb	Fatal	Fatal Error during boot - %s	Logged when a fatal error occurs during booting the controller to safe mode

Appendix B: 3ware CLI Commands to StorCLI Command Conversion

B.1 System Commands

Table 63 System Commands

Description	3ware® CLI Command	StorCLI Command
Show a general summary of all detected controllers.	`tw_cli show`	`show` `show ctrlcount`

B.2 Controller Commands

Table 64 Controller Commands

Description	3ware CLI Command	StorCLI Command
Show all information about the adapter, such as cluster state, BIOS, alarm, firmware, version, and so on.	`tw_cli /cx show all`	`/cx show all`
Download the firmware to all compatible controllers that can be flashed with the image. By default, CLI checks for signature and version.	`/cx update fw=filename_with_path [force]`	`/cx download src=filepath [nosigchk] [noverchk]`

Show the status of properties related to the controllers.	`/cx show <PropertyName>`
	
	The following properties can be used with this command:
	`a0,1,2|-aALL`
	`achip`
	`AENs [reverse]`
	`alarms [reverse]`
	`allunitstatus`
	`autocarve`
	`autorebuild`
	`bios`
		`/cx show <PropertyName>`
		
		The following properties can be used with this command:
		`abortcconerror`
		`activityforlocate`
		`alarm`
		`autorebuild`
		`backplane`
		`batterywarning`
		`bgirate`
		`bootwithpinnedcache`

	`carvesize`
	`ctlbus diag`
	`dpmstat [type=<inst|ra|ext>`
	`driver`
	`drivestatus`
	`events [reverse]`
	`exportjbod firmware`
	`memory`
	`model`
		`cachebypass`
		`cacheflushint`
		`ccrate`
		
		`coercion`
		`copyback`
		`directpdmapping`
		`ds`
		`eccbucketleakrate`

Table 64 Controller Commands (Continued)

Description	3ware CLI Command	StorCLI Command
	monitor	eccbucketsize
	numdrives	enableeghsp
	numports	enableesmarter
	numunits	enableeug
	ondegrade	exposeencldevice
	pcb	jbod
	pchip	loadbalancemode
	phy	maintainpdfailhistory
	rebuild	migraterate
	rebuildmodel	ncq
	rebuildrate	perfmode
	selftest	pr
	serial	prcorrectunconfiguredareas
	spinup	prrate
	stagger	rebuildrate
	unitstatus	rehostinfo
	verify	restorehotspare
	verifymode	safeid
	verifyrate	smartpollinterval
		spinupdelay
		spinupdrivecount
		time
		usefdeonlyencrypt
Set properties on the selected controllers.	autocarve=<on|off>	abortcconerror=<on|off>
	autodetect=<on|off >	activityforlocate=<on|off>
	disk=<p:-p>|all	alarm=<on|off>
	autorebuild=<on|off>	autorebuild=<on|off>
	carvesize=<1024..32768>	backplane=<*value*>
	dpmstat=<on|off>	batterywarning=<on|off>
	ondegrade=<cacheoff|follow>	bgirate=<*value*>
	rebuild=<enable|disable|><1..5>	bootwithpinnedcache=<on|off>
	rebuildmode=<adaptive|lowlatency>	cachebypass=<on|off>
	rebuildrate=<1..5>	flush|flushcache
	selftest=<enable|disable>	cacheflushinterval=<*value*>
	spinup=<*value*>	ccrate=<*value*>
	stagger=<*value*>	coercion=<*value*>
	verify=advanced|basic|<1..5>	clusterenable=<*value*>

Table 64 Controller Commands (Continued)

Description	3ware CLI Command	StorCLI Command
	verify=basic [pref=*ddd*:*hh*] where *hh*=(00...23 and *ddd*={mon|tue|wed|thu|fri|sat|sun}	copyback=<on|off> type=<smartssd|smarthdd|all>
	verify=enable|disable|<1..5>	directpdmapping=<on|off>
	verifymode=<adaptive|lowlatency>	eccbucketleakrate=<*value*>
	verifyrate=<1..5>	eccbucketsize=<*value*>
		enableeghsp=<on|off>
		enableesmarter=<*value*>
		enableeug=<on|off>
		exposeencldevice=<on|off>
		foreignautoimport=<on|off>
		jbod=<on|off>
		loadbalancemode=<*value*>
		maintainpdfailhistory=<on|off>
		migraterate=<*value*>
		ncq=<on|off>
		perfmode=<*value*>
		prcorrectunconfiguredareas=<on|off>
		prrate=<*value*>
		rebuildrate=<*value*>
		restorehotspare=<on|off>
		smartpollinterval=<*value*>
		spinupdelay=<*value*>
		spinupdrivecount=<*value*>
		stoponerror=<on|off>
		usefdeonlyencrypt=<on|off>
		time=yyyymmddhh:mm:ss|systemtime
		usefdeonlyencrypt=<on|off>

B.3 Alarm Commands

Table 65 Alarm Commands

Description	3Ware CLI Command	StorCLI Command
Set alarm properties.	/cx/ex/almx set alarm=<mute|unmute|off> **NOTE** The 3ware® controllers have enclosure alarms.	/cx set alarm=<on|off|silence> **NOTE** The StorCLI controllers have controller alarms.
Show alarm properties.	/cx/ex show alarms **NOTE** This command applies for only 9750 and 9690SA controllers.	/cx show alarm

B.4 Patrol Read and Consistency Check Commands

Table 66 Patrol Read and Consistency Check Commands

Description	3ware CLI Command	StorCLI Command
Show patrol read status and patrol read parameters, if any in progress.	/cx/ux show	/cx show patrolRead
Set the patrol read options on a single adapter, multiple adapters, or all adapters (*x* = single controller).	/cx/ux start verify /cx/ux set autoverify=<on|off> /cx add verify=dddh:hh:duration	/cx set patrolread {=on mode=<auto|manual>}|{off} /cx set patrolread [starttime=<*yyyy/mm/dd hh* [maxconcurrentpd=<*value*>] [includessds=<on|off>] [uncfgareas=on|off] /cx set patrolread delay=<*value*>
Show consistency check status, if any in progress, and consistency check parameters.	/cx/ux show	/cx/vx show cc /cx show ccrate
Set consistency check options on a single adapter, multiple adapters, or all adapters (*x* = single controller).	/cx/ux start verify /cx/ux set autoverify=<on|off> /cx add verify=ddd:hh:duration	storcli /cx set consistencycheck|cc=[off|seq|conc] [delay=*value*] [starttime=*yyyy/mm/dd hh*] [excludevd=*x-y*,*z*]

NOTE The 3ware® CLI combines both patrol read and consistency check into a single command. The StorCLI has different commands for each.

B.5 BBU Commands

Table 67 BBU Commands

Description	3ware CLI Command	StorCLI Command
Show complete BBU information, such as status, capacity information, design information, and properties.	/*cx*/bbu show all	/*cx*/bbu show all
Show BBU summary information.	/*cx*/bbu show	/*cx*/bbu show
Show BBU properties.	/*cx*/bbu show batinst /*cx*/bbu show bootloader /*cx*/bbu show fw /*cx*/bbu show lasttest /*cx*/bbu show pcb /*cx*/bbu show serial /*cx*/bbu show status /*cx*/bbu show temp /*cx*/bbu show tempstat /*cx*/bbu show tempval /*cx*/bbu show volt	/*cx*/bbu show properties /*cx*/bbu show status **NOTE** Not all the properties shown in the 3ware CLI are shown in the StorCLI.
Show BBU capacity information.	/*cx*/bbu show cap	/*cx*/bbu show all
Start the learning cycle on the BBU.	/*cx*/bbu test [quiet]	/*cx*/bbu start learn

B.6 Virtual Drive Commands

Table 68 Virtual Drive Commands

Description	3Ware CLI Command	StorCLI Command
Create a RAID volume of the specified RAID type.	/*cx* add vd type=<*RaidType*> disk=<*p*:*p*|*p*-*p*|*p*:*p*-*p*>>> (where *p*=port or drive number) [strip=<*size*>] [nocache|nowrcache] [nordcache|rdcachebasic] [name=string (9000 series)] [ignoreECC] [autoverify|noautoverify] v0=*n*|vol=*a*:*b*:*c*:*d*] (*n*, *a*, *b*, *c*, *d*=size of volume in GB) [noqpolicy] [storsave=<protect|balance|perform>] [noscan] [rapidrecovery=<all|rebuild|disable>] [group=<3|4|5|6|7|8|9|10|11|12|13|14|15|16>] *RaidType*={raid0, raid1, raid5, raid10, raid50, single, spare, raid6}	/*cx* add vd type=raid[0|1|5|6|10|50|60] [[size=<*vd1_size*>,<*vd2_size*>,..|*all][name=<*vdname1*>,..] drives=*e*:*s*|*e*:*s*-*x*|*e*:*s*-*x*,*y*|*e*:*s*-*x*,*y*,*z* [pdperarray=*x*|*auto] [sed] [pdcache=on|off|*default] [pi][dimmerswitch|ds=default|automatic(auto)|*none|maximum(max) |maximumwithoutcaching(maxnocache)] [wt|*wb|awb] [nora|*ra] [*direct|cached] [strip=<8|16|32|64|128|256|512|1024] [aftervd=*x*] [spares=[*e*:]*s*|[*e*:]*s*-*x*|[*e*:]*s*-*x*,*y*|[*e*:]*s*-*x*,*y*,*z*>] [force] **NOTE** The supported strip size can vary from a minimum of 64 KB to 1 MB for MegaRAID controllers and only 64 KB for Integrated MegaRAID controllers. The LSISAS2108 controller supports strip size from 8 KB to 1 MB.
Delete virtual drives.	/*cx*/*ux* del [quiet] **NOTE** You can delete a single unit using this command.	/*cx*/*vx* [all] delete [force] [cachecade] **NOTE** You can delete one virtual disk, multiple virtual disks, or all the selected virtual disks on selected adapters using this command.
Show drive group information.	/*cx*/*ux* show [all] **NOTE** Information of each unit is shown individually.	/*cx*/dall show [cachecade]
Scan and show available foreign configurations, provide a preview of the imported foreign configuration, show or import foreign configuration.	/*cx* rescan	*cx*/*fall* [all] show [preview] [securityKey=*sssssssssss*] *cx*/*fall* [all] import [securityKey=*sssssssssss*]
Show VD information, including name, RAID level, RAID level qualifier, size in MBs, state, strip size, number of drives, span depth, cache policy, access policy, and any ongoing activity progress, which includes initialization, background initialization, consistency check, and reconstruction.	/*cx*/*ux* show [all]	/*cx*/*vx* show all

Table 68 Virtual Drive Commands (Continued)

Description	3Ware CLI Command	StorCLI Command
Show the virtual drive properties.	/cx/ux show autoverify /cx/ux show identify /cx/ux show ignoreECC /cx/ux show initializestatus /cx/ux show name /cx/ux show parit /cx/ux show qpolicy /cx/ux show rapidrecovery /cx/ux show rdcache /cx/ux show rebuildstatus /cx/ux show serial /cx/ux show status /cx/ux show storsave /cx/ux show verifystatus /cx/ux show volumes /cx/ux show wrcache	/cx/vx show all **NOTE** The StorCLI does not have commands to show individual virtual drive properties.
Set virtual drive properties.	/cx/ux set autoverify=on|off /cx/ux set cache=on|off [quiet] /cx/ux set identify=on|off /cx/ux set ignoreECC=on|off /cx/ux set name=string /cx/ux set qpolicy=on|off /cx/ux set rapidrecovery=all|rebuild|disable /cx/ux set rdcache=basic|intelligent|off /cx/ux set storsave=protect|balance|perform [quiet] /cx/ux set wrcache=on|off [quiet]	/cx/vx set accesspolicy=<rw|ro|blocked|rmvblkd> /cx/vx set iopolicy=<cached|direct> /cx/vx set name=<*namestring*> /cx/vx set pdcache=<on|off|default> /cx/vx set rdcache=<ra|nora|adra> /cx/vx set security=<on|off> /cx/vx|vall set ssdcaching=<on|off> /cx/vx set wrcache=<wt|wb|awb>
Show cache and access policies of the virtual drive.	/cx/ux show [all] /cx/ux show autoverify /cx/ux show cache /cx/ux show identify /cx/ux show ignoreECC /cx/ux show name /cx/ux show parit /cx/ux show qpolicy /cx/ux show rapidrecovery /cx/ux show rdcache /cx/ux show rebuildstatus /cx/ux show serial /cx/ux show status intializestatus /cx/ux show storsave /cx/ux show verify status /cx/ux show volumes /cx/ux show wrcache	/cx/vx show all **NOTE** The StorCLI does not have commands to show individual virtual drive properties.

Table 68 Virtual Drive Commands (Continued)

Description	3Ware CLI Command	StorCLI Command
Start initialization (writing 0s) on the virtual drive.	/c*x*/u*x* start verify **NOTE** Only the bios can do a foreground initialization. A background initialization does otherwise. A verify starts a back ground initialization.	/c*x*/v*x* start init [Full]
Stop an ongoing initialization on the virtual drive.	/c*x*/u*x* stop verify **NOTE** Only the bios can do a foreground initialization. A background initialization does otherwise. A verify starts a back ground initialization	/c*x*/v*x* stop init
Show a snapshot of the ongoing initialization, if any.	/c*x*/u*x* show [all] **NOTE** Only the bios can do a foreground initialization. A background initialization does otherwise. A verify starts a back ground initialization.	/c*x*/v*x* show init
Start a consistency check on the virtual drive.	/c*x*/u*x* start verify	/c*x*/v*x* start cc
Stop a consistency check on the virtual drive.	/c*x*/u*x* stop verify	/c*x*/v*x* stop cc
Reconstruct the selected virtual disk to a new RAID level.	/c*x*/u*x* migrate type=<*RaidType*> [disk=<*p*:-*p*..>] [strip=<*size*>] [noscan] [nocache] [autoverify] [*group*=<3|4|5|6|7|8|9|10|11|12|13|14|15|16>] *RaidType*={ raid0, raid1, raid5, raid10, raid50, single, raid6 }	/c*x*/v*x* start migrate <type=*raidlevel*> [option=<add | remove> disk=<e1:s1,e2:s2 ..>] /c*x*/v*x* show migrate
Change the power-saving setting on the virtual drive.	/c*x*/u*x* set powersavestandbytimer=<5 to 999>	/c*x*/v*x* set ds=<default | Auto | None | Max | MaxNoCache>

B.7 Physical Drive Commands

Table 69 Physical Drive Commands

Description	3ware CLI Command	StorCLI Command
Show physical disk information.	/c*x*/p*x* show [all]	/c*x*[/e*x*]/s*x* show [all]
Start, stop, suspend, or resume an ongoing rebuild operation.	/c*x*/u*x* start rebuild disk=<p:-p...> [ignoreECC] **NOTE** Rebuilds cannot be stopped or paused.	/c*x*[/e*x*]/s*x* start rebuild /c*x*[/e*x*]/s*x* stop rebuild /c*x*[/e*x*]/s*x* pause rebuild /c*x*[/e*x*]/s*x* resume rebuild
Mark the configured physical disk drive as missing for the selected adapter.	/c*x*/p*x* remove [quiet]	/c*x*[/e*x*]/s*x* set missing
Change the physical disk drive state to offline.	/c*x*/p*x* remove [quiet]	/c*x*[/e*x*]/s*x* set offline
Add jbod.	/cx add vd type=jbod disk=<*p*> (where *p* = port or drive number)	/c*x*[/e*x*]/s*x* set jbod
Change the physical disk drive hot spare state and associate the drive to an enclosure and virtual disk.	/cx add vd type=spare disk=<*p*:*p*|*p*-*p*|*p*:*p*-*p*> (where *p* = port or drive number)	/c*x*[/e*x*]/s*x* add hotsparedrive [{dgs=<N|0,1.2...*n*,,>] [EnclAffinity][nonRevertible]

Table 69 Physical Drive Commands (Continued)

Description	3ware CLI Command	StorCLI Command
Locate the physical disk drive and activate the physical disk activity LED.	/c*x*/p*x* set identify=on|off	/c*x*[/e*x*]/s*x* start | stop locate
Prepare the unconfigured physical drive for removal.	/c*x*/p*x* remove [quiet]	/c*x*[/e*x*]/s*x* spindown
Show information about all physical disk drives and other devices connected to the selected adapters; includes drive type, size, serial number, and firmware version.	/c*x*/p*x* show [all]	/c*x*/eall/sall show [all]
Download drive or expander firmware.	/c*x*/p*x* update fw=*image.name* [force]	/c*x*[/e*x*]/s*x* download src=filepath [satabridge]

B.8 Enclosure Commands

Table 70 Enclosure Commands

Description	3ware CLI Command	StorCLI Command
Show information about the enclosure for the selected adapter.	/c*x*/e*x* show [all]	/c*x*/e*x* show [all]
Show the status of the enclosure connected to the selected adapter.	/c*x*/e*x* show [all] /c*x*/e*x* show controllers /c*x*/e*x* show slots /c*x*/e*x* show fans /c*x*/e*x* show temp /c*x*/e*x* show pwrs /c*x*/e*x* show alms	/c*x*/e*x* show status
Download enclosure firmware.	/c*x*/e*x* update fw=*image.name* [force]	/c*x*/e*x* download src=*filepath* [offline] [forceActivate]

B.9 Events and Logs

Table 71 Events and Logs

Description	3ware CLI Command	StorCLI Command
Show the total number of events, newest and oldest sequence number, shutdown sequence number, reboot sequence number, clear sequence number.	/*cx* show alarms **NOTE** This command shows AENs since last controller reset.	/*cx* show eventloginfo
Show the total event entries available at the firmware since last clear, and details of each entries of error log.	/*cx* show alarms **NOTE** This command shows AENs since last controller reset.	/*cx* show events filter=<Info | warning| critical| fatal > file=<*path of the file*>
Show the count of events starting from specified seqNum and matching category and severity	/*cx* show alarms **NOTE** This command shows AENs since last controller reset.	/*cx* show events type=<sinceShutDown | sinceReboot |ccincon vd=<0,1,2…> | includeDeleted | latest=*x* filter=<Info | warning| critical| fatal > file=<*path of the file*>
Show TTY firmware terminal log entries with details on given adapters. The information is shown as total number of entries available on the firmware side.	/*cx* show diag	/*cx* show TermLog [type=contents|Config]

B.10 Miscellaneous Commands

Table 72 Miscellaneous Commands

Description	3ware CLI Command	StorCLI Command
Show version information.	tw_cli ?	ver
Show help for all show commands at server level.	tw_cli ? tw_cli /*cx* ? tw_cli /*cx*/u*x* ? tw_cli /*cx*/p*x* ? tw_cli /*cx*/phy*x* ? tw_cli /*cx*/bbu ? tw_cli /*cx*/e*x* ? tw_cli /e*x* **NOTE** The 3ware CLI shows context-sensitive help.	show help
Show PHY connection information for physical PHY medium on the adapters.	/*cx*/phy*x* show	/*cx*/p*x* show
Set PHY link speed.	/*cx*/phy*x* set link=<0|1.5|3.0|6.0|12.0>	/cx/px set linkspeed=0(auto)|1.5|3|6|12

Appendix C: MegaCLI Commands to StorCLI Command Conversion

C.1 System Commands

Table 73 System Commands

Description	MegaCLI Command	StorCLI Command
Show the software version.	`MegaCLI -v`	`storcli -v`
Show help information.	`MegaCLI -help|-h|?`	`storcli -help|-h|?`
Show the number of controllers connected.	`MegaCLI -adpCount`	`storcli show ctrlcount`

C.2 Controller Commands

Table 74 Controller Commands

Description	MegaCLI Command	StorCLI Command
Show the status of properties related to the controllers.	`MegaCli -AdpGetProp <PropertyName>-aN|-a0,1,2|-aALL`	`/cx show <propertyName>`
	The following properties can be used with this command:	The following properties can be used with this command:
	`abortcconerror`	`abortcconerror`
	`alarmdsply`	`alarm`
	`autodetectbackplanedsbl`	`backplane`
	`autoenhancedimportdsply`	`foreignautoimport`
	`autosnapshotspace`	
	`batwarndsbl`	`batterywarning`
	`bgirate`	`bgirate`
	`bootwithpinnedcache`	`bootwithpinnedcache`
	`cachebypass`	`cachebypass`
	`ccrate`	`ccrate`
	`clusterenable`	
	`coercionmode`	`coercion`
	`copybackdsbl`	`copyback`
	`defaultldpspolicy`	`ds`
	`defaultsnapshotspace`	
	`defaultviewspace`	

Table 74 Controller Commands (Continued)

Description	MegaCLI Command	StorCLI Command
	disableldpsinterval	ds
	disableldpstime	ds
	disableocr	ocr
	eccbucketcount	eccbucketsize
	eccbucketleakrate	eccbucketleakrate
	enableeghsp	eghs
	enableesmarter	eghs
	enableeug	eghs
	enablejbod	Jbod
	enblspindownunconfigdrvs	ds
	loadbalancemode	loadbalancemode
	maintainpdfailhistoryenbl	maintainpdfailhistory
	ncqdsply	ncq
	patrolreadrate	prrate
	perfmode	perfmode
	predfailpollinterval	smartpollinterval
	rebuildrate	rebuildrate
	reconrate	migraterate
	rstrhotspareoninsert	restorehotspare
	smartcpybkenbl	copyback
	spindowntime	ds
	spinupencdelay	ds
	spinupdelay	spinupdelay
	spinupencdrvcnt	spinupdrivecount
	ssdsmartcpybkenbl	copyback
	usediskactivityforlocate	activityforlocate
	usefdeonlyencrypt	usefdeonlyencrypt
Set properties on the selected controllers.	Megacli -AdpSetProp <*propertyname*>-an|-a0,1,2|-aall	/cx set <*property1*>
	The following properties can be set using this command:	The following properties can be set using this command:
	abortcconerror	abortcconerror=<on|off>
	alarmdsply	alarm=<on|off| silence>
	autodetectbackplanedsbl	backplane=<*value*>
	autoenhancedimportdsply	foreignautoimport=<on|off>
	batwarndsbl	batterywarning=<on|off>
	bgirate	bgirate=<*value*>
	bootwithpinnedcache	bootwithpinnedcache=<on|off>
	cachebypass	cachebypass=<on|off>

Table 74 Controller Commands (Continued)

Description	MegaCLI Command	StorCLI Command
	ccrate	ccrate=<*value*>
	clusterenable	
	coercionmode	coercion=<*value*>
	copybackdsbl	copyback=<on|off> type=<smartssd|smarthdd|all>
	defaultldpspolicy	ds=<*value*>
	defaultsnapshotspace	
	defaultviewspace	
	disableldpsinterval	ds=<*value*>
	disableldpstime	ds=<*value*>
	disableocr	ocr=<*value*>
	eccbucketcount	eccbucketsize=<*value*>
	eccbucketleakrate	eccbucketleakrate=<*value*>
	enableeghsp	eghs [state=<on|off>]
	enableesmarter	eghs [smarter=<on|off>]
	enableeug	eghs [eug=<on|off>
	enablejbod	jbod=<on|off>
	enblspindownunconfigdrvs	ds=<*value*>
	loadbalancemode	loadbalancemode=<*value*>
	maintainpdfailhistoryenbl	maintainpdfailhistory=<on|off>
	ncqdsply	ncq=<on|off>
	patrolreadrate	prrate=<*value*>
	perfmode	perfmode=<*value*>
	predfailpollinterval	smartpollinterval=<*value*>
	rebuildrate	rebuildrate=<*value*>
	reconrate	migraterate=<*value*>
	rstrhotspareoninsert	restorehotspare=<on|off>
	smartcpybkenbl	copyback=<on|off> type=<smartssd|smarthdd|all>
	spindowntime	ds=<on|off>
	spinupdelay	spinupdelay=<*value*>
	spinupdrivecount	spinupdrivecount=<*value*>
	spinupencdelay	ds
	spinupencdrvcnt	ds
	sdsmartcpybkenbl	copyback=<on|off> type=<smartssd|smarthdd|all>
	usediskactivityforlocate	activityforlocate=<on|off>
	usefdeonlyencrypt	usefdeonlyencrypt=<on|off>
Show the number of controllers connected.	MegaCLI -adpCount	storcli show ctrlcount

Table 74 Controller Commands (Continued)

Description	MegaCLI Command	StorCLI Command
Show all information about the adapter, such as cluster state, BIOS, alarm, firmware, version, and so on.	MegaCli -AdpAllInfo -aN|-a0,1,2|-aALL	storcli /c*x* show all
Show the freespace available in the controller.	MegaCLI -CfgFreeSpaceinfo -aN|-a0,1,2|-aALL	storcli /c*x* show freespace
Download the controller firmware.	MegaCli -AdpFwFlash -f *filename* [-NoSigChk] [-NoVerChk] [-ResetNow] -aN|-a0,1,2|-aALL	storcli /c*x* download file=<*filepath*> [fwtype=<*val*>] [nosigchk] [noverchk][resetnow]
Show the preserved cache status.	MegaCLI-GetPreservedCacheList -aN|-a0,1,2|-aALL	storcli /c*x* show preservedcache
Set the controller time	MegaCLI –AdpSetTime *yyyymmdd hh*:*mm*:*ss* -aN|-a0,1,2|-aALL	storcli /c(x|all) set time=<*yyyymmdd hh*:*mm*:*ss* | systemtime>
Show the controller time.	MegaCLI –AdpGetTime -aN	storcli /c*x* show time

C.3 Patrol Read Commands

Table 75 Patrol Read Commands

Description	MegaCLI Command	StorCLI Command
Show the patrol read status and patrol read parameters, if any in progress.	MegaCli -AdpPR -info -aN|-a0,1,2|-aALL	storcli/c*x* show patrolRead
Set the patrol read options on a single adapter, multiple adapters, or all adapters. (x = single controller).	MegaCli -AdpPR -Dsbl|EnblAuto|EnblMan|Start|Stop | Info|Suspend|Resume|Stop| SSDPatrolReadEnbl | SSDPatrolReadDsbl |{SetDelay *Val*}|{-SetStartTime *yyyymmdd hh*}|{maxConcurrentPD Val} -aN|-a0,1,2|-aALL	storcli /c*x* set patrolread {=on mode=<auto|manual>}|{off} storcli /c*x* set patrolread [starttime=<*yyyy*/*mm*/*dd hh*>] [maxconcurrentpd=<*value*>] [includessds=<on|off>] [uncfgareas=on|off] storcli /c*x* set patrolread delay=<*value*>
Disable patrol read.	MegaCli -AdpPR -Dsbl -aN|-a0,1,2|-aALL	storcli /c*x* set patrolread=off
Enable automatic patrol read.	MegaCli -AdpPR -EnblAuto -aN|-a0,1,2|-aALL	storcli /c*x* set patrolread=on mode=auto
Enable manual patrol read.	MegaCli -AdpPR -EnblMan -aN|-a0,1,2|-aALL	storcli /c*x* set patrolread=on mode=manual
Start patrol read.	MegaCli -AdpPR -Start -aN|-a0,1,2|-aALL	storcli /c*x* start patrolRead
Suspend a running patrol read.	MegaCli -AdpPR -Suspend -aN|-a0,1,2|-aALL	storcli /c*x* suspend patrolRead
Resume a suspended patrol read.	MegaCli -AdpPR -Resume -aN|-a0,1,2|-aALL	storcli /c*x* resume patrolRead
Stop a running patrol read.	MegaCli -AdpPR -Stop -aN|-a0,1,2|-aALL	storcli /c*x* stop patrolRead
Include SSD drives in patrol read.	MegaCli -AdpPR -SSDPatrolReadEnbl -aN|-a0,1,2|-aALL	storcli /c*x* set patrolRead includessds=on | onlymixed

Table 75 Patrol Read Commands (Continued)

Description	MegaCLI Command	StorCLI Command
Exclude SSD drives in patrol read.	MegaCli -AdpPR -SSDPatrolReadDsbl -aN|-a0,1,2|-aALL	storcli /cx set patrolRead includessds=off
Delay a patrol read,	MegaCli -AdpPR -SetDelay Val -aN|-a0,1,2|-aALL	storcli /cx set patrolread delay=<*value*>
Schedule a patrol read.	MegaCli -AdpPR -SetStartTime *yyyymmdd hh* -aN|-a0,1,2|-aALL	storcli /cx set patrolread=on starttime=*YYYY/MM/DD HH*
Set the value for maximum concurrent physical drives for the patrol read.	MegaCli -AdpPR -maxConcurrentPD Val -aN|-a0,1,2|-aALL	storcli /cx set patrolread maxconcurrentpd=*xx*

C.4 Consistency Check Commands

Table 76 Consistency Check Commands

Description	MegaCLI Command	StorCLI Command
Schedule a consistency check.	MegaCLI -AdpCcSched -Dsbl|-Info| {-ModeConc | -ModeSeq [-ExcludeLD -LN|-L0,1,2] [-SetStartTime *yyyymmdd hh*] [-SetDelay *val*] } -aN|-a0,1,2|-aALL	storcli /cx set consistencycheck|cc=[off|seq|conc] [delay=value] starttime=*yyyy/mm/dd hh* [excludevd=x-y,z]
Show consistency check status and consistency parameters, in progress, if any.	MegaCLI -AdpCcSched -Info	storcli /cx show cc/ConsistencyCheck

C.5 OPROM BIOS Commands

Table 77 OPROM BIOS Commands

Description	MegaCLI Command	StorCLI Command
Schedule a consistency check.	MegaCli -AdpBIOS -Dsply -aN|-a0,1,2|-aALL	storcli /cx show bios
Show consistency check status and consistency parameters, if any in progress.	MegaCli -AdpBootDrive -{-Set {-Lx | -physdrv[E0:S0]}} -aN|-a0,1,2|-aALL	storcli /cx/ex/sx set bootdrive=on|off storcli /cx/vx set bootdrive=on|off
Sets the BIOS properties for the controller.	MegaCli -AdpBIOS -Enbl | -Dsbl | -Dsply | SOE | BE EnblAutoSelectBootLd | DsblAutoSelectBootLd -aN|-a0,1,2|-aALL	storcli /cx set bios=<on|off> storcli /cx set stoponerror|SOE=<on|off> storcli /cx set autobootselect(abs)=<on|off>

C.6 Battery Commands

Table 78 Battery Commands

Description	MegaCLI Command	StorCLI Command
Show battery-related information.	`MegaCli -AdpBbuCmd -aN|-a0,1,2|-aALL`	`storcli /cx/bbu show` `storcli /cx/bbu show all`
Show the battery learn properties.	`MegaCli -AdpBbuCmd -GetBbuProperties -aN|-a0,1,2|-aALL`	`storcli /cx/bbu show properties`
Show the battery information, firmware status, and the gas gauge status.	`MegaCli -AdpBbuCmd -GetBbuStatus -aN|-a0,1,2|-aALL`	`storcli /cx/bbu show status`
Show battery capacity information.	`MegaCli -AdpBbuCmd -GetBbuCapacityInfo -aN|-a0,1,2|-aALL`	`storcli /cx/bbu show all`
Show battery design information.	`MegaCli -AdpBbuCmd -GetBbuDesignInfo -aN|-a0,1,2|-aALL`	`storcli /cx/bbu show all`

Set battery properties	`MegaCli -AdpBbuCmd -SetBbuProperties -f <fileName> -aN|-a0,1,2|-aALL`	`storcli /cx/bbu set learnDelayInterval=<value>`
		`storcli /cx/bbu set bbuMode=<value>`
		`storcli /cx/bbu set autolearnmode=<value>` where x= 0 – Enabled, 1 – Disabled, 2 – Warn though event.

Start battery learn cycle.	`MegaCli -AdpBbuCmd -BbuLearn -aN|-a0,1,2|-aALL`	`storcli /cx/bbu start learn`
Set the battery to low power storage mode.	`MegaCli -AdpBbuCmd -BbuMfgSleep -aN|-a0,1,2|-aALL`	`storcli /cx/bbu set powermode=sleep`
Seal the gas gauge EEPROM write access	`MegaCli -AdpBbuCmd -BbuMfgSeal -aN|-a0,1,2|-aALL`	`storcli /cx/bbu set writeaccess=sealed`

C.7 RAID Configuration Commands

Table 79 RAID Configuration Commands

Description	MegaCLI Command	StorCLI Command
Create a RAID configuration of RAID type 0, 1, 5, and 6.	MegaCli -CfgLDAdd -R0|-R1|-R5|-R6[E0:S0,E1:S1,...] [WT | WB] [NORA | RA | ADRA] [Direct | Cached] [CachedBadBBU|NoCachedBadBBU] [-sz*XXXXXXXX* [-sz*YYYYYYYY* [...]]] [-strpszM] [-Hsp[E5:S5,...]] [-afterLdX] -aN	storcli /c*x* add vd type=raid[0|1|5|6] [Size=<*VD1_Sz*>,< *VD2_Sz*>,..|*all] [name=<*VDNAME1*>,..] drives=e:*s*|e:*s-x*|e:*s-x*,*y*;e:*s-x*,*y*,*z* [PDperArray=*x*] [SED] [pdcache=on|off|*default][pi] [DimmerSwitch(ds)=default|automatic(auto)|*none|maximum(max) |MaximumWithoutCaching(maxnocache)] [wt|*wb|awb] [nora|*ra] [*direct|cached] [strip=<8|16|32|64|128|256|512|1024] [AfterVd=*X*] [Spares=[e:]*s*|[e:]*s-x*|[e:]*s-x*,*y*] [force] **NOTE** The supported strip size can vary from a minimum of 64 KB to 1 MB for MegaRAID controllers and only 64 KB for Integrated MegaRAID controllers. The LSISAS2108 controller supports strip size from 8 KB to 1 MB.
Create a CacheCade virtual drive.	MegaCLI -CfgCacheCadeAdd [-rX] -Physdrv[E0:S0,...] {-Name LdNamestring} [WT|WB|ForcedWB] [-assign -LX|L0,2,5..|LALL] -aN|-a0,1,2|-Aall	storcli /c*x* add vd cachecade|cc Type=raid[0,1,10] drives=[e:]*s*|[e:]*s-x*|[e:]*s-x*,*y* [< WT| WB>] [assignvds=0,1,2]*e*:]
Create a RAID configuration of RAID type 10, 50, and 60.	MegaCli -CfgSpanAdd -aN|-a0,1,2|-aALL -R10|-R50|R60 -Array0[E0:S0,E1:S1,...] -Array1[E0:S0,E1:S1,...] [...] [WT | WB] [NORA | RA | ADRA] [Direct | Cached] [CachedBadBBU|NoCachedBadBBU] [-sz*XXXXXXXX*[-sz*YYYYYYYY* [...]]] [-strpszM] [-afterLdX] -aN	storcli /c*x* add vd type=raid[10|50|60] [Size=<*VD1_Sz*>,<*VD2_Sz*>,..|*all] [name=<*VDNAME1*>,..] drives=e:*s*|e:*s-x*|e:*s-x*,*y*;e:*s-x*,*y*,*z* [PDperArray=*x*] [SED] [pdcache=on|off|*default][pi] [DimmerSwitch(ds)=default|automatic(auto)|*none|maximum(max) |MaximumWithoutCaching(maxnocache)] [wt|*wb|awb] [nora|*ra] [*direct|cached] [strip=<8|16|32|64|128|256|512|1024] [AfterVd=*X*] [Spares=[e:]*s*|[e:]*s-x*|[e:]*s-x*,*y*] [force] **NOTE** The supported strip size can vary from a minimum of 64 KB to 1 MB for MegaRAID controllers and only 64 KB for Integrated MegaRAID controllers. The LSISAS2108 controller supports strip size from 8 KB to 1 MB.
Clear the complete configuration.	MegaCli -CfgClr [-Force] -aN|-a0,1,2|-aALL	storcli /c*0*/vall delete [force]
Show the topology information of the drive group.	MegaCLI -CfgDsply -aN|-a0,1,2|-Aall	storcli /c*x*/dall show [all]
Show information for a CacheCade virtual drive.	MegaCLI -CfgCacheCadeDsply -aN|-a0,1,2|-Aall	storcli /c*x*/dall show CacheCade(cc)

Table 79 RAID Configuration Commands (Continued)

Description	MegaCLI Command	StorCLI Command
Delete a virtual drive hosting the operating system.	MegaCLI -CfgLdDel -LX|-L0,2,5...|-LALL [-Force] -aN|-a0,1,2|-aALL	storcli /c*x*/v/v*x* [all] delete -force
Delete a CacheCade virtual drive.	MegaCLI -CfgCacheCadeDel -LX|-L0,2,5...|-LALL -aN|-a0,1,2|-Aall	storcli /c*x*/v*x* [all] delete CacheCade(cc)
Show, delete, and import the foreign configuration commands.	MegaCli -CfgForeign -Scan | {-Preview | -Dsply| -Import | -Clear[FID]} -aN|-a0,1,2|-aALL"	storcli /c*x*/f(*x*|all) show [all] [securityKey=*xxx*] storcli /c*x*/f(*x*|all) del|delete [securityKey=*xxx*] storcli /c*x*/f(*x*|all) import [preview] [securityKey=*xxx*]"

C.8 Security Commands

Table 80 Security Commands

Description	MegaCLI Command	StorCLI Command
Set the key ID for the controller.	MegaCli -CreateSecurityKey -SecurityKey *sssssssssss* | [-Passphrase *sssssssssss*] |[-KeyID *kkkkkkkkkk*] -aN	storcli /c*x* set SecurityKey=*XXXXXX* [passphrase=*yyyyy*] [keyId=*zzzz*]
Change the security key for the controller.	MegaCli -ChangeSecurityKey -OldSecurityKey *sssssssssss* | -SecurityKey *sssssssssss* | [-Passphrase *sssssssssss*] | [-keyID *kkkkkkkkkk*] -aN	storcli /c*x* set SecurityKey=*XXXXXX* OldSecurityKey=*yyyyy*
Compare and verify the security key for the controller.	MegaCli -VerifySecurityKey -SecurityKey *sssssssssss* -aN	storcli /c*x* compare SecurityKey=*xxxxxx*
Delete the security key.	MegaCLI -DestroySecurityKey | [-Force] -aN	storcli /c*x* delete SecurityKey
Set the security key for the controller.	MegaCli -SetKeyID -KeyID *kkkkkkkkkk* -aN	storcli /c*x* set SecurityKey KeyId=*xxxx*

C.9 Virtual Drive Commands

Table 81 Virtual Drive Commands

Description	MegaCLI Command	StorCLI Command
Show the virtual drive information.	MegaCli -LDInfo -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) show storcli /c*x*/v(*x*|all) show all
Set virtual drive properties.	MegaCli -LDSetProp WT | WB|NORA |RA | ADRA|-Cached|Direct|CachedBadBBU| NoCachedBadBBU} | -RW|RO|Blocked | {-Name nameString} |-EnDskCache|DisDskCache -Lx| -L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) set wrcache=WT|WB|AWB storcli /c*x*/v(*x*|all) set rdcache=RA|NoRA storcli /c*x*/v(*x*|all) set iopolicy=Cached|Direct storcli /c*x*/v(*x*|all) set accesspolicy=RW|RO|Blocked|RmvBlkd storcli /c*x*/v(*x*|all) set pdcache=On|Off|Default storcli /c*x*/v(*x*|all) set name=<*NameString*>
Set power-saving (dimmer switch) properties.	MegaCli -LDSetPowerPolicy -Default| -Automatic| -None| -Maximum| -MaximumWithoutCaching -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) set ds=Default|Auto|None|Max|MaxNoCache
Show virtual drive expansion information.	MegaCli -getLdExpansionInfo -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) show expansion
Expand the virtual drive within the existing array; also use if you replace the drives with larger drives, beyond the size of the existing array.	MegaCli -LdExpansion -pN -dontExpandArray -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) expand Size=<*value*> [expandarray]
Secure the virtual drive.	MegaCLI --LDMakeSecure -Lx|-L0,1,2,...|-Lall -An	storcli /c*x*/v*x* set security=on
Show specific properties of virtual drives.	MegaCli -LDGetProp -Cache | -Access | -Name | -DskCache -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /c*x*/v*x* show
Start virtual drive initialization.	MegaCli -LDInit -Start [Fast|Full] -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) start init[Full]
Stop a running virtual drive initialization.	MegaCli -LDInit -Abort -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) stop init
Show the initialization progress.	MegaCli -LDInit -ShowProg -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) show init
Start a consistency check on an uninitialized virtual drive.	MegaCli -LDCC -Start -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) start cc[Force]

Table 81 Virtual Drive Commands (Continued)

Description	MegaCLI Command	StorCLI Command
Start, stop, suspend, resume, and show the progress of a consistency check operation.	MegaCli -LDCC -Start|-Abort| -Suspend|-Resume|-ShowProg| -ProgDsply -Lx|-L0,1,2|-LALL -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) start cc storcli /c*x*/v(*x*|all) stop cc storcli /c*x*/v(*x*|all) pause cc storcli /c*x*/v(*x*|all) resume cc storcli /c*x*/v(*x*|all) show cc
Enable/disable automatic background initialization. Show, stop, pause, resume, and show the progress of the background initialization.	MegaCLI -LDBI -Enbl|-Dsbl| -getSetting|-Abort|-Suspend| -Resume|-ShowProg|-ProgDsply -Lx|-L0,1,2|-LALL -aN|-a0,1,2|-Aall	storcli /c*x*/v(*x*|all) set autobgi=On|Off storcli /c*x*/v(*x*|all) show autobgi storcli /c*x*/v(*x*|all) stop bgi storcli /c*x*/v(*x*|all) pause bgi storcli /c*x*/v(*x*|all) resume bgi storcli /c*x*/v(*x*|all) show bgi
Start and show progress for a migrate operation.	MegaCli -LDRecon {-Start -Rx [Add | Rmv PhysDrv[E0:S0,E1:S1,...]] } | -ShowProg|-ProgDsply -Lx -aN	storcli /c*x*/v*x* start migrate type=raidx [option=add|remove drives=[e:]*s*|[e:]*s*-*x*|[e:]*s*-*x*,*y*] [Force] storcli /c*x*/v(*x*|all) show migrate
Delete preserved cache.	MegaCLI -DiscardPreservedCache -Lx|-L0,1,2|-Lall -force -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) delete preservedcache[force]
Assign the CacheCade virtual drive.	MegaCLI -Cachecade -assign|-remove -Lx|-L0,1,2|-LALL -aN|-a0,1,2|-aALL	storcli /c*x*/v*x*|all set ssdCaching=on|off

C.10 Physical Drive Commands

Table 82 Physical Drive Commands

Description	MegaCLI Command	StorCLI Command
Show drive information.	MegaCli -pdInfo -PhysDrv[E0:S0,E1:S1,...] -aN|-a0,1,2|-aALL	storcli /c*x*/e*x*/s*x* show storcli /c*x*/e*x*/s*x* show all
Start, stop, pause, resume, or show the progress of a rebuild operation.	MegaCLI PDRbld -Start|-Stop |-Suspend|-Resume|-ShowProg |-ProgDsply -PhysDrv [E0:S0,E1:S1,...] -aN|-a0,1,2|-aALL	storcli /c*x*/e*x*/s*x* start rebuild storcli /c*x*/e*x*/s*x* stop rebuild storcli /c*x*/e*x*/s*x* pause rebuild storcli /c*x*/e*x*/s*x* resume rebuild storcli /c*x*/e*x*/s*x* shnow rebuild
Start, stop, pause, resume, or show the progress of a copyback operation.	MegaCLI PDCpyBk -Start|-Stop |-Suspend|-Resume|-ShowProg |-ProgDsply -PhysDrv [E0:S0,E1:S1,...] -aN|-a0,1,2|-aALL	storcli /c*x*/e*x*/s*x* start copyback target = e*xx*:s*xx* storcli /c*x*/e*x*/s*x* stop copyback storcli /c*x*/e*x*/s*x* pause copyback storcli /c*x*/e*x*/s*x* resume copyback storcli /c*x*/e*x*/s*x* show copyback
Mark a drive as missing.	MegaCli -PdMarkMissing -physdrv[E0:S0,E1:S1,...] -aN|-a0,1,2|-aALL	storcli /c*x*/e*x*/s*x* set missing

Table 82 Physical Drive Commands (Continued)

Description	MegaCLI Command	StorCLI Command
Show missing drive information.	MegaCli -PdGetMissing -aN|-a0,1,2|-aALL	storcli /cx/ex/sx show all **NOTE** This information is shown as part of the show all command.
Replace the configured drive that is identified as missing, and then start an automatic rebuild.	MegaCli -PdReplaceMissing -physdrv[E0:S0] -arrayA, -rowB -aN	storcli /cx/ex/sx insert array=*x* row=*y*
Set the drive state to online	MegaCli -PDOnline -PhysDrv[E0:S0,E1:S1....] -aN|-a0,1,2	storcli /cx/ex/sx set online
Set the drive state to offline.	MegaCli -PDOffline -PhysDrv[E0:S0,E1:S1....] -aN|-a0,1,2|-aALL	storcli /cx/ex/sx set offline
Set the drive state to JBOD	MegaCli -PDMakeGood -PhysDrv[E0:S0,E1:S1....] -aN|-a0,1,2|-aALL	storcli /cx/ex/sx set good [force]
Set the drive state to JBOD	MegaCli -PDMakeJBOD -PhysDrv[E0:S0,E1:S1,...] -aN|-a0,1,2|-aALL	storcli /cx/ex/sx set jbod
Add and delete hot spare drives.	MegaCli -PDHSP {-Set [{-Dedicated -ArrayN |-Array0,1...}] [-EnclAffinity] [-nonRevertible] } | -Rmv -PhysDrv[E0:S0,E1:S1,...] -aN|-a0,1,2|-aALL	storcli /cx/ex/sx add hotsparedrive [dgs=<N|0,1,2..>] enclaffinity nonrevertible storcli /cx/ex/sx delete hotsparedrive
Start, stop, pause, resume or show the progress of an initialization process.	MegaCli -PDClear -Start |-Stop| -ShowProg |-ProgDsply -PhysDrv[E0:S0,E1:S1....] -aN|-a0,1,2|-aALL	storcli /cx/ex/sx start initialization storcli /cx/ex/sx stop initialization storcli /cx/ex/sx pause initialization storcli /cx/ex/sx resume initialization storcli /cx/ex/sx show initialization
Start a drive locate and activate the drive's LED or stop a drive locate and deactivate the drive's LED.	MegaCli -PDLocate {[-start] | -stop} -physdrv[E0:S0,E1:S1,...] -aN|-a0,1,2|-aALL	storcli /cx/ex/sx start locate storcli /cx/ex/sx stop locate
Spin down an unconfigured drive and prepare it for removal or spin up spun-down drive and mark the drive state as unconfigured good.	MegaCli -PDPrpRmv [-Undo] -PhysDrv[E0:S0,E1:S1....] -aN|-a0,1,2|-aALL	storcli /cx/ex/sx spindown storcli /cx/ex/sx spinup
Show physical drive information of all connected drives.	MegaCli -PDList -aN|-a0,1..|-aAll	storcli /cx/eall/sall show [all] **NOTE** This command does not show drives whose enclosure device ID is not available.
Flash the physical drive firmware.	MegaCLI PdFwDownload[offline] [ForceActivate] {[-SataBridge] -PhysDrv[0:1]}|{-EncdevId[*devId1*]} -f <*filename*> -aN|-a0,1,2|-Aall	storcli /cx[/ex]/sx download src=<*filepath*> [satabridge] [mode= 5|7] storcli /cx/ex download src=<*filepath*> [forceActivate]
Erase the drive's security configuration and securely erase data on a drive.	MegaCli -PDInstantSecureErase -PhysDrv[E0:S0,E1:S1,...] | [-Force] -aN|-a0,1,2|-aALL	storcli /cx/ex/sx secureerase [force]

Table 82 Physical Drive Commands (Continued)

Description	MegaCLI Command	StorCLI Command
Show the security key for secured physical drives	MegaCli -GetKeyID [-PhysDrv[E0:S0]] -aN	storcli /cx/ex/sx securitykey keyid
Start, stop, and show the progress of a secure erase operation	MegaCli -SecureErase Start[Simple| [Normal [|ErasePattern ErasePatternA|ErasePattern ErasePatternA ErasePattern ErasePatternB]]|[Thorough [|ErasePattern ErasePatternA|ErasePattern ErasePatternA ErasePattern ErasePatternB]]] | Stop| ShowProg| ProgDsply [-PhysDrv [E0:S0,E1:S1,...] | -Lx|-L0,1,2|-LALL] -aN|-a0,1,2|-aALL	storcli /cx[/ex]/sx start erase [simple | normal | thorough] [erasepatternA=<*val*>]\n[erasepatternB=<*val*>] Examples: storcli /cx/ex/sx start erase simple storcli /cx/ex/sx start erase normal erasepatterna=10101010 storcli /cx/ex/sx start erase thorough erasepatterna=10101010 erasepatternb=10101111 storcli /cx/ex/sx stop erase
Enable/disable the direct physical drive mapping mode.Show the current state of the direct physical drive mapping.	MegaCLI DirectPdMapping -Enbl|-Dsbl|-Dsply -aN|-a0,1,2|-Aall	storcli /cx set directpdmapping=<on | off> storcli /cx show directpdmapping

C.11 Enclosure Commands

Table 83 Enclosure Commands

Description	MegaCLI Command	StorCLI Command
Show enclosure information.	MegaCli –EncInfo -aN|-a0,1,2|-aALL	storcli /cx/ex show storcli /cx/ex show all
Show enclosure status.	MegaCli –EncStatus -aN|-a0,1,2|-aALL	storcli /cx/ex show status

C.12 PHY Commands

Table 84 PHY Commands

Description	MegaCLI Command	StorCLI Command
Show PHY information.	MegaCli –PHYInfo -phyM -aN|-a0,1,2|-aALL	storcli /cx/px(x|all) show storcli /cx/px(x|all) show all
Set PHY link speed.	MegaCLI PhySetLinkSpeed -phyM -speed -aN|-a0,1,2|-aALL	storcli /cx/px(x|all) set linkspeed=0(auto)|1.5|3|6|12
Show the PHY error counters.	Megacli PhyErrorCounters -An	storcli /cx/px(x|all) show storcli /cx/px(x|all) show all

C.13 Alarm Commands

Table 85 Alarm Commands

Description	MegaCLI Command	StorCLI Command
Show alarm properties.	MegaCli -AdpGetProp AlarmDsply -aN|-a0,1,2|-aALL	storcli /c*x*(*x*|all) show alarm
Set alarm properties.	MegaCli -AdpSetProp AlarmEnbl | AlarmDsbl | AlarmSilence -aN|-a0,1,2|-aALL	storcli /c*x*(*x*|all) set alarm=<on|off|silence>

C.14 Event Log Properties Commands

Table 86 Event Log Properties Commands

Description	MegaCLI Command	StorCLI Command
Show event logs.	MegaCli -AdpEventLog -GetEventLogInfo -aN|-a0,1,2|-aALL	storcli /c*x* show eventloginfo
Show the specified type of event logs.	MegaCli -AdpEventLog -GetEvents {-info -warning -critical -fatal} {-f <*fileName*>} -aN|-a0,1,2|-aALL	storcli /c*x* show events [[type=<sincereboot| sinceshutdown| includedeleted|latest=*x*| ccincon vd=<0,1,...>] filter=<info| warning| critical|fatal>] file=<*filepath*>
Show the specified event logs.	MegaCli -AdpEventLog -GetSinceShutdown {-info -warning -critical -fatal} {-f <*fileName*>} -aN|-a0,1,2|-aALL	storcli /c*x* show events [type=[latest=*x*|ccincon vd=|[sincereboot|sinceshutdown|includedeleted|latest|ccincon]]] [filter=[info|warning|critical|fatal]] file=*xyz*.txt
Delete the event logs.	MegaCli -AdpEventLog -Clear -aN|-a0,1,2|-aALL	storcli /c*x* delete events

C.15 Premium Feature Key Commands

Table 87 Premium Feature Key Commands

Description	MegaCLI Command	StorCLI Command
Show the Safe ID of the controller.	MegaCli -ELF -GetSafeId -a0	storcli /c*x*(*x*|all) show safeid
Show the Advanced Software Options that are enabled on the controller, including the ones in trial mode.	MegaCli -ELF -ControllerFeatures -a0	storcli /c*x*(*x*|all) show all **NOTE** This information shows as part of the controller show all.
Apply the Activation Key in preview mode.	MegaCli -ELF -Applykey key -val -preview -a0	storcli /c*x*(*x*|all) set aso key=<*key value*> preview
Apply the Activation Key.	MegaCli -ELF -Applykey key -val -a0	storcli /c*x*(*x*|all) set aso key=<*key value*>

Table 87 Premium Feature Key Commands (Continued)

Description	MegaCLI Command	StorCLI Command
Deactivate the trial key.	`MegaCli -ELF -DeactivateTrialKey -a0`	`storcli /cx(x|all) set aso deactivatetrialkey`
Show the re-host information and, if re-hosting is necessary, show the controller and key vault serial numbers.	`MegaCli -ELF -ReHostInfo -a0`	`storcli /cx(x|all) show rehostinfo`
Indicate to the controller that the re-host is complete.	`MegaCli -ELF -ReHostComplete -a0`	`storcli /cx(x|all) set aso rehostcomplete`

Appendix D: Unsupported Commands in Embedded MegaRAID

The commands in the following table are not supported in Embedded MegaRAID.

Table 88 Unsupported Commands in Embedded MegaRAID

Command Group	Command
Jbod	storcli /c0 set jbod=<on|off>
	storcli /c0/s2 set jbod
	storcli /c0/s2 set bootdrive=<on|off>
DS	storcli /c*x*(*x*|all) set ds=OFF type=1|2|3|4
	storcli /c*x*(*x*|all) set ds=ON type=1|2 [properties]
	storcli /c*x*(*x*|all) set ds=ON type=3|4 DefaultLdType=<*val*> [properties]
	storcli /c*x*(*x*|all) set ds [properties]
	storcli /c*x*/v(*x*|all) set ds=Default|Auto|None|Max|MaxNoCache
Security	storcli /c*x* delete security key
	storcli /c*x* set securitykey=*xxxxxxxx* {passphrase=*xxxx*} {keyid=*xxx*}
	storcli /c*x* set securitykey keyid=*xxx*
	storcli /c*x* compare securitykey=*xxxxxxxxxx*
	storcli /c*x* set securitykey=*xxxxxxxx* oldsecuritykey=*xxxxxxxx*
ASO	storcli /c*x*(*x*|all) set aso key=<*keyvalue*> preview
	storcli /c*x*(*x*|all) set aso key=<*key value*>
	storcli /c*x*(*x*|all) set aso transfertovault
	storcli /c*x*(*x*|all) set aso rehostcomplete
	storcli /c*x*(*x*|all) set aso deactivatetrialkey
	storcli /c*x*(*x*|all) show safeid
	storcli /c*x*(*x*|all) show rehostinfo
	storcli /c0 set time =<*yyyymmdd hh*:*mm*:*ss* | system>
	storcli /c0 show cc|consistencycheck
	storcli /c0/vall show expansion
	storcli /c0 set jbod
	storcli /c*x* download src=<*filepath*> [forceActivate]
Copy back	storcli /c*x*[/e*x*]/s*x* show copyback
	storcli /c*x*[/e*x*]/s*x* start copyback target=eID:sID
	storcli /c*x*[/e*x*]/s*x* stop copyback
	storcli /c*x*[/e*x*]/s*x* pause copyback
	storcli /c*x*[/e*x*]/s*x* resume copyback
Migrate	storcli /c*x*/v(*x*|all) show migrate
	storcli /c*x*/v*x* start migrate type=*raidx* [option=add|remove drives=[e:]*s*|[e:]*s-x*|[e:]*s-x,y*] [Force]
Cache	storcli /c*x*/v(*x*|all) set ssdcaching=on|off
	storcli /c*x*(*x*|all) show preservedcache
	storcli /c*x*/v(*x*|all) delete preservedcache[force]

Table 88 Unsupported Commands in Embedded MegaRAID (Continued)

Command Group	Command
BBU	storcli /c*x*/bbu show
	storcli /c*x*/bbu show all
	storcli /c*x*/bbu set [learnDelayInterval=<*val*> | bbuMode=<*val*>
	storcli /c*x*/bbu start learn
Secure ease	storcli /c*x*/s*x* secureerase [force]
	storcli /c*x*/s*x* start erase [simple| normal| thorough][erasepatternA=<*val*>]
	storcli /c*x*/s*x* stop erase
	storcli /c*x*/s*x* show erase
Consistency check	storcli /c*x* show cc/ConsistencyCheck
Controller	storcli /c*x* show cc

Appendix E: CLI Error Messages

This appendix lists the software error messages for the Storage Command Line Tool (StorCLI) and the MegaCLI Configuration Utility.

The Storage Command Line Tool (StorCLI) and the MegaCLI Configuration Utility are command line interface applications you can use to manage MegaRAID SAS RAID controllers.

E.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 the following table.

Table 89 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
0x17	LD operation not possible - CC is in progress
0x18	LD initialization in progress
0x19	LBA is out of range

Table 89 Error Messages and Descriptions (Continued)

Number	Event Text
0x1a	Maximum LDs are already configured
0x1b	LD is not OPTIMAL
0x1c	LD Rebuild is in progress
0x1d	LD is undergoing reconstruction
0x1e	LD RAID level is wrong for requested operation
0x1f	Too many spares assigned
0x20	Scratch memory not available - try command again later
0x21	Error writing MFC data to SEEPROM
0x22	Required HW is missing (i.e. Alarm or BBU)
0x23	Item not found
0x24	LD drives are not within an enclosure
0x25	PD CLEAR operation is in progress
0x26	Unable to use SATA(SAS) drive to replace SAS(SATA)
0x27	Patrol Read is disabled
0x28	Given row index is invalid
0x2d	SCSI command done, but non-GOOD status was received-see mf.hdr.extStatus for SCSI_STATUS
0x2e	IO request for MFI_CMD_OP_PD_SCSI failed - see extStatus for DM error
0x2f	Matches SCSI RESERVATION_CONFLICT
0x30	One or more of the flush operations failed
0x31	Firmware real-time currently not set
0x32	Command issues while firmware in wrong state (i.e., GET RECON when op not active)
0x33	LD is not OFFLINE - IO not possible
0x34	Peer controller rejected request (possibly due to resource conflict)
0x35	Unable to inform peer of communication changes (retry might be appropriate)
0x36	LD reservation already in progress
0x37	I2C errors were detected
0x38	PCI errors occurred during XOR/DMA operation
0x39	Diagnostics failed - see event log for details
0x3a	Unable to process command as boot messages are pending
0x3b	Returned in case if foreign configurations are incomplete
0x3d	Returned in case if a command is tried on unsupported hardware
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

Table 89 Error Messages and Descriptions (Continued)

Number	Event Text
0x45	Lock key backup cannot be verified
0x46	Lock key from backup failed verification
0x47	Rekey operation not allowed, unless controller already has a lock key
0x48	Lock key is not valid, cannot authenticate
0x49	Lock key from escrow cannot be used
0x4a	Lock key backup (pass-phrase) is required
0x4b	Secure LD exist
0x4c	LD secure operation is not allowed
0x4d	Reprovisioning is not allowed
0x4e	Drive security type (FDE or non-FDE) is not appropriate for requested operation
0x4f	LD encryption type is not supported
0x50	Cannot mix FDE and non-FDE drives in same array
0x51	Cannot mix secure and unsecured LD in same array
0x52	Secret key not allowed
0x53	Physical device errors were detected
0x54	Controller has LD cache pinned
0x55	Requested operation is already in progress
0x56	Another power state set operation is in progress
0x57	Power state of device is not correct
0x58	No PD is available for patrol read
0x59	Controller reset is required
0x5a	No EKM boot agent detected
0x5b	No space on the snapshot repository VD
0x5c	For consistency SET PiTs, some PiT creations might fail and some succeed
0x5d	Secondary iButton cannot be used and is incompatible with controller
0x5e	PFK doesn't match or cannot be applied to the controller
0x5f	Maximum allowed unconfigured (configurable) PDs exist
0x60	IO metrics are not being collected
0x61	AEC capture needs to be stopped before proceeding
0x62	Unsupported level of protection information
0x63	PDs in LD have incompatible EEDP types
0x64	Request cannot be completed because protection information is not enabled
0x65	PDs in LD have different block sizes
0x66	LD Cached data is present on a (this) SSCD
0x67	Config sequence number mismatch
0x68	Flash image is not supported
0x69	Controller cannot be online-reset
0x6a	Controller booted to safe mode, command is not supported in this mode

Table 89 Error Messages and Descriptions (Continued)

Number	Event Text
0x6b	SSC memory is unavailable to complete the operation
0x6c	Peer node is incompatible
0x6d	Dedicated hot spare assignment is limited to array(s) with same LDs.
0x6e	Signed component is not part of the image
0x6f	Authentication failure of the signed firmware image
0x70	Flashing was ok but FW restart is not required, ex: No change in FW from current
0x71	Firmware is in some form of restricted mode, example: passive in A/P HA mode
0x72	The maximum number of entries are exceed.
0x73	Cannot start the subsequent flush because the previous flush is still active.
0x74	Status is ok but a reboot is need for the change to take effect.
0x75	Cannot perform the operation because the background operation is still in progress.
0x76	Operation is not possible.
0x77	Firmware update on the peer node is in progress.
0x78	Hidden policy is not set for all of the virtual drives in the drive group that contains this virtual drive.
0x79	Indicates that there are one or more secure system drives in the system.
0x7a	Indicates that the virtual drive is a boot drive.
0x7b	Logical Drive count exceeds max transportable logical drive count.
0x7c	Dedicated Hot Spare is associated with more that one disk group.
0x7d	Configuration has one Logical Drive in Transport Ready state.
0x7e	IO request encountered a SCSI data under run error.
0x7f	Firmware flash not allowed in current mode.
0x80	Operation not possible since drive is in Transport ready state.
0x81	Operation not possible since Logical drive is in Transport ready state.
0x82	Operation not possible since Logical drive is not in Transport ready state.
0x83	Operation not possible since physical drive is in removal ready state.
0x84	Host reboot is required for changes to take effect.
0x85	Microcode update is pending on the device.
0x86	Microcode update is in progress on the device.
0xFF	Invalid status - used for polling command completion

Appendix F: 240 Virtual Drive Feature Limitations

This appendix provides information about limitations and known issues for the 240 virtual drives (VDs) feature in the MegaRAID 12Gb/s SAS RAID controller.

F.1 Host Software Utility Support

Table 90 Host Software Utilities Support Matrix

MegaRAID SAS RAID Utilities	0–63 VD Target ID's Support	240 VD Target ID's Support
StorCLI	Yes	Yes
MegaRAID Storage Manager™	Yes	No
SNMP	Yes	No
Providers	Yes	No
Human Interface Infrastructure (HII)	Yes	Yes
Preboot Utilities: **MegaRAID 6Gb/s SAS RAID Controller:** WebBIOS **MegaRAID 12Gb/s SAS RAID Controller:** Ctrl-R	Yes	Yes
StoreLib/StoreLib Test	Yes	Yes
StoreLib/StoreLib Test (OOB)	Yes	Yes
Legacy BIOS	Yes	Yes **NOTE** The Option ROM builds INT 13H for the boot VD, which is followed by INT 13H for the first 63 VDs reported in the VD list.

F.2 BIOS Known Limitations

The Legacy Option ROM displays only the first 64 VDs during the power-on self-test (POST). The following example describes the POST behavior when there are 90 VDs in the configuration.

Example:

- The Option ROM displays the first 64 VDs in the POST.
- 90 VDs are found on the host adapter.
- 64 VDs are handled by the BIOS.

Appendix G: Online Firmware Upgrade and Downgrade

This Appendix provides information about known issues when using the online firmware update feature of the MegaRAID 12Gb/s SAS RAID controller.

G.1 Online Firmware Upgrade Support

The following sections and table describe the known issues when using the OFU feature of the SAS-3 MegaRAID controller.

Known Limitations With Online Firmware Upgrade:

- For MR 6.7 Firmware GCA and later, any attempt to directly update the firmware to an older version using the online firmware update (OFU) process is not possible. The user must reboot the server for the older version to take effect. This is because of the product name rebranding effort that has resulted in changing the current VPD data to *AVAGO*, unlike the VPD data in the older firmware version (MR 6.6 Firmware GCA, and earlier), which is *LSI*. It is important that VPD data is presented the same to the operating system. Discrepancies in the VPD data results in an operating system crash since the operating system considers this critical data. Therefore, if any attempt to directly update the firmware to an older version using the online firmware update (OFU) process results in a change in VPD data (from *AVAGO* to *LSI*) and leads to an OS crash.
- MR 6.9 Firmware GCA supports 1 MB I/Os. The operating system driver presents this capability to the operating system during the initialization of the driver. However, the operating system driver cannot reinitialize the operating system with new values if there is an online firmware update (OFU) that does not support 1 MB I/Os. For example, OFU is not supported when you downgrade the firmware from MR 6.9 Firmware GCA to MR 6.8 Firmware GCA. Due to this operating system driver limitation, downgrading the firmware to an older version (for example, MR 6.8 Firmware GCA) using the OFU process is not possible when both the firmware and the driver have established 1 MB I/O support. However, firmware flash is allowed.
- If you are doing an online firmware update from a previous version to MR 6.9 Firmware GCA with large I/O support enabled, you need to reboot the system to enable large I/O support. Until you reboot the system, your operating system will be running with only those features that were available to it when it was initially booted.

Known Limitations With Reconstruction Operation

- From MR 6.6 Firmware GCA and later, you must back up the logical drive before initiating a reconstruction operation on the logical drive.
- You must not perform any firmware upgrade or downgrade when the reconstruction operation is in progress.
- When you flash a new firmware, you should not start a reconstruction operation until the system reboots or an Online Controller Reset (OCR) is performed.
- When more than 64 virtual drives are configured, downgrading the driver to an older version (for example, from MR 6.6 to MR 6.5) can cause the virtual drives with target IDs greater than 64 virtual drives to be masked to the host.
- When you upgrade from a non-240 virtual drive supported firmware (MR 6.5 and earlier) to a 240-virtual drive supported firmware (MR 6.6 and later), the auto-rebuild operation may not occur.
- When you downgrade from a 240-virtual drive supported firmware (MR 6.6 and later) to a non-240 virtual drive supported firmware (MR 6.5 and earlier), **Consistency Check**, **Background Initialization**, and **Secure Erase** operations might not resume.

NOTE The user must reboot the system for the flashed firmware to take effect.

Table 91 Online Firmware Upgrade and Downgrade Support Matrix

Release	OFU Downgrade Support	OFU Upgrade Support
MR 6.6 Firmware GCA and earlier	Yes (MR 6.6 and earlier)	Yes (MR 6.6 and later)
MR 6.7 Firmware GCA	No (MR 6.6 and earlier)	Yes (MR 6.7 and later)
MR 6.8 Firmware GCA	No (MR 6.7 and earlier)	Yes (MR 6.8 and later)
MR 6.9 Firmware GCA	No (MR 6.7 and earlier)	Yes (MR 6.8 and later)

Appendix H: Boot Messages and BIOS Error Messages

This appendix provides the boot messages and BIOS error messages present in the MegaRAID firmware.

H.1 Displaying Boot Messages

In platforms that load the UEFI driver first, the noncritical boot messages are discarded. To display a critical boot message, the platform should support driver health, and it should load the driver health formset when the Avago UEFI driver returns health status as `configuration required`.

In some systems, the platform supports the driver health protocol and calls the `GetHealthStatus` function automatically during boot time. In such platforms, if a critical boot problem exists, the platform shows a critical message dialog.

In some systems, you have to turn on the option in the system BIOS setup to enable the platform to call the `GetHealthStatus` function during boot time to check the health of the controller. To ensure that the platform supports driver health protocol and checks health during boot time, perform the following steps:

1. Set the controller's boot mode to SOE using CLI or RAID management/configuration application.
2. Connect one drive to the controller.
3. Create a RAID 0 volume.
4. Shut down the system, and remove the drive.
5. Boot the system.

   The following dialog should appear.

Figure 322 Driver Health Protocol Dialog

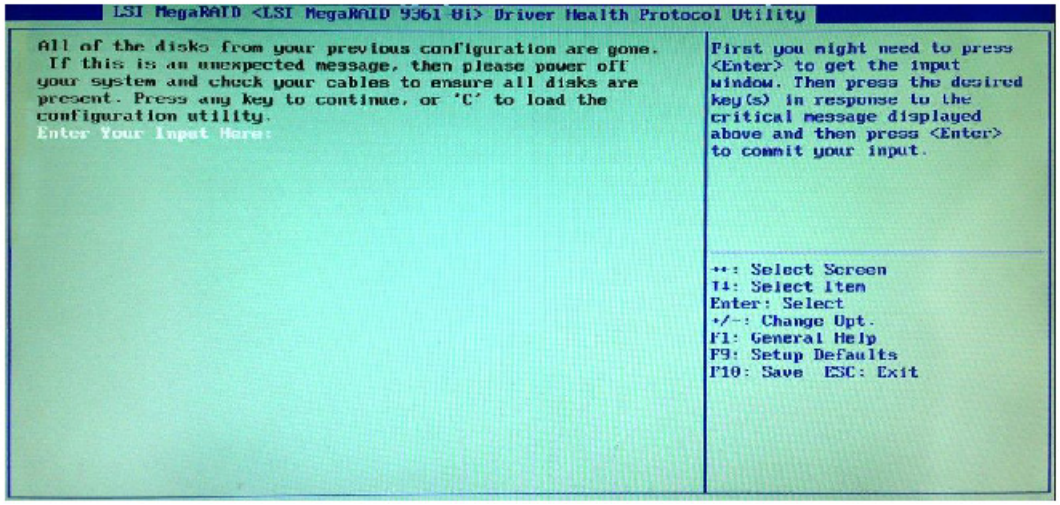

6. Press C.

   The following dialog appears.

Figure 323 Critical Message Completion Dialog

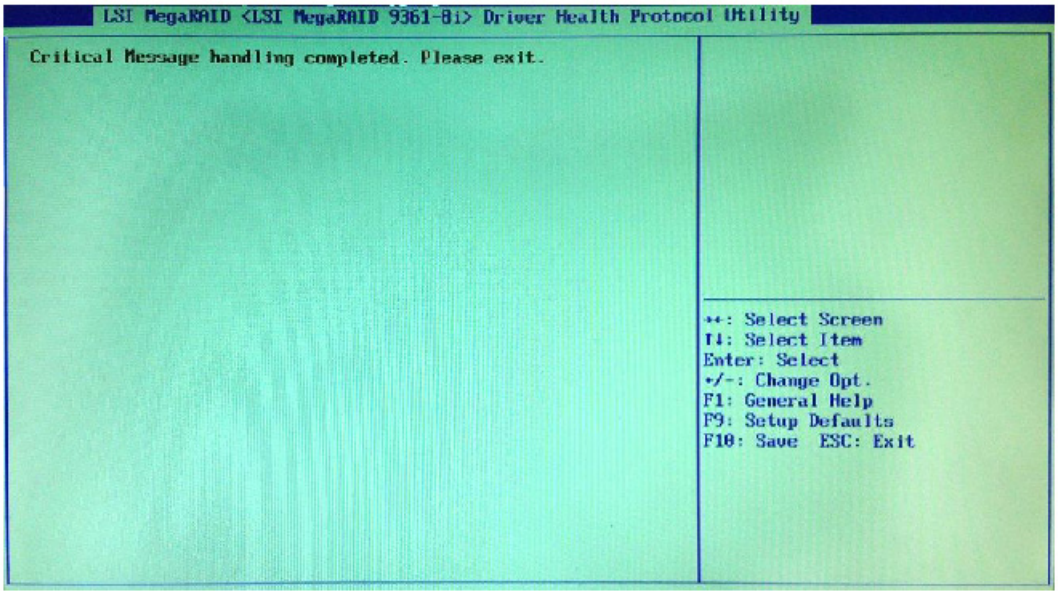

7. Press the Esc key to exit the browser.

H.2 Differences in the System Boot Mode

There is a behavioral differences in the controller boot mode (SOE, COE, HCOE, and HSM) and system boot mode (legacy or UEFI). Critical boot messages are reported through events for HSM. Both critical messages and warnings are reported in HCOE mode. The behavioral differences of system boot mode is because of the following:

- Some platforms might load both OpROMs (UEFI and legacy)
- Some platforms might load legacy first, and then the UEFI driver, or vice versa
- Some platforms might load only one OpROM depending upon the system boot mode (legacy versus UEFI)

On a hybrid system that loads the UEFI driver first, the noncritical boot messages are discarded and cannot be read if controller boot mode is set to SOE or COE. If the boot mode is set to HCOE or HSM, you can see the messages in the event log.

The following table describes the boot error messages present in the MegaRAID firmware.

- **Boot Message Type**: Name or type of the boot message on the firmware.
- **Wait Time**: A time value in seconds where the system waits for the user's input. If the wait time is elapsed, BIOS continues with default options.
  - For example, `BOOT_WAIT_TIME`, where the BIOS waits for the user's input for a default period of time (in seconds) and then continues with the default option if no user input is received.
  - For example, `BOOT_TIME_CRITICAL`, where the BIOS waits for the user's input until an input from the user is received.
- **Event Log**: When any event occurs, the firmware logs that particular event in its database.
- **Boot Message Description**: Boot message displayed on the console.
- **Comments**: Whether the message is associated with any specific controller settings or configuration settings related to the firmware.
- **Troubleshooting Actions**: If applicable, the user can take action to identify, diagnose, and resolve problems associated with the firmware. This can also be best practices, recommendations, and so on.

Table 92 Boot Messages

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
1	BOOT_MSG_CACHE_ DISCARD	BOOT_TIME_WAIT	MR_EVT_CTRL_CACHE _ DISCARDED	Memory or battery problems were detected. The adapter has recovered, but cached data was lost. Press any key to continue, or press C to load the configuration utility.	—	**Cause**: The cached data is lost and cannot be retrieved. **Action**: Perform memory and battery test. If needed, replace the memory card or the battery.
2	BOOT_MSG_TEST	5	Test boot message	This is a test message. You can press a key to ignore it, or you can wait five seconds. No further action is required. Press any key to continue, or press C to load the configuration utility.	—	N/A
3	BOOT_MSG_CACHE_V ERSION	BOOT_TIME_WAIT	MR_EVT_CTRL_CACHE _ VERSION_MISMATCH	Firmware version inconsistency was detected. The adapter has recovered, but cached data was lost. Press any key to continue, or press C to load the configuration utility.	—	**Causes**: The cached data is lost and cannot be retrieved. This boot message is displayed when dirty data needs to be flushed during boot. The version of the cache header with which dirty data was generated is different from the current version of the cache header. The version of the cache header is incremented when the cache layout is changed. On a single controller, during firmware upgrade, firmware ensures that there is no dirty data.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
						This message occurs only when dirty cache or pinned cache is migrated and is stored by ONFI from one controller to another controller where firmware versions on the both the controllers are different. **Action**: Ensure that the other controller also has the same firmware version.
4	BOOT_MSG_DDF_FOREIGN_ FOUND	10	MR_EVT_FOREIGN_CFG_ IMPORTED	Foreign configuration(s) found on adapter. Press any key to continue or press C to load the configuration utility or press F to import foreign configuration(s) and continue.	Use property autoEnhancedImport.	**Cause**: A storage device was inserted with the metadata that does not belong to any RAID volumes recognized by the controller. **Cause**: Either import the configuration settings of the inserted storage device or delete the RAID volume.
5	BOOT_MSG_DDF_IMPORT	10	NULL	Previous configuration cleared or missing. Importing configuration created on %02d/%02d %2d:%02d. Press any key to continue, or press C to load the configuration utility.	Not supported.	**Cause**: The controller is not able to recognize the current RAID volume configuration. **Action**: Either import the configuration settings or delete the foreign configuration found on storage device.
6	BOOT_MSG_PACKAGE_VERSION	0	MR_EVT_PACKAGE_ VERSION	Firmware package: %s	—	N/A
7	BOOT_MSG_FIRMWARE_ VERSION	0	NULL	Firmware version: %s	—	N/A

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
8	BOOT_MSG_FIRMWARE_TEST	1	NULL	This firmware is a TEST version. It has not completed any validation.	—	**Cause**: The controller is not able to recognize the current RAID volume configuration. **Action**: Update the firmware to the correct version.
9	BOOT_MSG_FIRMWARE_ALPHA	1	NULL	This firmware is an ALPHA version – It has not completed all validation. The validation stamp is: %s"""	—	**Cause**: The controller is not able to recognize the current RAID volume configuration. **Action**: Update the firmware to the correct version.
10	BOOT_MSG_FIRMWARE_BETA	1	NULL	This firmware is BETA version – It has not completed all validation. The validation stamp is: %s"""	—	**Cause**: The controller is not able to recognize the current RAID volume configuration. **Action**: Update the firmware to the correct version.
11	BOOT_MSG_SAS_SATA_MIXING_VIOLATION	BOOT_TIME_WAIT	MR_EVT_ENCL_SAS_SATA_MIXING_DETECTED	An enclosure was found that contains both SAS and SATA drives, but this controller does not allow mixed drive types in a single enclosure. Correct the problem then restart your system. Press any key to continue, or press C to load the configuration utility.	—	**Cause**: A single enclosure that has both SAS and SATA drives cannot be used as the controller does not support mixed drive types in a single enclosure. **Actions**: Use only one type of drive, either SAS or SATA drive. Replace the controller with a controller that supports mixed drive types in a single enclosure. Contact Technical Support to enable this feature.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
12	BOOT_MSG_SAS_NOT_ SUPPORTED	BOOT_TIME_WAIT	SAS drives are not supported.	SAS drives were detected, but this controller does not support SAS drives. Remove the SAS drives then restart your system. Press any key to continue, or press C to load the configuration utility.	—	**Cause**: This controller does not support SAS drives. **Action**: Replace the SAS drives with SATA drives and restart the system.
13	BOOT_MSG_SATA_NOT_ SUPPORTED	BOOT_TIME_WAIT	SATA drives are not supported.	SATA drives were detected, but this controller does not support SATA drives. Remove the SATA drives then restart your system. Press any key to continue, or press C to load the configuration utility.	—	**Cause**: This controller does not support SATA drives. **Action**: Replace the SATA drives with SAS drives and restart the system.
14	BOOT_MSG_ENCL_COUNT_ PER_PORT_EXCEEDED	BOOT_TIME_WAIT	MR_EVT_ENCL_MAX_ PER_PORT_EXCEEDED	There are %d enclosures connected to connector %s, but only maximum of %d enclosures can be connected to a single SAS connector. Remove the extra enclosures then restart your system.	—	**Cause**: This controller supports only a particular number of enclosures. **Action**: Remove extra enclosures or insert a controller that supports your enclosure requirements.
15	BOOT_MSG_SAS_TOPOLOGY_ ERROR	BOOT_TIME_WAIT	SAS discovery error	Invalid SAS topology detected. Check your cable configurations, repair the problem, and restart your system.	—	**Cause**: The controller has detected an invalid SAS topology. **Action**: Check the cables or reconfigure the attached devices to create a valid SAS topology.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
16	BOOT_MSG_BBU_BAD	10	NULL	The battery is currently discharged or disconnected. Verify the connection and allow 30 minutes for charging. If the battery is properly connected and it has not returned to operational state after 30 minutes of charging then contact technical support for additional assistance.	Not supported.	**Actions**: Check the battery cable to ensure that it is connected properly. Ensure that the battery is charging properly. Contact Technical Support to replace the battery if the battery is draining out.
17	BOOT_MSG_BBU_MSG_DISABLE	10	MR_EVT_BBU_NOT_PRESENT	The battery hardware is missing or malfunctioning, or the battery is unconnected, or the battery could be fully discharged. If you continue to boot the system, the battery-backed cache will not function. If battery is connected and has been allowed to charge for 30 minutes and this message continues to appear, contact technical support for assistance. Press D to disable this warning (if your controller does not have a battery)	Use property disableBatteryWarning	**Action**: Check the battery cable to ensure that it is connected properly. Ensure that the battery is charging properly. Contact Technical Support to replace the battery if the battery is draining out.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
18	BOOT_MSG_BAD_MFC_ SASADDRESS	10	MFC data error! Invalid SAS address	Invalid SAS Address present in MFC data. Program a valid SAS Address and restart your system.	—	**Cause**: Invalid SAS address may be present. **Actions**: 1. Power off the system and remove the controller. 2. Find the SAS address label and re-program the SAS address. Contact Technical Support if you are unable to re-program the SAS address. OEMs can access the StorCLI and re-program the SAS address.
19	BOOT_MSG_PDS_MISSING	BOOT_TIME_WAIT	MR_EVT_CTRL_BOOT_MISSING_PDS	Some configured disks have been removed from your system, or are no longer accessible. Check your cables and also make sure all disks are present. Press any key to continue, or press C to load the configuration utility.	—	**Cause**: The controller is unable to find the configured drives. **Actions**: Check if the configured drives are present and they are properly connected. Go to BIOS and check if the devices are displayed. Ensure that the drives are spun-up and have power supplied to them. If there is a backplane, check the connector to ensure that power is being supplied to the drive.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
20	BOOT_MSG_LDS_OFFLINE	BOOT_TIME_WAIT	MR_EVT_CTRL_BOOT_LDS_WILL_GO_OFFLINE	The following VDs have missing disks: %s. If you proceed (or load the configuration utility), these VDs will be marked OFFLINE and will be inaccessible. Check your cables and make sure all disks are present. Press any key to continue, or press C to load the configuration utility.	—	**Cause**: The controller is unable to find the configured drives. **Actions**: Check if the configured drives are present and they are properly connected. Go to BIOS and check if the devices are displayed. Ensure that the drives are spun-up and have power supplied to them. If there is a backplane, check the connector to ensure that power is being supplied to the drive.
21	BOOT_MSG_LDS_MISSING	BOOT_TIME_WAIT	MR_EVT_CTRL_BOOT_LDS_MISSING	The following VDs are missing: %s. If you proceed (or load the configuration utility), these VDs will be removed from your configuration. If you wish to use them at a later time, they will have to be imported. If you believe these VDs should be present, power off your system and check your cables to make sure all disks are present. Press any key to continue, or press C to load the configuration utility.	—	**Cause**: The controller is unable to find the configured drives. **Actions**: Check if the configured drives are present and they are properly connected. Go to BIOS and check if the devices are displayed. Ensure that the drives are spun-up and have power supplied to them. If there is a backplane, check the connector to ensure that power is being supplied to the drive.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
22	BOOT_MSG_LDS_MISSING_ SPANS	BOOT_TIME_WAIT	MR_EVT_CTRL_BOOT_LDS_MISSING	The following VDs are missing complete spans: %s. If you proceed (or load the configuration utility), these VDs will be removed from your configuration and the remaining drives marked as foreign. If you wish to use them at a later time, restore the missing span(s) and use a foreign import to recover the VDs. If you believe these VDs should be present, please power off your system and check your cables to make sure all disks are present. Press any key to continue, or press C to load the configuration utility.	—	**Cause**: The controller is unable to find the configured drives. **Actions**: Check if the configured drives are present and they are properly connected. Go to BIOS and check if the devices are displayed. Ensure that the drives are spun-up and have power supplied to them. If there is a backplane, check the connector to ensure that power is being supplied to the drive.
23	BOOT_MSG_CONFIG_MISSING	BOOT_TIME_WAIT	MR_EVT_CTRL_BOOT_CONFIG_MISSING	All of the disks from your previous configuration are gone. If this is an unexpected message, power off your system and check your cables to make sure all disks are present. Press any key to continue, or press C to load the configuration utility.	Headless mode – should not appear, if autoEnhancedImport is set.	**Cause**: The controller is unable to find the configured drives. **Actions**: Check if the configured drives are present and they are properly connected. Go to BIOS and check if the devices are displayed. Ensure that the drives are spun-up and have power supplied to them. If there is a backplane, check the connector to ensure that power is being supplied to the drive.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
24	BOOT_MSG_CACHE_FLUSH_ NOT_POSSIBLE	BOOT_TIME_CRITICAL	NULL	The cache contains dirty data, but some VDs are missing or will go offline, so the cached data can not be written to disk. If this is an unexpected error, power off your system and check your cables to make sure all disks are present. If you continue, the data in cache will be permanently discarded. Press X to acknowledge and permanently destroy the cached data.	Not supported	**Cause**: The controller is unable to find the configured drives. **Actions**: Check if the configured drives are present and they are properly connected. Go to BIOS and check if the devices are displayed. Ensure that the drives are spun-up and have power supplied to them. If there is a backplane, check the connector to ensure that power is being supplied to the drive.
25	BOOT_MSG_LDS_WILL_RUN_ WRITE_THRU	5	NULL	Your VDs that are configured for Write-Back are temporarily running in Write-Through mode. This is caused by the battery being charged, missing, or bad. Allow the battery to charge for 24 hours before evaluating the battery for replacement. The following VDs are affected: %s Press any key to continue.	No event is logged, information for the user	**Actions**: Check the battery cable to ensure that it is connected properly. Ensure that the battery is charging properly. Contact Technical Support to replace the battery if the current supplied by the battery is draining out.
26	BOOT_MSG_MEMORY_INVALID	BOOT_TIME_CRITICAL	NULL	Invalid memory configuration detected. Contact your system support. System has halted.	Not supported	**Action**: Reseat or replace the DIMM.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
27	BOOT_MSG_CACHE_ DISCARD_WARNING	BOOT_TIME_WAIT	MR_EVT_CTRL_CACHE _ DISCARDED	Cache data was lost due to an unexpected power-off or reboot during a write operation, but the adapter has recovered. This could be because of memory problems, bad battery, or you might not have a battery installed. Press any key to continue or C to load the configuration utility.	Posted only when disableBatteryWarning is set, same as BOOT_MSG _CACHE_DISCARD	**Actions**: Check the battery cable to ensure that it is connected properly. Ensure that the battery is charging properly. Contact Technical Support to replace the battery if power supplied by the battery is draining out.
28	BOOT_MSG_CONFIG_ CHANGE_WARNING	BOOT_TIME_CRITI CAL	NULL	Entering the configuration utility in this state will result in drive configuration changes. Press Y to continue loading the configuration utility or power off your system and check your cables to make sure all disks are present and reboot the system.	Posted from other messages like BOOT_MSG _LDS_MISSING, when the user clicks C.	**Cause**: The controller is unable to find the configured drives. **Actions**: Check if the configured drives are present and they are properly connected. Go to BIOS and check if the devices are displayed. Ensure that the drives are spun-up and have power supplied to them. If there is a backplane, check the connector to ensure that power is being supplied to the drive. If the controller is being used to create a new configuration by reusing the drives, purge the existing data and then continue.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
29	BOOT_MSG_EMBEDDED_ MULTIBIT_ECC_ERROR	BOOT_TIME_CRITICAL	Multibit ECC error - memory or controller needs replacement.	Multibit ECC errors were detected on the RAID controller. If you continue, data corruption can occur. Contact technical support to resolve this issue. Press X to continue, otherwise power off the system, replace the controller, and reboot.	OEM Specific, see BOOT_MSG_HBA_MULTIBIT_ECC_ERROR for Avago Generic message	**Action**: 1. Reseat or replace the DIMM. 2. Restart system. If the problem persists, contact Technical Support.
30	BOOT_MSG_EMBEDDED_ SINGLE_BIT_ECC_ERROR	BOOT_TIME_CRITICAL	MR_EVT_CTRL_MEM_ECC_SINGLE_BIT_CRITICAL or WARNING	Single-bit ECC errors were detected on the RAID controller. Contact technical support to resolve this issue. Press X to continue or else power off the system, replace the controller, and reboot.	OEM Specific, see BOOT_MSG_HBA_SINGLE_BIT_ECC_ERROR for Avago Generic message	**Action**: 1. Reseat or replace the DIMM. 2. Restart system. If the problem persists, contact Technical Support.
31	BOOT_MSG_EMBEDDED_ SINGLE_BIT_OVERFLOW_ECC_ ERROR	BOOT_TIME_CRITICAL	NULL	Single-bit overflow ECC errors were detected on the RAID controller. If you continue, data corruption can occur. Contact technical support to resolve this issue. Press X to continue or else power off the system, replace the controller, and reboot.	Not supported	**Action**: 1. Reseat or replace the DIMM. 2. Restart system. If the problem persists, contact Technical Support.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
32	BOOT_MSG_HBA_MULTIBIT_ ECC_ERROR	BOOT_TIME_CRITICAL	Multibit ECC error – memory or controller needs replacement.	Multibit ECC errors were detected on the RAID controller. The DIMM on the controller needs replacement. Contact technical support to resolve this issue. If you continue, data corruption can occur. Press X to continue, otherwise power off the system and replace the DIMM module and reboot. If you have replaced the DIMM press X to continue.	—	**Action**: 1. Reseat or replace the DIMM. 2. Restart system. If the problem persists, contact Technical Support.
33	BOOT_MSG_HBA_SINGLE_BIT_ ECC_ERROR	BOOT_TIME_CRITICAL	MR_EVT_CTRL_MEM_ECC_SINGLE_BIT_CRITICAL or WARNING	Single-bit ECC errors were detected during the previous boot of the RAID controller. The DIMM on the controller needs replacement. Contact technical support to resolve this issue. Press X to continue, otherwise power off the system and replace the DIMM module and reboot. If you have replaced the DIMM press X to continue.	—	**Action**: 1. Reseat or replace the DIMM. 2. Restart system. If the problem persists, contact Technical Support.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
34	BOOT_MSG_HBA_SINGLE_BIT_OVERFLOW_ECC_ERROR	BOOT_TIME_CRITICAL	NULL	Single-bit overflow ECC errors were detected during the previous boot of the RAID controller. The DIMM on the controller needs replacement. Contact technical support to resolve this issue. If you continue, data corruption can occur. Press X to continue, otherwise power off the system and replace the DIMM module and reboot. If you have replaced the DIMM press X to continue.	Not supported	**Action**: 1. Reseat or replace the DIMM. 2. Restart system. If the problem persists, contact Technical Support.
35	BOOT_MSG_ENCL_VIOLATION_MODE	BOOT_TIME_CRITICAL	MR_EVT_CTRL_CRASH	The attached enclosure does not support in controller's Direct mapping mode. Contact your system support. The system has halted because of an unsupported configuration.	Should be able to enter HSM	**Causes**: Too many chained enclosures may be present. May also be related to a security feature in the drive. **Actions**: Remove the drives that are not supported. Reduce the number of drives. Replace the enclosure with an other one. Ensure that the firmware version is updated. Contact Technical Support if the problem persists.
36	BOOT_MSG_EXP_VIOLATION_FORCE_REBOOT	10	MR_EVT_CTRL_CRASH	Expander detected in controller with direct mapping mode. Reconfiguring automatically to persistent mapping mode. Automatic reboot would happen in 10 seconds.	OEM Specific action, see BOOT_MSG_ENCL_VIOLATION_MODE for LSI generic	**Action**: No action required. The controller will configure itself to a persistent mapping mode and then reboot. Contact Technical Support if problem persists.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
37	BOOT_MSG_8033X_ATU_ISSUE	BOOT_TIME_CRITICAL	NULL	Your controller's I/O processor has a fault that can potentially cause data corruption. Your controller needs replacement. Contact your system support. To continue, press Y to acknowledge.	DEPRECATED	**Action**: Contact Technical Support for replacement of the controller.
38	BOOT_MSG_MAX_DISKS_ EXCEEDED	BOOT_TIME_CRITICAL	MR_EVT_PD_NOT_SUPPORTED	The number of disks exceeded the maximum supported count of %d disks. Remove the extra drives and reboot system to avoid losing data. Press Y to continue with extra drives.	—	**Actions**: Power off the system and remove the controller. Remove the extra drives to reduce the size of the topology. Replace the controller with a controller that supports a larger topology.
39	BOOT_MSG_MAX_DISKS_ EXCEEDED_PER_QUAD	BOOT_TIME_CRITICAL	NULL	The number of devices exceeded the maximum limit of devices per quad. Remove the extra drives and reboot the system to avoid losing data System has halted due to unsupported configuration.	Not supported	**Actions**: Power off the system and remove the controller. Remove the extra drives to reduce the size of the topology. Replace the controller with a controller that supports a larger topology.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
40	BOOT_MSG_DISCOVERY_ERROR	BOOT_TIME_CRITICAL	Discovery errors – power cycle system and drives, and try again.	A discovery error has occurred, power cycle the system and all the enclosures attached to this system.	—	**Actions**: Shutdown and restart the system as well as all the enclosures attached to the system. Ensure that all the cables are connected and connected properly. Reduce the topology in case of a bad drive. If the problem persists, collect the logs of the system, driver, and firmware and contact Technical Support.
41	BOOT_MSG_CTRL_SECRET_KEY_FIRST	BOOT_TIME_WAIT	NULL	Drive security is enabled on this controller and a pass phrase is required. Enter the pass phrase.	Requires user input, if undesired, change Security binding	**Action**: Enter the pass phrase.
42	BOOT_MSG_CTRL_SECRET_KEY_RETRY	BOOT_TIME_WAIT	NULL	Invalid pass phrase. Enter the pass phrase.	opRom must be enabled for user input, if undesired, change Security binding	**Action**: Enter the pass phrase.
43	BOOT_MSG_CTRL_LOCK_KEY_ INVALID	BOOT_TIME_WAIT	MR_EVT_CTRL_LOCK_KEY_FAILED	There was a drive security key error. All secure drives will be marked as foreign. Press any key to continue, or C to load the configuration utility.	—	**Action**: Check if the controller supports self-encrypting drives.
44	BOOT_MSG_KEY_MISSING_ REBOOT_OR_CONTINUE	BOOT_TIME_WAIT	MR_EVT_CTRL_LOCK_KEY_FAILED	Invalid pass phrase. If you continue, a drive security key error will occur and all secure configurations will be marked as foreign. Reboot the machine to retry the pass phrase or press any key to continue.	—	**Action**: Restart the system to retry the pass phrase or press any key to continue.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
45	BOOT_MSG_KEY_EKMS_ FAILURE	BOOT_TIME_WAIT	MR_EVT_CTRL_LOCK_KEY_EKM_FAILURE	Unable to communicate to EKMS. If you continue, there will be a drive security key error and all secure configurations will be marked as foreign. Check the connection with the EKMS, reboot the machine to retry the EKMS or press any key to continue.	—	**Action**: Check the connection of EKMS, restart the system to re-establish the connection to EKMS.
46	BOOT_MSG_REKEY_TO_EKMS_ FAILURE	BOOT_TIME_WAIT	MR_EVT_CTRL_LOCK_KEY_REKEY_FAILED	Unable to change security to EKMS as not able to communicate to EKMS. If you continue, the drive security will remain to existing security mode. Check the connection with the EKMS, reboot the machine to retry the EKMS or press any key to continue.	—	**Action**: Check the connection of EKMS, restart the system to re-establish the connection to EKMS.
47	BOOT_MSG_KEY_EKMS_ FAILURE_MERCURY	20	MR_EVT_CTRL_LOCK_KEY_EKM_FAILURE	DKM existing key request failed; existing secure configurations will be labeled foreign and will not be accessible. Reboot the server to retry.	OEM Specific, see BOOT_MSG_KEY_EKMS_FAILURE for Avago generic	**Action**: Check the connection of EKMS, restart the system to re-establish the connection to EKMS.
48	BOOT_MSG_REKEY_TO_EKMS_ FAILURE_MERCURY	BOOT_TIME_CRITICAL	MR_EVT_CTRL_LOCK_KEY_REKEY_FAILED	DKM new key request failed; controller security mode transition was not successful. Reboot the server to retry request, or press any key to continue.	OEM Specific, see BOOT_MSG_REKEY_TO_EKMS_FAILURE for Avago generic	**Action**: Check the connection of EKMS, restart the system to re-establish the connection to EKMS.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
49	BOOT_MSG_NVDATA_ IMAGE_ MISSING	BOOT_TIME_WAIT	NVDATA image is invalid – reflash NVDATA image	Firmware did not find valid NVDATA image. Program a valid NVDATA image and restart your system. Press any key to continue.	—	**Actions**: Flash the correct firmware package that has proper NV Data image. Check the current firmware version, and if needed, updated to the latest firmware version. Updating to the latest firmware version may require importing foreign volumes.
50	BOOT_MSG_IR_MR_M IGRATION_FAILED	BOOT_TIME_WAIT	IR to MR migration failed.	IR to MR Migration failed. Press any key to continue with MR defined NVDATA values	—	N/A
51	BOOT_MSG_DUAL_BA T_PRSNT	10	NULL	Two BBUs are connected to the adapter. This is not a supported configuration. Battery and caching operations are disabled. Remove one BBU and reboot to restore battery and caching operations. If dirty cache is lost in this boot, that could have been because of dual battery presence.	Not supported	**Actions**: Remove one BBU and restart the system to restore battery and caching operations. Due to the presence of a dual battery, you may lose the data in dirty cache while restarting the system.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
52	BOOT_MSG_LDS_CACHE_ PINNED	10	MR_EVT_CTRL_BOOT_LDS_CACHE_PINNED	Offline or missing virtual drives with preserved cache exist. Check the cables and make sure that all drives are present. Press any key to continue, or C to load the configuration utility.	Use property allowBootWithPinnedCache	**Cause**: The controller is unable to find the configured drives. **Actions**: Check if the configured drives are present and they are properly connected. Go to BIOS and check if the devices are displayed. Ensure that the drives are spun-up and have power supplied to them. If there is a backplane, check the connector to ensure that power is being supplied to the drive. Cache offload occurs if the missing drive is restored.
53	BOOT_MSG_LDS_CACHE_ PINNED_HALT	BOOT_TIME_CRITICAL	MR_EVT_CTRL_BOOT_LDS_CACHE_PINNED	Offline or missing virtual drives with preserved cache exist. Check the cables and make sure that all drives are present. Press any key to enter the configuration utility.	If property allowBootWithPinnedCache is disabled	**Cause**: The controller is unable to find the configured drives. **Actions**: Check if the configured drives are present and they are properly connected. Go to BIOS and check if the devices are displayed.Ensure that the drives are spun-up and have power supplied to them. If there is a backplane, check the connector to ensure that power is being supplied to the drive. Cache offload occurs if the missing drive is restored.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
54	BOOT_MSG_BAD_SBR_ SASADDRESS	BOOT_TIME_CRITICAL	NULL	Invalid SAS Address present in SBR. Contact your system support. Press any key to continue with Default SAS Address.	Not supported	**Cause**: Invalid SAS address present in the SBR. **Action**: Contact Technical Support to restore to the factory default values.
55	BOOT_MSG_INCOMPATIBLE_ SECONDARY_IBUTTON	BOOT_TIME_CRITICAL	Incompatible secondary iButton detected	Incompatible secondary iButton present! Insert the correct iButton and restart the system. Press any key to continue but OEM specific features will not be upgraded!	—	**Actions**: Insert the correct iButton or key-vault and restart the system. If problem persists, contact Technical Support for replacement of the iButton or key-vault.
56	BOOT_MSG_CTRL_ DOWNGRADE_DETECTED	BOOT_TIME_CRITICAL	NULL	Upgrade Key Missing! An upgrade key was present on a previous power cycle, but it is not connected. This can result in inaccessible data unless it is addressed. Re-attach the upgrade key and reboot.	Not supported	**Cause**: An upgrade key that was present on a previous power cycle may not be connected. **Actions**: Reattach the upgrade key and restart the system. If the problem persists, contact Technical Support for replacement of the upgrade key.
57	BOOT_MSG_DDF_MFC_ INCOMPATIBLE	BOOT_TIME_WAIT	Native configuration is not supported, check MFC.	The native configuration is not supported by the controller. Check the controller, iButton or key-vault. If you continue the configuration will be marked foreign. Press any key to continue.	—	**Actions**: Insert the correct iButton or key-vault and restart the system. If problem persists, contact Technical Support for replacement of the iButton or key-vault.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
58	BOOT_MSG_BBU_MSG_DISABLE_PERC	10	MR_EVT_BBU_NOT_PRESENT or REMOVED	The battery is currently discharged or disconnected. Verify the connection and allow 30 minutes for charging. If the battery is properly connected and it has not returned to operational state after 30 minutes of charging, contact technical support for additional assistance. Press D to disable this warning (if your controller does not have a battery).	Use property disableBatteryWarning, OEM Specific, also see BOOT_MSG_BBU_MSG_DISABLE	**Actions**: Check the battery cable to ensure that it is connected properly. Ensure that the battery is charging properly. Contact Technical Support to replace the battery if power supplied by the battery is draining out.
59	BOOT_MSG_LDS_WILL_RUN_WRITE_THRU_PERC	5	NULL	The battery is currently discharged or disconnected. VDs configured in Write-Back mode will run in Write-Through mode to protect your data and will return to the Write-Back policy when the battery is operational. If VDs have not returned to Write-Back mode after 30 minutes of charging then contact technical support for additional assistance. The following VDs are affected: %s. Press any key to continue.	No event is logged, information for the user	**Actions**: Check the battery cable to ensure that it is connected properly. Ensure that the battery is charging properly. Contact Technical Support to replace the battery if the battery is draining out.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
60	BOOT_MSG_CACHE_DISCARD_WARNING_PERC	BOOT_TIME_WAIT	MR_EVT_CTRL_CACHE_ DISCARDED	Cache data was lost, but the controller has recovered. This could be because your controller had protected cache after an unexpected power loss and your system was without power longer than the battery backup time. Press any key to continue or C to load the configuration utility.	Property disableBatteryWarning is set	**Actions**: Check the memory and the battery. Check the voltage levels and cache offload timing in case of power loss. If necessary, replace the memory or battery.
61	BOOT_MSG_ROLLBACK_ACTIVE	BOOT_TIME_CRITICAL	NULL	A snapshot rollback is in progress on VDs %s, the controller cannot boot until the rollback operation completes. Press any key to enter the configuration utility.	opRom must be enabled, if undesired, do not request rollback. Not supported in 12G	**Actions**: Wait for some time until the rollback is complete.
62	BOOT_MSG_ROLLBACK_ACTIVE_REPOSITORY_MISSING	BOOT_TIME_CRITICAL	Rollback requested, but repository is missing	The following VDs: %s have Rollback active and the corresponding Repository is missing. If you continue to boot the system or enter the configuration utility, these VDs will become unusable. Press any key to Continue.	Not supported in 12G	**Cause**: This may be related to the snapshot feature, which is not supported on 12G controllers. **Action**: Wait for some time until the rollback is complete.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
63	BOOT_MSG_REPOSITORY_ MISSING	BOOT_TIME_WAIT	Snapshot repository is missing, snapshot disabled	Snapshot Repository VDs %s have been removed from your system, or are no longer accessible. Check the cables and make sure all disks are present. If you continue to boot the system, the snapshot related data will be lost. Press any key to continue, or C to load the configuration utility.	Not supported in 12G	**Cause**: The controller is unable to find the configured drives. **Actions**: Check if the configured drives are present and they are properly connected. Go to BIOS and check if the devices are displayed. Ensure that the drives are spun-up and have power supplied to them. If there is a backplane, check the connector to ensure that power is being supplied to the drive.
64	BOOT_MSG_CFG_CMD_LOST	BOOT_TIME_WAIT	MR_EVT_CFG_CMD_LOST	The most recent configuration command could not be committed and must be retried. Press any key to continue, or C to load the configuration utility.	—	N/A
65	BOOT_MSG_CFG_CHANGES_ LOST	10	Configuration command was not committed, please retry	Firmware could not synchronize the configuration or property changes for some of the VD's/PD's. Press any key to continue, or C to load the configuration utility.	—	**Actions**: If the same problem persists, contact Technical Support.
66	BOOT_MSG_CFG_ONBOARD_ EXP_NOT_DETECTED	BOOT_TIME_CRITICAL	On-board expander FW or mfg image is corrupted – reflash image	On-board expander firmware or manufacturing image is corrupted. The flash expander firmware and manufacturing image use the recovery tools.	—	**Actions**: Contact Technical Support for factory-only tools to assist in recovery of the expander.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
67	BOOT_MSG_PFK_INCOMPATIBLE	BOOT_TIME_WAIT	MFC record not found, ensure you have the correct FW version	The native configuration is not supported by the current firmware. Make sure that the correct controller firmware is being used. If you continue, the configuration will be marked as foreign. Press any key to continue.	—	**Actions**: Collect the logs of the system, driver, and firmware. Ensure that the firmware version corrected and is updated to the latest version. Contact Technical Support if the problem persists.
68	BOOT_MSG_INVALID_FOREIGN_CFG_IMPORT	5	MR_EVT_FOREIGN_CFG_AUTO_IMPORT_NONE	Foreign configuration import did not import any drives. Press any key to continue.	—	**Actions**: Check the firmware version of the controller. Replace the controller and try again. If the problem persists, contact Technical Support.
69	BOOT_MSG_UPGRADED_IMR_ TO_MR	2	Reboot required to complete the iMR to MR upgrade	Valid memory detected. Firmware is upgraded from iMR to MR. Reboot the system for the MR firmware to run.	—	N/A
70	BOOT_MSG_PFK_ENABLED_AT_BOOT_TIME	BOOT_TIME_WAIT	BOOT_MSG_EVENT_USE_ BOOT_MSG	Advanced software options keys were detected, features activated – %s.	—	N/A
71	BOOT_MSG_PFK_DISABLED_AT_BOOT_TIME	BOOT_TIME_WAIT	BOOT_MSG_EVENT_USE_ BOOT_MSG	Advanced software options keys were missing, features deactivated – %s.	—	**Actions**: Check the cable connection. Check for the Advanced Software Options key. If the problem persists, contact Technical Support.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
72	BOOT_MSG_EEPROM _ERROR_ FEATURES_DISABLED	BOOT_TIME_CRITICAL	Cannot communicate with iButton, possible extreme temps.	Cannot communicate with iButton to retrieve premium features. This is probably because of extreme temperatures. The system has halted!	—	**Actions**: Check the cable connection. Ensure that iButton is present. Check the ambient temperature near the iButton. If the problem persists, contact Technical Support.
73	BOOT_MSG_DC_ON_ DEGRADED_LD	BOOT_TIME_CRITICAL	Multiple power loss detected with I/O transactions to non optimal VDs.	Consecutive power loss detected during I/O transactions on non-optimal write-back volumes. This might have resulted in data integrity issues. Press 'X' to proceed.	—	**Actions**: Check if the controller is securely locked in the PCI slot. Check the power supply, battery, and Supercap. If you find any hardware defect, contact Technical Support.
74	BOOT_MSG_USB_DEVICE_ ERROR	BOOT_TIME_CRITICAL	USB cache device is not responding.	USB cache device is not responding. Power down the system for 2 minutes to attempt recovery and avoid cache data loss, and then power-on.	Not supported in 12G.	This message is not applicable to 12G because the 3108 controller supports ONFI-based cache offload. **Actions**: The 2208 controller supports USB cache offload. Ensure that USB cache is present and secure. Reseat and replace the USB cache. Power off the system for 2 minutes to attempt recovery and avoid cache data loss, then power on the system.

Table 92 Boot Messages (Continued)

No	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
75	BOOT_MSG_DOWNGRADE_ MR_TO_IMR	BOOT_TIME_CRITICAL	Bad or missing RAID controller memory module detected.	Bad or missing RAID controller memory module detected. Press D to downgrade the RAID controller to iMR mode. **Warning!** Downgrading to iMR mode, might result in incompatible Logical drives. Press any other key to continue, controller shall boot to safe mode.	—	**Actions**: 1. Reseat or replace the DIMM. 2. Restart system. If the problem persists, contact Technical Support for repair or replacement.
76	BOOT_MSG_HEADLESS_DUMMY	0	NULL	—	—	N/A
77	BOOT_MSG_LIST_TERMINATOR	0	NULL	—	—	N/A

Appendix I: Transport Ready Pre-requisites

These are some of the important information that you must consider before you begin the transport ready operation.

Virtual Drive Group:

- Any virtual drive group is marked as Transport Ready only when all the virtual drives in that group are in optimal state. During the transport ready process, if you try to remove or insert a transport ready physical device, the virtual drive remains on-line and does not change it's state.
- The maximum number of drive groups supported by Transport Ready is 16.
- When you set the Transport Ready policy on a drive group, all the VD Erase operations running in that drive group are aborted.

Background Operations:

- When you move a transport ready virtual drive to a different system whose firmware supports the transport ready feature, the background operations such as Consistency Check and Background initialization which were in paused state, resume their function.

The following table provides the status of the Background Operations when you apply the Transport Ready policy on a virtual drive group:

Table 93 Background Operations

-	Reboot	FC Import	TR Set	TR Clear	Reboot with TR	TR Clear after Reboot	FC Import with TR on the same or other system
Recon	resume	NA	NA	NA	NA	NA	NA
Rebuild	resume	restart	NA (Option grayed out) VD is not in optimal state.	NA (Option grayed out)	NA (Option grayed out)	NA (Option grayed out)	NA
Copyback	resume	restart	suspended	resume	none	resume	restart
FGI	aborted	none	aborted	none	none	none	none
VD Erase	resume	none	aborted	none	none	none	none
BGI	resume	restart/ NOTE (Unplug and replug the VDs in the same system)	suspended	resume	none	resume	resume/ NOTE
CC	resume	restart	suspended	resume	none	resume	resume

NOTE When BGI condition (RAID 5 has more than 4 drives and RAID 6 has more than 6 drives) is met, BGI can start on foreign configuration import even though BGI was not running on the previous system.

Where:

- TR - Transport Ready
- NA - Operation Not Available
- FC - Foreign Configuration
- FGI - Foreground Initialization
  - FGI- When TR is set on the VD, FGI operation has to be ABORTED.

  - FGI - Since FGI has to be aborted after the VD is set in TR, therefore, Clearing TR will remain in NONE state.
- BGI - Background Initialization
- VD - Virtual Disk
- CC - Consistency Check
  - CC - Import FC after CC operation starts, this will ensure the CC operation RESTARTS from zero percent.
- none - nothing to do. TR VDs have spun down physical drives. I/O activity is not allowed on them. So "nothing to do" means background operations will retain their state when TR was set. If a background operation is aborted at TR Set, it will remain aborted or if it is suspended at TR set, it will remain suspended. If an operation is suspended on TR Set, it will resume on TR Clear. If an operation is aborted on TR Set, it will not start on TR Clear, the user has to manually start it.

NOTE When all the physical drives of a TR drive group are unplugged, they will become foreign configuration. If these drives are inserted back on the same system or any other system, behavior should be same as "Foreign Configuration import with TR on the same or other system."

Firmware Upgrade:

- The Transport Ready policy will be persistent across reboots, shutdowns, and firmware upgrades which have Transport Ready enabled.

Firmware Downgrade

- If the firmware is downgraded to legacy firmware or Transport Ready is disabled in the new firmware, Transport Ready will be cleared in NVRAM. However, if the firmware is not a legacy firmware or it does not have Transport Ready feature, then Transport Ready will not be cleared. In this case, if the Transport Ready aware firmware is flashed again, Transport Ready drive groups will reappear; you have to manually clear the Transport Ready.
- On a system that supports Transport Ready, if you downgrade the firmware to a version that does not support Transport Ready, then the background operations like Consistency check and Background Initialization will be suspended and the user has to manually resume them. The background operations resume their function from where they had stopped.
- When the Transport ready virtual drives are moved to a different system which does not support the Transport Ready feature, then the background operations like Consistency check and Background Initialization will start from 0.

SSD Drives:

- SSD drives respond immediately to SSU (Start Stop unit) command because there is no spin down needed for SSD drives. SSD drives will move to power save mode immediately. This behavior is different from HDD drives.

Transport Ready is not enabled in the following scenarios:

- Reconstruction operation is in progress on any virtual drive of the drive group.
- Pinned Cache present on any virtual drive of the drive group
- Any virtual drive of the drive group is marked as Cacheable or associated with Cachecade virtual drive.
- Any virtual drive of the drive group is a Cachecade virtual drive.
- Any virtual drive of the drive group is not on-line.
- Any virtual drive of the drive group is a boot virtual drive.

Supported Operations on transport ready virtual drives

The following operations can be performed on transport ready virtual drives:

- Clear - When you clear the transport ready policy of a virtual drive, the virtual drive moves back to the previous state.
- Locate - You can locate the virtual drives whose policy is set as transport ready.
- Force clear - You can use the force clear option to clear the transport ready policy set on a virtual drive.

The following operations can be performed on transport ready physical devices:

- You can locate the physical drives whose policy is set as transport ready.

NOTE The Transport Ready feature does not support operations other than the ones that are listed above.

Moving transport ready physical devices

When you move transport ready physical devices from virtual drives of one controller to another controller, the physical devices become foreign configurations and they have to be imported like the regular configurations.

Removing or Including transport ready physical devices

When you insert or remove transport ready physical devices on the source system, the physical devices will retain their transport ready state, and the drive state of the transport ready virtual drive does not change. When all the physical drives of a transport ready virtual drives are removed, virtual drives are deleted from the controller and the virtual group GUID and DHSP byte in NVRAM is cleared.

Appendix J: Glossary

This glossary defines the terms used in this document.

A

Absolute state of charge	Predicted remaining battery capacity expressed as a percentage of Design Capacity. Note that the Absolute State of Charge operation can return values greater than 100 percent.
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*.
Auto learn mode	The controller performs the learn cycle automatically in this mode. This mode offers the following options: ■ BBU Auto Learn: Firmware tracks the time since the last learn cycle and performs a learn cycle when due. ■ BBU Auto Learn Disabled: Firmware does not monitor or initiate a learn cycle. You can schedule learn cycles manually. ■ BBU Auto Learn Warn: Firmware warns about a pending learn cycle. You can initiate a learn cycle manually. After the learn cycle is complete, the firmware resets the counter and warns you when the next learn cycle time is reached.
Auto learn period	Time between learn cycles. A learn cycle is a battery calibration operation performed periodically by the controller to determine the condition of the battery.
Average time to empty	One-minute rolling average of the predicted remaining battery life.
Average time to full	Predicted time to charge the battery to a fully charged state based on the one minute rolling average of the charge current.

B

Battery module version	Current revision of the battery pack module.
Battery replacement	Warning issued by firmware that the battery can no longer support the required data retention time.
Battery retention time	Time, in hours, that the battery can maintain the contents of the cache memory.
Battery status	Operating status of the battery. Possible values are Missing, Optimal, Failed, Degraded (need attention), and Unknown.
Battery type	Possible values are intelligent Battery Backup Unit (BBU), intelligent Battery Backup Unit (iBBU), intelligent Transportable Battery Backup Unit (iTBBU®), and ZCR Legacy.
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.
BGI	Background Initialization Operation
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.

C

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.
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	Measure of the current flowing to (+) or from (-) the battery, reported in milliamperes.
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.
Cycle count	The count is based on the number of times the near fully charged battery has been discharged to a level below the cycle count threshold.

D

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.
Design capacity	Designed charge capacity of the battery, measured in milliampere-hour units (mAh).
Design charge capacity remaining	Amount of the charge capacity remaining, relative to the battery pack design capacity.
Design voltage	Designed voltage capacity of the battery, measured in millivolts (mV).
Device chemistry	Possible values are NiMH (nickel metal hydride) and LiON (lithium ion).
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 physical drive or a virtual drive property indicating the status of the appropriate drive. **Physical Drive State** A physical drive can be in any 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. In the output of the StorCLI commands, **Unconfigured Good** is displayed as **UGood**. ■ **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. In the output of the StorCLI commands, **Online** is displayed as **onln**. ■ **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. In the output of the StorCLI commands, **Unconfigured Bad** is displayed as **UBad**. ■ **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. In the output of the StorCLI commands, **Offline** is displayed as **offln**. ■ **None** – A drive with an unsupported flag set. An Unconfigured Good or Offline drive that has completed the prepare for removal operation.

Virtual Drive State

A virtual drive can be in any one of the following states:

- **Optimal** – A virtual drive whose members are all online.

  In the output of the StorCLI commands, **Optimal** is displayed as **optl**.
- **Partially Degraded** – A virtual drive with a redundant RAID level that is capable of sustaining more than one member drive failure. This state also applies to the virtual drive's member drives. Currently, a RAID 6 or RAID 60 virtual drive is the only virtual drive that can be partially degraded.

  In the output of the StorCLI commands, **Partially Degraded** is displayed as **Pdgd**.
- **Degraded** – A virtual drive with a redundant RAID level with one or more member failures and can no longer sustain a subsequent drive failure.

  In the output of the StorCLI commands, **Degraded** is displayed as **dgrd**.
- **Offline** - A virtual drive with on e or more member failures that make the data inaccessible.

  In the output of the StorCLI commands, **Offline** is displayed as **OfLn**.

Drive state 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.

E

EKM	External Key Management
Estimated time to recharge	Estimated time necessary to complete recharge of the battery at the current charge rate.
Expected margin of error	Indicates how accurate the reported battery capacity is in terms of percentage.

F

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. Avago 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.
FGI	Foreground Initialization Operation.
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.
Full charge capacity	Amount of charge that can be placed in the battery. This value represents the last measured full discharge of the battery. This value is updated on each learn cycle when the battery undergoes a qualified discharge from nearly full to a low battery level.

G

Gas gauge status	Hexadecimal value that represents the status flag bits in the gas gauge status register.

H

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.

I

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.)

L

LDBBM	Logical drive bad block management
Learn delay interval	Length of time between automatic learn cycles. You can delay the start of the learn cycles for up to 168 hours (seven days).
Learning cycle	A battery calibration operation performed by a RAID controller periodically to determine the condition of the battery. You can start battery learn cycles manually or automatically
Learn mode	Mode for the battery auto learn cycle. Possible values are Auto, Disabled, and Warning.
Learn state	Indicates that a learn cycle is in progress.
LKM	Local Key Management
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.
Low-power storage mode	Storage mode that causes the battery pack to use less power, which save battery power consumption.

M

Manufacturing date	Date on which the battery pack assembly was manufactured.
Manufacturing name	Device code that indicates the manufacturer of the components used to make the battery assembly.

Max error	Expected margin of error (percentage) in the state of charge calculation. For example, when Max Error returns 10 percent and Relative State of Charge returns 50 percent, the Relative State of charge is more likely between 50 percent and 60 percent. The gas gauge sets Max Error to 100 percent on a full reset. The gas gauge sets Max Error to 2 percent on completion of a learn cycle, unless the gas gauge limits the learn cycle to the +512/–256-mAh maximum adjustment values. If the learn cycle is limited, the gas gauge sets Max Error to 8 percent unless Max Error was already below 8 percent. In this case Max Error does not change. The gas gauge increments Max Error by 1 percent after four increments of Cycle Count without a learn cycle.
Maximum learn delay from current start time	Maximum length of time between automatic learn cycles. You can delay the start of a learn cycle for a maximum of 168 hours (7 days).
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.
N	
Name	A virtual drive property indicating the user-assigned name of the virtual drive.
Next learn time	Time at which the next learn cycle starts.
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 nonvolatile 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.
O	
Offline	A drive is offline when it is part of a virtual drive but its data is not accessible to the virtual drive.
P	
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 before 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.
Predicted battery capacity status (hold 24hr charge)	Indicates whether the battery capacity supports a 24-hour data retention time.

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.

R

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 1E	Uses two-way mirroring on two or more drives. RAID 1E provides better performance than a traditional RAID 1 array.
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. Avago SAS RAID controllers support RAID levels 0, 1, 5, 6, 10, 50, and 60.
RAID Migration	A feature in RAID subsystems that allows changing a RAID level to another level without powering down the system.
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.
Relative state of charge	Predicted remaining battery capacity expressed as a percentage of Full Charge Capacity.
Remaining capacity	Amount of remaining charge capacity of the battery as stated in milliamp hours. This value represents the available capacity or energy in the battery at any given time. The gas gauge adjusts this value for charge, self-discharge, and leakage compensation factors.
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.
Run time to empty	Predicted remaining battery life at the present rate of discharge in minutes.
S	
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.
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 1 MB of drive space and has 64 KB of data residing on each drive in the stripe. In this case, the stripe size is 1 MB and the strip size is 64 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.
Strip size	The portion of a stripe that resides on a single drive in the drive group.
Subvendor ID	A controller property that lists additional vendor ID information about the controller.
T	
Temperature	Temperature of the battery pack, measured in Celsius.
TR	Transport Ready
U	

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.

V

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.
VD	Virtual Disk.
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.

W

Write-back	In Write-Back Caching mode, the controller sends a data transfer completion signal to the host when the controller cache has received all 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.

History of Technical Changes

This appendix lists all the technical changes made to this guide for this release.

Version and Date	Description of Changes
Version 1.0, August 08, 2016	■ Added Section 4.14.22, Performing a Transport Ready Operation ■ Added Section 5.6.2, Performing a Transport Ready Operation ■ Updated Section 6.6.7.1, Drive Group Show Commands ■ Added Section 9.13, Performing a Transport Ready Operation ■ Updated Section E.1, Error Messages and Descriptions

AVAGO
TECHNOLOGIES