



ServeRAID-M Software

User Guide

October 2013

Part Number: 00AH215

Tenth Edition (October 2013)

© Copyright IBM Corporation 2007, 2013.

US Government Users Restricted Rights -- Use, duplication or disclosure restricted by
GSA ADP Schedule Contract with IBM Corp.

Table of Contents

Chapter 1: Overview	14
1.1 SAS Technology	14
1.2 Serial-attached SCSI Device Interface	14
1.3 Serial ATA Features	15
1.4 Solid State Drive Features	15
1.5 Solid State Drive Guard	16
1.6 Integrated MegaRAID Mode and MegaRAID Mode	16
1.7 Feature on Demand (FoD) Upgrades	17
1.8 Feature on Demand: iMR RAID 5 + Self-Encrypting Disk Upgrade	17
1.8.1 Feature on Demand: ServeRAID RAID 6/60 Upgrade	17
1.9 Feature on Demand: FastPath Upgrade	18
1.10 Feature on Demand: CacheCade 2 Upgrade	19
1.11 UEFI 2.0 Support	19
1.12 Configuration Scenarios	19
1.12.1 Valid Drive Mix Configurations	21
Chapter 2: Introduction to RAID	22
2.1 RAID Description	22
2.2 RAID Benefits	22
2.3 RAID Functions	22
2.4 Components and Features	22
2.5 Physical Array	23
2.6 Virtual Drive	23
2.7 RAID Drive Group	23
2.8 Fault Tolerance	23
2.8.1 Consistency Check	24
2.8.2 Copyback	24
2.8.3 Background Initialization	25
2.8.4 Patrol Read	25
2.8.5 Disk Striping	25
2.8.6 Disk Mirroring	26
2.8.7 Parity	27
2.8.8 Disk Spanning	27
2.8.9 Hot Spares	28
2.8.10 Disk Rebuilds	29
2.8.11 Rebuild Rate	30
2.8.12 Hot Swap	30
2.8.13 Drive States	30
2.8.14 Virtual Drive States	31
2.8.15 Enclosure Management	31
2.9 RAID Levels	31
2.9.1 Summary of RAID Levels	31
2.9.2 Selecting a RAID Level	32
2.9.3 RAID 0	32
2.9.4 RAID 1	33
2.9.5 RAID 5	34
2.10 RAID 6	34

2.10.1 RAID 10	35
2.10.2 RAID 50	36
2.10.3 RAID 60	37
2.11 RAID Configuration Strategies	38
2.11.1 Maximizing Fault Tolerance	39
2.11.2 Maximizing Performance	40
2.11.3 Maximizing Storage Capacity	41
2.12 RAID Availability	42
2.12.1 RAID Availability Concept	42
2.13 Configuration Planning	43
2.13.1 Number of Drives	43
2.13.2 Drive Group Purpose	43
Chapter 3: Self Encrypting Disk	44
3.1 Overview	44
3.2 Purpose	44
3.3 Terminology	44
3.4 Workflow	45
3.4.1 Enable Security	45
3.4.2 Change Security	46
3.4.3 Import a Foreign Configuration	47
3.5 Instant Secure Erase	47
Chapter 4: Starting the HII Configuration Utility	49
Chapter 5: Managing Configurations	52
5.1 Creating a Virtual Drive from a Profile	53
5.2 Manually Creating a Virtual Drive	56
5.3 Creating a CacheCade Virtual Drive	60
5.4 Viewing Drive Group Properties	61
5.5 Viewing Global Hot Spare Drives	62
5.6 Clearing a Configuration	63
5.7 Make Unconfigured Good and Make JBOD	64
5.7.1 Make Unconfigured Good	64
5.7.2 Make JBOD	65
5.8 Managing Foreign Configurations	66
5.8.1 Previewing and Importing a Foreign Configuration	66
5.8.2 Clearing a Foreign Configuration	69
Chapter 6: Managing Controllers	70
6.1 Viewing Advanced Controller Management Options	71
6.2 Scheduling a Consistency Check	76
6.3 Saving or Clearing Controller Events	78
6.4 Enabling or Disabling Drive Security	79
6.5 Changing a Security Key	82
6.6 Saving the TTY Log	84
6.7 Managing and Changing Link Speeds	85
Chapter 7: Managing Virtual Drives	86
7.1 Selecting Virtual Drive Operations	88
7.1.1 Locating Physical Drives in a Virtual Drive	88
7.1.2 Deleting a Virtual Drive	89
7.1.3 Reconfiguring a Virtual Drive	89

7.1.4 Initializing a Virtual Drive	91
7.1.5 Erasing a Virtual Drive	92
7.1.6 Enabling and Disabling SSD Caching	93
7.1.7 Securing a Virtual Drive	94
7.1.8 Running a Consistency Check	95
7.1.9 Expanding a Virtual Drive	97
7.2 Managing CacheCade Virtual Drives	98
7.3 Viewing Associated Drives	99
7.4 Viewing and Managing Virtual Drive Properties and Options	100
Chapter 8: Managing Physical Drives	103
8.1 Performing Drive Operations	104
8.1.1 Locating a Drive	105
8.1.2 Making a Drive Unconfigured Bad, Unconfigured Good, or JBOD	105
8.1.3 Replacing a Drive	105
8.1.4 Placing a Drive Offline	106
8.1.5 Placing a Drive Online	107
8.1.6 Assigning a Global Hot Spare Drive	108
8.1.7 Assigning a Dedicated Hot Spare Drive	109
8.1.8 Unassigning a Hot Spare Drive	110
8.1.9 Initializing or Erasing a Drive	111
8.1.10 Rebuilding a Drive	112
8.1.11 Securely Erasing a Drive	113
8.2 Viewing Advanced Drive Properties	115
Chapter 9: Managing Hardware Components	118
9.1 Managing Batteries	120
9.1.1 Setting Automatic Learn Cycle Properties	122
9.2 Managing Enclosures	123
Chapter 10: ServeRAID StorCLI	125
10.1 Overview	125
10.2 Support for MegaCLI Commands	125
10.3 Installation	125
10.3.1 Installing StorCLI on Microsoft Windows Operating Systems	126
10.3.2 Installing StorCLI on Linux Operating Systems	126
10.3.3 Installing StorCLI on Ubuntu Operating System	126
10.3.4 Installing StorCLI on VMware Operating Systems	126
10.3.5 Installing StorCLI on FreeBSD Operating Systems	126
10.3.6 Installing StorCLI on Microsoft EFI	127
10.3.7 Installing StorCLI on Solaris Operating Systems	127
10.4 StorCLI Command Syntax	127
10.5 Working with the Storage Command Line Tool	129
10.5.1 System Commands	129
10.5.2 Controller Commands	129
10.5.3 Drive Commands	140
10.5.4 Virtual Drive Commands	148
10.5.5 Foreign Configurations Commands	157
10.5.6 BIOS-Related Commands	158
10.5.7 Drive Group Commands	159
10.5.8 Dimmer Switch Commands	160
10.5.9 BBU Commands	161
10.5.10 Enclosure Commands	163
10.5.11 PHY Commands	164
10.5.12 Logging Commands	165
10.6 Frequently Used Tasks	166

10.6.1 Showing the Version of the Storage Command Line Tool	166
10.6.2 Showing StorCLI Help	166
10.6.3 Showing System Summary Information	166
10.6.4 Showing Free Space in a Controller	166
10.6.5 Adding Virtual Drives	166
10.6.6 Setting the Cache Policy in a Virtual Drive	167
10.6.7 Showing Virtual Drive Information	167
10.6.8 Deleting Virtual Drives	167
10.6.9 Flashing Controller Firmware	167
Chapter 11: MegaRAID Storage Manager Overview	169
11.1 Overview	169
11.1.1 Creating Storage Configurations	169
11.1.2 Monitoring Storage Devices	169
11.1.3 Maintaining Storage Configurations	169
11.2 Hardware and Software Requirements	169
11.3 Installing MegaRAID Storage Manager	170
11.3.1 Prerequisite for MegaRAID Storage Manager Installation	170
11.3.2 Installing MegaRAID Storage Manager Software on Microsoft Windows	171
11.3.3 Uninstalling the MegaRAID Storage Manager Software on Microsoft Windows	176
11.3.4 Installing MegaRAID Storage Manager Software for Solaris 10 x86	176
11.3.5 Uninstalling MegaRAID Storage Manager Software on the Solaris 10 x86 Operating System	177
11.3.6 Installing MegaRAID Storage Manager Software for Solaris 10 SPARC	177
11.3.7 Uninstalling MegaRAID Storage Manager Software on the Solaris SPARC Operating System	177
11.3.8 Prerequisites for Installing MegaRAID Storage Manager on the RHEL6.X x64 Operating System	178
11.3.9 Installing MegaRAID Storage Manager Software on RHEL or SLES/SuSE Linux	178
11.3.10 Linux Error Messages	179
11.3.11 Kernel Upgrade	180
11.3.12 Uninstalling MegaRAID Storage Manager Software on RHEL or SLES/SuSE Linux	180
11.3.13 MegaRAID Storage Manager Software Customization	181
11.3.14 Stopping the Pop-Up Notification Process	181
11.3.15 Restarting the Pop-Up Notification Process	182
11.4 Installing and Supporting MegaRAID Storage Manager Software on VMware	182
11.4.1 Prerequisites for Installing MegaRAID Storage Manager for VMware	182
11.4.2 Installing MegaRAID Storage Manager on VMware ESX (VMware Classic)	182
11.4.3 Uninstalling MegaRAID Storage Manager for VMware	183
11.4.4 MegaRAID Storage Manager Support on the VMware ESXi Operating System	183
11.4.5 Limitations of Installation and Configuration	184
11.5 Installing and Configuring a CIM Provider	185
11.5.1 Installing a CIM SAS Storage Provider on the Linux Operating System	185
11.5.2 Running the CIM SAS Storage Provider on Pegasus	185
11.5.3 Installing a CIM SAS Storage Provider on Windows	186
11.6 Installing and Configuring an SNMP Agent	186
11.6.1 Prerequisite for IBM SNMP Agent RPM Installation	186
11.6.2 Prerequisite for Installing SNMP Agent on Linux Server	187
11.6.3 Installing and Configuring an SNMP Agent on Linux	187
11.6.4 Installing and Configuring an SNMP Agent on Solaris	188
11.6.5 Installing a SNMP Agent on Windows	191
11.7 MegaRAID Storage Manager Remotely Connecting to VMware ESX	193
11.8 Prerequisites to Running MegaRAID Storage Manager Remote Administration	193
Chapter 12: MegaRAID Storage Manager Screens	194
12.1 Starting the MegaRAID Storage Manager Software	194
12.2 Discovery and Login	194
12.3 LDAP Support	198
12.4 Configuring LDAP Support Settings	200

12.5 MegaRAID Storage Manager Main Menu	201
12.5.1 Dashboard / Physical View/ Logical View	201
12.5.2 Physical Drive Temperature	203
12.5.3 Shield State	204
12.5.4 Shield State Physical View	204
12.5.5 Logical View Shield State	205
12.5.6 Viewing the Physical Drive Properties	205
12.5.7 Viewing Server Profile of a Drive in Shield State	206
12.5.8 Displaying the Virtual Drive Properties	207
12.5.9 Emergency Spare	210
12.5.10 SSD Disk Cache Policy	213
12.5.11 Non-SED Secure Erase Support	216
12.5.12 Rebuild Write Cache	221
12.5.13 Background Suspend or Resume Support	221
12.5.14 Enclosure Properties	223
12.6 GUI Elements in the MegaRAID Storage Manager Window and Menus	223
12.6.1 Device Icons	224
12.6.2 Properties and Graphical View Tabs	225
12.6.3 Event Log Panel	226
12.6.4 Menu Bar	226
Chapter 13: Configurations	228
13.1 Creating a New Configuration	228
13.1.1 Selecting Virtual Drive Settings	228
13.1.2 Optimum Controller Settings for CacheCade	229
13.1.3 Optimum Controller Settings for Fast Path	229
13.1.4 Creating a Virtual Drive Using Simple Configuration	229
13.1.5 Creating a Virtual Drive Using Advanced Configuration	233
13.2 Converting JBOD Drives to Unconfigured Good	239
13.2.1 Converting JBOD to Unconfigured Good from the MegaRAID Storage Manager Main Menu	240
13.3 Adding Hot Spare Drives	240
13.4 Changing Adjustable Task Rates	241
13.5 Changing Power Settings	242
13.6 Recovering and Clearing Punctured Block Entries	243
13.7 Changing Virtual Drive Properties	244
13.8 Changing a Virtual Drive Configuration	246
13.8.1 Accessing the Modify Drive Group Wizard	246
13.8.2 Adding a Drive or Drives to a Configuration	248
13.8.3 Removing a Drive from a Configuration	251
13.8.4 Replacing a Drive	251
13.8.5 Migrating the RAID Level of a Virtual Drive	252
13.8.6 New Drives Attached to a ServeRAID Controller	255
13.9 Deleting a Virtual Drive	256
Chapter 14: Monitoring System Events and Storage Devices	257
14.1 Alert Delivery Methods	257
14.1.1 Vivaldi Log/MegaRAID Storage Manager Log	257
14.1.2 System Log	258
14.1.3 Pop-up Notification	259
14.1.4 Email Notification	259
14.2 Configuring Alert Notifications	260
14.3 Editing Alert Delivery Methods	262
14.4 Changing Alert Delivery Methods for Individual Events	263
14.5 Changing the Severity Level for Individual Events	264
14.6 Roll Back to Default Individual Event Configuration	264

14.7 Entering or Editing the Sender Email Address and SMTP Server	265
14.8 Authenticating the SMTP Server	266
14.9 Adding Email Addresses of Recipients of Alert Notifications	266
14.10 Testing Email Addresses of Recipients of Alert Notifications	267
14.11 Removing Email Addresses of Recipients of Alert Notifications	267
14.12 Saving Backup Configurations	268
14.13 Loading Backup Configurations	268
14.14 Monitoring Server Events	268
14.15 Monitoring Controllers	269
14.16 Monitoring Drives	270
14.17 Running a Patrol Read	271
14.17.1 Patrol Read Task Rates	272
14.18 Monitoring Virtual Drives	272
14.19 Monitoring Enclosures	274
14.20 Monitoring Battery Backup Units	274
14.21 Battery Learn Cycle	275
14.21.1 Setting Automatic Learn Cycle Properties	276
14.21.2 Starting a Learn Cycle Manually	277
14.22 Monitoring Rebuilds and Other Processes	277
Chapter 15: Maintenance	279
15.1 Initializing a Virtual Drive	279
15.1.1 Running a Group Initialization	279
15.2 Running a Consistency Check	280
15.2.1 Setting the Consistency Check Settings	281
15.2.2 Scheduling a Consistency Check	281
15.2.3 Running a Group Consistency Check	283
15.3 Scanning for New Drives	284
15.4 Rebuilding a Drive	284
15.4.1 New Drives Attached to a ServeRAID Controller	285
15.5 Making a Drive Offline or Missing	285
15.6 Removing a Drive	286
15.7 Upgrading Firmware	286
Chapter 16: CacheCade 2.0	288
16.1 Logical Drive Property Settings Required for CacheCade	288
16.2 Viewing a Logical Drive with CacheCade	289
16.3 WebBIOS Configuration for CacheCade	289
16.4 MegaRAID Storage Manager Configuration for CacheCade	292
16.5 Modifying the CacheCade Virtual Drive Properties	294
16.6 FastPath Advanced Software	298
Appendix A: Events and Messages	300
A.1 Error Levels	300
A.2 Event Messages	300
Appendix B: 3Ware CLI Commands to StorCLI Command Conversion	319
B.1 System Commands	319
B.2 Controller Commands	319
B.3 Alarm Commands	322
B.4 Patrol Read and Consistency Check Commands	322
B.5 BBU Commands	323

B.6 Virtual Drive Commands	324
B.7 Physical Drive Commands	326
B.8 Enclosure Commands	327
B.9 Events and Logs	328
B.10 Miscellaneous Commands	328
Appendix C: MegaCLI Commands to StorCLI Command Conversion	329
C.1 System Commands	329
C.2 Controller Commands	329
C.3 Patrol Read Commands	332
C.4 Consistency Check Commands	333
C.5 OPROM BIOS Commands	333
C.6 Battery Commands	334
C.7 RAID Configuration Commands	335
C.8 Security Commands	336
C.9 Virtual Drive Commands	337
C.10 Physical Drive Commands	338
C.11 Enclosure Commands	340
C.12 PHY Commands	341
C.13 Alarm Commands	341
C.14 Event Log Properties Commands	341
C.15 Premium Feature Key Commands	342
Appendix D: Unsupported Commands in Embedded MegaRAID	343

[illegible]

Important:

Each caution and danger statement in this document is labeled with a number. This number is used to cross reference an English-language caution or danger statement with translated versions of the caution or danger statement in the *Safety Information* document.

For example, if a caution statement is labeled "Statement 1," translations for that caution statement are in the *Safety Information* document under "Statement 1."

Be sure to read all caution and danger statements in this document before you perform the procedures. Read any additional safety information that comes with the server or optional device before you install the device.

This device is intended for use with UL Listed IBM devices.

Statement 1:



DANGER

Electrical current from power, telephone, and communication cables is hazardous.

To avoid a shock hazard:

- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- Connect all power cords to a properly wired and grounded electrical outlet.
- Connect to properly wired outlets any equipment that will be attached to this product.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.
- Connect and disconnect cables as described in the following table when installing, moving, or opening covers on this product or attached devices.

To Connect:	To Disconnect:
1. Turn everything OFF.	1. Turn everything OFF.
2. First, attach all cables to devices.	2. First, remove power cords from outlet.
3. Attach signal cables to connectors.	3. Remove signal cables from connectors.
4. Attach power cords to outlet.	4. Remove all cables from devices.
5. Turn device ON.	

Statement 3:



CAUTION:

When laser products (such as CD-ROMs, DVD drives, fiber optic devices, or transmitters) are installed, note the following:



DANGER

Some laser products contain an embedded Class 3A or Class 3B laser diode. Note the following.

Laser radiation when open. Do not stare into the beam, do not view directly with optical instruments, and avoid direct exposure to the beam.



Class 1 Laser Product
Laser Klasse 1
Laser Klass 1
Luokan 1 Laserlaite
Appareil À Laser de Classe 1

Statement 8:



CAUTION:

Never remove the cover on a power supply or any part that has the following label attached.



Hazardous voltage, current, and energy levels are present inside any component that has this label attached. There are no serviceable parts inside these components. If you suspect a problem with one of these parts, contact a service technician.

Statement 28:



CAUTION:

The battery is a lithium ion battery. To avoid possible explosion, do not burn the battery. Exchange it only with the approved part. Recycle or discard the battery as instructed by local regulations.

Chapter 1: Overview

This guide documents the utilities used to configure, monitor, and maintain IBM ServeRAID-M[®] Serial-attached SCSI (SAS)/Serial-ATA (SATA) controllers with RAID control capabilities and the storage-related devices connected to them. This guide explains how to use the MegaRAID Storage Manager™ software and the WebBIOS utility. In addition, it documents self-encrypting disks (SED), SAS technology, SATA technology, Solid State Drive (SSD) technology, configuration scenarios, and drive types.

1.1 SAS Technology

The ServeRAID controllers are high-performance, Serial-attached-SCSI/Serial ATA controllers with RAID control capabilities. ServeRAID 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 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) devices, and SSD devices within the same storage infrastructure. These devices bring individual characteristics that make each one a more suitable choice depending on your storage needs. ServeRAID-M 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 mix SAS drives and SATA drives within the same *virtual drives*. Although you can mix drives, IBM strongly discourages the practice. This applies to both HDDs and SSDs.

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

The controllers support the SAS protocol as described in the *Serial Attached SCSI Standard, Version 2.0*. In addition, the controller supports the SATA II protocol defined by the *Serial ATA specification, Version 2.0*, and the SATA III protocol defined by the *Serial ATA Specification, Version 3.0*. SATA III is an extension to SATA II.

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

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

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



NOTE The ServeRAID SAS/SATA controllers support SATA, SATA II, and SATA III technologies.

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

- Supports SATA III data transfers of 6.0 Gbits/s
- Supports STP data transfers of 6.0 Gbits/s
- Provides a serial, point-to-point storage interface
- Simplifies cabling between devices
- Eliminates the master-slave construction used in parallel ATA
- Allows addressing of multiple SATA 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 firmware supports Solid State Drives attached to ServeRAID-M SAS controllers. The features and operations for SSDs are the same as for hard disk drives (HDDs), and these drives are expected to behave like SATA HDDs or SAS HDDs. The major advantages of SSDs include:

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

The WebBIOS Configuration Utility and the MegaRAID Storage Manager utility display the SSDs by the type, either SAS or SATA. For example, a SATA SSD drive displays as “SSD (SATA)”. HDDs are identified simply as “SAS” or “SATA”.



NOTE ServeRAID-M implements support for only those SATA SSD drives which support ATA-8 ACS compliance.

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



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

1.5 Solid State Drive Guard

SSDs are known for their reliability and performance. SSD Guard™ 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 very reliable, 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 S.M.A.R.T. (Self-Monitoring, Analysis and Reporting Technology) error log. If errors indicate that an SSD failure is imminent, ServeRAID-M starts a rebuild to preserve the data on the SSD and sends appropriate warning event notifications.

1.6 Integrated MegaRAID Mode and MegaRAID Mode

Some ServeRAID SAS/SATA controllers function in either integrated MegaRAID (iMR) mode or in MegaRAID (MR) mode.

Integrated MegaRAID is a highly integrated, low-cost RAID solution made possible by Fusion-MPT™ architecture. Integrated MegaRAID is a processor-based, hardware RAID solution designed for system environments requiring redundancy and high availability where a full-featured RAID implementation is not desired or might be cost prohibitive.

The major advantage of Integrated MegaRAID is that iMR provides RAID at the processor level, so it does not burden the CPU, which allows for more efficient operation.

The major advantage of MegaRAID mode is that the MR mode supports more RAID levels than iMR mode. iMR mode supports RAID levels 0, 1, 5, 10, and 50. MR mode supports RAID levels 0, 1, 5, 6, 10, 50, and 60.



NOTE For the ServeRAID M1100 SAS/SATA controllers and the ServeRAID M5100 SAS/SATA controllers, iMR RAID 5 requires purchase of the Feature on Demand (FoD) upgrade



NOTE For the ServeRAID M1100 SAS/SATA controllers and the ServeRAID M5100 SAS/SATA controllers, MegaRAID RAID 5/50 requires a transportable memory module (3 options) or the Feature on Demand upgrade



NOTE For the ServeRAID M1100 SAS/SATA controllers and the ServeRAID M5100 SAS/SATA controllers, MegaRAID RAID 6/60 requires a transportable memory module (3 options) and the Feature on Demand upgrade.

See [Section 1.7, Feature on Demand \(FoD\) Upgrades](#) for information about these upgrades.

See [Section 2.9.1, Summary of RAID Levels](#) for information about the supported RAID levels.

1.7 Feature on Demand (FoD) Upgrades

To use RAID levels 5, 6, 50, or 60, the FastPath feature, or the CacheCade 2.0 feature with selected controllers, you must install a Feature on Demand (FoD) upgrade and/or a transportable memory module (TMM). The following sections describe these upgrades, required installations, and supported controllers.

The following table lists the FoD upgrades available.

Table 1.1 List of Feature on Demand Upgrades

IBM Option	Official Name	Functionality
81Y4542	ServeRAID M1100 Series Zero Cache/RAID 5 Upgrade for IBM System x	iMR RAID 5 + SED
81Y4544	ServeRAID M5100 Series Zero Cache/RAID 5 Upgrade for IBM System x	iMR RAID 5 + SED
81Y4546	ServeRAID M5100 Series RAID 6 Upgrade for IBM System x	MegaRAID RAID 6
90Y4273	ServeRAID M5100 Series Performance Accelerator for IBM System x	FastPath
90Y4318	ServeRAID M5100 Series SSD Caching Enabler for IBM System x	CacheCade 2.0

1.8 Feature on Demand: iMR RAID 5 + Self-Encrypting Disk Upgrade

The ServeRAID M5100 Series Zero Cache/RAID 5 Upgrade for IBM System x supports iMR RAID levels 5 and 50, and self-encrypting disks for the following ServeRAID SAS/SATA controllers:

- ServeRAID M5110 SAS/SATA controller for IBM System x
- ServeRAID M5110e SAS/SATA controller for IBM System x
- ServeRAID M5120 SAS/SATA controller for IBM System x

The ServeRAID M1100 Series Zero Cache/RAID 5 Upgrade for IBM System x supports iMR RAID levels 5 and 50, and self-encrypting disks for the following ServeRAID SAS/SATA controller:

- The ServeRAID M1115 SAS/SATA controller for IBM System x

The SED feature offers the ability to encrypt data on disks and use disk-based key management to provide data security. With the SED feature, data is encrypted by the drives. You can designate which data to encrypt at the individual virtual drive (VD) level.

This solution provides data protection in the event of theft or loss of physical drives. With self-encrypting disks, if you remove a drive from its storage system or the server it is housed in, the data on that drive is encrypted and useless to anyone who attempts to access without the appropriate security authorization.

See [Section Chapter 3, Self Encrypting Disk](#) for more information about the self-encrypting disk feature.

See [Section 2.9.1, Summary of RAID Levels](#) for information about the supported RAID levels.

1.8.1 Feature on Demand: ServeRAID RAID 6/60 Upgrade

The ServeRAID M5100 Series RAID 6 Upgrade for System x is a Feature on Demand that supports ServeRAID levels 6 and 60 for the following ServeRAID SAS/SATA controllers:

- ServeRAID M5110 SAS/SATA controller for IBM System x

- ServeRAID M5110e SAS/SATA controller for IBM System x
- ServeRAID M5120 SAS/SATA controller for IBM System x

Install the ServeRAID M5100 Series RAID 6 Upgrade for System FoD *and* any of the transportable memory modules in the following table to upgrade to support RAID 6 and 60.

Table 1.2 Transportable Memory Modules

IBM Option	Official Name	Functionality
81Y4584	ServeRAID M5100 Series 512MB Cache/RAID 5 Upgrade for IBM System x	ServeRAID RAID 5 + SED
81Y4587	ServeRAID M5100 Series 512MB Flash/RAID 5 Upgrade for IBM System x	ServeRAID RAID 5 + SED
81Y4559	ServeRAID M5100 Series 1GB Flash/RAID 5 Upgrade for IBM System x	ServeRAID RAID 5 + SED

1.9 Feature on Demand: FastPath Upgrade

ServeRAID M5100 Series Performance Accelerator for IBM System x is a Feature on Demand that supports the FastPath feature for the following ServeRAID SAS/SATA controllers:

- ServeRAID M5110 SAS/SATA controller for IBM System x
- ServeRAID M5110e SAS/SATA controller for IBM System x
- ServeRAID M5120 SAS/SATA controller for IBM System x

The FastPath feature is a high-performance IO accelerator for SSD drive groups connected to a ServeRAID controller card. FastPath software combined with SSDs delivers a performance advantage over HDD installations and consumes less power. This feature dramatically boosts storage subsystem bandwidth and overall transactional application performance when used with a 12Gb/s/s ServeRAID SAS/SATA controller.

The FastPath feature delivers optimization of SSD virtual disk groups to enable read and write IOPS three times greater than ServeRAID controllers not using FastPath technology. In addition, the FastPath advanced software is faster and more cost-effective than current flash-based adapter card solutions. This feature is faster and more cost-effective than current Flash-based adapter card solutions.

To use the FastPath feature, you must install the ServeRAID M5100 Series Performance Accelerator for IBM System x FoD and a transportable memory module. See [Section 1.8.1, Feature on Demand: ServeRAID RAID 6/60 Upgrade](#) for the list of transportable memory modules.

The following table lists the option number and functionality for the ServeRAID M5100 Series Performance Accelerator for IBM System x.

Table 1.3 ServeRAID M5100 Series Performance Accelerator for IBM System x Upgrade

IBM Option	Official Name	Functionality
81Y4544	ServeRAID M5100 Series Performance Accelerator for IBM System x	FastPath

1.10 Feature on Demand: CacheCade 2 Upgrade

The ServeRAID M5100 Series SSD Caching Enabler for IBM System x is a Feature on Demand that supports the CacheCade 2 feature for the following ServeRAID SAS/SATA controllers:

- ServeRAID M5110 SAS/SATA controller for IBM System x
- ServeRAID M5110e SAS/SATA controller for IBM System x
- ServeRAID M5120 SAS/SATA controller for IBM System x

The CacheCade 2 Software feature improves I/O performance to meet the needs of high-performing Solid State Drives (SSDs). In addition, this feature benefits hard disk drives (HDDs). To support full-throughput for multiple direct-attached SSDs, the CacheCade 2 Software feature reduces I/O-processing overhead for the ServeRAID controllers. CacheCade 2 Software offers performance equivalent to Flash-based controllers.

To use the CacheCade 2.0 feature, you must install the ServeRAID M5100 Series SSD Caching Enabler for IBM System x FoD and a transportable memory module. See [Section 1.8.1, Feature on Demand: ServeRAID RAID 6/60 Upgrade](#) for the list of transportable memory modules.

The following table lists the option number and functionality for the ServeRAID M5100 Series SSD Caching Enabler for IBM System x.

Table 1.4 ServeRAID M5100 Series SSD Caching Enabler for IBM System x Upgrade

IBM Option	Official Name
90Y4318	ServeRAID M5100 Series SSD Caching Enabler for IBM System x

1.11 UEFI 2.0 Support

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

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

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

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

1.12 Configuration Scenarios

You can use the SAS RAID controllers in the following three main scenarios:

- **Low-end, internal SATA configurations:** In this configuration, use the RAID controller as a high-end SATA compatible controller that connects up to eight disks either directly or through a port expander. This configuration is mostly for low-end or entry servers. Enclosure management is provided through out-of-band I²C bus. Side bands of both types of internal SAS connectors support the SFF-8485 (SGPIO) interface.
- **Midrange internal SAS configurations:** This configuration is like the internal SATA configurations, but with high-end disks. This configuration is more suitable for low-range to midrange servers.

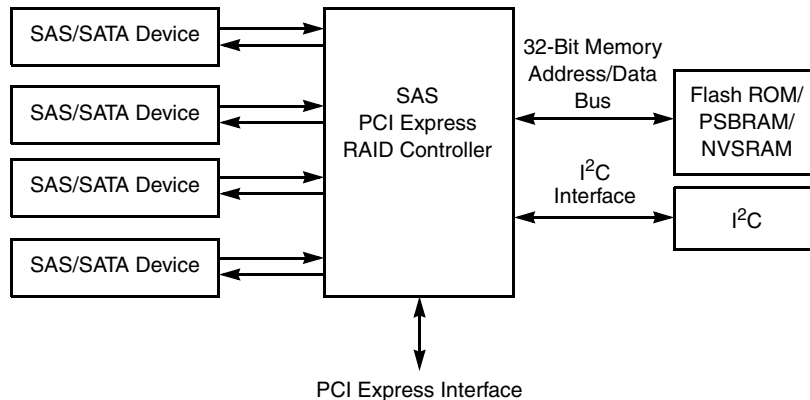
- **High-end external SAS/SATA configurations:** This configuration is for both internal connectivity and external connectivity, using SATA drives, SAS drives, or both. External enclosure management is supported through in-band, SCSI-enclosed storage. The configuration must support STP and SMP.

The following figure shows a direct-connect configuration. The Inter-IC (I²C) interface communicates with peripherals. The external memory bus provides a 32-bit memory bus, parity checking, and chip select signals for pipelined synchronous burst static random access memory (PSBRAM), nonvolatile static random access memory (NVSRAM), and Flash ROM.



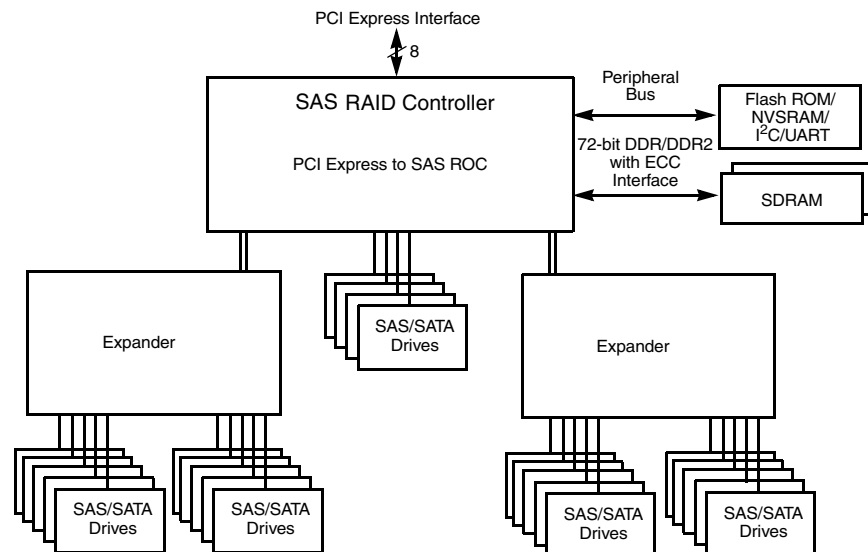
NOTE The external memory bus is 64-bit for the ServeRAID-MR10il RAID controller and the ServeRAID-MR10M RAID controller.

Figure 1.1 Example of a SAS Direct-Connect Application



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

Figure 1.2 Example of a SAS RAID Controller Configured with an Expander



1.12.1 Valid Drive Mix Configurations

You *cannot* have both SSDs and HDDs in a virtual drive. For a virtual drive that consists of SSDs only, you can choose whether to allow both SAS SSD drives and SATA SSD drives in that virtual drive.



NOTE The valid drive mix applies to hot spares, also. For hot spare information, see [Section 2.8.9, Hot Spares](#).

Chapter 2: Introduction to RAID

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

2.1 RAID Description

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

2.2 RAID Benefits

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

2.3 RAID Functions

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

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

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

2.4 Components and Features

RAID levels describe a system for ensuring the availability and redundancy of data stored on large disk subsystems. See [Section 2.9, "RAID Levels,"](#) for detailed information about RAID levels. The following subsections describes the components of RAID drive groups and RAID levels.

2.5 Physical Array

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

2.6 Virtual Drive

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

2.7 RAID Drive Group

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

2.8 Fault Tolerance

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

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

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



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

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

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

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

2.8.0.1 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. Alerts about drives contain correct information about the "enclosure", when the drives are connected by multiple paths. The enclosure module detects partner ESMs and issue events appropriately.

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

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

2.8.1 Consistency Check

The Consistency Check operation verifies correctness of the data in virtual drives that use RAID levels 1, 5, 6, 10, 50, and 60. (RAID 0 does not provide data redundancy). For example, in a system with parity, checking consistency means computing the data on one drive and comparing the results to the contents of the parity drive.



NOTE Perform a consistency check at least once a month.

2.8.2 Copyback

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

Typically, when a drive fails or is expected to fail, the data is rebuilt on a hot spare. The failed drive is replaced with a new 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.

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



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

2.8.2.1 Order of Precedence

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

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

2.8.3 Background Initialization

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

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

2.8.4 Patrol Read

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

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


You can use the MegaRAID Command Tool or the MegaRAID Storage Manager to select the patrol read options, which you can use to set automatic or manual operation, or disable patrol read.

2.8.5 Disk Striping

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

For example, in a four-disk system using only disk striping (used in RAID level 0), segment 1 is written to disk 1, segment 2 is written to disk 2, and so on. Disk striping enhances performance because multiple drives are accessed simultaneously, but disk striping does not provide data redundancy.

Figure 2.1 Example of Disk Striping (RAID 0)



2.8.5.1 Stripe Width

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

2.8.5.2 Stripe Size

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

2.8.5.3 Strip Size


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

2.8.6 Disk Mirroring

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

Disk mirroring provides 100 percent redundancy, but is expensive because each drive in the system must be duplicated.

Figure 2.2 Example of Disk Mirroring (RAID 1)



2.8.7 Parity


Parity generates a set of redundancy data from two or more parent data sets. The redundancy data can be used to reconstruct one of the parent data sets in the event of a drive failure. Parity data does not fully duplicate the parent data sets, but parity generation can slow the write process. In RAID, this method is applied to entire drives or stripes across all of the drives in a drive group. [Table 1](#) describes the types of parity.

Table 1 Types of Parity

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

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

Figure 2.3 Example of Distributed Parity (RAID 5)



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

2.8.8 Disk Spanning

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

Spanning alone does not provide reliability or performance enhancements. Spanned virtual drives must have the same stripe size and must be contiguous. In [Figure 2.4](#), RAID 1 drive groups are turned into a RAID 10 drive group.




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



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

Figure 2.4 Example of Disk Spanning



2.8.8.1 Spanning for RAID 10, RAID 50, and RAID 60

Table 2.1 describes how to use spanning to configure RAID 10, 50, and 60 virtual drives. The virtual drives must have the same stripe size and the maximum number of spans is eight. The full drive capacity is used when you span virtual drives; you cannot specify a smaller drive capacity.

See <hyperactive>Chapter 7, “Configuration,” for detailed procedures for configuring drive groups and virtual drives, and spanning the drives.

Table 2.1 Spanning for RAID 10, RAID 50, and RAID 60

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

2.8.9 Hot Spares

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



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

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

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

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



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

There are two types of hot spares:

- Global hot spare
- Dedicated hot spare

2.8.9.1 Global Hot Spare

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

2.8.9.2 Dedicated Hot Spare

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

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

Observe the following parameters when using hot spares:

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

2.8.10 Disk Rebuilds

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

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

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



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



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

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

2.8.11 Rebuild Rate

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

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

2.8.12 Hot Swap

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

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

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

2.8.13 Drive States

A drive state is a property indicating the status of the drive. The drive states are described in <hyperactive>Table 2.2.

Table 2.2 Drive States

State	Description
Online	A drive that can be accessed by the RAID controller and is part of the virtual drive.
Unconfigured Good	A drive that is functioning normally but is not configured as a part of a virtual drive or as a hot spare.
Hot Spare	A drive that is powered up and ready for use as a spare in case an online drive fails.
Failed	A drive that was originally configured as Online or Hot Spare, but on which the firmware detects an unrecoverable error.
Rebuild	A drive to which data is being written to restore full redundancy for a virtual drive.

Table 2.2 Drive States (Continued)

State	Description
Unconfigured Bad	A drive on which the firmware detects an unrecoverable error; the drive was Unconfigured Good or the drive could not be initialized.
Missing	A drive that was Online but which has been removed from its location.
Offline	A drive that is part of a virtual drive but which has invalid data as far as the RAID configuration is concerned. Note: When a virtual drive with cached data goes offline, the cache for the virtual drive is discarded. Because the virtual drive is offline, the cache cannot be saved.

2.8.14 Virtual Drive States

The virtual drive states are described in <hyperactive>Table 2.3.

Table 2.3 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.
Partially Degraded	The operating condition in a RAID 6 virtual drive is not optimal. One of the configured drives has failed or is offline. RAID 6 can tolerate up to two drive failures.
Failed	The virtual drive has failed.
Offline	The virtual drive is not available to the RAID controller.

2.8.15 Enclosure Management

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

2.9 RAID Levels

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

2.9.1 Summary of RAID Levels

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

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

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

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

RAID 10, a combination of RAID 0 and RAID 1, consists of striped data across mirrored spans. A RAID 10 drive group is a spanned drive group that creates a striped set from a series of mirrored drives. RAID 10 allows a maximum of eight spans. You must use an even number of drives in each RAID virtual drive in the span. The RAID 1 virtual drives must have the same stripe size. RAID 10 provides high data throughput and complete data redundancy but uses a larger number of spans.

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

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



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

2.9.2 Selecting a RAID Level

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

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

2.9.3 RAID 0

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




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

By breaking up a large file into smaller segments, the RAID controller can use both SAS drives and SATA drives to read or write the file faster. RAID 0 involves no parity calculations to complicate the write operation. This makes RAID 0 ideal for applications that require high bandwidth but do not require fault tolerance. <hyperactive>Table 2.4 provides an overview of RAID 0. The following figure provides a graphic example of a RAID 0 drive group.

Table 2.4 RAID 0 Overview

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

Figure 2.5 RAID 0 Drive Group Example with Two Drives




2.9.4 RAID 1

In RAID 1, the RAID controller duplicates all data from one drive to a second drive in the drive group. RAID 1 supports an even number of drives from 2 to 32 in a single span. RAID 1 provides complete data redundancy, but at the cost of doubling the required data storage capacity. <hyperactive>Table 2.5 provides an overview of RAID 1. The following figure provides a graphic example of a RAID 1 drive group.

Table 2.5 RAID 1 Overview

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

Figure 2.6 RAID 1 Drive Group



2.9.5 RAID 5

RAID 5 includes disk striping at the block level and parity. Parity is the data's property of being odd or even, and parity checking is used to detect errors in the data. In RAID 5, the parity information is written to all drives. RAID 5 is best suited for networks that perform a lot of small input/output (I/O) transactions simultaneously.


RAID 5 addresses the bottleneck issue for random I/O operations. Because each drive contains both data and parity, numerous writes can take place concurrently.

<hyperactive>Table 2.6 provides an overview of RAID 5. The following figure provides a graphic example of a RAID 5 drive group.

Table 2.6 RAID 5 Overview

Uses	Provides high data throughput, especially for large files. Use RAID 5 for transaction processing applications because each drive can read and write independently. If a drive fails, the RAID controller uses the parity drive to recreate all missing information. Use also for office automation and online customer service that requires fault tolerance. Use for any application that has high read request rates but low write request rates.
Strong Points	Provides data redundancy, high read rates, and good performance in most environments. Provides redundancy with lowest loss of capacity.
Weak Points	Not well-suited to tasks requiring lot of writes. Suffers more impact if no cache is used (clustering). Drive performance will be reduced if a drive is being rebuilt. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.
Drives	3 to 32

Figure 2.7 RAID 5 Drive Group with Six Drives



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

2.10 RAID 6

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

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


<hyperactive>Table 2.5 provides a graphic example of a RAID 6 drive group.

Table 2.7 RAID 6 Overview

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

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

Figure 2.8 Example of Distributed Parity across Two Blocks in a Stripe (RAID 6)



Parity is distributed across all drives in the drive group.


2.10.1 RAID 10

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

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

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

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



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


<hyperactive>Table 2.8 provides an overview of RAID 10.

Table 2.8 RAID 10 Overview

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

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

Figure 2.9 RAID 10 Level Virtual Drive



2.10.2 RAID 50

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

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


RAID level 50 can support up to eight spans and tolerate up to eight drive failures, though less than total drive capacity is available. Though multiple drive failures can be tolerated, only one drive failure can be tolerated in each RAID 5 level drive group.

<hyperactive>Table 2.9 provides an overview of RAID 50. The following figure shows an example of a RAID 50 array.

Table 2.9 RAID 50 Overview

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

Figure 2.10 RAID 50 Level Virtual Drive



2.10.3 RAID 60

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

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


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

Table 2.10 RAID 60 Overview

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

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

Figure 2.11 RAID 60 Level Virtual Drive



Parity is distributed across all drives in the drive group.

2.11 RAID Configuration Strategies

The most important factors in RAID drive group configuration are:

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

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

2.11.1 Maximizing Fault Tolerance

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

A hot swap is the manual substitution of a replacement unit in a disk subsystem for a defective one, where the substitution can be performed while the subsystem is running hot swap drives. Auto-Rebuild in the WebBIOS Configuration Utility allows a failed drive to be replaced and automatically rebuilt by “hot-swapping” the drive in the same drive bay. The RAID drive group continues to handle requests while the rebuild occurs, providing a high degree of fault tolerance and zero downtime. <hyperactive>Table 2.11 describes the fault tolerance for each RAID level.

Table 2.11 RAID Levels and Fault Tolerance

RAID Level	Fault Tolerance
0	Does not provide fault tolerance. All data lost is if any drive fails. Disk striping writes data across multiple drives instead of just one drive. It involves partitioning each drive storage space into stripes that can vary in size. RAID 0 is ideal for applications that require high bandwidth but do not require fault tolerance.
1	Provides complete data redundancy. If one drive fails, the contents of the other drive in the drive group can be used to run the system and reconstruct the failed drive. The primary advantage of disk mirroring is that it provides 100 percent data redundancy. Since the contents of the drive are completely written to a second drive, no data is lost if one of the drives fails. Both drives contain the same data at all times. RAID 1 is ideal for any application that requires fault tolerance and minimal capacity.
5	Combines distributed parity with disk striping. Parity provides redundancy for one drive failure without duplicating the contents of entire drives. If a drive fails, the RAID controller uses the parity data to reconstruct all missing information. In RAID 5, this method is applied to entire drives or stripes across all drives in a drive group. Using distributed parity, RAID 5 offers fault tolerance with limited overhead.
6	Combines distributed parity with disk striping. RAID 6 can sustain two drive failures and still maintain data integrity. Parity provides redundancy for two drive failures without duplicating the contents of entire drives. If a drive fails, the RAID controller uses the parity data to reconstruct all missing information. In RAID 6, this method is applied to entire drives or stripes across all of the drives in a drive group. Using distributed parity, RAID 6 offers fault tolerance with limited overhead.

Table 2.11 RAID Levels and Fault Tolerance (Continued)

RAID Level	Fault Tolerance
10	Provides complete data redundancy using striping across spanned RAID 1 drive groups. RAID 10 works well for any environment that requires the 100 percent redundancy offered by mirrored drive groups. RAID 10 can sustain a drive failure in each mirrored drive group and maintain drive integrity.
50	Provides data redundancy using distributed parity across spanned RAID 5 drive groups. RAID 50 includes both parity and disk striping across multiple drives. If a drive fails, the RAID controller uses the parity data to recreate all missing information. RAID 50 can sustain one drive failure per RAID 5 drive group and still maintain data integrity.
60	Provides data redundancy using distributed parity across spanned RAID 6 drive groups. RAID 60 can sustain two drive failures per RAID 6 drive group and still maintain data integrity. It provides the highest level of protection against drive failures of all of the RAID levels. RAID 60 includes both parity and disk striping across multiple drives. If a drive fails, the RAID controller uses the parity data to recreate all missing information.

2.11.2 Maximizing Performance

A RAID disk subsystem improves I/O performance. The RAID drive group appears to the host computer as a single storage unit or as multiple virtual units. I/O is faster because drives can be accessed simultaneously. <hyperactive>Table 2.12 describes the performance for each RAID level.

Table 2.12 RAID Levels and Performance

RAID Level	Performance
0	RAID 0 (striping) offers excellent performance. RAID 0 breaks up data into smaller blocks, then writes a block to each drive in the drive group. Disk striping writes data across multiple drives instead of just one drive. It involves partitioning each drive storage space into stripes that can vary in size from 8 KB to 128 KB. These stripes are interleaved in a repeated sequential manner. Disk striping enhances performance because multiple drives are accessed simultaneously.
1	With RAID 1 (mirroring), each drive in the system must be duplicated, which requires more time and resources than striping. Performance is impaired during drive rebuilds.
5	RAID 5 provides high data throughput, especially for large files. Use this RAID level for any application that requires high read request rates, but low write request rates, such as transaction processing applications, because each drive can read and write independently. Since each drive contains both data and parity, numerous writes can take place concurrently. In addition, robust caching algorithms and hardware based exclusive-or assist make RAID 5 performance exceptional in many different environments. Parity generation can slow the write process, making write performance significantly lower for RAID 5 than for RAID 0 or RAID 1. Drive performance is reduced when a drive is being rebuilt. Clustering can also reduce drive performance. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.
6	RAID 6 works best when used with data that requires high reliability, high request rates, and high data transfer. It provides high data throughput, data redundancy, and very good performance. However, RAID 6 is not well suited to tasks requiring a lot of writes. A RAID 6 virtual drive has to generate two sets of parity data for each write operation, which results in a significant decrease in performance during writes. Drive performance is reduced during a drive rebuild. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.

Table 2.12 RAID Levels and Performance (Continued)

RAID Level	Performance
10	RAID 10 works best for data storage that need the enhanced I/O performance of RAID 0 (striped drive groups), which provides high data transfer rates. Spanning increases the capacity of the virtual drive and improves performance by doubling the number of spindles. The system performance improves as the number of spans increases. (The maximum number of spans is eight.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans and RAID performance degrades to that of a RAID 1 or RAID 5 drive group.
50	RAID 50 works best when used with data that requires high reliability, high request rates, and high data transfer. It provides high data throughput, data redundancy, and very good performance. Spanning increases the capacity of the virtual drive and improves performance by doubling the number of spindles. The system performance improves as the number of spans increases. (The maximum number of spans is eight.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans and RAID performance degrades to that of a RAID 1 or RAID 5 drive group.
60	RAID 60 works best when used with data that requires high reliability, high request rates, and high data transfer. It provides high data throughput, data redundancy, and very good performance. Spanning increases the capacity of the virtual drive and improves performance by doubling the number of spindles. The system performance improves as the number of spans increases. (The maximum number of spans is eight.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans and RAID performance degrades to that of a RAID 1 or RAID 6 drive group. RAID 60 is not well suited to tasks requiring a lot of writes. A RAID 60 virtual drive has to generate two sets of parity data for each write operation, which results in a significant decrease in performance during writes. Drive performance is reduced during a drive rebuild. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.

2.11.3 Maximizing Storage Capacity

Storage capacity is an important factor when selecting a RAID level. There are several variables to consider. Striping alone (RAID 0) requires less storage space than mirrored data (RAID 1) or distributed parity (RAID 5 and RAID 6). RAID 5, which provides redundancy for one drive failure without duplicating the contents of entire drives, requires less space than RAID 1. <hyperactive>Table 2.13 explains the effects of the RAID levels on storage capacity.

Table 2.13 RAID Levels and Capacity

RAID Level	Capacity
0	RAID 0 (striping) involves partitioning each drive storage space into stripes that can vary in size. The combined storage space is composed of stripes from each drive. RAID 0 provides maximum storage capacity for a given set of drives.
1	With RAID 1 (mirroring), data written to one drive is simultaneously written to another drive, which doubles the required data storage capacity. This is expensive because each drive in the system must be duplicated.
5	RAID 5 provides redundancy for one drive failure without duplicating the contents of entire drives. RAID 5 breaks up data into smaller blocks, calculates parity by performing an exclusive-or on the blocks, then writes the blocks of data and parity to each drive in the drive group. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set.
6	RAID 6 provides redundancy for two drive failures without duplicating the contents of entire drives. However, it requires extra capacity because it uses two parity blocks per stripe. This makes RAID 60 more expensive to implement.

Table 2.13 RAID Levels and Capacity (Continued)

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

2.12 RAID Availability

2.12.1 RAID Availability Concept

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

2.12.1.1 Spare Drives

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

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

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



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

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

2.12.1.2 Rebuilding

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

2.13 Configuration Planning

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

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

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

2.13.1 Number of Drives

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

2.13.2 Drive Group Purpose

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

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

Fill out <hyperactive>Table 2.14 to help you plan the drive group configuration. Rank the requirements for your drive group, such as storage space and data redundancy, in order of importance, and then review the suggested RAID levels.

Table 2.14 Factors to Consider for Drive Group Configuration

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

Chapter 3: Self Encrypting Disk

3.1 Overview

The SED feature offers the ability to encrypt data on disk and use disk-based key management to provide data security. With the SED feature, data is encrypted by the drives. You can designate which data to encrypt at the individual virtual drive (VD) level.

This solution provides data protection in the event of theft or loss of physical drives. With self-encrypting disks, if you remove a drive from its storage system or the server it is housed in, the data on that drive is encrypted and useless to anyone who attempts to access without the appropriate security authorization.

Any encryption solution requires management of the encryption keys. The security feature provides a way to manage these keys. You can change the encryption key for all ServeRAID controllers that are connected to SED drives. All SED drives, whether locked or unlocked, always have an encryption key. This key is set by the drive and is always active. When the drive is unlocked, the data to host from the drive (on reads) and from the host to the drive cache (on writes) is always provided. However, when resting on the drive platters, the data is always encrypted by the drive.

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

The WebBIOS Configuration Utility ([Section Figure 4.5, “WebBIOS Virtual Drive Definition Screen”](#)) and MegaRAID Storage Manager ([Chapter 8, “Monitoring System Events and Storage Devices,”](#)) offer procedures that you can use to manage the security settings for the drives.

3.2 Purpose

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

3.3 Terminology

[Table 3.1](#) describes the terminology related to the SED feature.

Table 3.1 SED Terminology

Option	Description
Authenticated Mode	The RAID configuration is keyed to a user passphrase. The passphrase must be provided on system boot to authenticate the user and facilitate unlocking the configuration for user access to the encrypted data.
Blob	A blob is created by encrypting a key(s) using another key. There are two types of blob in the system – encryption key blob and security key blob.
Key backup	You need to provide the controller with a lock key if the controller is replaced or if you choose to migrate secure virtual drives. To do this, you must back up the security key.
Passphrase	An optional authenticated mode is supported in which you must provide a passphrase on each boot to make sure the system boots only if the user is authenticated. Firmware uses the user passphrase to encrypt the security key in the security key blob stored on the controller.

Table 3.1 SED Terminology (Continued)

Option	Description
Re-provisioning	Re-provisioning disables the security system of a device. For a controller, it involves destroying the security key. For self-encrypting drives, when the drive lock key is deleted, the drive is unlocked and any user data on the drive is securely deleted. This does not apply to controller-encrypted drives, because deleting the virtual drive destroys the encryption keys and causes a secure erase. See <hyperactive>Section 3.5, “Instant Secure Erase” for information about the instant secure erase feature.
Security Key	A key based on a user-provided string. The controller uses the security key to lock and unlock access to the secure user data. This key is encrypted into the security key blob and stored on the controller. If the security key is unavailable, user data is irretrievably lost. You must take all precautions to never lose the security key.
Un-Authenticated Mode	This mode allows controller to boot and unlock access to user configuration without user intervention. In this mode, the security key is encrypted into a security key blob, stored on the controller, but instead of a user passphrase, an internal key specific to the controller is used to create the security key blob.
Volume Encryption Keys (VEK)	The controller uses the Volume Encryption Keys to encrypt data when a controller-encrypted virtual drive is created. These keys are not available to the user. The firmware (FW) uses a unique 512-bit key for each virtual drive. The VEK for the VDs are stored on the physical drives in a VEK blob.

3.4 Workflow

3.4.1 Enable Security

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

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

- Create the security key identifier
- Create the security key
- Create a pass phrase (optional)

See <hyperactive>Section 4.5, “Selecting Self-Encrypting Disk Security Options” for the procedures used to enable security in WebBIOS or <hyperactive>Section 8.1, “Monitoring System Events,” for the procedures used to enable security in MegaRAID Storage Manager.

3.4.1.1 Create the Security Key Identifier

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

3.4.1.2 Create the Security Key

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



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

3.4.1.3 Create a Passphrase (Optional)

The pass phrase provides additional security. The pass phrase should be different from the security key. If you choose this option, you must enter it whenever you boot your server.



NOTE If you forget the pass phrase, you will lose access to your data.

When you use the specified security key identifier, security key, and pass phrase, security will be enabled on the controller.

3.4.2 Change Security

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

There are three procedures you can perform to change the security settings on the controller:

- Change the security key identifier
- Change the security key
- Change a pass phrase

See <hyperactive>Section 4.5, “Selecting Self-Encrypting Disk Security Options” for the procedures used to change security options in WebBIOS or <hyperactive>Section 7.2, “Selecting Self-Encrypting Disk Security Options” for the procedures used to change security options in MegaRAID Storage Manager.

3.4.2.1 Change the Security Key Identifier

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

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

3.4.2.2 Change the Security Key

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

3.4.2.3 Add or Change the Pass Phrase

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

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

3.4.2.4 Create Secure Virtual Drives

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

3.4.2.5 Simple Configuration

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

See <hyperactive>Section 7.1.2, “Creating a Virtual Drive Using Simple Configuration” for the procedures used to select the redundancy type and drive security method for a configuration.

3.4.2.6 Advanced Configuration

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

See <hyperactive>Section 7.1.3, “Creating a Virtual Drive Using Advanced Configuration” for the procedures used to import a foreign configuration.

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

3.4.3 Import a Foreign Configuration

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

See <hyperactive>Section 4.5.4, “Importing Foreign Configurations” for the procedure used to import a foreign configuration in WebBIOS or <hyperactive>Section 7.2.4, “Importing or Clearing a Foreign Configuration” for the procedure in MegaRAID Storage Manager.

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

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

3.5 Instant Secure Erase

The Instant Secure Erase feature offers a way to erase data that you can use with SED drives. After the initial investment into a SED drive, there is no additional cost in dollars or time to erase data using the Instant Secure Erase feature.

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 SED over other technologies that exists today, both in terms of the security provided and time saved.

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

There are three major reasons for using instant secure erase.

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

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

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

Chapter 4: Starting the HII Configuration Utility

This chapter explains how to use this the Human Interface Infrastructure (HII) Configuration Utility. HII configuration utility is a tool for configuring controllers, physical disks, and virtual disks, and for performing other configuration tasks in a pre-boot, Unified Extensible Firmware Interface (UEFI) environment.

Follow these steps to start the utility and to access the main configuration menu.

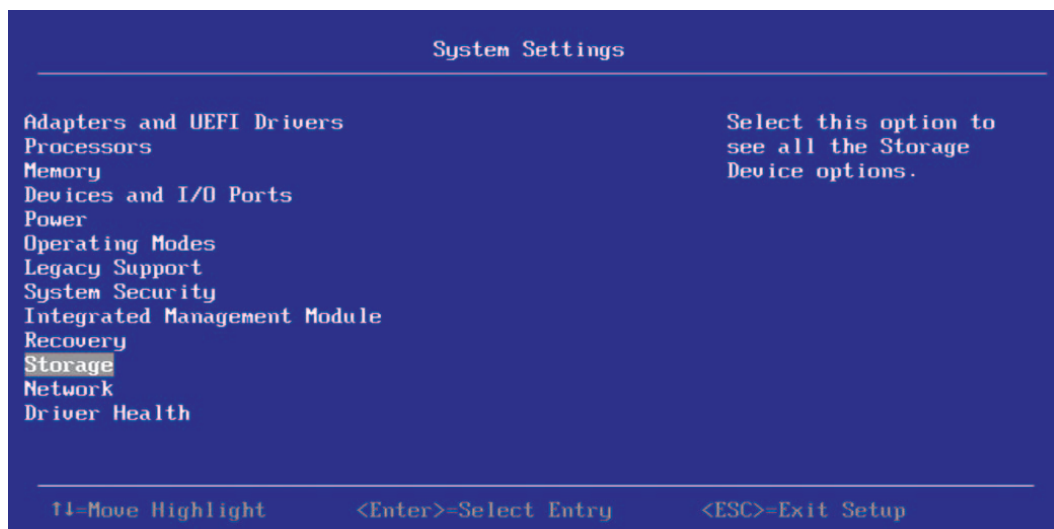
1. Boot the computer and press the appropriate key to start the setup utility during bootup.



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.

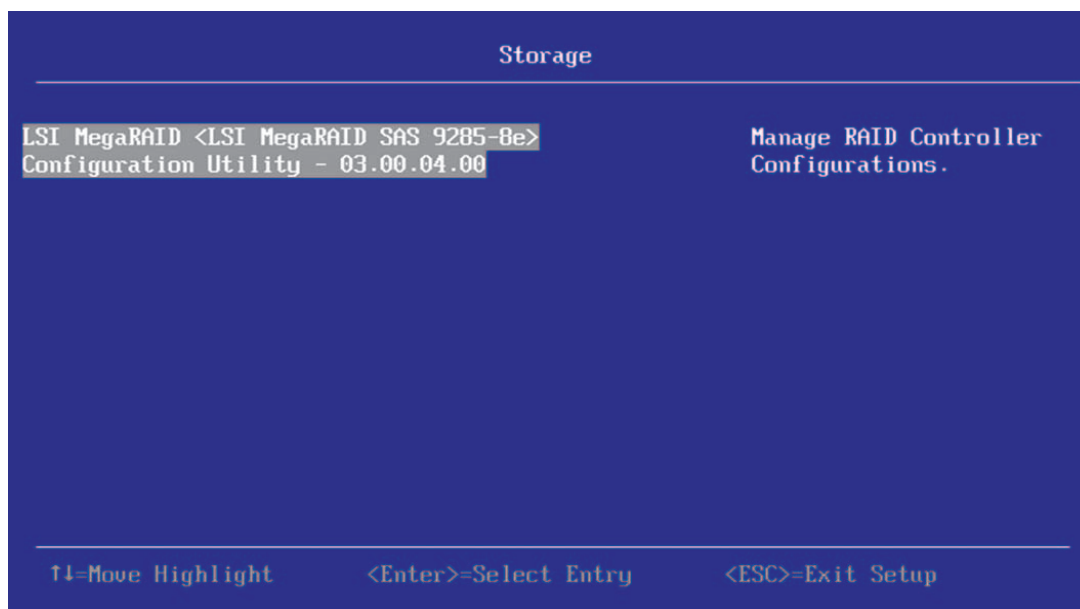
2. When the initial window appears, press **System Settings**. The following window appears.

Figure 4.1 System Settings Window



3. Highlight **Storage** and press Enter. The **Controller Selection** menu appears. The following figure shows the a sample **Controller Selection** menu.

Figure 4.2 Controller Selection Menu

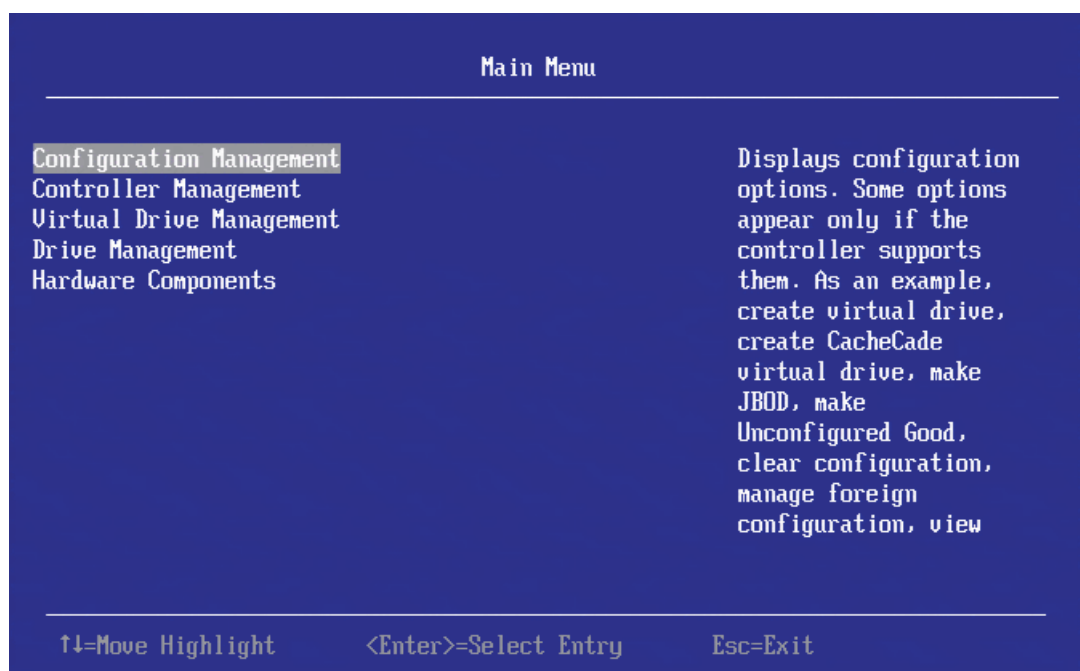


This window lists the RAID controllers installed in your computer system. Use the PCI slot number to differentiate between controllers of the same type.

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

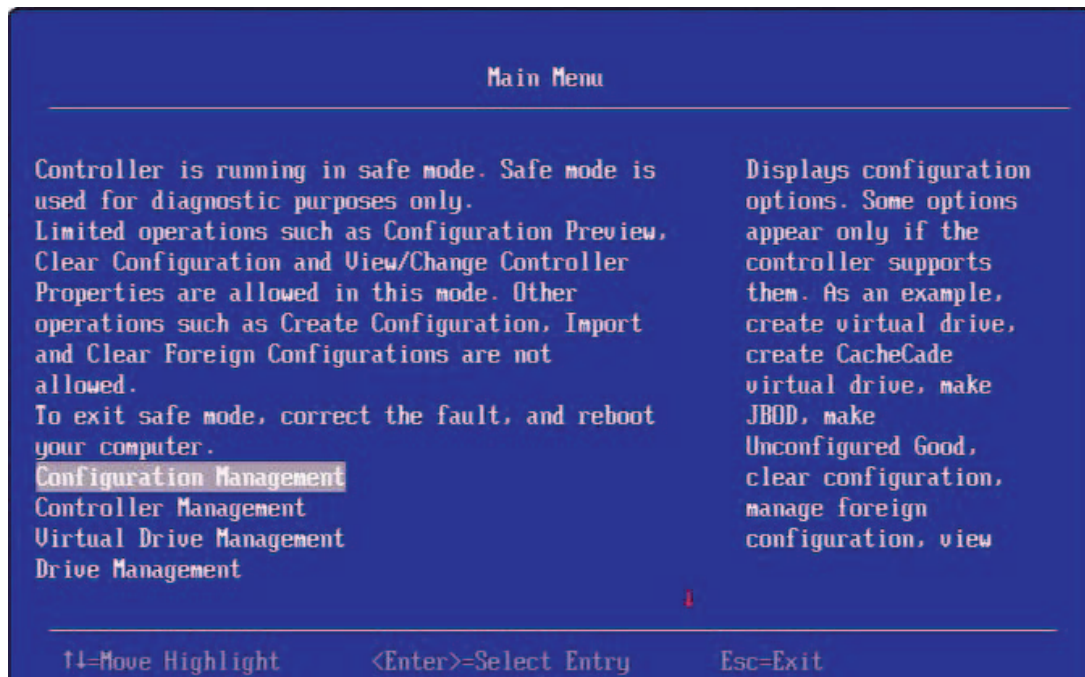
The **Main Menu** appears, as shown in the following figure.

Figure 4.3 Main Menu



When the controller is running in Safe Mode the main menu includes the warning message as shown in the following figure.

Figure 4.4 Main Menu – Safe Mode



5. Select one of the following menu options:
 - **Configuration Management** to perform tasks such as creating virtual drives, viewing drive group properties, viewing hotspare information, and clearing a configuration. See [Chapter 5, Managing Configurations](#).
 - **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. See [Chapter 6, Managing Controllers](#).
 - **Virtual Drive Management** to perform tasks such as viewing virtual drive properties, locating virtual drives, and running a consistency check. See [Chapter 7, Managing Virtual Drives](#).
 - **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. See [Chapter 8, Managing Physical Drives](#).
 - **Hardware Components** to view battery properties, manage batteries, and manage enclosures. See [Chapter 9, Managing Hardware Components](#).

Chapter 5: Managing Configurations

When you select **Configuration Management** on the **Main Menu**, the **Configuration Management** menu appears, as shown in the following figure.

Figure 5.1 Configuration Management Menu



The following sections explain each configuration management option.



NOTE Only the first two menu options appear if no virtual drives have been created on this controller. The **Make JBOD** and **Make Unconfigured Good** options are included for some controllers. (See Section 5.7, [Make Unconfigured Good and Make JBOD](#).) The Manage Foreign Configuration option is included for some configurations. (See Section 5.8, [Managing Foreign Configurations](#).)

5.1 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 Virtual Drive** from the **Configuration Management** menu.
3. Select **Generic RAID 0** from the **Create Virtual Drive** menu.
4. Choose an option from the **Drive Selection Criteria** field (if there is more than one option).
5. Select **Save Configuration** to create the chose profile.

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

Figure 5.2 Create Virtual Drive Window

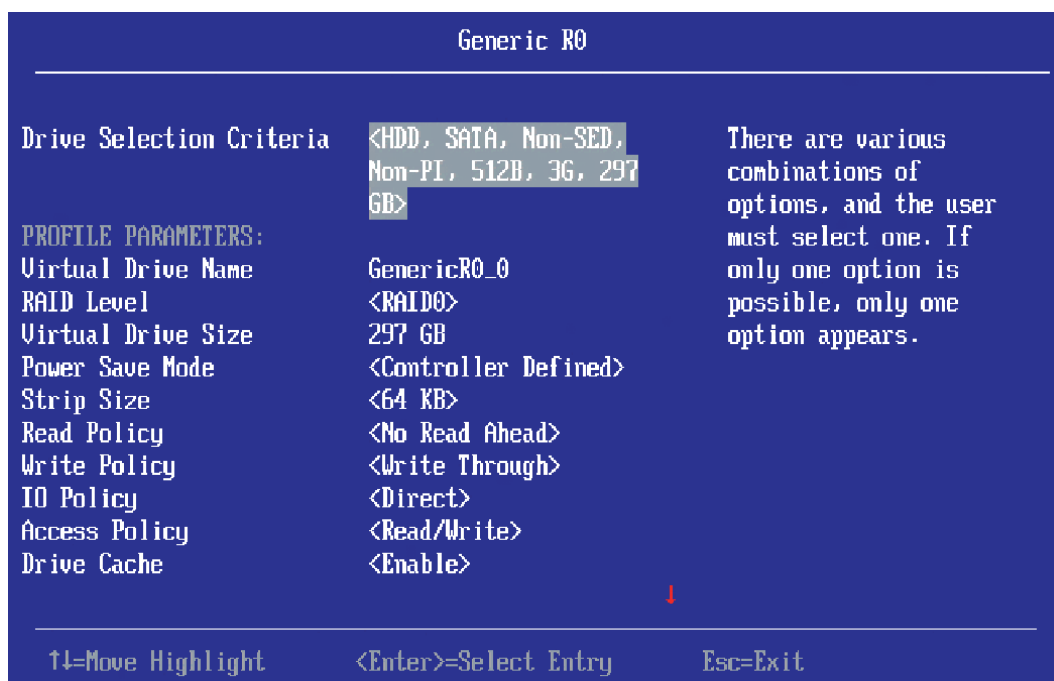


Select a RAID level profile for the virtual drive from the list.

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 5.3](#) for brief descriptions of the RAID levels.

The following window appears after selecting a RAID level profile.

Figure 5.3 Generic R0 Window



The small red arrow at the bottom of the window 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.

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

Table 5.1 Virtual Drive Creation Profile Options

Option	Description
Drive Selection Criteria	There are various combinations of options, and the user must select one. 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, <i>RAID 0</i> is displayed.
Virtual Drive Size	Displays the amount of virtual drive storage space. By default, it displays the maximum capacity available for the virtual drive.
Power Save Mode	Displays the Power Save Mode out of the five available options: <i>None</i> , <i>Auto</i> , <i>Max</i> , <i>Max without Cache</i> , and <i>Controller Defined</i> .
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: <i>8 KB</i> , <i>16 KB</i> , <i>32 KB</i> , <i>64 KB</i> , <i>128 KB</i> , <i>256 KB</i> , <i>512 KB</i> , <i>1 MB</i> .

Table 5.1 Virtual Drive Creation Profile Options (Continued)

Option	Description
Read Policy	<p>Displays the read cache policy for the virtual drive. For any profile, if the drive is an SSD drive, the <i>No Read Ahead</i> option is displayed. Otherwise, the <i>Default</i> option is displayed. The possible options are:</p> <ul style="list-style-type: none"> ■ <i>Default</i> – A virtual drive property that indicates whether the default read policy is <i>Read Ahead</i> or <i>No Read Ahead</i>. ■ <i>Read Ahead</i> – Allows the controller to read requested data and store the additional data in cache memory, anticipating that the data is required soon. ■ <i>No Read Ahead</i> – Specifies that the controller does not use <i>Read Ahead</i> for the current virtual drive.
Write Policy	<p>Displays the write cache policy for the virtual drive. For any profile, if the drive is an SSD drive, the <i>Write Through</i> option is displayed. Otherwise, the <i>Always Write Back</i> option is displayed. The possible options are:</p> <ul style="list-style-type: none"> ■ <i>Default</i> – A virtual drive property that indicates whether the default write policy is <i>Write Through</i> or <i>Write Back</i>. ■ <i>Write Through</i> – Eliminates the risk of losing cached data in case of power failure. However, it might result in slower performance. ■ <i>Write Back with BBU</i> – In <i>Write-Back Caching</i> 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. ■ <i>Force Write Back</i> – A data transfer completion signal is sent to the host when the controller cache has received all of the data in a transaction.
I/O Policy	<p>Displays the Input/Output policy for the virtual drive. For any profile, if the drive is an SSD drive, the <i>Direct</i> option is displayed. The possible options are:</p> <ul style="list-style-type: none"> ■ <i>Default</i> – A virtual drive property that indicates whether the default I/O policy is <i>Direct IO</i> or <i>Cached IO</i>. ■ <i>Direct IO</i> – Data reads 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.) ■ <i>Cached IO</i> – All reads are buffered in cache.
Access Policy	The access policy for the virtual drive. The options are <i>Read/Write</i> and <i>Read Only</i> .
Disk Cache Policy	Displays the virtual drive cache setting. The possible options are <i>Unchanged</i> , <i>Enable</i> , and <i>Disable</i> .
Default Initialization	<p>Displays the virtual drive initialization setting. The <i>Default Initialization</i> displays the following options:</p> <ul style="list-style-type: none"> ■ <i>No</i> – Do not initialize the virtual drive. ■ <i>Fast</i> – Initializes the first 100 MB on the virtual drive. ■ <i>Full</i> – Initializes the entire virtual drive.
Save Configuration	Saves the configuration that the wizard created.

5.2 Manually Creating a Virtual Drive

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

Figure 5.4 Create Configuration Window

Create Configuration

Save Configuration		Submits the changes made to the entire form and creates a virtual drive with the specified parameters.
Select RAID Level	<RAID0>	
Secure Virtual Drive	[]	
Protect Virtual Drive	[]	
Select Drives From	<Unconfigured Capacity>	
Select Drives		
CONFIGURE VIRTUAL DRIVE PARAMETERS:		
Virtual Drive Name	-	
Virtual Drive Size Unit	<GB>	
Strip Size	<128 KB>	
Read Policy	<Adaptive>	
Write Policy	<Write Back>	
I/O Policy	<Direct>	

↓

↑↓=Move Highlight <Enter>=Select Entry Esc=Exit

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



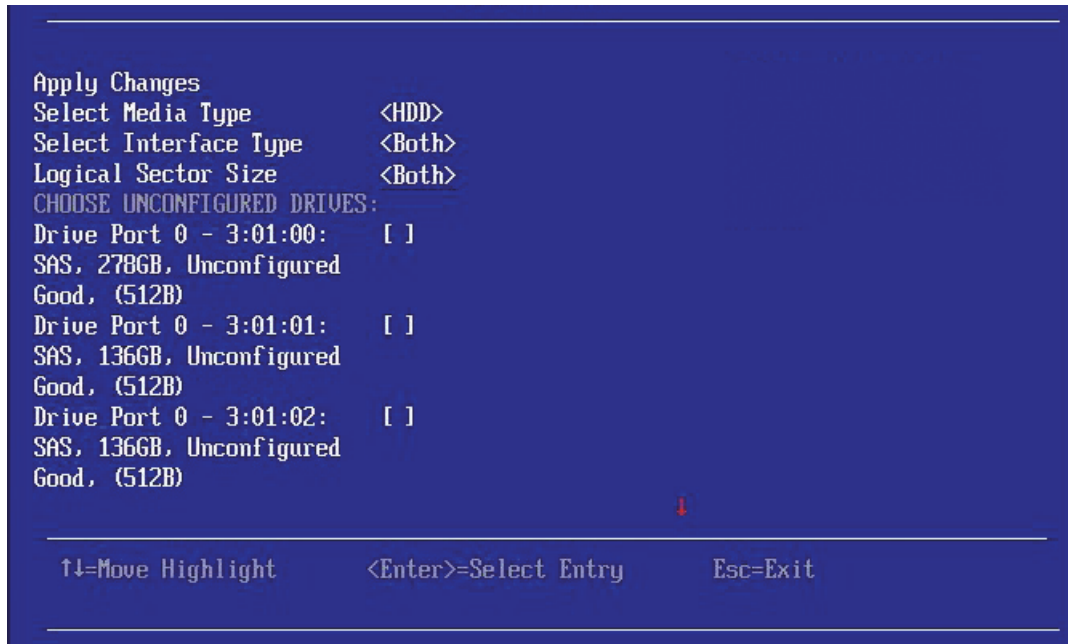
NOTE The red arrow appears when there is too much information to display in one window. The amount of information that can be displayed in one window depends on the capabilities of the HII browser.

Follow 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 popup 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 5.3](#) for brief descriptions of the RAID levels.
3. To view the **Secure Virtual Drive** field, enable security and an FDE drive must be attached. If either is missing, the field is grayed out.
4. To view the **Protect Virtual Drive** field, enable protection and a protected drive must be attached. If either is missing, the field is grayed out.
5. If the security key is enabled, highlight the **Protect 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** field to protect the new virtual drive with a password.
This field is not available unless protection is already enabled on 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 popup 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 5.2](#) for descriptions of these options.
11. Highlight **Select Drives** and press Enter. The following window appears.

Figure 5.5 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 popup menu.
The choices are *HDD*, *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 popup 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. (Optional) Change the Default **Logical Sector Size** by highlighting this field, pressing enter, and selecting an option from the popup menu.
This field allows you to choose between a 512 byte and 4K logical sector size. To use 4K logical sector size, the physical drives that are being configured must be 4K compatible (Advanced Format Drives). Otherwise, if 4K is selected as the desired sector size, no drives will display in the drive list. 512 byte sector size is supported by all drives.
4. 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 window. The amount of information that can be displayed in one window 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 5.3](#) for more information.

5. When you have selected all 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 displays a message warning you that the remaining free capacity on the larger drives will be unusable.

6. 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** window.

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

A message appears confirming that the configuration is being created.

8. 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 5.2 Virtual Drive Policies

Property	Description
Strip Size	The virtual drive strip size per DDF. The possible values are as follows: 0: 512 bytes 1: 1 KB 2: 2 KB 3: 4 KB 4: 8 KB ... 8: 1 MB
Logical Sector Size	The size of logical sector blocks used by the virtual drive when writing data. The choices are 4 K and 512 bytes. NOTE To use 4K as the logical sector block size, the drives used by the controller must be 4k (also known as Advanced Format) compatible. If the drives are not compatible, they will not appear as selections when choosing drives, or the 4K option will not be selectable.
Read Policy	The read cache policy for the virtual drive. The possible values are as follows: <ul style="list-style-type: none"> ■ <i>Ahead</i>: 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 there is little improvement when accessing random data. ■ <i>Normal</i>: Read-ahead capability is disabled.
Write Policy	The write cache policy for the virtual drive. The possible values are as follows: <ul style="list-style-type: none"> ■ <i>Write-through (WThru)</i>: 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. ■ <i>Write-back (WBack)</i>: 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. ■ <i>Force Write Back</i>.

Table 5.2 Virtual Drive Policies (Continued)

Property	Description
I/O Policy	The I/O policy for the virtual drive. The possible values are as follows: <ul style="list-style-type: none"> ■ <i>Direct</i>: 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.) ■ <i>Cached</i>: All reads are buffered in cache.
Access Policy	The access policy for the virtual drive. The options are <i>Read/Write</i> , <i>Read Only</i> , and <i>Blocked</i> .
Drive Cache	The disk cache policy for the virtual drive. The possible values are <i>Unchanged</i> , <i>Enable</i> , and <i>Disable</i> .
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 <i>No</i> , <i>Fast</i> , and <i>Slow</i> .

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.

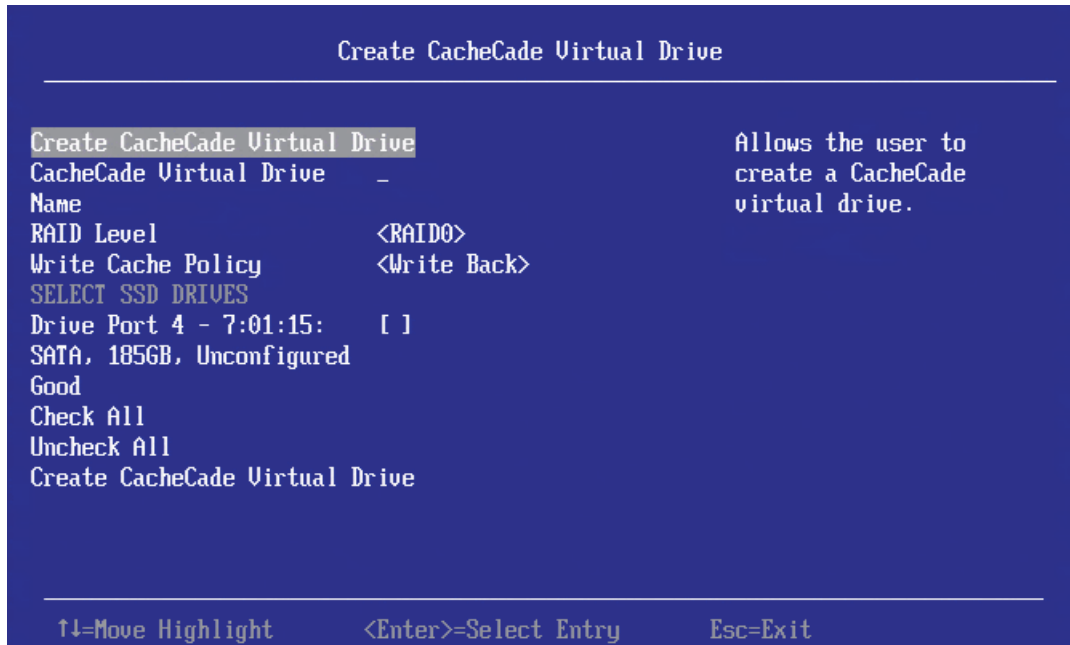
Table 5.3 RAID Levels

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 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.

5.3 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 5.6 Create CacheCade Virtual Drive Window



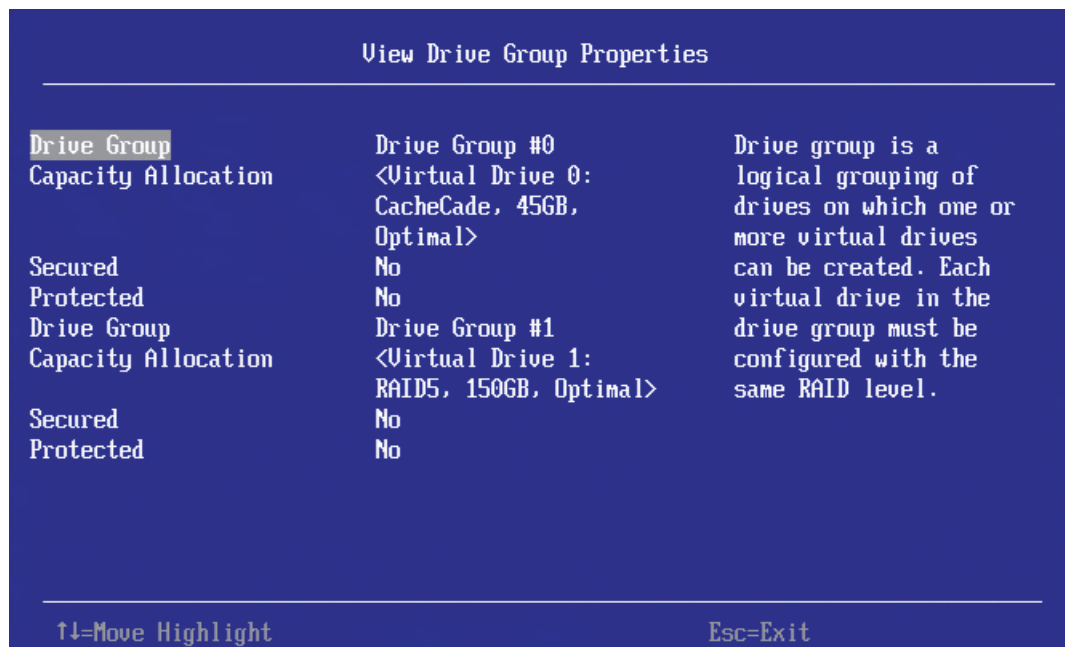
Follow these steps to create a CacheCade virtual drive.

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 popup 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 5.3](#) 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 Viewing Drive Group Properties

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

Figure 5.7 View Drive Group Properties Window



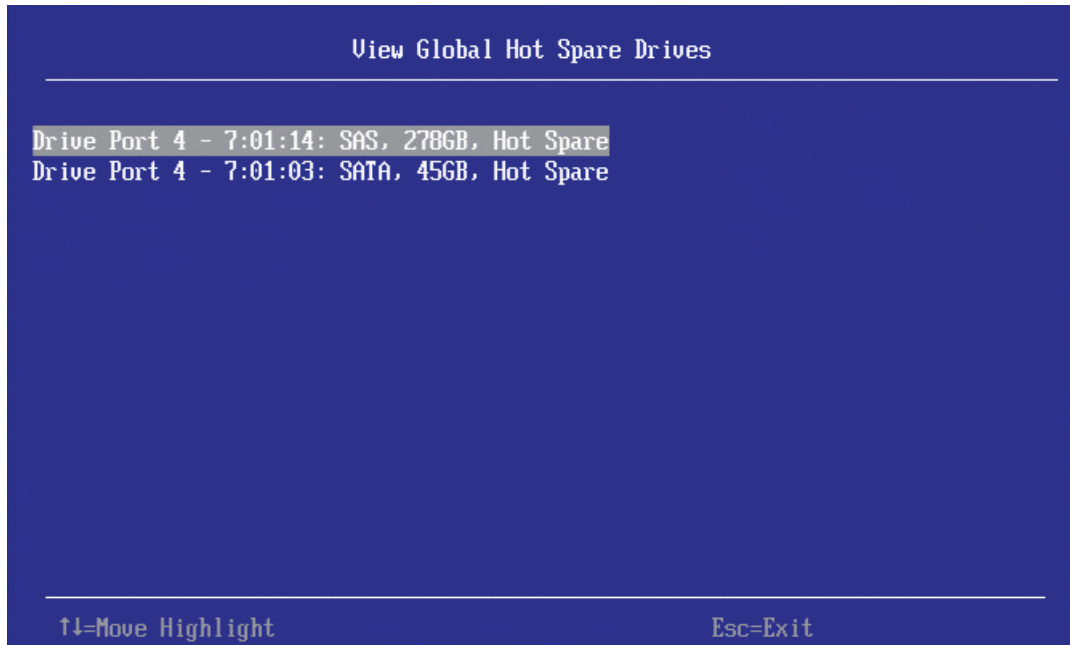
A drive group is a logical grouping of drives attached to a RAID controller on which one or more virtual drives can be created. Each virtual drive in the drive group must be configured with the same RAID level. This figure shows information for one drive group.

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 popup window.

5.5 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** window.

Figure 5.8 View Global Hot Spare Drives Window

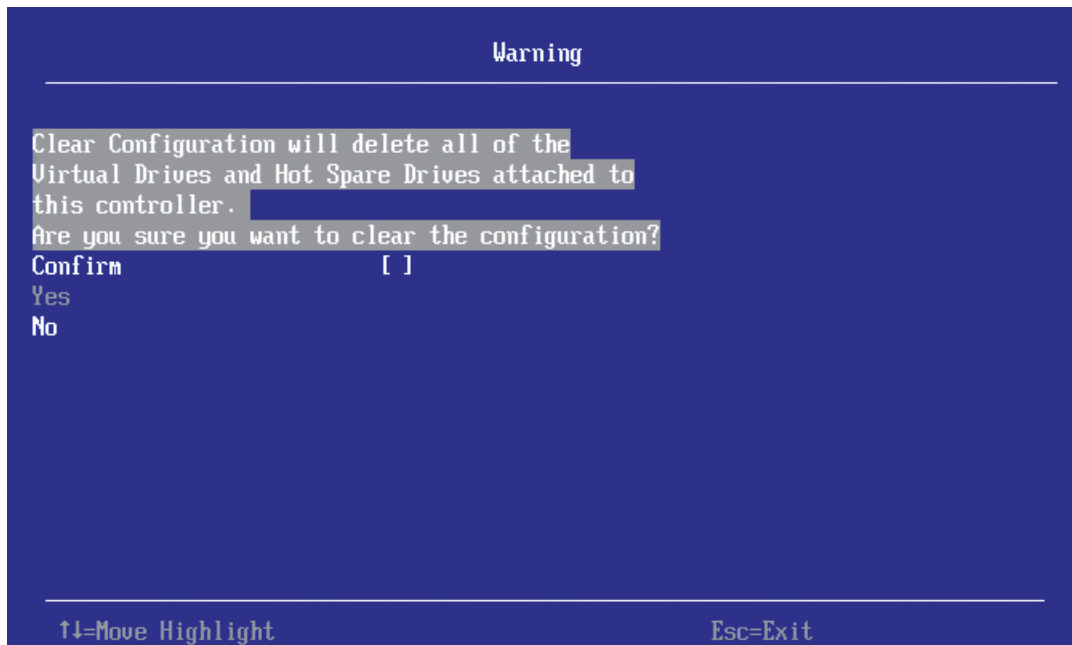


Press Esc to exit this window when you are finished viewing information.

5.6 Clearing a Configuration

The following warning appears when you select **Clear Configuration** from the **Configuration Management** menu.

Figure 5.9 Clear Configuration Warning



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.

5.7 Make Unconfigured Good and Make JBOD

When you power down a ServeRAID system 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 Drives) when you power the system again. If the drive does contain 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 Unconfigured Good.



NOTE iMR controllers support JBOD drives, but ServeRAID controllers do not.

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.



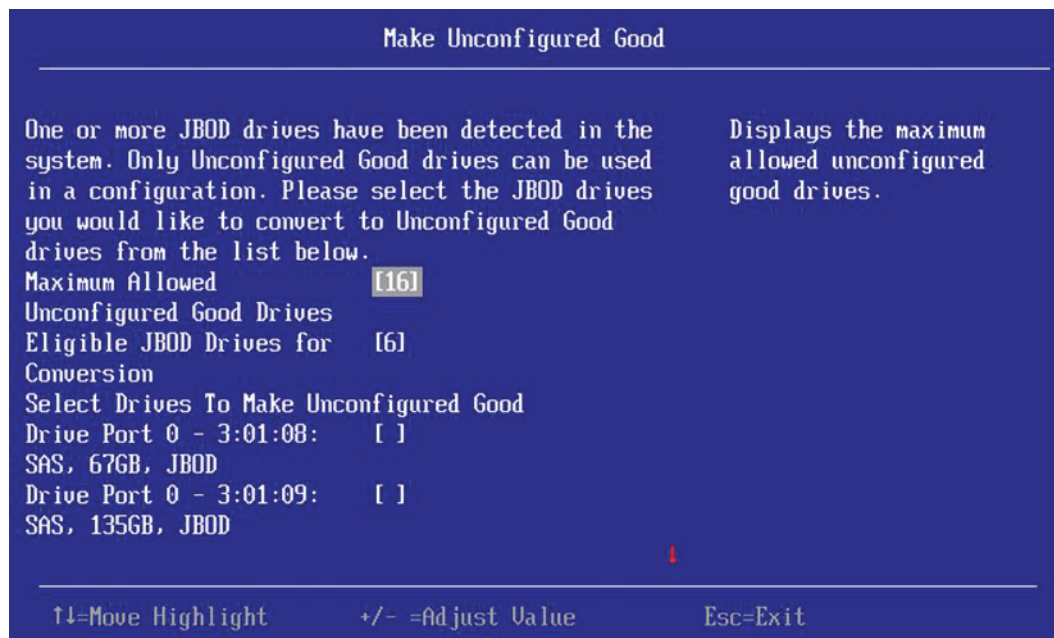
NOTE If the controller supports JBOD drives, you can also change the status of JBOD drives to Unconfigured Good when you create a new configuration using the **Create Configuration** option.

5.7.1 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 window appears, listing information about the JBOD drives currently connected to the controller.

Figure 5.10 Make Unconfigured Good Window



Scroll down, if necessary, to view other drives listed in the window.

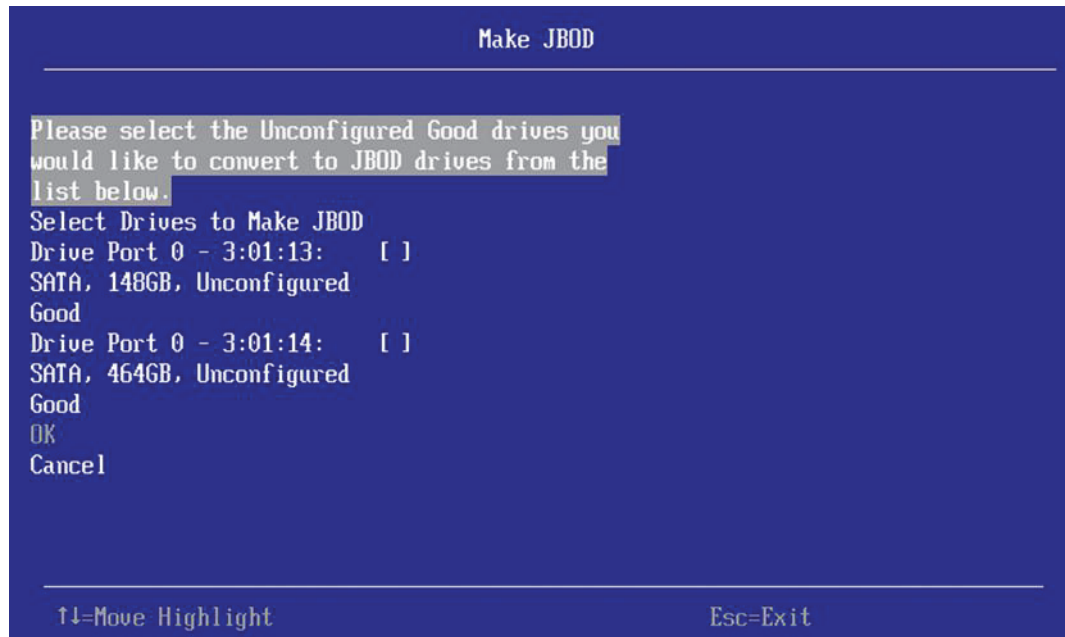
2. Highlight each JBOD drive you want to make Unconfigured Good and press the spacebar to select it.
3. Highlight **OK** (at the bottom of the JBOD drive list) and press Enter to convert the JBOD drives to Unconfigured Good status.

5.7.2 Make JBOD

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

1. Highlight **Make JBOD** on the **Configuration Management** menu and press Enter.
A window appears listing the Unconfigured Good drives currently connected to the controller.

Figure 5.11 Make JBOD Window

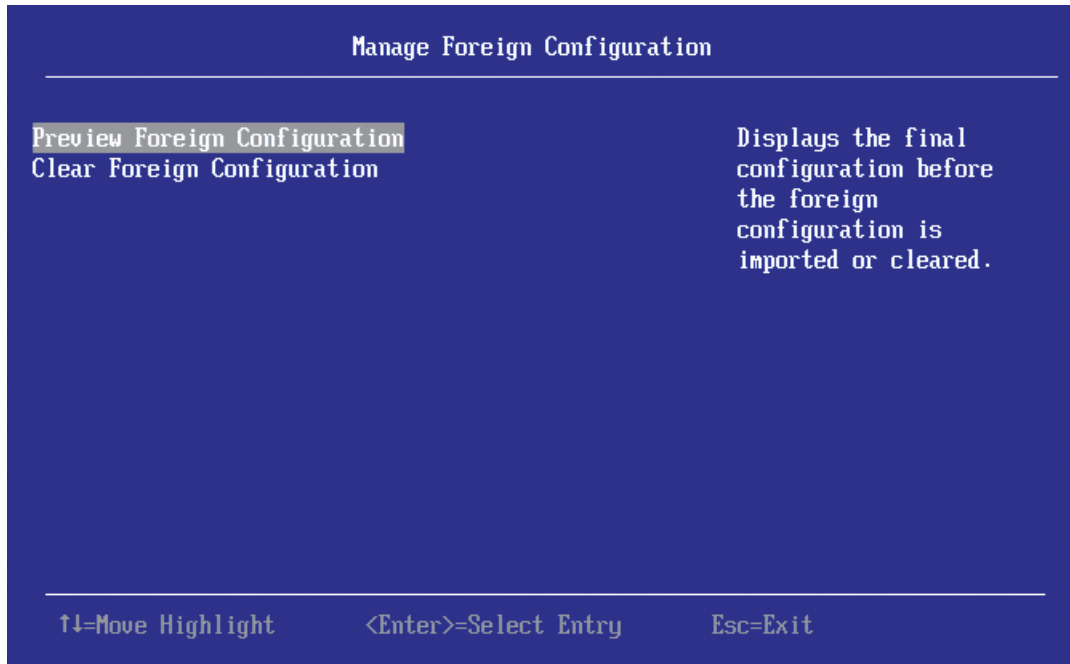


2. Highlight each drive you want to convert to JBOD status and press the spacebar to select it.
3. Highlight **OK** and press Enter to convert the Unconfigured Good drives to JBOD status.

5.8 Managing Foreign Configurations

The following window appears when you select **Manage Foreign Configuration** from the **Configuration Management** menu.

Figure 5.12 Manage Foreign Configuration Window



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.

5.8.1 Previewing and Importing a Foreign Configuration

You preview a foreign configuration prior to 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 need to 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 5.13 Enter Security Key Message

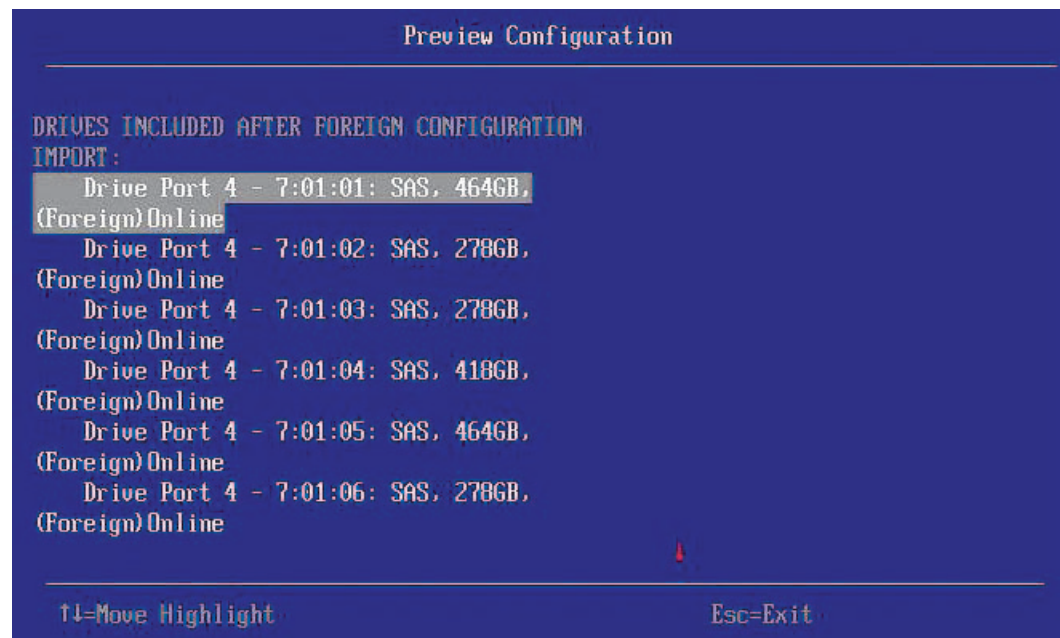


To successfully import the foreign configuration, follow the directions in the message.

Follow these steps to preview and import a foreign configuration.

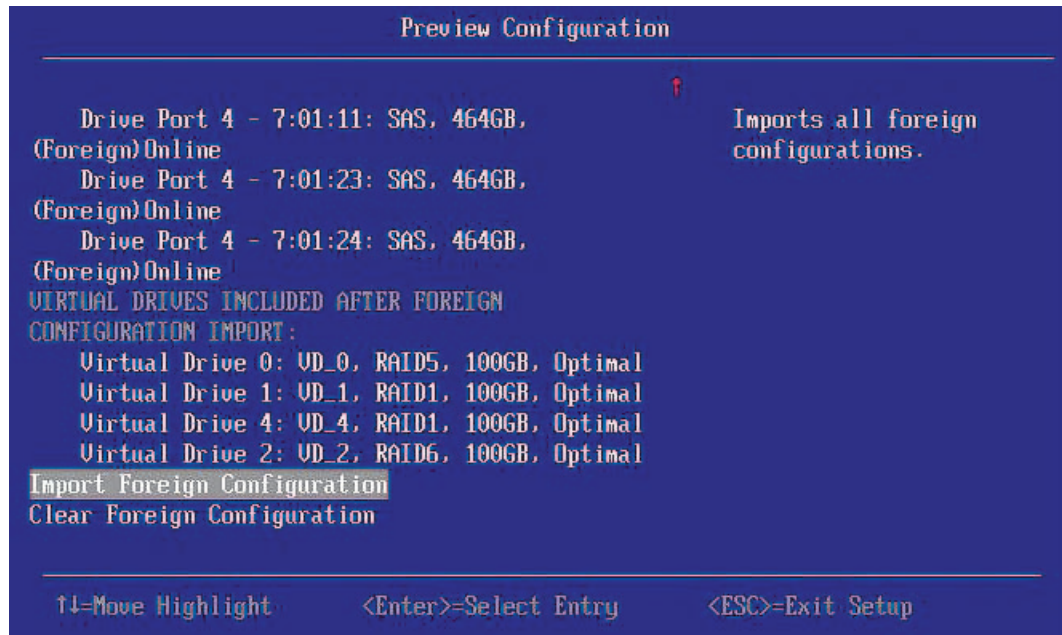
1. Highlight **Preview Foreign Configuration** on the **Manage Foreign Configuration** menu and press Enter.
The following window appears, listing information about the physical drives in the foreign configuration.

Figure 5.14 Preview Configuration Window 1



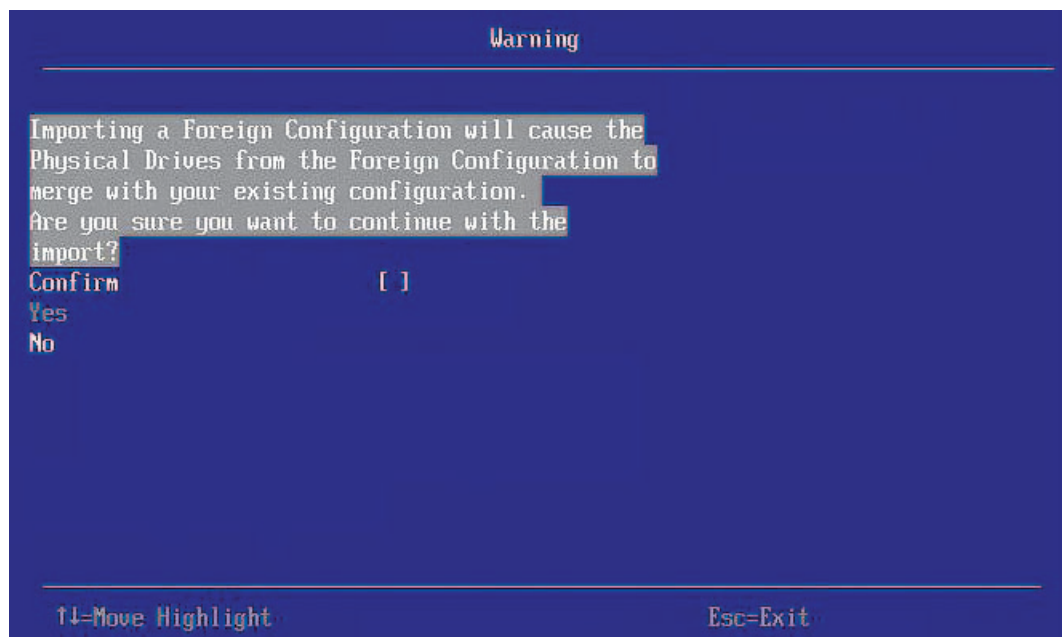
2. Scroll down, if needed, to view more information about the drives in the foreign configuration, as shown in the following window.

Figure 5.15 Preview Configuration Window 2



3. Review the information listed in the window.
4. Highlight **Import Foreign Configuration** and press Enter.
The following warning message appears.

Figure 5.16 Warning for Import Foreign Configuration



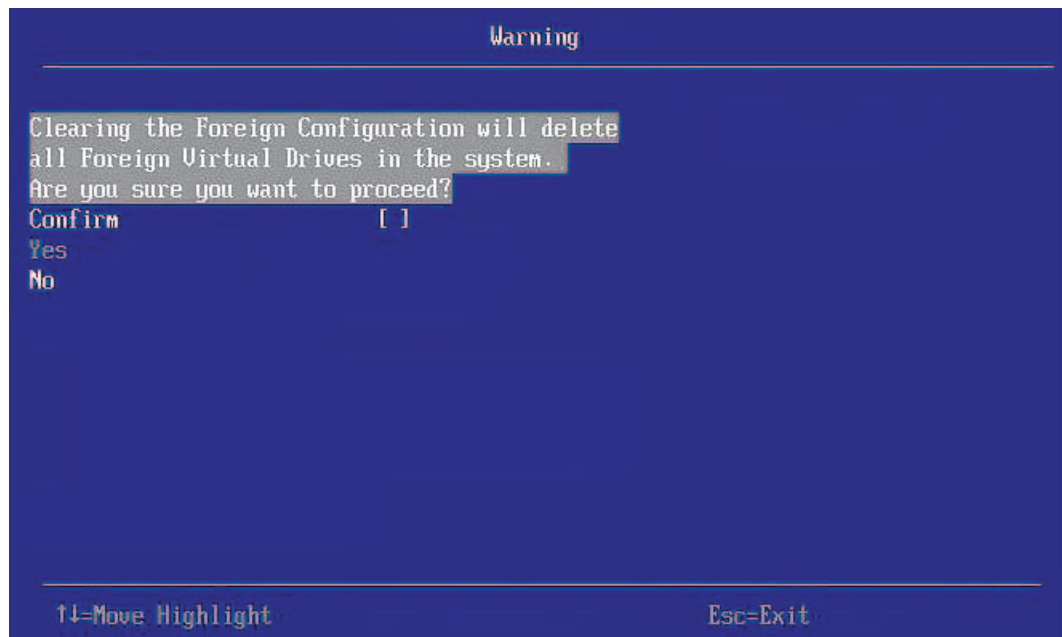
5. To confirm the import, highlight **Confirm** and press the spacebar.
6. Highlight **Yes** and press Enter.
The foreign configuration is imported.

5.8.2 Clearing a Foreign Configuration

Follow these steps to clear a foreign configuration.

1. Highlight **Clear Foreign Configuration** on the **Manage Foreign Configuration** menu and press Enter.
The following warning message appears.

Figure 5.17 Clear Foreign Configuration Warning



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 previewing the configuration, as shown in [Figure 5.15](#).

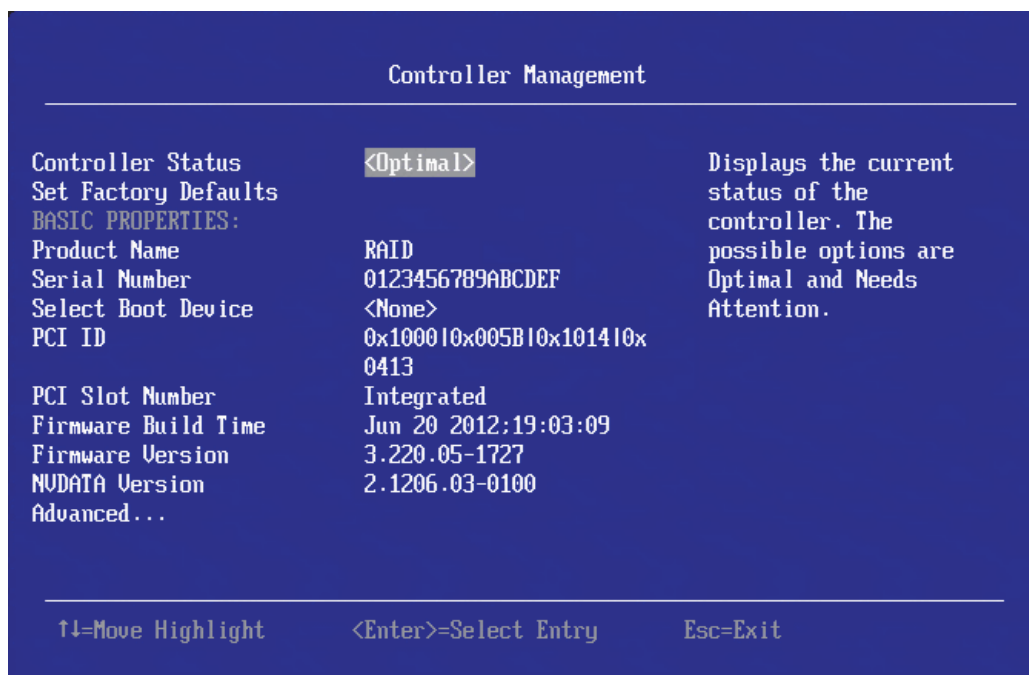
Chapter 6: Managing Controllers

When you select **Controller Management** on the **Main Menu**, the **Controller Management** menu appears, as shown in the following figure.

The top-level **Controller Management** window lists basic information about the attributes and condition of the controller.

- To reset the controller to its factory settings, highlight **Set Factory Defaults** and press Enter.
- To select a boot device, highlight **Select Boot Device**, press Enter, and select the boot device from the popup menu.
- To view advanced controller properties, highlight **Advanced** and press Enter.

Figure 6.1 Controller Management Menu



The **Controller Management** window lists the following basic controller properties:

Table 6.1 Basic Controller Properties

Property	Description
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: <i>Optimal</i> , if all components are operating normally. <i>Needs Attention</i> , if any component needs attention. <i>Safe Mode</i> , if the controller encountered critical errors. Most features are disabled and the controller requires user attention.
Product Name	The marketing name of the controller.
Serial Number	The serial number of the controller.
Select Boot Device	This field selects the primary boot device.
PCI ID	The PCI ID of the controller.

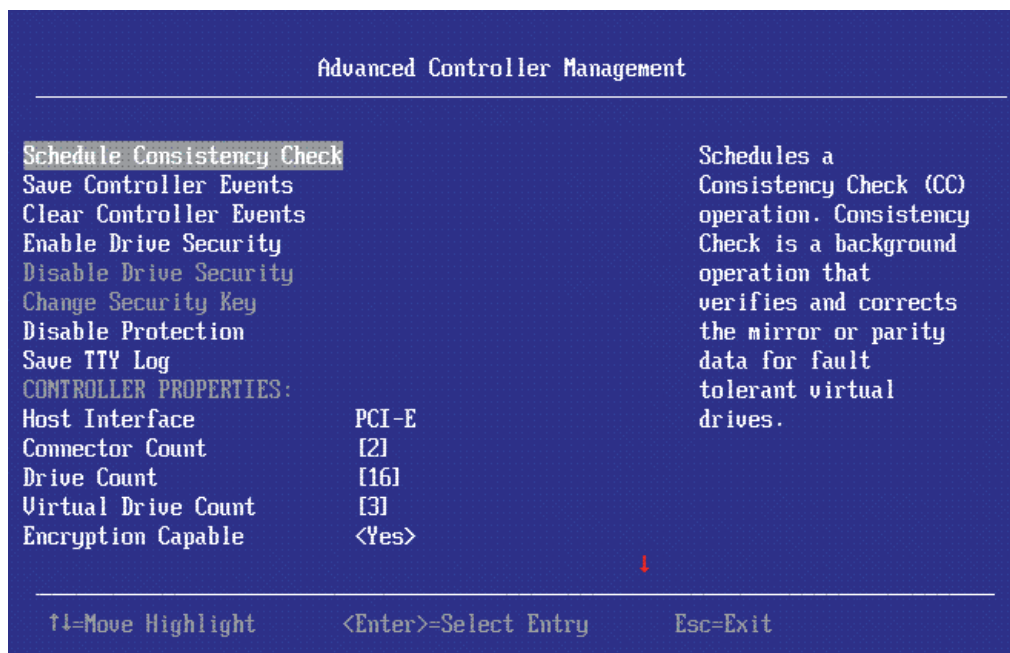
Table 6.1 Basic Controller Properties (Continued)

Property	Description
PCI Slot Number	The slot ID number of the PCI slot where the controller is installed.
Firmware Build Time	The build time of the controller firmware.
Firmware Version	The version number of the controller firmware.
NVDATA Version	The version number of the controller NVDATA.

6.1 Viewing Advanced Controller Management Options

The **Advanced Controller Management** window lists a number of controller properties and also includes options for performing various actions on the controller.

Figure 6.2 Advanced Controller Management



This window lists numerous properties that cannot all be displayed in one window. Scroll down to view all the options.



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.

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

1. Move the highlight to the value for any option and press Enter.
A popup menu of the available options appears.

2. Highlight the value you want and press Enter. For options like **Rebuild Rate** or **Reconstruction Rate** that require a number, use the + and - keys on the keypad to increase or decrease the number. Press Enter when you are done.



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

3. When you finish changing controller properties, scrolling up and down on the menu as needed, move the highlight to **Apply Changes** (at the bottom of the list of options) and press Enter.

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

The following table describes all of the entries on the **Advanced Controller Management** window, including the ones that are not visible in [Figure 6.2](#).

Table 6.2 Controller Properties and Controller Management Options

Property	Description
Schedule Consistency Check	Select this option to schedule a consistency check operation to verify and correct the mirror and parity data for fault tolerant virtual drives. For more information, see Section 6.2, Scheduling a Consistency Check .
Save Controller Events Clear Controller Events	Select these options to save the controller log entries to a file or to clear entries from the log. For more information, see Section 6.3, Saving or Clearing Controller Events .
Enable Drive Security Disabling Drive Security	Select these options to enable drive security to protect the data on your system from unauthorized access or use, or to disable security. For more information, see Section 6.4, Enabling or Disabling Drive Security .
Change Security Key	Select this option to change the security key or to switch between drive security modes on the controller. For more information, see Section 6.5, Changing a Security Key .
Save TTY Log	Select this option to save a copy of the firmware's terminal log entries for the controller. For more information, see Section 6.6, Saving the TTY Log .
Host Interface	The type of interface used by the computer host system, such as PCI Express (PCIe).
Connector Count	The number of host data ports, connectors, or both currently in use on this controller.
Drive Count	The number of physical drives attached to this controller.
Virtual Drive Count	The number of virtual drives defined on this controller
Encryption Capable	Indicates whether the controller supports the encryption of data.
Encryption Enabled	Indicates whether controller encryption is currently enabled or disabled.
Protection Capable	Indicates whether the controller supports data protection.
Protection Enabled	Indicates whether data protection is enabled or disabled.
ROC Temperature (C)	The 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.
Memory Properties	
Memory Size (MB)	The size of the controller's cache memory.
NVSRAM Size (MB)	The size of the NVRAM.
Metadata Size (MB)	The total size used for storing metadata.

Table 6.2 Controller Properties and Controller Management Options (Continued)

Property	Description
Minimum Strip Size	The minimum strip size per DDF. The possible values are as follows: 0: 512 bytes 1: 1 KB 2: 2 KB 3: 4 KB 4: 8 KB ... 8: 1 MB
Maximum Strip Size	The maximum strip size per DDF. The possible values are as follows: 0: 512 bytes 1: 1 KB 2: 2 KB 3: 4 KB 4: 8 KB ... 8: 1 MB
CacheCade - SSD Caching	Indicates whether the controller supports CacheCade solid state disk (SSD) caching.
Write Cache Capable	Indicates whether the controller supports write caching.
Configured Cache Size (GB)	The total, currently configured cache size, in GB.
Maximum Allowed Cache Size (GB)	The maximum allowed cache size, in GB.
Other Properties	
Drive Replace	Enables or disables the option to copy data back from a hot spare drive to a physical drive.
Drive Replace on SMART Error	Enables or disables the option to start a Drive Replace operation if a Self-Monitoring Analysis and Report Technology (SMART) error is detected on a physical drive.
Cluster Mode	Enables or disables cluster mode on the controller, if supported.
Rebuild Rate	Displays or changes the percentage of system resources dedicated to rebuilding data on a new drive after a storage configuration drive has failed. NOTE 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.
Background Initialization (BGI) Rate	Displays or changes the percentage of system resources dedicated to performing a BGI on a redundant virtual drive. NOTE 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 and 100 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.

Table 6.2 Controller Properties and Controller Management Options (Continued)

Property	Description
Consistency Check Rate	Displays or changes the percentage of system resources dedicated to performing a consistency check operation on a redundant virtual drive. NOTE 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.
Reconstruction Rate	Displays or changes the percentage of system resources dedicated to performing a RAID level Migration (RLM) or an Online Capacity Expansion (OCE) on a virtual drive. NOTE 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 and 100 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.
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.
Coercion Mode	Drive coercion forces physical drives of varying capacities to the same size so they can be used in a virtual drive. The coercion mode options are as follows: <ul style="list-style-type: none"> ■ None (default) ■ 120 MB ■ 1 GB
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 allow you to edit the numeric value directly, without using the + and - keys.
Alarm Control	Enables or disables the controller alarm.
Boot Error Handling	Specifies the option for handling errors that the firmware might encounter during boot. The errors might require you to take action or to acknowledge the error and allow the boot process to continue. The options are <i>Stop on error</i> , <i>Pause on error</i> , <i>Ignore errors</i> , and <i>Safe mode</i> .
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.
Maintain Drive Fail History	Enables or disables the option of tracking bad physical drives across reboot.
Load Balance	Enables or disables the ability to load balance I/O transactions when redundant paths are detected.
Enable Auto Import	Enables or disables the automatic import of foreign configurations without any user intervention.
Persistent Hot Spare	Enables or disables the ability to have drive slots in the system backplane or in a storage enclosure dedicated as hot spare slots. If enabled, replacement of a hot spare drive in the same slot automatically configures the drive as a hot spare.
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 Section 6.7, Managing and Changing Link Speeds .
Patrol Read Properties	
Patrol Read State	The current state of the patrol read operation. The possible values are <i>Stopped</i> , <i>Ready</i> , <i>Active</i> , and <i>Aborted</i> .
Patrol Read Iterations	The number of patrol read iterations.

Table 6.2 Controller Properties and Controller Management Options (Continued)

Property	Description
Patrol Read Mode	Displays or changes the patrol read mode for the controller. The patrol read operation scans and resolves potential problems on configured physical drives. The possible settings are as follows: <ul style="list-style-type: none"> ■ 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 is disabled.
Patrol Read Rate	Indicates the percentage of system resources dedicated to performing a patrol read operation on configured physical drives. <p>NOTE 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.</p>
Start Patrol Read	Use this option to start a patrol read operation on the selected controller, if the current patrol read mode is <i>Manual</i> .
Suspend/Resume/Stop Patrol Read	Use these options to suspend, resume, or stop a patrol read on the selected controller, if the current patrol read mode is <i>Manual</i> .
Patrol Read Setting for Unconfigured Space	Displays or changes the patrol read setting for unconfigured space on the disks.
Cache Properties	
Cache Flush Interval	Displays or changes the interval, in seconds, at which the contents of the onboard data cache are flushed. To change this value, highlight it, press Enter, type the new value, and press Enter again.
Preserved Cache	Indicates whether the controller cache is preserved because of missing or offline virtual drives; the cache is preserved until the virtual drive is imported or the cache is discarded.
Discard Preserved Cache	Select this option to discard the preserved cache for the selected controller. If any foreign configurations exist, import them before discarding the preserved cache. Otherwise, you may lose data that belongs with the foreign configuration.
Emergency Spare Properties	
Emergency Spare	Indicates whether it is acceptable to commission otherwise incompatible global hot spare drive and/or unconfigured good drives as emergency hot spare drive.
Emergency for SMARTer	Indicates whether it is acceptable to commission emergency hot spare drive(s) for PFA events.

6.2 Scheduling a Consistency Check

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

Figure 6.3 Schedule Consistency Check Window

Schedule Consistency Check		
Consistency Check Frequency	<Weekly>	Selects the frequency of the consistency check runs. This setting can be set to disabled, hourly, daily, or weekly.
Consistency Check Start	[03/10/2012]	
Consistency Check Start	[03:01:43]	
Consistency Check Mode	<Concurrent>	
SELECT VIRTUAL DRIVES TO CHECK:		
Exclude Virtual Drives		
Apply Changes		
Apply Changes		
↑↓=Move Highlight <Enter>=Select Entry Esc=Exit		

Use this window to schedule consistency checks on the redundant virtual drives configured on the controller. The non-selectable 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.
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, but 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 popup 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. (Optional) To exclude specified virtual drives from consistency checks, highlight the **Exclude Virtual Drives** field and press Enter.

The following window appears, listing the virtual drives defined on this controller.

Figure 6.4 Exclude Virtual Drives Window



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.

6. 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 will not undergo a consistency check.

7. Highlight the **Apply Changes** field and press Enter.

The program returns you to the **Schedule Consistency Check** window.

8. Highlight the **Apply Changes** field on the **Schedule Consistency Check** window and press Enter.

The consistency check changes are now registered.

6.3 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 6.5 Save Controller Events Window

Save Controller Events		
Select File System	⏏	Enables you to choose the appropriate file system to save the controller logs.
Select Directory	<Current Dir>	
Enter Filename	ctrlEvents.txt	
Save Events		
↑↓=Move Highlight <Enter>=Select Entry Esc=Exit		

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

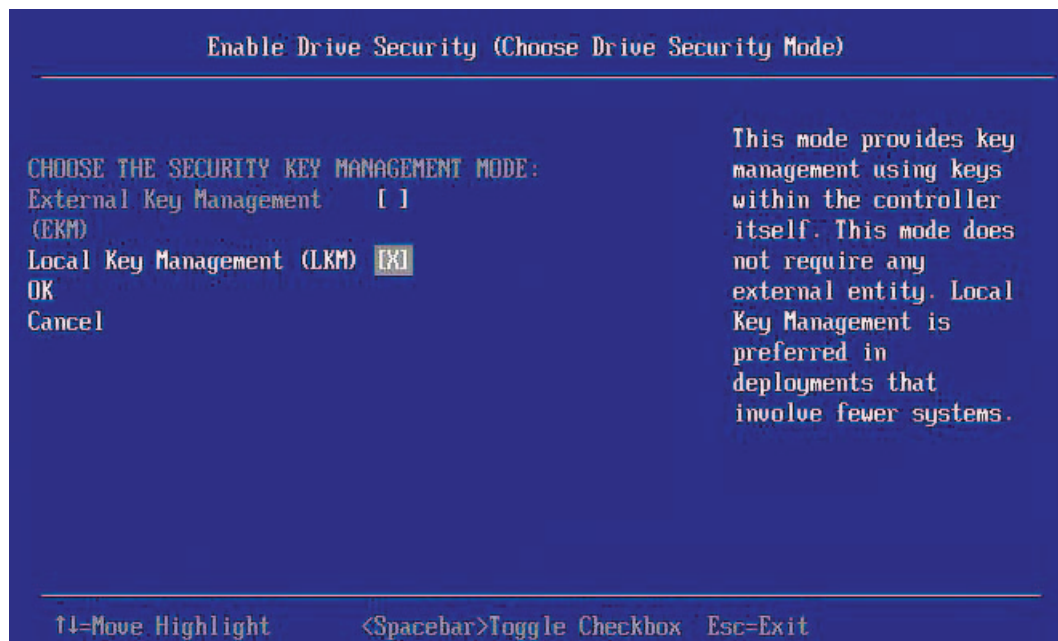
1. To select a different file system from the one listed in the **Select File System** field, highlight the current file system name and press Enter.
An error message appears if there is no file system.
2. Select a file system from the popup menu and press Enter.
3. To save the controller events file 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 popup 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 popup window 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** window. When the confirmation message appears, highlight **OK** and press Enter.

6.4 Enabling or Disabling Drive Security

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

Figure 6.6 Enable Drive Security Window



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 needed, press the spacebar to enter an X in this field.
2. Highlight **OK** and press Enter. The following window appears.

Figure 6.7 Enable Drive Security Window



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 reenter 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.



NOTE If you forget the security key, you will lose access to the data on the drives. Be sure to record your security key and password information. You might need to enter the security key to perform certain operations.

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 6.8 Disable Drive Security Warning

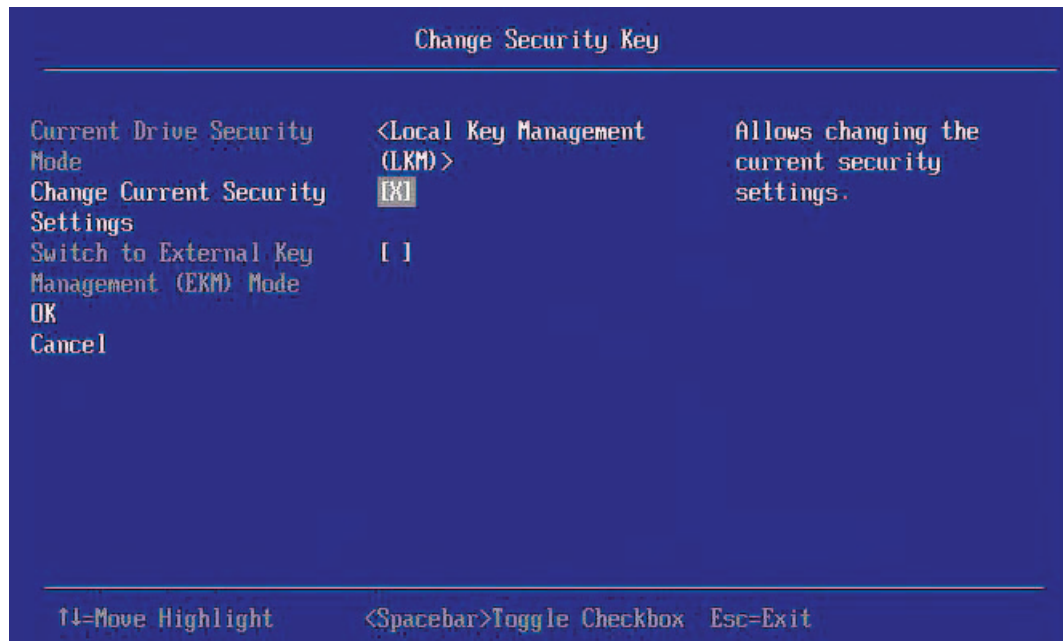


2. Read the warning and be sure you understand what will happen if you disable the drive security.
3. Highlight **Confirm** and press the spacebar to select it.
4. Highlight **Yes** and press Enter.
Drive security is disabled.

6.5 Changing a Security Key

The following window appears when you select **Change Security Key** from the **Advanced Controller Management** menu.

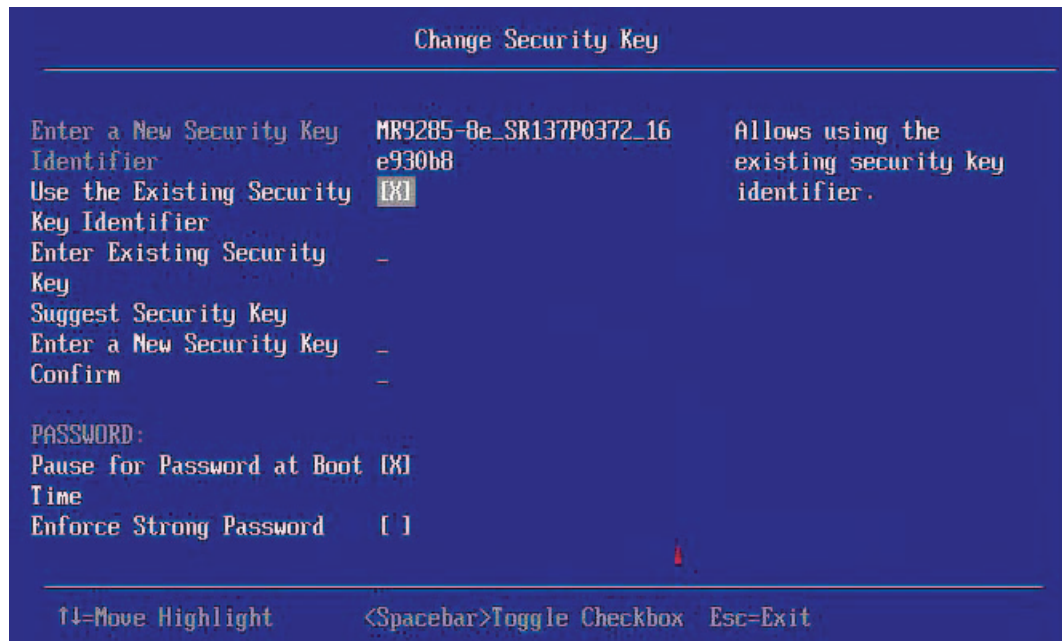
Figure 6.9 Change Security Key Window



Follow these steps to change the security key settings.

1. Highlight **OK** and press Enter.
The following window appears.

Figure 6.10 Security Key Settings Window



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. To request the controller to suggest a new drive security key, highlight **Suggest Security Key** and press Enter.
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.



NOTE If you forget the security key, you will lose access to the data on the drives. Be sure to record your new security key and password information. You might need to enter the security key to perform certain operations.

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.

6.6 Saving the TTY Log

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

Figure 6.11 Save TTY Log Window

Save TTY Log		
File Systems	<HANTOOL>	Enables you to choose the appropriate directory to save the controller logs. The default (root) directory will be selected upon entering this form.
Select File System		
Directories	<DOS>	
Select Directory		
Enter Filename	ttyLog.txt	
Entries to Save	<All>	
Save Log		
↑↓=Move Highlight <Enter>=Select Entry Esc=Exit		

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 popup 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 popup 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.

6.7 Managing and Changing Link Speeds

The Manage 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. The following window appears when you select **Manage Link Speed** on the **Advanced Controller Management** window. The default setting for all phys is *Auto*.

Figure 6.12 Manage Link Speed



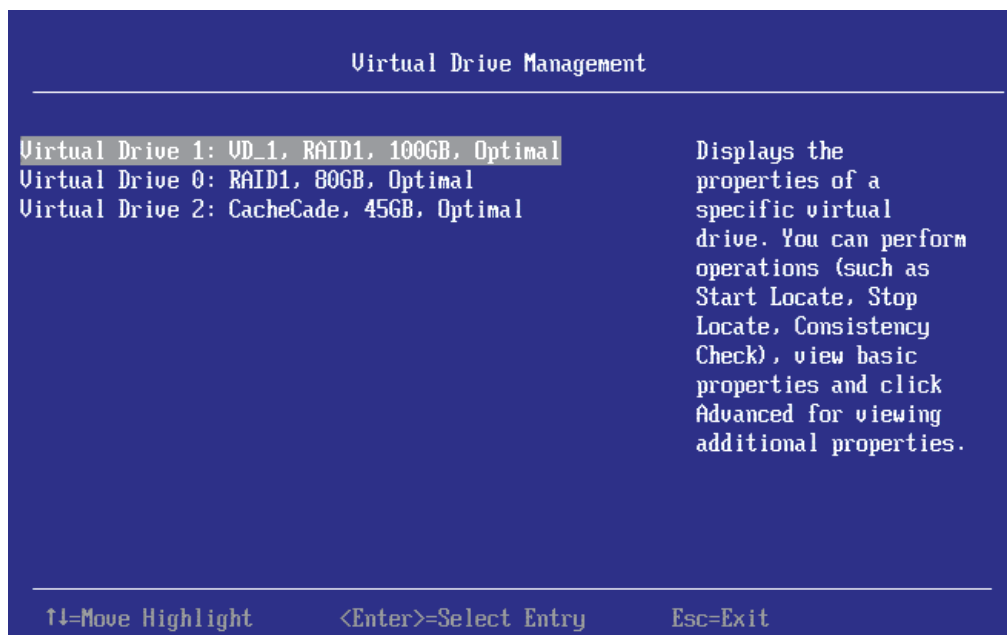
Follow these steps to change the link speed for one or more phys:

1. Highlight the field to the right of the phy number and press Enter.
2. Select an option from the popup menu.
The choices are *1.5 Gbps*, *3 Gbps*, *6 Gbps*, and *Auto*.
3. Scroll to the bottom of the phy list, highlight **OK**, and press Enter.

Chapter 7: Managing Virtual Drives

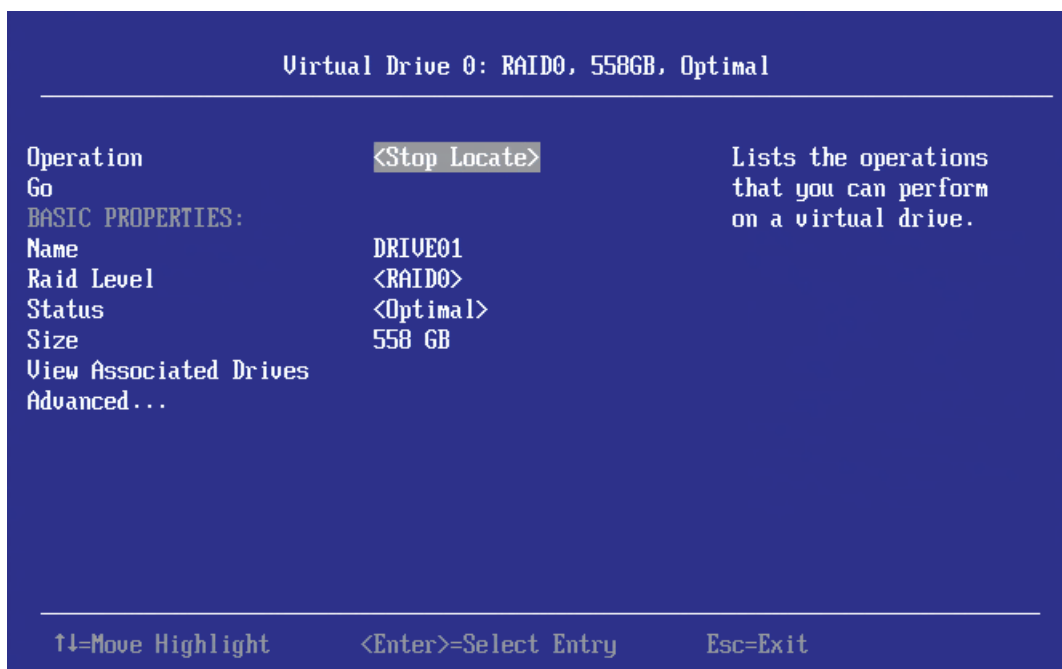
When you select **Virtual Drive Management** on the **Main Menu**, the **Virtual Drive Management** window appears, as shown in the following figure.

Figure 7.1 Virtual Drive Management Menu



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 window appears.

Figure 7.2 Virtual Drive Management Menu



This window lists the following basic virtual drive properties.

Table 7.1 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.

For information on how to perform virtual drive operations, see Section 7.1, [Selecting Virtual Drive Operations](#).

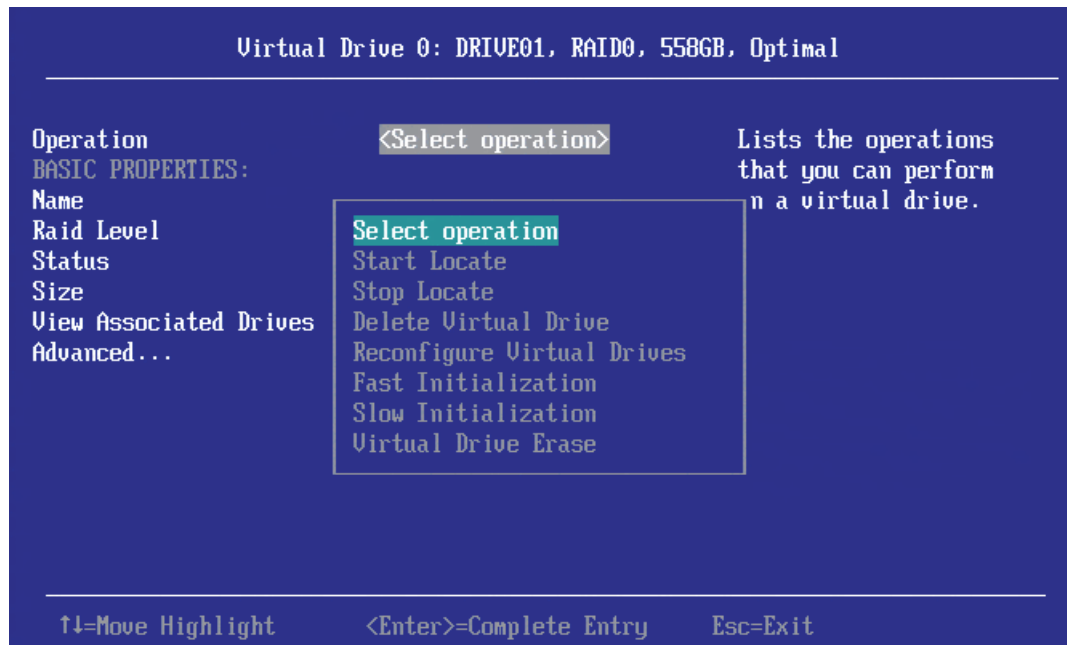
For information on how to view the physical drives associated with the virtual drive, see Section 7.3, [Viewing Associated Drives](#).

For information on how to view and change advanced virtual drive settings, see Section 7.4, [Viewing and Managing Virtual Drive Properties and Options](#).

7.1 Selecting Virtual Drive Operations

The following popup menu appears when you highlight **Operation** in the **Virtual Drive** window and press Enter.

Figure 7.3 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.

7.1.1 Locating Physical Drives in a Virtual Drive

Follow these steps to locate the physical drives in a virtual drive by flashing their LEDs.

1. Highlight **Start Locate** on the popup 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** window.
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** window.
The LEDs on the physical drives stop flashing.

7.1.2 Deleting a Virtual Drive



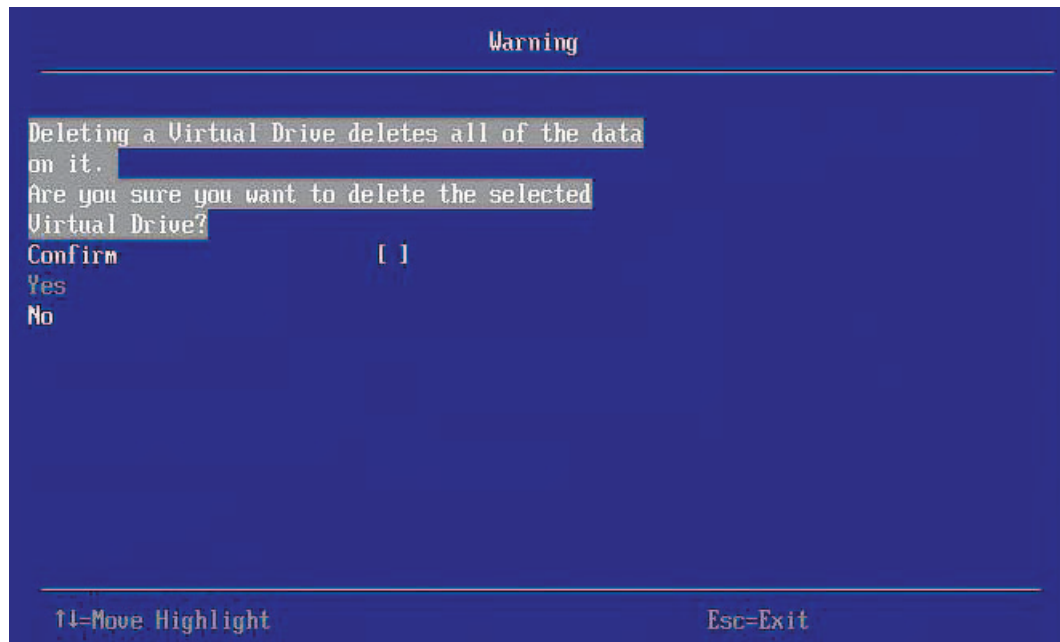
ATTENTION All data on a virtual drive is lost when you delete it. Be sure to back up data you want to keep before you delete a virtual drive.

Follow these steps to delete a virtual drive. The 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** window and select the virtual drive.

1. Highlight **Delete Virtual Drive** on the popup menu and press Enter.
2. Highlight the word **Go** that appears beneath **Operation** and press Enter.

The following warning message appears.

Figure 7.4 Delete Virtual Drive Warning



3. Highlight **Confirm** and press the spacebar to confirm the deletion, then highlight **Yes** and press Enter.

The virtual drive is deleted.

7.1.3 Reconfiguring a Virtual Drive

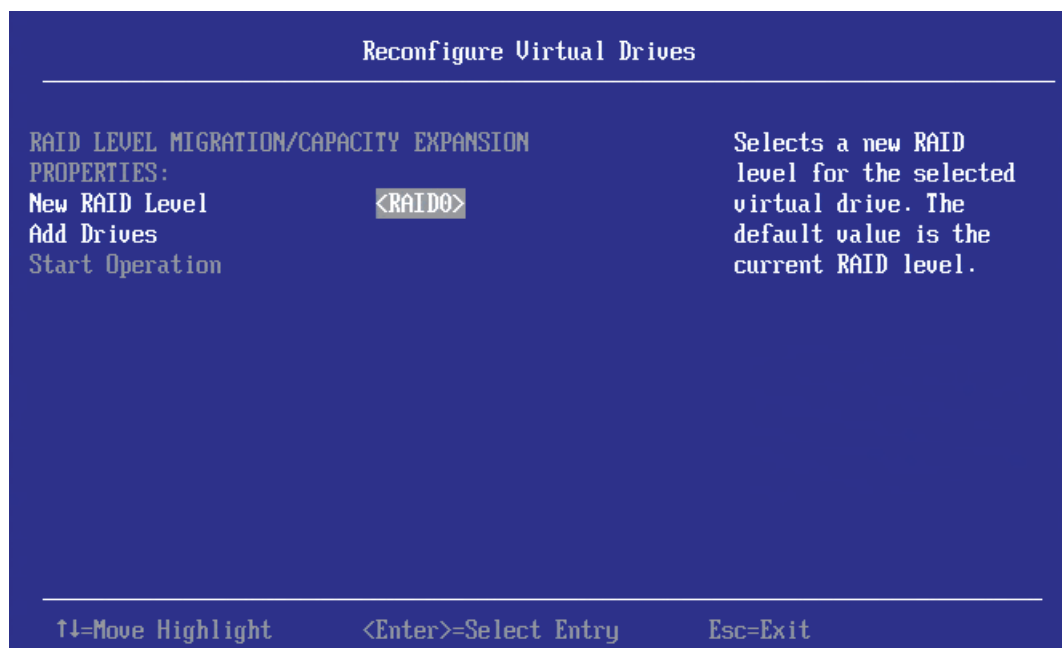
You can reconfigure a virtual drive by changing its RAID level, or by adding physical drives to it, or by doing both of these actions. When performing these changes, however, you must observe the maximum drive and minimum drive restrictions for the various RAID levels. See [Chapter 2: Introduction to RAID](#) for more information.

Follow these steps to reconfigure a virtual drive.

1. Highlight **Reconfigure Virtual Drive** on the popup menu and press Enter.
2. Highlight the word **Go** that appears beneath **Operation** and press Enter.

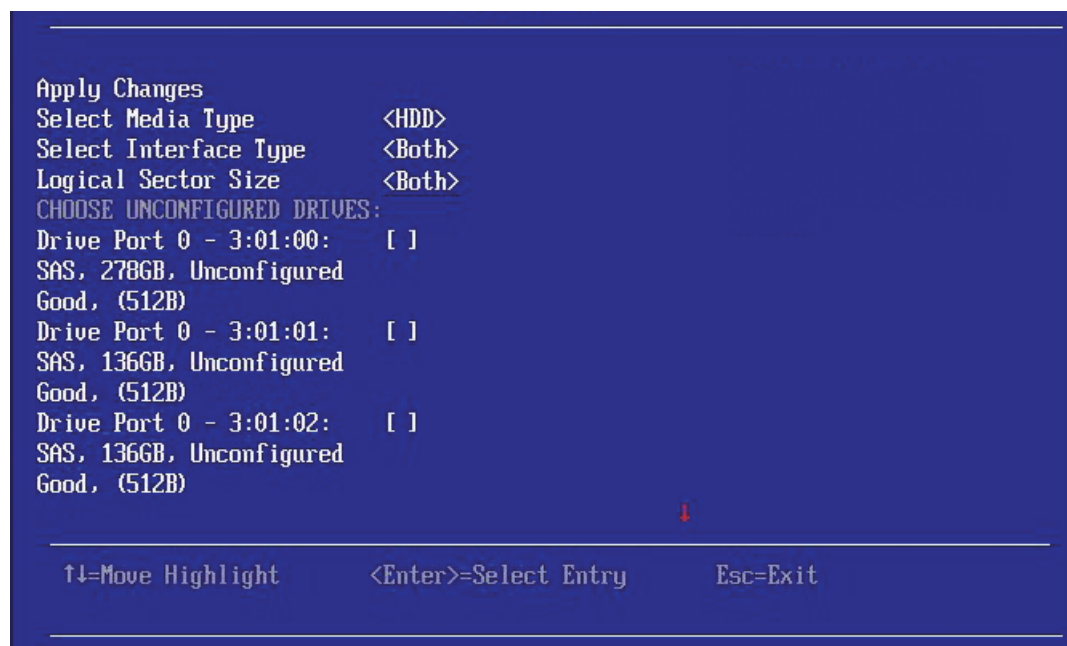
The following window appears.

Figure 7.5 Reconfigure Virtual Drives Window



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 popup menu.
5. To add physical drives to the selected virtual drive, highlight **Add Drives** and press Enter.
The following window appears.

Figure 7.6 Select Drives Window



6. (Optional) Change the default **Select Media Type** value by highlighting this field, pressing Enter, and selecting an option from the popup menu.

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

7. (Optional) Change the default **Select Interface Type** value by highlighting this field, pressing Enter, and selecting an option from the popup 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.

8. Optional) Change the Default **Logical Sector Size** by highlighting this field, pressing enter, and selecting an option from the popup menu.

This field allows you to choose between a 512 byte and 4K logical sector size. To use 4K logical sector size, the physical drives that are being configured must be 4K compatible (Advanced Format Drives). Otherwise, if 4K is selected as the desired sector size, no drives will display in the drive list. 512 byte sector size is supported by all drives.

9. Select physical drives to add to the virtual drive by highlighting drives and pressing the spacebar to select them. 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 window. The amount of information that can be displayed in one window 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, or else 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 5.3](#) for more information.

-
10. When you have selected all the drives to add to the virtual drive, highlight **Apply Changes** and press Enter.



NOTE If you 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** window.

11. Highlight **Start Operation** and press Enter to implement the changes to the virtual drive.

7.1.4 Initializing a Virtual Drive

Follow these steps to initialize a virtual drive.



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.

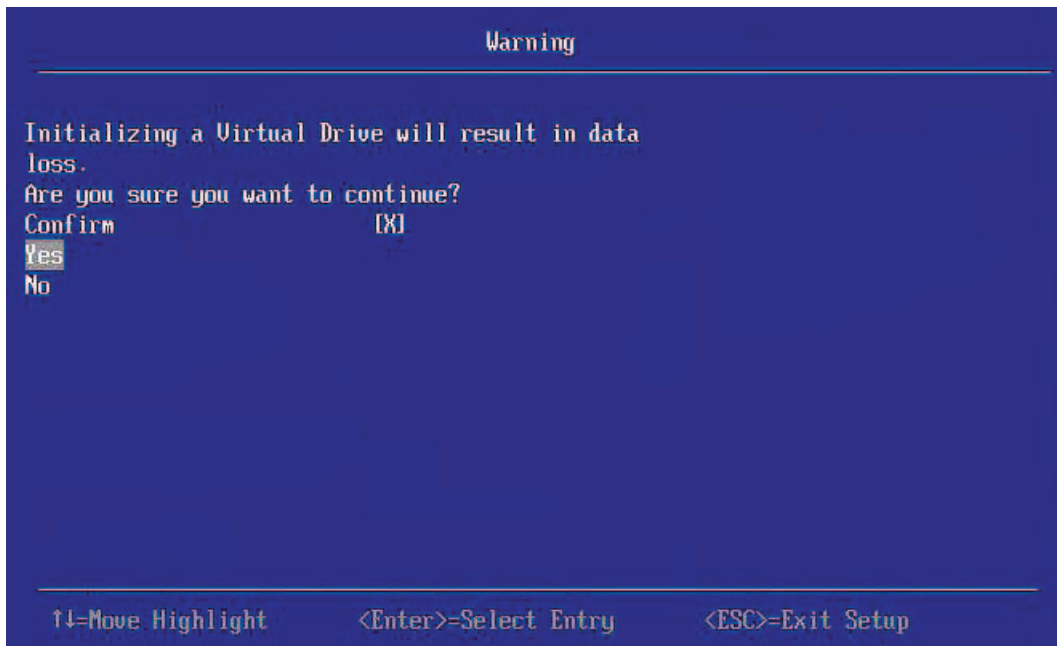
-
1. Highlight **Fast Initialization** or **Slow Initialization** on the popup menu and press Enter.

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 following warning appears.

Figure 7.7 Initialize Virtual Drive Warning



3. Highlight **Confirm** and press the spacebar to confirm the operation, then highlight **Yes** and press Enter.
The initialization process begins on the virtual drive.

7.1.5 Erasing a Virtual Drive

Follow these steps to erase the data on a virtual drive. After the data is erased, you have the option of keeping the blank virtual drive, which you can use to store other data, or deleting the virtual drive completely.



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.

1. Highlight **Virtual Drive Erase** on the popup menu and press Enter.
Two additional fields appear.
2. Highlight **Erase Mode** and press Enter.
3. Select *Simple*, *Normal*, or *Thorough* from the popup 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 following warning message appears.

Figure 7.8 Virtual Drive Erase Warning



6. Highlight **Confirm** and press the spacebar to confirm the operation, then highlight **Yes** and press Enter.
The virtual drive is erased.

7.1.6 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 popup menu and press Enter.
2. Highlight **Go** and press Enter.
The following warning message appears.

Figure 7.9 Enable SSD Caching Warning



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.

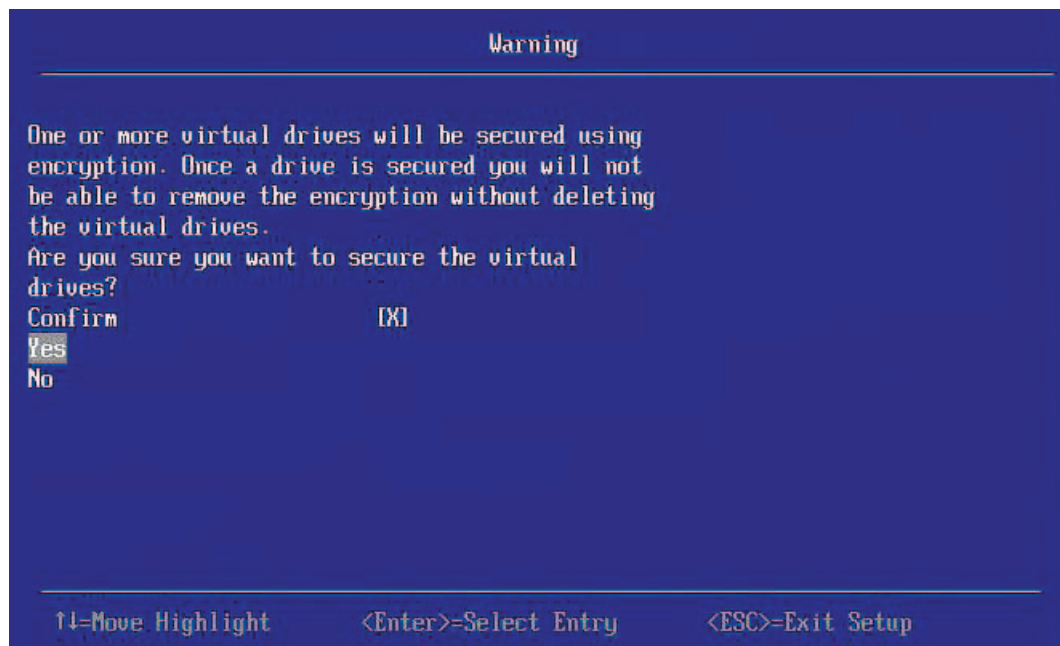
The warning is similar when you disable SSD caching.

7.1.7 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. Follow these steps to secure a virtual drive.

1. Highlight **Secure Virtual Drive** on the popup menu and press Enter.
The following warning appears.

Figure 7.10 Secure Virtual Drive Warning



2. Highlight **Confirm** and press the spacebar to confirm the operation, then highlight **Yes** and press Enter.
The virtual drive is secured.

7.1.8 Running a Consistency Check

Follow these steps to run a consistency check on the currently selected redundant virtual drive.

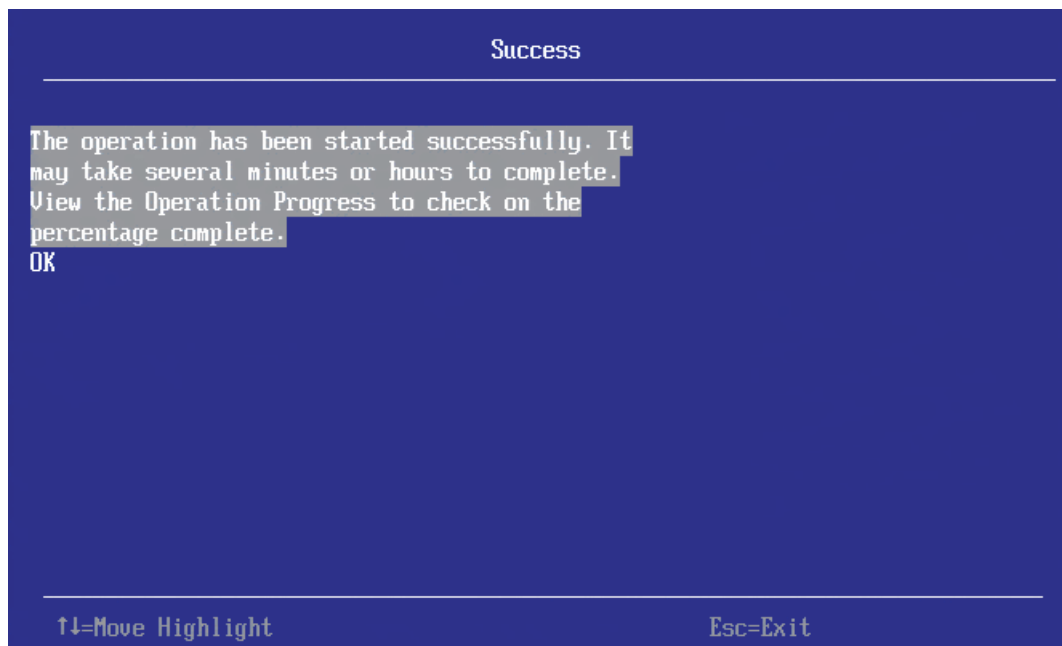
1. Highlight **Check Consistency** on the popup menu and press Enter.



NOTE The Check Consistency option does not appear on the menu if the currently selected virtual drive is RAID 0 (non-redundant).

2. Highlight **Go** and press Enter.
The following window appears.

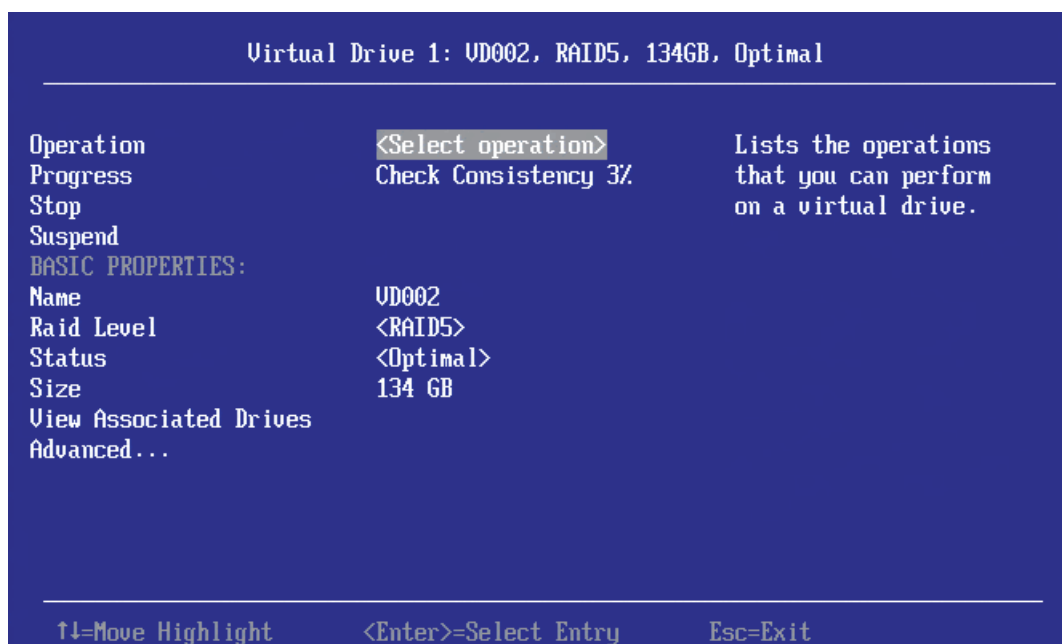
Figure 7.11 Consistency Check Success Window



As the message indicates, the consistency check is now running.

3. Highlight **OK** and press Enter to return to the previous window, as shown in the following figure.

Figure 7.12 Consistency Check Progress Indicator



The Progress indicator in the window shows the percentage progress of the consistency check. To refresh the indicator, exit the window and reenter 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.

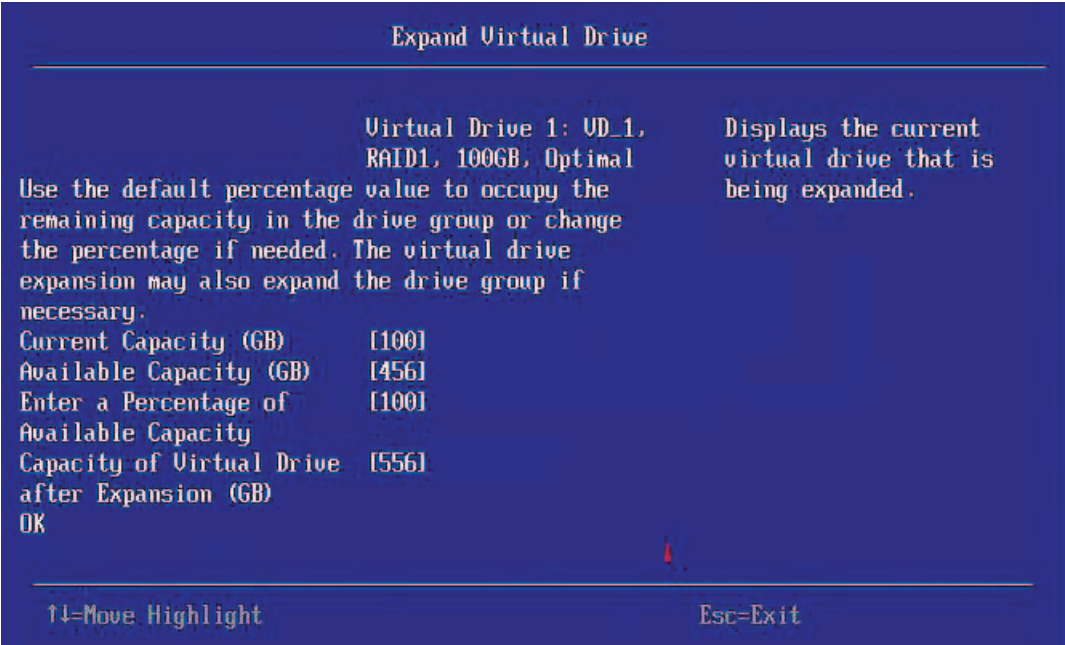
For more information about consistency checks, see Section 6.2, [Scheduling a Consistency Check](#).

7.1.9 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 popup menu.
The following window appears.

Figure 7.13 Expand Virtual Drive Window



- The window 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.

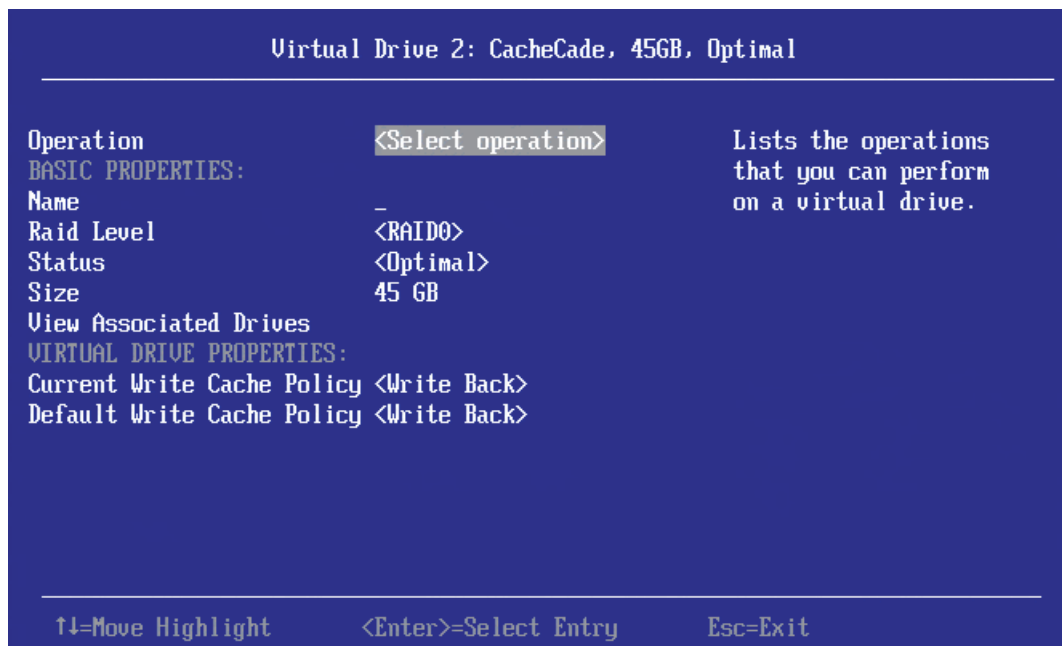
3. When you have set the capacity to the desired level, highlight **OK** and press Enter.
The capacity of the virtual drive is expanded.

7.2 Managing CacheCade Virtual Drives

After you create a CacheCade virtual drive, as described in Section 5.3, [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 7.14 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 Section 7.1.1, [Locating Physical Drives in a Virtual Drive](#).
- **Delete Virtual Drive:** Use this option to delete the CacheCade virtual drive. For more information, see Section 7.1.2, [Deleting a Virtual Drive](#).

To assign a name to the CacheCade virtual drive, highlight **Name**, press Enter, type the name, and press Enter again.

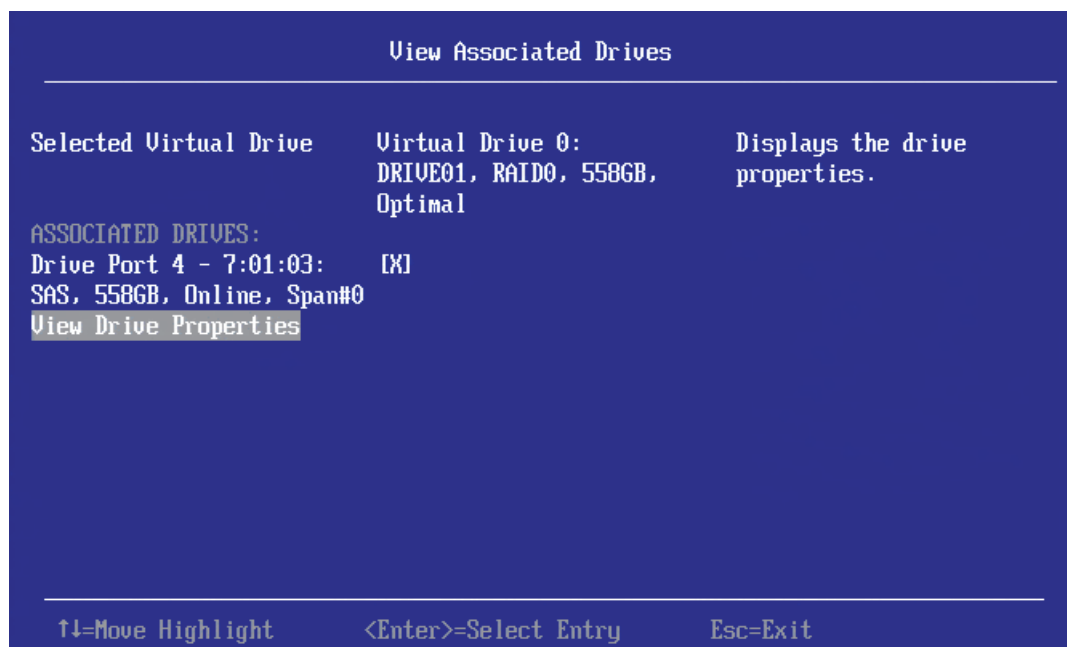
To change the default write cache policy, highlight **Default Write Cache Policy**, press Enter, and select an option from the popup menu. Options are *Write Through*, *Write Back*, and *Force Write Back*.

To view the drives associated with the CacheCade virtual drive, highlight **View Associated Drives** and press Enter. For more information, see Section 7.1.6, [Enabling and Disabling SSD Caching](#).

7.3 Viewing Associated Drives

The following window appears when you select **View Associated Drives** at the bottom of the **Virtual Drive** window.

Figure 7.15 View Associated Drives Window



The window 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 Section 8.2, [Viewing Advanced Drive Properties](#).

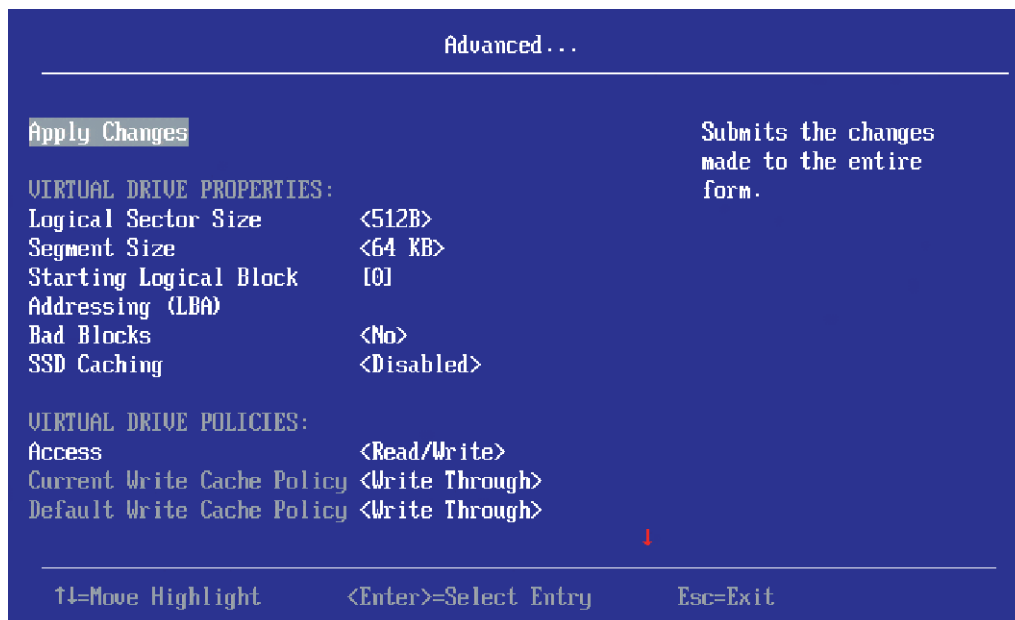
7.4 Viewing and Managing Virtual Drive Properties and Options

The following window appears when you select **Advanced** from the **Virtual Drive** window. (The second figure shows the rest of the options that are visible when you scroll down.)



NOTE The properties and options shown in the window 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 window.

Figure 7.16 Advanced Virtual Drive Properties 1

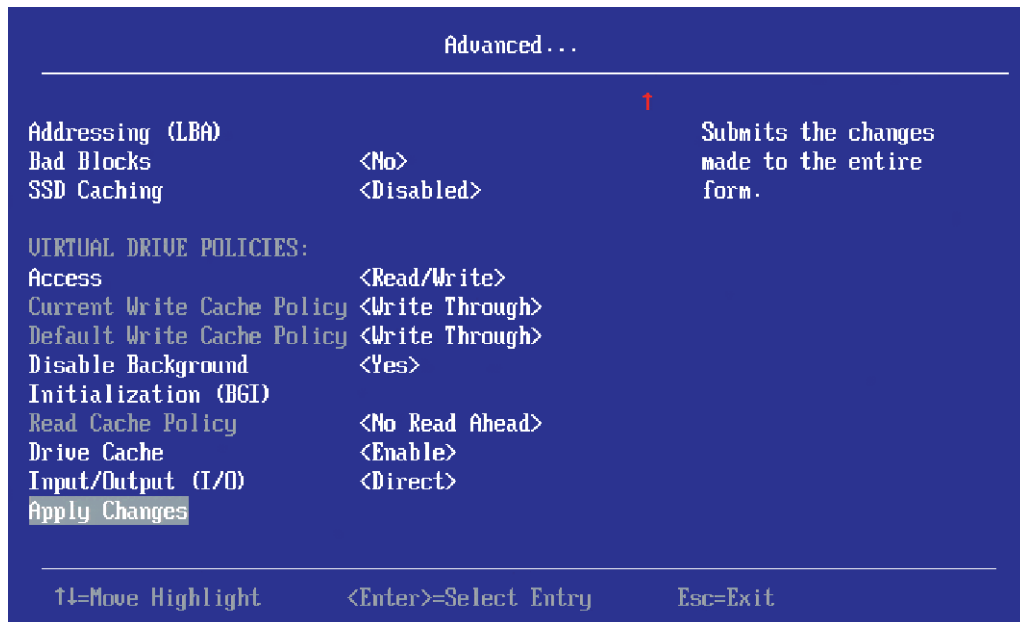


The small red arrow at the bottom of the window indicates that you can scroll down to view more virtual drive properties and virtual drive policies, as shown in the following figure.



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.

Figure 7.17 Advanced Virtual Drive Properties 2



The following table describes all of the virtual drive properties listed in this window.

Table 7.2 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) Secured	Indicates whether the virtual drive is secured or not 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 window 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 popup 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 7.3 Virtual Drive Policies

Property	Description
Access	The access policy for the virtual drive. The options are <i>Read/Write</i> , <i>Read Only</i> , and <i>Blocked</i> .
Current Write Cache Policy	Displays the current write cache policy. The possible values are as follows: <ul style="list-style-type: none"> 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.

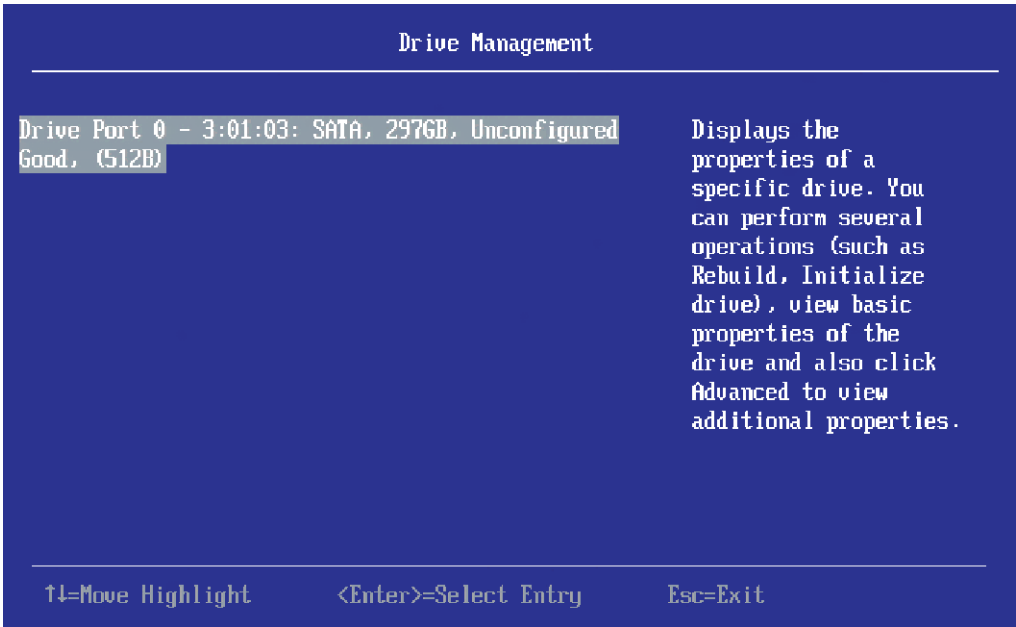
Table 7.3 Virtual Drive Policies (Continued)

Property	Description
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	<p>The read cache policy for the virtual drive. The possible values are as follows:</p> <ul style="list-style-type: none"> ■ <i>Ahead</i>: 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 there is little improvement when accessing random data. ■ <i>Normal</i>: Read-ahead capability is disabled.
Drive Cache	The disk cache policy for the virtual drive. The possible values are <i>Unchanged</i> , <i>Enable</i> , and <i>Disable</i> .
Input/Output (I/O)	<p>The I/O policy for the virtual drive. The possible values are as follows:</p> <ul style="list-style-type: none"> ■ <i>Direct</i>: 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.) ■ <i>Cached</i>: All reads are buffered in cache.

Chapter 8: Managing Physical Drives

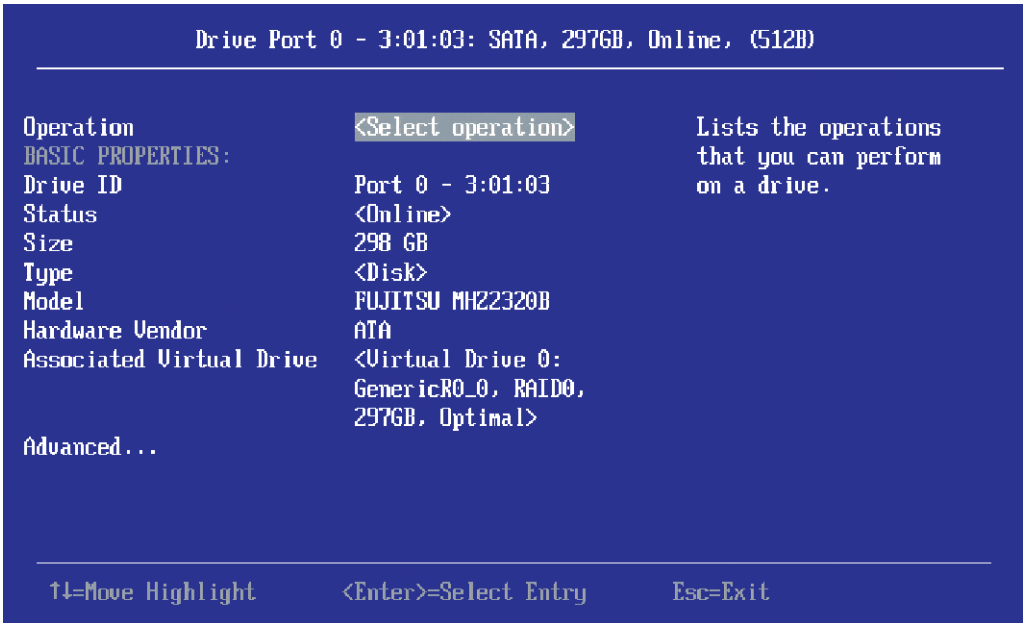
When you select **Drive Management** on the **Main Menu**, the **Drive Management Selection** menu appears, as shown in the following figure.

Figure 8.1 Drive Management Selection Menu



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 window appears.

Figure 8.2 Drive Management Menu



This window lists the following basic drive properties for the selected drive:

Table 8.1 Basic Physical Drive Properties

Property	Description
Drive ID	The ID of the currently selected drive. The format of the ID is <i>Connector: Port wired order: Slot</i> . If the drive is not installed in an enclosure, the format of the ID is <i>Connector: Port wired order</i> .
Status	The status of the drive, such as <i>Online</i> , <i>Ready</i> , <i>Available</i> , or <i>Failed</i> .
Size	The drive capacity, in GB.
Type	The device type of the drive, which is normally <i>Disk</i> .
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.

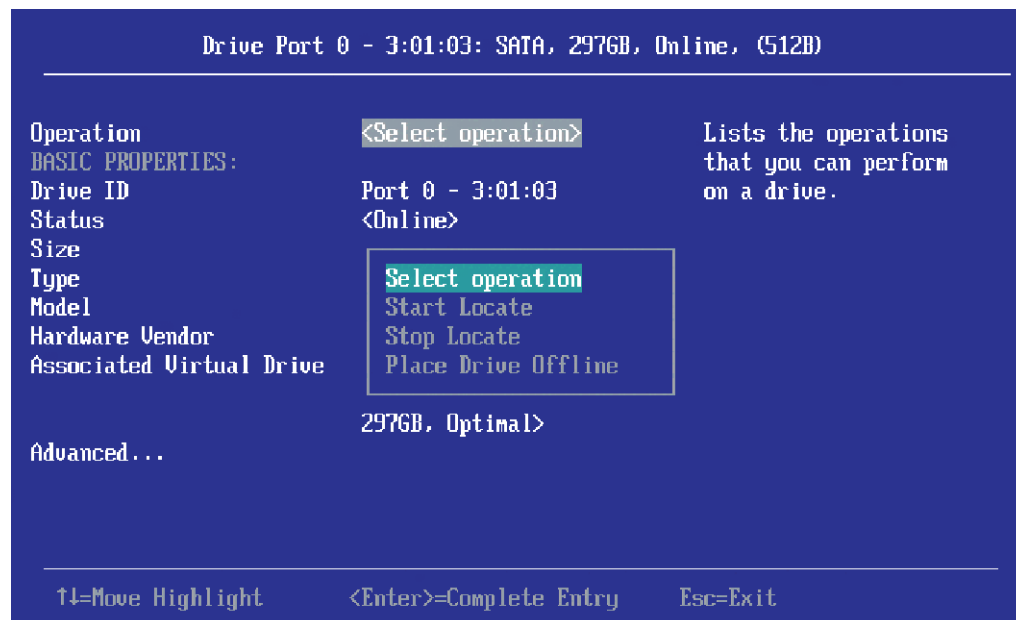
For information on performing drive operations, see Section 8.1, [Performing Drive Operations](#).

For information on viewing and changing drive settings and properties, see Section 8.2, [Viewing Advanced Drive Properties](#).

8.1 Performing Drive Operations

As the following figure shows, a popup drive operations menu appears when you highlight the **Select operations** field and press Enter.

Figure 8.3 Select Drive Operations Menu



Start Locate and *Stop Locate* are available options for any selected drive. The other menu options vary depending 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, and press Enter to select it and return to this window.

8.1.1 Locating a Drive

Follow these steps to locate a physical drive by flashing its LED.

1. Open the popup 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 window.

8.1.2 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 ServeRAID system and insert a new physical drive, if the inserted drive does not contain valid DDF metadata, the drive status is *JBOD* (Just a Bunch of Drives) 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 *JBOD* drives to create a RAID configuration, because they do not have valid DDF records. First, the drives must be converted to *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*, *Unconfigured Good*, or *JBOD* drive.

1. Open the popup 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.
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 window, exit back to the **Main Menu** and then reenter the **Drive Management** window.

8.1.3 Replacing a Drive

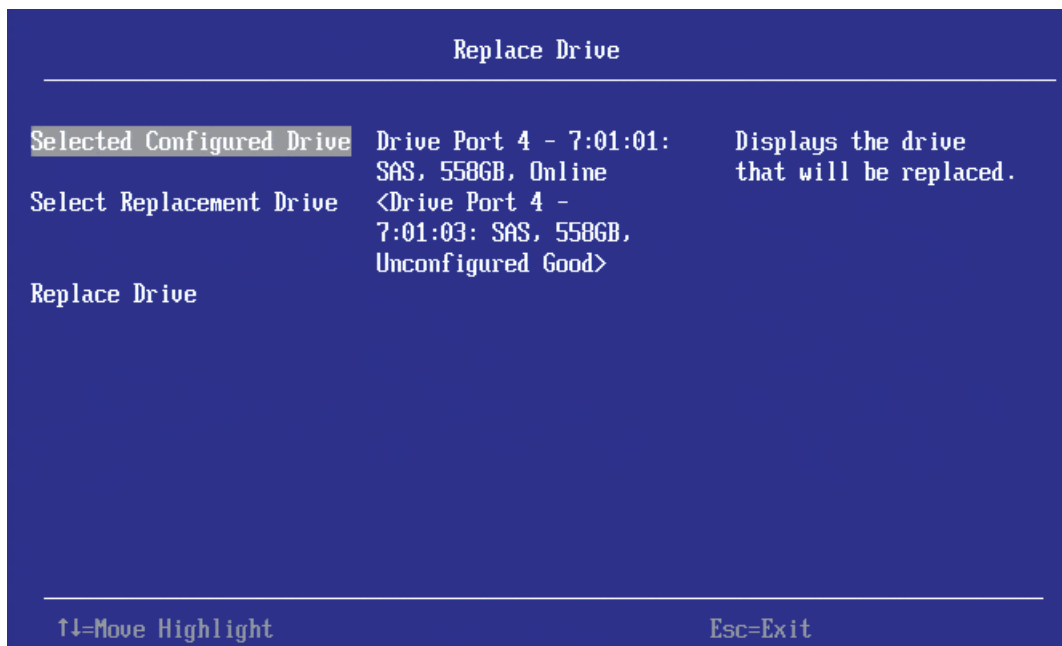
You might need to replace a drive that is a member of a redundant virtual drive connected to the controller if the drive shows signs of failing. Before you start this operation, be sure that an available *Unconfigured Good* replacement drive is available. The replacement drive must have at least as much capacity as the drive you are replacing.

Follow these steps to replace a drive.

1. Open the popup drive operations menu, highlight **Replace Drive**, and press Enter.
2. Highlight **Go**, which appears beneath Operation, and press Enter.

The following window appears.

Figure 8.4 Replace Drive Window



3. Highlight **Select Replacement Drive** and press Enter.
A popup 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.

8.1.4 Placing a Drive Offline

Follow 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 popup drive operations menu, highlight **Place Drive Offline**, and press Enter.
2. Highlight **Go**, which appears beneath Operation, and press Enter.

The following message appears.

Figure 8.5 Place Drive Offline Warning



3. Highlight **Confirm** and press the spacebar to confirm the operation.
4. Highlight **Yes** and press Enter.
The selected drive is forced offline.

8.1.5 Placing a Drive Online

Follow these steps to force a selected member drive of a virtual drive online after it been forced offline.

1. Open the popup drive operations menu, highlight **Place Drive Online**, and press Enter.
2. Highlight **Go** and press Enter.
The following warning appears.

Figure 8.6 Place Drive Online Warning



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 window.
The drive is now online.

8.1.6 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 of assigning it as a global hot spare drive. Follow these steps to assign a global hot spare.

1. Open the popup drive operations menu, highlight **Assign Hot Spare Drive**, and press Enter.
2. Highlight **Go**, which appears beneath Operation, and press Enter.
A hot spare selection window 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 window, exit back to the **Main Menu** and then reenter the **Drive Management** window.

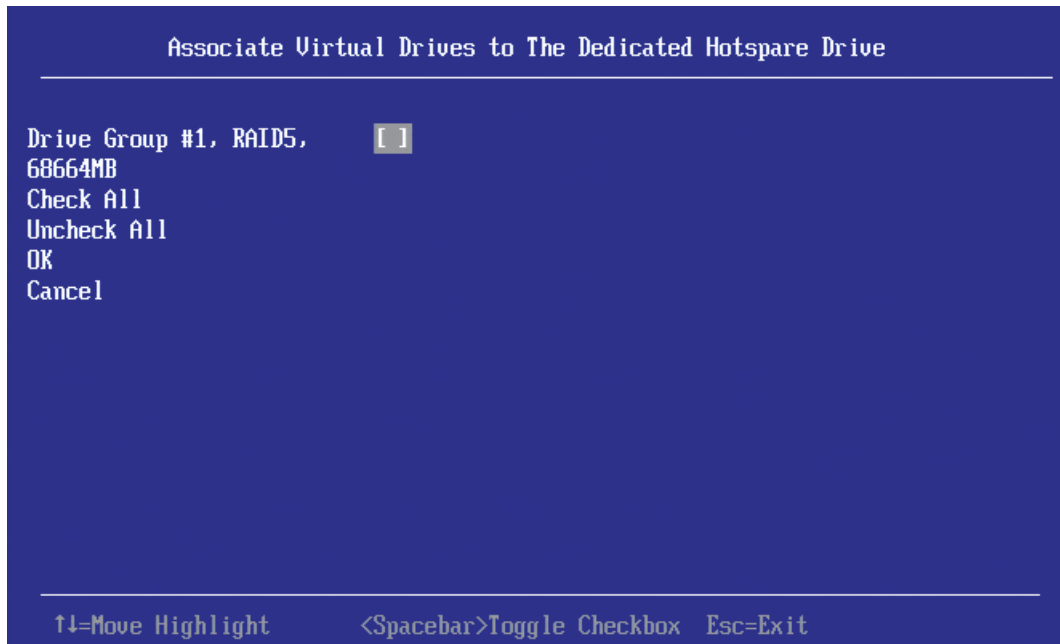
8.1.7 Assigning a Dedicated Hot Spare Drive

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 of assigning it as a dedicated spare drive. Follow these steps to assign a dedicated hot spare.

1. Open the popup drive operations menu, highlight **Assign Dedicated Spare Drive**, and press Enter.
2. Highlight **Go**, which appears beneath Operation, and press Enter.

The following selection window appears.

Figure 8.7 Assign Virtual Drives to Dedicated Hot Spare Drive Window



The window lists a single entry for each existing drive group. If you create a partial virtual drive on the same drive group, you will see a single entry with the cumulative size.

3. Select the drive groups to which this hot spare drive will be dedicated by highlighting each drive group and pressing the spacebar.

Alternatively, use the **Check All** or **Uncheck All** commands to select or deselect all drive groups.

4. When your selection is complete, highlight **OK** and press Enter.

When you return to the previous window, the status of the selected drive changes to Hot Spare.



NOTE To refresh the status of the drive displayed in the window, exit back to the **Main Menu** and then reenter the **Drive Management** window.

8.1.8 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.

Follow these steps to unassign a hot spare drive.

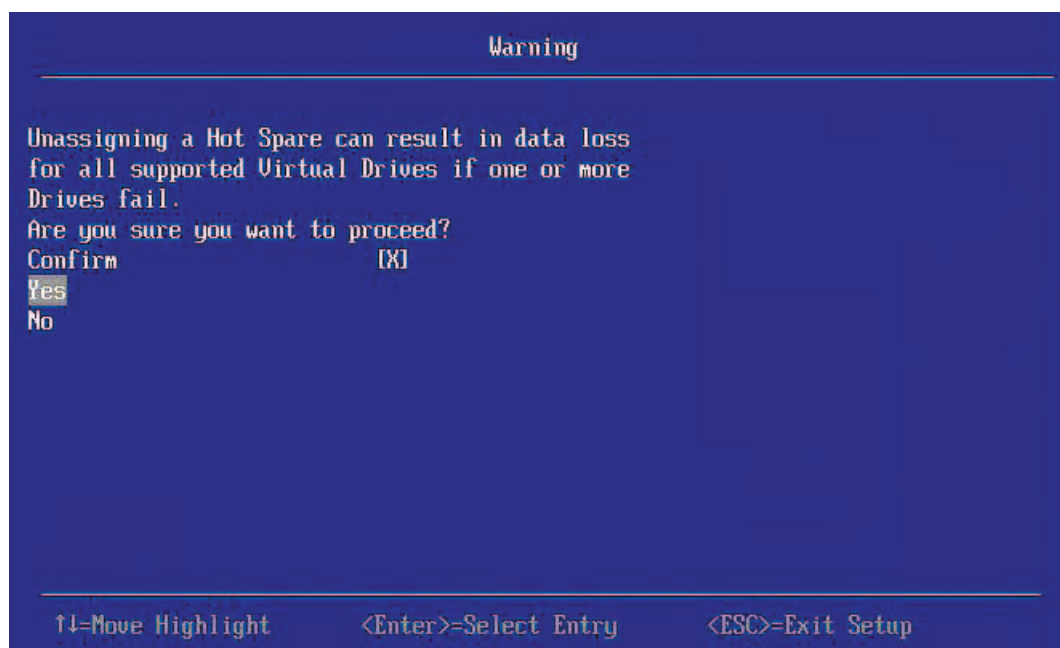


ATTENTION If you unassign a global hot spare drive or a dedicated hot spare drive, you reduce the protection level of the data on the virtual drives.

1. Open the popup drive operations menu, highlight **Unassign Hot Spare Drive**, and press Enter.
2. Highlight **Go**, which appears beneath Operation, and press Enter.

The following warning message appears.

Figure 8.8 Unassign Hot Spare Drive Warning



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 window, exit back to the **Main Menu** and then reenter the **Drive Management** window.

8.1.9 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 zeroes and ones.

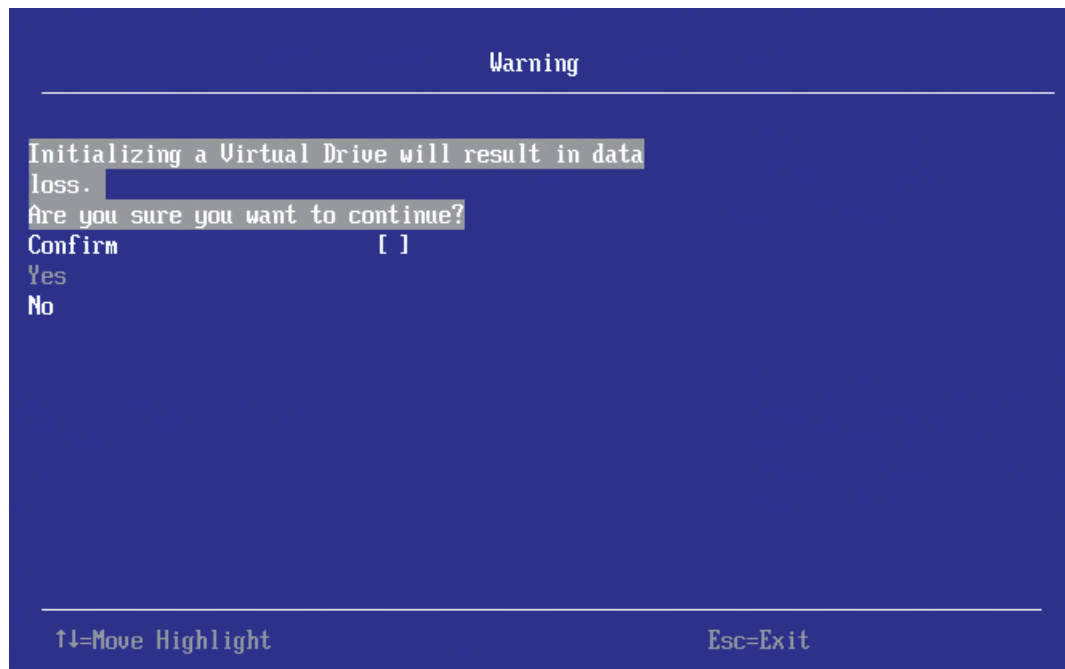


ATTENTION All data on the drive is lost when you initialize it or erase it. Before starting these operations, back up any data that you want to keep.

1. Open the popup 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 popup menu and press Enter.
4. Highlight **Go** and press Enter.

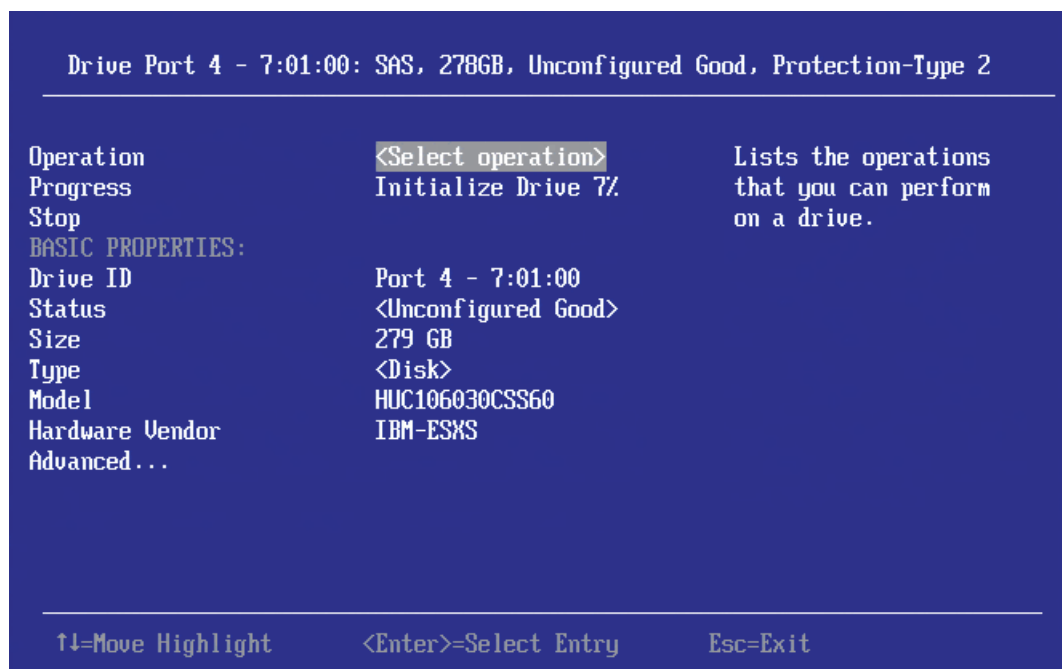
The following message appears. (The message is similar for erasing a drive.)

Figure 8.9 Initialize Virtual Drive Warning



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 window now displays a progress field and a Stop command, as shown in the following figure.

Figure 8.10 Initialize Progress Indicator



- To stop the initialization or erase process, highlight **Stop** and press Enter.



NOTE To refresh the progress indicator, press Esc to exit this window, then open it again.

8.1.10 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*, the HLL firmware automatically uses an Unconfigured Good drive to rebuild a failed or offline drive if no hot spares are available.

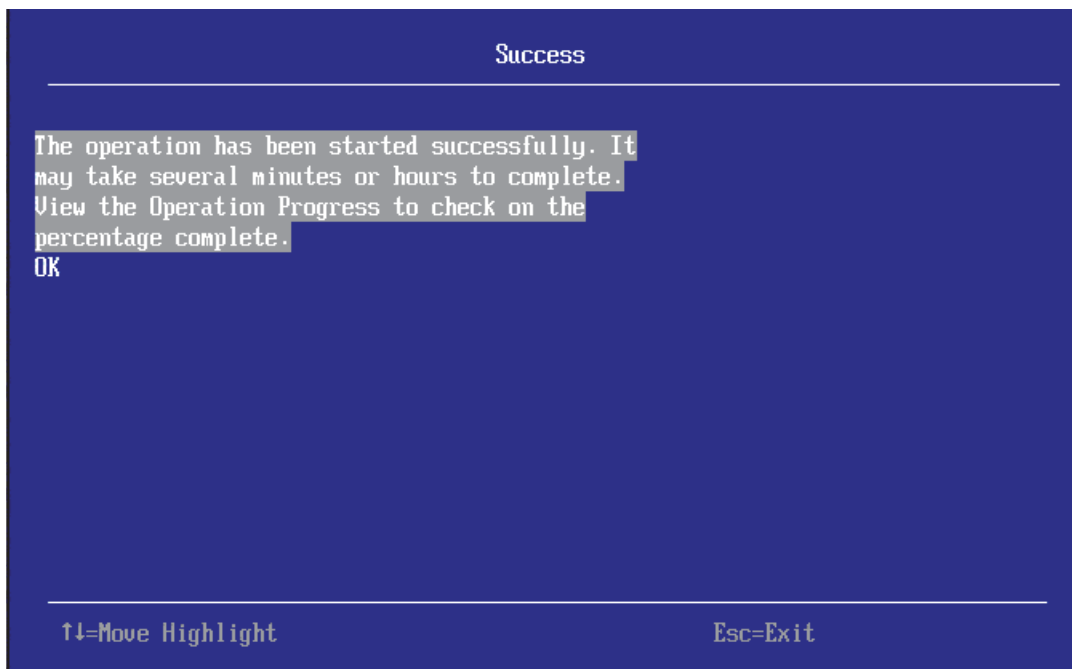
The manual Rebuild option is available only if a member drive of a virtual drive fails, there are no available hot spare drives, and the Emergency Spare controller property is set to *None*.

Follow these steps to start a manual Rebuild operation on an Unconfigured Good drive.

- Open the popup drive operations menu, highlight **Rebuild**, and press Enter.
- Highlight **Go** and press Enter.

The rebuild operation begins, and the following success message appears.

Figure 8.11 Rebuild Drive Success Message



8.1.11 Securely Erasing a Drive

Follow 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.



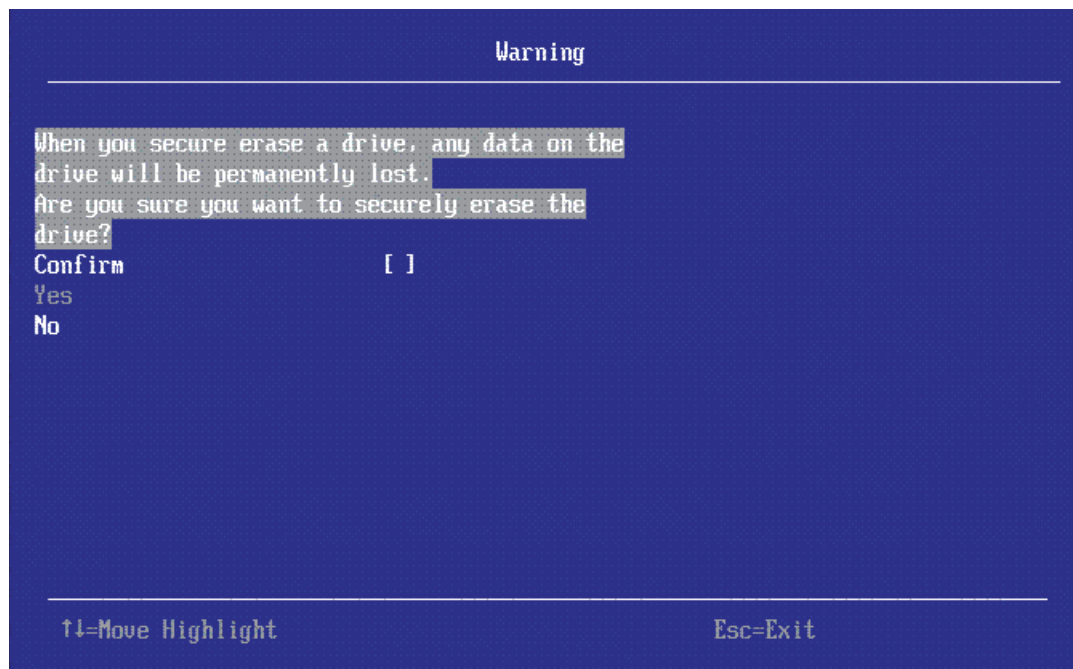
ATTENTION All data on the drive is lost when you erase it. Before starting these operations, back up any data that you want to keep.

Follow these steps to securely erase an FDE capable drive:

1. Open the popup drive operations menu, highlight **Secure Erase**, and press Enter.
2. Highlight **Go** and press Enter.

The following warning message appears.

Figure 8.12 Secure Erase Warning



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 window. This window now displays a progress field and a Stop command.
6. To stop the secure erase process, highlight **Stop** and press Enter.

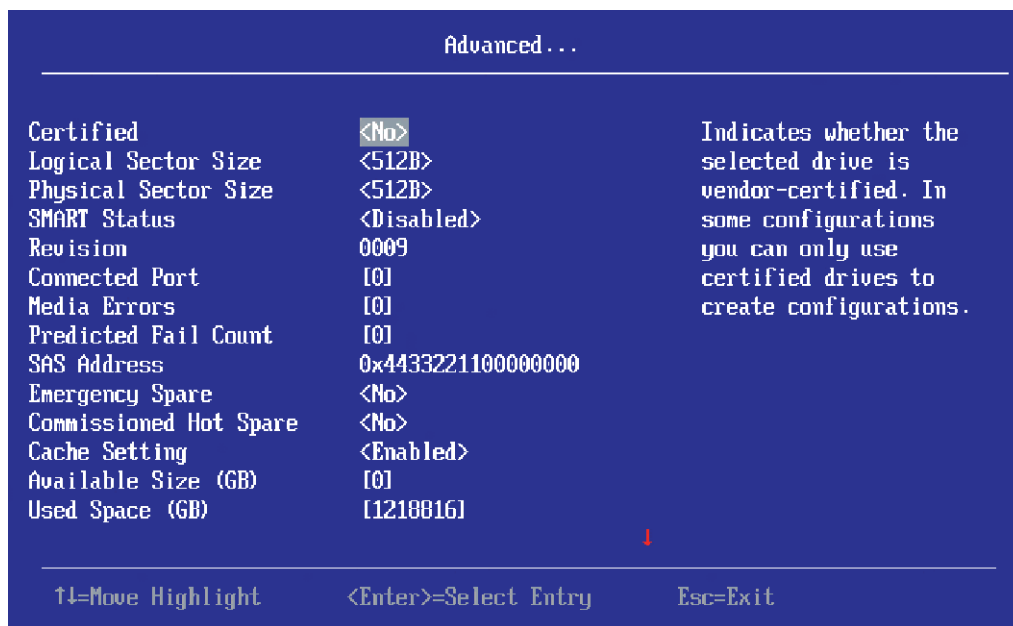


NOTE To refresh the progress indicator, press Esc to exit this window, then open it again.

8.2 Viewing Advanced Drive Properties

The following window appears when you select **Advanced** on the **Drive Management** window. The property information in this window is view-only and cannot be modified.

Figure 8.13 Advanced Drive Properties 1

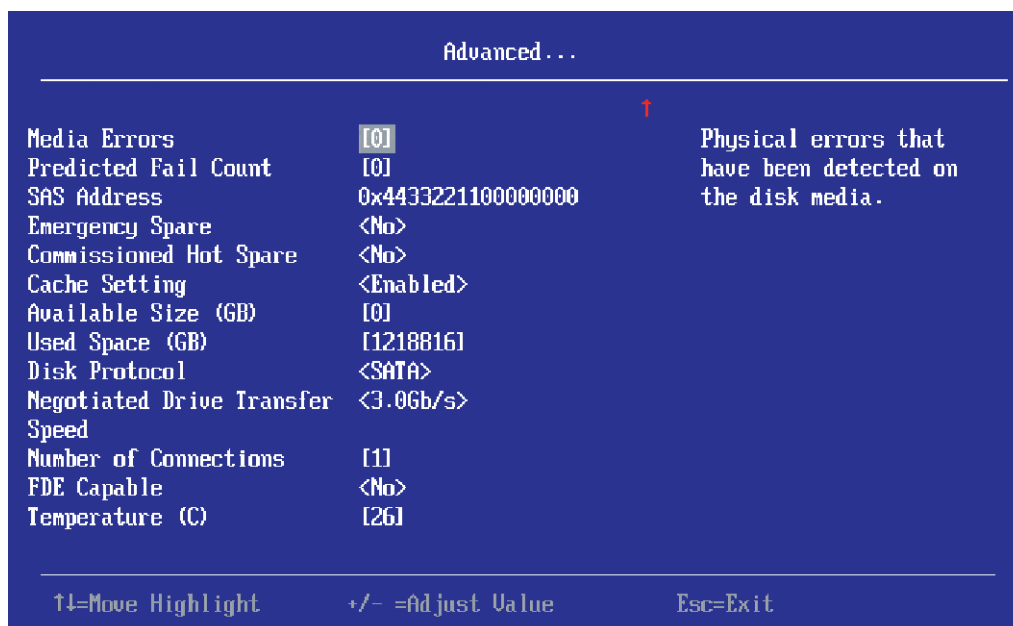


The small red arrow at the bottom of the window indicates that you can scroll down to view more physical drive properties, as shown in the following figure.



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 HLL browser.

Figure 8.14 Advanced Drive Properties 2



The following table describes all of the entries listed on the **Advanced Drive Properties** window.

Table 8.2 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 <i>4 KB</i> or <i>512 B</i> .
Physical Sector Size	The physical sector size of this drive. The possible options are <i>4 KB</i> or <i>512 B</i> .
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.
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.

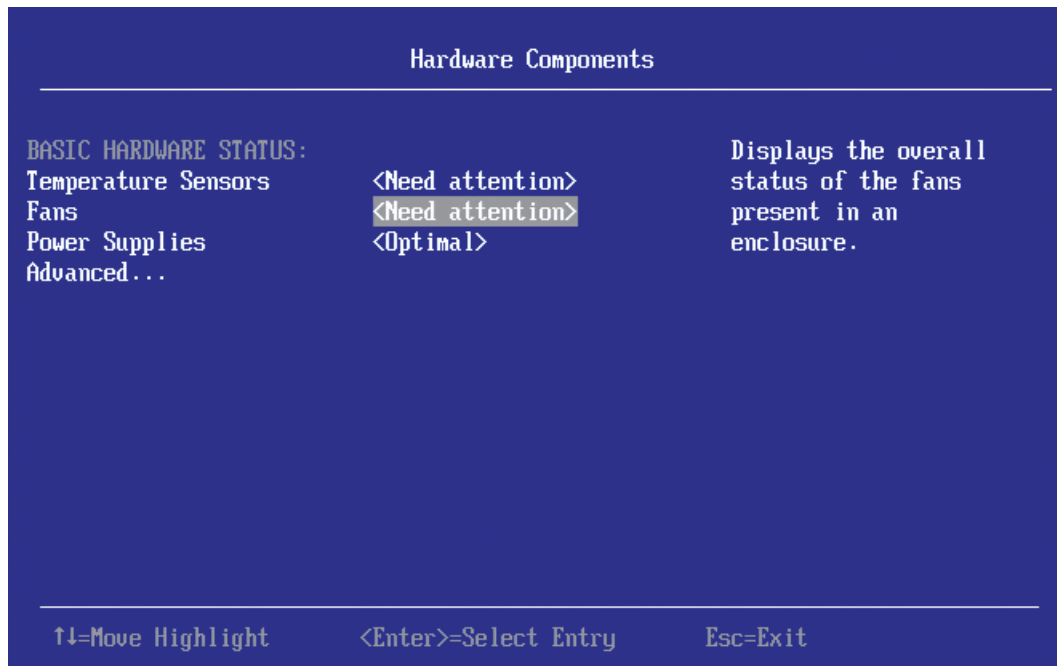
Table 8.2 Advanced Drive Properties (Continued)

Property	Description
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.

Chapter 9: Managing Hardware Components

When you select **Hardware Components** on the **Main Menu**, the **Hardware Components** menu appears, as shown in the following figure.

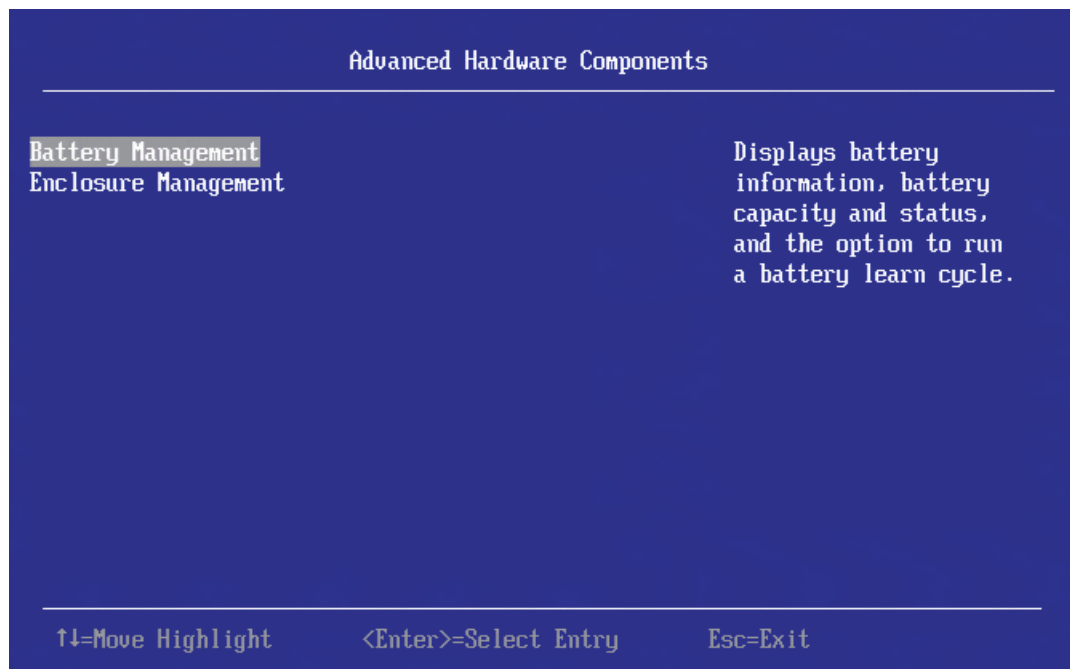
Figure 9.1 Hardware Components Menu



The window 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 window appears.

Figure 9.2 Advanced Hardware Components Menu

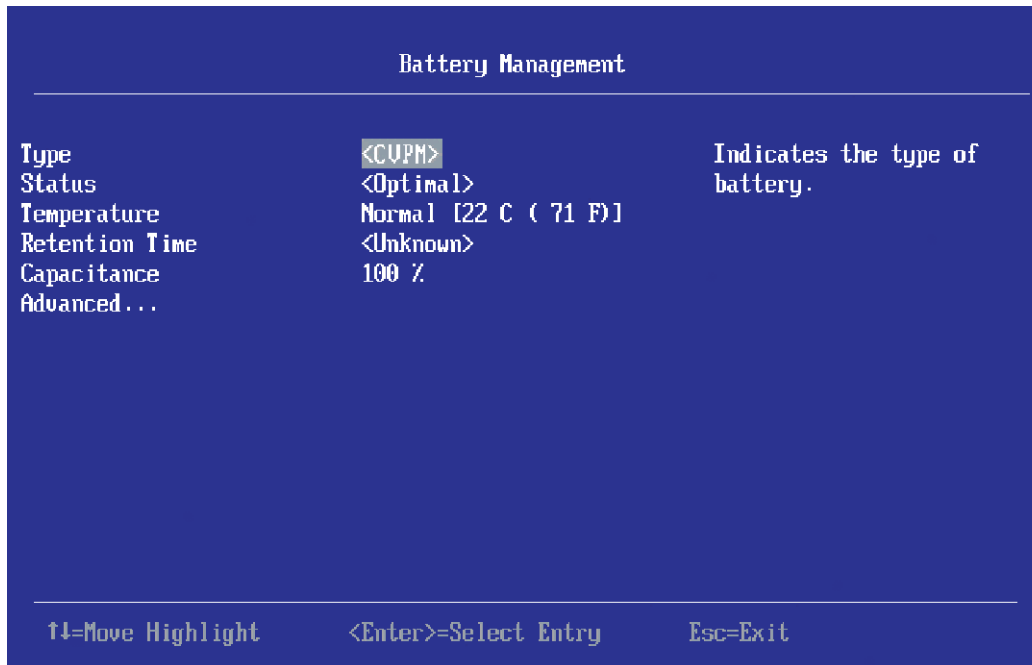


Select **Battery Management** or **Enclosure Management** to view more detailed information.

9.1 Managing Batteries

The following window appears when you select **Battery Management** on the **Advanced Hardware Components** menu.

Figure 9.3 Battery Management Window



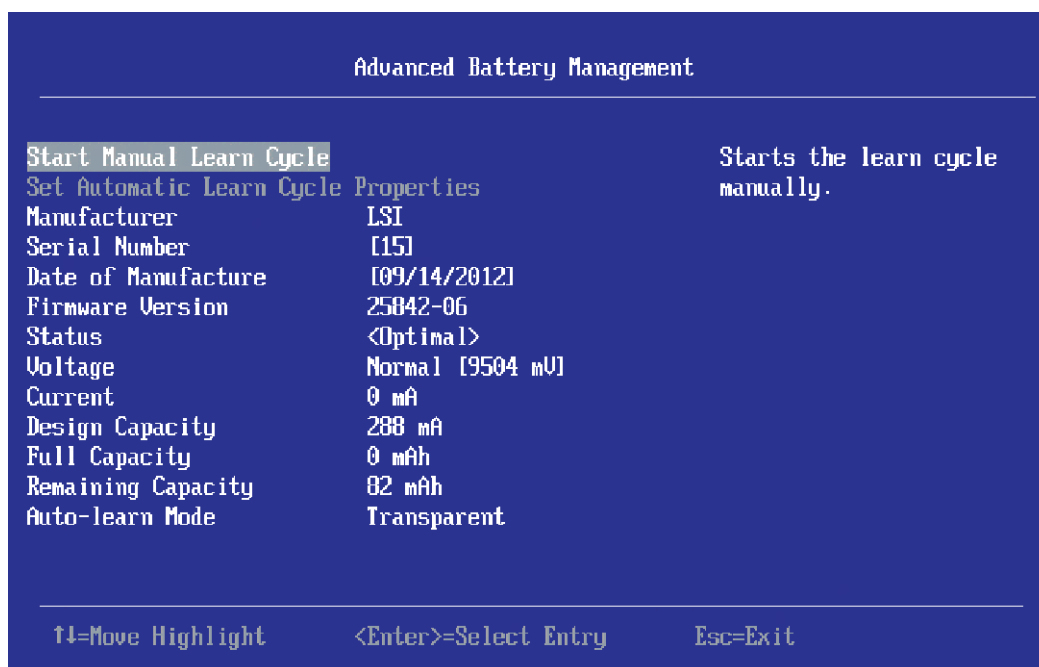
The following table describes the basic battery properties.

Table 9.1 Basic Battery Management Properties

Property	Description
Type	Type of the battery, such as Super Capacitor.
Battery Status	Current status of the battery, such as <i>Optimal</i> .
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 <i>48+ hours</i> , <i>Unknown</i> , 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 window appears.

Figure 9.4 Advanced Battery Management Window



The small red arrow at the bottom of the window 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 window. The amount of information that can be displayed in one window depends on the capabilities of the HII browser.

The following table describes the advanced battery properties and the other options on this window. Properties marked with an asterisk are user-selectable. All other properties are display only

Table 9.2 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. For details, see Section 9.1.1, Setting Automatic Learn Cycle Properties .
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.

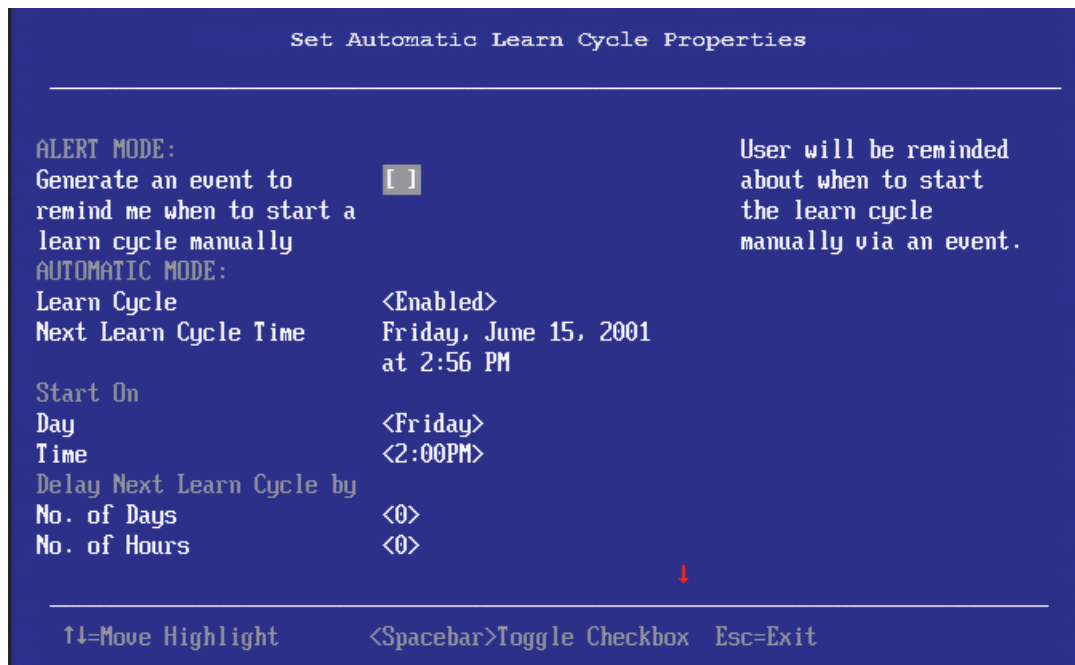
Table 9.2 Advanced Battery Management Properties (Continued)

Property	Description
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.

9.1.1 Setting Automatic Learn Cycle Properties

The following window appears when you select **Set Automatic Learn Cycle Properties** on the **Advanced Battery Management** window.

Figure 9.5 Set Automatic Learn Cycle Properties Window



The small red arrow at the bottom of the window 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 window. The amount of information that can be displayed in one window 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...**, as shown in the figure, 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 popup menu.

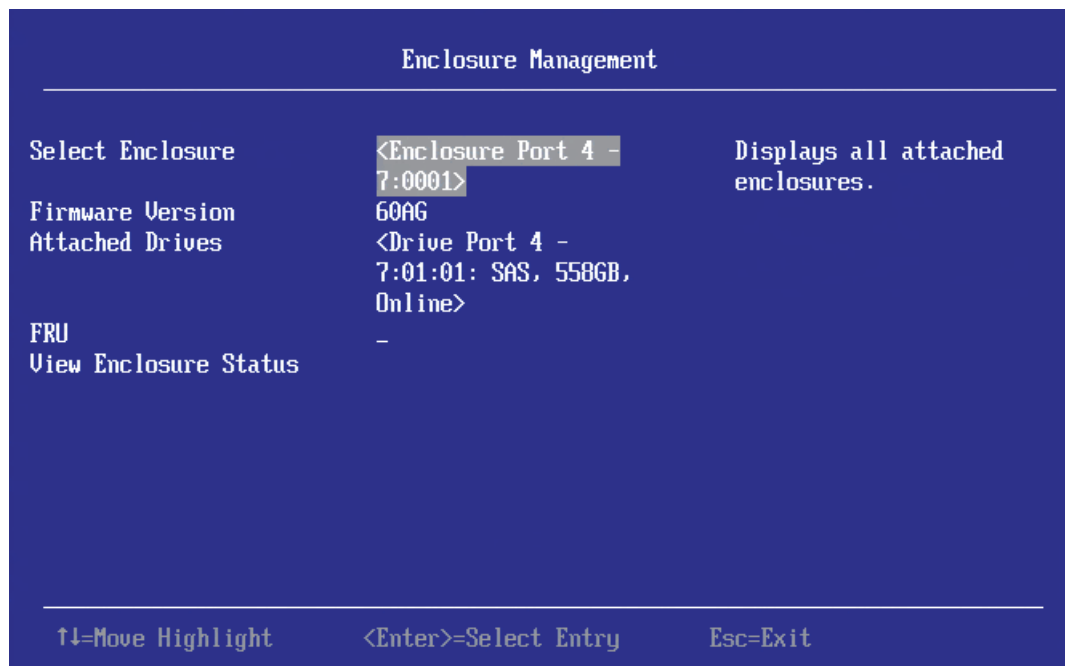
The **Day**, **Time**, **No. of Days**, and **No. of Hours** fields are also user-selectable through popup menus. The **Next Learn Cycle Time** field shows the time of the next learn cycle.

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.

9.2 Managing Enclosures

To manage enclosures and view enclosure properties, select **Enclosure Management** from the **Advanced Hardware Components** menu. The following window appears.

Figure 9.6 Enclosure Management Menu

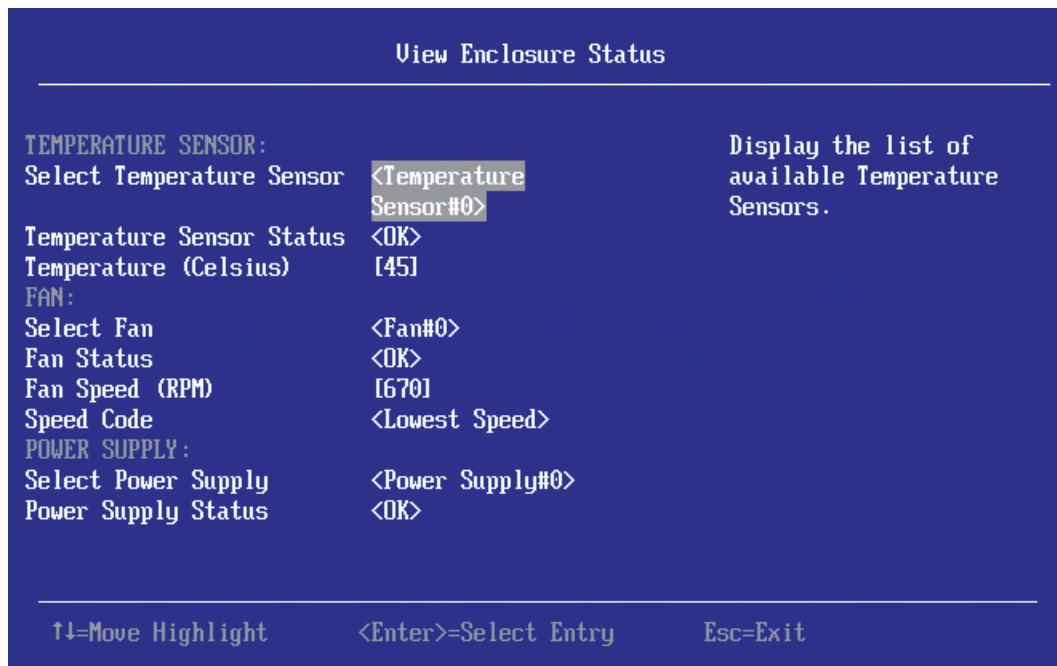


This window displays the firmware version and field replacement unit (FRU) number for the selected enclosure. To select a different enclosure, highlight the **Select Enclosure** field, press Enter, and select the enclosure from the popup menu.

To view a popup 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 window appears.

Figure 9.7 View Enclosure Status Window



The **View Enclosure Status** window shows information about the temperature sensors, fans, and power supplies installed in the selected enclosure. To view a selectable popup menu of all the installed sensors, fans, or power supplies, highlight the appropriate **Select** field and press Enter.

Chapter 10: ServeRAID StorCLI

10.1 Overview

The Storage Command Line Tool (StorCLI) is the command line management software designed for the ServeRAID® product line. The StorCLI is a command line interface that is designed to be easy to use, consistent, and easy to script. This document is the reference manual for installing and using the Storage Command Line Tool, and it explains the various features of the Storage Command Line Tool.



NOTE The legacy commands are deprecated from this guide.

10.2 Support for MegaCLI Commands

The MegaCLI commands can be executed on the Storage Command Line (StorCLI) tool. A single binary is output for the StorCLI commands and its equivalent MegaCLI commands. See [Appendix C: MegaCLI Commands to StorCLI Command Conversion](#) for the information for conversion from MegaCLI commands to StorCLI commands.

10.3 Installation

The ServeRAID 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 (32/64 bit)
- Red Hat® Enterprise Linux® 5.8 (32/64 bit)
- Red Hat Enterprise Linux 6.1
- Red Hat Enterprise Linux 6.2 (32/64 bit)
- SUSE® Linux Enterprise Server 11 SP2 (32/64 bit)
- SUSE Linux Enterprise Server 10 SP4 (32/64 bit)
- Fedora Core Linux 15
- VMware® ESX 4.0
- VMware ESX 4.1 U2
- VMware ESXi 4.1 U2
- VMware ESXi 5.0 U1
- Solaris
- FreeBSD (32/64 bit)
- EFI
- Ubuntu



NOTE The LSI SAS2208 and LSI SAS2108 controllers provide support for Microsoft Windows 8 and Microsoft Windows Server 2012 operating systems.

10.3.1 Installing StorCLI 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 IBM website.
2. Place the binary file in the directory from which you want to run the Storage Command Line Tool, and run the tool.



NOTE StorCLI must be run with the administrator privileges.

10.3.2 Installing StorCLI on Linux Operating Systems

To install StorCLI on Linux operating systems (except Ubuntu), perform the following steps:

1. Unzip the StorCLI package.
2. To install the StorCLI RPM, run the `rpm -ivh <StorCLI-x.xx-x.noarch.rpm>` command.
3. To upgrade the StorCLI RPM, run the `rpm -Uvh <StorCLI-x.xx-x.noarch.rpm>` command.

10.3.3 Installing StorCLI on Ubuntu Operating System



NOTE All commands must be run using the super user (sudo) login.

1. Run the command `sudo dpkg -i storcli_1.0_all.deb` to install the debian package.
2. Run the command `dpkg -l | grep -i storcli` for verifying if the debian package was installed successfully or not. If not, go back to step one.
3. To uninstall the debian package, run the command `sudo dpkg -r storcli`.

10.3.4 Installing StorCLI on VMware Operating Systems

To install StorCLI on VMware operating systems, run the following syntax 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
```

10.3.5 Installing StorCLI 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 IBM website.
2. Place the binary file in the directory from which you want to run the Storage Command Line Tool, and run the tool.

10.3.6 Installing StorCLI 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 IBM website.
2. Place the binary file in the directory from which you want to run the Storage Command Line Tool, and run the tool.

10.3.7 Installing StorCLI on Solaris Operating Systems

To install StorCLI on Solaris operating systems, run the following command:

```
pkgadd -d storcli.pkg
```

10.4 StorCLI 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.

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 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 10.1 Object Identifiers in the StorCli Command Syntax

Object Identifier	Description
No object identifier specified	If there is no object identifier, the command is a system command.
/cx	This object identifier is for controller x.
/cx/vx	This object identifier is for a virtual drive x on controller x.
/cx/vall	This object identifier is for all virtual drives on controller x.
/cx/ex	This object identifier is for an enclosure x on controller x.
/cx/eall	This object identifier is for all enclosures on controller x.
/cx/fx	This object identifier is for a foreign configuration x on controller x.
/cx/fall	This object identifier is for all foreign configurations on controller x.
/cx/ex/sx	This object identifier is for the drive is slotx on enclosure x on controller x.
/cx/sx	This object identifier represents the drives that are directly attached to controller x.
/cx/ex/sall	This object identifier is for all the drives on enclosure x on controller x.

Object Identifier	Description
/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 10.2 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.
suspend	This verb suspends an ongoing operation. A suspended operation cannot be resumed.
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.
shutdown	This verb will shutdown the adapter. All background operations are put on hold for resume. The controller cache is flushed, all disk drive caches are flushed.

- `<[adverb | attributes | properties] >` – Specifies what the verb modifies or displays.
- `<[key=value]>` – Specifies a value, if a value is required by the command.

10.5 Working with the Storage Command Line Tool

This chapter describes the commands supported by the Storage Command Line Tool..



NOTE The Storage Command Line Tool is not case sensitive.



ATTENTION The order in which you specify the command options should be the same as in the User Guide; otherwise, the command will fail.



NOTE The Storage Command Line Tool does not support the Snapshot feature.

10.5.1 System Commands

10.5.1.1 System Show Commands

The Storage Command Line 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 Tool.

10.5.2 Controller Commands

Controller commands provide information and perform actions related to the specified controller, such as the /c0 controller. The Storage Command Line Tool supports the controller commands described in this section.

10.5.2.1 Show and Set Controller Properties Commands

Table 10.3 Controller Commands Quick Reference Table

Commands	Value Range	Description
show <properties>	See Table 10.4	Shows specific controller properties.
set <properties>	See Table 10.4	Sets controller properties.
show	all: Shows all properties of the virtual drive. freespace: Shows the freespace in the controller. See Section 10.5.2.2 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 abortconerror
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 clusterenable
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 enableeeghsp
storcli /cx show enableesmarter
storcli /cx show enableeug
storcli /cx show exposeencldevice
storcli /cx show jbod
storcli /cx show loadbalancemode
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 preventpiimport
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 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 abortcconererror=<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 backplane=<value>
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 clusterenable=<value>
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 enableeeghsp=<on|off>
storcli /cx set enableesmarter=<value>
storcli /cx set enableeug=<on|off>
storcli /cx set exposeencldevice=<on|off>
storcli /cx set dimmerswitch|ds=<on|off type=1|2|3|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=<on|off>
storcli /cx set preventpiimport=<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>
```

The following table lists and describes the properties for the show and set commands.

Table 10.4 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.
backplane	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.
batterywarning	on off	Enables/disables battery warnings.
bgirate	0 to 100	Sets background initialization rate in percentage.
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 10.5.2.3.3 Consistency Check .	See 10.5.2.3.3 Consistency Check .

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 leak rate of the single-bit bucket in minutes (one entry removed per leak-rate).
eccbucketsize	0 to 255	Sets size of ECC single-bit-error bucket (logs event when full).
enableeghsp	on off	Enables/disables the commissioning of otherwise incompatible global hot spare drives as Emergency Hot Spare (EHSP) drives.
enableesmarter	on off	Enables/disables the commissioning of Emergency Hot Spare (EHSP) drives for Predictive Failure (PFA) events.
enableeug	on off	Enables/disables the commissioning of Unconfigured Good drives as Emergency Hot Spare (EHSP) drives.
exposeencldevice	on off	Enables/disables device drivers to expose enclosure devices; for example, expanders, SEPs.
dimmerswitch ds	See 10.5.8 Dimmer Switch Commands	See 10.5.8 Dimmer Switch Commands .
foreignautoimport	on off	Imports foreign configuration automatically, at boot.
jbod	on off	Enables/disables JBOD mode; by default, drives become system drives. Not supported by all controllers.
loadbalancemode	on off	Enables/disables automatic load balancing between SAS phys or ports in a wide port configuration.
maintainpdfailhistory	on off	Maintains the physical drive fail history.
migraterate	0 to 100	Sets data migration rate in percentage.
patrolread pr	See 10.5.2.3.2 Patrol Read .	See 10.5.2.3.2 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.
preventpiimport	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 patrol read rate of the virtual drives in percentage.
rebuildrate	0 to 100	Sets rebuild rate of the drive in percentage.
reconrate	0 to 100	Sets reconstruction rate for a drive in percentage.
restorehotspare	on off	Becomes a hot spare on insertion of a failed drive.

Property Name	Set Command Range	Description
smartpollinterval	0 to 65535	Set time for polling of SMART errors in seconds.
spinupdrivecount	0 to 255	Sets number of drives that are spun up at a time.
spinupdelay	0 to 255	Sets spin-up delay between a group of drives or a set of drives, in seconds.
stoponerror	on off	Stops the ServeRAID BIOS during POST, if any errors are encountered.
time	Valid time in <i>yy:mm:dd hh:mm:ss</i> format or <i>systemtime</i>	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.

10.5.2.2 Controller Show Commands

The Storage Command Line Tool supports the following show commands:

```
storcli /cx show
storcli /cx show all
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

This command shows all 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.

Input example:

```
storcli /c0 show all
```



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
```

10.5.2.3 Controller Background Tasks Operation Commands

10.5.2.3.1 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 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
```

10.5.2.3.2 Patrol Read

The Storage Command Line 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 suspend 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

Table 10.5 Set Patrolread 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.
uncfgareas	—	Include the areas not configured in the patrol read.

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 /cx set patrolread [delay=<value>]

This command delays the scheduled patrol read in hours.

Input example:

```
storcli /c0 set patrolread delay=30
```

storcli /cx show patrolRead

This command shows the progress on the current patrol read in percentage.

Input example:

```
storcli /c0 show patrolread
```

storcli /cx start patrolread

This command starts the patrol read operation. This command starts a patrol read immediately.

Input example:

```
storcli /c0 start patrolread
```

storcli /cx 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 /cx suspend patrolread

This command pauses a running patrol read operation.

Input example:

```
storcli /c0 suspend patrolread
```



NOTE You can run this command only when a patrol read operation is running on the controller.

10.5.2.3.3 Consistency Check

The Storage Command Line 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 10.6 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 (Vx,Vy ... format) or the range of virtual drives that you want to exclude from a consistency check (Vx-Vy 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.

10.5.2.4 Premium Feature Key Commands

The Storage Command Line 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 10.7 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
```

10.5.2.5 Controller Security Commands

The Storage Command Line 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 /cx compare securitykey=ssssss

This command compares and verifies the security key of the controller.

storcli /cx delete securitykey

This command deletes the security key of the controller.

Input example:

```
storcli /c0 delete securitykey
```

storcli /cx set securitykey keyId=kkkk

This command sets the key ID for the controller. The key ID is unique for every controller.

storcli /cx 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 10.8 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=IBM@12345 passphrase=IBM@123456 keyid=1
```

storcli /cx set securitykey=sssss oldsecuritykey=ssss [passphrase=sssss][keyid=sssss]

This command changes the security key for the controller.

Input example:

```
storcli /c0 set securitykey=IBM@12345 oldsecuritykey=pass123  
passphrase=IBM@123456 keyid=1 ?
```

10.5.2.6 Flashing Controller Firmware Command



NOTE The Flashing Controller Firmware command is not supported in Embedded MegaRAID.

The following command flash the controller firmware.

storcli /cx download file=filepath [fwtype=<value>] [nosigchk] [noverchk] [resetnow]

This command flashes the firmware to the specified adapter from the given file location (*filepath* is the absolute file path). You can use the following options when you flash the firmware:

Table 10.9 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	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.

10.5.2.7 Controller Cache Command

The following command flushes the controller cache.

storcli /cx flush|flushcache

This command flushes the controller cache.

Input example:

```
storcli /c0 flushcache
```

10.5.3 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 10.10 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 10.5.3.1 Drive Show Commands .	Shows virtual drive information.

10.5.3.1 Drive Show Commands

The Storage Command Line Tool supports the following drive show commands:

```
storcli /cx[/ex]/sx show
storcli /cx[/eall]/sall show
storcli /cx[/ex]/sx|sall show all
```

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/s4,5 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.

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.

10.5.3.2 Missing Drives Commands

The Storage Command Line 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*.

10.5.3.3 Set Drive State Commands

The Storage Command Line 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 drive changes the drive state to unconfigured good. If the drive has the operating system in it, use the *force* option.

Input example:

```
storcli /c1/e56/s3 set good
```

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
```

10.5.3.4 Drive Initialization Commands

When you initialize drives, all the data from the drives is cleared. The Storage Command Line 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.

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
```

10.5.3.5 Drive Firmware Download Commands

The Storage Command Line Tool supports the following command to download drive firmware:

storcli /cx[/ex]/sx download src=filepath [satabridge]

This command flashes the firmware with the specified file. The `satabridge` option lets you download the SATA bridge firmware in online mode.

Input example:

```
storcli /c0/e56/s1 download src=c:\file1.bin
```

10.5.3.6 Locate Drives Commands

The Storage Command Line 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
```

10.5.3.7 Prepare to Remove Drives Commands

The Storage CLI 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
```

10.5.3.8 Drive Security Command

The Storage Command Line Tool supports the following drive security command:

```
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
```

10.5.3.9 Drive Secure Erase Commands

The Storage Command Line 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] [erasepatternA=<value1>]  
[erasepatternB=<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] [erasepatternA=<val1>] [erasepatternB=<val2>]

This command securely erases non-SED drives. The drive is written with erase patterns to ensure that the data is securely erased. You can use the following options with the start erase command:

Table 10.11 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.
erasepatternA	8-bit value	Erase pattern A to overwrite the data.
erasepatternB	8-bit value	Erase pattern B to overwrite the data.

Input example:

```
storcli /c0/e25/s1 start erase thorough erasepatternA=10010011  
erasepatternB=11110000
```

10.5.3.10 Rebuild Drives Commands

The following commands rebuild drives in the Storage Command Line 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 /cx[/ex]/sx 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 /cx[/ex]/sx 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 /cx[/ex]/sx show rebuild

This command shows the progress of the rebuild process in percentage.

Input example:

```
storcli /c0/s5 show rebuild
```

storcli /cx[/ex]/sx start rebuild

This command starts a rebuild operation for a drive.

Input example:

```
storcli /c0/s4 start rebuild
```

storcli /cx[/ex]/sx 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
```

10.5.3.11 Drive Copyback Commands

The Storage Command Line 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.

storcli /cx[/ex]/sx 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 /cx[/ex]/sx 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 /cx[/ex]/sx show copyback

This command shows the progress of the copyback operation in percentage.

Input example:

```
storcli /c0/e25/s4 show copyback?
```

storcli /cx[/ex]/sx 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 /cx[/ex]/sx 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.

10.5.3.12 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 /cx[/ex]/sx 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 10.12 Add Hotsparedrive 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
```

10.5.4 Virtual Drive Commands

The Storage Command Line Tool supports the following virtual drive commands. The following table describes frequently used virtual drive commands.

Table 10.13 Virtual Drives Commands Quick Reference Table

Commands	Value Range	Description
add	See 10.5.4.1 Add Virtual Drives Commands and	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 10.5.3.3 Set Drive State Commands and 10.5.4.5 Change Virtual Properties Commands	Sets virtual drive properties.
show	all: Shows all properties of the virtual drive. cc: Shows properties of CacheCade virtual drives. See 10.5.4.3 Virtual Drive Show Commands .	Shows virtual drive information.

10.5.4.1 Add Virtual Drives Commands

The Storage Command Line Tool supports the following commands to add virtual drives:

```
storcli /cx add vd type=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] [nora|ra] [direct|cached] [CachedBadBBU|NoCachedBadBBU] [cachevd]
[Strip=<8|16|32|64|128|256|1024>] [AfterVd=X] [Spares = [e:]s|[e:]s-x|[e:]s-x,y]
[force] [ExclusiveAccess]
```

```
storcli /cx add vd each type=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] [nora|ra]
[direct|cached]
[CachedBadBBU|NoCachedBadBBU] [Strip=<8|16|32|64|128|256|1024>] [ExclusiveAccess]
```

```
storcli /cx add VD cachecade|nytrocache Type = [raid|r][0,1,10, R1EC] drives =
[e:]s|[e:]s-x|[e:]s-x,y [WT| WB] [assignvds = 0,1,2 [BOOTVOLSIZE=x]
```

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 type=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][nora|*ra] [*direct|cached] [CachedBadBBU|*NoCachedBadBBU] [Strip=<8|16|32|64|128|256|1024>]
[AfterVd=X] [Spares = [e:s|e:s-x|e:s-x,y] [force]
```

Table 10.14 Add RAID Configuration Input Options

Option	Value Range	Description
type	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: <ul style="list-style-type: none"> 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.
pdperarray	1-16.	Specifies the number of physical drives per array. The default value is automatically chosen.
sed	—	Creates security-enabled drives.
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.
wt wb	wt: Write through.wb: 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.
cachedbadbbu nocachedbadbbu	cachedbadbbu: Enable bad BBU caching. nocachedbadbbu: Disable bad BBU caching.	Enables caching when BBU is not functioning. Disabled is the default.
cachevd	—	Enables SSD caching on the created virtual drive.
strip	8, 16, 32, 64, 128, 256, 512, 1024.	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 type=raid10 size=2gb,3gb,4gb names=tmp1,tmp2,tmp3
drives=252:2-3,5,7 pdperarray=2
```

storcli /cx add vd cc|cachecade type=[0,1,10] drives=[e:s][e:s-x][e:s-x,y [[wt]*wb]] [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 10.15 Add RAID Configuration Input Options

Option	Value Range	Description
cachecade	—	Creates a CacheCade virtual drive.
type	0, 1, 10	Sets the RAID type of the CacheCade virtual drive.
drives	Valid enclosure number and valid slot number	See the <code>drives</code> row in the previous table for format.
wt *wb	wt: Enables write through. wb: Enables 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 type=raid10 size=2gb,3gb,4gb names=tmp1,tmp2,tmp3
drives=252:2-3, 7
```

10.5.4.2 Delete Virtual Drives Commands

The Storage Command Line 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
```



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
```



NOTE 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
```



NOTE 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
```



NOTE 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.

10.5.4.3 Virtual Drive Show Commands

The Storage Command Line Tool supports the following virtual drive show commands:

```
storcli /cx/vx show  
storcli /cx/vx show all
```

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

This command shows all virtual drive information, which includes virtual drive information, physical drives used for the virtual drives, and virtual drive properties.

Input example:

```
storcli /c0/v0 show all
```

10.5.4.4 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 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
```

10.5.4.5 Change Virtual Properties Commands

The Storage Command Line 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
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>
```

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 /cx/vx 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 /cx/vx 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 /cx/vx 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 /cx/vx set rdcache=<ra|nora>

This command sets the read cache policy on a virtual drive to read ahead, no read ahead, or adaptive read ahead.

Input example:

```
storcli /c0/v0 set rdcache=nora
```

storcli /cx/vx|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, then 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
```

10.5.4.6 Virtual Drive Initialization Commands

The Storage Command Line 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.

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
```

10.5.4.7 Virtual Drive Erase Commands

The Storage Command Line 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
```

10.5.4.8 Virtual Drive Migration Commands



NOTE The virtual drive migration commands are not supported in Embedded MegaRAID.

The Storage Command Line Tool supports the following commands for virtual drive migration (reconstruction):

```
storcli /cx/vx show migrate  
storcli /cx/vx start migrate <type=raidlevel>  
[option=<add|remove> disk=<e1/s1,e2/s2 ...> ]
```

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.

Input example:

```
storcli /c0/v0 show migrate
```

storcli /cx/vx start migrate <type=raidlevel> [option=<add | remove> disk=<e1:s1,e2:s2 ...>]

This command starts the reconstruction on a virtual drive to the specified RAID level by adding or removing disks from the existing virtual drive. You can use the following options with the start migrate command:

Table 10.16 Virtual Drive Migration Command Options

Options	Value Range	Description
<code>type =RAID level</code>	RAID [0 1 5 6]	The RAID level to which the virtual drive must be migrated.
<code>[option=<add remove> disk=<e1:s1,e2:s2, ...>]</code>	add: Adds disks to the virtual drive and starts reconstruction. remove: Removes disks from the virtual drive and starts reconstruction. disk: The enclosure number and the slot number of the disks to be added to the virtual drive.	Adds or removes disks from the virtual drive.

Virtual drive migration can be done between the following RAID levels.

Table 10.17 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:

```
storcli /c0/v3 start migrate type=r5 option=add disk=e5:s2,e5:s3
```

10.5.4.9 Virtual Drive Consistency Check Commands

The Storage Command Line 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.

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.

10.5.4.10 Background Initialization Commands

The Storage Command Line 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.

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.

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
```

10.5.4.11 Virtual Drive Expansion Commands

The Storage Command Line 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. The value of the expand size is in GB. 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
```

10.5.5 Foreign Configurations Commands

The Storage Command Line 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
```

10.5.6 BIOS-Related Commands

The Storage Command Line Tool supports the following BIOS commands:

```
storcli /cx set autobootselect(abs)=<on|off>
storcli /cx set bios=<on|off>
storcli /cx set BIOSMode=<SOE|BE|HCOE|HSM>
```

The detailed description for each command follows.

storcli /cx set autobootselect|abs=<on|off>

This command enables the BIOS to select the best logical drive as the boot drive.

Input example:

```
storcli /cx set autobootselect=on
```

storcli /cx set bios=<on|off>

This command enables or disables the ServeRAID controller's BIOS.



NOTE The legacy BIOS can load a limited number of the PCI device's BIOS. Disable the ServeRAID BIOS to avoid issues during POST.

Input example:

```
storcli /c0 set bios=enable
```

storcli /cx set BIOSMode=<SOE|BE|HCOE|HSM

This command sets the BIOS boot mode.

Input example:

```
storcli /c0/ set BIOSMode=SOE
```

10.5.6.1 OPRM BIOS Commands

The Storage Command Line Tool supports the following OPRM 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
```

10.5.7 Drive Group Commands

This section describes the drive group commands.

10.5.7.1 Drive Group Show Commands

The Storage Command Line 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/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
```

10.5.8 Dimmer Switch Commands

10.5.8.1 Change Virtual Drive Power Settings Commands

The Storage Command Line Tool supports the following command to change the Dimmer Switch® setting. The Dimmer Switch is the power-saving policy for the virtual drive.

storcli /cx/vx set ds=<default | auto | none | max | maxnocache>

This command changes the power-saving properties on a virtual drive. See `dimmerswitch` in the following table for values.

Input example:

```
storcli /cx/vx set ds=default
```



NOTE Only the ds3 dimmer switch option cannot be selected in the Storage Command Line Tool.

You can use the following combinations for the dimmer switch commands:

```
storcli /cx set ds=off type=1|2|3|4
storcli /cx set ds=on type=1|2 [properties]
storcli /cx set ds=on type=3|4 defaultldtype=<value> [properties]
storcli /cx set ds=on [properties]
```

The following table describes the power-saving options.

Table 10.18 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 3: Virtual drive 4: All	Specifies the type of drives that the dimmer switch feature is applicable. By default, it is activated for unconfigured drives, hot spare drives and virtual 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 <i>hh:mm</i> 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
```

10.5.9 BBU Commands

The Storage Command Line Tool supports the following battery backup unit (BBU) commands:

```
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 show modes
storcli /cx/bbu show properties
storcli /cx/bbu show status
storcli /cx/bbu start learn
```

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 10.19 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 show modes

This command shows the bbu mode information that includes the bbu mode number, retention time, service life, maximum temperature, and battery learn information.

Input example:

```
storcli /c0/bbu show modes
```

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
```

10.5.10 Enclosure Commands

The Storage Command Line 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 10.20 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
```

10.5.11 PHY Commands

The Storage Command Line 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
```

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
```

10.5.12 Logging Commands

The Storage Command Line 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
```

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

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.

Input example:

```
storcli /c0 show termlog type=contents
```

10.6 Frequently Used Tasks

10.6.1 Showing the Version of the Storage Command Line Tool

The following command shows the version of the command line tool:

```
Storcli -v
```

10.6.2 Showing StorCLI Help

The following command shows the command line tool help:

```
Storcli -h
```

Help appears for all the StorCLI commands.

10.6.3 Showing System Summary Information

The following command shows the summary of all the controller information:

```
Storcli -show [all]
```

10.6.4 Showing Free Space in a Controller

The following command shows the free space available in the controller:

```
Storcli /cx show freespace
```

10.6.5 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] [nora|*ra]  
[*direct|cached] [CachedBadBBU|*NoCachedBadBBU]  
[strip=<8|16|32|64|128|256|512|1024] [AfterVd=x] [Spares=[e:]s|[e:]s-x|[e:]s-x,y]  
[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), orMaximumWithoutCaching(maxnocache).
- The wt option disables write back.
- The nora option disables read ahead.
- The cached option enables the cached memory.
- The CachedBadBBU option enables caching when bbu is not functional.
- The strip option sets the strip size. It can take the values 8, 16, 32, 64, 128, 256, 512, 1024.
- 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.

10.6.6 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.

10.6.7 Showing Virtual Drive Information

The following command shows the virtual drive information for all the virtual drives in the controller:

```
storcli /cx show [all]
```

10.6.8 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.

10.6.9 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 [10.5.2.6 Flashing Controller Firmware Command](#).

Chapter 11: MegaRAID Storage Manager Overview

This chapter provides a brief overview of the MegaRAID Storage Manager software and explains how to install it on the supported operating systems.

11.1 Overview

The MegaRAID Storage Manager software enables you to configure, monitor, and maintain storage configurations on IBM SAS controllers. The MegaRAID Storage Manager graphical user interface (GUI) makes it easy for you to create and manage storage configurations.

11.1.1 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.

11.1.2 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.

11.1.3 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.

11.2 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 optimal)
- 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, and Microsoft Windows 8
- Red Hat Linux 3.0, 4.0, 5.0, 5.8, and 6.0. The MegaRAID Storage Manager software supports 64-bit environment from RHEL 6 onwards.
- Solaris 10 x86, Solaris SPARC, Solaris 11 x86, Solaris 11 SPARC
- SUSE Linux/SLES 9, 10, 11, and 11 SP2 with the latest updates and service packs
- VMware ESX 4.0 and 4.1
- VMware ESXi 4.0, 4.1, 5.0, and 5.1
- Citrix XenServer 6.0

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 is supported in the Network Address Translation (NAT) environment also. 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. Please do a clean un-installation and only then install the MegaRAID Software Manager to avoid any conflicts.

11.3 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.

11.3.1 Prerequisite for MegaRAID Storage Manager Installation

The MegaRAID Storage Manager software installation script also installs the IBM SNMP agent, Red Hat Package Manager (RPM). The IBM 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.

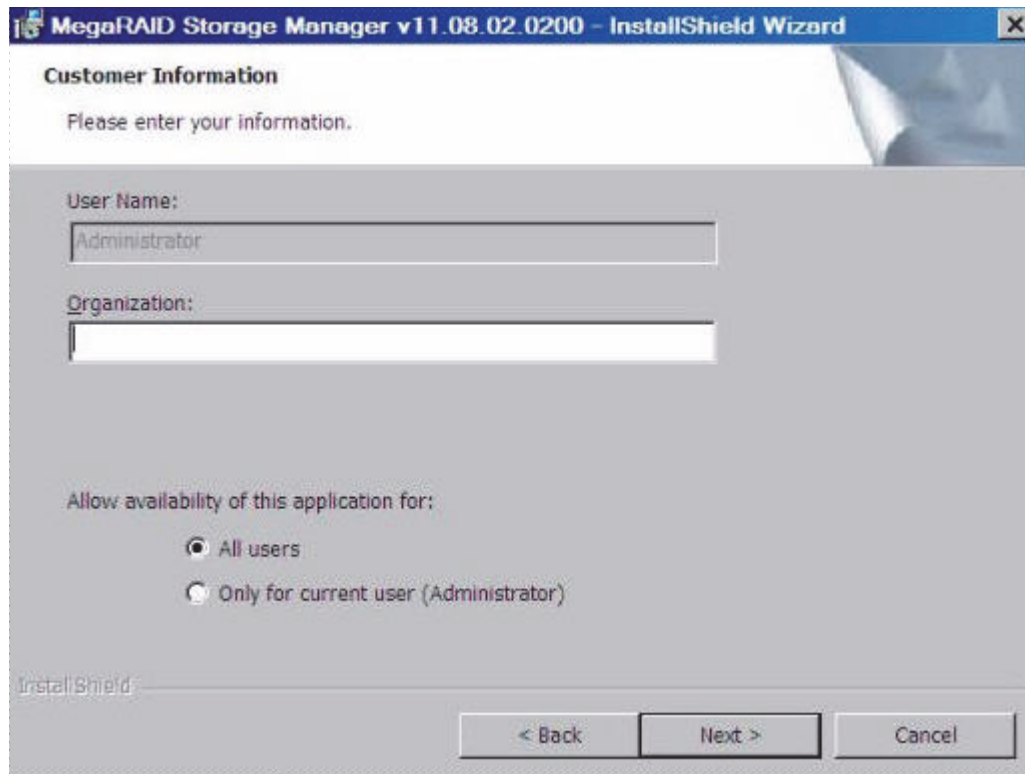
The SNMP-Util package includes the net-snmp-libs and the net-snmp-utils RPMs and additional dependent RPMs. Make sure that these RPMs are installed from the operating system media before you install the MegaRAID Storage Manager software.

11.3.2 Installing MegaRAID Storage Manager Software on Microsoft Windows

To install the MegaRAID Storage Manager software on a system running the Microsoft Windows Server 2003, Microsoft Windows Server 2008, Microsoft Server 2008 R2, Microsoft Windows XP, Microsoft Windows Vista, Microsoft Windows 7, or Microsoft Windows 8 operating system, perform the following steps:


1. Insert the MegaRAID Storage Manager software installation CD in the CD-ROM drive.
If necessary, find and double-click the `setup.exe` file to start the installation program.
2. In the **Welcome** screen that appears, click **Next**.
If the MegaRAID Storage Manager software is already installed on this system, then an upgraded installation occurs.
3. Read and accept the user license and click **Next**.
The **Customer Information** window appears, as shown in the following figure.

Figure 11.1 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.

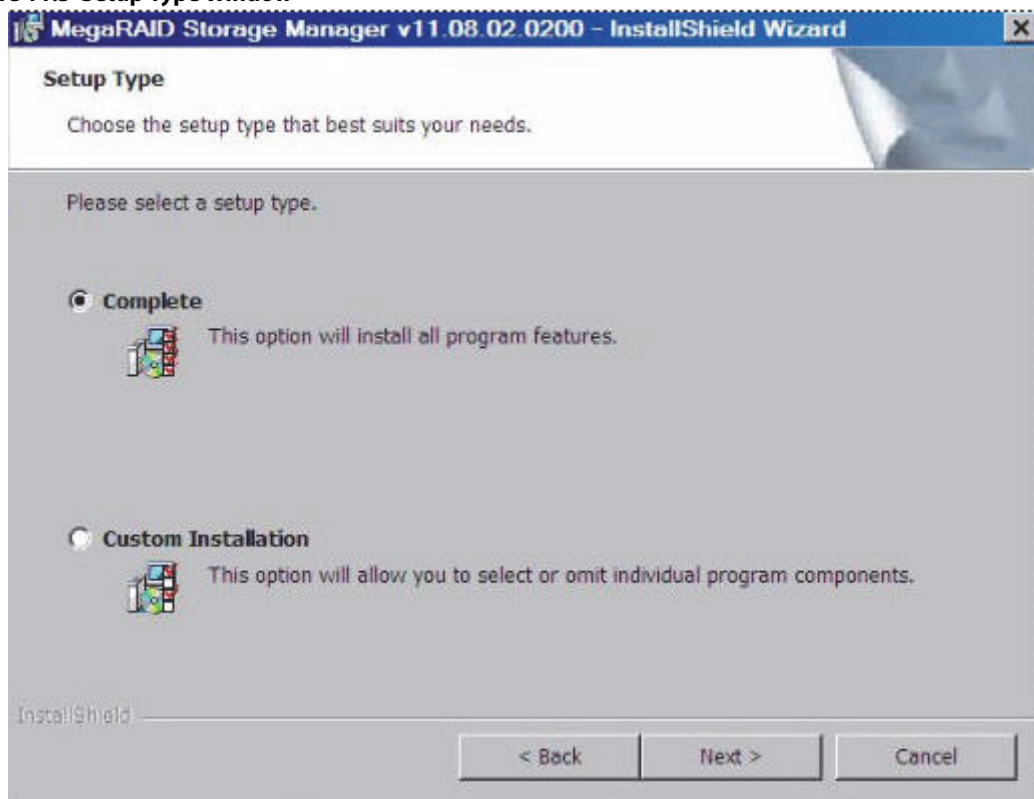
Figure 11.2 Destination Folder Window



7. Click **Next** to continue.

The **Setup Type** window appears, as shown in the following figure.

Figure 11.3 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 [Figure 11.3](#). If you select **Complete** as your setup option, the LDAP Login Information appears.

Figure 11.4 LDAP Logon Information

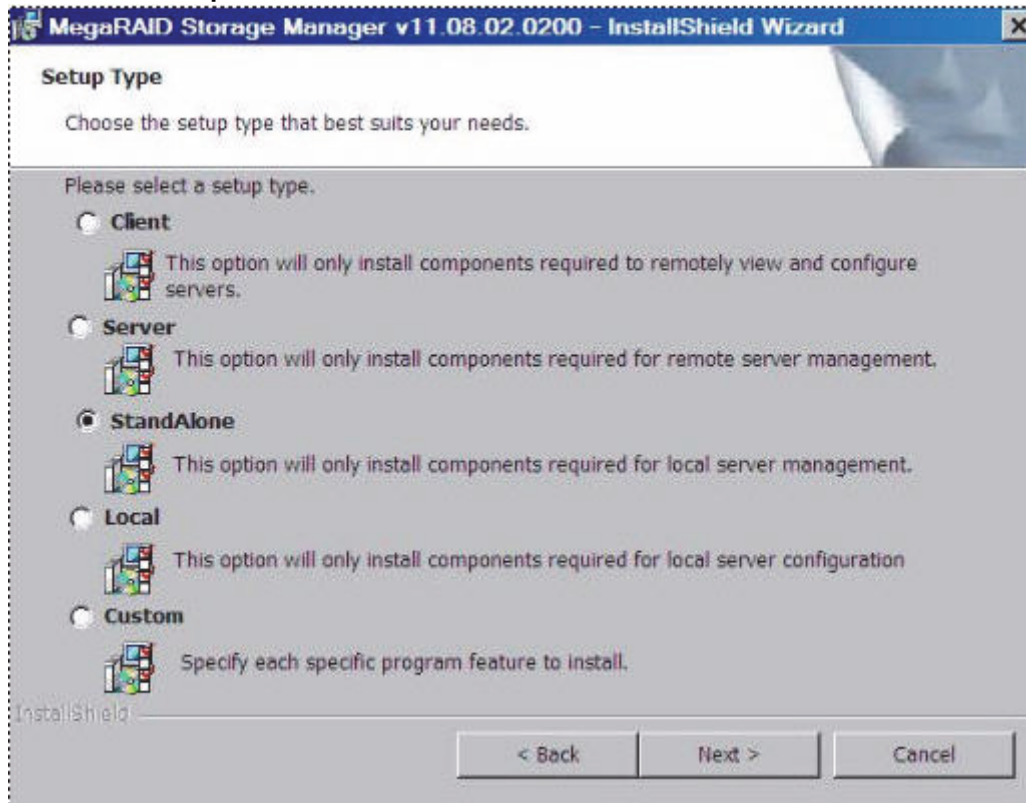
The screenshot shows a Windows-style dialog box titled "MegaRAID Storage Manager v11.12.00.0100 - InstallShield Wizard". The main heading is "LDAP Logon Information" with the subtitle "Specify LDAP Login Details". Inside the dialog, there is a section asking "Do you wish to specify ldap configuration details?" with two radio buttons: "Yes" (which is selected) and "No". Below this are four text input fields: "Server IP:", "User name:", "Distinguished User name:", and "Port:". To the right of the "Port:" field is a checkbox labeled "Use LDAP as default Login". At the bottom left, the "InstallShield" logo is visible. At the bottom right, there are three buttons: "< Back", "Next >", and "Cancel".

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**.
12. In the dialog that appears, click **Install** to begin the installation.
13. Select one of the setup options. See [Section 11.3.2.1, Setup Options](#) for specific information.

Figure 11.5 Custom Setup Window



14. Click **Next** to proceed.

15. Click **Install** to install the program.

16. 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.

11.3.2.1 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 on **Install** on the next window that appears.

- Select the **StandAlone** radio button if you will use the MegaRAID Storage Manager software to create and manage storage configurations on a stand-alone workstation. To begin installation, click on **Install** on the next window that appears.



NOTE If you select Client or Standalone as your setup option, the LDAP Logon Information dialog 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.

11.3.3 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 via the Control Panel, the Command Prompt, or the MegaRAID Storage Manager Uninstallation Utility.

11.3.3.1 Uninstalling MegaRAID Storage Manager Software through 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**.

11.3.3.2 Uninstalling MegaRAID Storage Manager Software Using 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).

11.3.3.3 Uninstalling 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.

11.3.4 Installing 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.

11.3.5 Uninstalling MegaRAID Storage Manager Software on the Solaris 10 x86 Operating System

Follow these steps to uninstall the MegaRAID Storage Manager software on a system running the Solaris 10 x86 operating system:

1. Execute the `Uninstaller.sh` file located in `/opt/MegaRaidStorageManager` directory.
2. When prompted by the uninstallation scripts, select Y to complete the installation.

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.

11.3.6 Installing 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 IBM ServeRAID Pro 2.0 software is not applicable in SPARC.

11.3.7 Uninstalling MegaRAID Storage Manager Software on the Solaris SPARC Operating System

To uninstall the MegaRAID Storage Manager software on a system running Solaris 10 SPARC, perform the following steps:

1. Execute the `Uninstaller.sh` file located in `/opt/MegaRaidStorageManager` directory.
2. When prompted by the uninstallation scripts, select Y to complete the installation.

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.

11.3.8 Prerequisites for Installing MegaRAID Storage Manager on the RHEL6.X x64 Operating System

Before installing the MegaRAID Storage Manager software on RHEL 6.X x64 system, install the following RPMs. Without these RPMs the MegaRAID Storage Manager software might not install properly or might not work as expected.

- `libstdc++-4.4.4-13.el6.i686.rpm`
- `compat-libstdc++-33-3.2.3-69.i686.rpm`
- `libXau-1.0.5-1.el6.i686.rpm`
- `libxcb-1.5-1.el6.i686.rpm`
- `libX11-1.3-2.el6.i686.rpm`
- `libXext-1.1-3.el6.i686.rpm`
- `libXi-1.3-3.el6.i686.rpm`
- `libXtst-1.0.99.2-3.el6.i686.rpm`

The RHEL6.X x64 complete operating system installation is required for the MegaRAID Storage Manager software to work. The above mentioned RPMs come as part of RHEL6.X x64 Operating System DVD. These RPMs might need additional dependent RPMs as well, and you must install all the dependent RPMs on the target system.



NOTE The RPMs versions mentioned above may get changed in the future RHEL6.x releases. Install the corresponding RPMs from the operating system installation media



NOTE The MegaRAID Storage Manager software now provides an additional binary to run it in a native 64-bit Linux environment.

11.3.9 Installing 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 Red Hat Linux 3.0, 4.0, 5.0, 6.0 or SUSE Linux/SLES 9, 10, and 11:



NOTE For installing the MegaRAID Storage Manager software on a SLES 64-bit platform, you need to create certain symbolic links that are mentioned in [Section 11.3.12.1, Executing 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:

```
csch install.csh -a
```

The preceding command works only if 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 installs all the features of the product.
- **Client Components Only** – The storelib feature of the product is not installed in this type of installation. As a result, the resident system can only administer and configure all of the servers in the subnet, but it cannot serve as a server.
- **StandAlone** – Only the networking feature will not be installed in this case. But the system can discover other servers in the subnet and can be discovered by the other servers in the subnet.
- **Local** – This option enables you to view only the workstation that has the MegaRAID Storage Manager software installed. You will not be able to discover other remote servers and other remote servers will also not be able to connect to your workstation. In a local mode installation, you will be using the loopback address instead of the IP address.
- **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 need to disable SELinux.

11.3.10 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 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.

11.3.11 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
```

11.3.12 Uninstalling MegaRAID Storage Manager Software on RHEL or SLES/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.

11.3.12.1 Executing a CIM Plug-in on Red Hat Enterprise Linux 5

To execute a Common Information Model (CIM) plug-in on Red Hat Enterprise Linux 5, you must create the following symbolic links:

1. `cd /usr/lib` on RHEL 5
2. Search for `libcrypto`, `libssl`, and `libsysfs` libraries as follows:

```
ls -lrt libcrypto*, ls -lrt libssl*, ls -lrt libsysfs*
```
3. If the files `libcrypto.so.4`, `libssl.so.4`, and `libsysfs.so.1` are missing, manually create sym links as follows:

```
ln -s libcrypto.so libcrypto.so.4
ln -s libssl.so libssl.so.4
ln -s libsysfs.so libsysfs.so.1
```

For more information about CIM, see [Section 11.4.4, MegaRAID Storage Manager Support on the VMware ESXi Operating System](#).

If the `.so` files are not present in the `/usr/lib` directory, create a link with the existing version of the library. For example, if `libcrypto.so.6` is present and `libcrypto.so` is not, create the link as follows:

```
ln -s libcrypto.so.6 libcrypto.so.4
```

On a 64-bit operating system, the system libraries are present in the `/usr/lib64` directory by default. However, for supporting CIM Plug-in, make sure that the libraries are also present in `/usr/lib` by installing the appropriate RPMs.

11.3.13 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:

- a. CHANNELLOGO=
- b. CHANNELSPLASHSCREEN=

No default values are assigned for these keys; therefore, the MegaRAID Storage Manager uses the default IBM Logo and splash screen.

To customize the Logo and splash screen, enter the Logo and Splash screen file name against these entries.

To enter the file names follow these steps:

1. Open the `msm.properties` file in the `<installation-directory\MegaRAID Storage Manager>` folder.
2. Enter the value for the logo file against the key CHANNELLOGO.
3. Enter the value for the splash screen file against the key CHANNELSPLASHSCREEN.
4. Save the file.
5. Place these two images in the `<installation-directory\MegaRAID Storage Manager>` folder.
6. Start the application.

Following are some of important points that you need to keep in mind:

- File names for both entries should not have any spaces. For example, the valid file name would be: `logo_test_1.png`, `LogoTest1.png`, or `TEST_SPLASH_FILE.jpg`.
- The logo image dimensions should not exceed 160 pixels x 85 pixels (width x height).
- The splash screen image dimensions should not exceed 390 pixels x 260 pixels (width x height).

After making the changes mentioned previously, when you log into the MegaRAID Storage Managers software, you will be able to view the changes with the new splash screen and logo in the MegaRAID Storage Manager software.

11.3.14 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 need to follow certain steps based on your operating system.

11.3.14.1 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. Go to the `<MSM_INSTALLATION_FOLDER>\MegaPopup` folder.
3. Run the command, `popup -stop`.

After running the preceding command, the pop-up process stops.

11.3.14.2 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.

11.3.15 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 popup process.

- For the Windows operating system, you must first stop the pop-up process (see [Section 11.3.14.1, 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 [Section 11.3.14.2, 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 [Section 11.3.14.2, 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.

11.4 Installing and Supporting 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.

11.4.1 Prerequisites for Installing MegaRAID Storage Manager for VMware

For the VMware 3.5 operating system, it is necessary to install `libstdc++34-3.4.0-1.i386.rpm` before installing the MegaRAID Storage Manager software. You can download the `rpm` file from:

<http://rpm.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.

11.4.2 Installing MegaRAID Storage Manager on VMware ESX (VMware Classic)

The VMware operating system does not support any graphics components. To install the MegaRAID Storage Manager software on the VMware operating system, run the script `./vmware_install.sh` from the installation disk.



NOTE Ensure 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 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.

11.4.3 Uninstalling MegaRAID Storage Manager for VMware

To uninstall the Server Component of the MegaRAID Storage Manager software on VMware, either use the `Uninstall` command in the Program menu, or run the script `/usr/local/MegaRAID Storage Manager/uninstaller.sh`.

You need to keep in mind the following points:

- A MegaRAID Storage Manager upgrade is supported in this release. Future releases can update this release.
- To shut down the MegaRAID Storage Manager Framework service, run the following command:

```
/etc/init.d/vivaldiframeworkd stop
```

The Linux RPM of the MegaRAID Storage Manager software works under the console with minimal changes. Hardware RAID is currently supported in ESX 4.x.



NOTE There is a known limitation that virtual drives that are created or deleted will not be reflected to the kernel. The workaround is to reboot the server or to run `esxcfg-rescan <vmhba#>` from COS shell.

11.4.4 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 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 IBM management software. Please make sure that the network settings are correct by making the following changes:

- Provide a proper host name and an IP address while doing 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
```

11.4.5 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 Windows/Linux/Solaris. Events are collected from the time a client logs in to an ESXi machine for the first time, and it continues to be collected as long as the Framework is running.

11.4.5.1 Differences in MegaRAID Storage Manager for the VMware ESXi System

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.
- 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 support for retrieving initial logs is limited to 30 events. Only those events that occur after a client logs in for the first time to an ESXi server appear in the Event Logger dialog.
 - The System logs are not displayed.
 - The "Save log" feature is not supported; however, the "Save Log as Text" is supported.
 - 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

11.5 Installing and Configuring a CIM Provider

This section describes the installation and configuration of the IBM MegaRAID Common Information Model (CIM) provider. The Common Information Model offers common definitions of management information for networks, applications, and services, and allows you to exchange management information across systems throughout a network.

On a VMware ESXi system, management is possible only through a CIM provider, and it is performed through the MegaRAID Storage Manager software installed on a remote machine running a Linux or Windows operating system.

The VMware ESXi system comes with the Small Footprint CIM Broker (SFCB) CIM Object Manager (or CIMOM). A CIMOM manages communication between providers, which interact with the hardware, and a CIM client, where the administrator manages the system.

SFCB supports Common Manageability Programming Interface (CMPI)-style providers. CMPI defines a common standard used to interface manageability instrumentation (providers, instrumentation) to management brokers (CIM Object Manager). CMPI standardizes manageability instrumentation, which allows you to write and build instrumentation once and run it in different CIM environments (on one platform).

11.5.1 Installing a CIM SAS Storage Provider on the Linux Operating System

The following procedure documents how to install and uninstall the CIM SAS Storage Provider on a system running on the Linux operating system.



NOTE Uninstall all the previous versions of LSI SAS Provider before you install this version. You can check all of the installed versions of LSI SAS Provider by running the `rpm -qa | grep LsiSASProvider` command.

- To install a CIM SAS Storage Provider on a Linux system, install the SAS Provider using the Red Hat Package Manager (RPM) by entering the following command:

```
rpm -ivh
```

The RPM installs all of the necessary files and the Managed Object Format (MOF), and it registers the libraries. The SAS Provider is now ready to use.



NOTE After you install 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 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.

11.5.2 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 SAS Pegasus provider, verify that the `libLsiSASProvider.so` file and the `libLsiSASProvider.so.1` file are in `/usr/lib/Pegasus/providers` directory.

If these files are not present, copy the `libLsiSASProvider.so.1` file from `/opt/tog-pegasus/providers/lib` to `/usr/lib/Pegasus/providers`, and create a symbolic link `libLsiSASProvider.so` to `/usr/lib/Pegasus/providers/libLsiSASProvider.so.1` at `/usr/bin/Pegasus/providers`.

2. Restart the Pegasus CIM Server and IBMServer by performing the following steps:
 - To start the tog-pegasus server, run the following command:

```
# /etc/init.d/tog-pegasus restart
```
 - To start IBMSAS Sever, run the following command:

```
# /etc/init.d/LsiSASd restart
```

11.5.3 Installing a CIM SAS Storage Provider on Windows

The following procedure describes how to install and uninstall the IBM 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 IBM 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.

11.6 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 IBM ServeRAID 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

11.6.1 Prerequisite for IBM SNMP Agent RPM Installation

The IBM SNMP agent application depends upon the standard SNMP Utils package. Make sure that the SNMP-Util package is present in the system before you install IBM 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 IBM SNMP agent RPM.

11.6.2 Prerequisite for Installing SNMP Agent on Linux Server

The SNMP application requires the standard library libsysfs. Make sure that this library is present in the system before installing the SNMP RPM.

The minimum library versions required for installing SNMP server are as follows.

- libsysfs version 2.0. This library is available in the rpm `<Lib_Utils-1.xx-xx.noarch.rpm>`. `<Lib_Utils-1.xx-xx.noarch.rpm>` is packaged in the SNMP zip file.
- libstdc++.so.6. This library is present in `/usr/lib` directory.
You can install the SNMP application from the Linux software component RPM that provides these libraries. These RPMs are available in the Linux OS DVD.

11.6.3 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 IBM SAS SNMP Agent using the `rpm -ivh <sas rpm>` command.



NOTE Before installation, check whether there is any pass command exists that starts with 1.3.6.1.4.1.3582 OID in `snmpd.conf`. If so, delete all of the old pass commands that start with 1.3.6.1.4.1.3582 OID. (This situation could occur if an earlier version of IBM 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 `lsi_mrdsnmpd.conf`. For reference, a sample configuration file (`lsi_mrdsnmpd.conf`) is in the `/etc/lsi_mrdsnmp` directory.

2. To run an SNMP query from a remote machine, add the IP address of that machine in the `snmpd.conf` file, as in this example:

```
com2sec    snmpclient    172.28.136.112    public
```

Here, the IP address of the remote machine is 172.28.136.112.

3. To receive an SNMP trap to a particular machine, add the IP address of that machine in the `com2sec` section of the `snmpd.conf` file.

For example, to get a trap in 10.0.0.144, add the following to `snmpd.conf`.

```
#          sec.name      source          community
com2sec    snmpclient    10.0.0.144     public
```

4. To 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/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 `disableAuthorization yes` in `snmptrapd.conf` and then run `sudo snmptrapd -P -F "%02.2h:%02.2j TRAP%w.%q from %A %v\n"`.



NOTE It is assumed that `snmpd.conf` is located in `/etc/snmp` for the Red Hat operating system and `/etc` for the SLES operating system. You can change the file location from the `/etc/init.d/lsi_mrdsnmpd` file.

You can install SNMP without the trap functionality. To do so, set the `TRAPIND` environment variable to "N" before running RPM.

Before you install a new version, you must uninstall all previous versions.

For the SLES 10 operating system, perform the following steps to run SNMP:

1. Copy `/etc/snmp/snmpd.conf` to `/etc/snmpd.conf`.
2. Modify the `/etc/init.d/snmpd` file, and change `SNMPDCONF=/etc/snmp/snmpd.conf` entry to `SNMPDCONF=/etc/snmpd.conf`.
3. Run `LSI SNMP rpm`.

11.6.4 Installing and Configuring an SNMP Agent on Solaris

This section explains how to install and configure SAS SNMP Agent for the Solaris operating system.

11.6.4.1 Prerequisites

This package requires that you have Solaris System Management Agent installed on the Solaris machine.



NOTE While installing the SAS SNMP Agent on Solaris 11, the `net-snmp` package needs to be installed on the machine.

11.6.4.2 Installing SNMP on Solaris

To install SNMP for the Solaris operating system, perform the following steps:

1. Unzip the IBM SAS SNMP Agent package.
2. Run the install script by using the following command:


```
# ./install.sh
```

The installation exits if any existing versions of storelib and sassnmp are installed on the Solaris machine. Uninstall the existing version by using the following commands:

```
# pkgrm sassnmp (to uninstall the IBM SAS SNMP Agent)
# pkgrm storelib (to uninstall storelib library)
```

11.6.4.3 IBM SAS SNMP MIB Location

After you install the IBM SAS SNMP Agent package, the MIB file `LSI-AdapterSAS.mib` is installed under `/etc/lsi_mrdsnmp/sas` directory.

11.6.4.4 Starting, Stopping, and Checking the Status of the IBM 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`

11.6.4.5 Configuring snmpd.conf

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 above excerpt.

- b. Add a readwrite community name and client, hostname, ipaddress under SECTION: Access Control Setup in `/etc/sma/snmp/snmpd.conf` file, as shown in the following excerpt.

```
#####
# SECTION: Access Control Setup
# This section defines who is allowed to talk to your
# running snmp agent.
# 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 above 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 need to start the net-snmpd daemon service, by executing the following command: `# svcadm enable svc:/application/management/net-snmp`

11.6.4.6 Configuring SNMP Traps

To receive SNMP traps, perform the following steps:

1. Stop the IBM SAS SNMP Agent by using the following command:


```
#/etc/init.d/lssi_mrdsnmpd stop
```
2. Edit the `/etc/lssi_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 in case, 'TrapDestInd' above is set to 1, IP addresses shall be taken from `/etc/sma/snmp/snmpd.conf` file in the following format: 'com2sec snmpclient 172.28.157.149 public' 'Trapsink' and 'TrapCommunity' tokens are supported for sending customised SNMP traps



NOTE In Solaris 11 x86, the file will be taken from `/etc/net-snmp/snmp/snmpd.conf`.

4. Start the IBM SAS SNMP Agent by entering the following command:
`#/etc/init.d/lsi_mrdsnmpd start`

11.6.4.7 Uninstalling the SNMP Package

The `uninstall.sh` script is located under the `/etc/lsi_mrdsnmp/sas` directory. Use the following command to uninstall the package:

```
# cd /etc/lsi_mrdsnmp/sas
# ./uninstall.sh
```

11.6.5 Installing a SNMP Agent on Windows

This section explains how to install and configure SAS SNMP Agent for the Windows operating system.

11.6.5.1 Installing SNMP Agent

Perform the following steps to install SNMP Agent:

1. Run `setup.exe` from DISK1.
2. Use SNMP Manager to retrieve the SAS data (it is assumed that you have compiled `LSI-AdapterSAS.mib` file already).
The `LSI-AdapterSAS.mib` file is available under the `%ProgramFiles%\LSI Corporation\SNMPAgent\SAS` directory.
3. Use a trap utility to get the traps.



NOTE Before you install the Agent, make sure that SNMP Service is already installed in the system.

11.6.5.2 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.

11.6.5.3 Configuring SNMP Service on the Server Side

Perform the following steps to configure SNMP Service on the server side.

1. Select **Administrative Tools** from the **Control Panel**.
2. Select **Services** in the **Administrative Tools** window.
3. Select **SNMP Service** in the **Services** window.
4. Open **SNMP Service**.
5. Click the **Security** tab, and make sure that **Accept SNMP Packets from any host** is selected.
6. Click the **Traps** tab, and select the list of host IP addresses to which you want the traps to be sent with the community name.

11.6.5.4 Installing SNMP Service for the Windows 2008 Operating System

Before you install the IBM Agent, make sure that SNMP Service is already installed in the system.

If you do not have SNMP Service installed on your system, perform the following steps to install SNMP Service for Windows 2008 system.

1. Select **Program and Features** from the **Control Panel**.
2. Click **Turn windows feature on/off** to select the windows components to install.
3. Select **Features** from the menu.
4. Click **Add Features**.
5. Select **SNMP Services**.
6. Click **Next**.
7. Click **Install**, and the SNMP installation starts. You will be prompted for the Windows 2008 CD during the installation.
8. Insert the CD, and click **Ok**.

The installation resumes.

After the installation is finished, the system displays a message saying that the installation is successful.

11.6.5.5 Configuring SNMP Service on the Server Side for the Windows 2008 Operating System

To configure SNMP service on the server side for Windows 2008 operating system, perform the following steps:

1. Select **Administrative Tools** from the **Control Panel**.
2. Select **Services** from **Administrative Tools** window.
3. Select **SNMP Service** from the **Services** window.
4. Open **SNMP Service**, and go to its properties.
5. Go to the **Security** tab, and make sure that **Accept SNMP Packets from any host** is selected.
6. Click the **Traps** tab, and select the list of host IP addresses to which you want the traps to be sent with the community name.

11.7 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 /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.

11.8 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.

Chapter 12: MegaRAID Storage Manager Screens

This chapter explains how to start the MegaRAID Storage Manager software and describes the MegaRAID Storage Manager window and menus.

12.1 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 JavaTechnology.)

- To start the MegaRAID Storage Manager software on a Red Hat Linux operating system, select **Applications > System Tools > MegaRAID Storage Manager StartupUI**.
- To start MegaRAID Storage Manager software on a SUSE Linux or SLES operating system, select **Start > System > More Programs > MegaRAID Storage Manager**.
- To start MegaRAID Storage Manager software on a Solaris X86 and Solaris SPARC operating system, select **Launch > Applications > Utilities > MegaRAID Storage Manager StartupUI**.

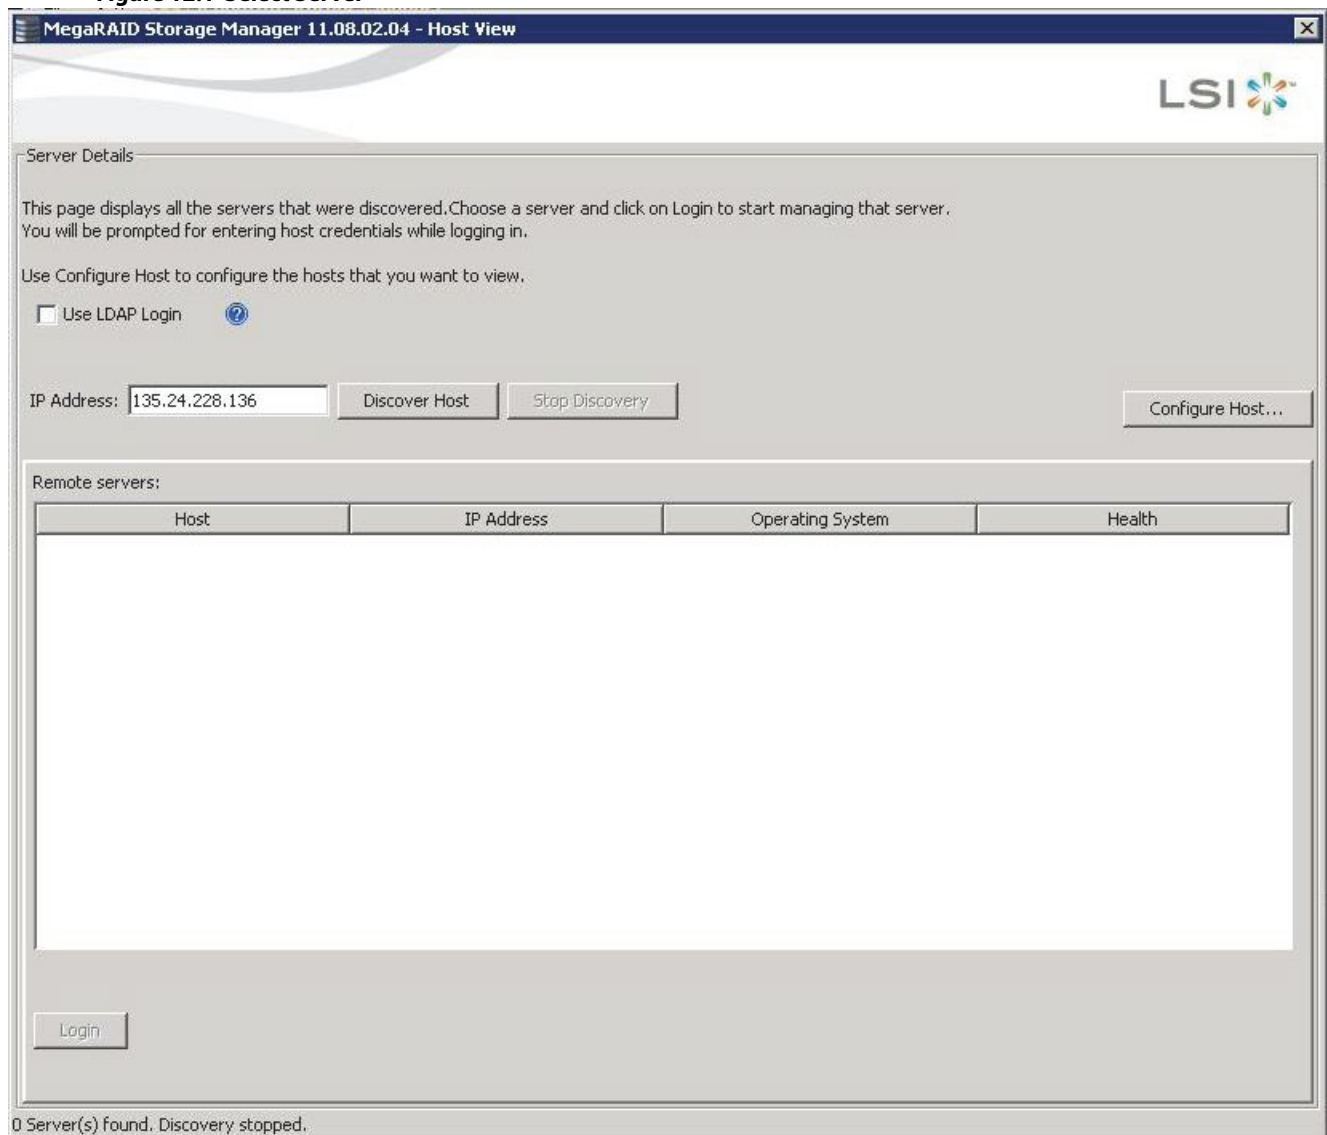
12.2 Discovery and Login

You can start the MSM software from a remote Windows/ Linux machine that has the MSM software installed in complete mode. When the program starts, the **Select Server** 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 MSM software on Windows, the following figure does not show. It directly prompts you to the login dialog as shown in the Server Login.

Figure 12.1 Select Server



The **Select Server** 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:

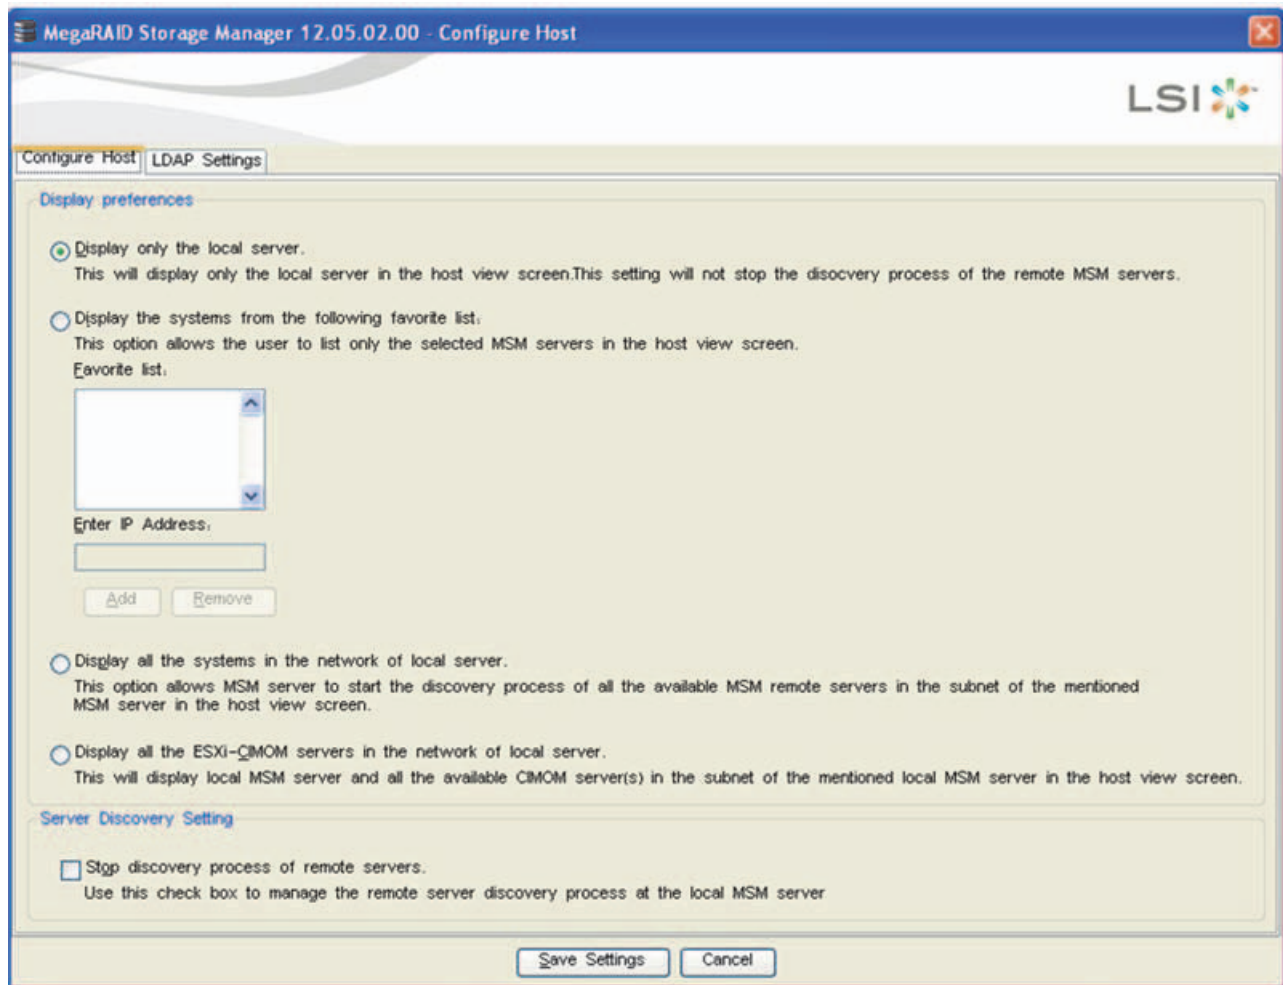
- Green: The server is operating properly.
- Yellow: The server is running in a partially degraded state (possibly because a drive in a virtual drive has failed).
- Orange: The server is running in a degraded state.
- Red: The server storage configuration has failed.



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.

Figure 12.2 Configure Host



Select any one of the following options from **Display preferences**.

- **Display only the local server** – Discovers only the local MegaRAID Storage Manager server.
- **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 On some Windows machines, the discovery of VMware ESXi servers fail as a result of a bug in the third-party application that is used for discovery. This is caused by one of the Windows servers in the network that contains a service called IBM SLP SA, which gets installed along with the IBM Director. If we stop this service on all the Windows servers in the network, the MegaRAID Storage Manager can discover all the ESXi servers.

2. Click **Save Settings**. A confirmation dialog appears asking you to confirm your settings.
 3. Click **OK** in the confirmation dialog to start the discovery process.
- After the discovery process is completed, the servers appear in the **Select Server** dialog.

To abort the discovery process which has already begun, select the **Stop discovery process of remote servers** check box and click **Save Settings**. This check box is enabled only when there is an active discovery process.



NOTE For the VMware ESXi, the server icon does not denote the health of the server. The icon is always green regardless of the health of the system. The VMware server does not show the system health and the operating system labels. It shows only the host name and the IP address of the server. 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.

4. Double-click the icon of the server that you want to access.
The **Server Login** window appears.

Figure 12.3 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, LSI\abc, where LSI 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.

12.3 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 12.4 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.

12.4 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 12.5 Configure Host LDAP

The screenshot shows the 'Configure Host' dialog box with the 'LDAP Settings' tab selected. The dialog has a title bar 'MegaRAID Storage Manager 11.08.02.04 - Configure Host' and the LSI logo in the top right. The 'LDAP Settings' tab is active, showing two checkboxes: 'Use LDAP login as default login mode' and 'Do not prompt for credentials when connecting to LDAP'. Below these are two sections: 'Server' and 'Connection'. The 'Server' section contains 'IP Address', 'Port', and 'Distinguished Name' fields. The 'Connection' section contains 'User Name' and 'Password' fields. At the bottom are 'Save Settings' and 'Cancel' buttons.

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.

12.5 MegaRAID Storage Manager Main Menu

This section describes the MegaRAID Storage Manager main menu window:

- [Section 12.5.1, Dashboard / Physical View/ Logical View](#)
- [Section 12.6.2, Properties and Graphical View Tabs](#)
- [Section 12.6.3, Event Log Panel](#)

12.5.1 Dashboard / Physical View/ Logical View

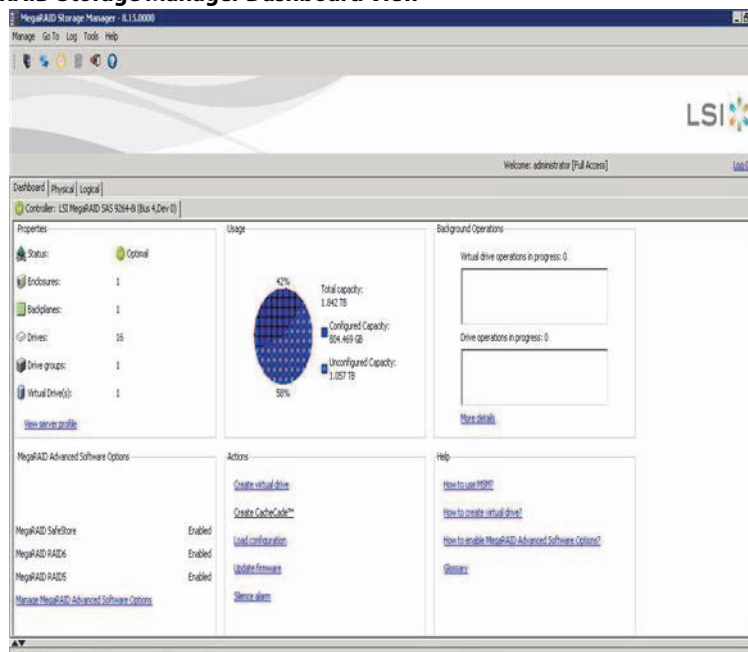
The left panel of the **MegaRAID Storage Manager** window displays the *Dashboard* view, the *Physical* view, or the *Logical* view of the system and the attached devices, depending on which tab is selected.

Dashboard View

The *Dashboard* view shows an overview of the system and covers the following features:

- Properties of the virtual drives and the physical drives
- Total capacity, configured capacity, and unconfigured capacity
- Background operations in progress
- 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

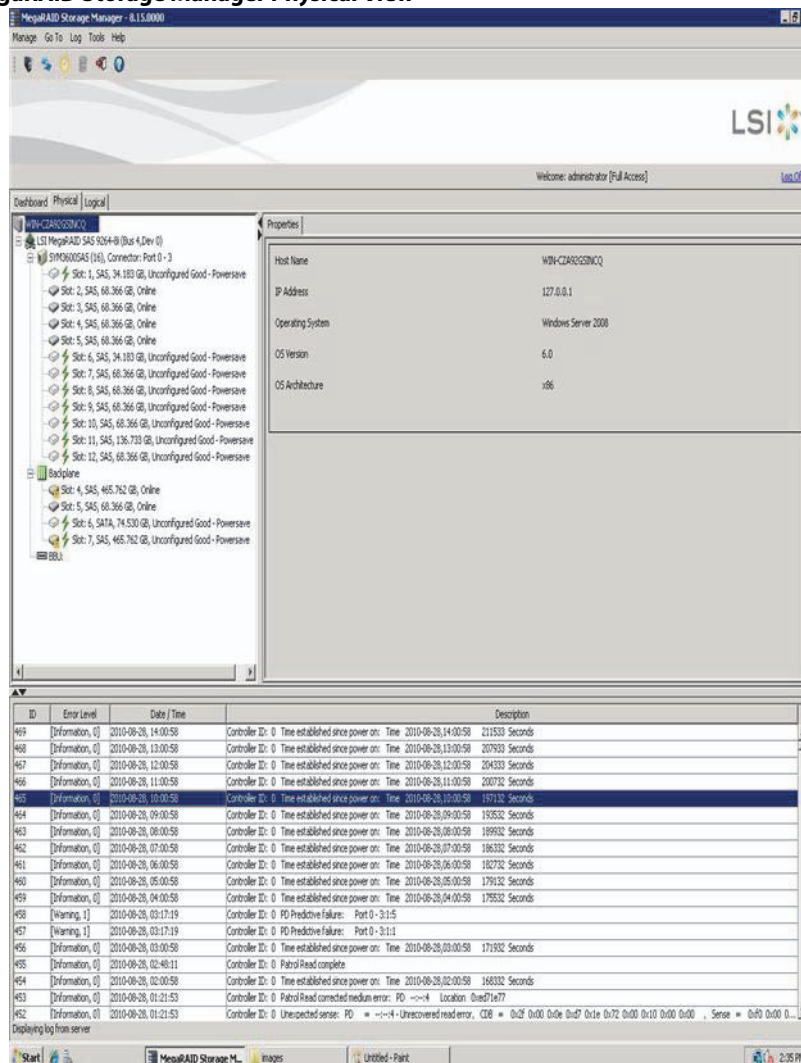
Figure 12.6 MegaRAID Storage Manager Dashboard View



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 ServeRAID controller so that you can easily differentiate between multiple controllers. Each controller has one or more ports. Drives and other devices are attached to the ports. The properties for each item appear in the right panel of the screen.

Figure 12.7 MegaRAID Storage Manager Physical View

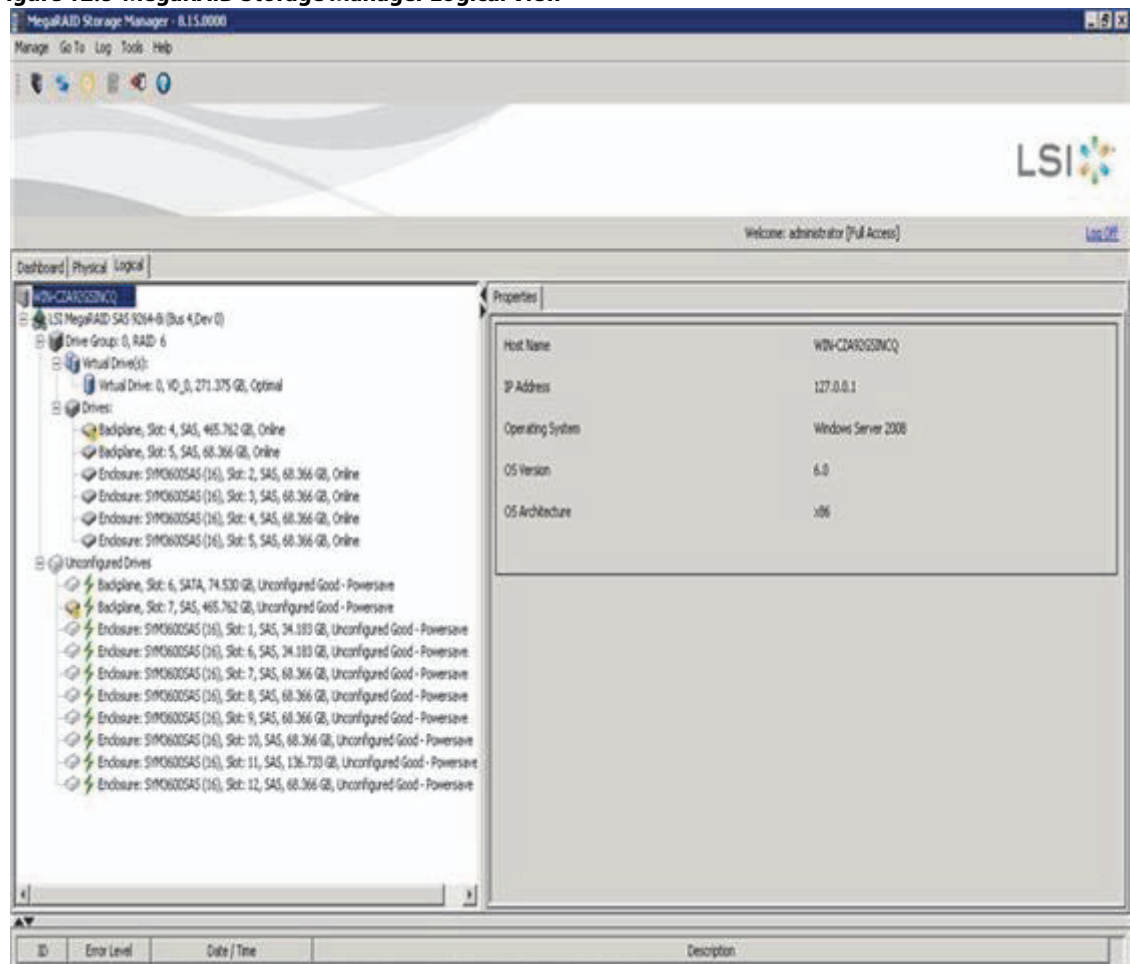


Logical View

The *Logical* view shows the hierarchy of controllers, virtual drives, and the drives and drive groups that make up the virtual drives. The properties for these components appear in the right panel.

The following figure shows the Logical view.

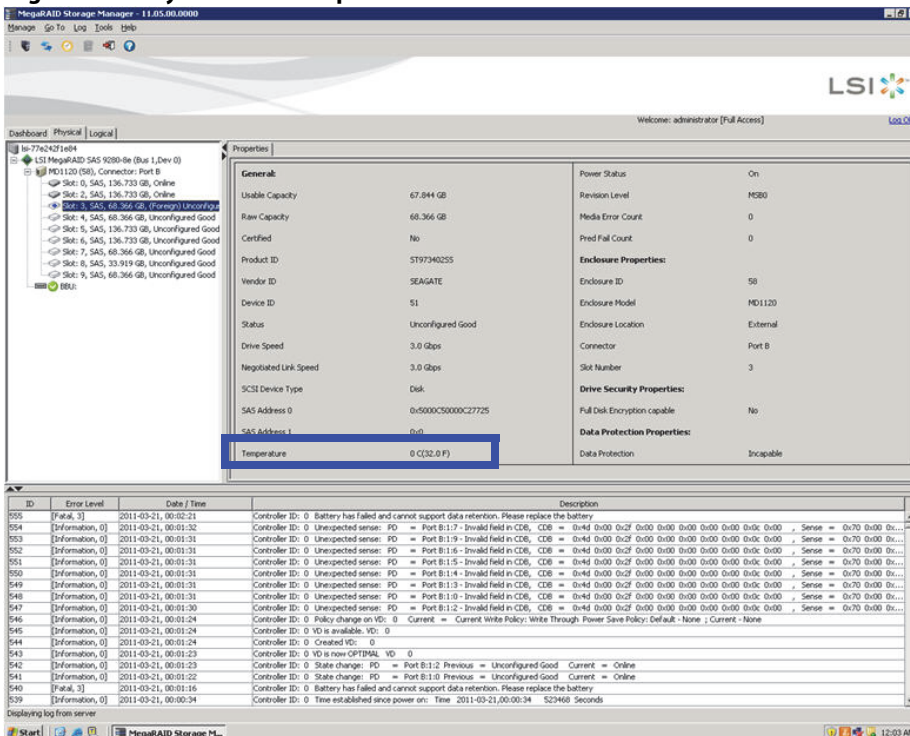
Figure 12.8 MegaRAID Storage Manager Logical View



12.5.2 Physical Drive Temperature

The temperature for the physical drive appears in the following figure. You can scroll down to view the **Temperature** property.

Figure 12.9 Physical Drive Temperature



12.5.3 Shield State

This section describes the Shield state in the MegaRAID Storage Manager software.

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

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

The three possible Shield states are **Unconfigured - Shielded**, **Configured - Shielded**, and **Hotspare - Shielded**.

12.5.4 Shield State Physical View

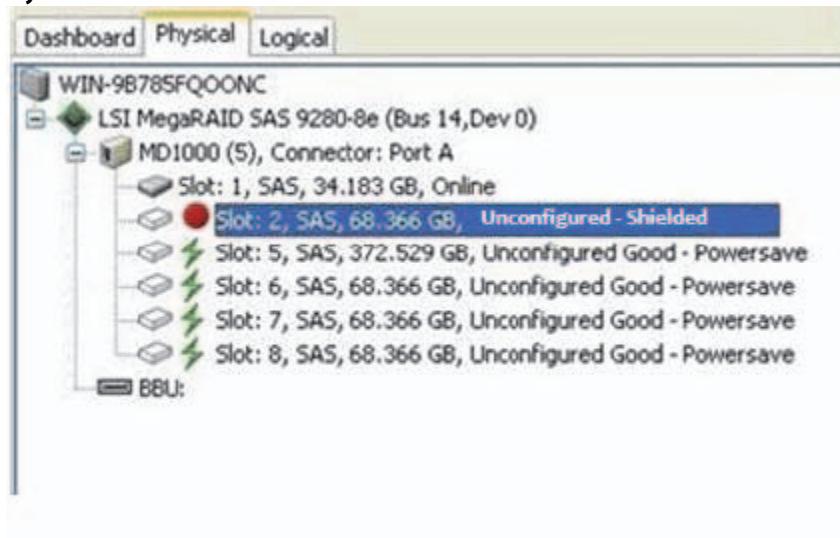
Follow these steps to view the Shield state under the **Physical** view tab.

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 12.10 Physical View Shield State

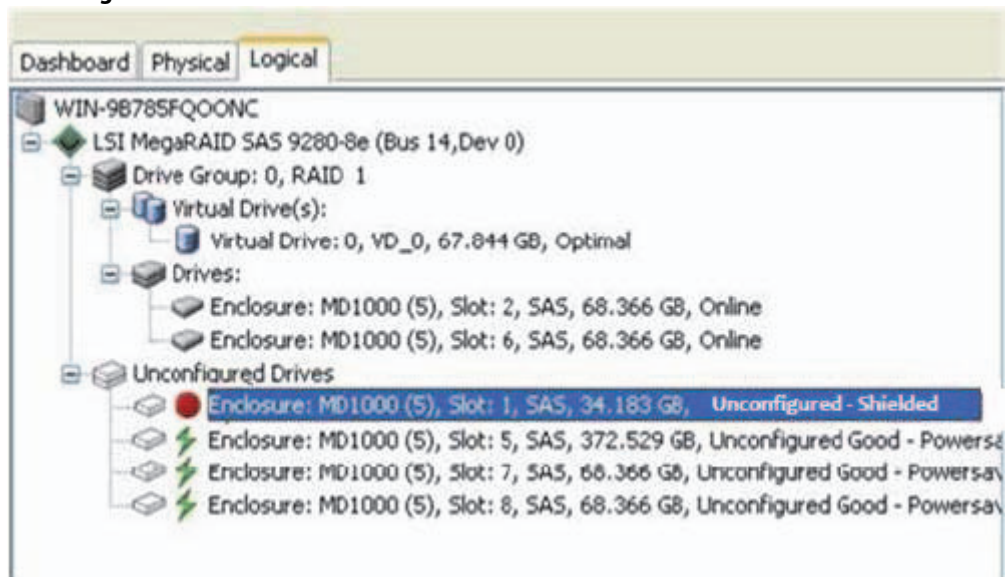


12.5.5 Logical View Shield State

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 12.11 Logical View Shield State



12.5.6 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.

- 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.



NOTE The Status of the drive must be of the Shield type.

Figure 12.12 Physical Drive Properties of a Drive in Shield State

Properties			
General:		SAS Address 0	0x4433221107000000
SSD Life Left	100 % - Optimal	Temperature	36 C(96.8 F) - Critical
Current Location of SSD		Commissioned Hotspare	No
Usable Capacity	90.656 GB	Emergency Spare	No
Raw Capacity	93.160 GB		
Certified	No	Revision Level	TI35
Product ID	TX43E10100GB0LSI	Media Error Count	0
Vendor ID	ATA	Pred Fail Count	0
Serial Number	5L0010ZE	Slot Number	4
Device ID	46	Drive Security Properties:	
Status	Online	Full Disk Encryption capable	No
Drive Speed	6.0 Gbps	Data Protection Properties:	
Negotiated Link Speed	6.0 Gbps	Data Protection	Incapable
SCSI Device Type	Disk	Shield Counter	0

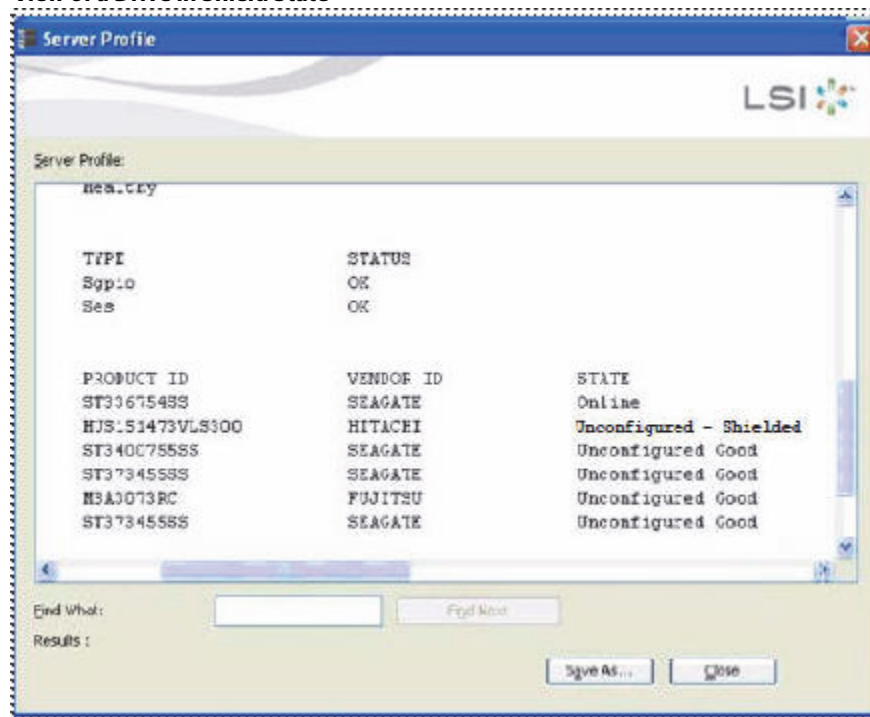
12.5.7 Viewing Server Profile of a Drive in Shield State

Perform these steps to view the server properties of the drive in Shield state.

- Click the **Dashboard** tab in the device tree.
- Click the **View Server Profile** link in the dashboard view.

The server profile information is displayed, as shown in the following figure.

Figure 12.13 View of a Drive in Shield State



12.5.8 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

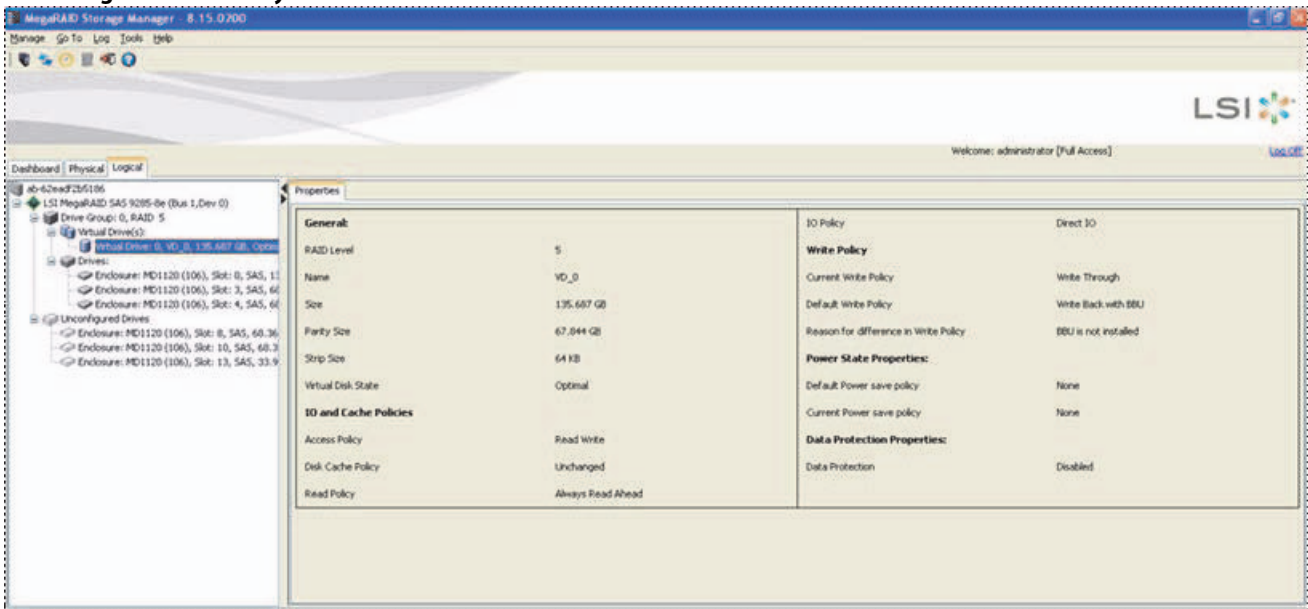
12.5.8.1 Parity Size

Parity size is used for storing parity information on RAID 5, RAID 6, RAID 50, and RAID 60 virtual drives.

Follow these steps to view the Parity Size.

1. In the Logical view, click the **Virtual Drive** node.
2. For RAID 5, RAID 6, RAID 50, and RAID 60, the **Parity Size** is displayed, as shown in the following figure.

Figure 12.14 Parity Size



12.5.8.2 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.


 **NOTE** The parity size and mirror data size are not displayed for RAID 0 and RAID 00 volumes.

Figure 12.15 Mirror Data Size




12.5.8.3 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.



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.

Figure 12.16 Metadata Size

Properties			
Alarm Present	Yes	Backend SAS Address 2	0x0
Alarm Enabled	No	Correctable Error Count	0
Cache Flush Interval	4 sec	Nonretry uncorrectable count	0
Coercion Mode	None	Cluster Enable	No
BBU Present	Yes	Cluster Active	No
BBU Present	Yes	SSD Guard	Enabled
BBU Size	32,000 KB	Drive Security Properties:	
BBU Version	3.18.00_4.09.05.00_0x020000	Drive security capable	No
Native Command Queuing	Enabled	Background Operation Properties:	
Flash Size	8,000 MB	Rebuild Rate	55
Memory Size	512,000 MB	Patrol Read Rate	38
Metadata Size	500 MB	Reconstruction Rate	30
Power State Properties:		B0L Rate	34
Power savings on unconfigured drives	Enabled	Consistency Check Rate	35
Power savings on hot spares	Enabled	MegaRAID Recovery Properties:	
Power Save Policy for Configured Drives	Auto	MegaRAID Recovery	Enabled
Drive Standby Time	43mins		
Firmware Properties:			

12.5.9 Emergency Spare

When a drive within a redundant virtual drive fails or is removed, the ServeRAID 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.

12.5.9.1 Emergency Spare for Physical Drives

The Emergency Spare property determines whether a particular drive can be 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 12.17 Emergency Spare- Physical Drive Properties

Properties			
General:			
Usable Capacity	278.875 GB	Revision Level	FT00
Raw Capacity	279.397 GB	Media Error Count	0
Certified	No	Pred Fail Count	0
Product ID	ST9300503SS	Enclosure Properties:	
Vendor ID	SEAGATE	Enclosure ID	252
Serial Number	3SE0F0PJ	Enclosure Model	Backplane
Device ID	30	Enclosure Location	External
Status	Online	Connector	Port A
Drive Speed	6.0 Gbps	Slot Number	5
Negotiated Link Speed	6.0 Gbps	Drive Security Properties:	
SCSI Device Type	Disk	Full Disk Encryption capable	Yes
SAS Address 0	0x5000C500126F77ED	Data Protection Properties:	
SAS Address 1	0x0	Data Protection	Incapable
Temperature	38 C(100.4 F)	Shield Counter	0
Commissioned Hotspare	Yes	Diagnostics Complete Date	0-0-0
Emergency Spare	Yes		

12.5.9.2 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 12.18 Emergency Spare Properties for Controllers

Properties			
Cache Flush Interval	4 sec	Drive Security Properties:	
Coercion Mode	None	Drive security capable	No
BBU Present	Yes	Background Operation Properties:	
NVRAM Present	Yes	Rebuild Rate	60
NVRAM Size	32.000 KB	Patrol Read Rate	30
BIOS Version	5.32.00_4.12.05.00_0x05150000	Reconstruction Rate	30
Native Command Queuing	Enabled	BGI Rate	30
Flash Size	16.000 MB	Consistency Check Rate	30
Memory Size	1.000 GB	MegaRAID Recovery Properties:	
Chip Temperature	65535 C(117995.0 F)	MegaRAID Recovery	Enabled
Shield State Supported	Yes	CacheCade™ Properties:	
Power State Properties:		CacheCade™ - SSD Caching	Enabled
Power savings on unconfigured drives	Enabled	Write Cache Capable	No
Power savings on hot spares	Enabled	Total Cache Size	0 Bytes
Drive Standby Time	30mins	Maximum Cache Size	512.000 GB
Firmware Properties:		Emergency Spare Properties:	
Firmware Package Version		Emergency Spare	Unconfigured Good & Global Hotspare
Firmware Version	3.152.35-1593	Emergency for SMARTer	Enabled
Firmware Build Time	Apr 04 2012 21:38:45		

12.5.9.3 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 12.19 Commissioned Hotspare

Properties			
General:			
Usable Capacity	33.656 GB	Revision Level	A130
Raw Capacity	34.183 GB	Media Error Count	0
Certified	No	Pred Fail Count	0
Product ID	HUS151436VLS300	Enclosure Properties:	
Vendor ID	HITACHI	Enclosure ID	252
Serial Number	J3VPJL6K	Enclosure Model	Backplane
Device ID	45	Enclosure Location	External
Status	Online	Connector	Port A
Drive Speed	3.0 Gbps	Slot Number	6
Negotiated Link Speed	3.0 Gbps	Drive Security Properties:	
SCSI Device Type	Disk	Full Disk Encryption capable	No
SAS Address 0	0x5000CCA00349CA0F	Data Protection Properties:	
SAS Address 1	0x0	Data Protection	Incapable
Temperature	27 C(80.6 F)	Shield Counter	0
Commissioned Hotspare	No	Diagnostics Complete Date	0-0-0
Emergency Spare	No		

12.5.10 SSD Disk Cache Policy

The ServeRAID firmware provides support to change the write-cache policy for SSD media of individual physical drives.

The ServeRAID 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 12.20 Controller Properties – SSD Disk Cache Policy

Properties			
Host Port Count	0	Backend SAS Address 6	0x0
FLU		Backend SAS Address 7	0x0
Alarm Present	Yes	Correctable Error Count	0
Alarm Enabled	Yes	Memory uncorrectable count	0
Cache Flush Interval	4 sec	Cluster Enable	No
Coercion Mode	None	Cluster Active	No
Batter Present	No	SSD Guard	Enabled
NRAM Present	Yes	SSD Disk Cache Setting	Disabled
NRAM Size	32,000 KB	Drive Security Properties:	
BIOS version	3.18.00_4.09.05.00_0cdH16A000	Drive security enabled	No
Native Command Queuing	Enabled	Drive security method	PGE Only
Flash Size	8,000 MB	Drive security capable	Yes
Memory Size	256,000 MB	EDM Supported	Yes
Power State Properties:		Key Management Mode	N/A
Power savings on unconfigured drives	Enabled	Background Operation Properties:	
Power savings on hot spares	Enabled	Rebuild Rate	30
Drive Standby Time	30mins	Patrol Read Rate	30
Firmware Properties:		Reconstruction Rate	30
Firmware Package Version	11-10-9-0015	Bit Rate	30
		Consistency Check Rate	30

12.5.10.1 Virtual Drive Settings

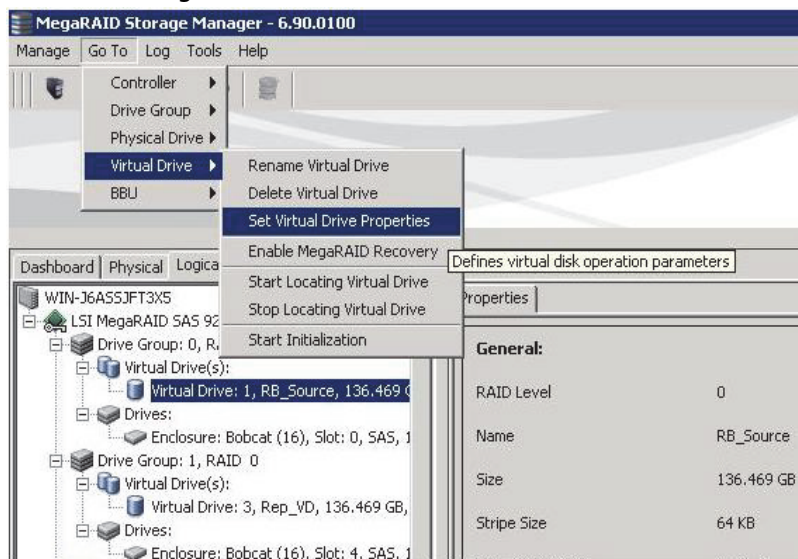
If the SSD cache property is enabled in the controller properties screen as shown in [Figure 12.21](#), then you cannot select the disk cache policy for the virtual drives having only SSD drives or a mix of SSD drives and HDD drives during virtual drive creation. The value of the disk cache policy is 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 12.21 Virtual Drive Settings



The value of the disk cache policy is unchanged, and the drop-down list is disabled.

12.5.10.2 Set Virtual Drive Properties

Follow these steps to set virtual drive properties.

1. Right-click on virtual drive node in the logical view of the device tree.
2. Select **Set Virtual Drive Properties**.

The **Set Virtual Drive Properties** dialog appears, as shown in the following figure.



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.

Figure 12.22 Virtual Drive Properties

Set Virtual Drive Properties

LSI

Description : Defines virtual disk operation parameters

Name:

Read Policy:

Write Policy:

IO Policy:

Access Policy:

Disk Cache Policy:

Background Initialization:

12.5.11 Non-SED Secure Erase Support

This section describes the firmware changes required to securely erase data on non-SEDs (normal HDDs).

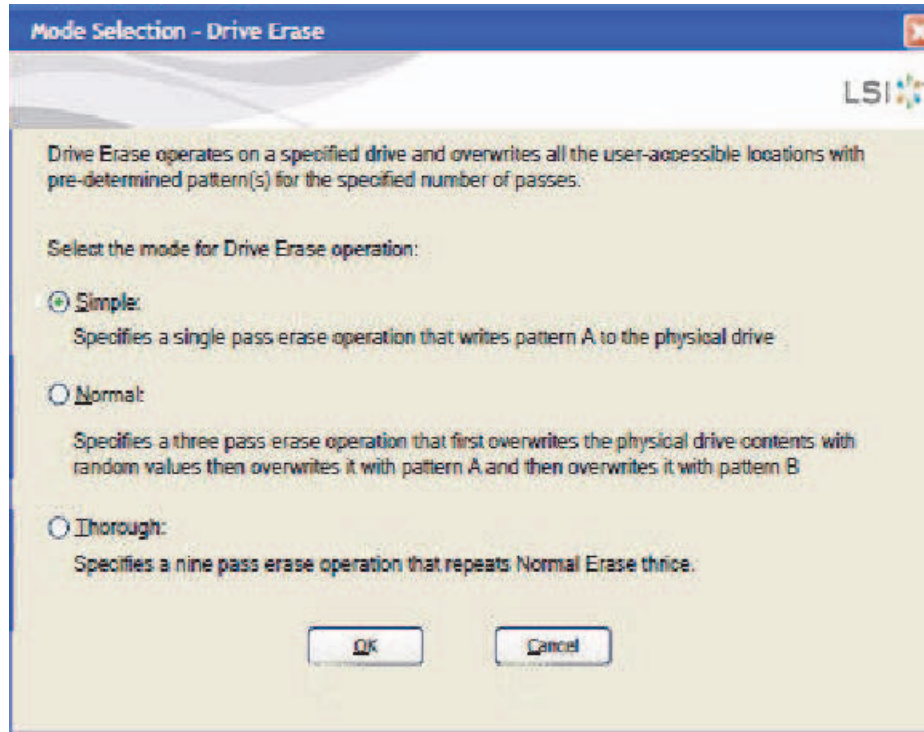
SEDs securely erase their internal encryption keys, effectively destroying all of the data present on the drive. For Non-SED drives, the erase operation consists of a series of write operations to a drive that overwrites every user-accessible sector of the drive with specified patterns. It can be repeated 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).

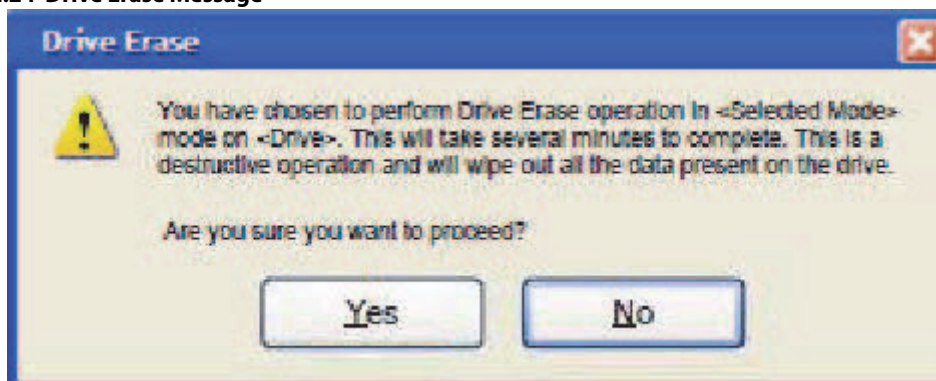
The **Mode Selection - Drive Erase** dialog appears.

Figure 12.23 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 12.24 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.

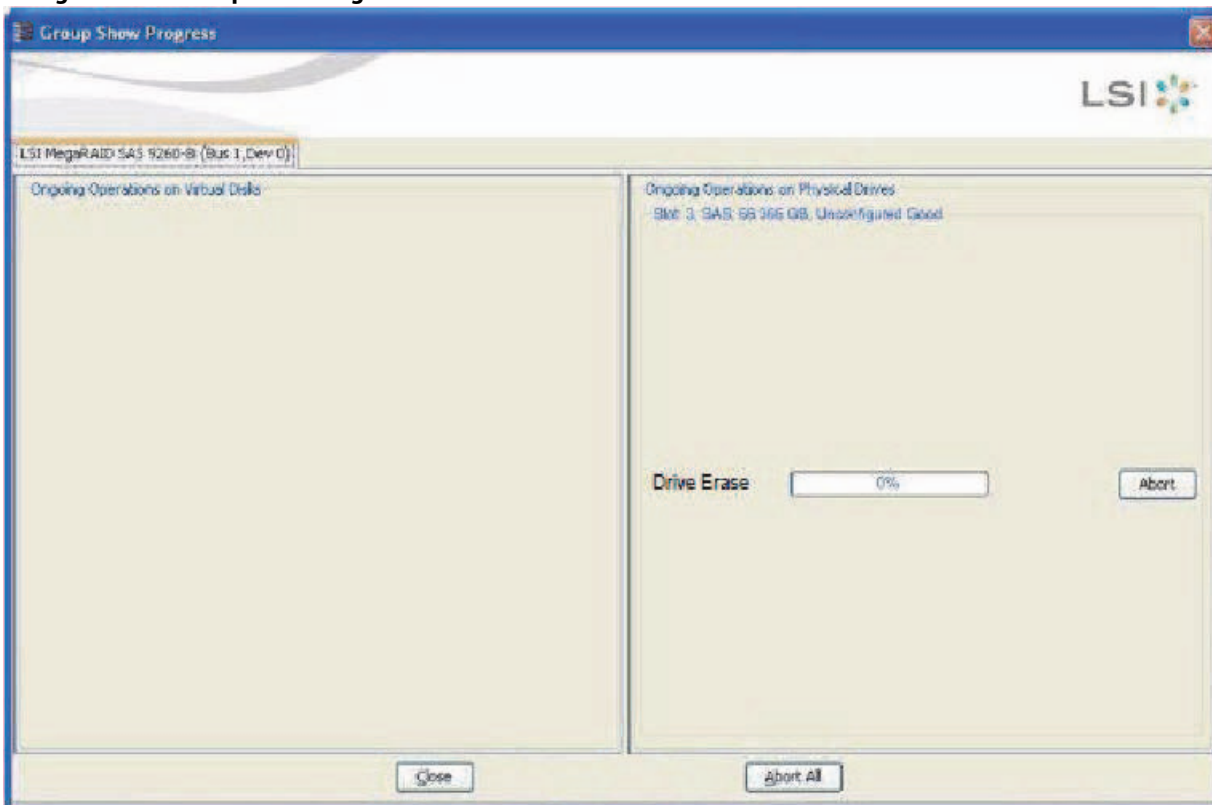
12.5.11.1 Group Show Progress for Drive Erase

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 12.25 Group Show Progress



When you click the **Abort All** button, all Drive Erase operations stop, and the progress bar is not displayed.

12.5.11.2 Virtual Drive Erase

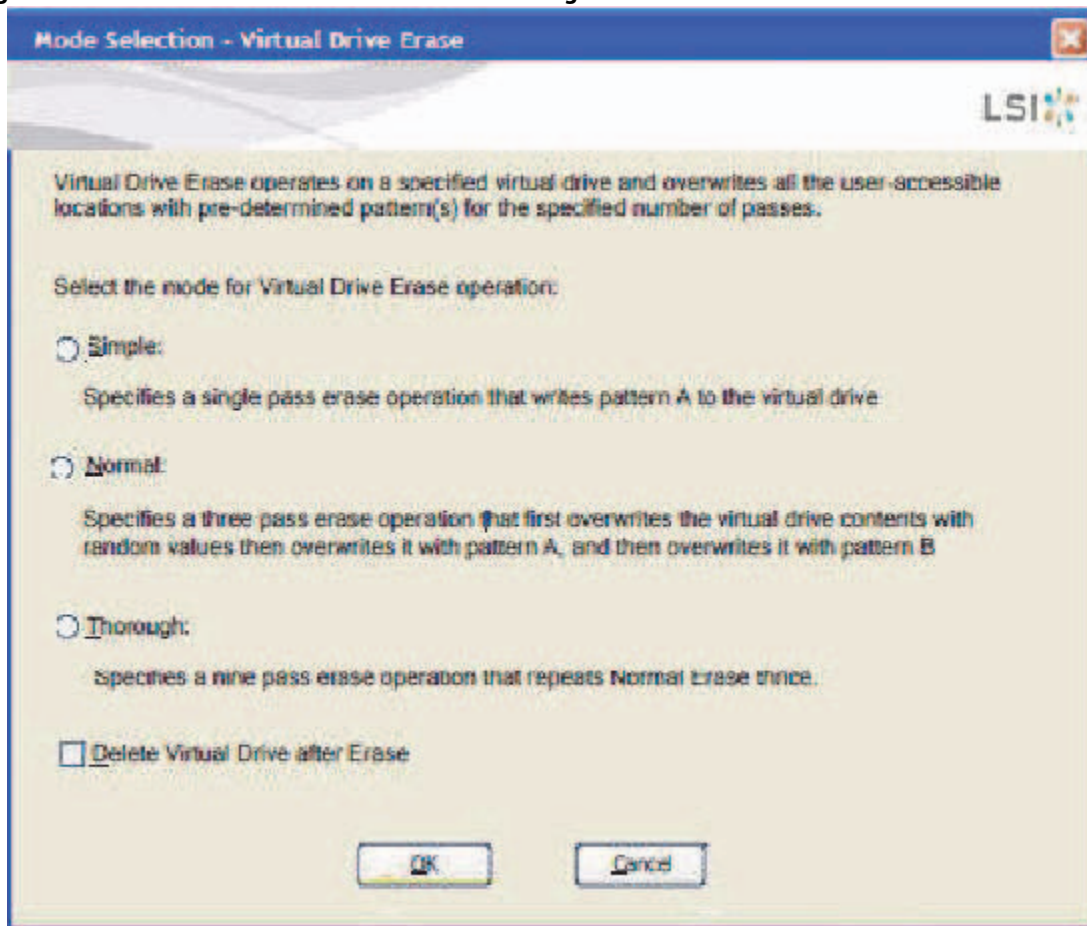
Virtual drive erase operates on a specified virtual drive and overwrites all user-accessible locations. It supports non-zero patterns and multiple passes. Virtual drive erase optionally deletes the virtual drive and erases the data within the virtual drive's LBA range. Virtual drive erase is a background 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.

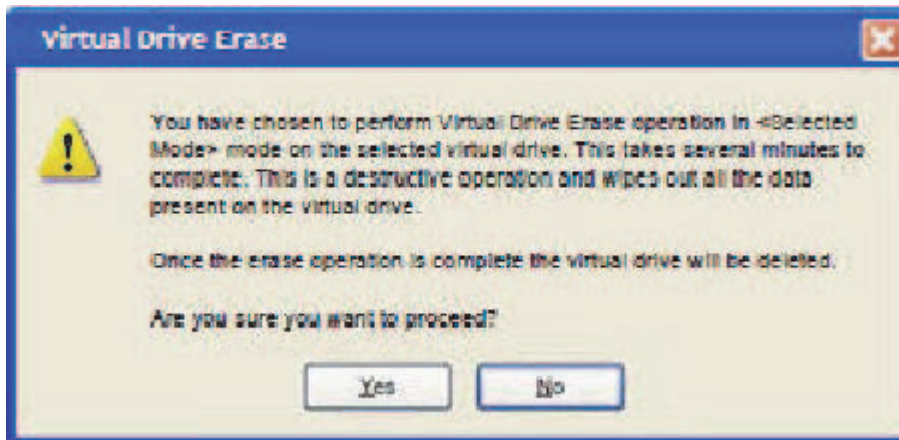
Figure 12.26 Mode Selection – Virtual Drive Erase Dialog



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, Figure 12.27 is displayed; otherwise, Figure 12.28 is displayed.
- **Normal** – (Alt + N) – After you select this option and click **OK**, and if **Delete Virtual Drive after Erase** is selected, Figure 12.27 is displayed; otherwise, Figure 12.28 is displayed.
- **Thorough** – (Alt + T) – After you select this option and click **OK** and if **Delete Virtual Drive after Erase** is selected, Figure 12.27 is displayed; otherwise, Figure 12.28 is displayed.
- **Delete Virtual Drive after Erase** – (Alt + D) – When you select this option, the virtual drive is erased and Figure 12.27 is displayed; otherwise, Figure 12.28 is displayed.
- **OK** – (Alt + O) – Click **OK** and if **Delete Virtual Drive after Erase** is checked, Figure 12.27 is displayed; otherwise, Figure 12.28 is displayed.
- **Cancel** – (Alt + C) – When you select this option, the dialog closes, and the MegaRAID Storage Manager navigates back to Physical view.

Figure 12.27 Warning Message for Virtual Drive Erase



- Click **Yes** to erase the virtual drive.
- Click **No** to cancel the erase and close the dialogue.

Figure 12.28 Warning Message for Virtual Drive Erase without Virtual Drive Delete



- Click **Yes** to erase the virtual drive.
- Click **No** to cancel the erase and close the dialogue.

12.5.11.3 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 12.29 Group Show Progress – Virtual Drive



12.5.12 Rebuild Write Cache

ServeRAID firmware supports drive cache properties during a rebuild operation. The ServeRAID 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 ServeRAID software automatically changes the setting for a drive that is being rebuilt. If the PD_CACHE for the rebuilt drive is already set, the firmware does not need to do anything extra.

The firmware identifies and sets the cache policy of the drives whenever a rebuild operation starts and the cache policy is reflected in the event logs. The firmware also makes sure to flush the cache just before committing the drive to the disk group.

12.5.13 Background Suspend or Resume Support

ServeRAID 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.

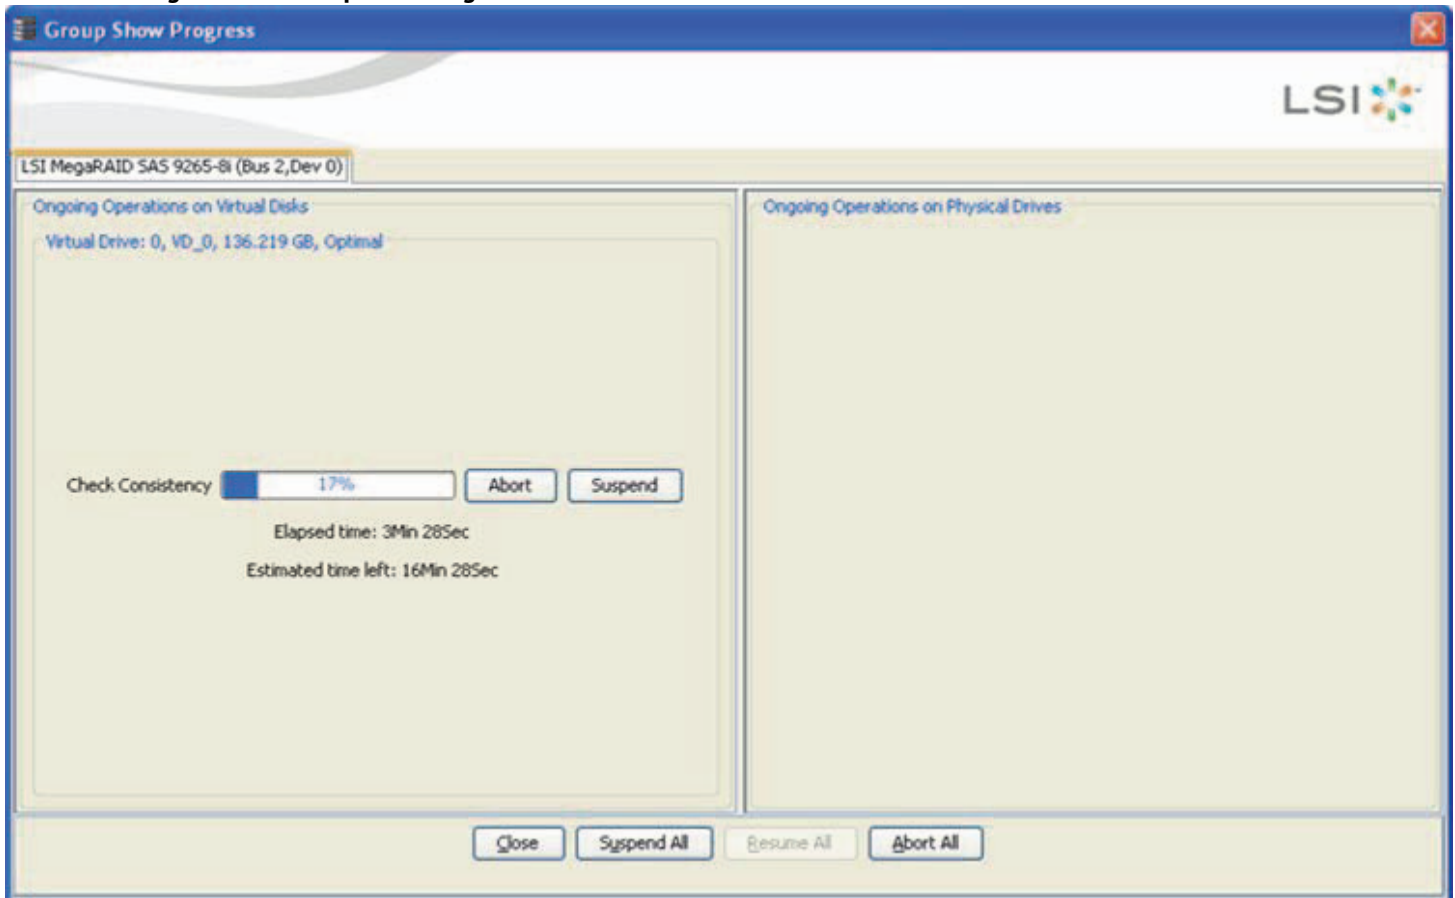
The background operations, including consistency-check, rebuild, replace, and background initialization are supported by an abort operation. If any operation is stopped before completion, it is considered to be aborted. An aborted operation cannot be resumed from the place where it was stopped.

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 12.30 Group Show Progress

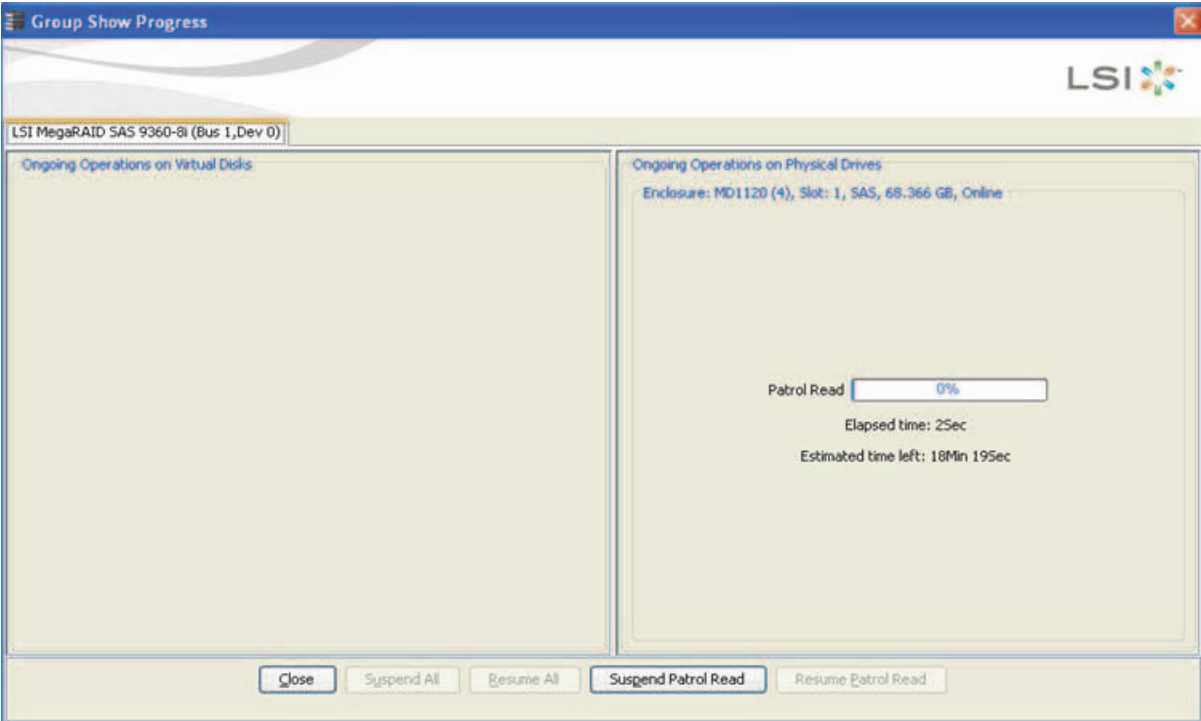


- **Suspend** (Alt + S) – Click the **Suspend** button to suspend the background operation taking place at that particular point of time. When the operations gets suspended, the **Resume** button appears instead of the **Suspend** button.
- **Resume** (Alt + E) – Click the **Resume** button to resume 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 12.31 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.

12.5.14 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 12.32 Enclosure Properties




















Vendor ID	DELL	FRU Number	41R3133
Enclosure ID	5	Part Number	CP-111-006-020
Enclosure Type	SES	Component Properties	
Enclosure Model	MD1120	Number of Temperature Sensors	4
Enclosure Location	External	Number of Fans	4
Firmware Version	A.04	Number of Power Supplies	2
Serial Number	0802V16NTE	Number of Voltage Sensors	0
Connector	Port A		
Number of Slots	15		

12.6 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.


12.6.1 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.

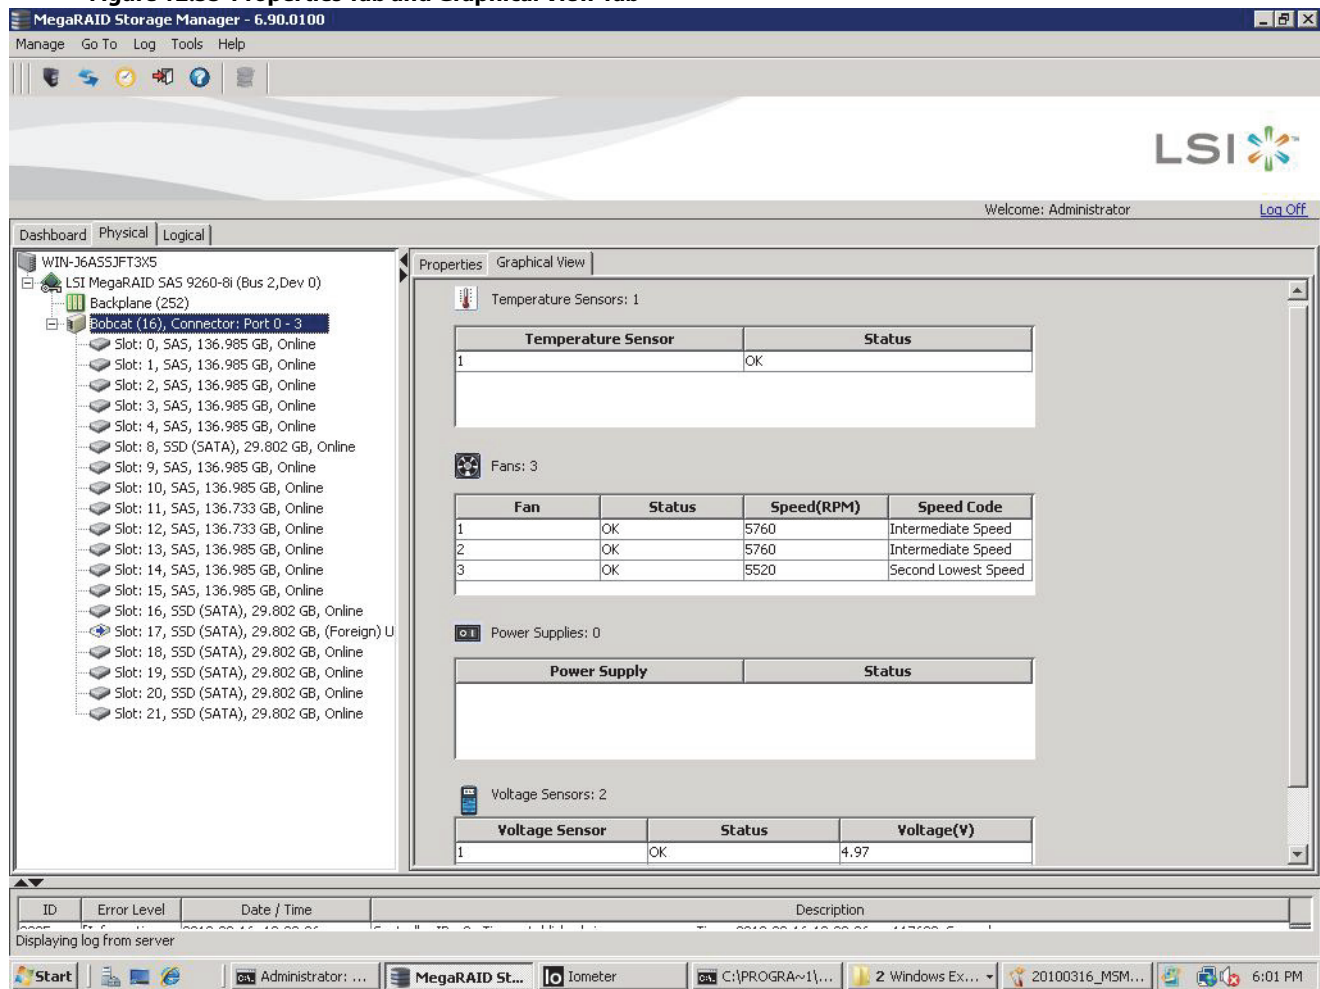
An orange circle to the right of an icon indicates that a device is running in a degraded state.

12.6.2 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. For more information, see [Section 14.15, Monitoring Controllers](#), [Section 14.16, Monitoring Drives](#), and [Section 14.18, Monitoring Virtual Drives](#).
- The **Graphical View** tab displays information about the temperature, fans, power supplies, and voltage sensors. To display a graphical view of a drive, click an enclosure icon in the left panel of the **MegaRAID Storage Manager** window, and click the **Graphical View** tab.

Figure 12.33 Properties Tab and Graphical View Tab



12.6.3 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 Appendix [Section Appendix A; Events and Messages](#).

12.6.4 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 [Section Appendix A; Events and Messages](#).

Tools Menu

On the Tools menu, you can select **Tools > Configure Alerts** to access the **Configure Alerts** dialog, where you can set the alert delivery rules, event severity levels, exceptions, and e-mail settings. For more information, see [Section 14.2, Configuring Alert Notifications](#).

Help Menu

On the Help menu, you can select **Help > Contents** to view the MegaRAID Storage Manager online help file. You can select **Help > About MegaRAID Storage Manager** to view version information for the MegaRAID Storage Manager software.



NOTE When you use the MegaRAID Storage Manager online help, you might see a warning message that Internet Explorer has restricted the file from showing active content. If this warning appears, click on the active content warning bar, and enable the active content.



NOTE If you are using the Linux operating system, you must install Firefox® browser or Mozilla® browser for the MegaRAID Storage Manager online help to display.



NOTE When connected to the VMware server, only the IP address and the host name information appear. The other information, such as the operating system name, version, and architecture do not appear.

Chapter 13: Configurations

This chapter explains how to use MegaRAID Storage Manager software to create and modify storage configurations on IBM.

The IBM 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 [Section Chapter 2: Introduction to RAID](#).



NOTE You cannot create or modify a storage configuration unless you are logged on to a server with administrator privileges.

13.1 Creating a New Configuration

You can use the MegaRAID Storage Manager software to create new storage configurations on systems with IBM 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.

13.1.1 Selecting Virtual Drive Settings

This section describes the virtual drive settings that you can select when you use the advanced configuration procedure to create virtual drives. You should change these parameters only if you have a specific reason for doing so. It is usually best to leave them at their default settings.

- **Initialization state:** Initialization prepares the storage medium for use. Specify the initialization status:
 - **No Initialization:** (the default) The new configuration is not initialized, and the existing data on the drives is not overwritten.
 - **Fast Initialization:** The firmware quickly writes 0s to the first and last 8-MB regions of the new virtual drive and then completes the initialization in the background. This allows you to start writing data to the virtual drive immediately.
 - **Full Initialization:** A complete initialization is done on the new configuration. You cannot write data to the new virtual drive until the initialization is complete. This process can take a long time if the drives are large.



NOTE BGI is supported only for RAID 5 and RAID 6 and not for any other RAID levels. New RAID 5 virtual drives require at least five drives for a background initialization to start. New RAID 6 virtual drives require at least seven drives for a background initialization to start. If there are fewer drives, the background initialization does not start.

- **Strip size:** Strip sizes of 8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB, and 1024 KB are supported. The default is 64 KB. For more information, see the *striping* entry in the **Glossary**.
- **Read policy:** Specify the read policy for this virtual drive:
 - **Always read ahead:** Read ahead capability allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This process speeds up reads for sequential data, but little improvement occurs when accessing random data.

- **No read ahead:** (the default) Disables the read ahead capability.
- **Write policy:** Specify the write policy for this virtual drive:
 - **Write Through:** In this mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This option eliminates the risk of losing cached data in case of a power failure.
 - **Always Write Back:** In this mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction.
 - **Write Back with BBU:** (the default) In this mode, the controller enables write back caching when the battery backup unit (BBU) is installed and charged. This option provides a good balance between data protection and performance.



NOTE The write policy depends on the status of the BBU. If the BBU is not present, is low, is failed, or is being charged, the current write policy switches to write through, which provides better data protection.

- **I/O policy:** The I/O policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.
 - **Cached IO:** In this mode, all reads are buffered in cache memory.
 - **Direct IO:** (the default) In this mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory.

Cached IO provides faster processing, and **Direct IO** ensures that the cache and the host contain the same data.
- **Access policy:** Select the type of data access that is allowed for this virtual drive.
 - **Read/Write:** (the default) Allow read/write access. This setting is the default value.
 - **Read Only:** Allow read-only access.
 - **Blocked:** Do not allow access.
- **Disk cache policy:** Select a cache setting for this drive:
 - **Enabled:** Enable the disk cache.
 - **Disabled:** Disable the disk cache.
 - **Unchanged:** (the default) Leave the current disk cache policy unchanged.

13.1.2 Optimum Controller Settings for CacheCade

Write Policy: Write Back/Write Through/Always Write Back

13.1.3 Optimum Controller Settings for Fast Path

Write Policy: Write Through

IO Policy: Direct IO

Read Policy: No Read Ahead

Stripe Size: 64 KB

13.1.4 Creating a Virtual Drive Using Simple Configuration

Simple configuration is the quickest and easiest way to create a new storage configuration. When you select the simple configuration mode, the system creates the best configuration possible using the available drives.

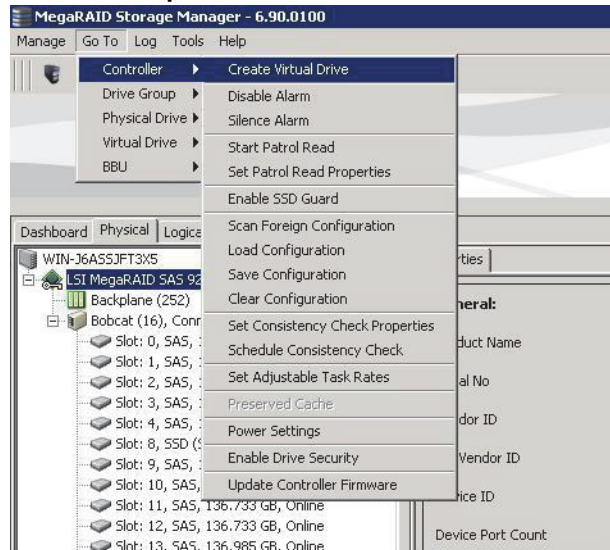


NOTE You cannot create spanned drives using the simple configuration procedure. To create spanned drives, use the advanced configuration procedure described in Section, [Section 13.1.5, Creating a Virtual Drive Using Advanced Configuration](#).

Follow these steps to create a new storage configuration in simple configuration mode.

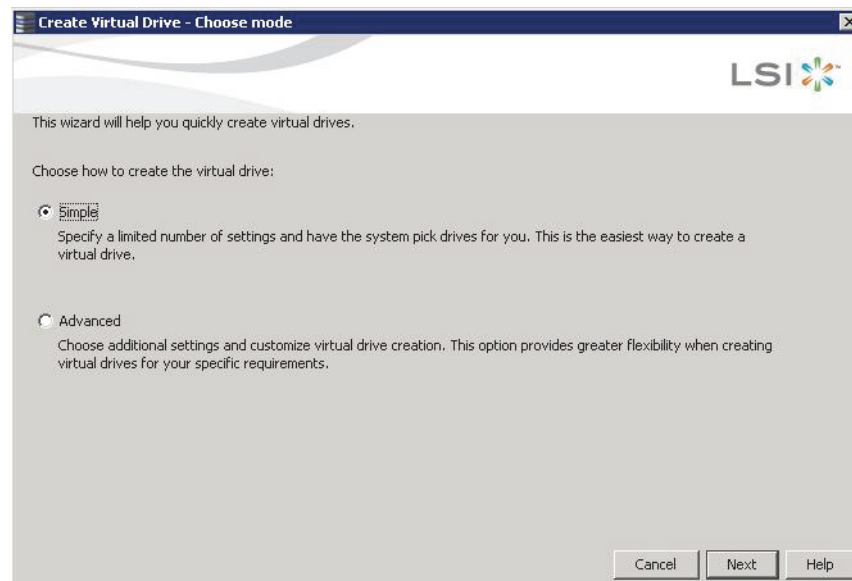
1. Perform either of the following steps:
 - Right-click the controller node in the device tree in the left frame of the **MegaRAID Storage Manager** window, and select **Create Virtual Drive**.
 - Select the controller node, and select **Go To > Controller > Create Virtual Drive** in the menu bar, as shown in the following figure.

Figure 13.1 Create Virtual Drive Menu Option



The dialog for the configuration mode (simple or advanced) appears, as shown in the following figure.

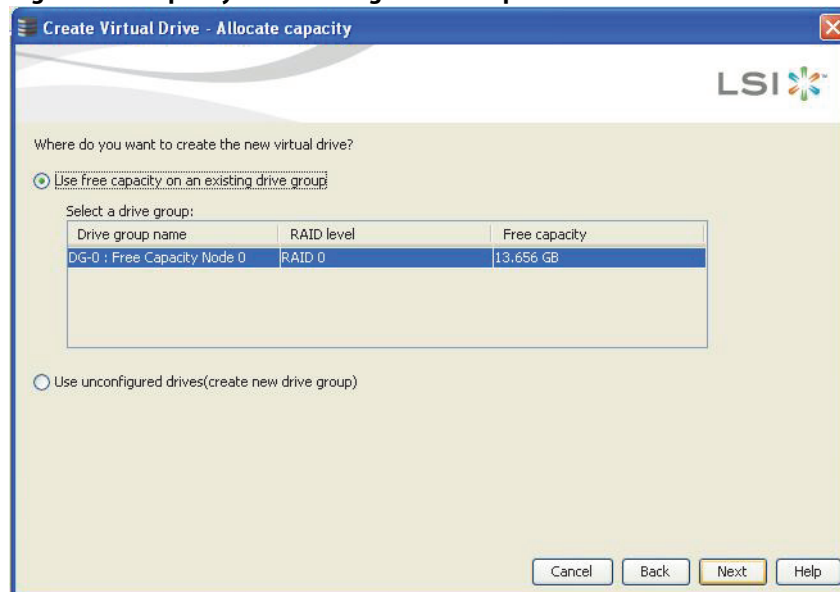
Figure 13.2 Create Virtual Drive - Choose mode



2. Select the **Simple** radio button, and click **Next**.

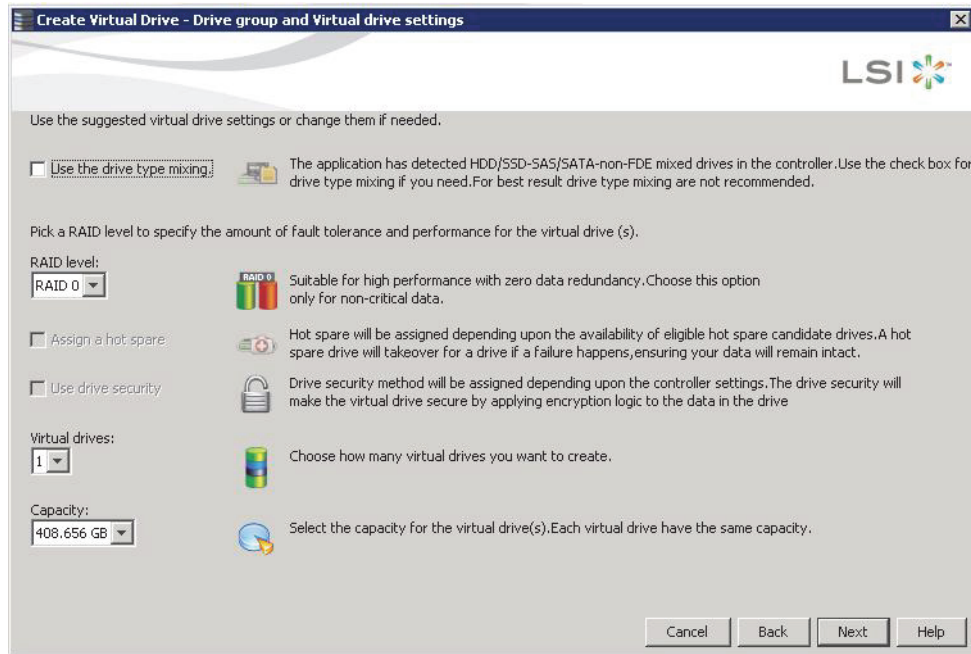
The **Create Virtual Drive - Allocate capacity** dialog appears, as shown in the following figure. If unconfigured drives are available, you have the option to use those unconfigured drives. If unconfigured drives are available, the **Create Drive Group Settings** window appears, and you can go to step 4.

Figure 13.3 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, SSD, SAS, and SATA, an option appears to allow drive type mixing.
 - If unconfigured drives are available, select the radio button to use the unconfigured drives, and click **Next**. Continue with step 10. The Summary window appears as shown in [Figure 13.5](#).

Figure 13.4 Create Virtual Drive - Drive group and Virtual drive settings Dialog



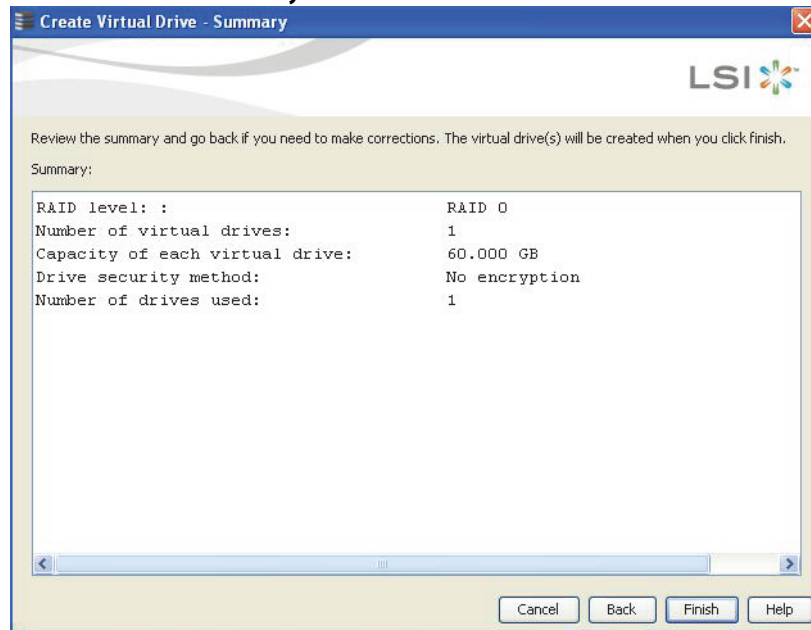
4. If you want to allow different types of drives in a configuration, select the **Use the drive type mixing** check box.



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 spare. Hot spares are drives that are available to replace failed drives automatically in a redundant virtual drive (RAID 1, RAID 5, or RAID 6).
7. Select the **Use drive security** check box if you want to set a drive security method.
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 13.5 Create Virtual Drive - Summary Window



11. Either click **Back** to return to the previous window to change any selections, or click **Finish** to accept and complete the configuration.

The new virtual drive is created after you click **Finish**. After the configuration is completed, a dialog box notifies you that the virtual drives were created successfully.



NOTE If you create a large configuration using drives that are in Power-Save mode, it could take several minutes to spin up the drives. A progress bar appears as the drives spin up. If any of the selected unconfigured drives fail to spin up, a dialog box that identifies these drives appears.

13.1.5 Creating a Virtual Drive Using Advanced Configuration

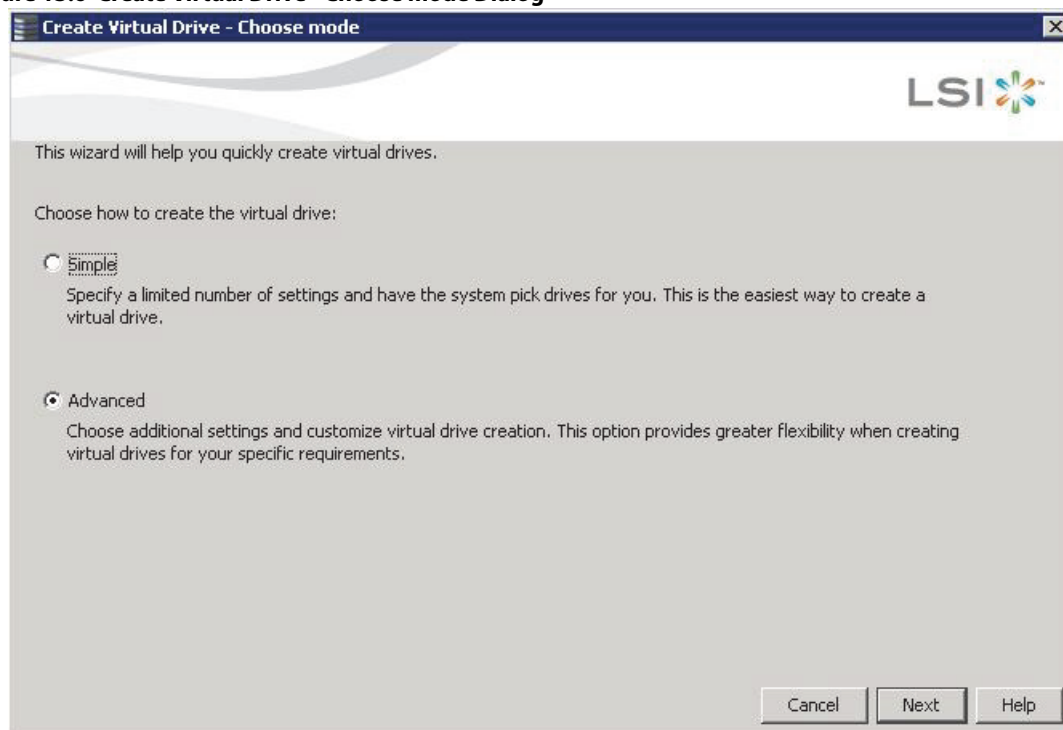
The advanced configuration procedure provides an easy way to create a new storage configuration. Advanced configuration gives you greater flexibility than simple configuration because you can select the drives and the virtual drive parameters when you create a virtual drive. In addition, you can use the advanced configuration procedure to create spanned drive groups.

Follow these steps to create a new storage configuration in the advanced configuration mode. This example shows the configuration of a spanned drive group.

1. Perform either of the following steps to bring up the **Configuration** wizard:
 - Right-click the controller node in the device tree in the left frame of the **MegaRAID Storage Manager** window, and select **Create Virtual Drive**.
 - Select the controller node, and select **Go To > Controller > Create Virtual Drive** in the menu bar.

The dialog for the choosing the configuration mode (simple or advanced) appears, as shown in the following figure.

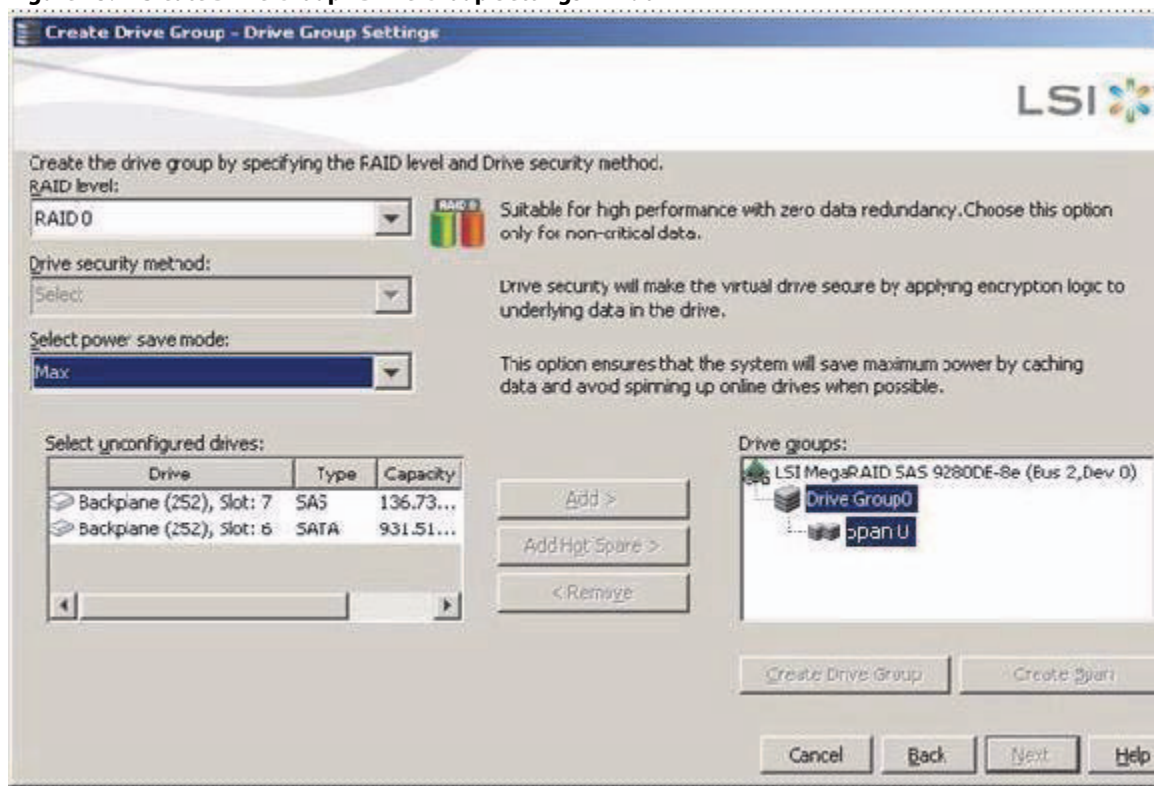
Figure 13.6 Create Virtual Drive - Choose mode Dialog



2. Select the **Advanced** radio button, and click **Next**.

The **Create Drive Group Settings** window appears, as shown in the following figure.

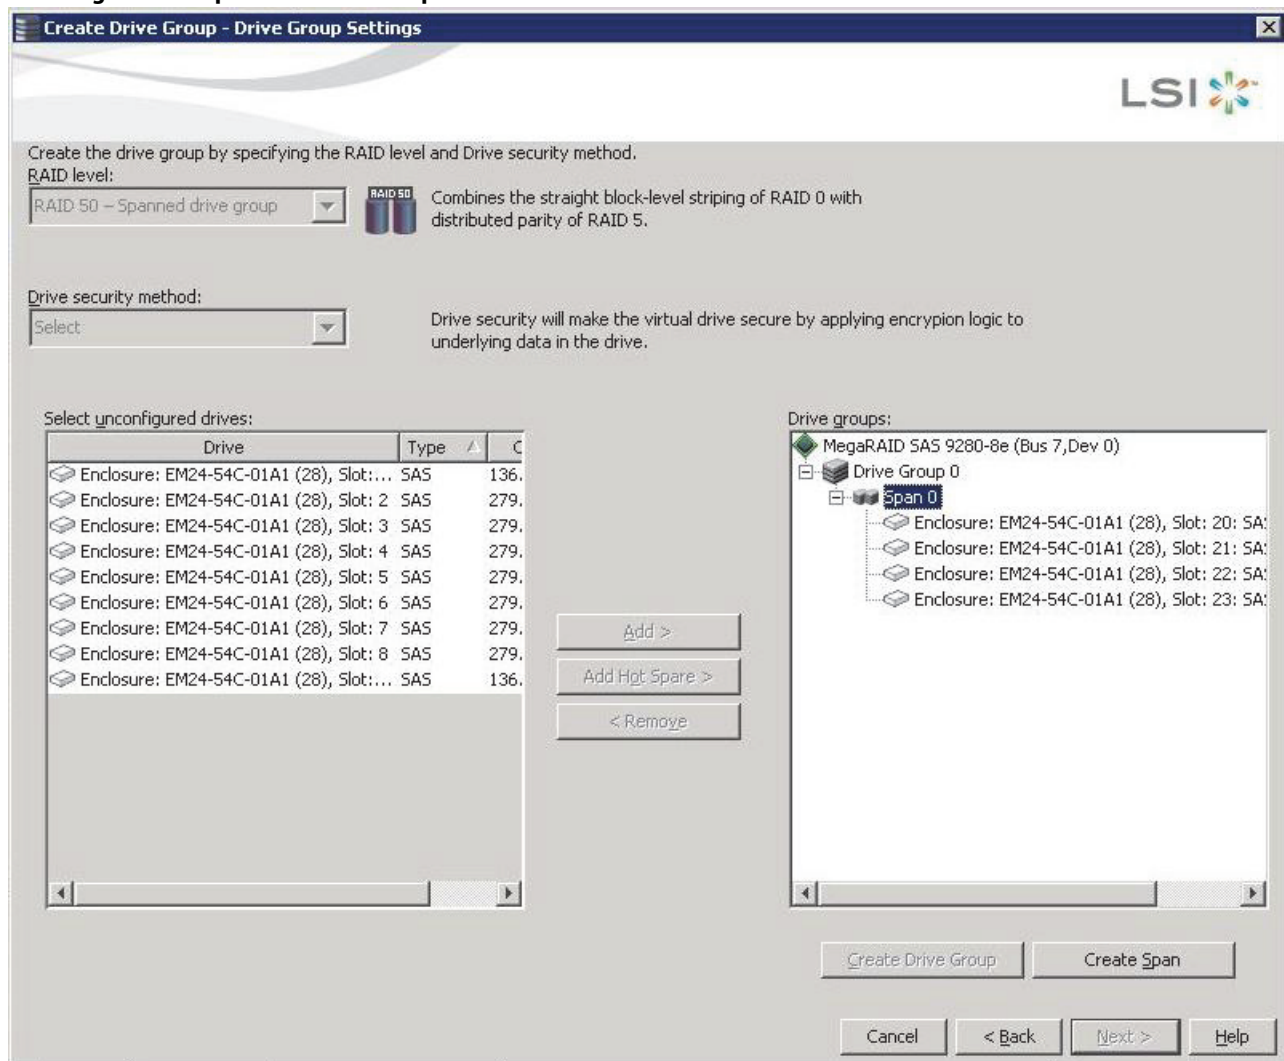
Figure 13.7 Create Drive Group - Drive Group Settings Window



3. Select the following items on the **Create Drive Group - Drive Group Settings** window:

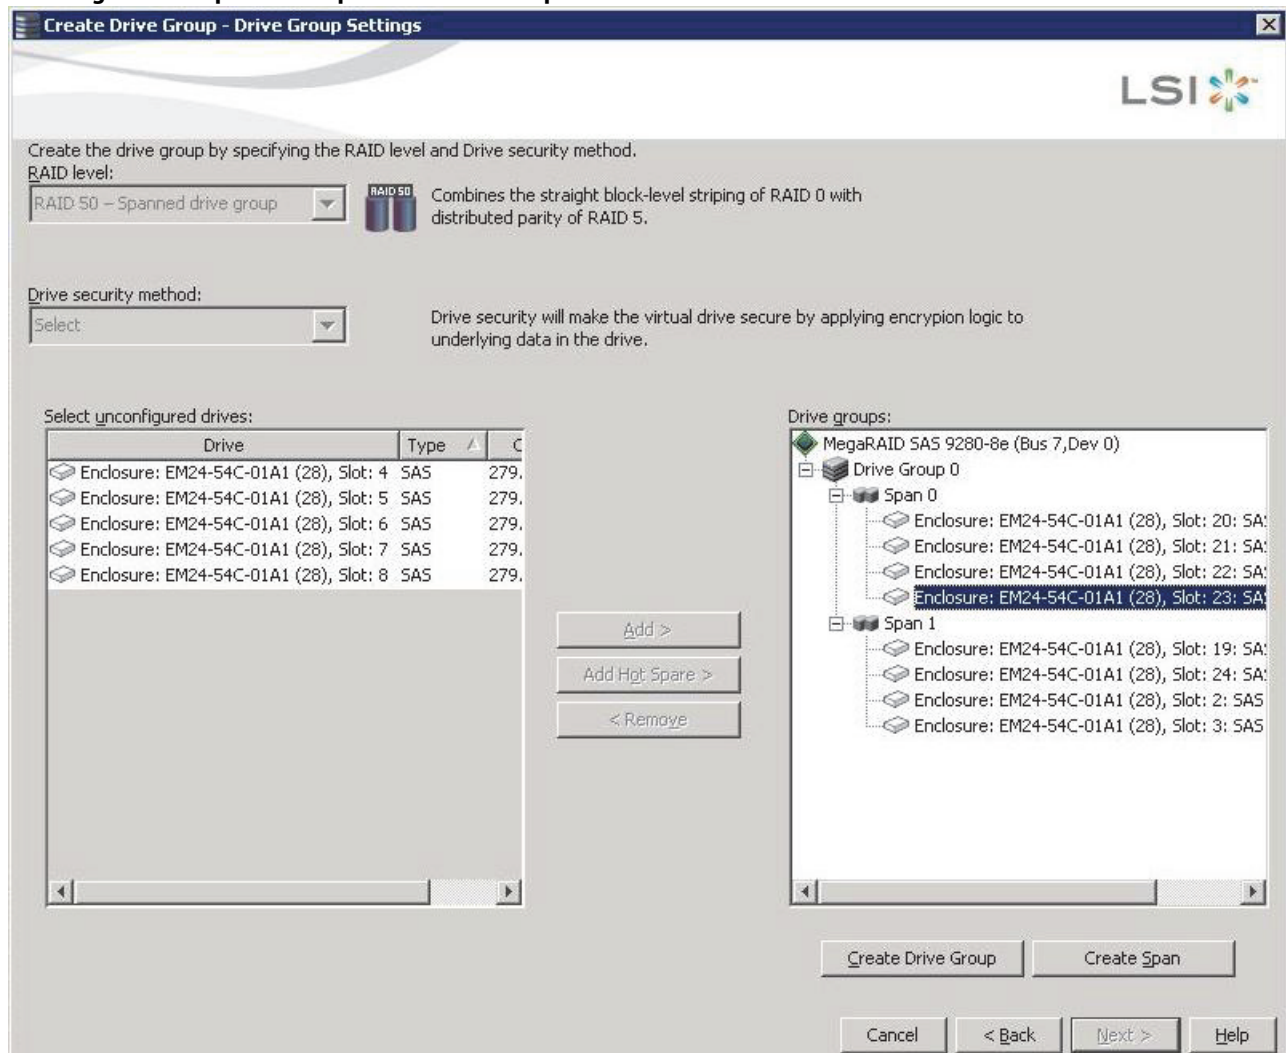
- a. Select the RAID level desired for the drive group from the drop-down menu. To make a spanned drive, select **RAID 10, RAID 50, or RAID 60** in the **RAID level** field.
Drive Group 0 and **Span 0** appear in the **Drive groups** field when you select RAID 10, 50, or 60.
The RAID controller supports 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.
- 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 13.8 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 13.9 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..

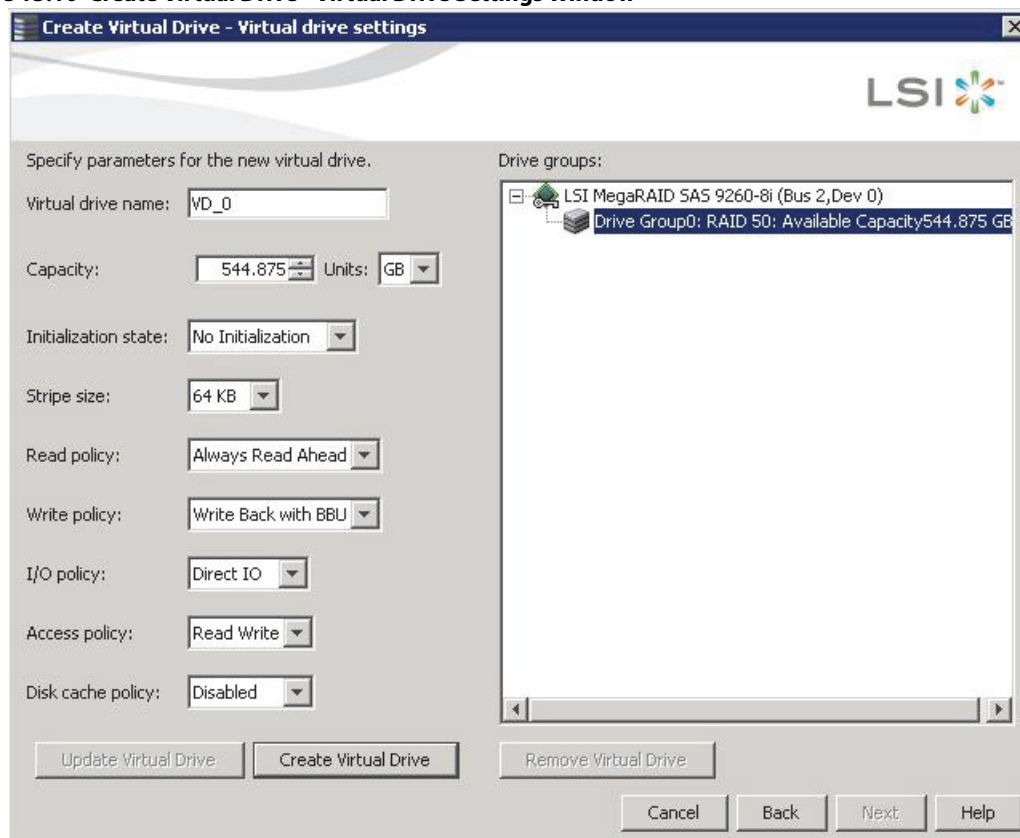


NOTE The parameters in the **Create Virtual Drive - Virtual drive settings** window display in Disabled mode (grayed out) for SAS-Integrated RAID (IR) controllers because these parameters do not apply to SAS-IR controllers.



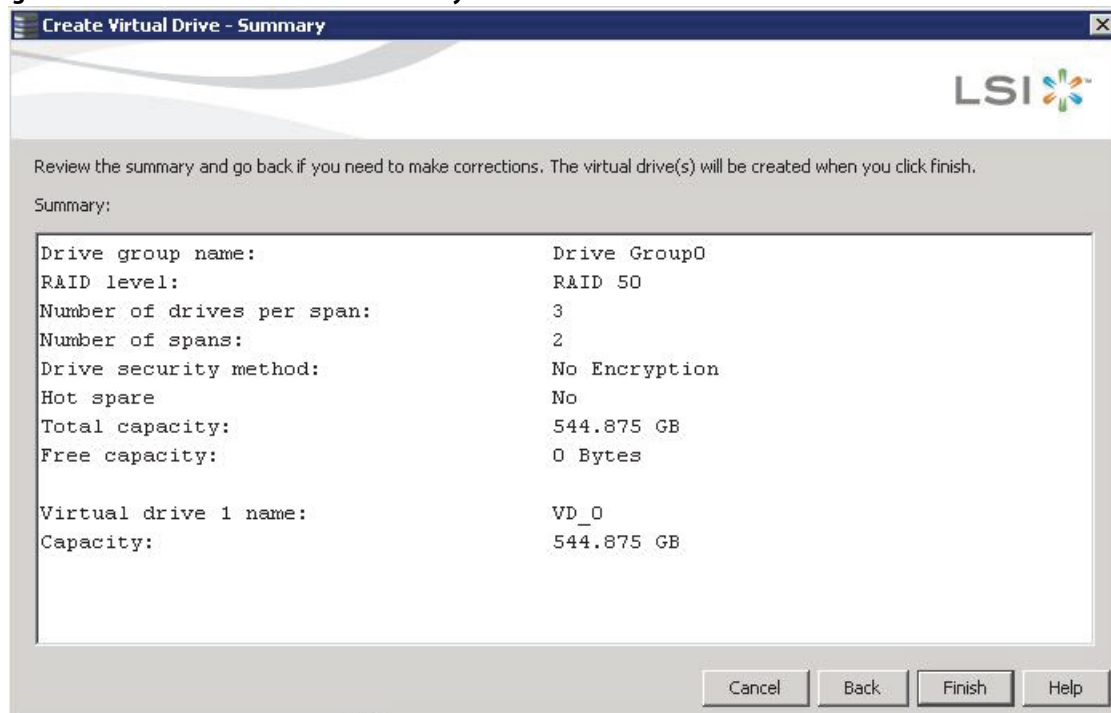
NOTE If you select **Write Back with BBU** as the write policy, and no battery exists, the battery is low or failed, or the battery is running through a re-learn cycle, the write policy switches to **Write Through**. This setting eliminates the risk of data loss in case of a power failure. A message window notifies you of this change.

Figure 13.10 Create Virtual Drive - Virtual Drive Settings Window



4. Change any virtual drive settings, if desired.
See Section [Section 13.1.1, Selecting Virtual Drive Settings](#), for more information about the virtual drive settings.
5. Click **Create Virtual Drive**.
The new virtual drive appears under the drive group. The options **Update Virtual Drive** and **Remove Virtual Drive** are available. **Update Virtual Drive** allows you to change the virtual drive settings, and **Remove Virtual Drive** allows you to delete the virtual drive.
6. Click **Next**.
The **Create Virtual Drive - Summary** window appears, as shown in the following figure. This window shows the selections you made for advanced configuration.

Figure 13.11 Create Virtual Drive - Summary Window



- Click **Back** to return to the previous window to change any selections, or click **Finish** to accept and complete the configuration.

After you click **Finish**, the new storage configuration is created and initialized according to the selected options.



NOTE If you create a large configuration using drives that are in Power-Save mode, it could take several minutes to spin up the drives. A progress bar appears as the drives spin up. If any of the selected unconfigured drives fail to spin up, a dialog appears that identifies the drives.

After the configuration is completed, a dialog notifies you that the virtual drives were created successfully.

- Click **OK**.
The **Enable SSD Caching on New Virtual Drives** dialog appears.
The newly created virtual drive is enabled for SSD caching by default.
- 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.
- Select either **Yes** or **No** to indicate whether you want to create additional virtual drives.
If you select **Yes**, the system takes you to the **Create Virtual Drive** window, as shown in [Figure 13.1](#). If you select **No**, the utility asks whether you want to close the wizard.
- 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.

13.2 Converting JBOD Drives to Unconfigured Good

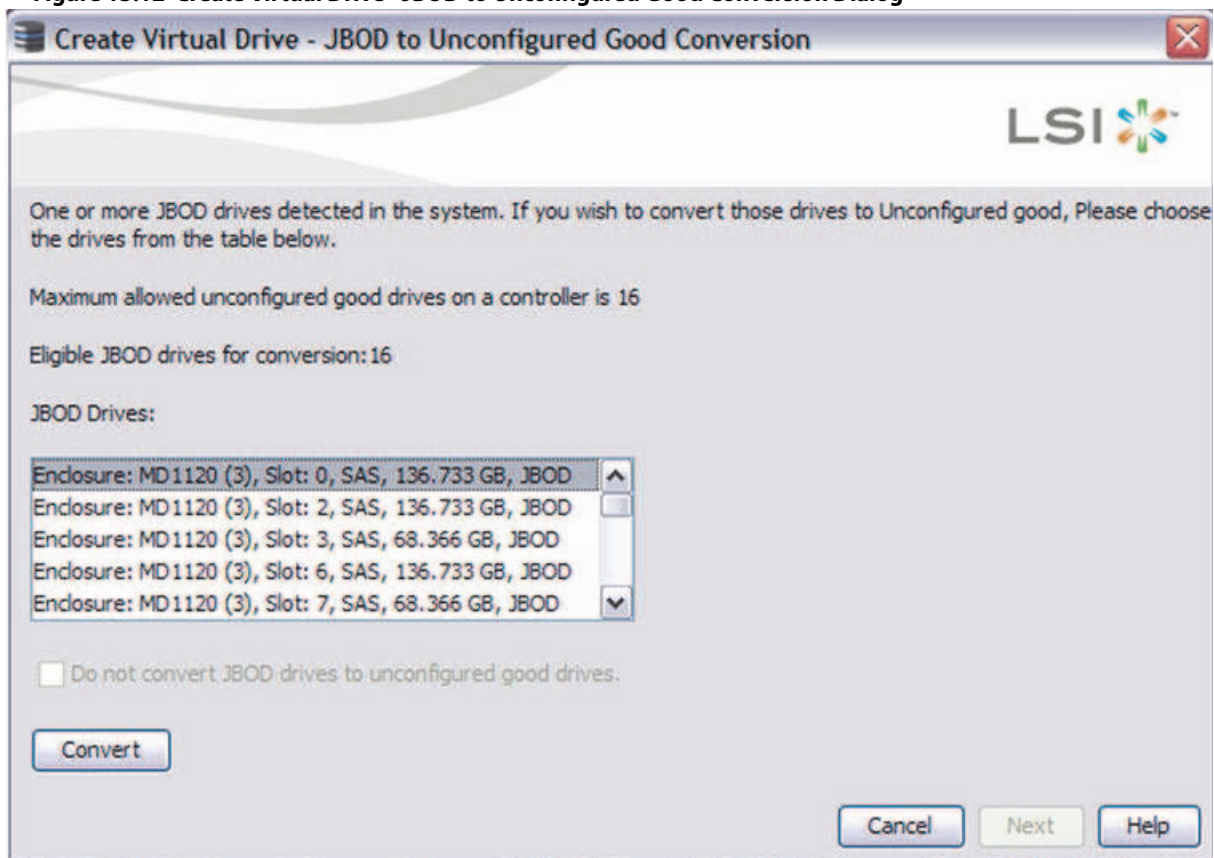
You can convert JBOD drives to Unconfigured Good using the **Create Virtual Drive** option or **Make Unconfigured Good** drive option with a single configuration.

Perform the following steps to configure JBOD to Unconfigured Good drives:

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 13.12 Create Virtual Drive - JBOD to Unconfigured Good Conversion Dialog



The **JBOD Drives** field displays the available JBOD drives available in the system.

2. Select the drives which you want configured as Unconfigured Good and then click **Convert**. Clicking on **Convert** configures the selected JBODs to Unconfigured Good Drives.



NOTE If you do not 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.

3. Click **Next**.

The **Create Virtual Drive - Drive group and Virtual drive settings** dialog appears.

13.2.1 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 13.13 Make Configured Good Dialog



2. Select the JBOD drives to be configured as unconfigured good.
3. Click **OK**.

The selected JBOD drives are configured as unconfigured good.

13.3 Adding Hot Spare Drives

Hot spares are drives that are available to automatically replace failed drives in a RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, or RAID 60 virtual drive. *Dedicated hot spares* can be used to replace failed drives in a selected drive group only. *Global hot spares* are available to any virtual drive on a specific controller.

To add a dedicated or global hot spare drive, follow these steps:

1. Select the **Physical** tab in the left panel of the MegaRAID Storage Manager main menu, and click the icon of an unused drive.
For each drive, the window displays the port number, enclosure number, slot number, drive state, drive capacity, and drive manufacturer.
2. Either select **Go To > Physical Drive > Assign Global Hot Spare**, or select **Go To > Physical Drive > Assign Dedicated Hot Spare**.
3. Perform one of these actions:
 - If you selected **Assign Dedicated Hot Spare**, select a drive group from the list that appears. The hot spare is dedicated to the drive group that you select.

- If you selected **Assign Global Hot Spare**, 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.

13.4 Changing Adjustable Task Rates

If you want to change the Rebuild rate and other task rates for a controller, you must first log onto the server in Full Access mode.



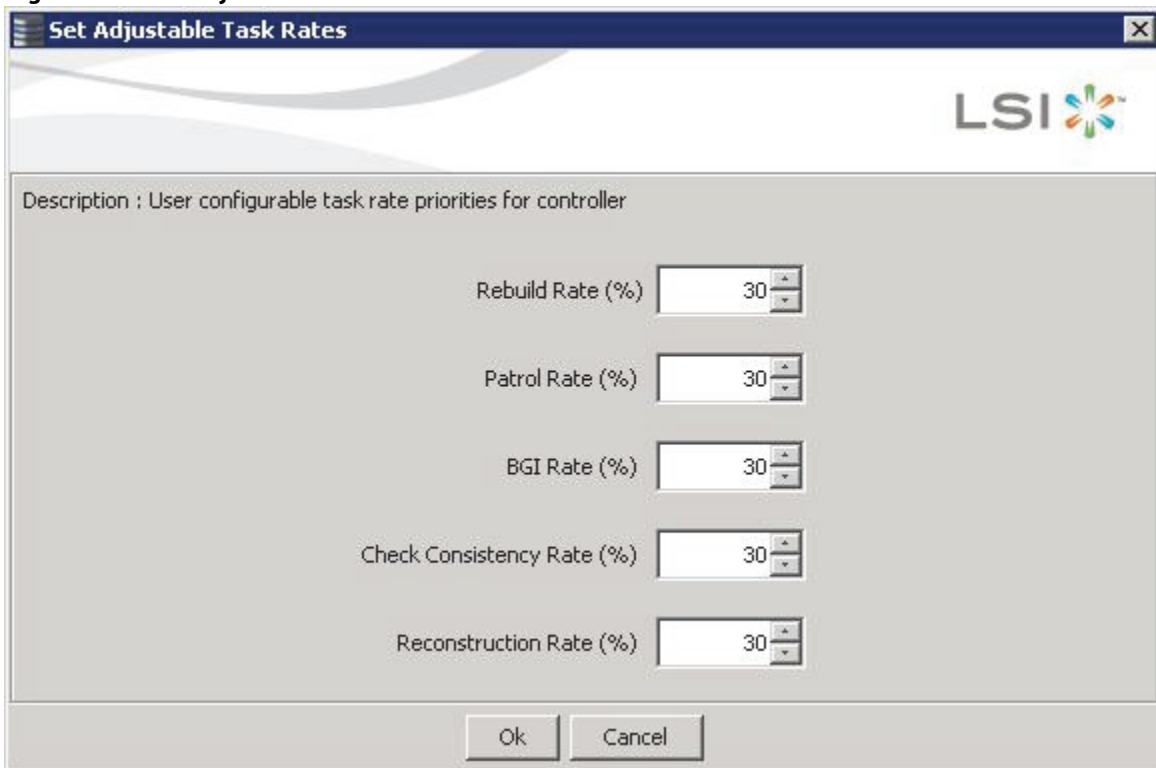
NOTE Leave the adjustable task rates at their default settings to achieve the best system performance. If you raise the task rates above the defaults, foreground tasks will run more slowly and it might seem that the system is not responding. If you lower the task rates below the defaults, rebuilds and 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 13.14 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.

13.5 Changing Power Settings

The RAID controller includes Dimmer Switch technology that conserves energy by placing certain unused drives into Power-Save mode. In Power-Save mode, the drives use less energy, and the fan and the enclosure require less energy to cool and house the drives, respectively. Also, this technology helps avoid application timeouts caused by spin-up delays and drive wear caused by excessive spin-up/down cycles.

You can use the **Power Settings** field in the MegaRAID Storage Manager software to choose whether to allow unconfigured drives or Commissioned Hotspares to enter Power-Save mode.



NOTE The Dimmer Switch technology is enabled by default.

When they are in the Power-Save mode, unconfigured drives and drives configured as Commissioned Hotspares (dedicated or global) can be spun down. When spun down, the drives stay in Power-Save mode except for periodic maintenance, which includes the following:

- Periodic background media scans (Patrol Read) to find and correct media defects to avoid losing data redundancy (hot spare drives only)
- Use of a Commissioned Hotspare to rebuild a degraded drive group (Commissioned Hotspare drives only)
- Update of disk data format (DDF) and other metadata when you make changes to RAID configurations (Commissioned Hotspare drives and unconfigured drives)



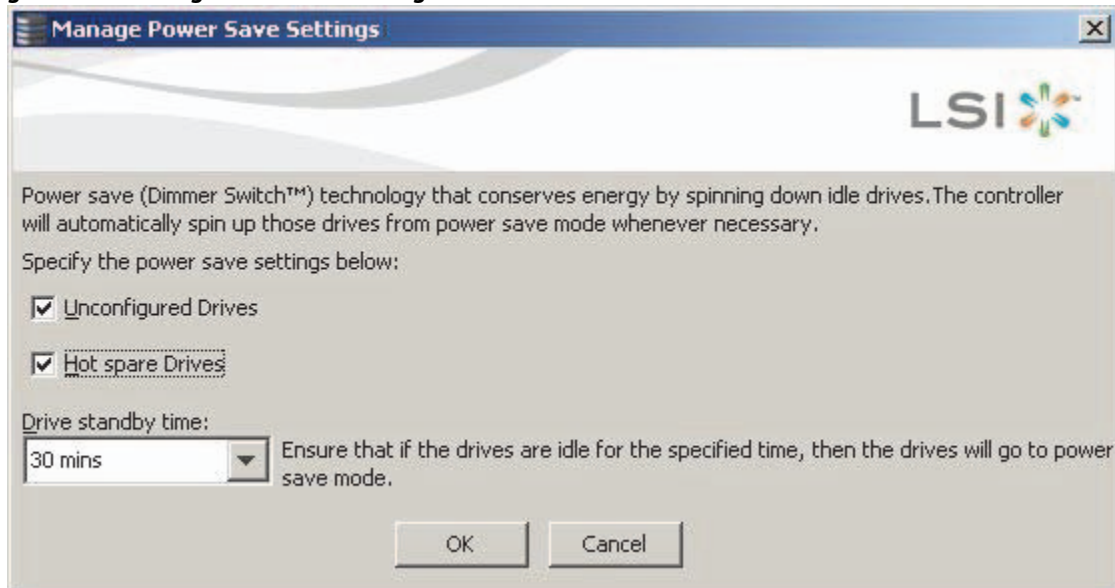
NOTE If your controller does not support this option, the **Power Settings** field does not appear.

Follow these steps to change the power-save setting.

1. Select a controller icon in the **Physical** tab or the **Logical** tab in the left panel of the **MegaRAID Storage Manager** window.
2. Select **Go To > Controller > Manage Power Settings** from the menu bar.

The **Manage Power Save Settings** dialog appears.

Figure 13.15 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.

13.6 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 [Section 13.1.1, *Selecting Virtual Drive Settings*](#).

13.7 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.

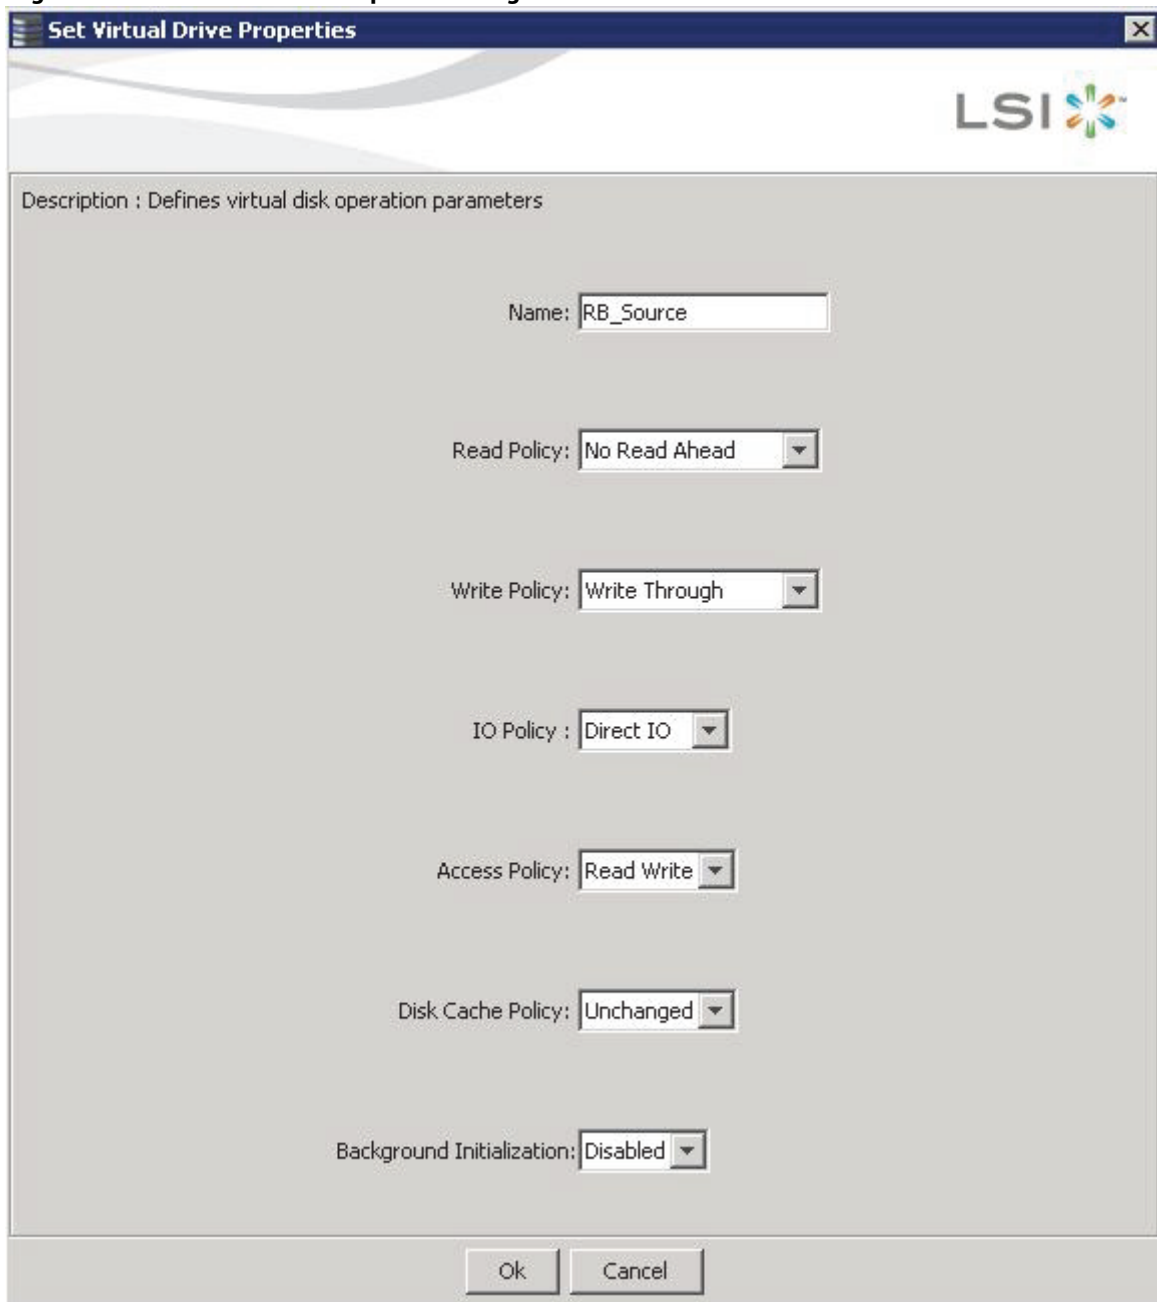


NOTE For virtual drives with SAS drives only, set the drive write cache policy set to **Disabled**, by default. For virtual drives with SATA drives only, set the drive write cache policy to **Enabled**, by default.

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.

Figure 13.16 Set Virtual Drive Properties Dialog



The dialog box is titled "Set Virtual Drive Properties" and features the LSI logo in the top right corner. Below the title bar, a description reads: "Description : Defines virtual disk operation parameters". The main area contains several settings, each with a label and a text or dropdown field:

- Name: RB_Source
- Read Policy: No Read Ahead
- Write Policy: Write Through
- IO Policy : Direct IO
- Access Policy: Read Write
- Disk Cache Policy: Unchanged
- Background Initialization: Disabled

At the bottom of the dialog are "Ok" and "Cancel" buttons.

3. Change the virtual drive properties as required.
For information about these properties, see [Section 13.1.1, *Selecting Virtual Drive Settings*](#).
4. Click **Ok** to accept the changes.
The virtual drive settings are updated.

13.8 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.)

13.8.1 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 13.17 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.

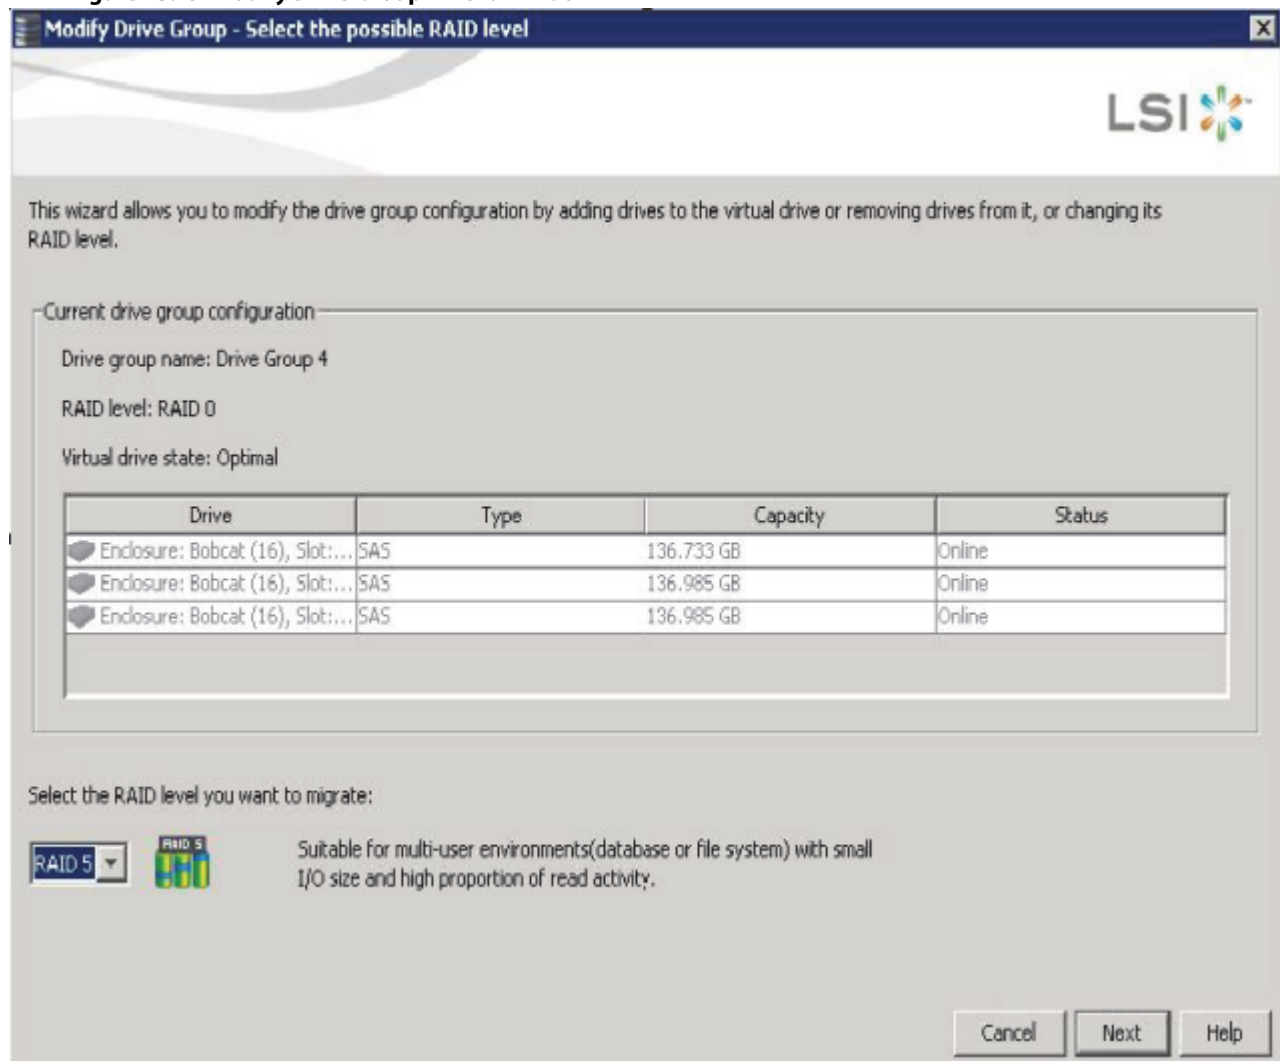
Figure 13.18 Warning to Back Up Data Message



5. Select the **Confirm** check box, and click **Yes**.

The **Modify Drive Group** wizard window appears, as shown in the following figure.

Figure 13.19 Modify Drive Group Wizard Window



The following sections explain the **Modify Drive Group** wizard options.

13.8.2 Adding a Drive or Drives to a Configuration



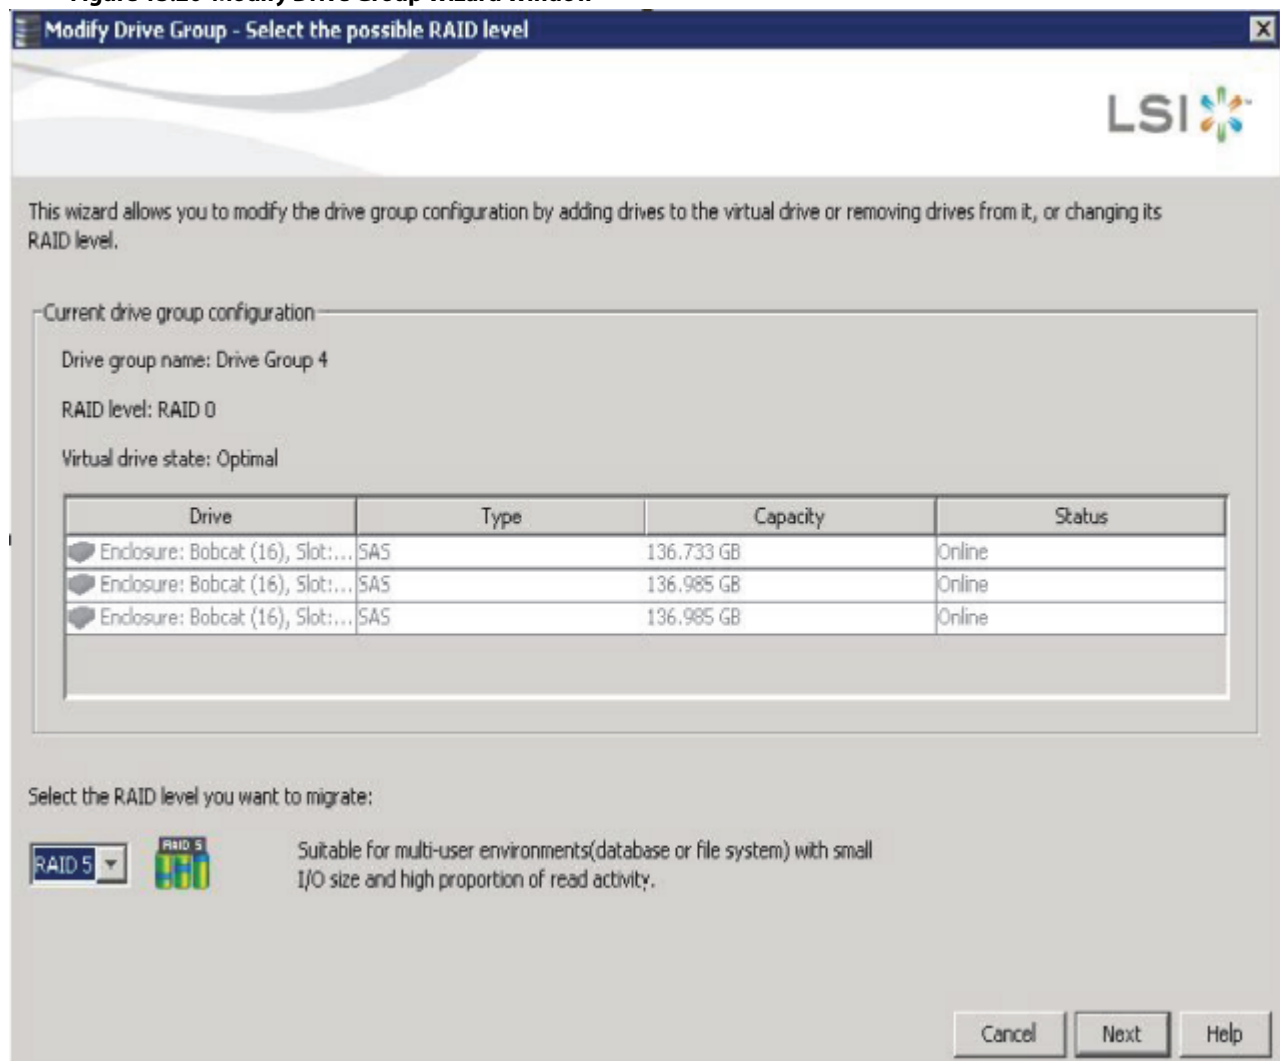
ATTENTION Be sure to back up the data on the virtual drive before you add a drive to it.

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.

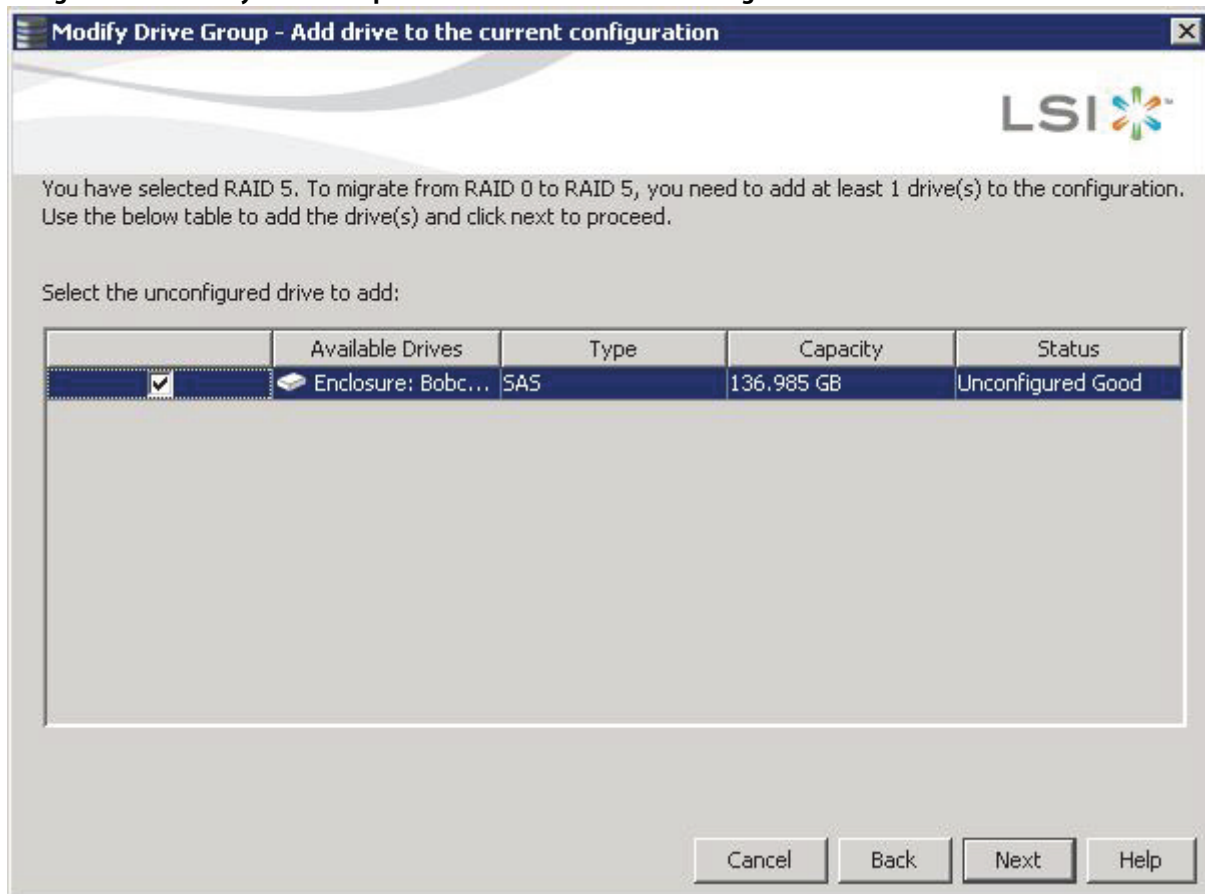
Figure 13.20 Modify Drive Group Wizard Window



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 13.21 Modify Drive Group – Add Drives to the Current Configuration Window



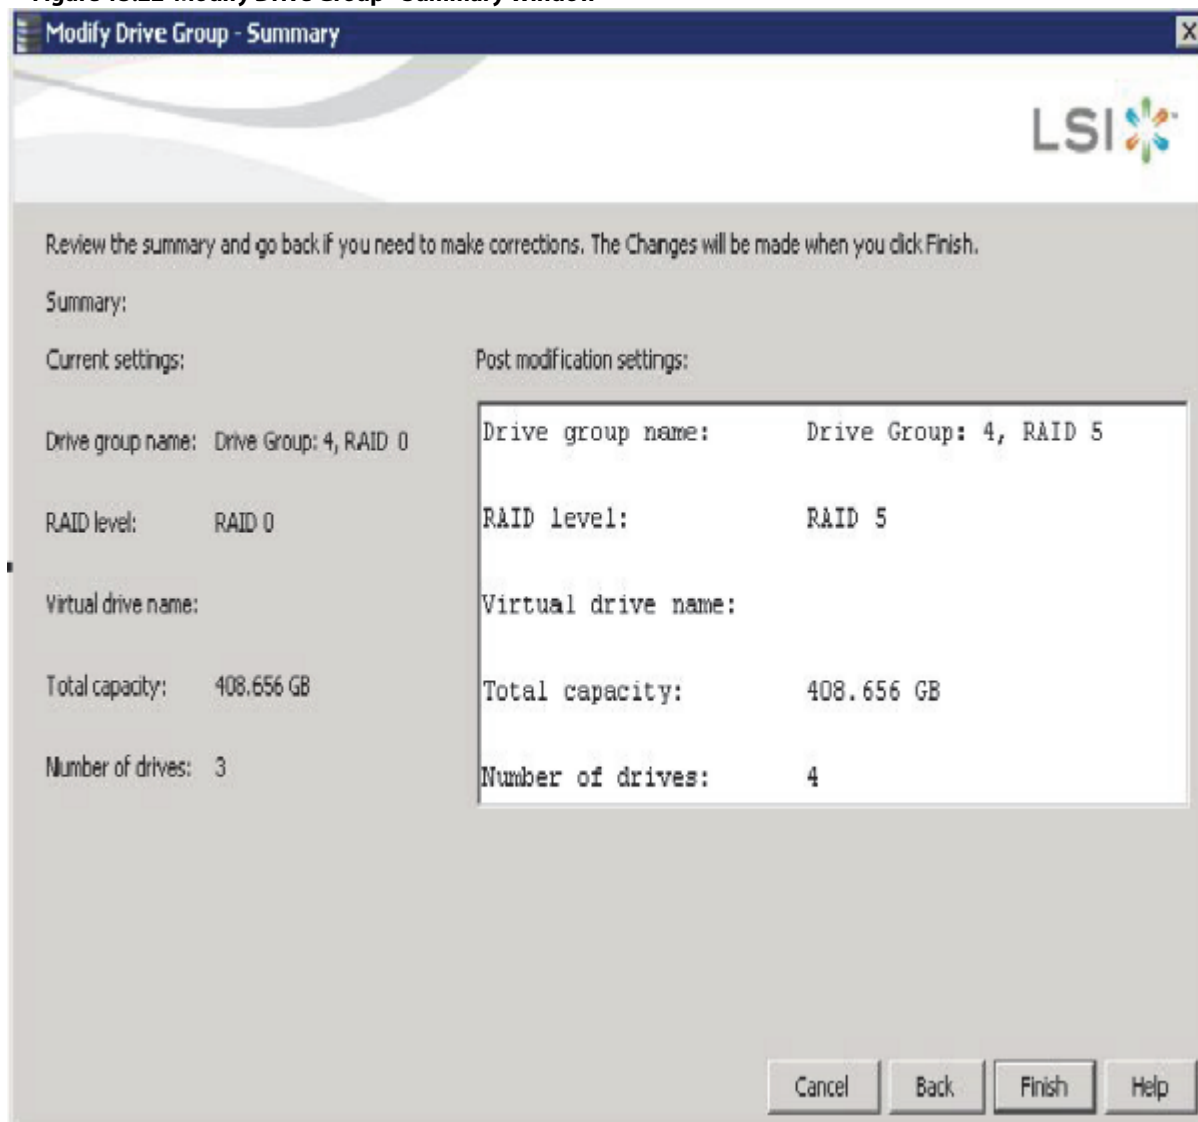
5. Click the check box next to any unconfigured drives that you want to add, and then click **Next**.



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 13.22 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.

13.8.3 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.

13.8.4 Replacing a Drive

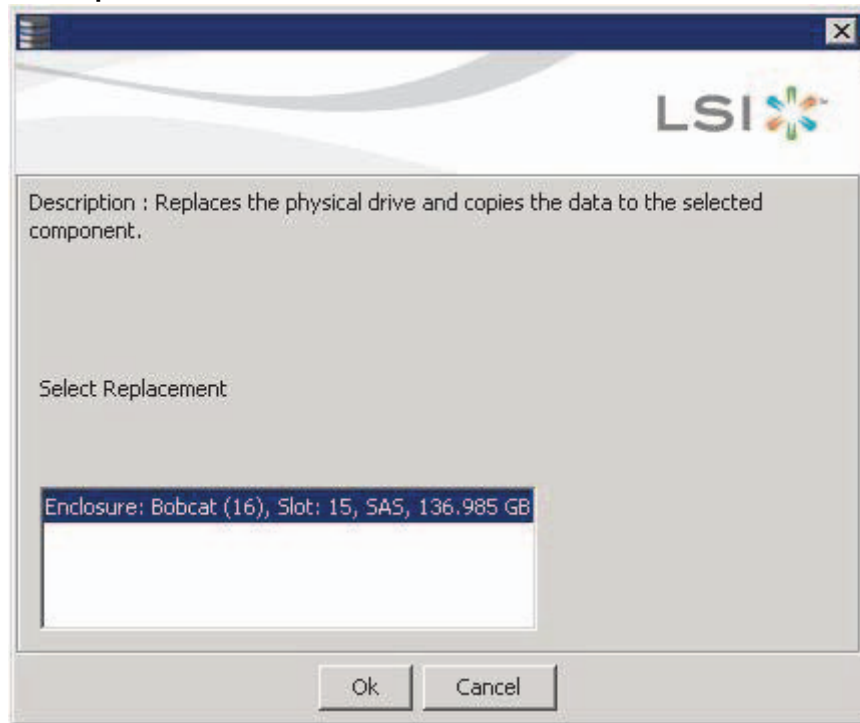


ATTENTION Be sure to back up the data on the virtual drive before you replace a drive.

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 13.23 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.

13.8.5 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.



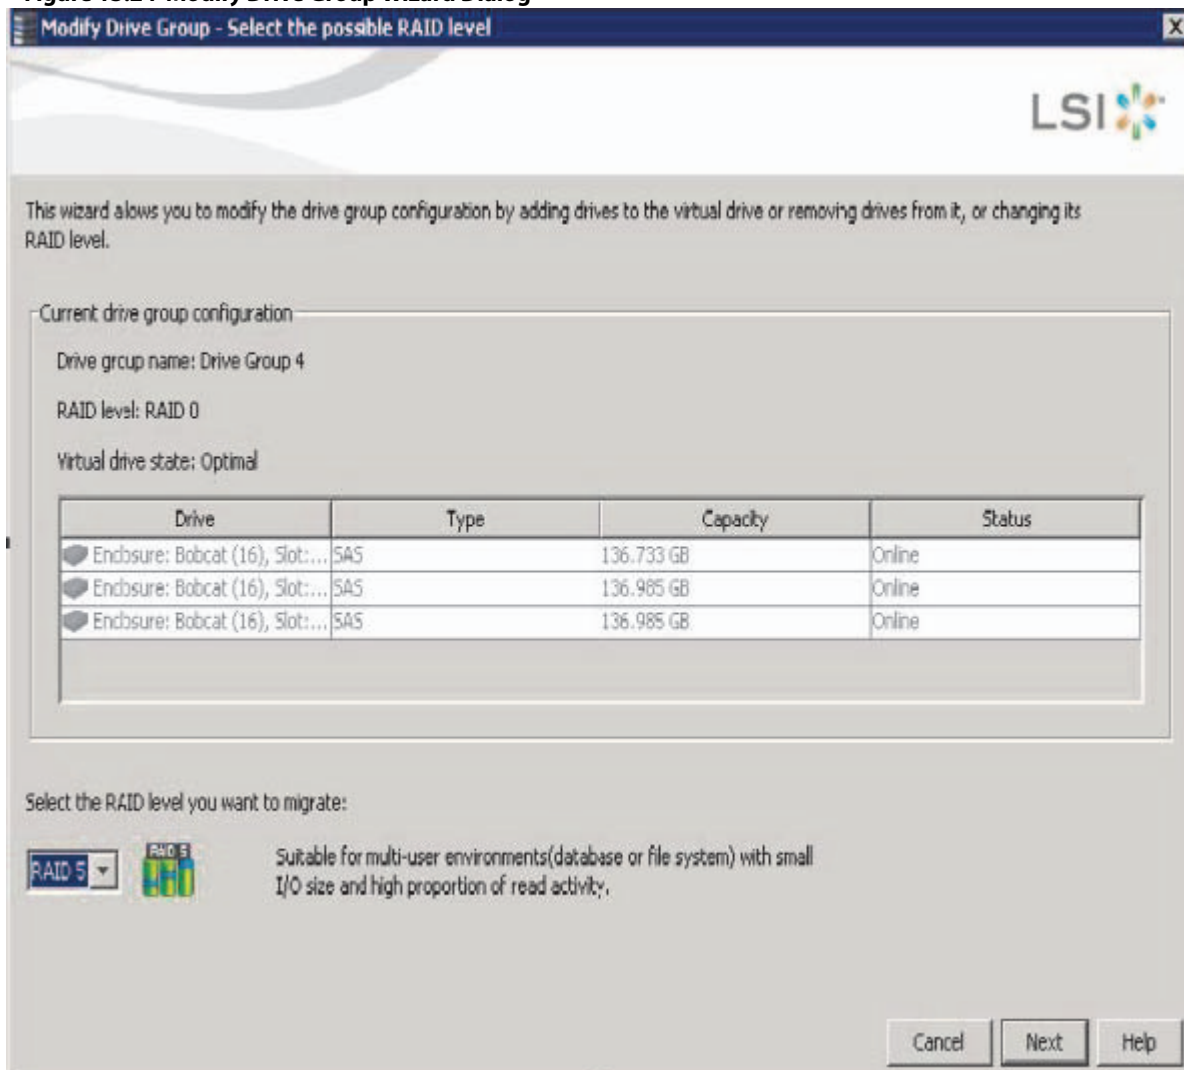
ATTENTION Be sure to back up the data on the virtual drive before you change the RAID level.

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.

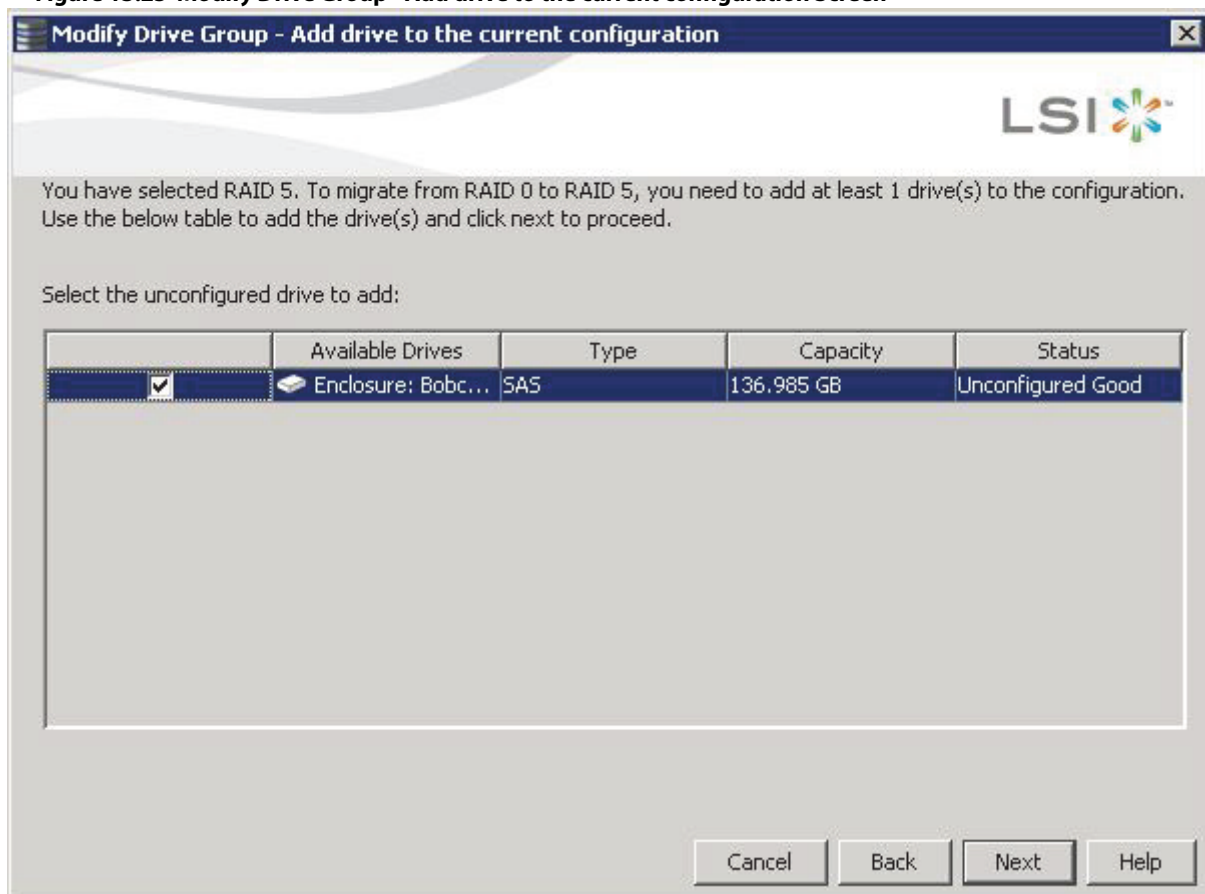
Figure 13.24 Modify Drive Group Wizard Dialog



- On the **Modify Drive Group Wizard** dialog, select the RAID level to which you want to change ("migrate") the drive group to, and click **Next**.

The following dialog appears. The dialog states the number of drives that you have to add to change the RAID level from the current level to a new RAID level that requires more drives.

Figure 13.25 Modify Drive Group - Add drive to the current configuration Screen



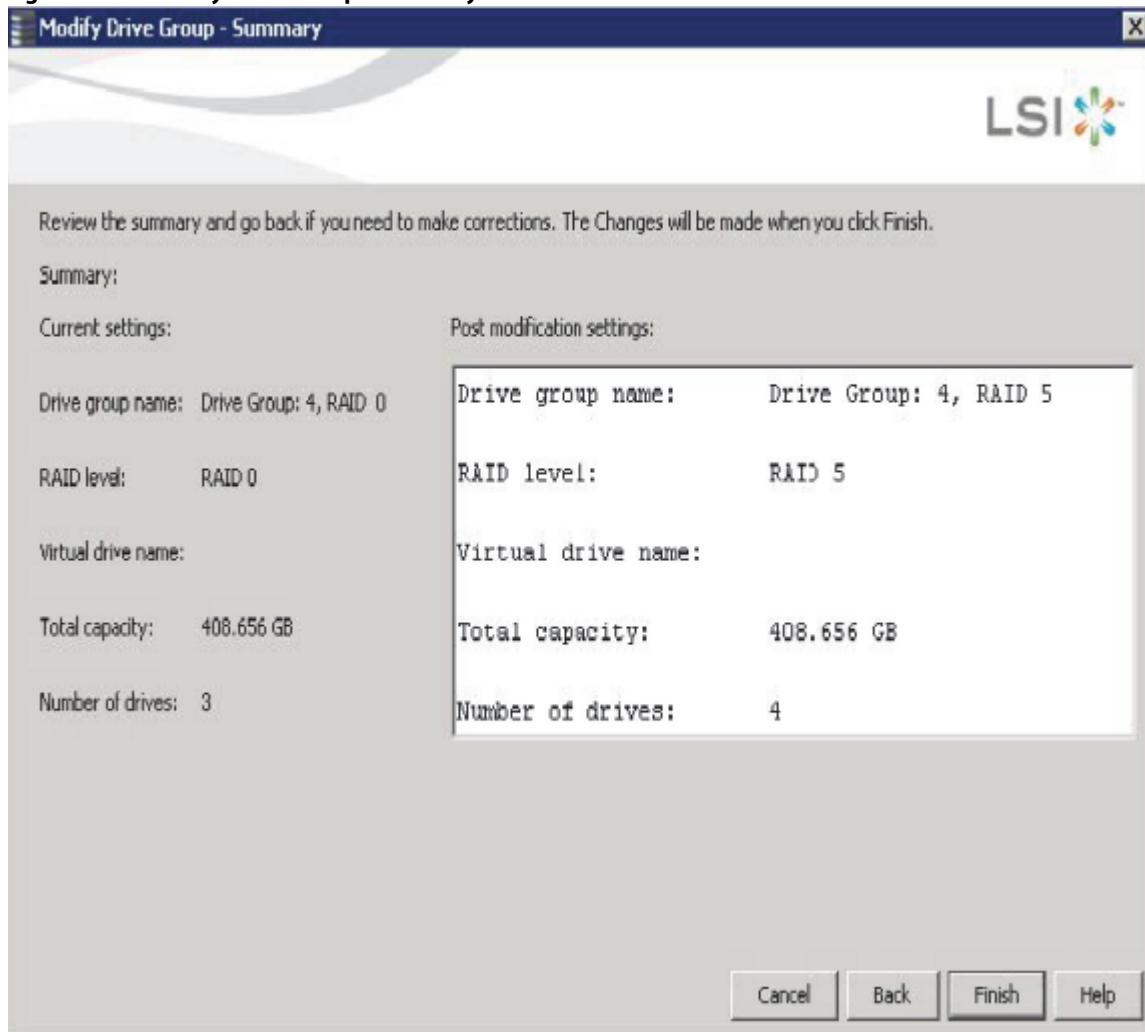
5. Select the unconfigured drive or drives to add, and click **Next**.



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 13.26 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.

13.8.6 New Drives Attached to a ServeRAID Controller

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

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

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

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

See the section **Making a Drive Offline or Missing** for the procedure to change a drive to the Unconfigured Good drive state.

13.9 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.

Chapter 14: Monitoring System Events and Storage 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 synch.

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.

14.1 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 [Section 14.2, Configuring Alert Notifications](#). The different alert delivery methods are as follows:

- Vivaldi Log/MegaRAID Storage Manager Log
- System Log
- Pop-up Notification
- Email Notification

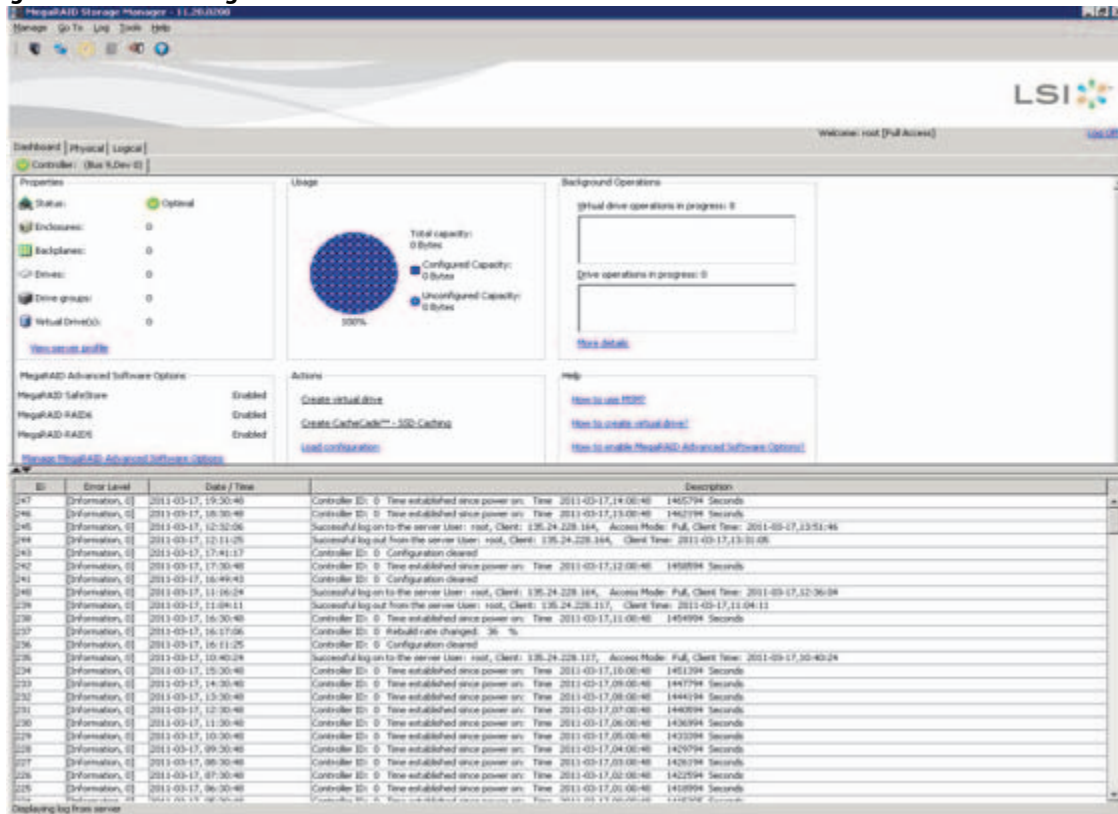
Table 14.1 Severity Level and Default Alert Delivery Methods

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.

14.1.1 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 14.1 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. For a list of all events, see [Appendix A: Events and Messages](#). The status bar at the bottom of the screen indicates whether the log is a MegaRAID Storage Manager server log or a locally stored log file.

When a Vivaldi log/MegaRAID Storage Manager log appears, the Log menu has the following options:

- **Save Log:** Saves the current log to a .log file.
- **Save Log Text:** Saves the current log in .txt format.
- **Load:** Enables you to load a local .log file in the bottom of the MegaRAID Storage Manager main menu window. If you select the **Load** menu, you will not be able to view the current log.
- **Rollback to Current Log:** This menu appears if we have loaded the logs from a local .log file. Once you select this menu, you can view the current log.
- **Clear Log:** Clears the current log information, if you have full access (versus view-only access). You have the option to save the log first.

14.1.2 System Log

By default, all the severity events are logged in the local syslog. Based on the operating system you are using, the system log is logged in the following syslog locations:

- In Windows, the system log is logged in **Event Viewer > Application**.

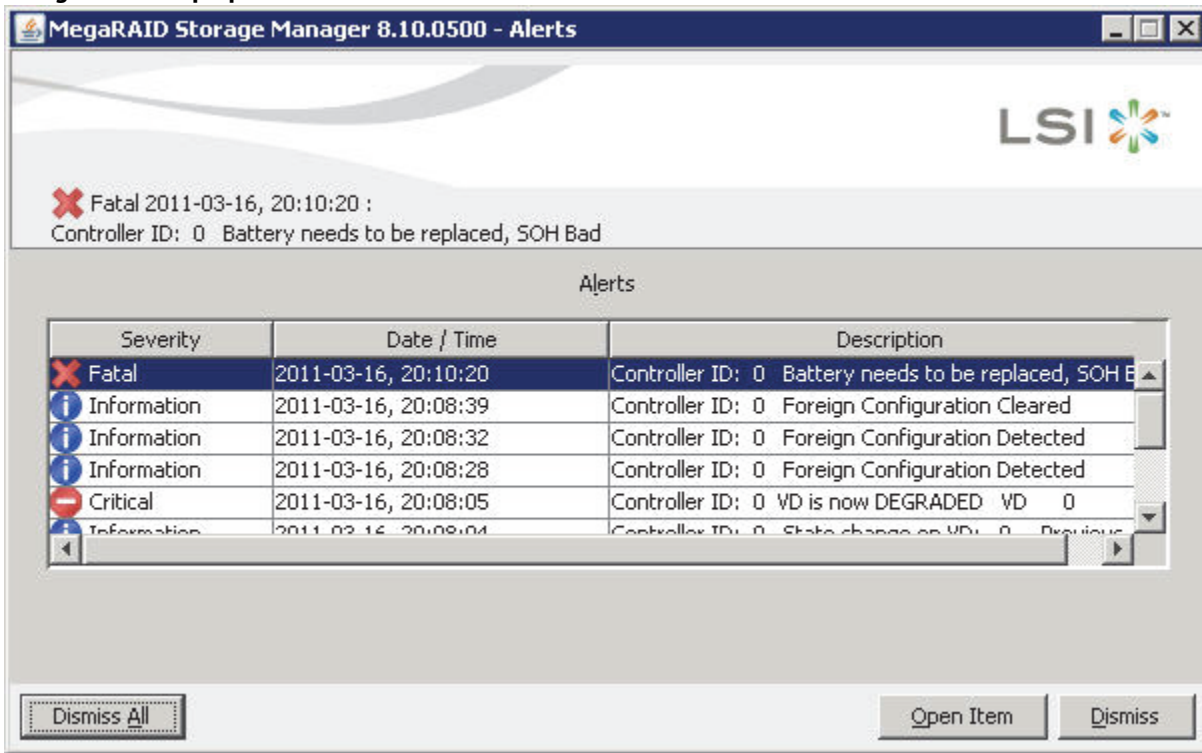
- In Linux, the system log is logged in `/var/log/messages`.
- In Solaris, the system log is logged in `/var/adm/messages`.

14.1.3 Pop-up Notification

By default, fatal and critical events are displaying in a pop-up notification. Pop-up notification is started automatically when you 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 14.2 Pop-up Notification



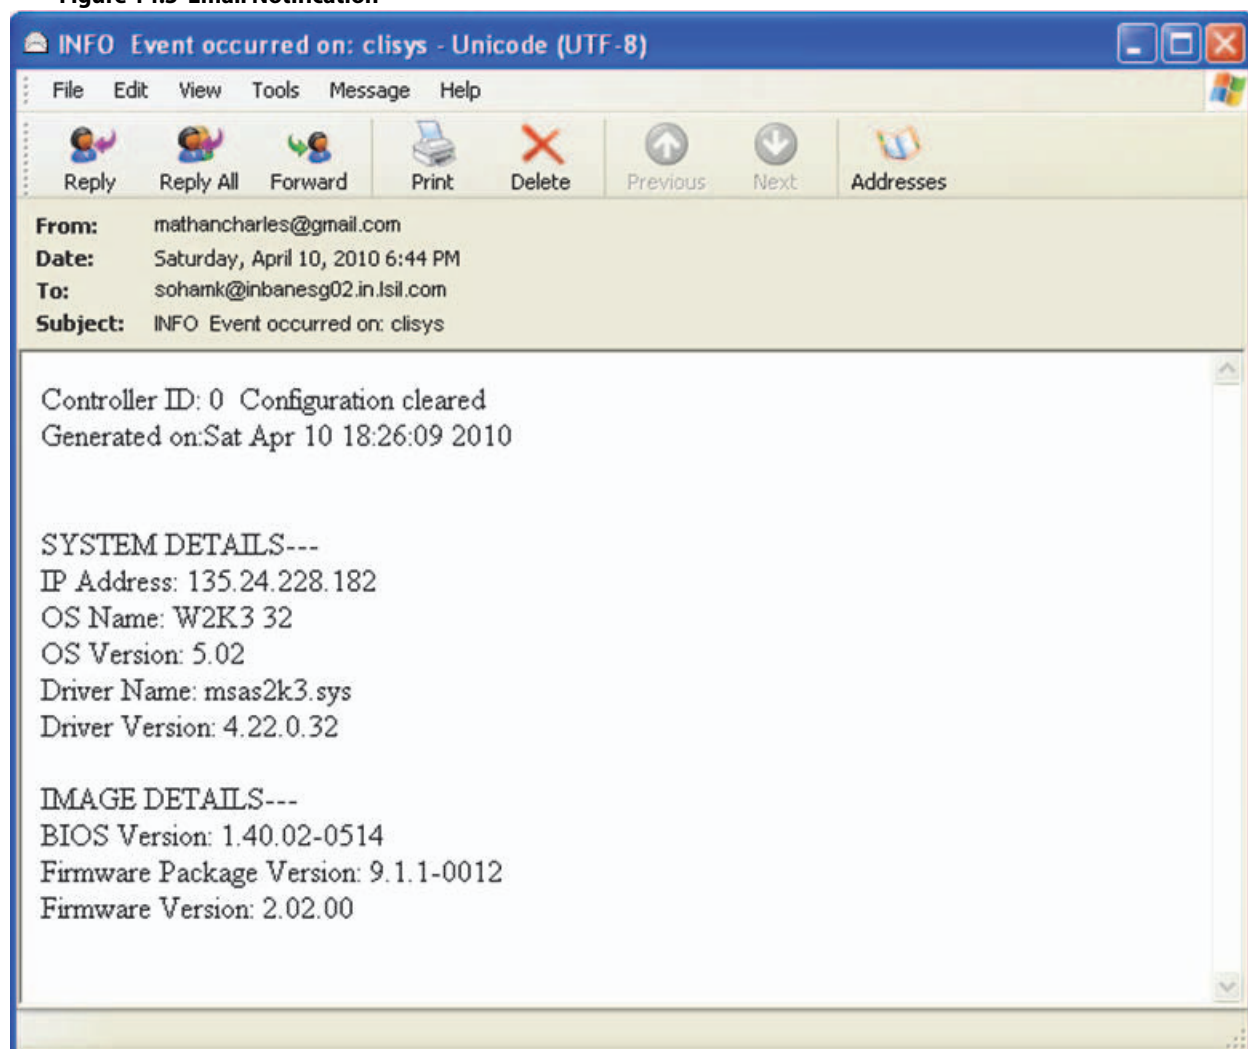
14.1.4 Email Notification

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 14.3 Email Notification



14.2 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.

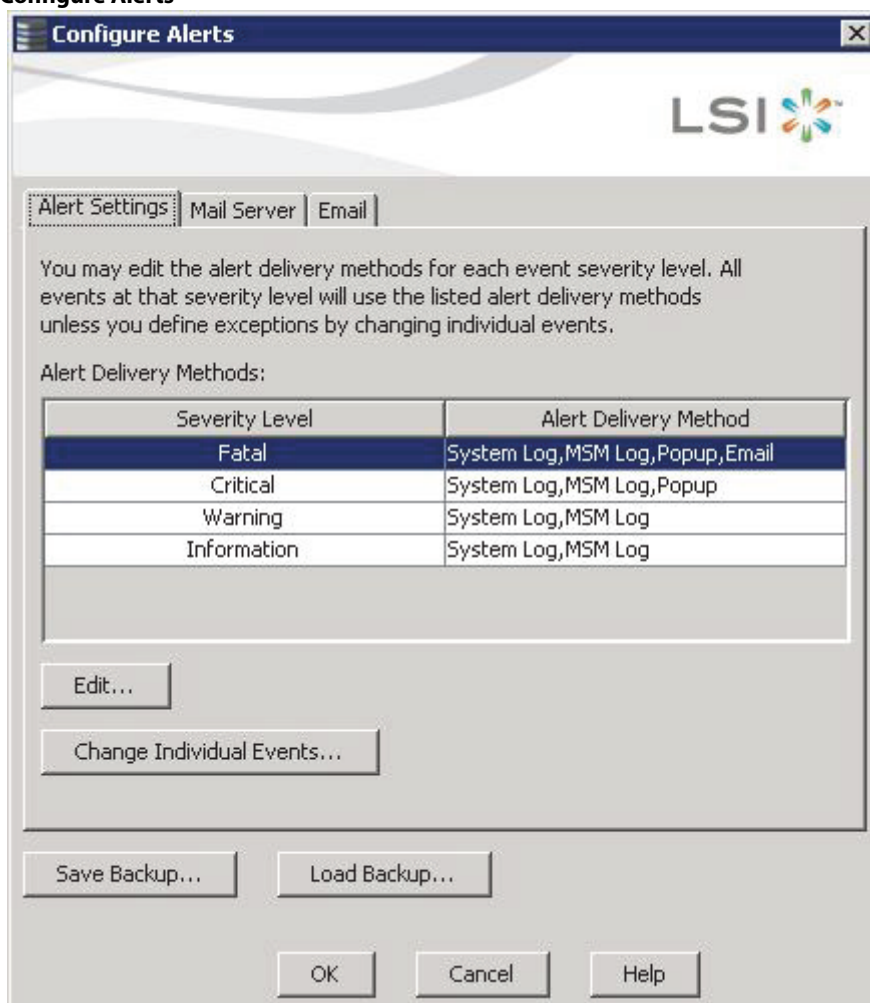
Select **Tools > Configure Alerts** on the main menu screen.



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 14.4 Configure Alerts



You can select the **Alert Settings** tab to perform the following actions:

- Edit the alert delivery method for different severity levels.
- Change the method of delivery for each individual event.
- Change the severity level of each individual event.
- Save an .xml backup file of the entire alert configuration.
- Load all the values from a previously saved backup into the dialog to edit or save these values as the current alert notification configuration.



NOTE When you load a saved backup file, all unsaved changes made in the current session will be lost.

You can select the **Mail Server** tab to perform the following actions:

- Enter or edit the sender email address.
- Enter the SMTP server name or the IP address.
- Enter the SMTP server authentication related information (user name and password).



NOTE These fields are optional and are filled only when the SMTP server requires authentication.

- Save an .xml backup file of the entire alert configuration.
- Load all of the values from a previously saved backup into the dialog to edit or save these values as the current alert notification configuration.



NOTE When you load a saved backup file, all unsaved changes made in the current session will be lost.

You can select the **Email** tab to perform the following actions:

- Add new email addresses for recipients of alert notifications.
- Send test messages to the recipient email addresses.
- Remove email addresses of recipients of alert notifications.
- Save an .xml backup file of the entire alert configuration.
- Load all of the values from a previously saved backup into the dialog to edit or save these values as the current alert notification configuration.



NOTE When you load a saved backup file, all unsaved changes made in the current session will be lost.

14.3 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 14.5 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.

14.4 Changing Alert Delivery Methods for Individual Events

You can change the alert delivery options for an event without changing the severity level.

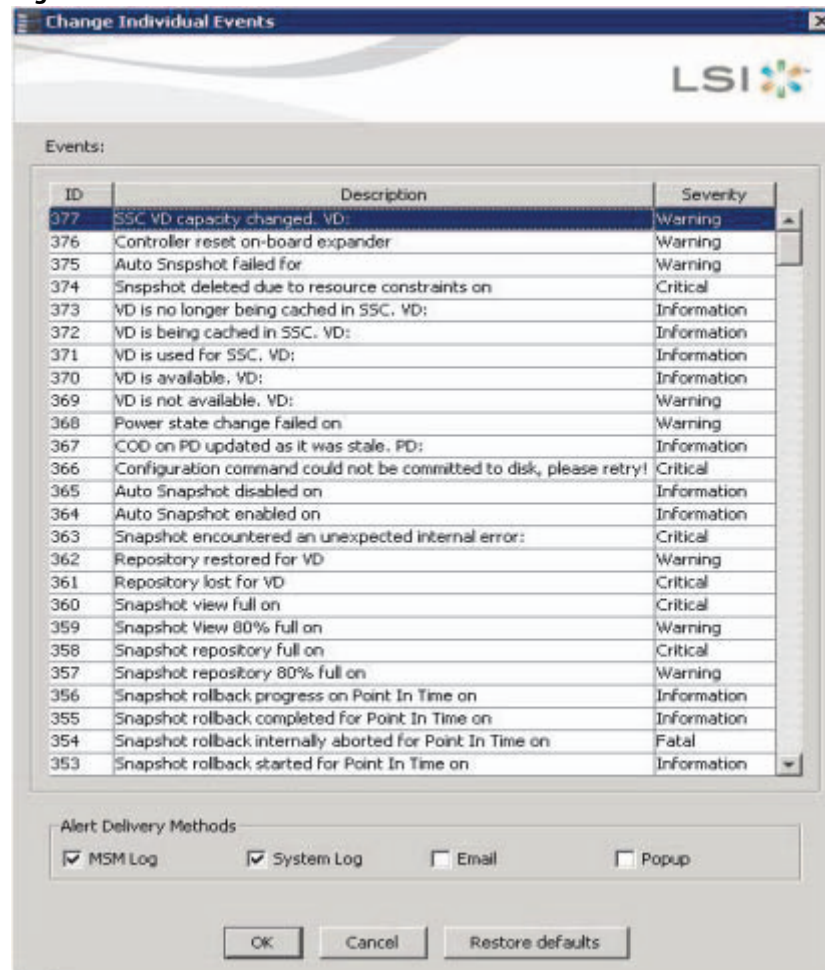
1. On the **Configure Alerts** window, click the **Alerts Setting** tab.

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 14.6 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**.



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 [Section 14.6, Roll Back to Default Individual Event Configuration](#).

14.5 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 14.7 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.

14.6 Roll Back to Default Individual Event Configuration

To revert back to the default alert delivery method and the default severity level of an individual event, perform the following steps:

1. On the **Configure Alerts** window, click the **Alerts Setting** tab.
The **Alerts Setting** portion of the window appears.
2. Click **Change Individual Events**.
The **Change Individual Events** dialog appears. 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.

14.7 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 14.8 Mail Server Options

The screenshot shows the 'Configure Alerts' dialog box with the 'Mail Server' tab selected. The dialog has a title bar with a close button (X) and the LSI logo. Below the title bar are three tabs: 'Alert Settings', 'Mail Server' (selected), and 'Email'. The 'Mail Server' tab contains the following fields and options:

- Sender email address:** A text field containing 'monitor@server.com'.
- SMTP Server:** A text field containing '127.0.0.1'.
- Port:** A text field containing '25'.
- ☒ **Use Default**
- ☒ **This server requires authentication**
- User name:** An empty text field.
- Password:** An empty text field.

At the bottom of the dialog are three buttons: 'Save Backup...', 'Load Backup...', and a group of 'OK', 'Cancel', and 'Help' buttons.

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**.

14.8 Authenticating the SMTP Server

The MegaRAID Storage Manager software supports a SMTP authentication mechanism called *Login*. This feature provides an extra level of security, while sending an email from the MegaRAID Storage Manager server.

To enter or modify the SMTP server authentication information, perform the following steps:

1. On the **Configure Alerts** window, click the **Mail Server** tab.
The Mail Server options appear, as shown in [Figure 14.8](#).
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**.

14.9 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 14.9 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.

14.10 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 [Figure 14.4](#).
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.

14.11 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 [Figure 14.4](#).
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.

14.12 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.

14.13 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.


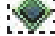
1. On the **Configure Alerts** window, click the **Alert Setting** tab, the **Mail Server** tab, or the **Email** tab.
2. Click **Load Backup**.
You are prompted to confirm your choice. The drive directory appears from which you can select a backup configuration to load.
3. Select the backup configuration 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.

14.14 Monitoring Server Events

The MegaRAID Storage Manager software enables you to monitor the activity of MegaRAID Storage Manager users in the network.

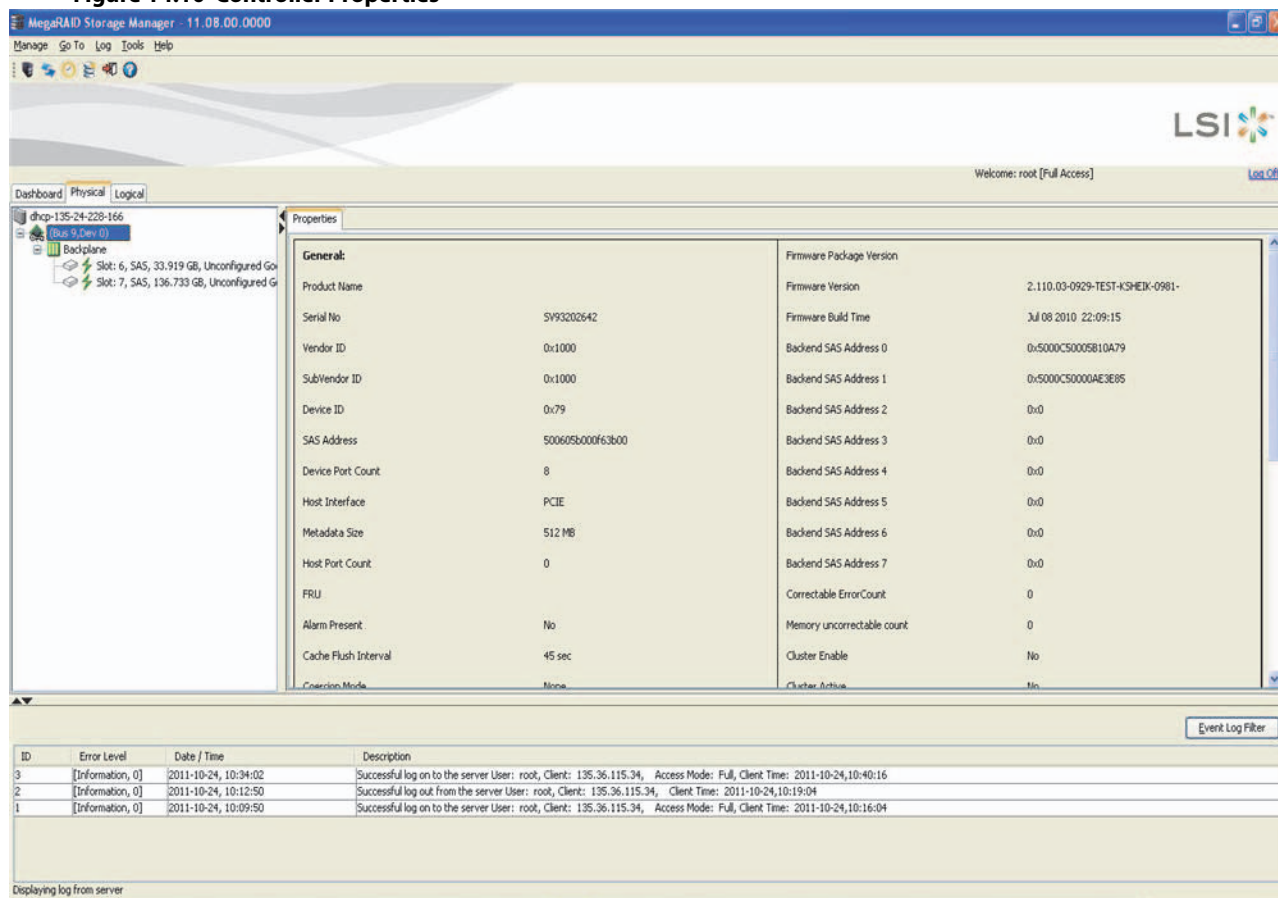
When a user logs on/logs off from the application, the event message appears in the log displayed at the bottom of the MegaRAID Storage Manager screen (the Vivaldi log/MegaRAID Storage Manager Log). These event message have a severity level, a date and timestamp (User log on / log off time), and a brief description that contains a user name, client IP address, an access mode (full/view only) and a client system time.

14.15 Monitoring Controllers

When the MegaRAID Storage Manager software is running, you can see the  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 14.10 Controller Properties




The Rebuild rate, Patrol read rate, Reconstruction rate, Consistency check rate, and BGI rate (background initialization) are all user selectable. For more information, see [Section 13.4, Changing Adjustable Task Rates](#).

The **BBU Present** field indicates whether a battery backup unit is installed.

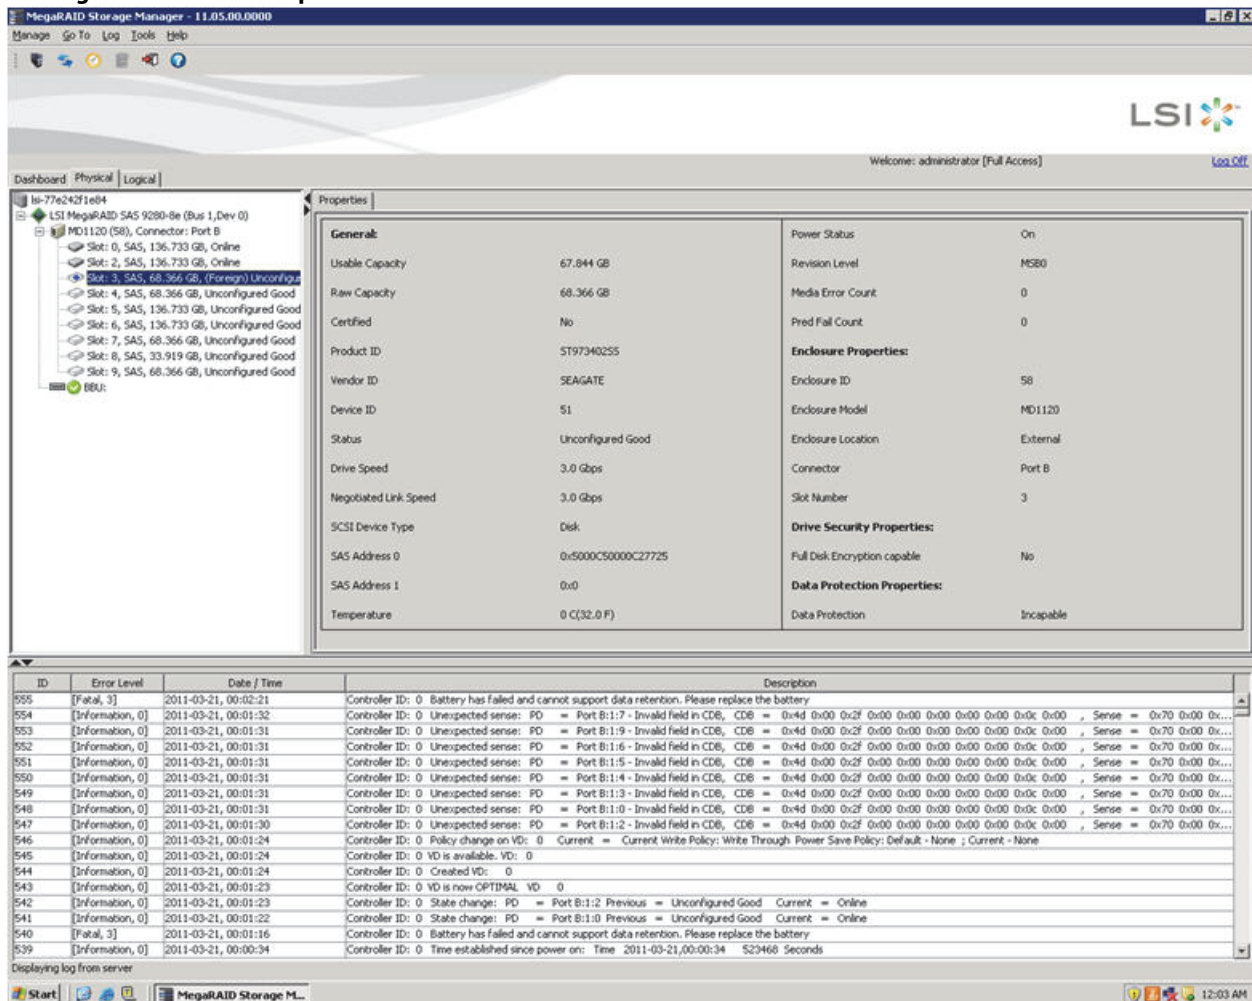
The **Alarm Enabled** field indicates whether the controller has an alarm to alert the user with an audible tone when there is an error or a problem on the controller. Options are available for disabling or silencing the alarm by right clicking on a controller icon or by selecting **Go To > Controller** menu.

14.16 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 14.11 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**.

14.17 Running a Patrol Read

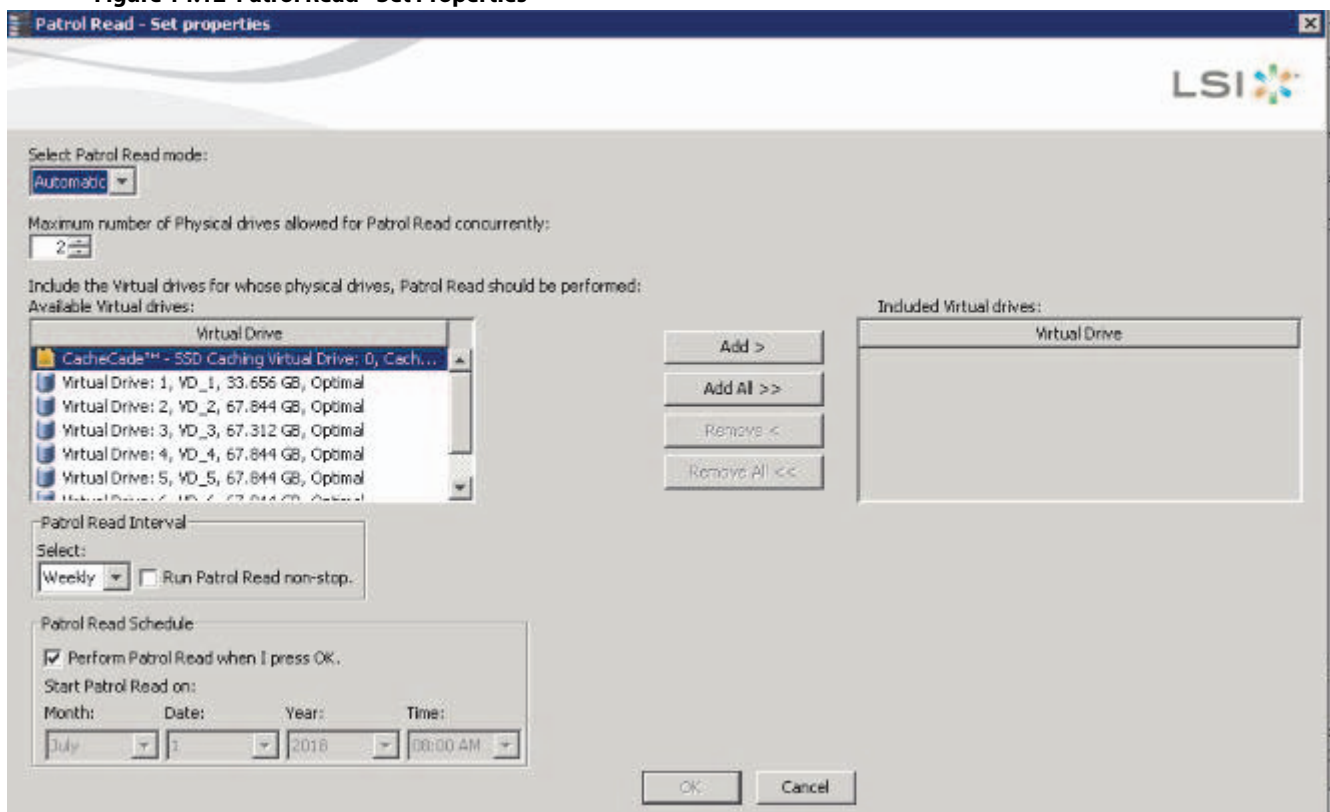
A patrol read periodically verifies all sectors of the drives connected to a controller, including the system reserved area in the RAID configured drives. You can run a patrol read for all RAID levels and for all hot spare drives. A patrol read is initiated only when the controller is idle for a defined 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 14.12 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.

14.17.1 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.

14.18 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 14.13 Virtual Drive Properties

The screenshot shows the MegaRAID Storage Manager interface. The left sidebar displays a tree view of the storage configuration, including the RAID controller, virtual drives, and physical drives. The main panel shows the properties for the selected virtual drive, VD_0. The properties are organized into several sections: General, IO Policy, Write Policy, Access Policy, and Power State Properties. The bottom of the interface shows a log window with system events.

General:		IO Policy	
RAID Level	1	IO Policy	Direct IO
Name	VD_0	Write Policy:	
Size	6,000 GB	Current Write Policy	Write Through
Mirror Data Size	6,000 GB	Default Write Policy	Write Back with BBU
Strip Size	64 KB	Reason for difference in Write Policy	BBU in re-learn cycle
Virtual Disk State	Optimal	Access Policy:	
IO and Cache Policies:		Current Access Policy	Read Write
Access Policy	Read Write	Default Access Policy	Read Write
Disk Cache Policy	Unchanged	Power State Properties:	
Read Policy	Always Read Ahead	Default Power save policy	None
		Current Power save policy	No power saving

ID	Error Level	Date / Time	Description
571	[Fatal, 3]	2011-03-21, 00:11:01	Controller ID: 0 Battery has failed and cannot support data retention. Please replace the battery
570	[Fatal, 3]	2011-03-21, 00:09:56	Controller ID: 0 Battery has failed and cannot support data retention. Please replace the battery
569	[Fatal, 3]	2011-03-21, 00:08:51	Controller ID: 0 Battery has failed and cannot support data retention. Please replace the battery
568	[Fatal, 3]	2011-03-21, 00:07:46	Controller ID: 0 Battery has failed and cannot support data retention. Please replace the battery
567	[Fatal, 3]	2011-03-21, 00:06:41	Controller ID: 0 Battery has failed and cannot support data retention. Please replace the battery
566	[Information, 0]	2011-03-21, 00:06:31	Controller ID: 0 Unexpected sense: PD = Port 8:1:7 - Failure prediction threshold exceeded, CDB = 0x03 0x00 0x00 0x00 0x40 0x00 , Sense = 0x70 0x00 0x00 0x...
565	[Warning, 1]	2011-03-21, 00:06:31	Controller ID: 0 PD Predictive failure: Port 8:1:7
564	[Information, 0]	2011-03-21, 00:06:31	Controller ID: 0 Unexpected sense: PD = Port 8:1:5 - Failure prediction threshold exceeded, CDB = 0x03 0x00 0x00 0x00 0x40 0x00 , Sense = 0x70 0x00 0x00 0x...
563	[Warning, 1]	2011-03-21, 00:06:31	Controller ID: 0 PD Predictive failure: Port 8:1:5
562	[Information, 0]	2011-03-21, 00:06:31	Controller ID: 0 Unexpected sense: PD = Port 8:1:4 - Failure prediction threshold exceeded, CDB = 0x03 0x00 0x00 0x00 0x40 0x00 , Sense = 0x70 0x00 0x00 0x...
561	[Warning, 1]	2011-03-21, 00:06:31	Controller ID: 0 PD Predictive failure: Port 8:1:4
560	[Information, 0]	2011-03-21, 00:06:30	Controller ID: 0 Unexpected sense: PD = Port 8:1:3 - Failure prediction threshold exceeded, CDB = 0x03 0x00 0x00 0x00 0x40 0x00 , Sense = 0x70 0x00 0x00 0x...
559	[Warning, 1]	2011-03-21, 00:06:30	Controller ID: 0 PD Predictive failure: Port 8:1:3
558	[Fatal, 3]	2011-03-21, 00:06:30	Controller ID: 0 Battery has failed and cannot support data retention. Please replace the battery
557	[Fatal, 3]	2011-03-21, 00:04:31	Controller ID: 0 Battery has failed and cannot support data retention. Please replace the battery
556	[Fatal, 3]	2011-03-21, 00:03:26	Controller ID: 0 Battery has failed and cannot support data retention. Please replace the battery
555	[Fatal, 3]	2011-03-21, 00:02:21	Controller ID: 0 Battery has failed and cannot support data retention. Please replace the battery

You can change the read policy, write policy, and other virtual drive properties. To change these properties, see [Section 13.7, 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.

14.19 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 14.14 Enclosure Properties

Vendor ID	DELL	FRU Number	41R3133
Enclosure ID	5	Part Number	CP-111-006-020
Enclosure Type	SES	Component Properties	
Enclosure Model	PM1000	Number of Temperature Sensors	4
Enclosure Location	External	Number of Fans	4
Firmware Version	A.04	Number of Power Supplies	2
Serial Number	0802V16NTE	Number of Voltage Sensors	0
Connector	Port A		
Number of Slots	15		

14.20 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.

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 14.15 Battery Properties

Type	BBU-09
Status	Optimal
Design Mode	48+ Hrs Retention with a Non-Transparent learn cycle and moderate service life
Temperature	Normal [21.0 C (69.8 F)]
Retention Time	48+ Hours
Charge	100 %
Charging Status	Charging
Advanced 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.

- Click **Advanced Properties** to view additional BBU properties
The **Advanced Properties** dialog appears.

Figure 14.16 Advanced Properties

Advanced Properties

LSI

Properties

Manufacturer	LSI101000G	Design Capacity	1350 mAh
Serial Number	1024	Full Capacity	n/a
Date of Manufacture	Thu, 01 Jan 0001 at 09:46:14	Remaining Capacity	n/a
Firmware Version	<value>	Expected Margin of Error	25 %
Status	Failed	Completed Discharge Cycles	63
	The battery has been failed. Please replace the battery pack.	Automatic Learn Mode	Enabled (Auto Learn Period: 30 Days)
Voltage	4035 mV	Next Learn Cycle Time	Fri, June 29, 2012 at 00:45:26
Current	0 mA		

Settings

Automatic Learn Mode: This option allows you to start a battery learn cycle automatically. You can either schedule a learn cycle or delay an existing scheduled learn cycle.

Next learn cycle time: Friday, June 29, 2012 At 12:45 AM

Delay next learn cycle by: day(s) hour(s) (Note: Please enter a value between 0 to 23 hours.)

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.

14.21 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.

14.21.1 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 Automatic Learn Cycle Properties** dialog appears, as shown in the following figure.

Figure 14.17 Set Automatic Learn Cycle Properties

4. Select the **Generate an event to remind me when to start a learn cycle manually** check box if you want an event to be generated to remind you to start a learn cycle manually.
5. Select **Enable** or **Disable** from the **Learn cycle** drop-down list to enable or disable an automatic learn cycle, respectively.

If you select **Disable**, the **Start on** and **Delay next learn cycle by** fields are disabled.

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.

6. In the **Start on** field, specify a day and time to start the automatic learn cycle.

7. 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.
8. If changes are made to the **Set Automatic Learn Cycle Properties** dialog, click **Apply** to refresh the dialog with the updated settings, without closing the dialog.
9. Click **OK** to save the settings and close the dialog.
If you selected **Disable** in the **Learn cycle** drop-down list, and click **OK** or **Apply**, a warning dialog appears asking for your confirmation to disable the automatic learn cycle.
If you selected the **Generate an event to remind me when to start a learn cycle manually** check box and click **OK** or **Apply**, an information dialog appears informing you about event generation.

14.21.2 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.

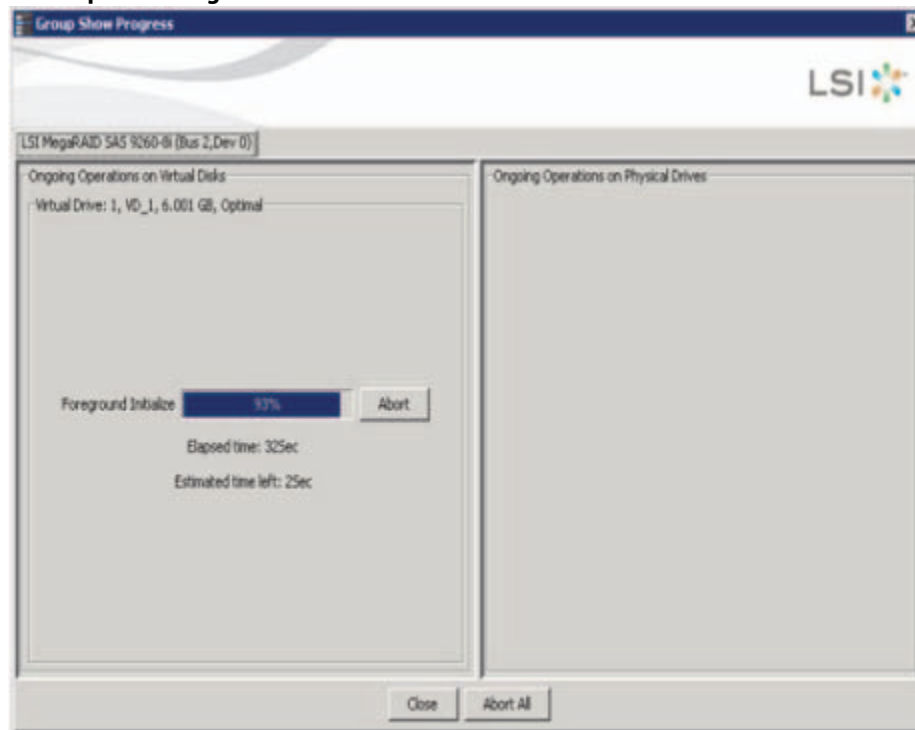
14.22 Monitoring Rebuilds and Other Processes

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 14.18 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 physical drives appear in the right panel. The type of operations that appear in this window are as follows:

- Background Initialization (BGI)
- PD Clear
- PD Erase
- Consistency Check
- LD Initialization
- LD Reconstruction
- Patrol Read
- Rebuild
- Replace
- LD Secure Erase
- LD disassociate

A Modify Drive Group process cannot be aborted. To abort any other ongoing process, click the **Abort** button next to the status indicator. Click **Abort All** to abort all ongoing processes. Click **Close** to close the window.

Chapter 15: Maintenance

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.

15.1 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.

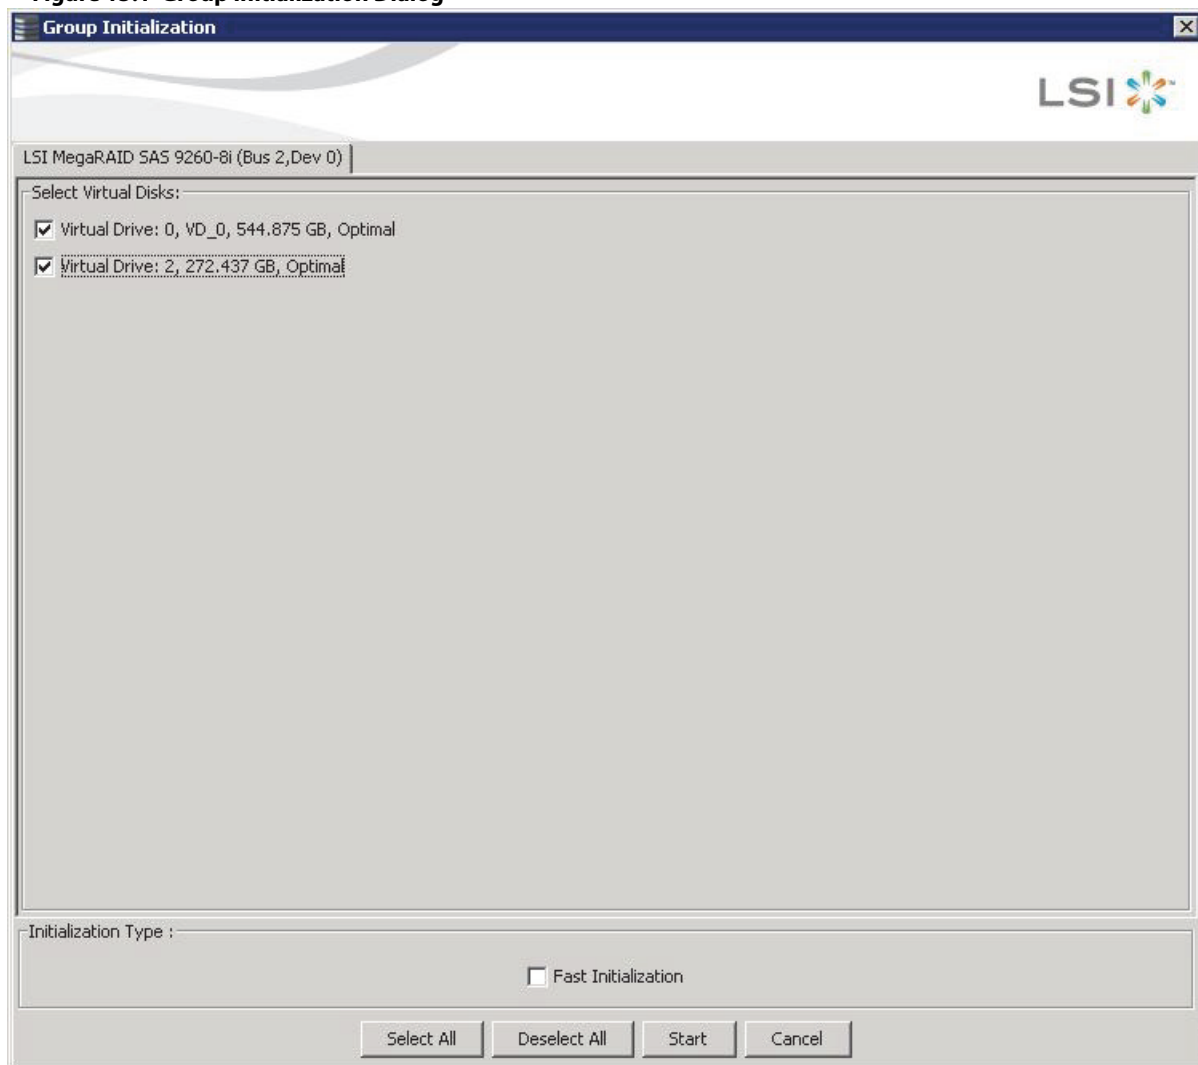
4. Select the **Fast Initialization** check box if you want to use this option.
If you leave the box unselected, the MegaRAID Storage Manager software runs a Full Initialization on the virtual drive. (For more information, see [Section 13.1.1, Selecting Virtual Drive Settings](#).)
5. Click **Start** to begin the initialization.
You can monitor the progress of the initialization. See [Section 14.22, Monitoring Rebuilds and Other Processes](#) for more information.

15.1.1 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 15.1 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 [Section 14.22, Monitoring Rebuilds and Other Processes](#) for more information.

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 does not provide data redundancy). A consistency check scans the virtual drive to determine whether the data has become corrupted and needs to be restored.

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



NOTE Make sure to back up the data before running a consistency check if you think the data might be corrupted.

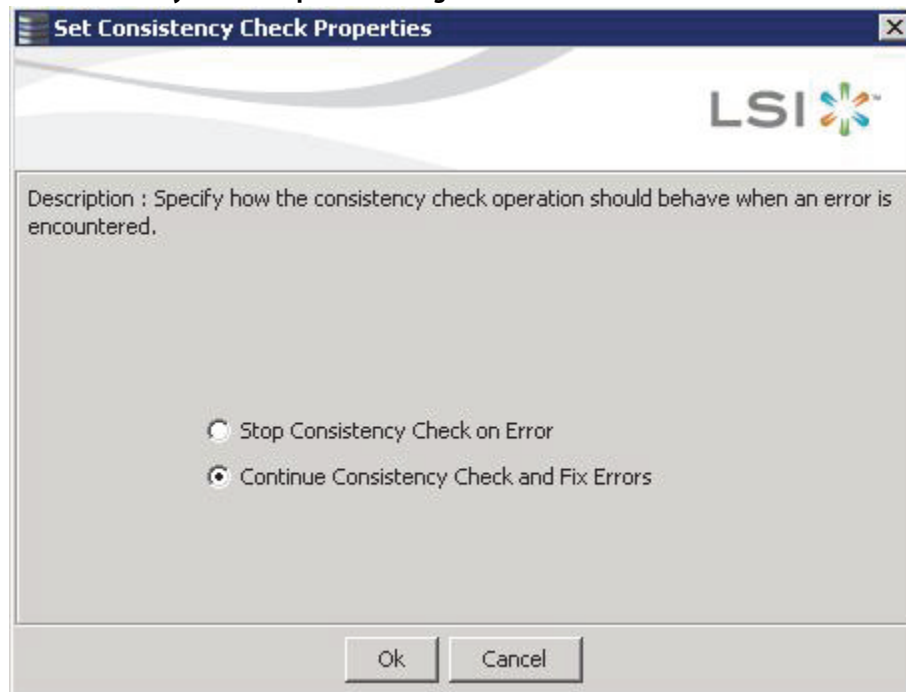
To run a consistency check, first set the consistency check properties, and then schedule the consistency check. This section explains how to set the properties, schedule the check, and run the consistency check.

15.2.1 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 15.2 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 an error, and then fixes the errors.
4. Click **Ok**.

15.2.2 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 15.3 Schedule Consistency Check Dialog

Schedule Consistency Check

LSI

Description : Establish schedule for consistency check operation.

Run consistency check:

Weekly

☐ Run consistency check continuously

Start on:

March 20 2010

Start time:

03:00 AM

Ok Cancel

3. Perform the following steps to schedule the consistency check:
 - a. Select how often to run the consistency check from the drop-down list.
You can click **Advanced** for more detailed date options.
 - b. (Optional) Select the **Run consistency check continuously** check box.
 - c. Select the month, day, and year on which to start the consistency check.
 - d. Select the time of day to start the consistency check.
4. Click **Ok**.
You can monitor the progress of the consistency check. See [Section 14.22, Monitoring Rebuilds and Other Processes](#) for more information.

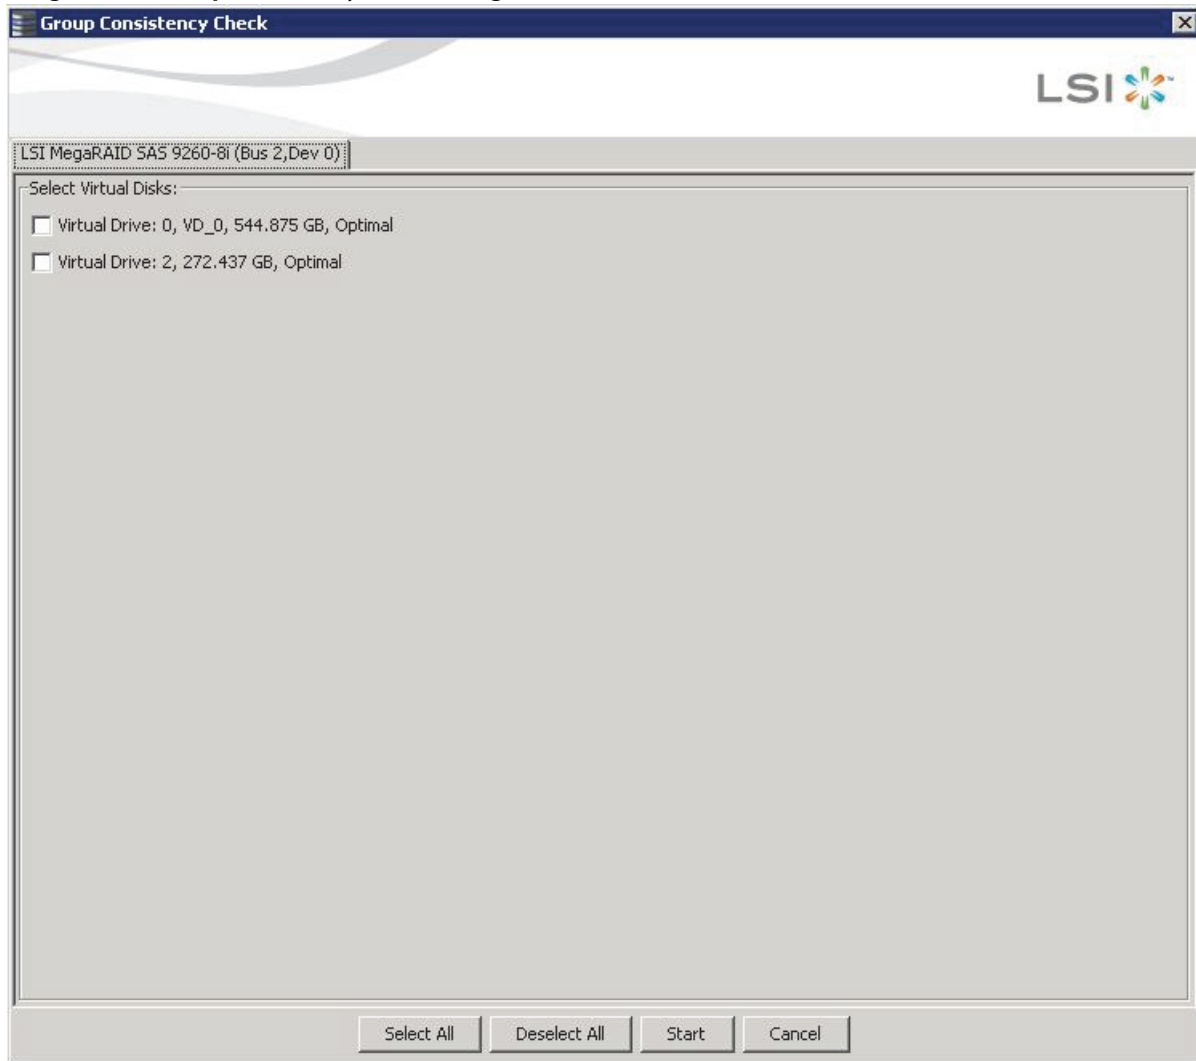
15.2.3 Running a Group Consistency Check

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 15.4 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 [Section 14.22, Monitoring Rebuilds and Other Processes](#) for more information.

15.3 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.

15.4 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.

15.4.1 New Drives Attached to a ServeRAID Controller

When you insert a new drive on a ServeRAID system and if the inserted drive does not contain valid DDF metadata, the drive displays as JBOD for ServeRAID entry-level controllers. 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.

15.5 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.

15.6 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**.

15.7 Upgrading Firmware

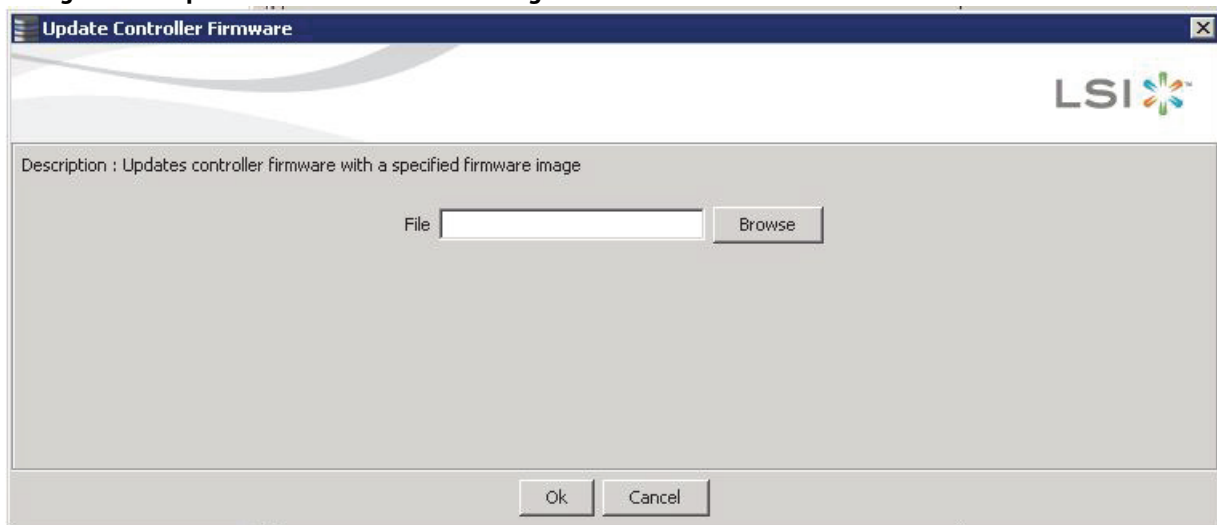
The MegaRAID Storage Manager software enables you to easily upgrade the controller firmware.

To avoid data loss because of dirty cache on the controller, the utility forces the virtual disks into Write Through mode after a firmware upgrade. It is in this mode until the server reboots. In Write Through mode, the controller sends a data transfer completion signal to the host when the disk subsystem has received all of the data in a transaction. This way, in case of a power outage, the controller does not discard the dirty cache.

Follow these steps to upgrade the firmware:

1. In the left panel of the **MegaRAID Storage Manager** window, click 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 15.5 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.

Chapter 16: CacheCade 2.0

CacheCade Pro 2.0 read/write software eliminates the need for manually configured hybrid arrays by intelligently and dynamically managing frequently-accessed data and copying it from HDD volumes to a higher performance layer of SSD cache. Copying the most accessed data to flash cache relieves the primary HDD array from time-consuming transactions, which allows for more efficient hard disk operation, reduced latency, and accelerated read and write speeds.

This feature provides significant improvements to overall system performance – two to twelve times that of HDD-only configurations – for a wide variety of server applications including web, file, online transaction processing (OLTP) database, data mining and other transaction-intensive applications.

16.1 Logical Drive Property Settings Required for CacheCade

For a logical drive to be valid to use as a CacheCade drive, the logical drive must be set to Write Back (WB) for write policy and Cached IO (CIO) for IO policy. The following screen shows the logical drive properties menu.

Figure 16.1 Virtual Drive Properties Menu

Set Virtual Drive Properties

LSI

Description : Defines virtual disk operation parameters

Name: RB_Source

Read Policy: No Read Ahead

Write Policy: Write Back

IO Policy : Cached IO

Access Policy: Read Write

Disk Cache Policy: Unchanged

Background Initialization: Disabled

Ok Cancel

16.2 Viewing a Logical Drive with CacheCade

If the logical drive properties are correctly set to use as the CacheCade logical drive (Write Back and Cached IO), then MSM correctly shows the associated logical drive when you click on the CacheCade drive.



NOTE If no logical drive exists with properties sufficient for the CacheCade virtual drive, this “Associated Virtual Drives” property does not appear. You can refresh the screen after you update the logical drive properties to make sure the properties are correctly updated.

16.3 WebBIOS Configuration for CacheCade

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



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

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

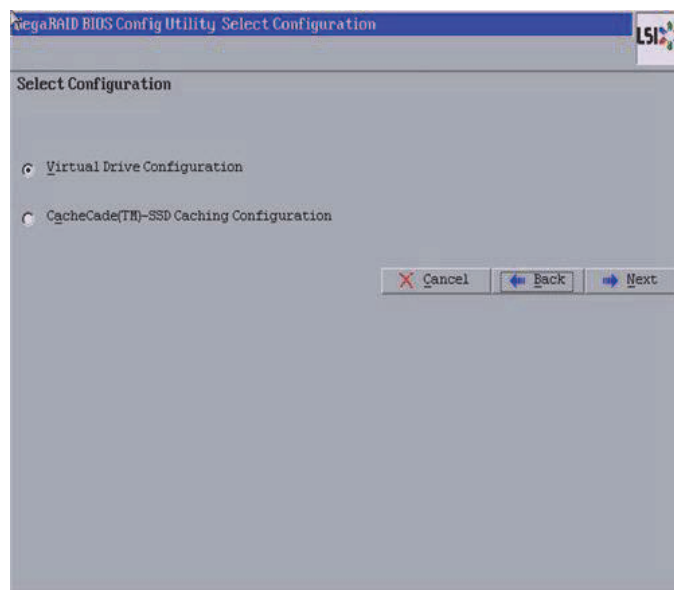
To support full-throughput for multiple direct-attached CacheCade software, this feature reduces I/O-processing overhead in the 2108-chip-based ServeRAID controllers. CacheCade offers performance equivalent to flash-based controllers and better performance for RAID 5 and RAID 6 when compared to Fusion I/O.

Follow these steps to create a CacheCade drive group.

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

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

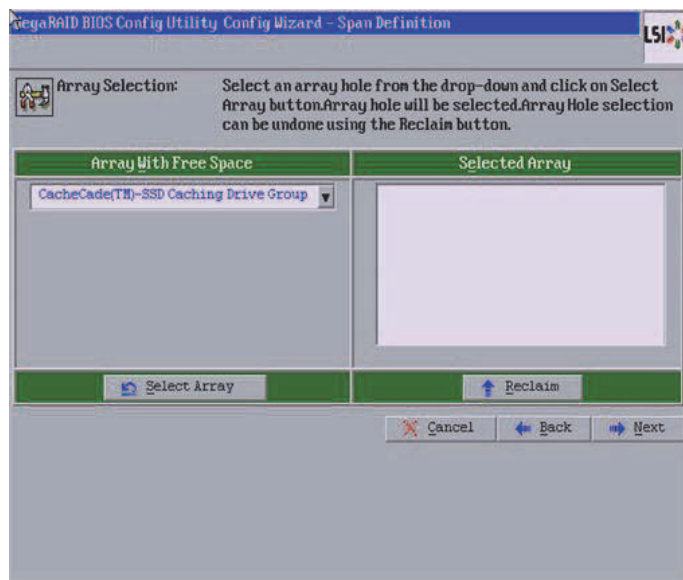
Figure 16.2 WebBIOS Configuration Wizard Screen



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


The Span Definition screen appears, as shown in the following figure.

Figure 16.3 CacheCade Array Selection Screen



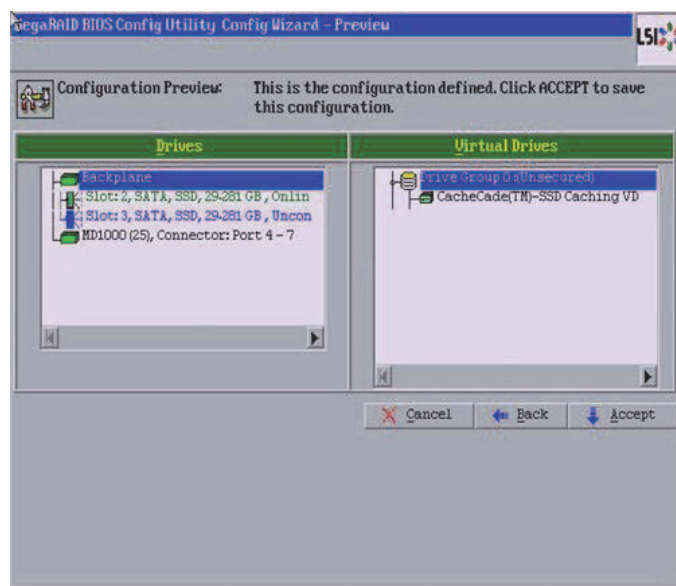
3. Select an array with free space from the drop-down list and click **Select Array**.
The selected array moves to the right frame under the heading **Selected Array**.
4. Click **Next**.
The Create CacheCade Disk screen appears, as shown in the following figure.

Figure 16.4 CacheCade Disk Screen



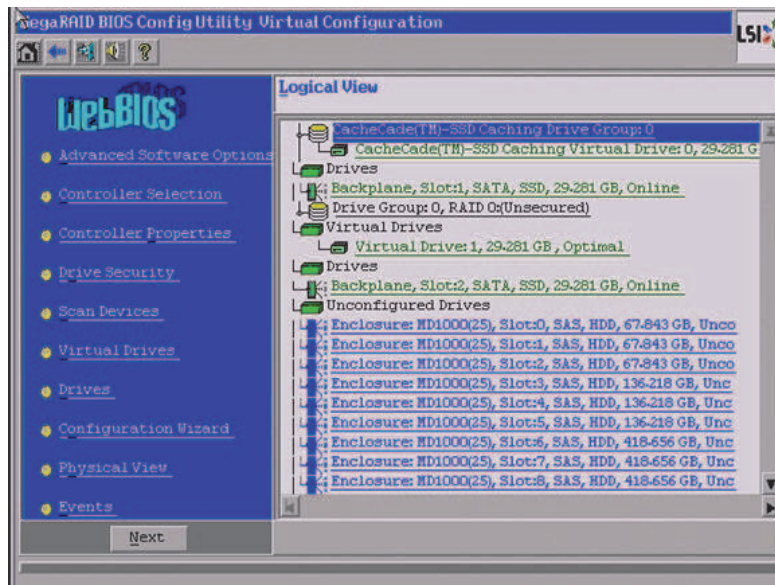
5. Click **Next** to accept the drive group.
If you need to undo the changes, click **Reclaim**.
The Config Wizard Preview screen appears, as shown in the following figure.

Figure 16.5 CacheCade Configuration Preview



6. Click **Accept** if the configuration is OK. Otherwise, or click **Back** to return to the previous screens and change the configuration.
7. If you accept the configuration, click **Yes** at the prompt to save the configuration.
The WebBIOS main menu screen appears, as shown in the following figure. It shows the CacheCade virtual drive.

Figure 16.6 WebBIOS Main Menu with a CacheCade Virtual Drive



16.4 MegaRAID Storage Manager Configuration for CacheCade

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



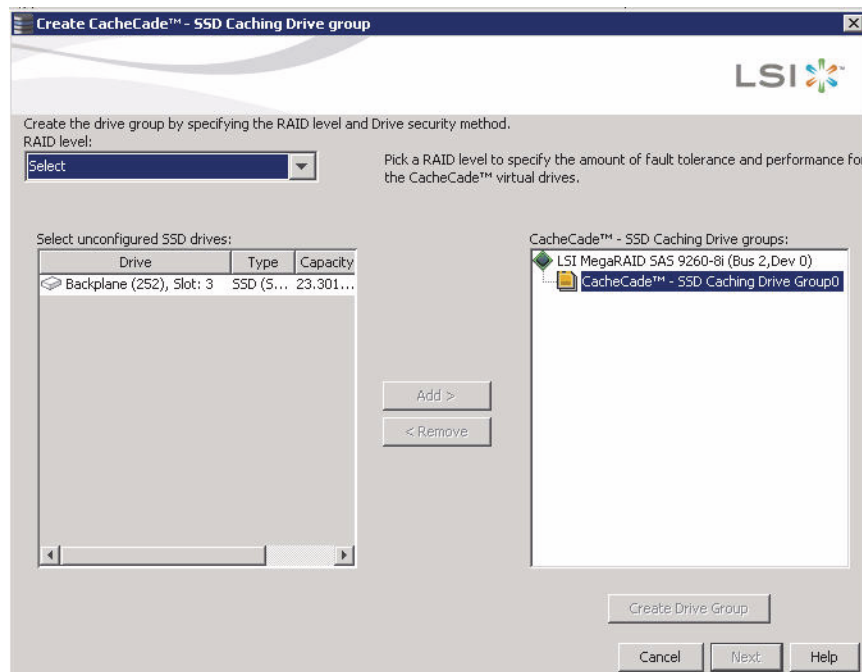
NOTE The ServeRAID firmware has the provision to monitor IO performance; changes have been made to accommodate the CacheCade Pro 2.0 software statistics. The CacheCade Pro 2.0 software metrics are captured for each logical drive that has CacheCade enabled. The CacheCade Pro 2.0 software gathers information about the cache windows allocated for a logical drive, the number of new windows allocated in this metrics collection period, the number of windows that are actively used, and the window hit rates.

Perform the following steps to use the CacheCade Pro 2.0 software:

1. Perform one of these actions:
 - Right-click 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 16.7 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.

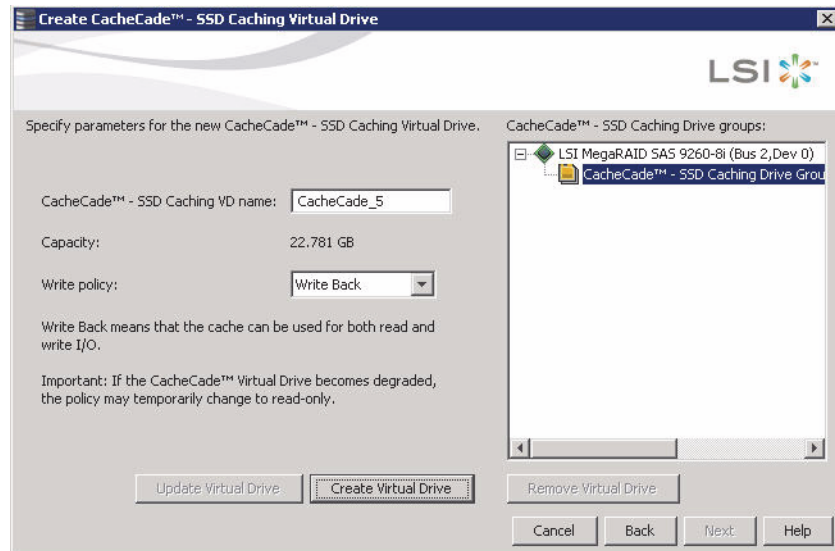
After you click **Add**, the Create Drive Group button is enabled.

5. Click Create Drive Group.

The newly created drive group appears in CacheCade SSD Caching Drive groups in the right frame.

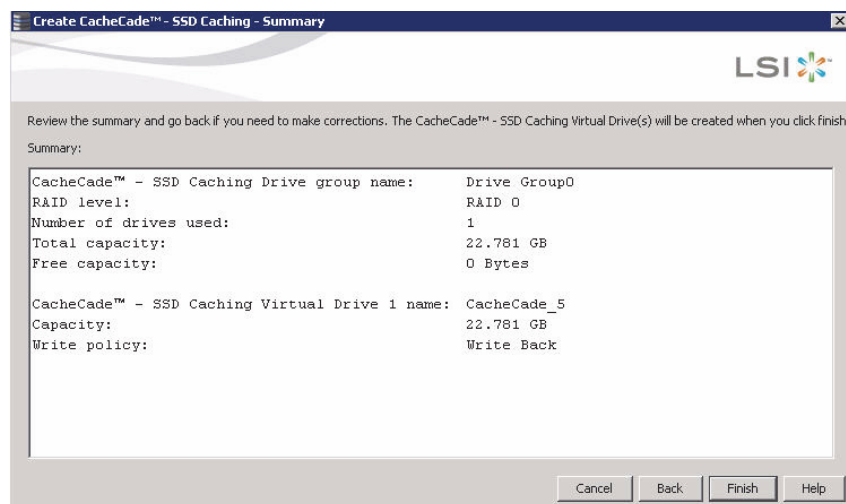
6. Click **Next**.
The next wizard screen appears.

Figure 16.8 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**.
10. 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.
11. Click **Next**.
The summary screen appears.

Figure 16.9 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.

12. 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.

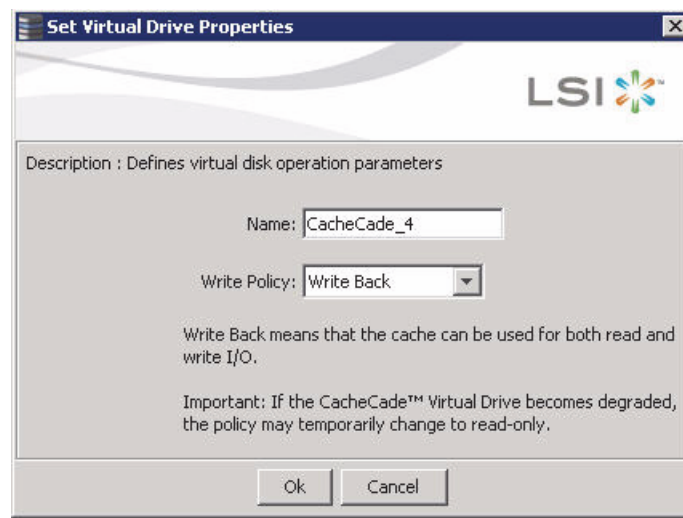
16.5 Modifying the CacheCade Virtual Drive Properties

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 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 16.10 Set Virtual Drive Properties Window



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**.

16.5.1 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 available only 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 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 16.11 Enable SSD Caching



2. Click **OK** to enable caching for that virtual drive.

16.5.2 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 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 16.12 Disable SSD Caching



2. Select the Confirm check box, and click OK to disable caching for that virtual drive.

16.5.3 Enabling or Disabling SSD Caching on Multiple Virtual Drives

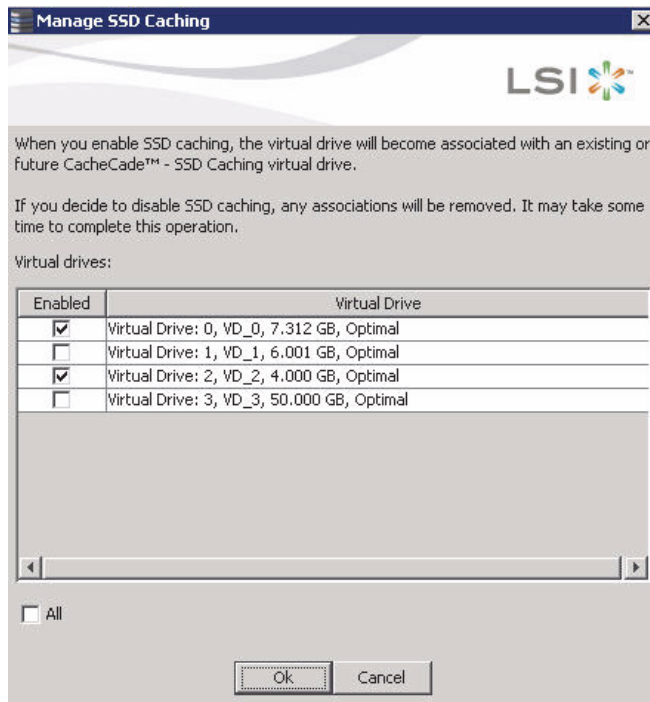
You can enable or disable SSD caching on multiple virtual drives at one go.

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 16.13 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 or disable SSD caching on the selected virtual drives.

16.5.4 Modifying a CacheCade Drive Group

To modify an existing CacheCade SSD caching drive group, you need to first delete the drive group and then create a new CacheCade drive group.

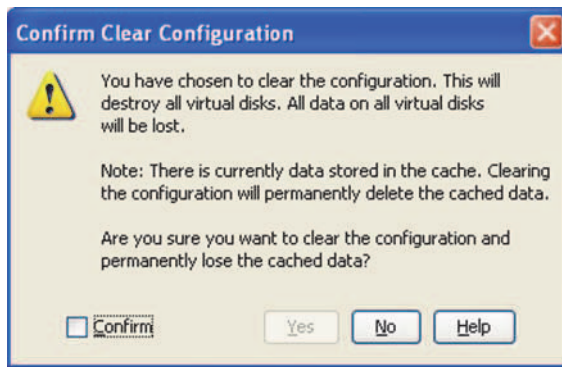
16.5.5 Clearing Configuration on CacheCade Pro 2.0 Virtual Drives

Perform the following steps to clear all existing configurations on a selected controller that has CacheCade Pro 2.0 virtual drives.

1. Perform one of these actions:
 - Right-click 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 16.14 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**.

16.5.6 Removing Blocked Access

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

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 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 16.15 Confirm Remove Blocked Access



2. Select the **Confirm** check box, and click **Yes**.

16.5.7 Deleting a Virtual Drive With SSD Caching Enabled

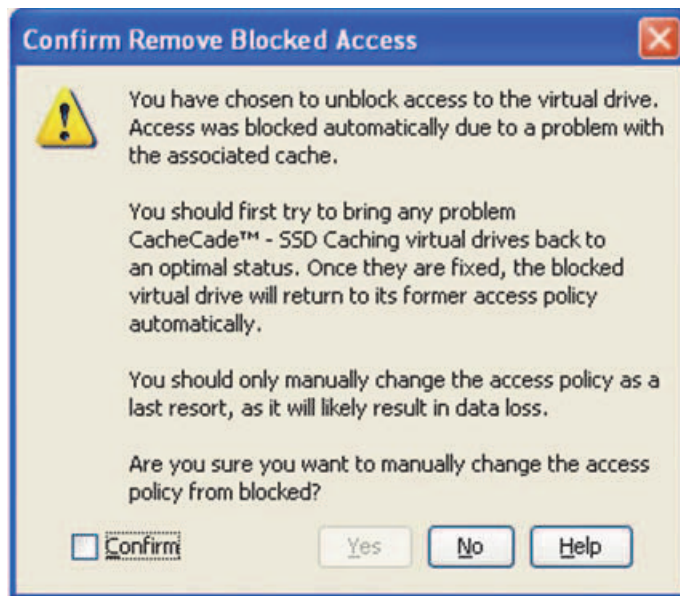
You can delete a virtual drive that has SSD caching enabled on it.

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 select **Go To >> Virtual Drive >> Delete Virtual Drive**.

The **Confirm Delete Virtual Disk** dialog appears, as shown in the following figure.

Figure 16.16 Confirm Delete Virtual Disk



2. Select the **Confirm** check box, and click **Yes**.



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

16.6 FastPath Advanced Software

Fast Path is a high-performance IO accelerator for the CacheCade software drive groups connected to a ServeRAID controller card. 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/s ServeRAID SATA+SAS controller.

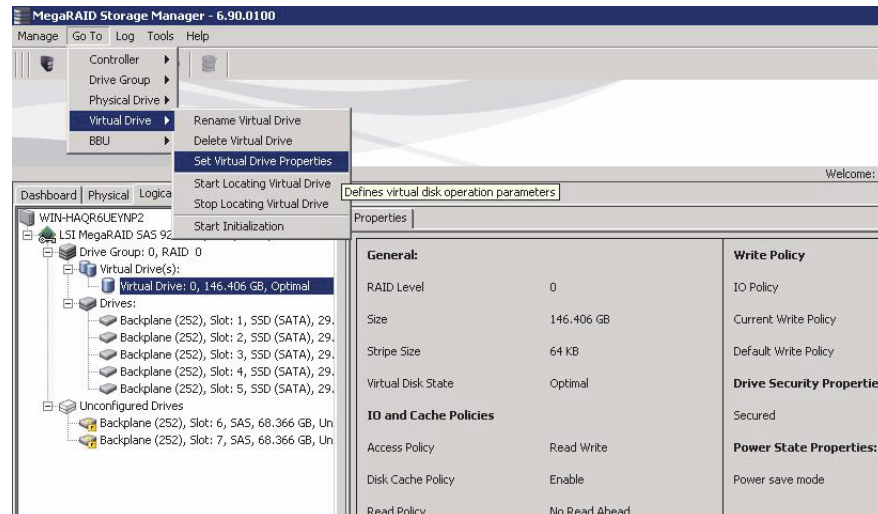
The Fast Path feature supports full optimization of the CacheCade software and hard disk drive (HDD) virtual disk groups to deliver an improvement in read and write IOPS that is three times greater than ServeRAID controllers not using FastPath technology. Also, the Fast Path advanced software is faster and more cost-effective than current flash-based adapter card solutions.

16.6.1 Setting Fast Path Options

Perform the following steps to use the FastPath advanced software:

1. Select the **Logical** tab on the MegaRAID Storage Manager main menu screen for the Logical view.
2. Select a virtual drive icon in the left frame.
3. Select **Virtual Drive >> Set Virtual Drive Properties** on the menu bar, as shown in the following figure.

Figure 16.17 Set Virtual Drive Properties Menu



The Set Virtual Drive Properties dialog appears and 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**.
The confirmation dialog appears.
 5. Select the **Confirm** check box, and click **Yes** to confirm that you want to set the virtual drive properties.

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 A.1 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 A.2 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	ServeRAID 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.

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 updatation 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.
0x0023	Information	Patrol Read complete	Logged when patrol read completes.

Number	Severity Level	Event Text	Generic Conditions when each event occurs
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, LBA 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.
0x003d	Critical	Consistency Check failed on %s	Logged when Consistency check failed as fatal error was found.

Number	Severity Level	Event Text	Generic Conditions when each event occurs
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 due to 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.
0x0056	Information	Drive Clear completed on %s	Logged when PD clear task is completed successfully on a PD.

Number	Severity Level	Event Text	Generic Conditions when each event occurs
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.
0x006f	Fatal	Unrecoverable medium error during recovery on %s at %lx	Logged when the recovery for a failed I/O encounters a medium error.

Number	Severity Level	Event Text	Generic Conditions when each event occurs
0x0070	Information	Drive removed: %s	Logged when a drive is removed from the controller.
0x0071	Warning	Unexpected sense: %s, CDB%, 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 removed 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 hot spare fails or is unconfigured by you.
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 restrictions) to cover certain arrays.
0x008a	Information	Created %s}	Logged as soon as the new logical drive created is added to the firmware configuration.

Number	Severity Level	Event Text	Generic Conditions when each event occurs
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.
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.

Number	Severity Level	Event Text	Generic Conditions when each event occurs
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.
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.

Number	Severity Level	Event Text	Generic Conditions when each event occurs
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.
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.

Number	Severity Level	Event Text	Generic Conditions when each event occurs
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.
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.

Number	Severity Level	Event Text	Generic Conditions when each event occurs
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.
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.

Number	Severity Level	Event Text	Generic Conditions when each event occurs
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.
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.

Number	Severity Level	Event Text	Generic Conditions when each event occurs
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.
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.

Number	Severity Level	Event Text	Generic Conditions when each event occurs
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.
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.

Number	Severity Level	Event Text	Generic Conditions when each event occurs
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.
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.

Number	Severity Level	Event Text	Generic Conditions when each event occurs
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.
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.

Number	Severity Level	Event Text	Generic Conditions when each event occurs
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.
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.

Number	Severity Level	Event Text	Generic Conditions when each event occurs
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.
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.

Number	Severity Level	Event Text	Generic Conditions when each event occurs
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 B.1 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 B.2 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= <i>filename_with_path</i> [force]	/cx download src= <i>filepath</i> [nosigchk] [noverchk]
Show the status of properties related to the controllers.	/cx show < <i>PropertyName</i> >	/cx show < <i>PropertyName</i> >
	The following properties can be used with this command:	The following properties can be used with this command:
	a0,1,2 -aALL	abortconerror
	achip	activityforlocate
	AENs [reverse]	
	alarms [reverse]	alarm
	allunitstatus	autorebuild
	autocarve	backplane
	autorebuild	batterywarning
	bios	bgirate
	carvesize	bootwithpinnedcache
	ctlbus diag	cachebypass
	dpmstat [type=<inst ra ext>]	cacheflushint
	driver	ccrate
	drivestatus	clusterenable
	events [reverse]	coercion
	exportjbod firmware	copyback

Description	3Ware CLI Command	StorCLI Command
	memory	directpdmapping
	model	ds
	monitor	eccbucketleakrate
	numdrives	eccbucketsize
	numports	enableeeghsp
	numunits	enableesmarter
	on degrade	enableeug
	pcb	exposeencldevice
	pchip	jbod
	phy	loadbalancemode
	rebuild	maintainpdfailhistory
	rebuildmodel	migraterate
	rebuildrate	ncq
	selftest	perfmode
	serial	pr
	spinup	prcorrectunconfiguredareas
	stagger	prrate
	unitstatus	rebuildrate
	verify	rehostinfo
	verifymode	restorehotspare
	verifystate	safeid
		smartpollinterval
		spinupdelay
		spinupdrivecount
		time
		usefdeonlyencrypt
Set properties on the selected controllers.	autocarve=<on off>	abortccconerror=<on off>
	autodetect=<on off> >disk=<p:-p> all	activityforlocate=<on off>
	autorebuild=<on off>	alarm=<on off>
	carvesize=<1024..32768>	autorebuild=<on off>
	dpmstat=<on off>	backplane=<value>
	on degrade=<cacheoff follow>	batterywarning=<on off>
	rebuild=<enable disable> <1..5>	bgirate=<value>
	rebuildmode=<adaptive lowlatency>	bootwithpinnedcache=<on off>
	rebuildrate=<1..5>	cachebypass=<on off>
	selftest=<enable disable>	flush flushcache
	spinup=<value>	cacheflushinterval=<value>

Description	3Ware CLI Command	StorCLI Command
	stagger=<value>	ccrate=<value>
	verify=advanced basic <1..5>	coercion=<value>
	verify=basic [pref=ddd:hh] where hh=(00...23 and ddd={mon tue wed thu fri sat sun}	clusterenable=<value>
	verify=enable disable <1..5>	copyback=<on off> type=<smartssd smarthdd all>
	verifymode=<adaptive lowlatency>	directpdmapping=<on off>
	verifystate=<1..5>	eccbucketleakrate=<value>
		eccbucketsize=<value>
		enableeeghsp=<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 B.3 Alarm Commands

Description	3Ware CLI Command	StorCLI Command
Set alarm properties.	<pre>/cx/ex/almx set alarm=<mute unmute off></pre> <p>NOTE The 3ware® controllers have enclosure alarms.</p>	<pre>/cx set alarm=<on off silence></pre> <p>NOTE The StorCLI controllers have controller alarms.</p>
Show alarm properties.	<pre>/cx/ex show alarms</pre> <p>NOTE This command applies for only 9750 and 9690SA controllers.</p>	<pre>/cx show alarm</pre>

B.4 Patrol Read and Consistency Check Commands

Table B.4 Patrol Read and Consistency Check Commands

Description	3Ware CLI Command	StorCLI Command
Show patrol read status and patrol read parameters, if any in progress.	<pre>/cx/ux show</pre>	<pre>/cx show patrolRead</pre>
Set the patrol read options on a single adapter, multiple adapters, or all adapters (x = single controller).	<pre>/cx/ux start verify /cx/ux set autoverify=<on off> /cx add verify=dddh:hh:duration</pre>	<pre>/cx set patrolread {=on mode=<auto manual>} {off} /cx set patrolread [starttime=<yyyy/mm/dd hh [max concurrentpd=<value>] [include ssds=<on off>] [uncfgareas=on off] /cx set patrolread delay=<value></pre>
Show consistency check status, if any in progress, and consistency check parameters.	<pre>/cx/ux show</pre>	<pre>/cx/vx show cc /cx show ccrate</pre>
Set consistency check options on a single adapter, multiple adapters, or all adapters (x = single controller).	<pre>/cx/ux start verify /cx/ux set autoverify=<on off> /cx add verify=ddd:hh:duration</pre>	<pre>storcli /cx set consistencycheck cc=[off seq c onc] [delay=value] [starttime=yyyy/ mm/dd hh] [excludevd=x-y,z]</pre>

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 B.5 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 B.6 Virtual Drive Commands

Description	3Ware CLI Command	StorCLI Command
Create a RAID volume of the specified RAID type.	<pre>/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) [nogppolicy] [storsave=<protect balance perform>] [noscan] [rapidrecovery=<all rebuild disable >] [group=<3 4 5 6 7 8 9 10 11 12 13 1 4 15 16>] RaidType={raid0, raid1, raid5, raid10, raid50, single, spare, raid6}</pre>	<pre>/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] [nora]*ra] [*direct cached] [cachedbadbbu]*nocachedbadbbu] [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]</pre>
Delete virtual drives.	<pre>/cx/ux del [quiet]</pre> <p>NOTE You can delete a single unit using this command.</p>	<pre>/cx/vx [all] delete [force] [cachecade]</pre> <p>NOTE You can delete one virtual disk, multiple virtual disks, or all the selected virtual disks on selected adapters using this command.</p>
Show drive group information.	<pre>/cx/ux show [all]</pre> <p>NOTE Information of each unit is shown individually.</p>	<pre>/cx/dall show [cachecade]</pre>
Scan and show available foreign configurations, provide a preview of the imported foreign configuration, show or import foreign configuration.	<pre>/cx rescan</pre>	<pre>cx/fall [all] show [preview] [securityKey=ssssssssss] cx/fall [all] import [securityKey=ssssssssss]</pre>
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.	<pre>/cx/ux show [all]</pre>	<pre>/cx/vx show all</pre>

Description	3Ware CLI Command	StorCLI Command
Show the virtual drive properties.	<pre> /cx/ux show autoverify /cx/ux show identify /cx/ux show ignoreECC /cx/ux show initializestatus /cx/ux show name /cx/ux show parity /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 wrccache </pre>	<pre> /cx/vx show all </pre> <p>NOTE The StorCLI does not have commands to show individual virtual drive properties.</p>
Set virtual drive properties.	<pre> /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 wrccache=on off [quiet] </pre>	<pre> /cx/vx set accesspolicy=<rw ro blocked rmvblk> /cx/vx set cachedbadbbu=<on off> /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 wrccache=<wt wb fwb> </pre>
Show cache and access policies of the virtual drive.	<pre> /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 parity /cx/ux show qpolicy /cx/ux show rapidrecovery /cx/ux show rdcache /cx/ux show rebuildstatus /cx/ux show serial /cx/ux show status initializestatus /cx/ux show storsave /cx/ux show verify status /cx/ux show volumes /cx/ux show wrccache </pre>	<pre> /cx/vx show all </pre> <p>NOTE The StorCLI does not have commands to show individual virtual drive properties.</p>

Description	3Ware CLI Command	StorCLI Command
Start initialization (writing 0s) on the virtual drive.	/cx/ux start verify NOTE Only the bios can do a foreground initialization. A background initialization does otherwise. A verify starts a back ground initialization.	/cx/vx start init [Full]
Stop an ongoing initialization on the virtual drive.	/cx/ux stop verify NOTE Only the bios can do a foreground initialization. A background initialization does otherwise. A verify starts a back ground initialization	/cx/vx stop init
Show a snapshot of the ongoing initialization, if any.	/cx/ux show [all] NOTE Only the bios can do a foreground initialization. A background initialization does otherwise. A verify starts a back ground initialization.	/cx/vx show init
Start a consistency check on the virtual drive.	/cx/ux start verify	/cx/vx start cc
Stop a consistency check on the virtual drive.	/cx/ux stop verify	/cx/vx stop cc
Reconstruct the selected virtual disk to a new RAID level.	/cx/ux 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 }	/cx/vx start migrate <type=raidlevel> [option=<add remove> disk=<e1:s1,e2:s2 ..>] /cx/vx show migrate
Change the power-saving setting on the virtual drive.	/cx/ux set powersavestandbytimer=<5 to 999>	/cx/vx set ds=<default Auto None Max MaxNoCache>

B.7 Physical Drive Commands

Table B.7 Physical Drive Commands

Description	3Ware CLI Command	StorCLI Command
Show physical disk information.	/cx/px show [all]	/cx[/ex]/sx show [all]
Start, stop, suspend, or resume an ongoing rebuild operation.	/cx/ux start rebuild disk=<p:-p..> [ignoreECC] NOTE Rebuilds cannot be stopped or paused.	/cx[/ex]/sx start rebuild /cx[/ex]/sx stop rebuild /cx[/ex]/sx pause rebuild /cx[/ex]/sx resume rebuild
Mark the configured physical disk drive as missing for the selected adapter.	/cx/px remove [quiet]	/cx[/ex]/sx set missing
Change the physical disk drive state to offline.	/cx/px remove [quiet]	/cx[/ex]/sx set offline
Add jbod.	/cx add vd type=jbod disk=<p> (where p = port or drive number)	/cx[/ex]/sx set jbod

Description	3Ware CLI Command	StorCLI Command
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)	/cx[/ex]/sx add hotsparedrive [{dgs=<N 0,1.2...n,,>] [EnclAffinity] [nonRevertible]]
Locate the physical disk drive and activate the physical disk activity LED.	/cx/p _x set identify=on off	/cx[/ex]/sx start stop locate
Prepare the unconfigured physical drive for removal.	/cx/p _x remove [quiet]	/cx[/ex]/sx 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.	/cx/p _x show [all]	/cx/eall/sall show [all]
Download drive or expander firmware.	/cx/p _x update fw=image.name [force]	/cx[/ex]/sx download src=filepath [satabridge]

B.8 Enclosure Commands

Table B.8 Enclosure Commands

Description	3Ware CLI Command	StorCLI Command
Show information about the enclosure for the selected adapter.	/cx/ex show [all]	/cx/ex show [all]
Show the status of the enclosure connected to the selected adapter.	/cx/ex show [all] /cx/ex show controllers /cx/ex show slots /cx/ex show fans /cx/ex show temp /cx/ex show pwrs /cx/ex show alms	/cx/ex show status
Download enclosure firmware.	/cx/ex update fw=image.name [force]	/cx/ex download src=filepath [offline] [forceActivate]

B.9 Events and Logs

Table B.9 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 B.10 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/ux ? tw_cli /cx/px ? tw_cli /cx/phyx ? tw_cli /cx/bbu ? tw_cli /cx/ex ? tw_cli /ex NOTE 3 Ware CLI shows context sensitive help.	show help
Show PHY connection information for physical PHY medium on the adapters.	/cx/phyx show	/cx/px show
Set PHY link speed.	/cx/phyx 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 C.1 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 C.2 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:
	abortconerror	activityforlocate
	alarmdsply	alarm
	autodetectbackplannedsbl	backplane
	autoenhancedimportdsply	batterywarning
	autosnapshotospace	bgirate
	batwarndsbl	bootwithpinnedcache
	bgirate	cachebypass
	bootwithpinnedcache	cacheflushint
	cachebypass	cc
	ccrate	ccrate
	clusterenable	clusterenable
	coercionmode	coercion
	copybackdsbl	copyback
	defaultldpspolicy	directpdmapping
	defaultsnapshotospace	ds
	defaultviewospace	eccbucketleakrate
	disableldpsinterval	eccbucketsize
	disableldpstime	enableeghsp
	disableocr	enableesmarter
	dsbl	enableeug
	eccbucketcount	exposeencldevice

Description	MegaCLI Command	StorCLI Command
	eccbucketleakrate	jbod
	eccbucketsize	loadbalancemode
	enableeeghsp	maintainpdfailhistory
	enableesmarter	migraterate
	enableeug	ncq
	enablejbod	perfmode
	enblspindownunconfigdrvs	pr
	loadbalancemode	prcorrectunconfiguredareas
	maintainpdfailhistoryenbl	prrate
	ncqdsply	rebuildrate
	patrolreadrate	rehostinfo
	perfmode	restorehotspare
	predfailpollinterval	safeid
	rebuildrate	smartpollinterval
	reconrate	spinupdelay
	rstrhotspareoninsert	spinupdrivecount
	smartcpybkenbl	time
	spindowntime	usefdeonlyencrypt
	spinupdelay	
	spinupencdrvnt	
	ssdsmartcpybkenbl	
	usediskactivityforlocate	
	usefdeonlyencrypt	
Set properties on the selected controllers.	Megacli -AdpSetProp <propertyname>-an -a0,1,2 -aa11	/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	activityforlocate=<on off>
	autodetectbackplanedsbl	alarm=<value>
	autoenhancedimportdsply	autorebuild=<on off>
	autosnapshotospace	backplane=<value>
	batwarndsbl	batterywarning=<on off>
	bgirate	bgirate=<value>
	bootwithpinnedcache	bootwithpinnedcache=<on off>
	cachebypass	cachebypass=<on off>
	ccrate	flush flushcache
	clusterenable	cacheflushinterval=<value>
	coercionmode	ccrate=<value>
	copybackdsbl	coercion=<value>
	defaultldpspolicy	clusterenable=<value>

Description	MegaCLI Command	StorCLI Command
	defaultsnapshotspace	copyback=<on off> type=<smartssd smarthdd all>
	defaultviewspace	dimmerswitch=<on off>
	disableldpsinterval	directpdmapping=<on off>
	disableldpstime	eccbucketleakrate=<value>
	disableocr	eccbucketsize=<value>
	dsbl	enableeghsp=<value>
	eccbucketcount	exposeencldevice=<on off>
	eccbucketleakrate	foreignautoimport=<on off>
	eccbucketsize	jbod=<on off>
	enableeghsp	loadbalancemode=<value>
	enableesmarter	maintainpdfailhistory=<on off>
	enableeug	migraterate=<value>
	enablejbod	ncq=<on off>
	enblspindownunconfigdrvs	perfmode=<value>
	loadbalancemode	prcorrectunconfiguredareas=<on off> >
	maintainpdfailhistoryenbl	prrate=<value>
	ncqdsply	rebuildrate=<value>
	patrolreadrate	restorehotspare=<on off>
	perfmode	smartpollinterval=<value>
	predfailpollinterval	spinupdelay=<value>
	rebuildrate	spinupdrivecount=<value>
	reconrate	stoponerror=<on off>
	rstrhotspareoninsert	usefdeonlyencrypt=<on off>
	smartcpybkenbl	time=yyyymmdd hh:mm:ss systemtime
	spindowntime	usefdeonlyencrypt=<on off>
	spinupdelay	
	spinupdrivecount	
	spinupencdelay	
	spinupencdrvnt	
	sdsmartcpybkenbl	
	usediskactivityforlocate	
	usefdeonlyencrypt	
Show the number of controllers connected.	MegaCLI -adpCount	storcli show ctrlcount
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 /cx show all
Show the freespace available in the controller.	MegaCLI -CfgFreeSpaceinfo -aN -a0,1,2 -aALL	storcli /cx show freespace

Description	MegaCLI Command	StorCLI Command
Download the controller firmware.	MegaCli -AdpFwFlash -f <i>filename</i> [-NoSigChk] [-NoVerChk] [-ResetNow] -aN -a0,1,2 -aALL	storcli /cx download file=<filepath> [fwtype=<val>] [nosigchk] [noverchk] [resetnow]
Show the preserved cache status.	MegaCLI -GetPreservedCacheList -aN -a0,1,2 -aALL	storcli /cx show preservedcache
Set the controller time	MegaCLI -AdpSetTime <i>yyyymmdd hh:mm:ss</i> -aN -a0,1,2 -aALL	storcli /c(x all) set time=<yyyymmdd hh:mm:ss systemtime>
Show the controller time.	MegaCLI -AdpGetTime -aN	storcli /cx show time

C.3 Patrol Read Commands

Table C.3 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/cx 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 /cx set patrolread {=on mode=<auto manual>} {off} storcli /cx set patrolread [starttime=<yyyy/mm/dd hh>] [maxconcurrentpd=<value>] [inclusssds=<on off>] [uncfgareas=on off] storcli /cx set patrolread delay=<value>
Disable patrol read.	MegaCli -AdpPR -Dsbl -aN -a0,1,2 -aALL	storcli /cx set patrolread=off
Enable automatic patrol read.	MegaCli -AdpPR -EnblAuto -aN -a0,1,2 -aALL	storcli /cx set patrolread=on mode=auto
Enable manual patrol read.	MegaCli -AdpPR -EnblMan -aN -a0,1,2 -aALL	storcli /cx set patrolread=on mode>manual
Start patrol read.	MegaCli -AdpPR -Start -aN -a0,1,2 -aALL	storcli /cx start patrolRead
Suspend a running patrol read.	MegaCli -AdpPR -Suspend -aN -a0,1,2 -aALL	storcli /cx suspend patrolRead
Resume a suspended patrol read.	MegaCli -AdpPR -Resume -aN -a0,1,2 -aALL	storcli /cx resume patrolRead
Stop a running patrol read.	MegaCli -AdpPR -Stop -aN -a0,1,2 -aALL	storcli /cx stop patrolRead
Include SSD drives in patrol read.	MegaCli -AdpPR -SSDPatrolReadEnbl -aN -a0,1,2 -aALL	storcli /cx set patrolRead includessds=on onlymixed

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 <i>yyyymmdd</i> <i>hh</i> -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 C.4 Consistency Check Commands

Description	MegaCLI Command	StorCLI Command
Schedule a consistency check.	MegaCLI -AdpCcSched -Dsbl -Info {-ModeConc -ModeSeq [-ExcludeLD -LN -L0,1,2] [-SetStartTime <i>yyyymmdd hh</i>] [-SetDelay <i>val</i>] } -aN -a0,1,2 -aALL	storcli /cx set consistencycheck cc=[off seq c onc] [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 C.5 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 C.6 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 C.7 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] [-szXXXXXXXX[-szYYYYYYYY [...]] [-strpszM] [-Hsp[E5:S5,...]] [-afterLdX] -aN	storcli /cx 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] [pdccache=on off *default][pi] [DimmerSwitch(ds)=default automatic(auto) *none maximum(max) MaximumWithoutCaching(maxnocache)] [wt *wb] [nora *ra] [*direct cached] [CachedBadBBU *NoCachedBadBBU] [strip=<8 16 32 64 128 256 512 1024] [AfterVd=X] [Spares=[e:]s [e:]s-x [e:]s-x,y] [force]
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 /cx add vd cachecade cc Type=[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] [-szXXXXXXXX[-szYYYYYYYY [...]] [-strpszM] [-afterLdX] -aN	storcli /cx 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] [pdccache=on off *default][pi] [DimmerSwitch(ds)=default automatic(auto) *none maximum(max) MaximumWithoutCaching(maxnocache)] [wt *wb] [nora *ra] [*direct cached] [CachedBadBBU *NoCachedBadBBU] [strip=<8 16 32 64 128 256 512 1024] [AfterVd=X] [Spares=[e:]s [e:]s-x [e:]s-x,y] [force]
Delete a virtual drive.	MegaCli -CfgClr [-Force] -aN -a0,1,2 -aALL	storcli /cx/vall delete
Show the topology information of the drive group.	MegaCLI -CfgDsply -aN -a0,1,2 -Aall	storcli /cx/dall show [all]
Show information for a CacheCade virtual drive.	MegaCLI -CfgCacheCadeDsply -aN -a0,1,2 -Aall	storcli /cx/dall show CacheCade(cc)

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 /cx/v/vx [all] delete -force
Delete a CacheCade virtual drive.	MegaCLI -CfgCacheCadeDel -LX -L0,2,5... -LALL -aN -a0,1,2 -Aall	storcli /cx/vx [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 /cx/f(x all) show [all] [securityKey=xxx] storcli /cx/f(x all) del delete [securityKey=xxx] storcli /cx/f(x all) import [preview] [securityKey=xxx] "

C.8 Security Commands

Table C.8 Security Commands

Description	MegaCLI Command	StorCLI Command
Set the key ID for the controller.	MegaCli -CreateSecurityKey -SecurityKey ssssssssss [-Passphrase ssssssssss] [-KeyID kkkkkkkkkk] -aN	storcli /cx set SecurityKey=XXXXXX [passphrase=yyy yy] [keyId=zzzz]
Change the security key for the controller.	MegaCli -ChangeSecurityKey -OldSecurityKey ssssssssss -SecurityKey ssssssssss [-Passphrase ssssssssss] [-keyID kkkkkkkkkk] -aN	storcli /cx set SecurityKey=XXXXXX OldSecurityKey= yyyyy
Compare and verify the security key for the controller.	MegaCli -VerifySecurityKey -SecurityKey ssssssssss -aN	storcli /cx compare SecurityKey=xxxxxx
Delete the security key.	MegaCLI -DestroySecurityKey [-Force] -aN	storcli /cx delete SecurityKey
Set the security key for the controller.	MegaCli -SetKeyID -KeyID kkkkkkkkkk -aN	storcli /cx set SecurityKey KeyId=xxxx

C.9 Virtual Drive Commands

Table C.9 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 /cx/v(x all) show storcli /cx/v(x all) show all
Set virtual drive properties.	MegaCli -LDSetProp WT WB NORA RA ADRA -Cached Direct CachedBad BBU NoCachedBadBBU} -RW RO Blocked {-Name nameString} -EnDskCache DisDskCache -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL	storcli /cx/v(x all) set wrcache=WT WB AWB storcli /cx/v(x all) set rdcache=RA NoRA storcli /cx/v(x all) set iopolicy=Cached Direct storcli /cx/v(x all) set accesspolicy=RW RO Blocked RmvBl kd storcli /cx/v(x all) set pdcache=On Off Default storcli /cx/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 /cx/v(x all) set ds=Default Auto None Max MaxNoCa che
Show virtual drive expansion information.	MegaCli -getLdExpansionInfo -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL	storcli /cx/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 /cx/v(x all) expand Size=<value> [expandarray]
Secure the virtual drive.	MegaCLI --LDMakeSecure -Lx -L0,1,2,... -Lall -An	storcli /cx/vx 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 /cx/vx show
Start virtual drive initialization.	MegaCli -LDInit -Start [Fast Full] -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL	storcli /cx/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 /cx/v(x all) stop init
Show the initialization progress.	MegaCli -LDInit -ShowProg -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL	storcli /cx/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 /cx/v(x all) start cc[Force]

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 /cx/v(x all) start cc storcli /cx/v(x all) stop cc storcli /cx/v(x all) pause cc storcli /cx/v(x all) resume cc storcli /cx/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 /cx/v(x all) set autobgi=On Off storcli /cx/v(x all) show autobgi storcli /cx/v(x all) stop bgi storcli /cx/v(x all) pause bgi storcli /cx/v(x all) resume bgi storcli /cx/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 /cx/vx start migrate type=raidx [option=add remove drives=[e:]s [e:]s-x [e:]s-x,y] [Force] storcli /cx/v(x all) show migrate
Delete preserved cache.	MegaCLI -DiscardPreservedCache -Lx -L0,1,2 -Lall -force -aN -a0,1,2 -aALL	storcli /cx/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 /cx/vx all set ssdCaching=on off

C.10 Physical Drive Commands

Table C.10 Physical Drive Commands

Description	MegaCLI Command	StorCLI Command
Show drive information.	MegaCli -pdInfo -PhysDrv[E0:S0,E1:S1,...] -aN -a0,1,2 -aALL	storcli /cx/ex/sx show storcli /cx/ex/sx 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 /cx/ex/sx start rebuild storcli /cx/ex/sx stop rebuild storcli /cx/ex/sx pause rebuild storcli /cx/ex/sx resume rebuild storcli /cx/ex/sx 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 /cx/ex/sx start copyback target = exx:sxx storcli /cx/ex/sx stop copyback storcli /cx/ex/sx pause copyback storcli /cx/ex/sx resume copyback storcli /cx/ex/sx show copyback

Description	MegaCLI Command	StorCLI Command
Mark a drive as missing.	MegaCli -PdMarkMissing -physdrv[E0:S0,E1:S1,...] -aN -a0,1,2 -aALL	storcli /cx/ex/sx set missing
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.

Description	MegaCLI Command	StorCLI Command
Flash the physical drive firmware.	MegaCLI PdFwDownload[offline] [ForceActivate] {[-SataBridge] -PhysDrv[0:1]} {-EncdevId[d evId1]} -f <filename> -aN -a0,1,2 -aAll	storcli /cx[/ex]/sx download src=<filepath> [satabridge] 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]
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[erasepatte rnB=<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 C.11 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 C.12 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 C.13 Alarm Commands

Description	MegaCLI Command	StorCLI Command
Show alarm properties.	MegaCli -AdpGetProp AlarmDsply -aN -a0,1,2 -aALL	storcli /cx(x all) show alarm
Set alarm properties.	MegaCli -AdpSetProp AlarmEnbl AlarmDsbl AlarmSilence -aN -a0,1,2 -aALL	storcli /cx(x all) set alarm=<on off silence>

C.14 Event Log Properties Commands

Table C.14 Event Log Properties Commands

Description	MegaCLI Command	StorCLI Command
Show event logs.	MegaCli -AdpEventLog -GetEventLogInfo -aN -a0,1,2 -aALL	storcli /cx show eventloginfo
Show the specified type of event logs.	MegaCli -AdpEventLog -GetEvents {-info -warning -critical -fatal} {-f <fileName>} -aN -a0,1,2 -aALL	storcli /cx 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 /cx show events [type=[latest=x ccincon vd=[sincereboot sinceshutdown inc ludedeleted latest ccincon]] [filter=[info warning critical fat al]] file=xyz.txt
Delete the event logs.	MegaCli -AdpEventLog -Clear -aN -a0,1,2 -aALL	storcli /cx delete events

C.15 Premium Feature Key Commands

Table C.15 Premium Feature Key Commands

Description	MegaCLI Command	StorCLI Command
Show the Safe ID of the controller.	MegaCli -ELF -GetSafeId -a0	storcli /cx(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 /cx(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 /cx(x all) set aso key=<key value> preview
Apply the Activation Key.	MegaCli -ELF -Applykey key -val -a0	storcli /cx(x all) set aso key=<key value>
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 D.1 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 /cx(x all) set ds=OFF type=1 2 3 4
	storcli /cx(x all) set ds=ON type=1 2 [properties]
	storcli /cx(x all) set ds=ON type=3 4 DefaultIdType=<val> [properties]
	storcli /cx(x all) set ds [properties]
	storcli /cx/v(x all) set ds=Default Auto None Max MaxNoCache
Security	storcli /cx delete security key
	storcli /cx set securitykey=xxxxxxxx {passphrase=xxxx} {keyid=xxx}
	storcli /cx set securitykey keyid=xxx
	storcli /cx compare securitykey=xxxxxxxx
	storcli /cx set securitykey=xxxxxxxx oldsecuritykey=xxxxxxxx
ASO	storcli /cx(x all) set aso key=<keyvalue> preview
	storcli /cx(x all) set aso key=<key value>
	storcli /cx(x all) set aso transfertovault
	storcli /cx(x all) set aso rehostcomplete
	storcli /cx(x all) set aso deactivatetrialkey
	storcli /cx(x all) show safeid
	storcli /cx(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 /cx download src=<filepath> [forceActivate]
Copy back	storcli /cx[/ex]/sx show copyback
	storcli /cx[/ex]/sx start copyback target=eID:sID
	storcli /cx[/ex]/sx stop copyback
	storcli /cx[/ex]/sx pause copyback
	storcli /cx[/ex]/sx resume copyback
Migrate	storcli /cx/v(x all) show migrate
	storcli /cx/vx start migrate type=raidx [option=add remove drives=[e:]s [e:]s-x [e:]s-x,y] [Force]
Cache	storcli /cx/v(x all) set ssdcaching=on off
	storcli /cx(x all) show preservedcache
	storcli /cx/v(x all) delete preservedcache[force]

Command Group	Command
BBU	storcli /cx/bbu show
	storcli /cx/bbu show all
	storcli /cx/bbu set [learnDelayInterval=<val> bbuMode=<val>
	storcli /cx/bbu start learn
Secure erase	storcli /cx/sx secureerase [force]
	storcli /cx/sx start erase [simple normal thorough][erasepatternA=<val>]
	storcli /cx/sx stop erase
	storcli /cx/sx show erase
Consistency check	storcli /cx show cc/ConsistencyCheck
Controller	storcli /cx show cc