

RAID Controller

User's Manual





© 2014 GeoVision, Inc. All rights reserved.

Under the copyright laws, this manual may not be copied, in whole or in part, without the written consent of GeoVision.

Every effort has been made to ensure that the information in this manual is accurate. GeoVision, Inc. makes no expressed or implied warranty of any kind and assumes no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages arising from the use of the information or products contained herein. Features and specifications are subject to change without notice.

GeoVision, Inc.
9F, No. 246, Sec. 1, Neihu Rd.,
Neihu District, Taipei, Taiwan
Tel: +886-2-8797-8377
Fax: +886-2-8797-8335
<http://www.geovision.com.tw>

Trademarks used in this manual: *GeoVision*, the *GeoVision* logo and *GV* series products are trademarks of GeoVision, Inc. *Windows* and *Windows XP* are registered trademarks of Microsoft Corporation.

March 2014

Contents

Chapter 1 Introduction to RAID	1
1.1 Components and Features	2
1.1.1 Drive Group	2
1.1.2 Virtual Drive	2
1.1.3 Fault Tolerance	2
1.1.4 Consistency Check	3
1.1.5 Replace	3
1.1.6 Background Initialization	4
1.1.7 Patrol Read	5
1.1.8 Disk Striping	5
1.1.9 Disk Mirroring	6
1.1.10 Parity	6
1.1.11 Disk Spanning	7
1.1.12 Hot Spares	8
1.1.13 Disk Rebuilds	10
1.1.14 Rebuild Rate	11
1.1.15 Hot Swap	11
1.1.16 Drive States	11
1.1.17 Virtual Drive States	12
1.1.18 Beep Codes	12
1.1.19 Enclosure Management	12
1.2 RAID Levels	13
1.2.1 Summary of RAID Levels	13
1.2.2 Selecting a RAID Level	14
1.2.3 RAID 0	14
1.2.4 RAID 1	15
1.2.5 RAID 5	16

1.2.6	RAID 6	17
1.2.7	RAID 00	18
1.2.8	RAID 10	18
1.2.9	RAID 50	19
1.2.10	RAID 60	21
1.3	RAID Configuration Strategies	22
1.3.1	Maximizing Fault Tolerance	22
1.3.2	Maximizing Performance	23
1.3.3	Maximizing Storage Capacity	24
1.4	RAID Availability	25
1.4.1	RAID Availability Concept	25
1.5	Configuration Planning	25
1.6	Number of Drives	26

Chapter 2 WebBIOS Configuration Utility 27

2.1	Overview	27
2.2	Starting the WebBIOS Configuration Utility	27
2.3	WebBIOS Configuration Utility Main Dialog Options	28
2.4	Creating a Storage Configuration	31
2.4.1	Using Automatic Configuration	33
2.4.2	Using Manual Configuration	33
2.4.2.1	Virtual Drive Options	33
2.4.2.2	Using Manual Configuration: RAID 0	36
2.4.2.3	Using Manual Configuration: RAID 1	39
2.4.2.4	Using Manual Configuration: RAID 5	41
2.4.2.5	Using Manual Configuration: RAID 6	43
2.4.2.6	Using Manual Configuration: RAID 00	44
2.4.2.7	Using Manual Configuration: RAID 10	48

2.4.2.8	Using Manual Configuration: RAID 50.	51
2.4.2.9	Using Manual Configuration: RAID 60.	54
2.5	Viewing and Changing Device Properties.	57
2.5.1	Viewing Controller Properties.	60
2.5.1.1	Controller Information Menu Options.	60
2.5.2	Viewing Virtual Drive Properties, Policies, and Operations	63
2.5.3	Viewing Drive Properties	64
2.5.4	Shield State.	66
2.5.4.1	Shield State Physical View.	66
2.5.4.2	Logical View Shield State.	67
2.5.4.3	Viewing the Physical Drive Properties of a Drive in Shield State.	67
2.5.4.4	Viewing if Shield State Is Enabled in a Controller.	68
2.5.5	SSD Disk Cache Policy	69
2.5.5.1	Viewing Cache Properties.	69
2.5.6	Emergency Spare.	69
2.5.6.1	Emergency Spare for Physical Drives.	69
2.5.7	Emergency Spare for Controllers	70
2.5.7.1	Setting Controller Emergency Spare Properties.	70
2.5.7.2	Spare Properties Viewing Controller Emergency Spare Properties.	70
2.5.7.3	Spare Properties Viewing Controller Emergency Spare Properties.	71
2.6	Viewing and Expanding a Virtual Drive.	72
2.7	Recovering and Clearing Punctured Block Entries	73
2.8	Suspending and Resuming Virtual Drive Operations.	73
2.9	Non-SED Secure Erase	75
2.9.1	Erasing a Non-SED Physical Drive	75
2.9.1.1	Drive Erase Progress.	77
2.9.2	Virtual Drive Erase.	77
2.9.2.1	Group Show Progress for Virtual Drive Erase.	79

2.10	Viewing System Event Information	80
2.11	Managing Configurations	81
2.11.1	Running a Consistency Check	81
2.11.2	Deleting a Virtual Drive	81
2.11.3	Importing or Clearing a Foreign Configuration	82
2.11.3.1	Group Show Progress for Virtual Drive Erase.....	84
2.11.3.2	Importing Foreign Configurations from Integrated RAID to MegaRAID	85
2.11.3.3	Troubleshooting Information	85
2.11.4	Importing Foreign Configurations	85
2.11.5	Migrating the RAID Level of a Virtual Drive	86
2.11.5.1	Additional Drives Required for RAID-Level Migration.....	86
2.11.5.2	Migrating the RAID Level	86
2.11.6	New Drives Attached to a MegaRAID Controller	87
Chapter 3: MegaRAID Storage Manager Overview		88
3.1	Overview	88
3.1.1	Creating Storage Configurations	88
3.1.2	Monitoring Storage Devices	88
3.1.3	Maintaining Storage Configurations	89
3.2	Prerequisites to Running MegaRAID Storage Manager Remote Administration	89
Chapter 4: MegaRAID Storage Manager Window and Menus		90
4.1	Starting the MegaRAID Storage Manager Software	90
4.2	Discovery and Login	91
4.3	High Availability DAS Support	95
4.4	LDAP Support	96
4.5	Configuring LDAP Support Settings	99
4.6	MegaRAID Storage Manager Main Menu	100
4.6.1	Dashboard / Physical View/ Logical View	100

4.6.2 Physical Drive Temperature	105
4.6.3 Shield State	105
4.6.4 Shield State Physical View	106
4.6.5 Logical View Shield State	106
4.6.6 Viewing the Physical Drive Properties	107
4.6.7 Viewing Server Profile of a Drive in Shield State	108
4.6.8 Displaying the Virtual Drive Properties	109
4.6.8.1 Parity Size.....	109
4.6.8.2 Minor Data Size.....	110
4.6.8.3 Metadata Size.....	111
4.6.9 Emergency Spare	112
4.6.9.1 Emergency Spare for Physical Drives.....	112
4.6.9.2 Emergency Spare Property for Controllers.....	113
4.6.9.3 Commissioned Hotspare.....	114
4.6.10 SSD Disk Cache Policy	115
4.6.10.1 Virtual Drive Settings.....	116
4.6.10.2 Set Virtual Drive Properties.....	117
4.6.11 Non-SED Secure Erase Support	118
4.6.11.1 Group Show Progress for Drive Erase.....	119
4.6.11.2 Virtual Drive Erase.....	120
4.6.11.3 Group Show Progress for Virtual Drive Erase.....	123
4.6.12 Rebuild Write Cache	123
4.6.13 Background Suspend / Resume Support	124
4.6.14 Enclosure Properties	126
4.7 GUI Elements in the MegaRAID Storage Manager Window and Menus	127
4.7.1 Device Icons	127
4.7.2 Properties and Graphical View Tabs	128
4.7.3 Event Log Panel	129
4.7.4 Menu Bar	129

Chapter 5: Configuration	131
5.1 Creating a New Configuration	131
5.1.1 Selecting Virtual Drive Settings	131
5.1.2 Optimum Controller Settings for Fast Path	133
5.1.3 Creating a Virtual Drive Using Simple Configuration	133
5.1.4 Creating a Virtual Drive Using Advanced Configuration	138
5.2 Creating Hot Spare Drives	144
5.3 Changing Adjustable Task Rates	144
5.4 Changing Power Settings	146
5.5 Recovering and Clearing Punctured Block Entries	148
5.6 Changing Virtual Drive Properties	148
5.7 Changing a Virtual Drive Configuration	150
5.7.1 Removing a Drive from a Configuration	150
5.7.2 Replacing a Drive	151
5.8 Deleting a Virtual Drive	152
Chapter 6: Monitoring Controllers and Their Attached Devices	153
6.1 Alert Delivery Methods	153
6.1.1 Vivaldi Log/MegaRAID Storage Manager Log	154
6.1.2 System Log	155
6.1.3 Pop-up Notification	156
6.1.4 Email Notification	156
6.2 Configuring Alert Notifications	157
6.3 Editing Alert Delivery Methods	160
6.4 Changing Alert Delivery Methods for Individual Events	161
6.5 Changing the Severity Level for Individual Events	162
6.6 Roll Back to Default Individual Event Configuration	163
6.7 Entering or Editing the Sender Email Address and SMTP Server	164
6.8 Authenticating the SMTP Server	165
6.9 Adding Email Addresses of Recipients of Alert Notifications	166
6.10 Testing Email Addresses of Recipients of Alert Notifications	166
6.11 Removing Email Addresses of Recipients of Alert Notifications	167
6.12 Saving Backup Configurations	167
6.13 Loading Backup Configurations	168

6.14 Monitoring Server Events	168
6.15 Monitoring Controllers	169
6.16 Monitoring Drives	171
6.17 Running a Patrol Read	172
6.17.1 Patrol Read Task Rates	174
6.18 Monitoring Virtual Drives	174
6.19 Monitoring Rebuilds and Other Processes	176
Chapter 7: Maintaining and Managing Storage Configurations	178
7.1 Initializing a Virtual Drive	178
7.1.1 Running a Group Initialization	179
7.2 Running a Consistency Check	180
7.2.1 Setting the Consistency Check Settings	180
7.2.2 Scheduling a Consistency Check	181
7.2.3 Running a Group Consistency Check	182
7.3 Scanning for New Drives	183
7.4 Rebuilding a Drive	184
7.5 Making a Drive Offline or Missing	185
7.6 Removing a Drive	186
7.7 Upgrading Firmware	186
Appendix A: Events and Messages	188
A.1 Error Levels	188
A.2 Event Messages	188
Appendix B: Glossary	207

Chapter 1 Introduction to RAID

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

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.

RAID Benefits

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

RAID Functions

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

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

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

1.1 Components and Features

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

1.1.1 Drive Group

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

1.1.2 Virtual Drive

A virtual drive is a partition in a drive group that is made up of contiguous data segments on the drives. A virtual drive can consist of these components:

- an entire drive group
- more than one entire drive group
- a part of a drive group
- parts of more than one drive group
- a combination of any two of these conditions

1.1.3 Fault Tolerance

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

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

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

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

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

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

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

1.1.4 Consistency Check

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

NOTE It is recommended that you perform a consistency check at least once a month.

1.1.5 Replace

Replace lets you copy data from a source drive into a destination drive that is not a part of the virtual drive. Replace often creates or restores a specific physical configuration for a drive group (for example, a specific arrangement of drive group members on the device I/O buses). You can run Replace automatically or manually.

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

Replace is also initiated when the first SMART error occurs on a drive that is part of a virtual drive. The destination drive is a hot spare that qualifies as a rebuild drive. The drive that has the SMART error is marked as *failed* only after the successful completion of the Replace. This situation avoids putting the drive group in Degraded status.

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

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

Order of Precedence

In the following scenarios, Rebuild takes precedence over Replace:

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

1.1.6 Background Initialization

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

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

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

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

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

1.1.7 Patrol Read

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

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

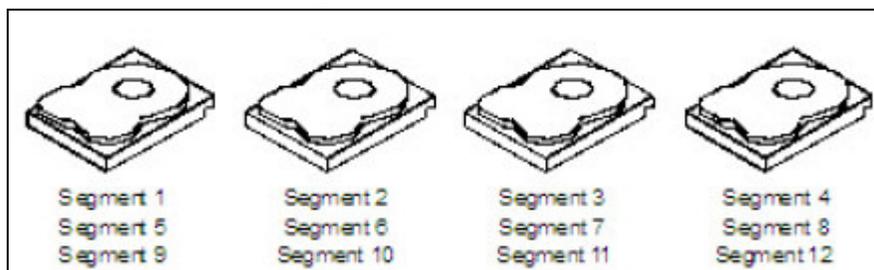
You can use the MegaRAID Storage Manager software to select the patrol read options, which you can use to set automatic or manual operation, or disable patrol read. See [Running a Patrol Read](#).

1.1.8 Disk Striping

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

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

Figure 1 Example of Disk Striping (RAID 0)



Stripe Width

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

Stripe Size

The stripe size is the length of the interleaved data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 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.

Strip Size

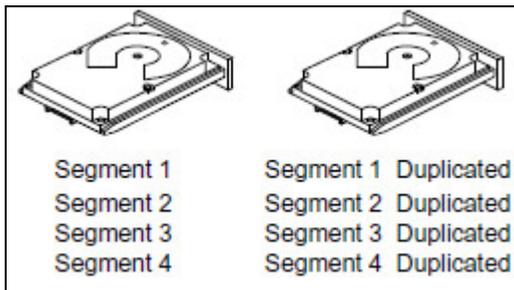
The strip size is the portion of a stripe that resides on a single drive.

1.1.9 Disk Mirroring

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

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

Figure 2 Example of Disk Mirroring (RAID 1)



1.1.10 Parity

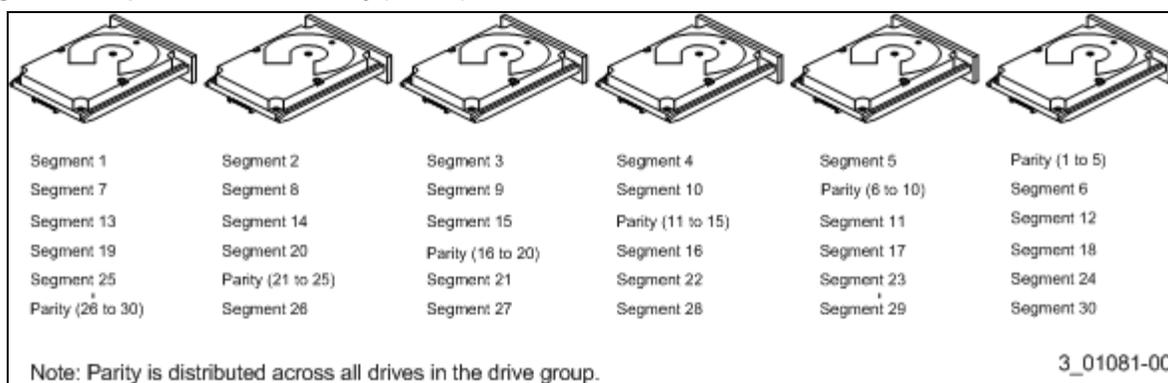
Parity generates a set of redundancy data from two or more parent data sets. The redundancy data can be used to reconstruct one of the parent data sets in the event of a drive failure. Parity data does not fully duplicate the parent data sets, but parity generation can slow the write process. In RAID,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

Table 1 Types of Parity

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

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

Figure 3 Example of Distributed Parity (RAID 5)



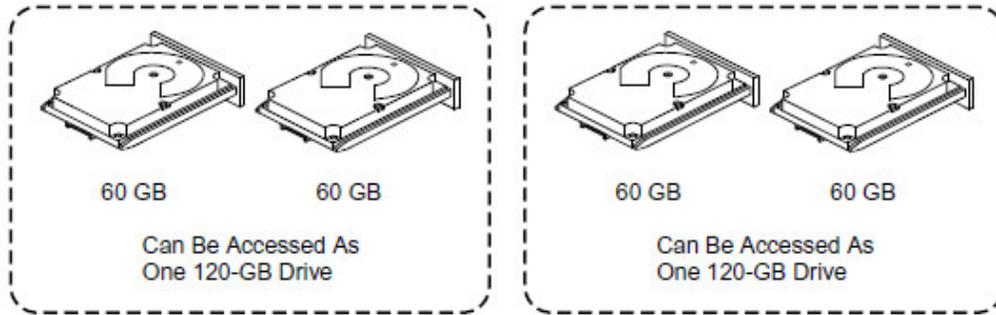
1.1.11 Disk Spanning

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

Spanning alone does not provide reliability or performance enhancements. Spanned virtual drives must have the same stripe size and must be contiguous. In the following figure, RAID 1 drive groups are turned into a RAID 10 drive group.

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

Figure 4 Example of Disk Spanning



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

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

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

See [Configuration](#) for detailed procedures for configuring drive groups and virtual drives, and spanning the drives.

Table 2 Spanning for RAID 10, RAID 50, and RAID 60

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

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

1.1.12 Hot Spares

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

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

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

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

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

The hot spare can be of two types:

- Global hot spare
- Dedicated hot spare

Global Hot Spare

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

Dedicated Hot Spare

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

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

Observe the following parameters when using hot spares:

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

- A hot spare must have free space equal to or greater than the drive it replaces. For example, to replace an 500-GB drive, the hot spare must be 500-GB or larger.

1.1.13 Disk Rebuilds

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

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

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

NOTE

When the rebuild to a hot spare begins, the failed drive is often removed from the virtual drive before management applications detect the failed drive. When this occurs, the events logs show the drive rebuilding to the hot spare without showing the failed drive. The formerly failed drive will be marked as “ready” after a rebuild begins to a hot spare. If a source drive fails during a rebuild to a hot spare, the rebuild fails, and the failed source drive is marked as offline. In addition, the rebuilding hot spare drive is changed back to a hot spare. After a rebuild fails because of a source drive failure, the dedicated hot spare is still dedicated and assigned to the correct drive group, and the global hot spare is still global.

An automatic drive rebuild will not start if you replace a drive during a RAID-level migration. The rebuild must be started manually after the expansion or migration procedure is complete. (RAID-level migration changes a virtual drive from one RAID level to another.)

1.1.14 Rebuild Rate

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

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

1.1.15 Hot Swap

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

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

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

1.1.16 Drive States

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

Table 3 Drive States

State	Description
Online	A drive that can be accessed by the RAID controller and is part of the virtual drive.
Unconfigured Good	A drive that is functioning normally but is not configured as a part of a virtual drive or as a hot spare.
Hot Spare	A drive that is powered up and ready for use as a spare in case an online drive fails.
Failed	A drive that was originally configured as Online or Hot Spare, but on which the firmware detects an unrecoverable error.
Rebuild	A drive to which data is being written to restore full redundancy for a virtual drive.
Unconfigured Bad	A drive on which the firmware detects an unrecoverable error; the drive was Unconfigured Good or the drive could not be initialized.
Missing	A drive that was Online but which has been removed from its location.
Offline	A drive that is part of a virtual drive but which has invalid data as far as the RAID configuration is concerned.

1.1.17 Virtual Drive States

The virtual drive states are described in the following table.

Table 4 Virtual Drive States

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

1.1.18 Beep Codes

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

Table 5 Beep Codes, Events, and Virtual Drive States

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

1.1.19 Enclosure Management

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

1.2 RAID Levels

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

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

1.2.1 Summary of RAID Levels

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

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

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

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

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

RAID 10, a combination of RAID 0 and RAID 1, consists of striped data across mirrored spans. A RAID 10 drive group is a spanned drive group that creates a striped set from a series of mirrored drives. RAID 10 allows a maximum of 8 spans. You must use an even number of drives in each RAID virtual drive in the span. The RAID 1 virtual drives must have the same stripe size. RAID 10 provides high data throughput and complete data redundancy but uses a larger number of spans.

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

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

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

NOTE The MegaSR controller supports the standard RAID levels – RAID 0, RAID 1, RAID 5, and RAID 10. The MegaSR controller comes in two variants, SCU and AHCI, both supporting a maximum of eight physical drives. A maximum of eight virtual drives can be created (using RAID 0, RAID 1, RAID 5, and RAID 10 only) and controlled by the MegaSR controller. One virtual drive can be created on an array (a maximum of eight if no other virtual drives are already created on the MegaSR controller), or you can create eight arrays with one virtual drive each. However, on RAID10, you can create only one virtual drive on a particular array.

1.2.2 Selecting a RAID Level

Select the optimal RAID level when you create a system drive. The optimal RAID level for your drive group depends on a number of factors:

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

1.2.3 RAID 0

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

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

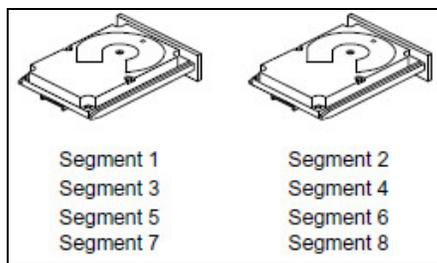
By breaking up a large file into smaller segments, the RAID controller can use SATA drives to read or write the file faster. RAID 0 involves no parity calculations to complicate the write operation. This situation makes RAID 0 ideal for applications that require high bandwidth but do not require fault tolerance.

The following table provides an overview of RAID 0. The following figure provides a graphic example of a RAID 0 drive group.

Table 6 RAID 0 Overview

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

Figure 5 RAID 0 Drive Group Example with Two Drives



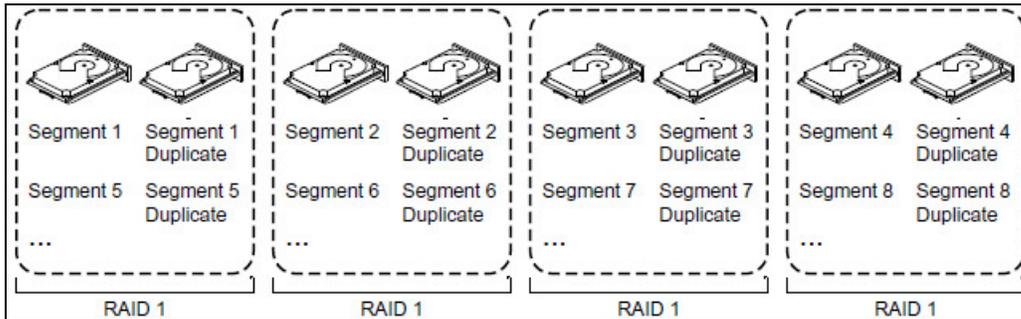
1.2.4 RAID 1

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

Table 7 RAID 1 Overview

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

Figure 6 RAID 1 Drive Group



1.2.5 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.

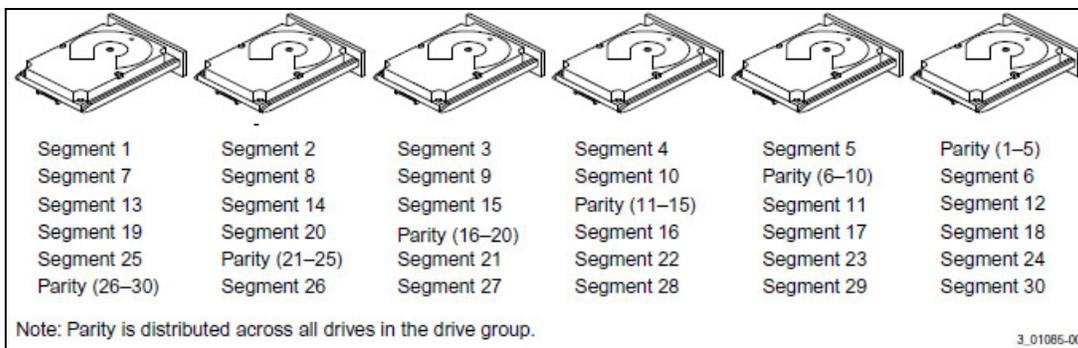
RAID 5 addresses the bottleneck issue for random I/O operations. Because each drive contains both data and parity, numerous writes can take place concurrently.

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

Table 8 RAID 5 Overview

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

Figure 7 RAID 5 Drive Group with Six Drives



1.2.6 RAID 6

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

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

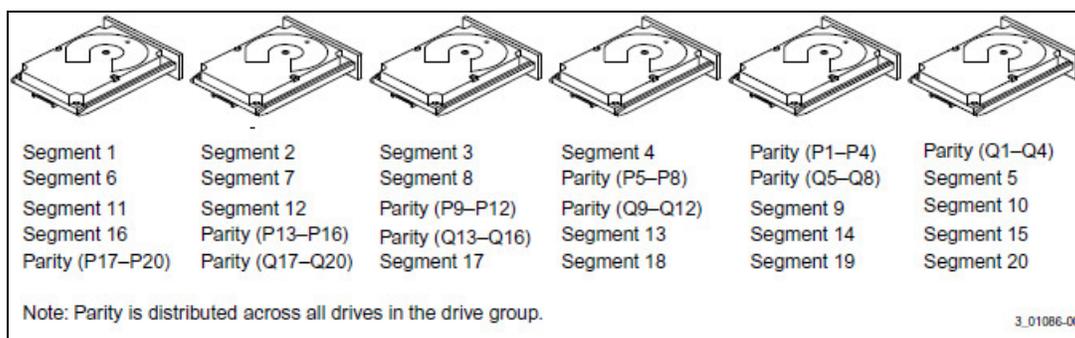
The following table provides an overview of a RAID 6 drive group.

Table 9 RAID 6 Overview

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

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

Figure 8 Example of Distributed Parity across Two Blocks in a Stripe (RAID 6)



1.2.7 RAID 00

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

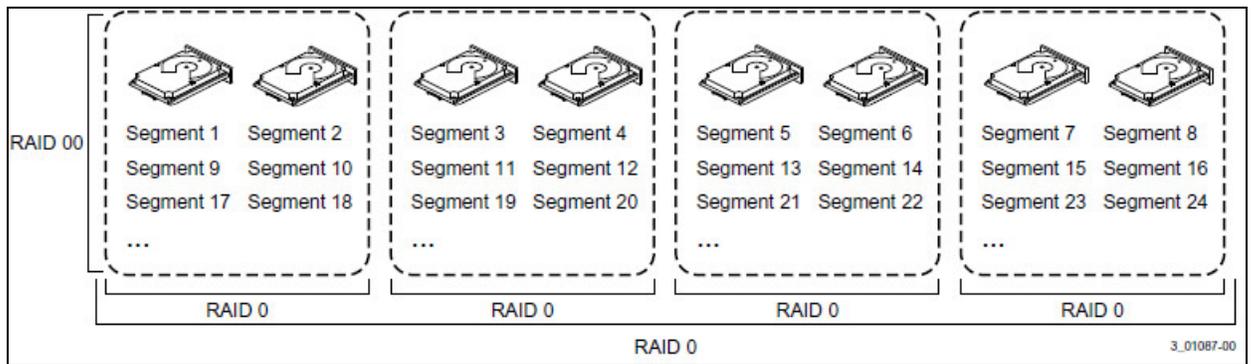
NOTE RAID level 00 is not fault tolerant. If a drive in a RAID 0 drive group fails, the entire virtual drive (all drives associated with the virtual drive) fails.

By breaking up a large file into smaller segments, the controller can use SATA drives to read or write the file faster. RAID 00 involves no parity calculations to complicate the write operation. This situation makes RAID 00 ideal for applications that require high bandwidth but do not require fault tolerance. The following table provides an overview of RAID 00. The following figure provides a graphic example of a RAID 00 drive group.

Table 10 RAID 00 Overview

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

Figure 9 RAID 00 Drive Group Example with Two Drives



1.2.8 RAID 10

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

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

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

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

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

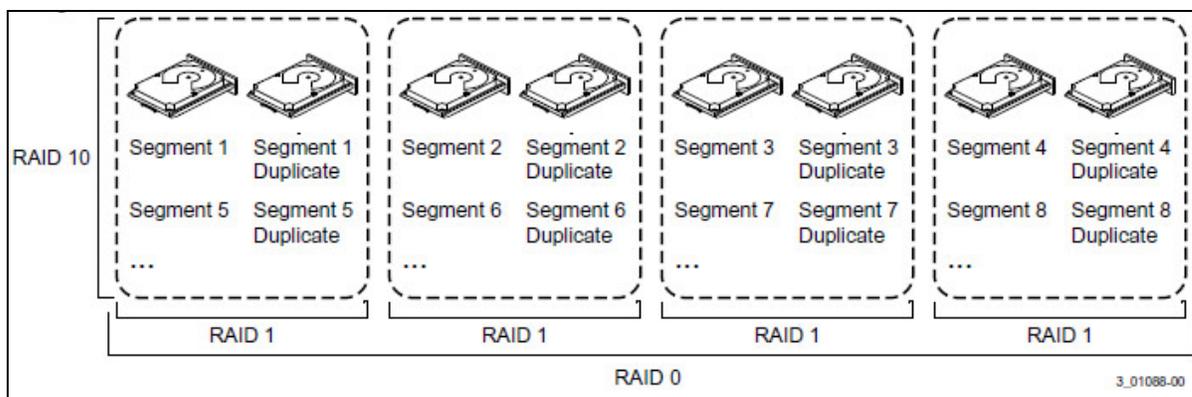
The following table provides an overview of RAID 10.

Table 11 RAID 10 Overview

Uses	Appropriate when used with data storage that needs 100 percent redundancy of mirrored drive groups and that also needs the enhanced I/O performance of RAID 0 (striped drive groups.) RAID 10 works well for medium-sized databases or any environment that requires a higher degree of fault tolerance and moderate-to-medium capacity.
Strong Points	Provides both high data transfer rates and complete data redundancy.
Weak Points	Requires twice as many drives as all other RAID levels except RAID 1.
Drives	4 to 32 in multiples of 4 — The maximum number of drives supported by the controller (using an even number of drives in each RAID 10 virtual drive in the span).

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

Figure 10 RAID 10 Level Virtual Drive



1.2.9 RAID 50

RAID 50 provides the features of both RAID 0 and RAID 5. RAID 50 includes both distributed parity and drive 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 operation 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.

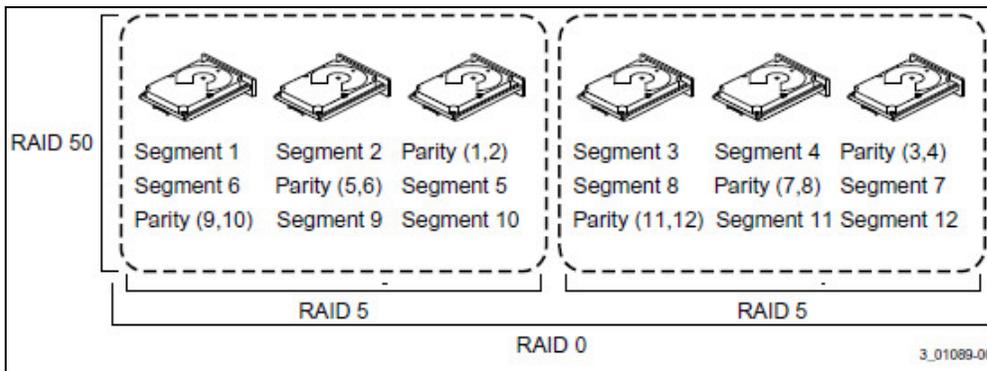
RAID level 50 can support up to eight spans and tolerate up to eight drive failures, though less than total drive capacity is available. Though multiple drive failures can be tolerated, only one drive failure can be tolerated in each RAID 5 level drive group.

The following table provides an overview of RAID 50.

Table 12 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 two times to eight times as many parity drives as RAID 5.
Drives	8 spans of RAID 5 drive groups containing 3 to 32 drives each (limited by the maximum number of devices supported by the controller)

Figure 11 RAID 50 Level Virtual Drive



1.2.10 RAID 60

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

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

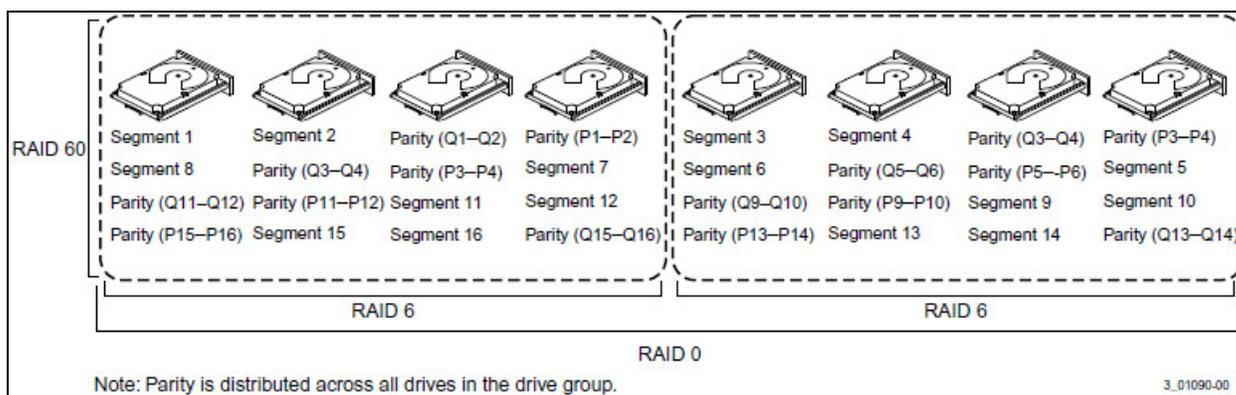
RAID 60 can support up to 8 spans and tolerate up to 16 drive failures, though less than total drive capacity is available. Two drive failures can be tolerated in each RAID 6 level drive group.

Table 13 RAID 60 Overview

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

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

Figure 12 RAID 60 Level Virtual Drive



1.3 RAID Configuration Strategies

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

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

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

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

1.3.1 Maximizing Fault Tolerance

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

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

Table 14 RAID Levels and Fault Tolerance

RAID Level	Fault Tolerance
0	Does not provide fault tolerance. All data is lost if any drive fails. Disk striping writes data across multiple drives instead of just one drive. It involves partitioning each drive storage space into stripes that can vary in size. RAID 0 is ideal for applications that require high performance but do not require fault tolerance.
1	Provides complete data redundancy. If one drive fails, the contents of the other drive in the drive group can be used to run the system and reconstruct the failed drive. The primary advantage of disk mirroring is that it provides 100 percent data redundancy. Because the contents of the drive are completely written to a second drive, no data is lost if one of the drives fails. Both drives contain the same data at all times. RAID 1 is ideal for any application that requires fault tolerance and minimal capacity.
5	Combines distributed parity with disk striping. Parity provides redundancy for one drive failure without duplicating the contents of entire drives. If a drive fails, the RAID controller uses the parity data to reconstruct all missing information. In RAID 5, this method is applied to entire drives or stripes across all drives in a drive group. Using distributed parity, RAID 5 offers fault tolerance with limited overhead.
6	Combines distributed parity with disk striping. RAID 6 can sustain two drive failures and still maintain data integrity. Parity provides redundancy for two drive failures without duplicating the contents of entire drives. If a drive fails, the RAID controller uses the parity data to reconstruct all missing information. In RAID 6, this method is applied to entire drives or stripes across all of the drives in a drive group. Using distributed parity, RAID 6 offers fault tolerance with limited overhead.
00	Does not provide fault tolerance. All data in a virtual drive is lost if any drive in that virtual drive fails. Disk striping writes data across multiple drives instead of just one drive. It involves partitioning each drive storage space into stripes that can vary in size. RAID 00 is ideal for applications that require high bandwidth but do not require fault tolerance.

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

1.3.2 Maximizing Performance

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

Table 15 RAID Levels and Performance

RAID Level	Performance
0	RAID 0 (striping) offers excellent performance. RAID 0 breaks up data into smaller blocks and then writes a block to each drive in the drive group. Disk striping writes data across multiple drives instead of just one drive. It involves partitioning each drive storage space into stripes that can vary in size from 8 KB to 1024 KB. These stripes are interleaved in a repeated sequential manner. Disk striping enhances performance because multiple drives are accessed simultaneously.
1	With RAID 1 (mirroring), each drive in the system must be duplicated, which requires more time and resources than striping. Performance is impaired during drive rebuilds.
5	RAID 5 provides high data throughput, especially for large files. Use this RAID level for any application that requires high read request rates, but low write request rates, such as transaction processing applications, because each drive can read and write independently. Because each drive contains both data and parity, numerous writes can take place concurrently. In addition, robust caching algorithms and hardware-based exclusive-or assist make RAID 5 performance exceptional in many different environments. Parity generation can slow the write process, making write performance significantly lower for RAID 5 than for RAID 0 or RAID 1. Drive performance is reduced when a drive is being rebuilt. Clustering can also reduce drive performance. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.
6	RAID 6 works best when used with data that requires high reliability, high request rates, and high data transfer. It provides high data throughput, data redundancy, and very good performance. However, RAID 6 is not well suited to tasks requiring a lot of writes. A RAID 6 virtual drive has to generate two sets of parity data for each write operation, which results in a significant decrease in performance during writes. Drive performance is reduced during a drive rebuild. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.
00	RAID 00 (striping in a spanned drive group) offers excellent performance. RAID 00 breaks up data into smaller blocks and then writes a block to each drive in the drive groups. Disk striping writes data across multiple drives instead of just one drive. Striping involves partitioning each drive storage space into stripes that can vary in size from 8 KB to 1024 KB. These stripes are interleaved in a repeated sequential manner. Disk striping enhances performance because multiple drives are accessed simultaneously.
10	RAID 10 works best for data storage that need the enhanced I/O performance of RAID 0 (striped drive groups), which provides high data transfer rates. Spanning increases the capacity of the virtual drive and improves performance by doubling the number of spindles. The system performance improves as the number of spans increases. (The maximum number of spans is 8.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans, and RAID performance degrades to that of a RAID 1 or RAID 5 drive group.

RAID Level	Performance
50	RAID 50 works best when used with data that requires high reliability, high request rates, and high data transfer. It provides high data throughput, data redundancy, and very good performance. Spanning increases the capacity of the virtual drive and improves performance by doubling the number of spindles. The system performance improves as the number of spans increases. (The maximum number of spans is 8.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans and RAID performance degrades to that of a RAID 1 or RAID 5 drive group.
60	RAID 60 works best when used with data that requires high reliability, high request rates, and high data transfer. It provides high data throughput, data redundancy, and very good performance. Spanning increases the capacity of the virtual drive and improves performance by doubling the number of spindles. The system performance improves as the number of spans increases. (The maximum number of spans is 8.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans, and RAID performance degrades to that of a RAID 1 or RAID 6 drive group. RAID 60 is not well suited to tasks requiring a lot of writes. A RAID 60 virtual drive has to generate two sets of parity data for each write operation, which results in a significant decrease in performance during writes. Drive performance is reduced during a drive rebuild. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.

1.3.3 Maximizing Storage Capacity

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

Table 16 RAID Levels and Capacity

RAID Level	Capacity
0	RAID 0 (striping) involves partitioning each drive storage space into stripes that can vary in size. The combined storage space is composed of stripes from each drive. RAID 0 provides maximum storage capacity for a given set of drives. The usable capacity of a RAID 0 array is equal to the number of drives in the array into the capacity of the smallest drive in the array.
1	With RAID 1 (mirroring), data written to one drive is simultaneously written to another drive, which doubles the required data storage capacity. This situation is expensive because each drive in the system must be duplicated. The usable capacity of a RAID 1 array is equal to the capacity of the smaller of the two drives in the array.
5	RAID 5 provides redundancy for one drive failure without duplicating the contents of entire drives. RAID 5 breaks up data into smaller blocks, calculates parity by performing an exclusive-or on the blocks and then writes the blocks of data and parity to each drive in the drive group. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set. The usable capacity of a RAID 5 array is equal to the number of drives in the array, minus one, into the capacity of the smallest drive in the array.
6	RAID 6 provides redundancy for two drive failures without duplicating the contents of entire drives. However, it requires extra capacity because it uses two parity blocks per stripe. This makes RAID 60 more expensive to implement. The usable capacity of a RAID 6 array is equal to the number of drives in the array, minus two, into the capacity of the smallest drive in the array.
00	RAID 00 (striping in a spanned drive group) involves partitioning each drive storage space into stripes that can vary in size. The combined storage space is composed of stripes from each drive. RAID 00 provides maximum storage capacity for a given set of drives.
10	RAID 10 requires twice as many drives as all other RAID levels except RAID 1. RAID 10 works well for medium-sized databases or any environment that requires a higher degree of fault tolerance and moderate-to-medium capacity. Disk spanning allows multiple drives to function like one large drive. Spanning overcomes lack of disk space and simplifies storage management by combining existing resources or adding relatively inexpensive resources.
50	RAID 50 requires two to four times as many parity drives as RAID 5. This RAID level works best when used with data that requires medium to large capacity.
60	RAID 60 provides redundancy for two drive failures in each RAID set without duplicating the contents of entire drives. However, it requires extra capacity because a RAID 60 virtual drive has to generate two sets of parity data for each write operation. This situation makes RAID 60 more expensive to implement.

1.4 RAID Availability

1.4.1 RAID Availability Concept

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

Spare Drives

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

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

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

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

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

Rebuilding

If a drive fails in a drive group that is configured as a RAID 1, 5, 6, 10, 50, or 60 virtual drive, you can recover the lost data by rebuilding the drive. If you have configured hot spares, the RAID controller automatically tries to use them to rebuild failed drives. Manual rebuild is necessary if hot spares with enough capacity to rebuild the failed drives are not available. You must insert a drive with enough storage into the subsystem before rebuilding the failed drive.

1.5 Configuration Planning

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

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

1.6 Number of Drives

Your configuration planning for the SAS RAID controller depends in part on the number of drives that you want to use in a RAID drive group.

The number of drives in a drive group determines the RAID levels that can be supported. Only one RAID level can be assigned to each virtual drive.

Drive Group Purpose

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

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

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

Table 17 Factors to Consider for Drive Group Configuration

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

Chapter 2 WebBIOS Configuration Utility

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

2.1 Overview

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

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

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

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

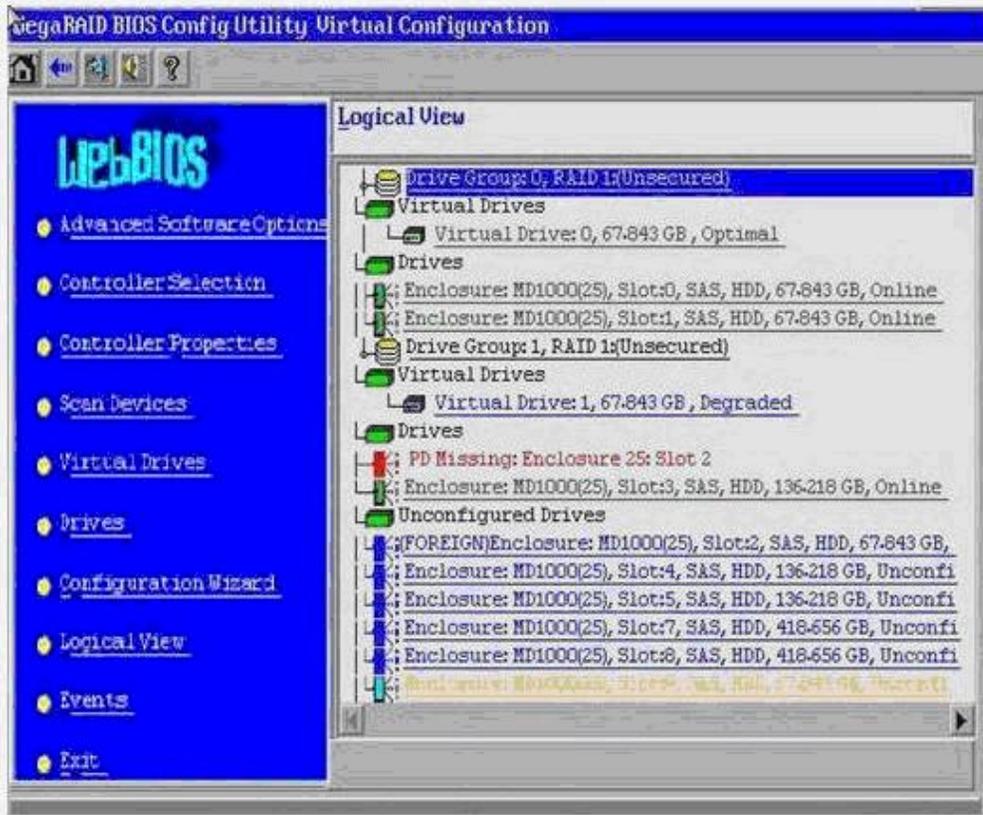
2.2 Starting the WebBIOS Configuration Utility

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

1. When the host computer is booting, press and hold down the Ctrl key and press the H key when the following text appears on the dialog:
`Copyright© LSI Corporation Press <Ctrl><H> for WebBIOS`
The **Controller Selection** dialog appears.
2. Click **Start** to continue.
The main **WebBIOS Configuration Utility** dialog appears.

2.3 WebBIOS Configuration Utility Main Dialog Options

Figure 13 WebBIOS Configuration Utility Main Dialog



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

NOTE In the list of virtual drives, the drive nodes are sorted based on the order in which you added the drives to the drive group, rather than the physical slot order that displays in the physical trees. The minimum dialog resolution for WebBIOS is 640 x 480.

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

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

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

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

The toolbar at the top of the WebBIOS configuration utility has the following buttons, as listed in the following table.

Table 18 WebBIOS Configuration Utility Toolbar Icons

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

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

- **Advanced Software Options** (Alt+a) (Not Supported): Select this option to enable the advanced features in the controller.
- **Controller Selection** (Alt+c): Select this option to view the **Controller Selection** dialog, where you can select a different SAS controller. You can also view information about the controller and the devices connected to it, or create a new configuration on the controller.

- **Controller Properties** (Alt+p): Select this option to view the properties of the currently selected SAS controller. For more information, see [Viewing Controller Properties](#).
- **Scan Devices** (Alt+s): Select this option to have the WebBIOS configuration utility re-scan the physical and virtual drives for any changes in the drive status or the physical configuration. The WebBIOS configuration utility displays the results of the scan in the physical and virtual drive descriptions.
- **Virtual Drives** (Alt+v): Select this option to view the **Virtual Drives** dialog, where you can change and view virtual drive properties, delete virtual drives, initialize drives, and perform other tasks. For more information, see [Viewing Virtual Drive Properties, Policies, and Operations](#).
- **Drives** (Alt+d): Select this option to view the **Drives** dialog, where you can view drive properties, create hot spares, and perform other tasks. For more information, see [Viewing Drive Properties](#).
- **Configuration Wizard** (Alt+o): Select this option to start the **Configuration Wizard** and create a new storage configuration, clear a configuration, or add a configuration. For more information, see [Creating a Storage Configuration](#).
- **Logical View/Physical View** (Alt+l for the Logical view; Alt+h for the Physical view): Select this option to toggle between the **Physical View** dialog and the **Logical View** dialog.
- **Events** (Alt+e): Select this option to view system events in the **Event Information** dialog. For more information, see [Viewing System Event Information](#).
- **Exit** (Alt+x): Select this option to exit the WebBIOS configuration utility and continue with system boot.

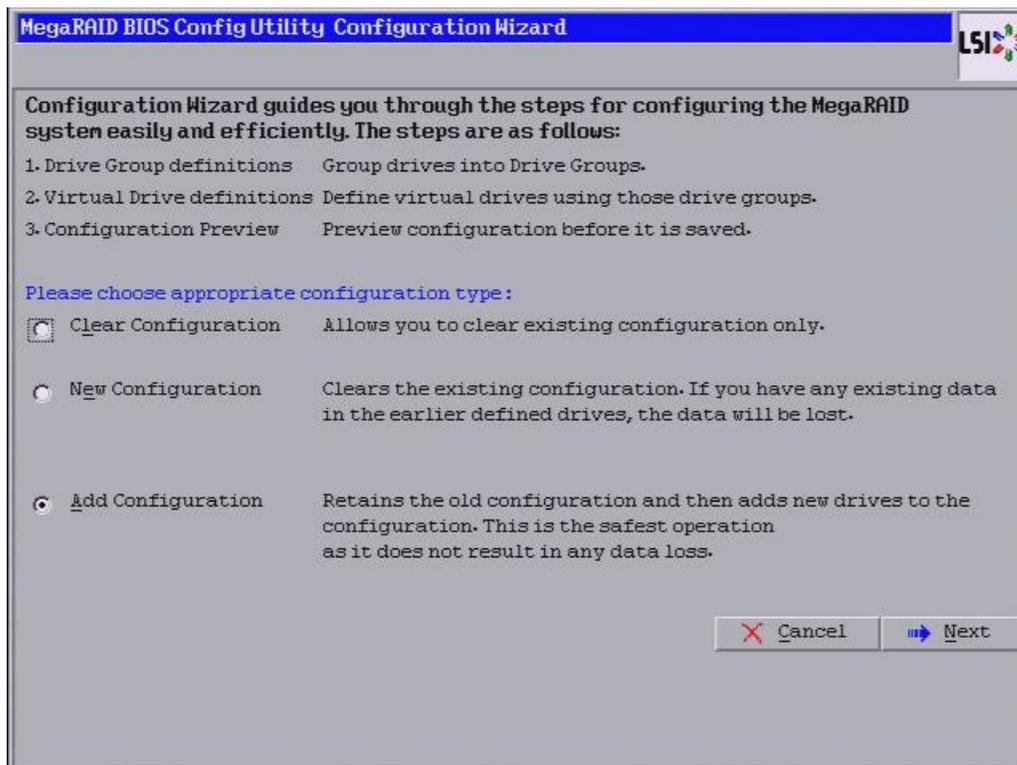
2.4 Creating a Storage Configuration

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

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

1. Click Configuration Wizard on the WebBIOS main dialog.

Figure 14 WebBIOS Configuration Wizard Dialog



2. Select a configuration option.

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

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

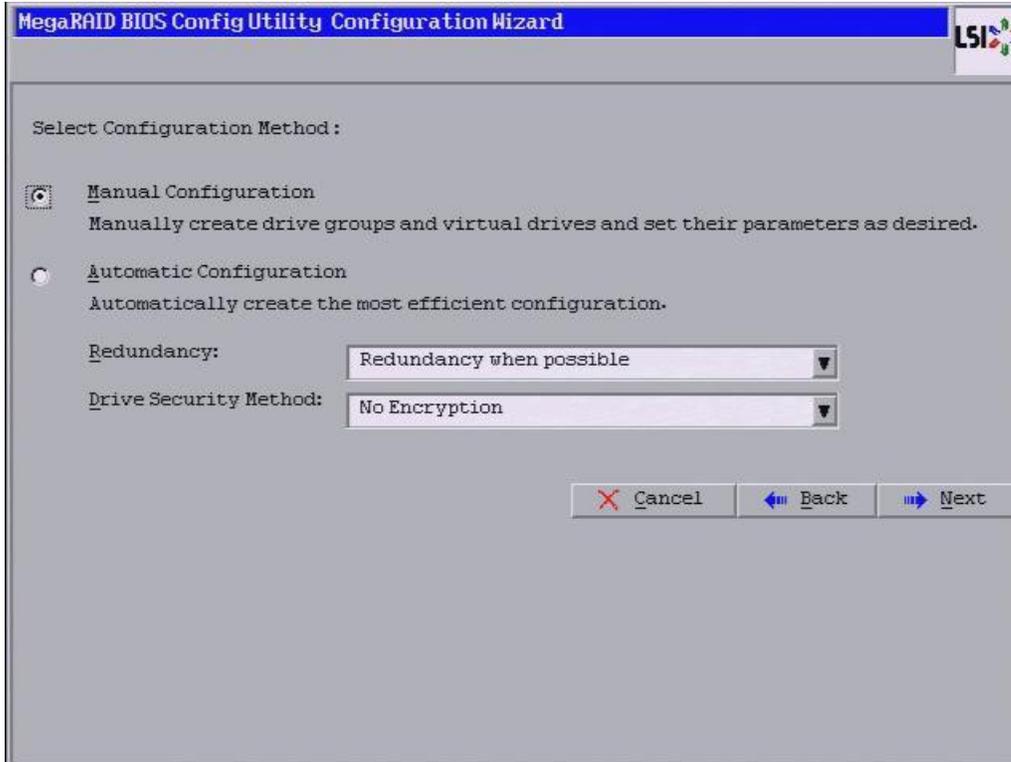
3. Click **Next**.

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

4. Click **Next**.

The **WebBIOS Configuration Method** dialog appears.

Figure 15 WebBIOS Configuration Method Wizard



5. Select a configuration mode:

- **Manual Configuration:** Allows you to control all attributes of the new storage configuration as you create drive groups and virtual drives, and set their parameters.
- **Automatic Configuration:** Automatically creates an optimal RAID configuration.

If you select **Automatic Configuration**, you can choose whether to create a redundant RAID drive group or a non-redundant RAID 0 drive group. Select one of the following options in the **Redundancy** drop-down list:

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

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

- **No Encryption**
- **Full Disk Encryption**

6. Click **Next** to continue.

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

2.4.1 Using Automatic Configuration

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

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

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

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

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

2.4.2 Using Manual Configuration

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

The following procedures include setting virtual drive options. These options are explained in the following Section, **Virtual Drive Options**, which appears before the manual configuration procedures.

2.4.2.1 Virtual Drive Options

This section explains the virtual drive options that are set using the manual procedures for creating RAID drive groups for RAID levels 0, 1, 5, 6, 00, 10, 50, and 60.

- **RAID Level:** The drop-down list shows the possible RAID levels for the virtual drive.
 - **RAID 0:** Select this option for RAID 0.
 - **RAID 1:** Select this option for RAID 1.
 - **RAID 5:** Select this option for RAID 5.
 - **RAID 6:** Select this option for RAID 6.

- **RAID 00:** Select this option for RAID 00.
 - **RAID 10:** Select this option for RAID 10.
 - **RAID 50:** Select this option for RAID 50.
 - **RAID 60:** Select this option for RAID 60.
-
- **Strip Size:** The strip size is the portion of a stripe that resides on a single drive in the drive group. The stripe consists of the data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB, and the strip size is 16 KB. You can set the strip size to **8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB, and 1024 KB**. A larger strip size produces higher read performance. If your computer regularly performs random read requests, choose a smaller strip size. The default is **64 KB**.

NOTE WebBIOS does not allow you to select **8 KB** as the strip size when you create a RAID 6 drive group with three drives or a RAID 60 drive group with six drives.

- **Access Policy:** Select the type of data access that is allowed for this virtual drive.
 - **RW:** Allow read/write access. This is the default.
 - **Read Only:** Allow read-only access.
 - **Blocked:** Do not allow access.

- **Read Policy:** Specify the read policy for this virtual drive.
 - **No Read Ahead:** This option disables the read ahead capability. This option is the default.
 - **Always Read Ahead:** This option enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This option speeds up reads for sequential data, but there is little improvement when accessing random data.

- **Write Policy:** Specify the write policy for this virtual drive.
 - **Always Write Back:** In Write back mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.
 - **Write Through:** In Write through mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This option is the default setting.
 - **Write Back with BBU:** Select this mode if you want the controller to use Write back mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Write through mode if it detects a bad or missing BBU.

NOTE Write back mode can be used with or without a BBU. Use *either* a battery to protect the controller cache, or an uninterruptible power supply (UPS) to protect the entire system. If you do not use a battery or a UPS, and a power failure occurs, you risk losing the data in the controller cache.

- **IO Policy:** The IO policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.
 - **Direct:** In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. This option is the default setting.
 - **Cached:** In Cached I/O mode, all reads are buffered in cache memory.
- **Drive Cache:** Specify the drive cache policy.
 - **Enable:** Enable the drive cache.
 - **Disable:** Disable the drive cache. This option is the default setting.
 - **Unchanged:** Leave the current drive cache policy as is.
- **Disable BGI:** Specify the Background Initialization (BGI) status.
 - **No:** Leave background initialization enabled, which means that a new configuration can be initialized in the background while you use WebBIOS to perform other configuration tasks. This option is the default setting.
 - **Yes:** Select **Yes** if you do not want to allow background initialization for configurations on this controller.

NOTE New RAID 5 virtual drives require at least five drives for a background initialization to start. New RAID 6 virtual drives require at least seven drives for a background initialization to start.

- **Select Size:** Specify the size of the virtual drive in MB, GB, or TB. Usually, this is the full size for RAID 0, RAID1, RAID 5, RAID 6, RAID 00, RAID 10, RAID 50, or RAID 60 shown in the **Configuration** panel on the right. You can specify a smaller size if you want to create other virtual drives on the same drive group.
- **Update Size:** Click **Update Size** to update the Select size value for the selected RAID levels.
- **Provide Shared Access:** Select this option if you want the virtual drive to be shared between the servers in a cluster. By default, drives are shared. This option appears only if the controller supports High Availability DAS.

2.4.2.2 Using Manual Configuration: RAID 0

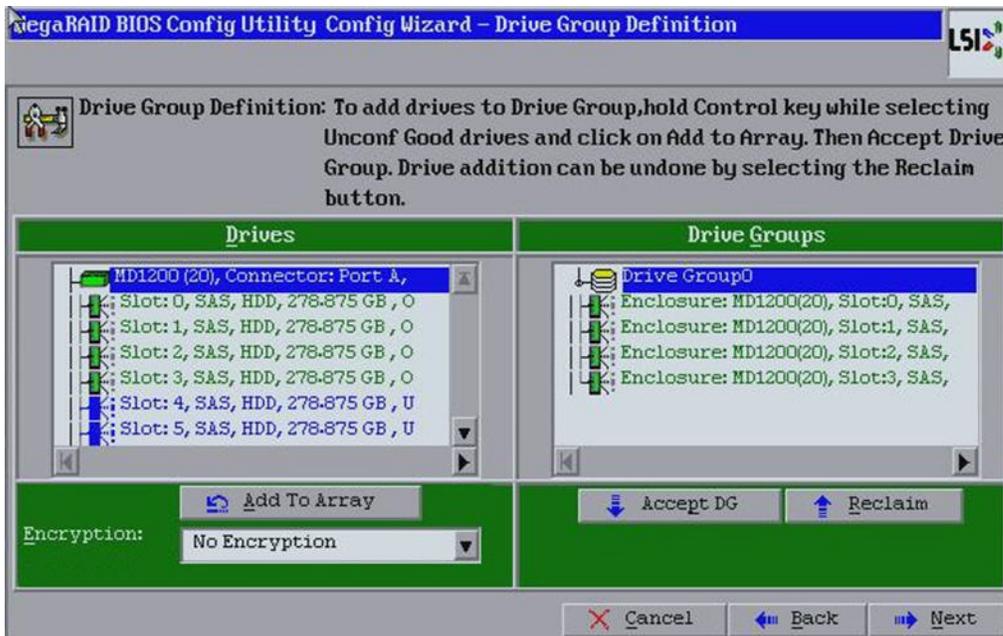
RAID 0 provides drive striping across all drives in the RAID drive group. RAID 0 does not provide any data redundancy but does offer excellent performance. RAID 0 is ideal for applications that require high bandwidth but do not require fault tolerance. RAID 0 also denotes an independent or single drive.

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

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

1. Press and hold the Ctrl key while selecting two or more unconfigured good drives in the **Drives** panel on the left until you have selected all desired drives for the drive group.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the **Drive Groups** panel on the right, as shown in the following figure.
If you need to undo the changes, select the drive and click **Reclaim**.
3. Choose whether to use drive encryption.

Figure 16 Drive Group Definition Dialog



4. After you finish selecting drives for the drive group, click **Accept DG**.
5. Click **Next**.
The **Span Definition** dialog appears. This dialog shows the drive group holes that you can select to add to a span.
6. Under the Array With Free Space frame, select a drive group, and click **Add to SPAN**.
The drive group you select appears in the right frame under Span.

- Click **Next**.

The **Virtual Drive Definition** dialog appears, as shown in the following figure. This dialog lists the possible RAID levels for the drive group.

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

Figure 17 Virtual Drive Definition

The screenshot shows the 'Virtual Drive Definition' dialog box. The title bar reads 'RAID BIOS Config Utility Config Wizard - Virtual Drive Definition' and includes the LSI logo. The dialog is divided into two main sections: configuration options on the left and a 'Virtual Drives' area on the right.

RAID Level	RAID 0
Strip Size	256 KB
Access Policy	RW
Read Policy	Always Read Ahead
Write Policy	Write Through
I/O Policy	Direct
Drive Cache	Unchanged
Disable BGI	No
Select Size	[] GB

At the bottom of the configuration section, there is a checked checkbox for 'Provide Shared Access', an 'Update Size' button, and 'Accept' and 'Reclaim' buttons.

The 'Virtual Drives' section contains a large empty box and the following text:

Next LD, Possible RAID Levels
R0:408.656 GB R5:272.437 GB R6:136.218 GB

At the bottom of the dialog, there are 'Cancel', 'Back', and 'Next' buttons.

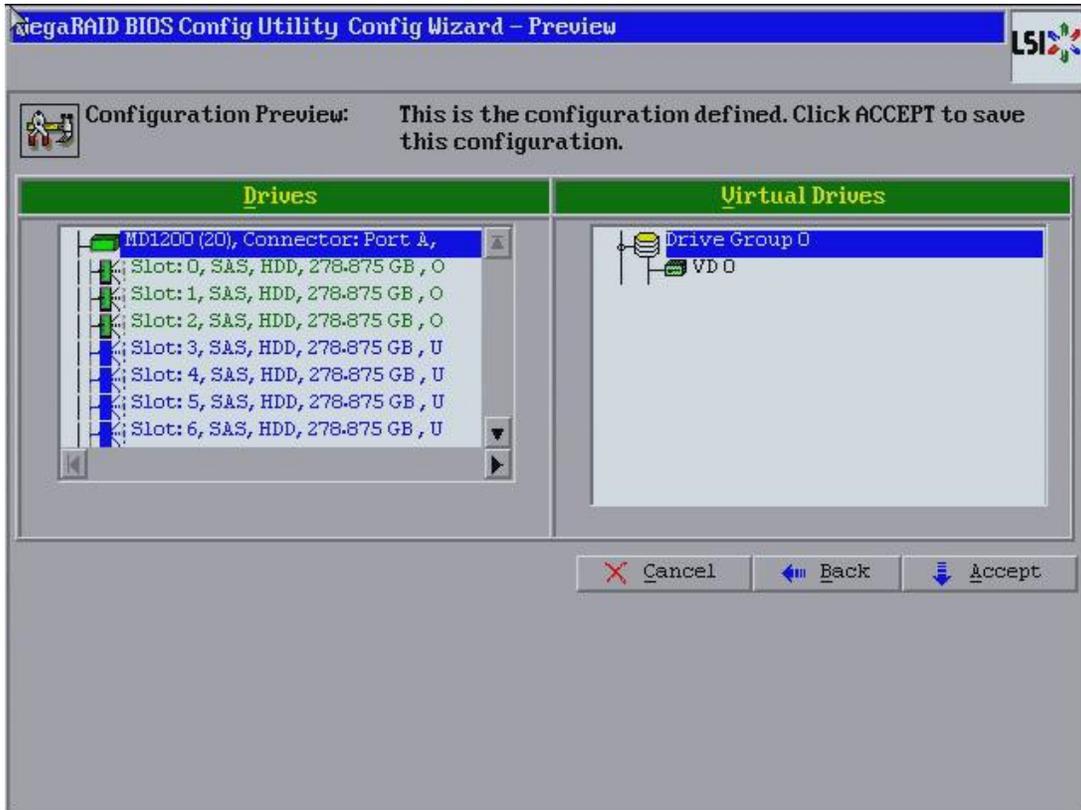
- Change the virtual drive options from the defaults listed on the dialog as needed.

NOTE For specific information about virtual drive options, see [Virtual Drive Options](#). The **Provide Shared Access** check box appears only if the controller supports High Availability DAS.

- Click **Accept** to accept the changes to the virtual drive definition.
A confirmation dialog appears.
- To confirm your changes, click **Yes**. Otherwise, to undo the changes, select the virtual drive and click **Reclaim**.

12. Click **Next** after you finish defining the virtual drives.
The **Configuration Preview** dialog appears, as shown in the following figure.

Figure 18 RAID 0 Configuration Preview Dialog



13. Check the information in the **Configuration Preview** dialog.
14. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Back** to return to the previous dialogs and change the configuration.
15. If you accept the configuration, click **Yes** at the prompt to save the configuration.
Another confirmation for initialization appears.
16. Click **No** to view the WebBIOS main menu. Otherwise, click **Yes** and the initialization process takes place, and the WebBIOS Config Utility Virtual Drive dialog appears.

2.4.2.3 Using Manual Configuration: RAID 1

In RAID 1, the RAID controller duplicates all data from one drive to a second drive. RAID 1 provides complete data redundancy, but at the cost of doubling the required data storage capacity. It is appropriate for small databases or any other environment that requires fault tolerance but small capacity.

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

1. Press and hold the Ctrl key while you select at least two unconfigured good drives in the **Drives** panel on the left. You must select an even number of drives.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the **Drive Groups** panel on the right, as shown in the following figure.
If you need to undo the changes, select the drive and click **Reclaim**.
3. Choose whether to use drive encryption.

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

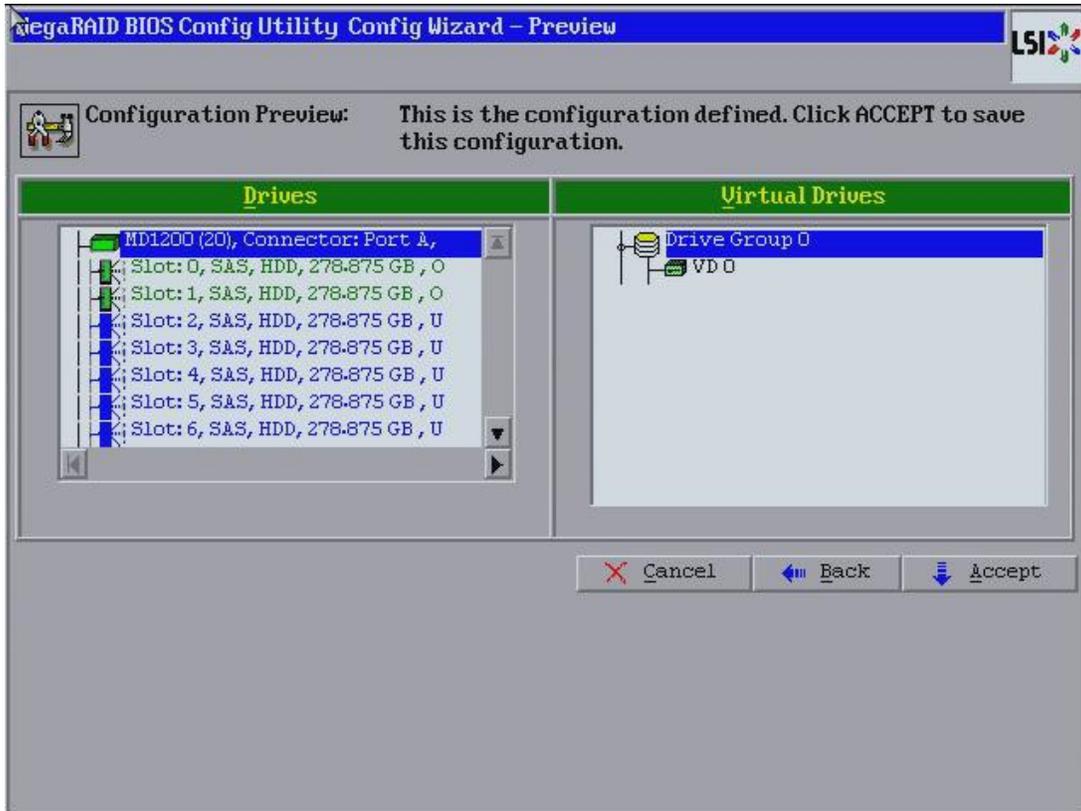
4. After you finish selecting drives for the drive group, click **Accept DG**.
5. Click **Next**.
The **Span Definition** dialog appears. This dialog shows the drive group holes that you can select to add to a span. You use this dialog to select the RAID level, strip size, read policy, and other attributes for the new virtual drives.
6. Under the Array With Free Space frame, select a drive group, and click **Add to SPAN**.
The drive group you select appears in the right frame under Span.
7. Click **Next**. The **Virtual Drive Definition** dialog appears.
8. Change the virtual drive options from the defaults listed on the dialog as needed.

NOTE For specific information about virtual drive options, see [Virtual Drive Options](#).

9. Click **Accept** to accept the changes to the virtual drive definition.
A confirmation dialog appears.
10. To confirm your changes, click **Yes**. Otherwise, to undo the changes, select the virtual drive and click **Reclaim**.

11. Click **Next** after you finish defining the virtual drives.
The **Configuration Preview** dialog appears, as shown in the following figure.

Figure 19 RAID 1 Configuration Preview Dialog



12. Check the information in the **Configuration Preview** dialog.
13. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Back** to return to the previous dialogs and change the configuration.
14. If you accept the configuration, click **Yes** at the prompt to save the configuration.
Another confirmation for initialization appears.
15. Click **No** to view the WebBIOS main menu. Otherwise, click **Yes** and the initialization process takes place, and the WebBIOS Config Utility Virtual Drive dialog appears.

2.4.2.4 Using Manual Configuration: RAID 5

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

RAID 5 provides high data throughput. RAID 5 is useful for transaction processing applications because each drive can read and write independently. If a drive fails, the RAID controller uses the parity drive to re-create all missing information. You can use RAID 5 for office automation and online customer service that require fault tolerance.

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

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

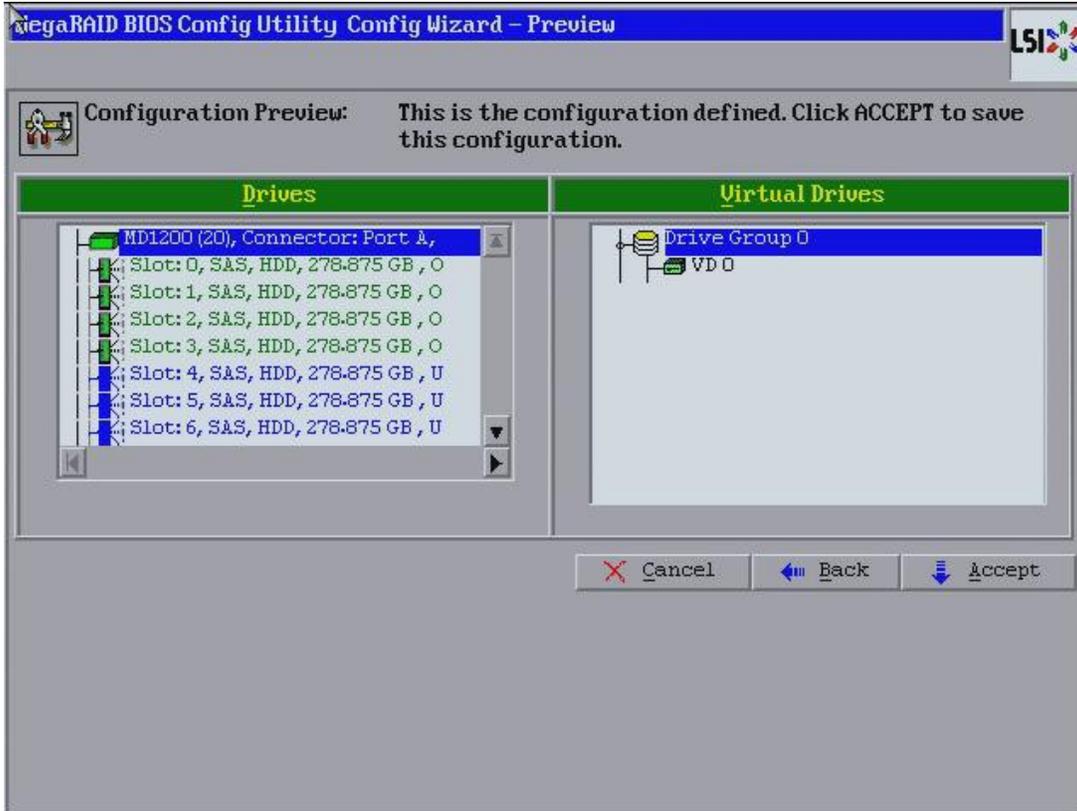
1. Press and hold the Ctrl key while you select at least three unconfigured good drives in the **Drives** panel on the left.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the **Drive Groups** panel on the right, as shown in the following figure.
If you need to undo the changes, select the drive and click **Reclaim**.
3. Choose whether to use drive encryption.
4. After you finish selecting drives for the drive group, click **Accept DG**.
5. Click **Next**
The **Span Definition** dialog appears. This dialog shows the drive group holes that you can select to add to a span.
6. Under the Array With Free Space frame, select a drive group, and click **Add to SPAN**.
The drive group you select appears in the right frame under Span.
7. Click **Next**.
The **Virtual Drive Definition** dialog appears.
8. Use this dialog to select the RAID level, strip size, read policy, and other attributes for the new virtual drives.
9. Change the virtual drive options from the defaults listed on the dialog as needed.

NOTE For specific information about virtual drive options, see [Virtual Drive Options](#).

10. Click **Accept** to accept the changes to the virtual drive definition.
A confirmation dialog appears.
11. To confirm your changes, click **Yes**. Otherwise, to undo the changes, select the virtual drive and click **Reclaim**.

12. Click **Next** after you finish defining the virtual drives.
The **Configuration Preview** dialog appears, as shown in the following figure.

Figure 20 RAID 5 Configuration Preview Dialog



13. Check the information in the **Configuration Preview** dialog.
14. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Cancel** to end the operation dialogs and return to the WebBIOS main menu, or click **Back** to return to the previous dialogs and change the configuration.
15. If you accept the configuration, click **Yes** at the prompt to save the configuration.
Another confirmation for initialization appears.
16. Click **No** to view the WebBIOS main menu. Otherwise, click **Yes** and the initialization process takes place, and the WebBIOS Config Utility Virtual Drive dialog appears.

2.4.2.5 Using Manual Configuration: RAID 6

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

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

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

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

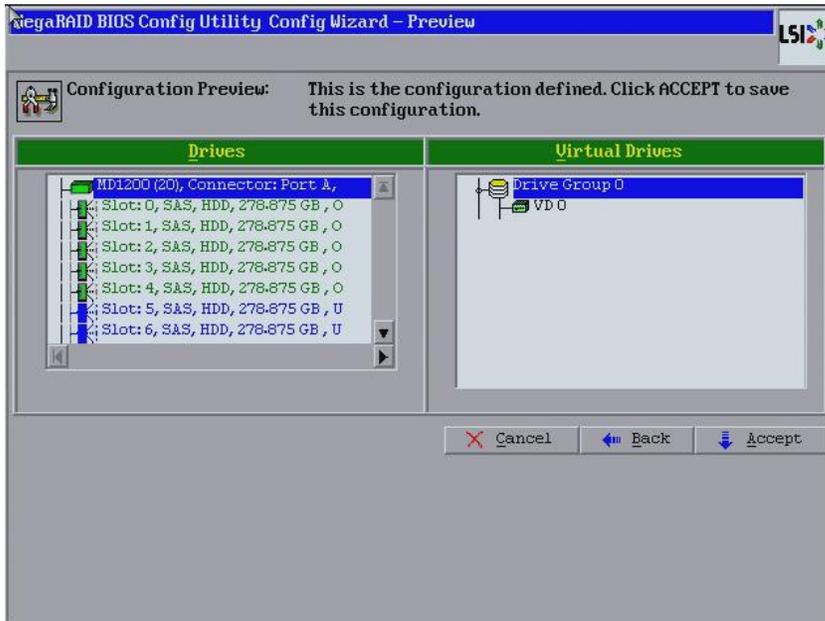
When you select **Manual Configuration**, and click **Next**, the **Drive Group Definition** dialog appears. You use this dialog to select drives to create drive groups.

1. Press and hold the Ctrl key while you select at least three unconfigured good drives in the **Drives** panel on the left.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the **Drive Groups** panel on the right, as shown in the following figure. If you need to undo the changes, select the drive and click **Reclaim**.
3. Choose whether to use drive encryption.
The drop-down list in the **Encryption** field lists the options.
4. After you finish selecting drives for the drive group, click **Accept DG** for each drive.
5. Click **Next**. The **Span Definition** dialog appears. This dialog shows the drive group holes that you can select to add to a span.
6. Under the Array With Free Space frame, select a drive group, and click **Add to SPAN**.
The drive group you select appears in the right frame under Span.
7. Click **Next**. The **Virtual Drive Definition** dialog appears.
8. Change the virtual drive options from the defaults listed on the dialog as needed.

NOTE For specific information about virtual drive options, see [Virtual Drive Options](#).

9. Click **Accept** to accept the changes to the virtual drive definition. A confirmation dialog appears.
10. To confirm your changes, click **Yes**. Otherwise, to undo the changes, select the virtual drive and click **Reclaim**.
11. Click **Next** after you finish defining the virtual drives. The **Configuration Preview** dialog appears, as shown in the following figure.

Figure 21 RAID 6 Configuration Preview Dialog



12. Check the information in the **Configuration Preview** dialog.
13. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Back** to return to the previous dialogs and change the configuration.
14. If you accept the configuration, click **Yes** at the prompt to save the configuration.
Another confirmation for initialization appears.
15. Click **No** to view the WebBIOS main menu. Otherwise, click **Yes** and the initialization process takes place, and the WebBIOS Config Utility Virtual Drive dialog appears.

2.4.2.6 Using Manual Configuration: RAID 00

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

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

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

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

1. Press and hold the Ctrl key while you select unconfigured good drives in the **Drives** panel on the left.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the **Drive Groups** panel on the right.

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

3. Click **Accept DG** to create a first drive group.

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

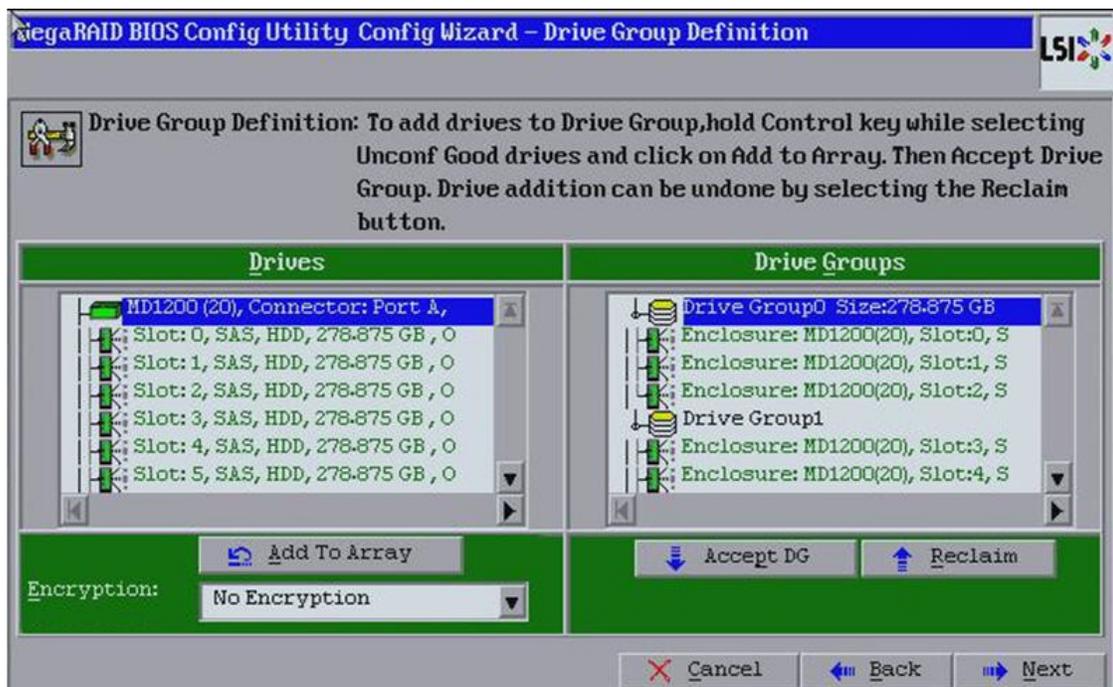
4. Press and hold the Ctrl key while you select the same number of unconfigured good drives in the **Drives** panel (that were selected for the first drive group) to create a second drive group.
5. Click **Add To Array** to move the drives to a second drive group configuration in the **Drive Groups** panel, as shown in the following figure.

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

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

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

Figure 22 Drive Group Definition Dialog

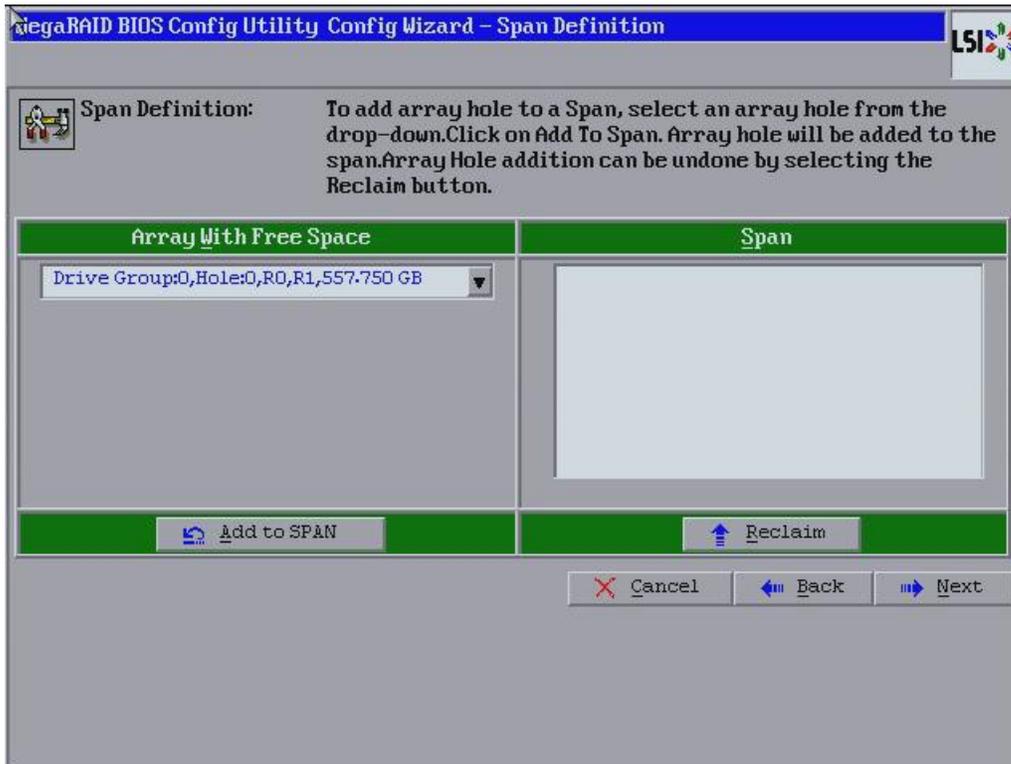


8. Repeat step 1 through step 5 until you have created all the required drive groups.

- Click **Next**.

The **Span Definition** dialog appears, as shown in the following figure. This dialog shows the drive group holes that you can select to add to a span.

Figure 23 Span Definition Dialog



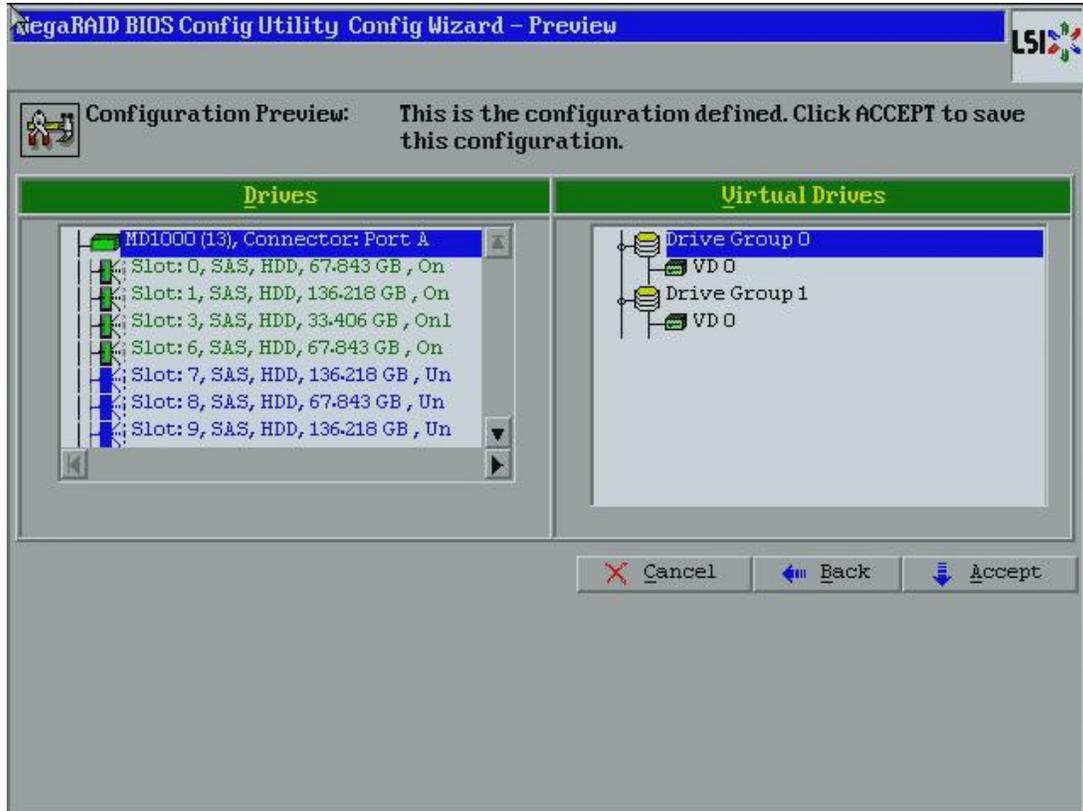
- Under the **Array With Free Space** frame, select a drive group, and then click **Add to SPAN**. The drive group you select appears in the right frame under **Span**.
- Repeat the previous steps until you have selected all of the drive groups that you want.
- Click **Next**. The **Virtual Drive Definition** dialog appears.
- Change the virtual drive options from the defaults listed on the dialog as needed.

NOTE For specific information about virtual drive options, see [Virtual Drive Options](#).

- Click **Accept** to accept the changes to the virtual drive definition. A confirmation dialog appears.
- To confirm your changes, click **Yes**. Otherwise, to undo the changes, select the virtual drive and click **Reclaim**.

16. After you finish defining the virtual drives, click **Next**.
The **Configuration Preview** dialog appears, as shown in the following figure.

Figure 24 RAID 00 Configuration Preview Dialog



17. Check the information in the **Configuration Preview** dialog.
18. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Cancel** to end the operation and return to the WebBIOS main menu, or click **Back** to return to the previous dialogs and change the configuration.
19. If you accept the configuration, click **Yes** at the prompt to save the configuration.
Another confirmation for initialization appears.
20. Click **No** to view the WebBIOS main menu. Otherwise, click **Yes** and the initialization process takes place, and the WebBIOS Config Utility Virtual Drive dialog appears.

2.4.2.7 Using Manual Configuration: RAID 10

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

RAID 10 provides both high data transfer rates and complete data redundancy. It works best for data storage that must have 100 percent redundancy of RAID 1 (mirrored drive groups) and that also needs the enhanced I/O performance of RAID 0 (striped drive groups); it works well for medium-sized databases or any environment that requires a higher degree of fault tolerance and moderate to medium capacity.

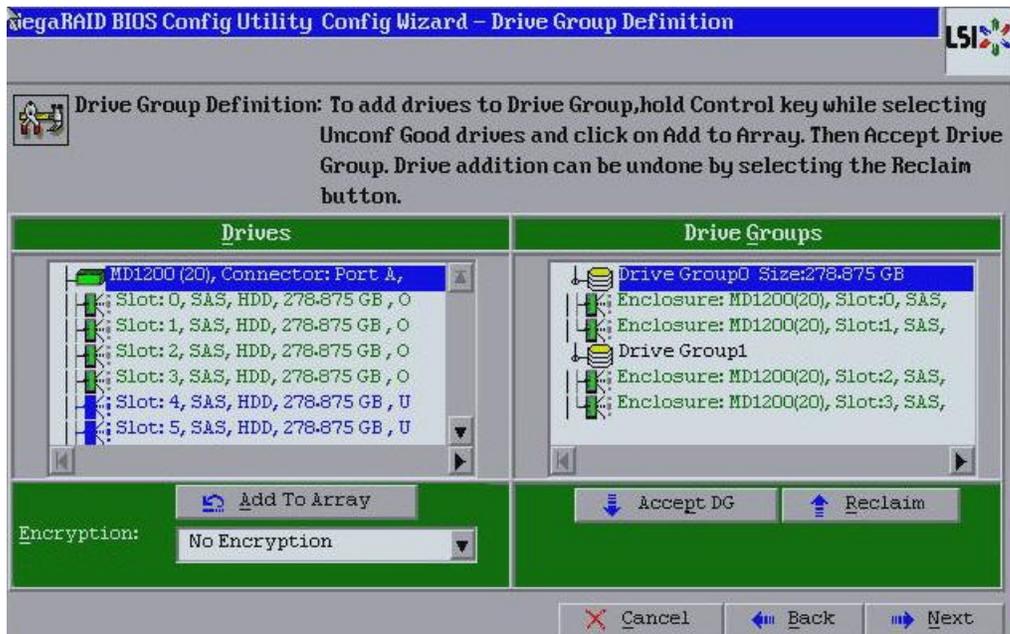
When you select **Manual Configuration** and click **Next**, the **Drive Group Definition** dialog appears.

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

1. Press and hold the Ctrl key while selecting two unconfigured good drives in the **Drives** panel on the left.
2. Click **Add To Array** to move the drives to a proposed two-drive group configuration in the **Drive Groups** panel on the right.
If you need to undo the changes, select the drive and click **Reclaim**.
3. Click **Accept DG** to create a first drive group.
An icon for the next drive group appears in the right panel.
4. Click the icon for the next drive group to select it.
5. Press and hold the Ctrl key while selecting same number of unconfigured good drives in the **Drives** panel to create a second RAID 1 drive group with two drives.
6. Click **Add To Array** to move the drives to a second drive group configuration in the **Drive Groups** panel, as shown in the following figure.
If you need to undo the changes, select the drive and click **Reclaim**.
7. Choose whether to use drive encryption.

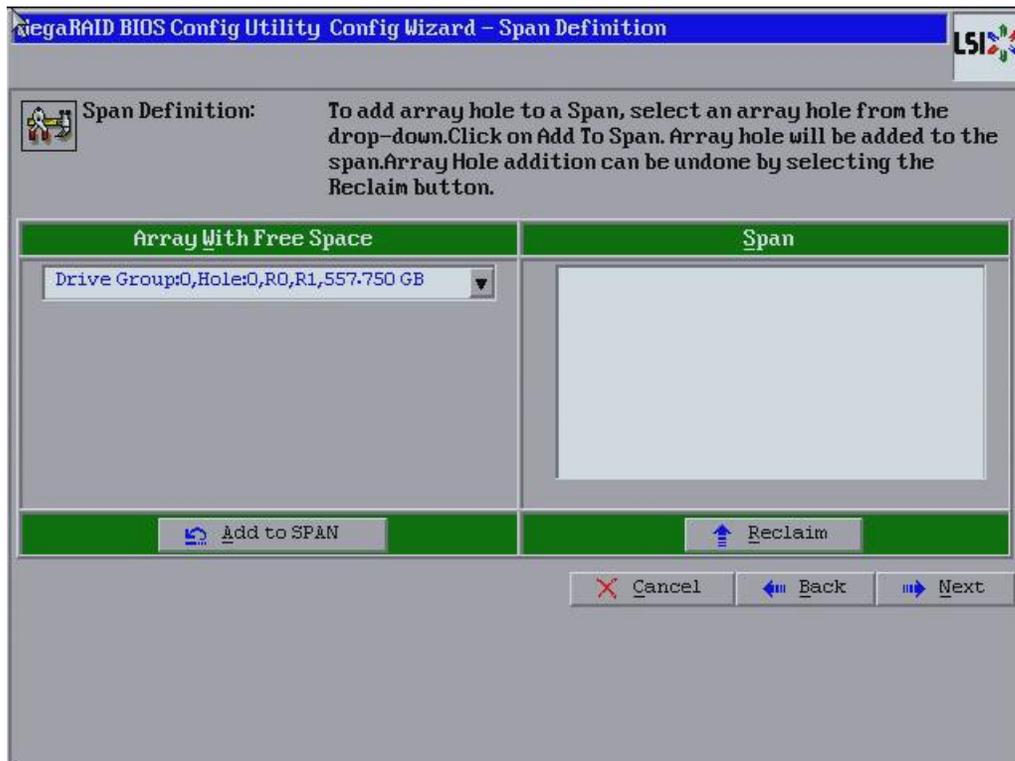
NOTE RAID 10 supports a maximum of eight spans, with a maximum of 32 drives per span. (Other factors, such as the type of controller, can limit the number of drives.) You must use an even number of drives in each RAID 10 drive group in the span.

Figure 25 Drive Group Definition Dialog



8. Repeat step 1 through step 6 until you have created all the required drive groups.
9. Click **Next**. The **Span Definition** dialog appears, as shown in the following figure. This dialog displays the drive group holes you can select to add to a span.

Figure 26 Span Definition Dialog



10. Under the **Array With Free Space** column, select a drive and click **Add to SPAN**.
The drive group you select displays in the right frame under the heading **Span**.
11. Select another drive group and click **Add to SPAN**.
If there are additional drive groups with two drives each, you can add them to the Span.
12. Click **Next**. The **Virtual Drive Definition** dialog appears.

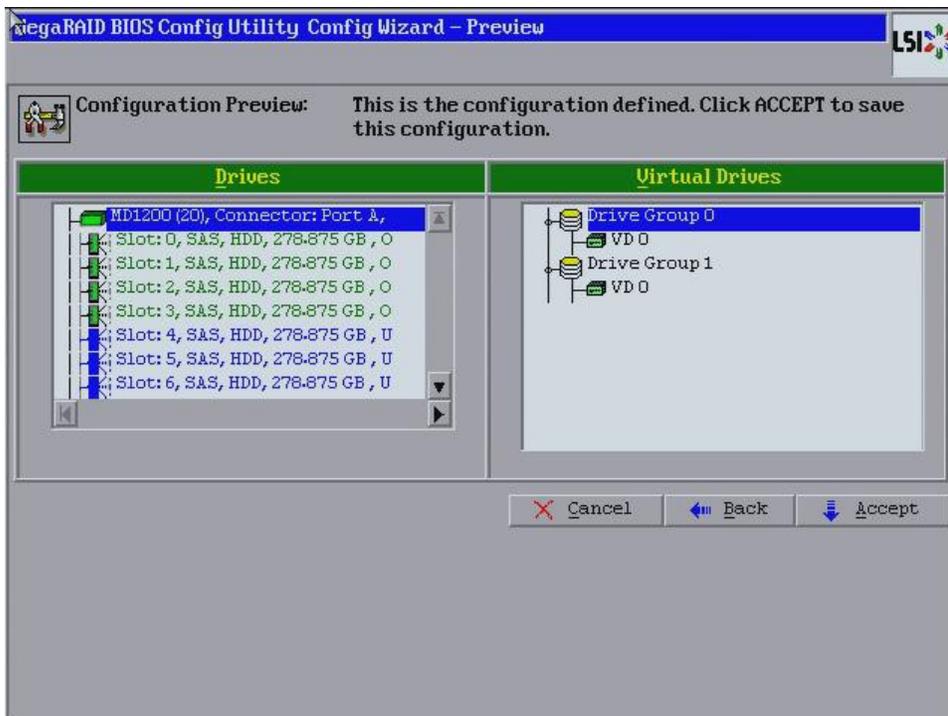
NOTE The WebBIOS Configuration Utility shows the maximum available capacity while creating the RAID 10 drive group. In version 1.03 of the utility, the maximum size of the RAID 10 drive group is the sum total of the two RAID 1 drive groups. In version 1.1, the maximum size is the size of the smaller drive group multiplied by 2.

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

NOTE For specific information about virtual drive options, see [Virtual Drive Options](#).

14. Click **Accept** to accept the changes to the virtual drive definition. A confirmation dialog appears.
15. To confirm your changes, click **Yes**. Otherwise, to undo the changes, select the virtual drive and click **Reclaim**.
16. After you finish defining the virtual drives, click **Next**. The **Configuration Preview** dialog appears, as shown in the following figure.

Figure 27 RAID 10 Configuration Preview Dialog



17. Check the information in the **Configuration Preview** dialog.
18. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Cancel** to end the operation and return to the WebBIOS main menu, or click **Back** to return to the previous dialogs and change the configuration.
19. If you accept the configuration, click **Yes** at the prompt to save the configuration.
Another confirmation for initialization appears.
20. Click **No** to view the WebBIOS main menu. Otherwise, click **Yes** and the initialization process takes place, and the WebBIOS Config Utility Virtual Drive dialog appears.

2.4.2.8 Using Manual Configuration: RAID 50

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

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

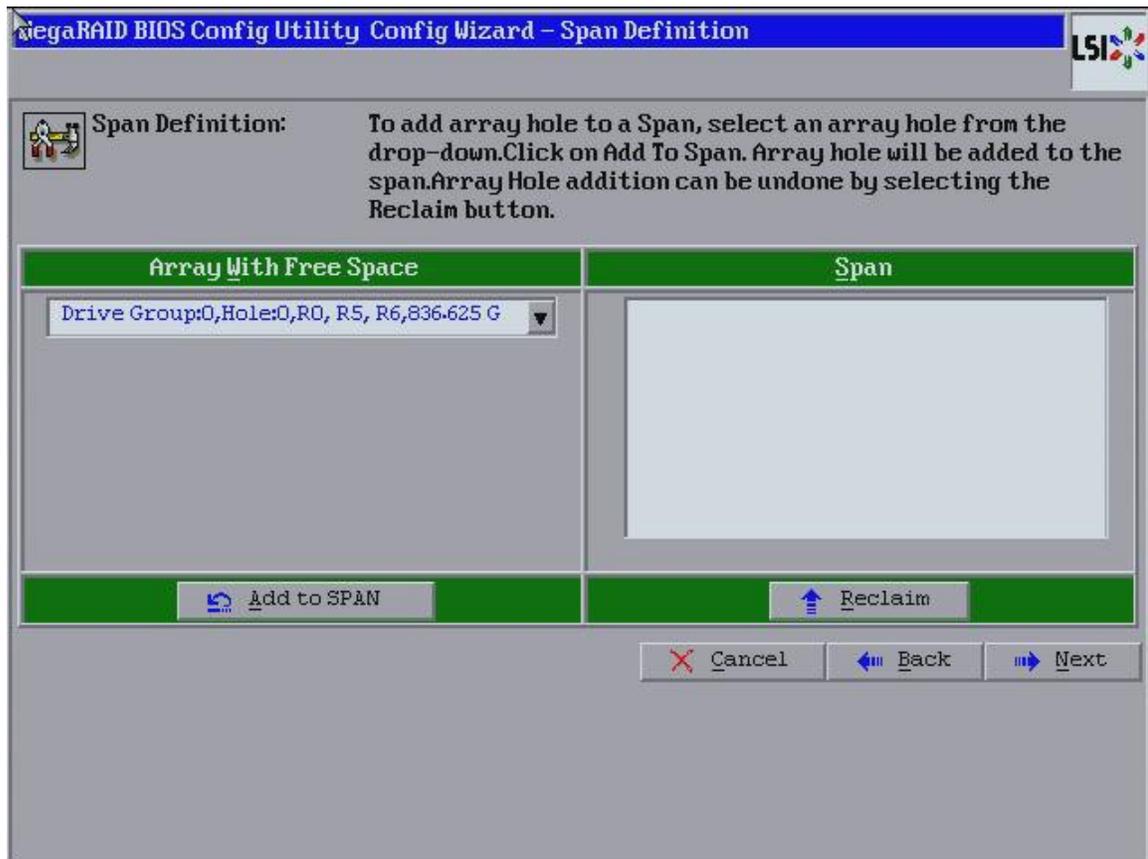
When you select **Manual Configuration** and click **Next**, the **Drive Group Definition** dialog appears. You use this dialog to select drives to create drive group.

1. Press and hold the Ctrl key while selecting at least three unconfigured good drives in the **Drives** panel on the left.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the **Drive Groups** panel on the right.
If you need to undo the changes, select the drive and click **Reclaim**.
3. Click **Accept DG** to create a first drive group.
An icon for a second drive group appears in the right panel.
4. Press and hold the Ctrl key while selecting the same number of unconfigured good drives in the **Drives** panel (that were selected for the first drive group) to create a second drive group.
5. Click **Add To Array** to move the drives to a proposed drive group configuration in the **Drive Groups** panel on the right, as shown in the following figure.
If you need to undo the changes, select the drive and click **Reclaim**.
6. Choose whether to use drive encryption.
7. Repeat step 1 through step 5 until you have created all the required drive groups.
8. After you finish selecting drives for the drive groups, select each drive group and click **Accept DG** for each drive group.

9. Click **Next**.

The **Span Definition** dialog appears, as shown in the following figure. This dialog displays the drive group holes you can select to add to a span.

Figure 28 Span Definition Dialog



10. Under the **Array With Free Space** column, select a drive group of three or more drives and click **Add to SPAN**. The drive group you select displays in the right frame under the heading **Span**.
11. Select another drive group and click **Add to SPAN**.
If there are additional drive groups with three drives each, you can add them to the span.
12. Click **Next**.
The Virtual Drive Definition dialog appears.
13. Change the virtual drive options from the defaults listed on the dialog as needed.

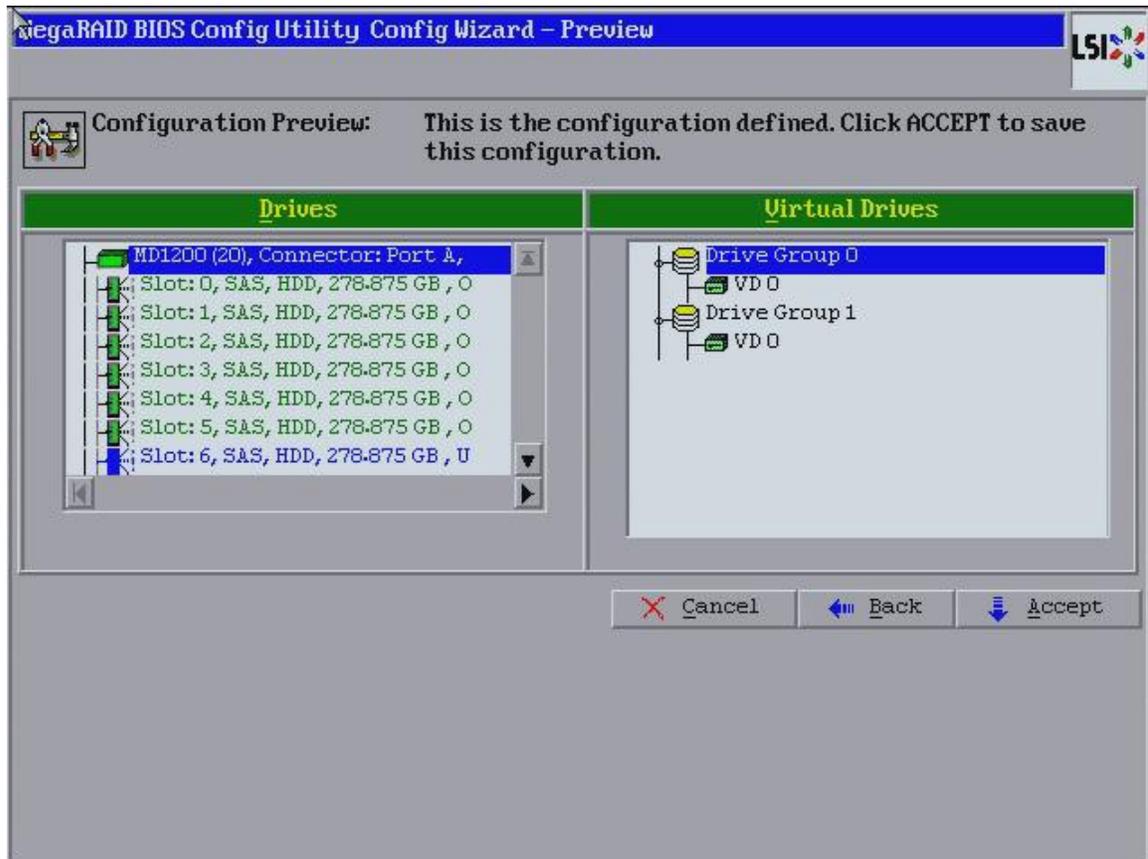
NOTE For specific information about virtual drive options, see [Virtual Drive Options](#).

14. Click **Accept** to accept the changes to the virtual drive definition.
A confirmation dialog appears.
15. To confirm your changes, click **Yes**. Otherwise, to undo the changes, select the virtual drive and click **Reclaim**.

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

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

Figure 29 RAID 50 Configuration Preview Dialog



17. Check the information in the **Configuration Preview** dialog.

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

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

Another confirmation for initialization appears.

20. Click **No** to view the WebBIOS main menu. Otherwise, click **Yes** and the initialization process takes place, and the WebBIOS Config Utility Virtual Drive dialog appears.

2.4.2.9 Using Manual Configuration: RAID 60

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

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

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

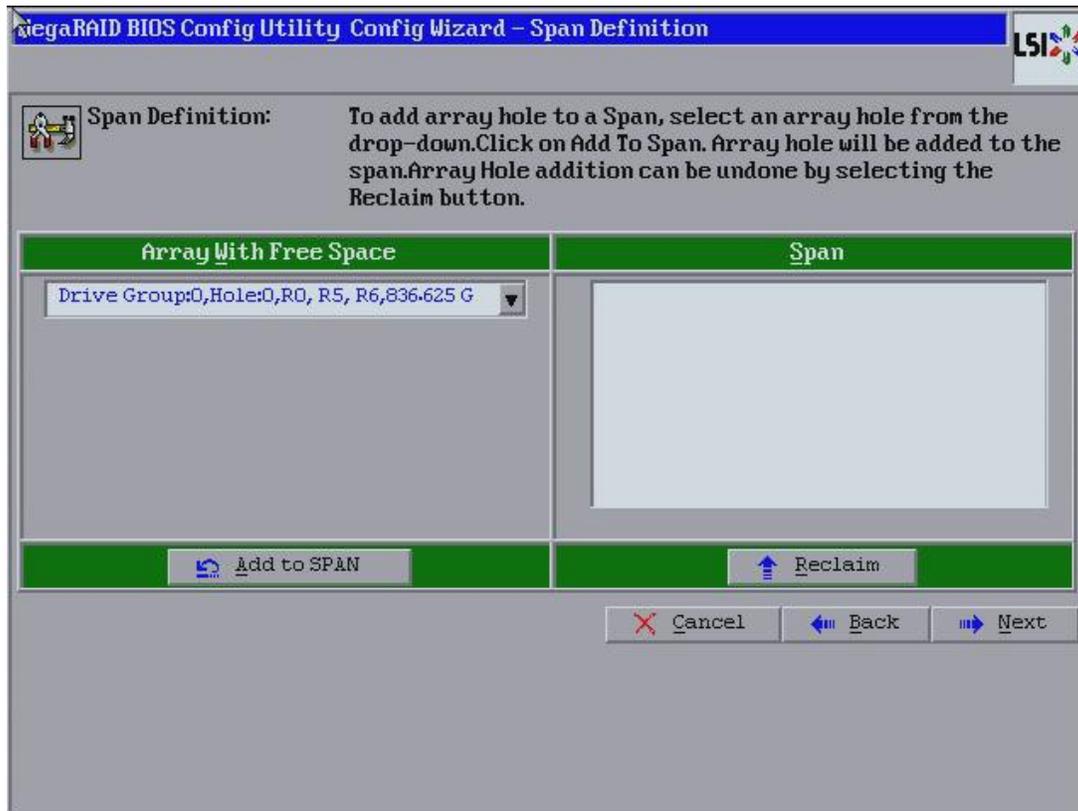
When you select **Manual Configuration** and click **Next**, the **Drive Group Definition** dialog appears. You use this dialog to select drives to create drive groups.

1. Press and hold the Ctrl key while selecting at least three unconfigured good drives in the **Drives** panel on the left.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the **Drive Groups** panel on the right. If you need to undo the changes, select the drive and click **Reclaim**.
3. Click **Accept DG** to create a first drive group.
An icon for a second drive group appears in the right panel.
4. Press and hold the Ctrl key while selecting the same number of unconfigured good drives in the **Drives** panel (that were selected for the first drive group) to create a second drive group.
5. Click **Add To Array** to move the drives to a proposed drive group configuration in the **Drive Groups** panel on the right, as shown in the following figure.
If you need to undo the changes, select the drive and click **Reclaim**.
6. Choose whether to use drive encryption.
7. Repeat step 1 through step 5 until you have created all the required drive groups.
8. After you finish selecting drives for the drive groups, select each drive group and click **Accept DG** for each drive group.

9. Click **Next**.

The **Span Definition** dialog appears, as shown in the following figure. This dialog displays the drive group holes you can select to add to a span.

Figure 30 WebBIOS Span Definition Dialog



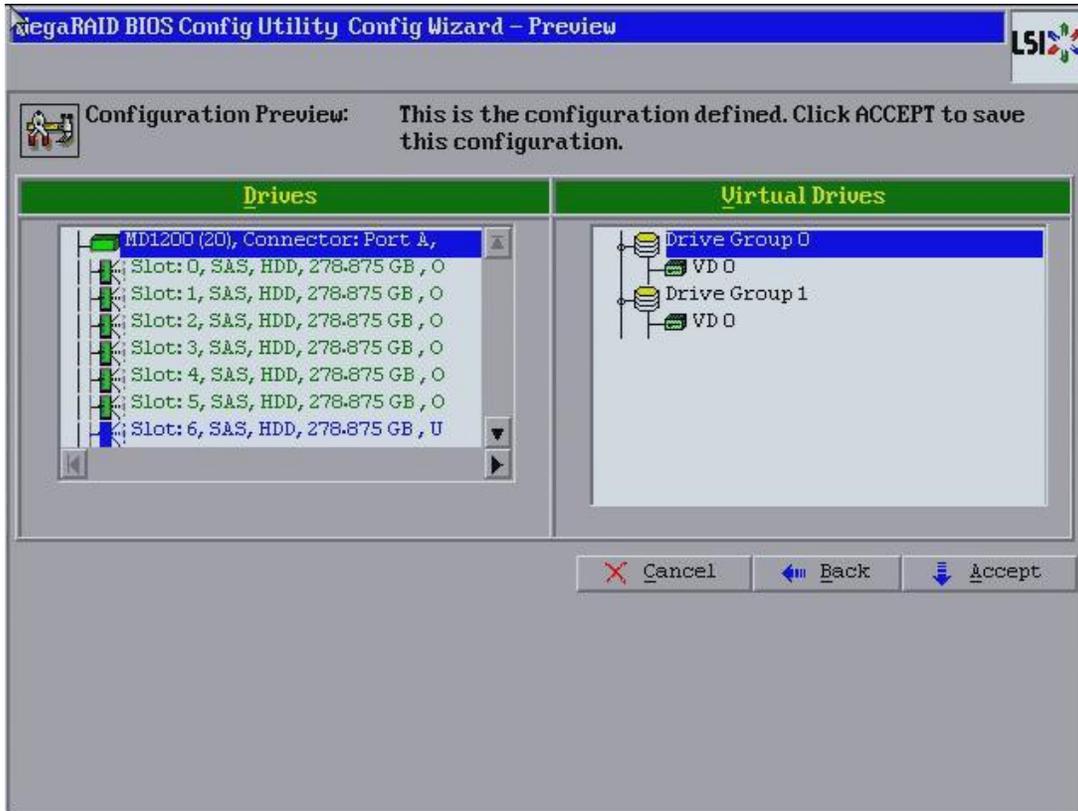
10. Under the heading **Array With Free Space**, select a drive group and click **Add to SPAN**.
The drive group you select displays in the right frame under the heading **Span**.
11. Select another drive group and click **Add to SPAN**.
If there are additional drive groups with three drives each, you can add them to the span.
12. Click **Next**.
The Virtual Drive Definition dialog appears.
13. Change the virtual drive options from the defaults listed on the dialog as needed.

NOTE For specific information about virtual drive options, see [Virtual Drive Options](#).

14. Click **Accept** to accept the changes to the virtual drive definition.
A confirmation dialog appears.
15. To confirm your changes, click **Yes**. Otherwise, to undo the changes, select the virtual drive and click **Reclaim**.

16. Click **Next** after you finish defining virtual drives.
The **Configuration Preview** dialog appears, as shown in the following figure.

Figure 31 RAID 60 Configuration Preview Dialog



17. Check the information in the **Configuration Preview** dialog.
18. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Back** to return to the previous dialogs and change the configuration.
19. If you accept the configuration, click **Yes** at the prompt to save the configuration.
Another confirmation for initialization appears.
20. Click **No** to view the WebBIOS main menu. Otherwise, click **Yes** and the initialization process takes place, and the WebBIOS Config Utility Virtual Drive dialog appears.

2.5 Viewing and Changing Device Properties

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

2.5.1 Viewing Controller Properties

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

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

1. Click **Controller Properties** on the main WebBIOS dialog.
There are four Controller Information dialogs. The following figure shows the first dialog.

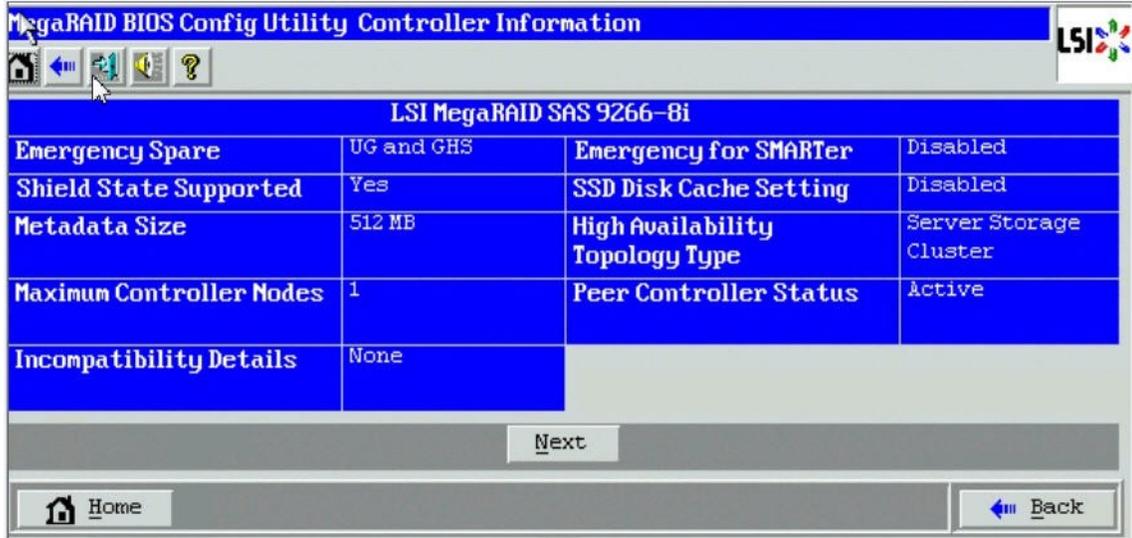
Figure 32 First Controller Information Dialog

LSI MegaRAID SAS 9266-8i			
Serial Number	SR12000571	FRU	021
SubVendorID	0x1000	Drive Security Capable	Yes
SubDeviceID	0x9266	Drive Security Enabled	Yes
PortCount	8	Drive Security Method	FDE Only
HostInterface	PCIE	NVRAMSize	32 KB
Firmware Version	3.190.05-1652	Memory Size	1024 MB
FW Package Version	23.7.0-0021	Min Strip Size	8 KB
Firmware Time	May 11 2012;17:43:04	Max Strip Size	1 MB
WebBIOS Version	6.1-45-Rel	Virtual Drive Count	0
Drive Count	24	MegaRAID Recovery	Enabled
Hot Spare Spin Down	Enabled	Unconf Good Spin Down	Enabled
Spin Down Time	30 minutes	Chip Temp	86 C

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

- Click **Next** to view the second Controller information dialog, as shown in the following figure.

Figure 33 Second Controller Information Dialog



NOTE If you are using CacheCade Pro 2.0, four additional fields appear in the **Second Controller Information** dialog – **CacheCade SSD Caching**, **Write Cache Capable**, **Total Cache Size**, and **Maximum Cache Size**.

NOTE If High Availability DAS is supported on the controller, four additional fields appear in the above dialog - **Maximum Controller Nodes**, **Incompatibility Details**, **High Availability Topology Type**, and **Peer Controller Status**.

- Click **Next** in the **Second Controller Information** dialog to view the third Controller information dialog, as shown in the following figure.

Figure 34 Third Controller Properties

The screenshot shows the 'MegaRAID BIOS Config Utility Controller Properties' dialog box. The title bar includes the LSI logo. Below the title bar is a navigation toolbar with icons for Home, Back, Forward, and Help. The main area is titled 'Properties' and contains a table of configuration options:

Battery Backup	Manage	Coercion Mode	None
Set Factory Defaults	No	S.M.A.R.T Polling	300 seconds
Cluster Mode	Disabled	Alarm Control	Enabled
Rebuild Rate	32	Patrol Read Rate	32
BGI Rate	32	Cache Flush Interval	5
CC Rate	33	Spinup Drive Count	4
Reconstruction Rate	31	Spinup Delay	13
NCQ	Enabled	Spin Down Delay Time	30 minutes

At the bottom of the dialog, there are buttons for 'Submit', 'Reset', and 'Next'. A 'Home' button is located in the bottom-left corner, and a 'Back' button is in the bottom-right corner.

- Click **Next** to view the fourth Controller Information dialog, as shown in the following figure.

Figure 35 Fourth Controller Properties Dialog

The screenshot shows the 'MegaRAID BIOS Config Utility Controller Properties' dialog box for the fourth controller. The title bar includes the LSI logo. Below the title bar is a navigation toolbar with icons for Home, Back, Forward, and Help. The main area is titled 'Properties' and contains a table of configuration options:

Stop CC On Error	No	Schedule CC	Supported
Maintain PD Fail History	Enabled	StopOnError	Enabled
Controller BIOS	Enabled	Disk Activity	Disabled
Emergency Spare	UG and GHS	Emergency for SMARTer	Disabled
Link Speed	Manage	Data Protection	Enabled
Drive Security	Change/Disable	SSD Caching	Manage

At the bottom of the dialog, there are buttons for 'Submit' and 'Reset'. A 'Home' button is located in the bottom-left corner, and a 'Back' button is in the bottom-right corner.

NOTE

If you are using CacheCade Pro 2.0, an additional field, **SSD Caching**

appears in the **Controller Properties** dialog.

NOTE If you have already enabled drive security, instead of the **Enable** link, the **Change/Disable** link appears in front of **Drive Security**.

The entries and options that appear in the second, third and fourth Controller Information dialogs are in the [Controller Information Menu Options](#).

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

2.5.1.1 Controller Information Menu Options

The following table describes the entries and options listed on the second, third, and fourth Controller Information dialogs. Leave these options at their default settings to achieve the best performance, unless you have a specific reason for changing them.

Table 19 Controller Information Menu Options

Option	Description
Battery Backup	Indicates if the selected controller has a BBU. If present, you can click Manage to view information about the BBU. For more information, see Viewing and Changing Battery Backup Unit Information .
Set Factory Defaults	Use this option to load the default MegaRAID WebBIOS configuration utility settings. The default is No .
Cluster Mode	Use this option to enable or disable Cluster mode. The default is Disabled . A cluster is a grouping of independent servers that can access the same data storage and provide services to a common set of clients. When Cluster mode is disabled, the system operates in Standard mode.
Rebuild Rate	Use this option to select the rebuild rate for drives connected to the selected controller. The default is 30 percent. The rebuild rate is the percentage of system resources dedicated to rebuilding a failed drive. The higher the number, the more system resources that are devoted to a rebuild.
BGI Rate	Use this option to select the amount of system resources dedicated to background initialization of virtual drives connected to the selected controller. The default is 30 percent.
CC Rate	Use this option to select the amount of system resources dedicated to consistency checks of virtual drives connected to the selected controller. The default is 30 percent.
Reconstruction Rate	Use this option to select the amount of system resources dedicated to reconstruction of drives connected to the selected controller. The default is 30 percent.
Controller BIOS	Use this option to enable or disable the BIOS for the selected controller. The default is Enabled . If the boot device is on the selected controller, the BIOS must be enabled; otherwise, the BIOS should be
NCQ	Native Command Queuing (NCQ) gives an individual drive the ability to optimize the order in which it executes the read and write commands. The default is Enabled .
Coercion Mode	Drive coercion forces drives of varying capacities to the same size so they can be used in a drive group. The coercion mode options are None , 128MB-way , and 1GB-way . The default is 1GB-way . The number you choose depends on how much the drives from various vendors vary in their actual size. Use the 1GB coercion mode option.

Option	Description
S.M.A.R.T. Polling	Use this option to determine how frequently the controller polls for drives reporting a predictive drive failure (self-monitoring analysis and reporting technology [SMART] error). The default is 300 seconds (5 minutes).
Alarm Control	Select this option to enable, disable, or silence the onboard alarm tone generator on the controller. The default is Enabled .
Patrol Read Rate	Use this option to select the rate for patrol reads for drives connected to the selected controller. The default is 30 percent. The patrol read rate is the percentage of system resources dedicated to running a patrol read.
Cache Flush Interval	Use this option to control the interval (in seconds) at which the contents of the onboard data cache are flushed. The default is 4 seconds.
Spinup Drive Count	Use this option to control the number of drives that spin up simultaneously. The default is 4 drives.
Spinup Delay	Use this option to control the interval (in seconds) between spin up of drives connected to this controller. The delay prevents a drain on the system's power supply that would occur if all drives spun up at the same time. The default is 12 seconds.
StopOnError	Enable this option if you want the boot process to stop when the controller BIOS encounters an error during boot-up. The default is Enabled .
Stop CC on Error	Enable this option if you want to stop a consistency check when the controller BIOS encounters an error. The default is No .
Maintain PD Fail History	Enable this option to maintain the history of all drive failures. This option is used to keep track of drives that the RAID controller believes have failed. With this feature enabled, the RAID controller will track bad drives and mark them as Unconfigured bad if they return from disconnect or failure. Drives can be marked Unconfigured bad if they are failing or if the RAID controller loses communication with the drive while it is part of a configuration (a virtual drive member or a hot spare). The HBA will lose communication with drives if they are removed while the system is turned on or if SIMs are removed while the system is turned on. The default is Enabled .
Schedule CC	Indicates whether the option to schedule the date and time for a consistency check is supported.
Snapshot	Use this option to create a snapshot of a volume. MegaRAID Recovery, also known as Snapshot, offers a simplified way to recover data and provides automatic protection for the boot volume. You can use the Recovery feature to take a snapshot of a volume and to restore a volume or file. Snapshot functionality allows you to capture data changes to the volume, and, if data is deleted accidentally or maliciously, restore the data from the view or roll back to a snapshot at a previous point-in-time (PIT). MegaRAID Recovery supports up to eight snapshots of PITs for each volume.
Disk Activity	Enable this property if you want to locate a particular disk. This disk can be identified with a continuous blinking of green activity LED. This works only if the disks are installed in an enclosure.
Manage JBOD	Converting the multiple JBOD drives to unconfigured drive at single selection.
Emergency Spare	Use this option to specify if it is acceptable to commission unconfigured good drives or global hotspares as emergency spare drives.
Emergency for SMARTer	Use this option to specify if it is acceptable to commission emergency hot spare drives for predictive failure analysis (PFA) events.
Drive Security	Use this option to encrypt data on the drives and use disk-based key management for the data security solution. This solution protects your data in case of theft or loss of physical drives.
Manage Powersave	Use this option to reduce the power consumption of drives that are not in use, by spinning down the unconfigured drives, hot spares, and configured drives.
Link Speed	Use this option to change the link speed between the controller and an expander or between the controller and a drive that is directly connected to the controller.

Option	Description
SSD Caching	Click on this link to invoke the Manage SSD Caching dialog to enable and disable SSD caching on <u>multiple virtual drives at one time</u> .
Maximum Controller Nodes	Indicates the total number of servers in a cluster.
Incompatibility Details	Indicates the reason for incompatibility between the servers in a cluster. The possible values are FW Level Mismatch, HW Incompatibility, Controller Property Mismatch, Premium Features Mismatch, or None.
High Availability Topology Type	Indicates whether clustering is supported or not on the controller. The possible values are Server Storage Cluster or None .
Peer Controller Status	Indicates if both the servers in a cluster are running or not. The possible values are: <ul style="list-style-type: none"> ■ Active— Both the servers in a cluster are running. ■ Inactive— Only one server in a cluster is running. ■ Incompatible— There is incompatibility between the servers.

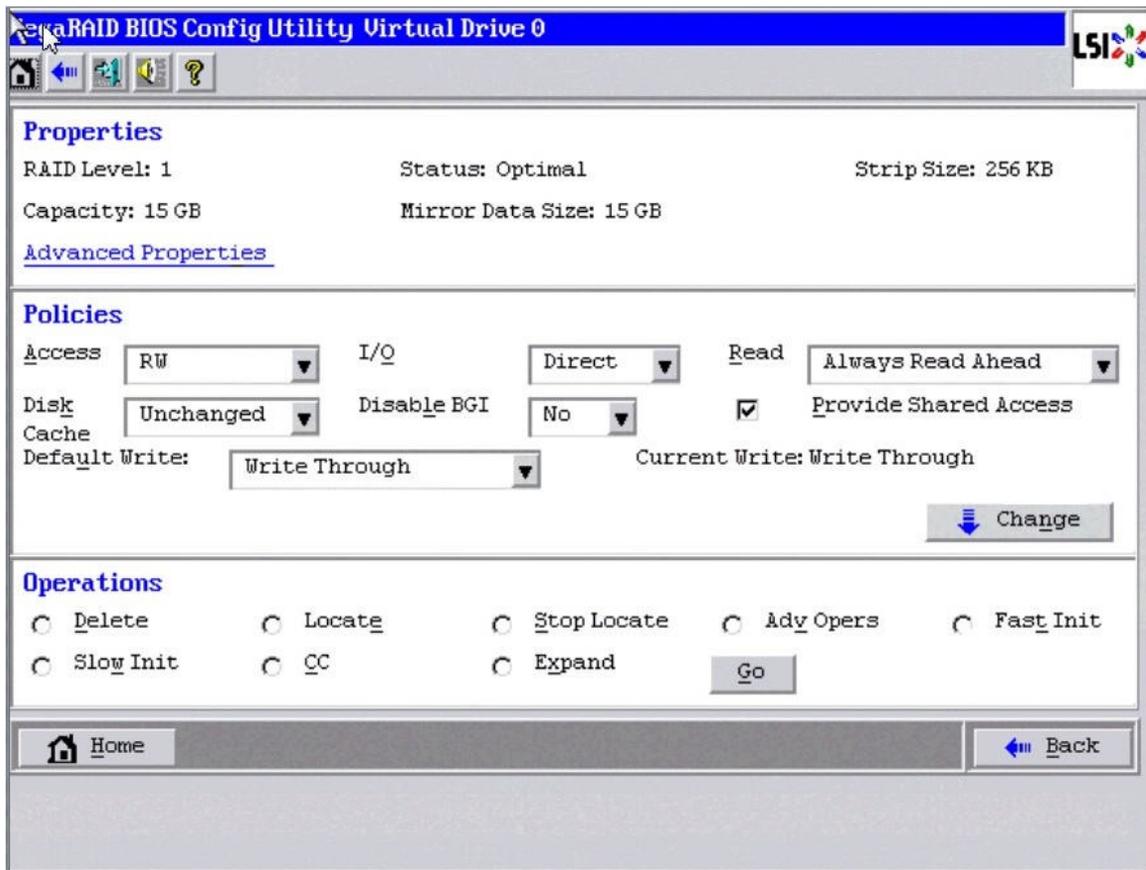
2.5.2 Viewing Virtual Drive Properties, Policies, and Operations

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

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

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

Figure 36 Virtual Drive Dialog



The Properties panel of this dialog displays the virtual drive's RAID level, state, capacity, strip size, and metadata size.

NOTE If High Availability DAS is supported on the controller, click **Advanced Properties** to display two additional properties, **GUID** and **Host Access Policy**. The value for the **GUID** property is hexadecimal and the values for the **Host Access Policy** property are **Shared**, **Exclusive**, or **Exclusive to Peer Controller**.

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

NOTE The **Provide Shared Access** check box appears if High Availability DAS is supported on the controller.

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

- Select **Delete** to delete this virtual drive. For more information, see [Deleting a Virtual Drive](#).
- Select **Locate** to make the LEDs blink on the drives used by this virtual drive. This action works only if the drives are installed in a drive enclosure that supports SCSI-Accessed-Fault-Tolerant-Enclosure (SAFTE).
- Select **Fast Init** or **Slow Init** to initialize this virtual drive. A fast initialization quickly writes zeroes to the first and last 10-MB regions of the new virtual drive and then completes the initialization in the background. A slow initialization is not complete until the entire virtual drive has been initialized with zeroes. It is seldom necessary to use this option, because the virtual drive was already initialized when you created it.

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

- Select **CC** to run a consistency check on this virtual drive. For more information, see [Running a Consistency Check](#). (This option is not available for RAID 0 virtual drives.)
- Select **Stop Locate** to stop the LED flash on the drive. This works only if the drive is installed in a drive enclosure.
- Select **Adv Opers** to access dialogs to remove drives, migrate RAID levels (that is, change the virtual drive configuration by adding a drive and changing the RAID level), virtual drive erase, enable/disable SSD Caching, and to remove blocked access.
See [Migrating the RAID Level of a Virtual Drive](#), for information about adding a drive to a virtual drive or migrating its RAID level. See [Using MegaRAID Recovery](#), for the MegaRAID Recovery procedure.
- Select **Expand** to increase the size of a virtual drive to occupy the remaining capacity in the drive group.
See [Viewing and Expanding a Virtual Drive](#), for the procedure you can use to expand a virtual drive.

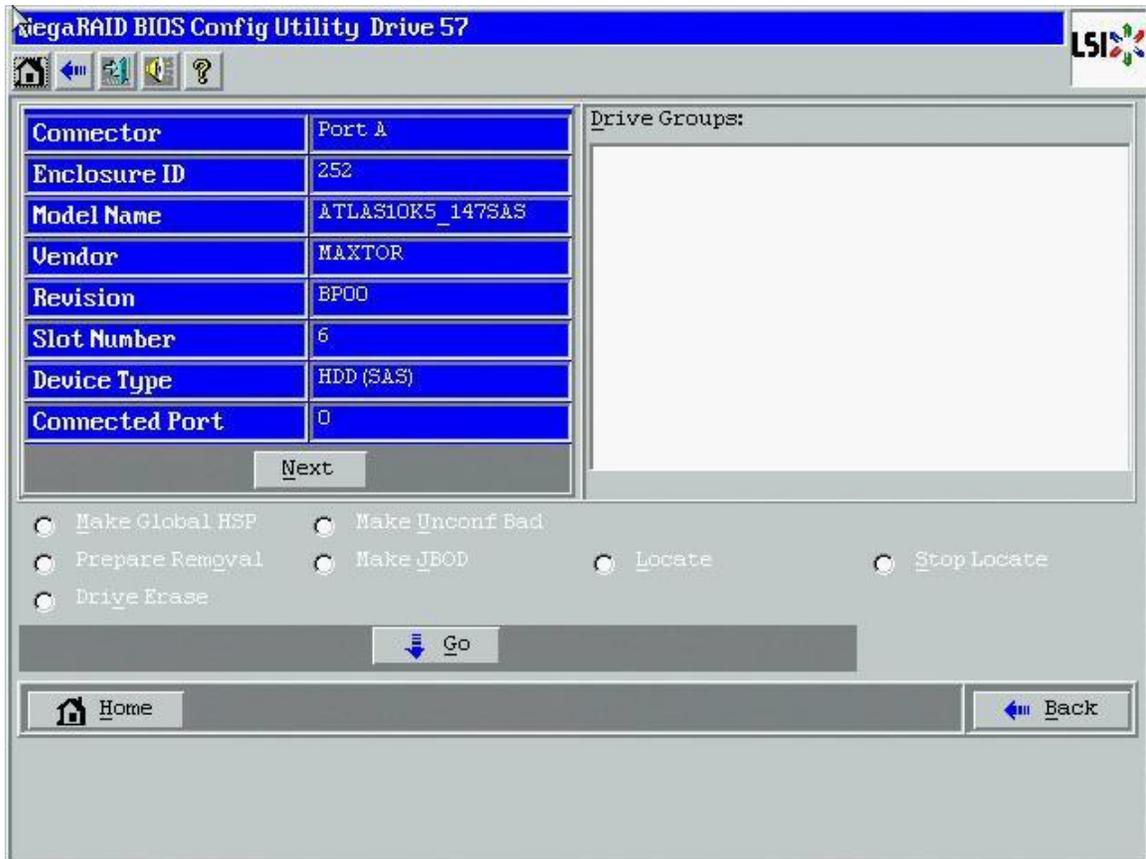
2.5.3 Viewing Drive Properties

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

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

The following figure shows the **Physical Drive** dialog.

Figure 37 Physical Drive Dialog



The drive properties are read-only and are self-explanatory. Note that the properties include the state of the drive. Operations you can perform are listed at the bottom of the dialog. After you select an operation, click **Go** to start the operation. The operations vary depending on the drive state. If the drive state is Online, the following operations appear.

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

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

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

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

- Select **Make Global HSP** to make a global hot spare, which is available to all of the virtual drives.
- Select **Make Dedicated HSP** to make a hot spare dedicated to a specific virtual drive.
WebBIOS displays the global hot spare as `Global` and the dedicated hot spare as `Ded`. The icon for the dedicated hot spare appears under its associated virtual drive. The drive number, drive state, drive capacity, and drive manufacturer appear.

NOTE If High Availability DAS is supported on the controller, dedicated hot spares can be assigned to only one drive group. If you try to assign a dedicated hot spare to multiple drive groups, an error message appears.

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

2.5.4 Shield State

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

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

2.5.4.1 Shield State Physical View

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

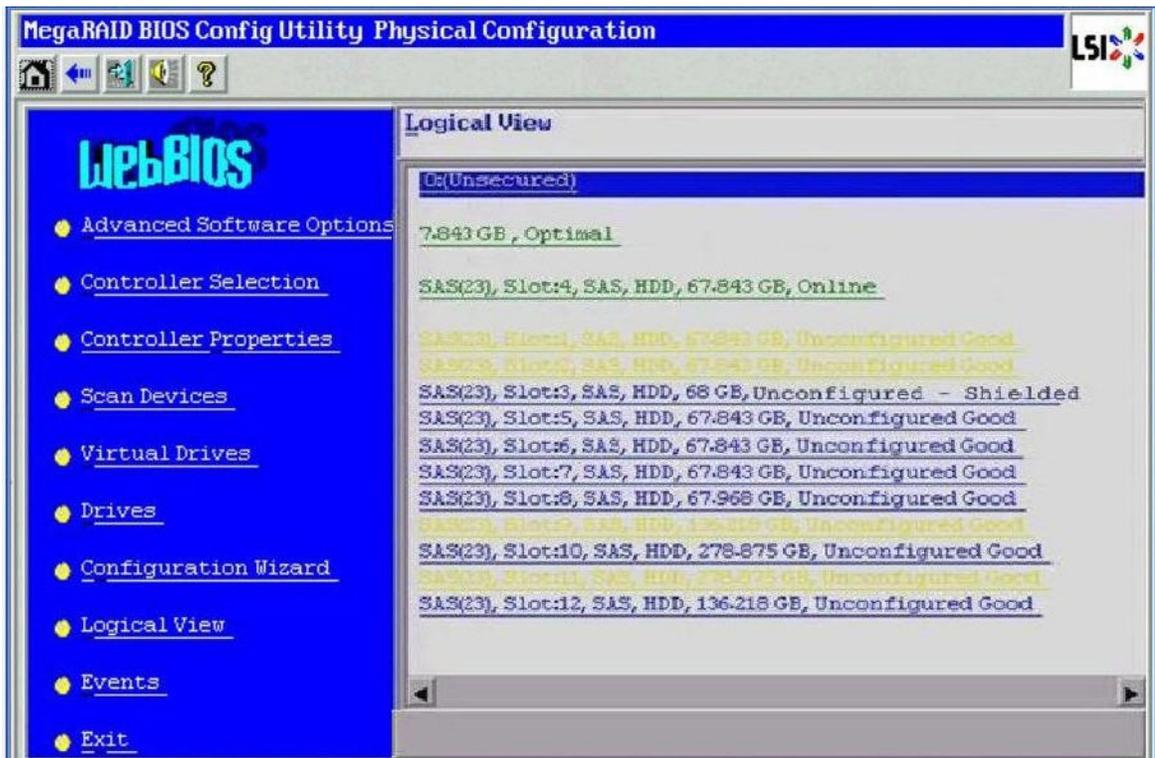
1. Click **Physical View** in the main dialog.
The physical drive that is in a shield state is marked as Shielded.

2.5.4.2 Logical View Shield State

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

1. Click **Logical View** in the main page.
The physical drive that is in a shield state is marked as Shielded.
The Logical view shield state is shown in the following figure.

Figure 38 Logical View Shield State

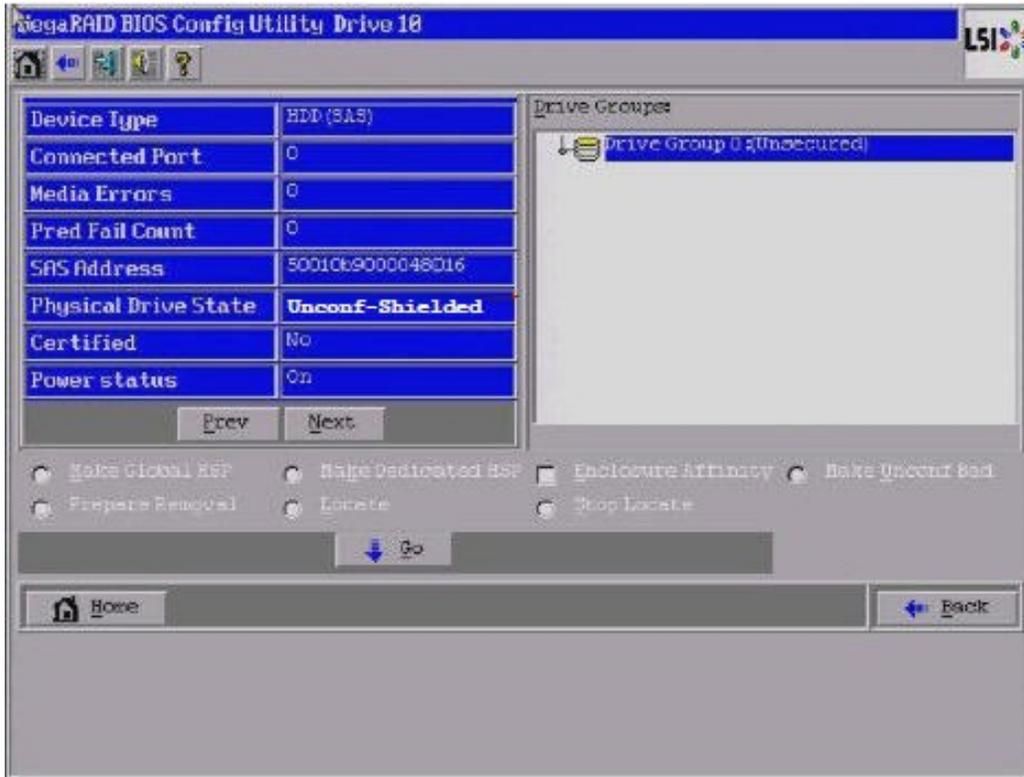


2.5.4.3 Viewing the Physical Drive Properties of a Drive in Shield State

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

1. Click on the **Physical view** tab or the **Logical view** tab in the device tree.
2. Click the physical drive that is in shield state on the physical or logical view of device tree to view the properties.
The device properties of the drive are displayed as shown in the following figure.

Figure 39 Physical Drive Properties of a Drive in Shield State

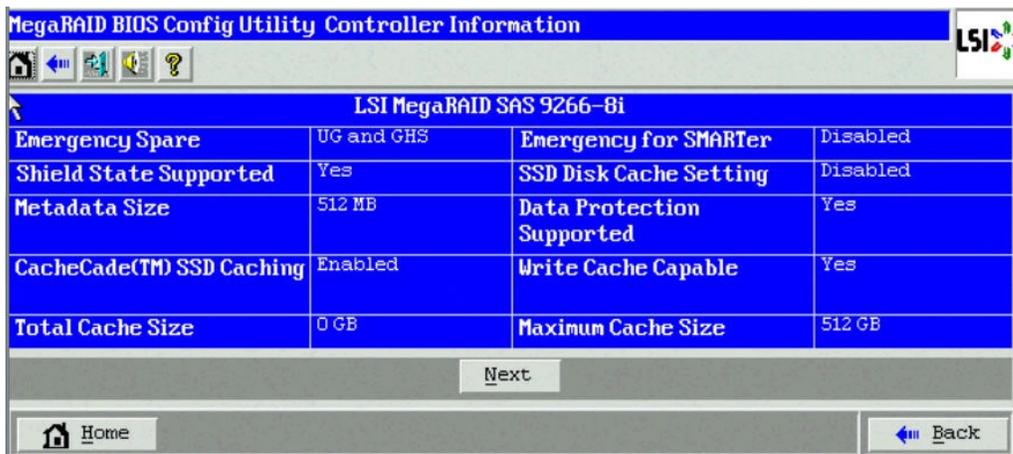


2.5.4.4 Viewing if Shield State Is Enabled in a Controller

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

1. Click **Controller Properties** on the WebBIOS main menu.
The **Shield State Supported** column is displayed, as shown in the following figure.

Figure 40 Shield State Support



2.5.5 SSD Disk Cache Policy

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

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

2.5.5.1 Viewing Cache Properties

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

1. Click the controller properties link in the main menu.
The First Controller Information dialog appears.
2. Click **Next** to view the **SSD Disk Cache Setting** property in the second Controller Properties dialog.

2.5.6 Emergency Spare

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

2.5.6.1 Emergency Spare for Physical Drives

The Emergency Spare property indicates whether the drive is currently commissioned as a emergency spare or not. You can select from the options None, UG (Unconfigured Good), GHS (Global Hotspare), or UG and GHS (Unconfigured Good and Global Hotspare).

Follow these steps to view a emergency spare for a drive.

1. Click the physical drive node in the right panel on the WebBIOS main dialog. The Emergency spare property of the drive is displayed, as shown in the following figure.

Figure 41 Emergency Spare



2.5.7 Emergency Spare for Controllers

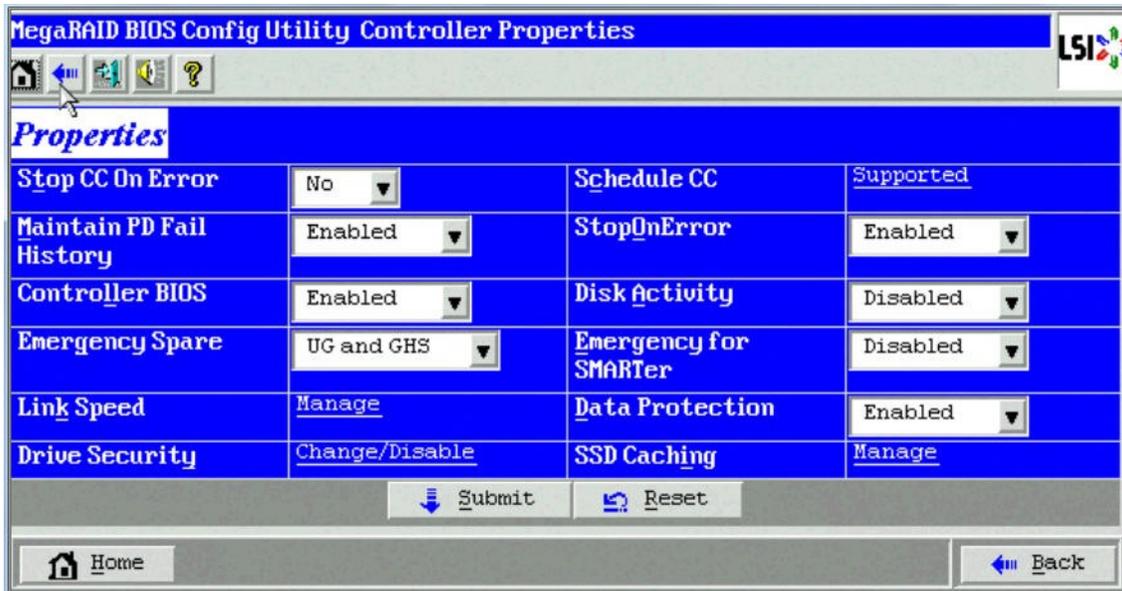
The Emergency Spare properties are configured in the controller properties. You can choose from the four options: **Global Hotspare (GHS)**, **Unconfigured Good (UG)**, **Unconfigured Good and Global Hotspare (UG AND GHS)**, and **None**. You can also enable or disable the **Emergency for SMARTer** property.

2.5.7.1 Setting Controller Emergency Spare Properties

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

1. From the WebBIOS main menu, click **Controller Properties**.
2. Keep clicking **Next** till you reach the last controller properties page.
The controller properties dialog appears, as shown in the following figure. You can choose the options (None, UG, GHS, and UG and GHS) from the **Emergency Spare** drop down list.

Figure 42 Setting Controller Hotspare Properties



MegaRAID BIOS Config Utility Controller Properties			
Properties		LSI	
Stop CC On Error	No	Schedule CC	Supported
Maintain PD Fail History	Enabled	StopOnError	Enabled
Controller BIOS	Enabled	Disk Activity	Disabled
Emergency Spare	UG and GHS	Emergency for SMARTer	Disabled
Link Speed	Manage	Data Protection	Enabled
Drive Security	Change/Disable	SSD Caching	Manage
Submit		Reset	
Home		Back	

2.5.7.2 Viewing Controller Emergency Spare Properties

Follow these steps to view the controllers' Emergency Spare properties.

1. Click the **Controller properties** link in the WebBIOS main menu.
The First Controller Information dialog appears.
2. Click **Next**.
The second Controller Properties dialog appears.
You can view the controller's emergency spare properties in this dialog.

2.5.7.3 Commissioned Hotspare

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

Click the online physical drive node in the right panel on the WebBIOS main dialog to view the Commissioned Hotspare property.

Figure 43 Commissioned Hotspare

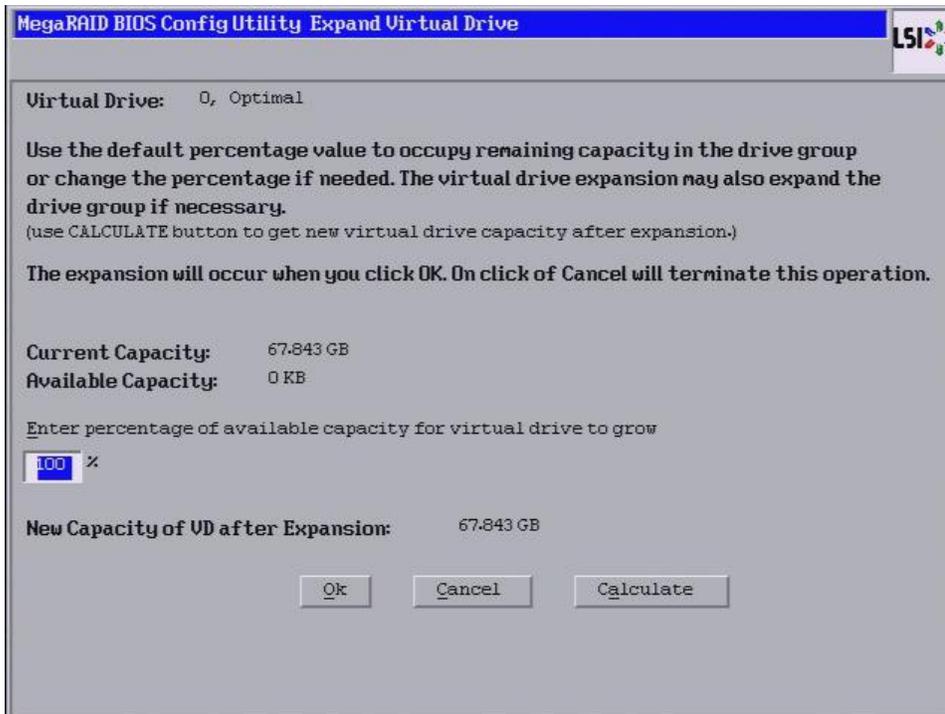


2.6 Viewing and Expanding a Virtual Drive

Follow these steps to view virtual drive properties:

1. In the Logical view of the device tree, click the **Virtual Drive Node**.
The **Virtual Drive** dialog appears.
You can increase the size of a virtual drive to occupy the remaining capacity in a drive group.
2. Select the **Expand** radio button, and click **Go**.
The **Expand Virtual Drive** dialog appears, as shown in the following figure.

Figure 44 Expand Virtual Drive Dialog



3. Enter the percentage of the available capacity that you want the virtual drive to use.
For example, if there are 100 GB of capacity available and you want to increase the size of the virtual drive by 30 GB, select 30 percent.
4. Click **Calculate** to determine the capacity of the virtual drive after expansion.
5. Click **Ok**.
The virtual drive expands by the selected percentage of the available capacity.

2.7 Recovering and Clearing Punctured Block Entries

You can recover and clear the punctured block area of a virtual drive.

ATTENTION This operation removes any data stored on the physical drives. Back up the good data on the drives before making any changes to the configuration.

When a Patrol Read or a Rebuild operation encounters a media error on the source drive, it punctures a block on the target drive to prevent the use of the data with the invalid parity. Any subsequent read operation to the punctured block completes but with an error. Consequently, the puncturing of a block prevents any invalid parity generation later while using this block.

To recover or clear the punctured block area of a virtual drive, run a Slow (or Full) Initialization to zero out and regenerate new parity causing all bad block entries to be removed from the bad block table.

To run a Slow (or Full) Initialization, see [Viewing Virtual Drive Properties, Policies, and Operations](#).

2.8 Suspending and Resuming Virtual Drive Operations

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

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

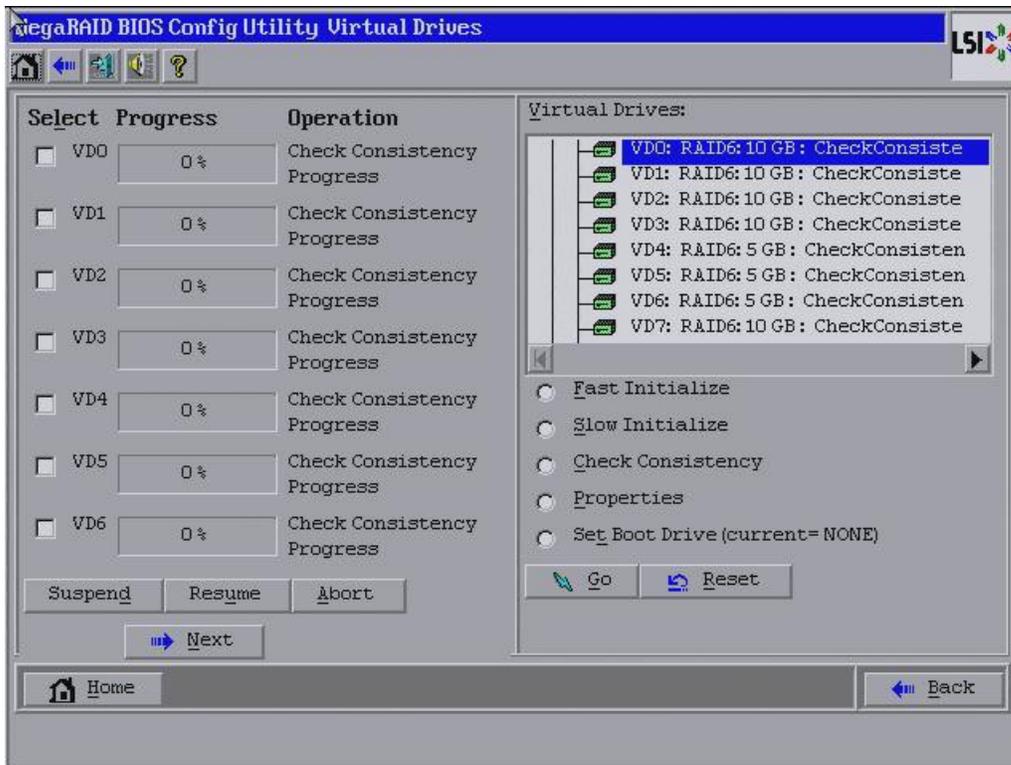
NOTE Suspend and resume are applicable for all the background operations, such as background initialization, rebuild and consistency check notes.

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

1. Perform one of these actions:
 - From the WebBIOS main menu, click the **Virtual Drives** link.
 - From the task bar, click the **VD Progress Info** button.

The **Virtual Drives** main dialog appears, as shown in the following figure.

Figure 45 Virtual Drives Dialog



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

2.9 Non-SED Secure Erase

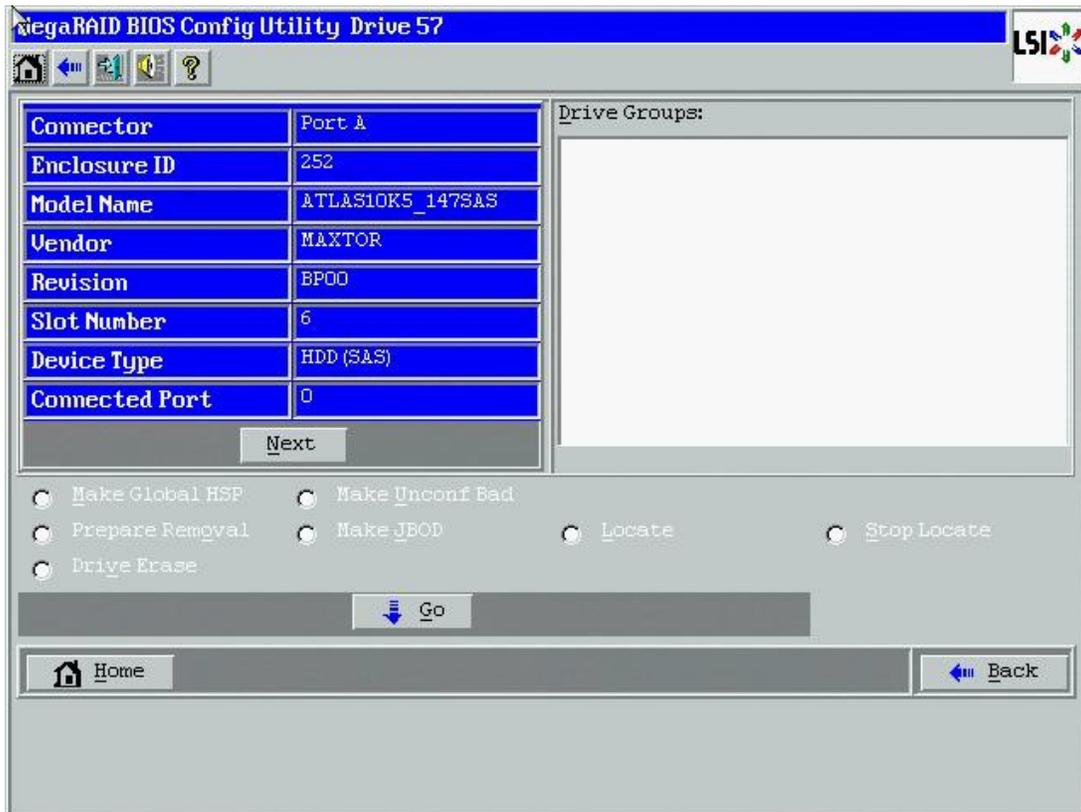
This section describes the procedure used to securely erase data on non self-encrypting drives (Non-SEDs), which are normal HDDs.

2.9.1 Erasing a Non-SED Physical Drive

Follow these steps for non-SED secure erase.

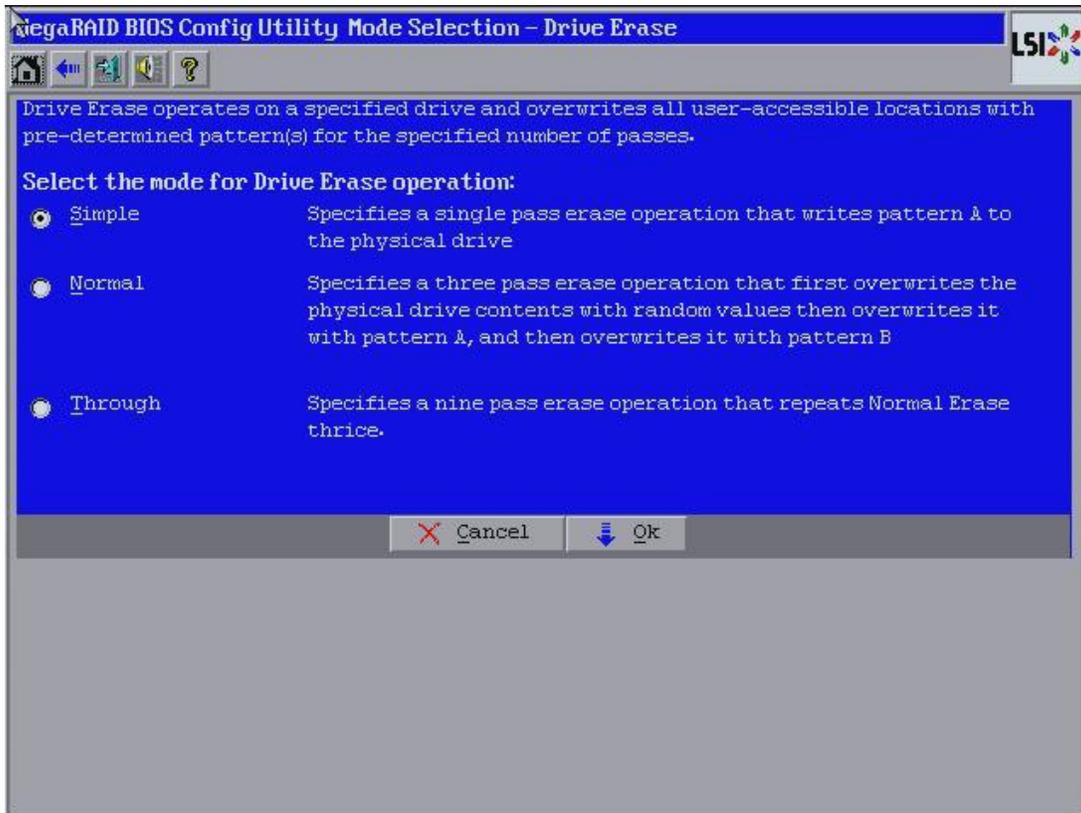
1. Go to the Physical view in the WebBIOS main menu.
2. Click the physical drive node.
3. Select the **Drive Erase** radio button, as shown in the following figure, and click **Go**.

Figure 46 Physical Drive Dialog



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

Figure 47 Mode Selection - Drive Erase



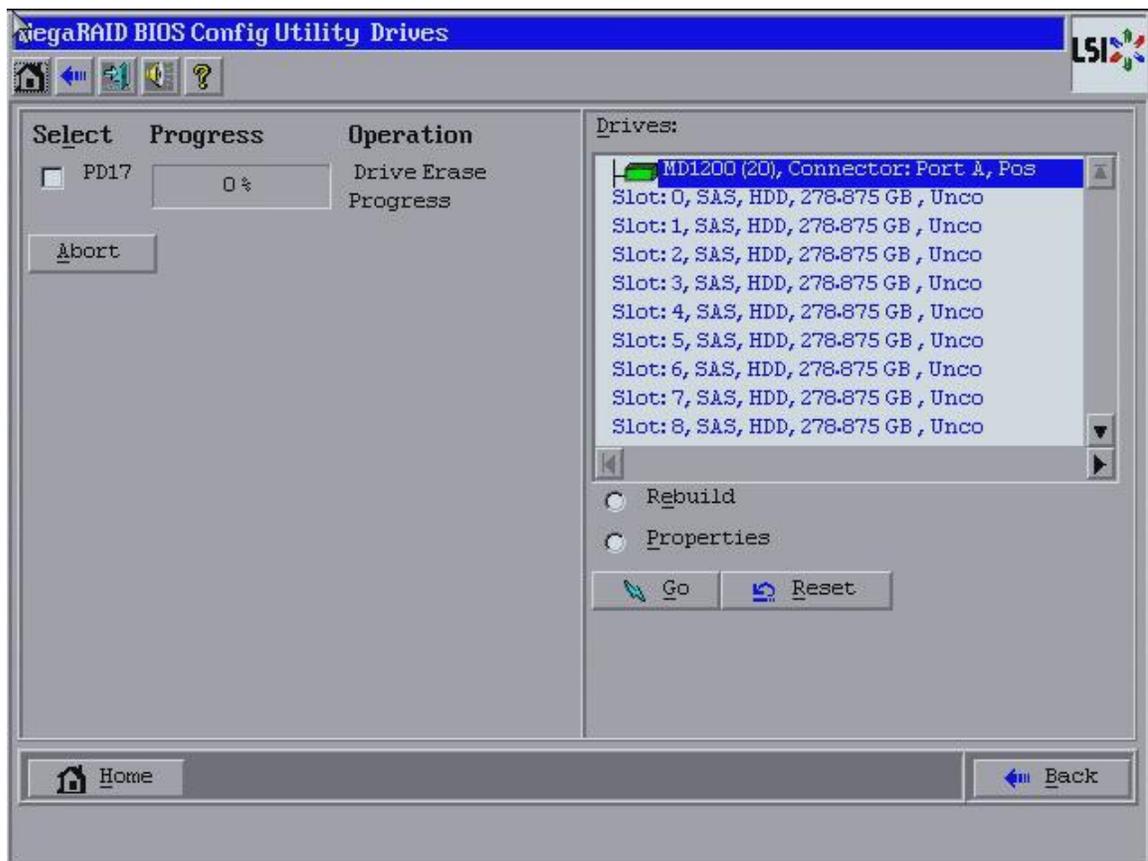
4. Select any of the modes available under the **Select the mode for Drive Erase Operation**
 - **Simple** – (Alt + S)
 - **Normal** – (Alt + N)
 - **Thorough** – (Alt + T)
5. Click **OK**.
A confirmation message dialog appears.
6. Click **Yes** to proceed.

2.9.1.1 Drive Erase Progress

Physical drives, erase operation is generally a time-consuming operation and is performed as a background task. Follow these steps to check the progress of a physical drive erase operation.

1. Click the **Drives** link in the left panel on the WebBIOS main dialog.
The **Drive Erase Progress** appears, as shown in the following figure.

Figure 48 Drive Erase Progress



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

2.9.2 Virtual Drive Erase

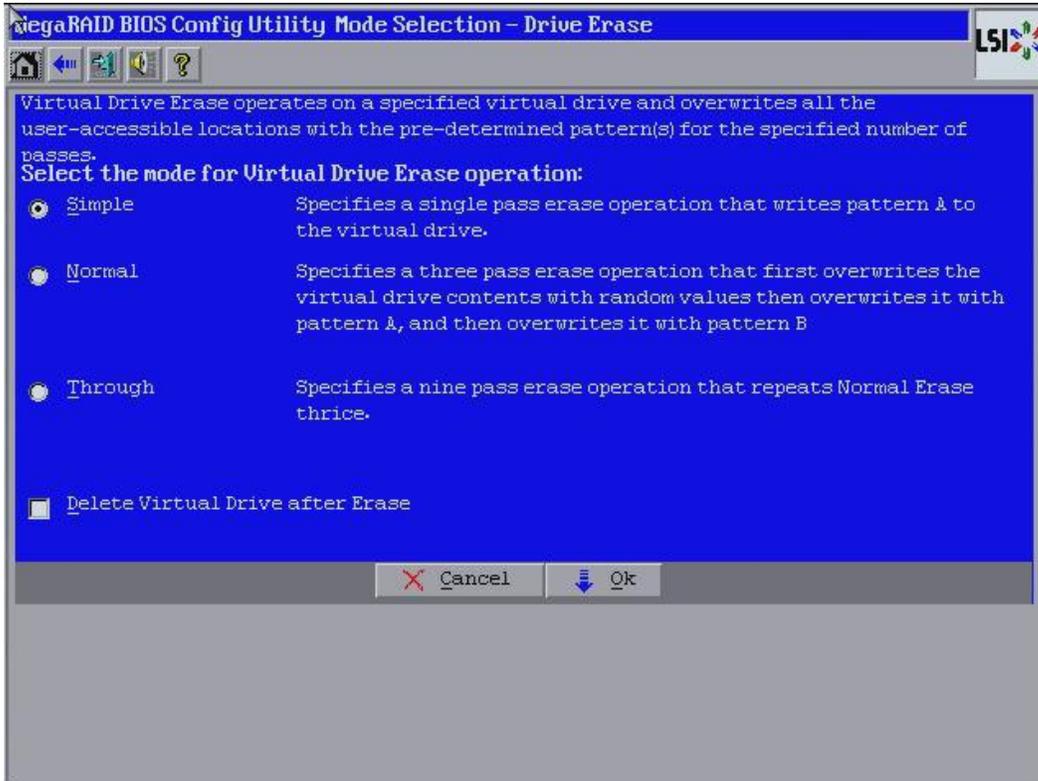
Virtual drive erase is a background operation.

Follow these steps to perform the virtual drive erase operation.

1. Go to the **Logical view**.

2. Click on the Virtual Drive node.
The **Virtual Drive** dialog appears.
3. Click **Adv Opers** and click **Go**.
The **Advanced Operations** dialog appears.
4. Select the **Virtual Drive Erase** radio button, and click **Go**.
The **Mode Selection - Drive Erase** dialog appears.

Figure 49 Mode Selection-Drive Erase



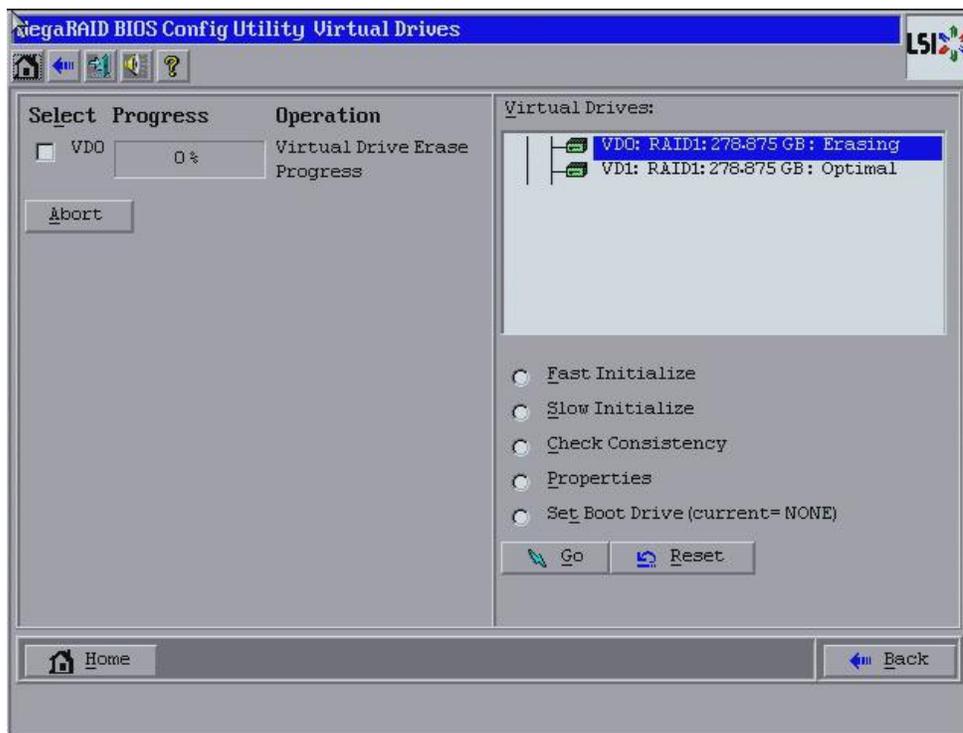
5. Select any of the following options.
 - **Simple** (Alt + S) – After you select this option and click **OK**, if the **Delete Virtual Drive after Erase** check box is selected, a confirmation dialog appears.
 - **Normal** (Alt + N) – After you select this option and click **OK**, if the **Delete Virtual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

2.9.2.1 Group Show Progress for Virtual Drive Erase

The virtual drive erase operation is a time-consuming operation, and it is performed as a background task. Follow these steps to view the progress of virtual drive erase.

1. Click the **Virtual Drives** link on the WebBIOS main menu.
The **Virtual Drives** dialog appears, as shown in the following figure.

Figure 50 Virtual Drive Dialog



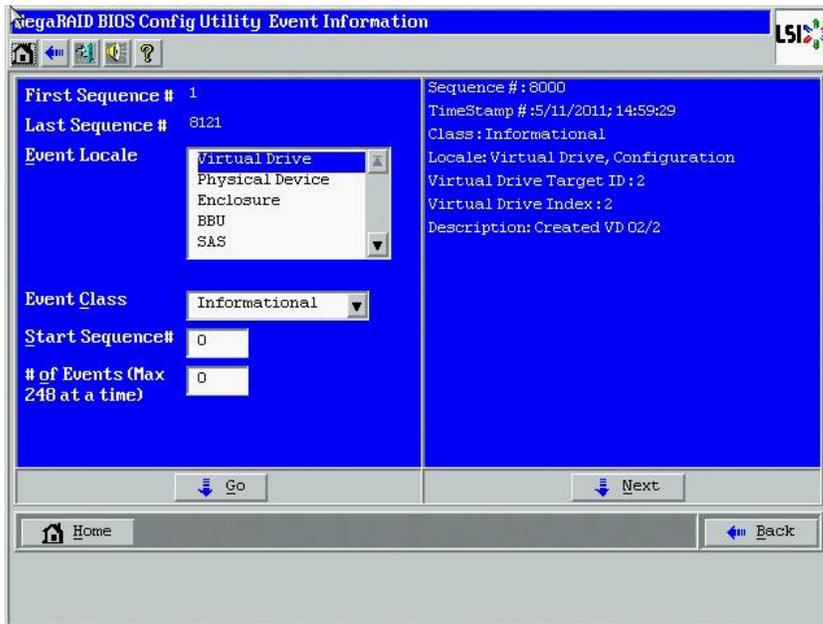
2. To abort the virtual drive erase, select the check box of the operation you want to abort, click **Abort**.

2.10 Viewing System Event Information

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

You can use the WebBIOS configuration utility to view these event messages. To do this, click Events on the main WebBIOS configuration utility dialog. The Event Information dialog appears, as shown in the following figure.

Figure 51 Event Information Dialog



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

To view event information, follow these steps:

1. Select an event locale from the **Event Locale** drop-down list. For example, select **Enclosure** to view events relating to the drive enclosure.
2. Select an event class: **Information, Warning, Critical, Fatal, or Dead**.
3. Enter a start sequence number, between the first sequence and the last sequence numbers. The higher the number, the more recent the event.
4. Enter the number of events of this type that you want to view, and click **Go**. The first event in the sequence appears in the right panel.
5. Click **Next** to page forward or **Prev** to page backward through the sequence of events.
6. Optionally, select different event criteria in the left panel, and click **Go** again to view a different sequence of events.

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

2.11 Managing Configurations

This section includes information about maintaining and managing storage configurations.

2.11.1 Running a Consistency Check

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

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

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

2.11.2 Deleting a Virtual Drive

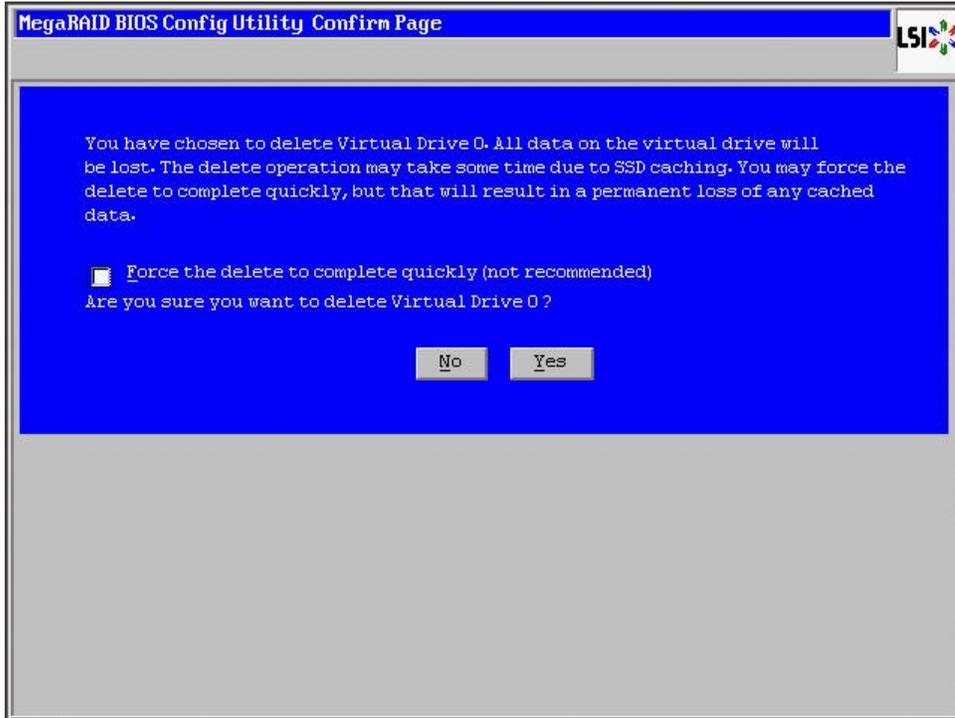
You can delete any virtual drive on the controller if you want to reuse that space for a new virtual drive. The WebBIOS configuration utility provides a list of configurable drive groups where there is a space to configure. If multiple virtual drives are defined on a single drive group, you can delete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

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

To delete a virtual drive, follow these steps.

1. Access the **Virtual Drive** dialog by clicking a virtual drive icon in the right panel on the WebBIOS configuration utility main dialog.
The **Virtual Drive** dialog appears.
2. Select **Delete** in the bottom panel under the heading Operations, and click **Go**.
3. When the message appears, confirm that you want to delete the virtual drive.
If a virtual drive is associated with a CacheCade virtual drive with a write policy, the following confirmation dialog appears. If a virtual drive is not associated with a CacheCade virtual drive, a different confirmation dialog appears.

Figure52 WebBIOS CU Confirmation Dialog



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

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

2.11.3 Importing or Clearing a Foreign Configuration

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

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

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

If WebBIOS configuration utility detects a foreign configuration, the **Foreign Configuration** dialog appears, as shown in the following figure.

Figure 53 Foreign Configuration Dialog

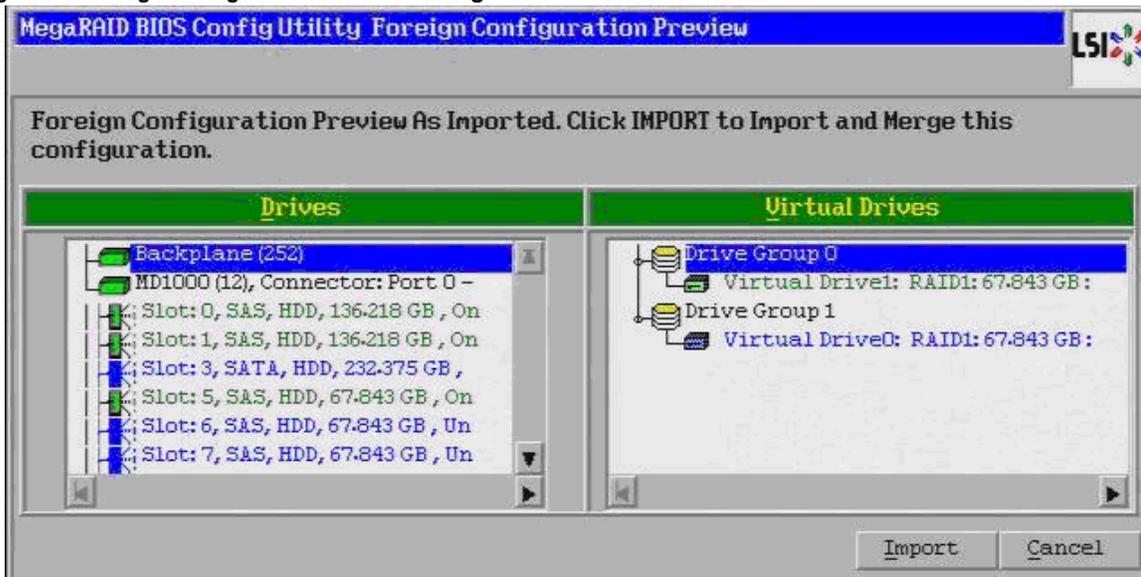


Follow these steps to import or clear a foreign configuration.

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

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

Figure 54 Foreign Configuration Preview Dialog



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

4. Click **Import** to import these foreign configurations and use them on this controller.
If you click **Cancel**, you return to the [Foreign Configuration Dialog](#).

2.11.3.1 Foreign Configurations in Cable Pull and Drive Removal Scenarios

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

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

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

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

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

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

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

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

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

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

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

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

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

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

2.11.3.2 Importing Foreign Configurations from Integrated RAID to MegaRAID

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

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

NOTE For more information about Integrated RAID, refer to the *Integrated RAID for SAS User's Guide*. You can find this document on the LSI website.

2.11.3.3 Troubleshooting Information

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

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

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

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

2.11.4 Importing Foreign Configurations

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

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

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

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

2.11.5 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 restart the system. When you migrate a virtual drive, you can keep the same number of drives, or you can add drives. You can use the WebBIOS configuration utility to migrate the RAID level of an existing virtual drive.

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

Migrations are allowed for the following RAID levels:

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

2.11.5.1 Additional Drives Required for RAID-Level Migration

The following table lists the number of additional drives required when you change the RAID level of a virtual drive.

Table 20 Additional Drives Required for RAID-Level Migration

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

2.11.5.2 Migrating the RAID Level

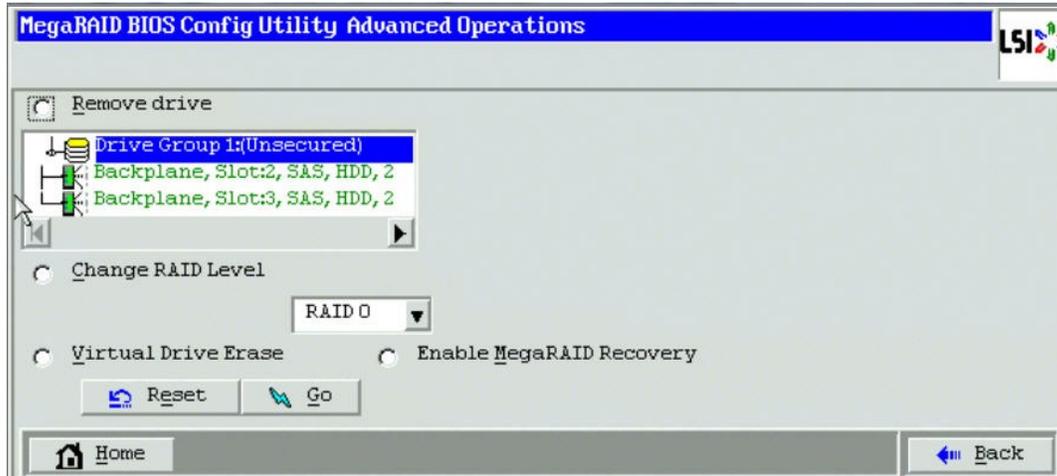
Follow these steps to migrate the RAID level:

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

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

3. From the **Virtual Drives** menu, select **Properties**.
4. From the **Properties** menu, select **Adv Opers** under the **Advanced Operations** heading.
The **Advanced Operations** dialog appears, as shown in the following figure.

Figure 55 Advanced Operations Dialog



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

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

6. Click **Go**.
7. When the message appears, confirm that you want to migrate the RAID level of the virtual drive.

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

2.11.6 New Drives Attached to a MegaRAID Controller

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

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

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

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

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

Chapter 3 MegaRAID Storage Manager Overview

This chapter provides a brief overview of the MegaRAID Storage Manager software.

3.1 Overview

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

3.1.1 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.

3.1.2 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.

3.1.3 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.

3.2 Prerequisites to Running MegaRAID Storage Manager Remote Administration (Not Supported)

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 4 MegaRAID Storage Manager Window and Menus

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

4.1 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. To start the MegaRAID Storage Manager software on a Microsoft Windows operating system, select **Start > Programs > MegaRAID Storage Manager > StartupUI**, or double-click the MegaRAID Storage Manager shortcut on the desktop.

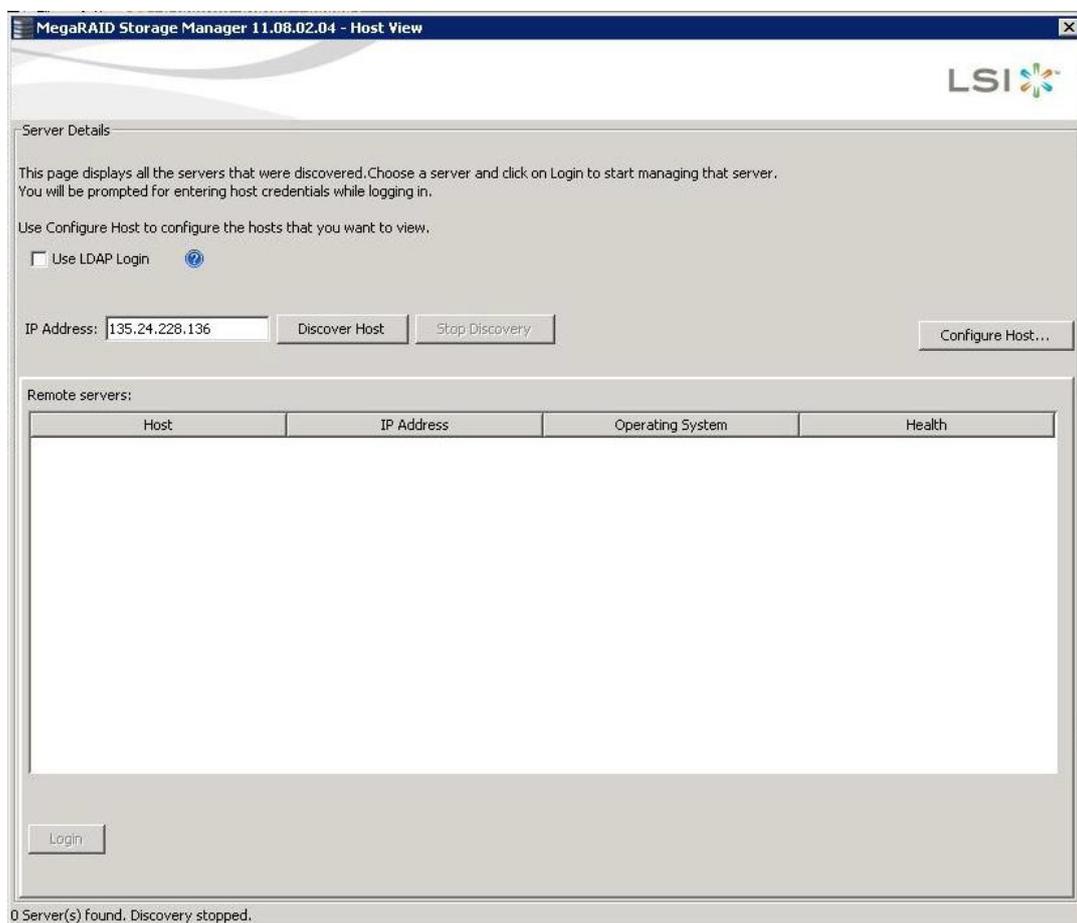
NOTE If a warning appears stating that Windows firewall has blocked some features of the program, click Unblock to allow the MegaRAID Storage Manager software to start. (The Windows firewall sometimes blocks the operation of programs that use Java Technology.)

4.2 Discovery and Login

You can start the MegaRAID Storage Manager software from a remote Windows machine that has the MegaRAID Storage Manager software installed in complete mode. When the program starts, the **Host View** dialog appears, as shown in the following figure. The remote servers are displayed, along with their IP addresses, operating system, and health status.

NOTE If you do a local mode installation, as shown in Section Installing MegaRAID Storage Manager Software on Microsoft Windows, the following figure will not be displayed. It will directly prompt you to the login dialog as shown in the Server Login.

Figure 56 Host View



If High Availability DAS is supported on the controller, instead of the above **Host View** dialog, the **Host View – High Availability DAS** dialog appears, as shown in [High Availability DAS Support](#).

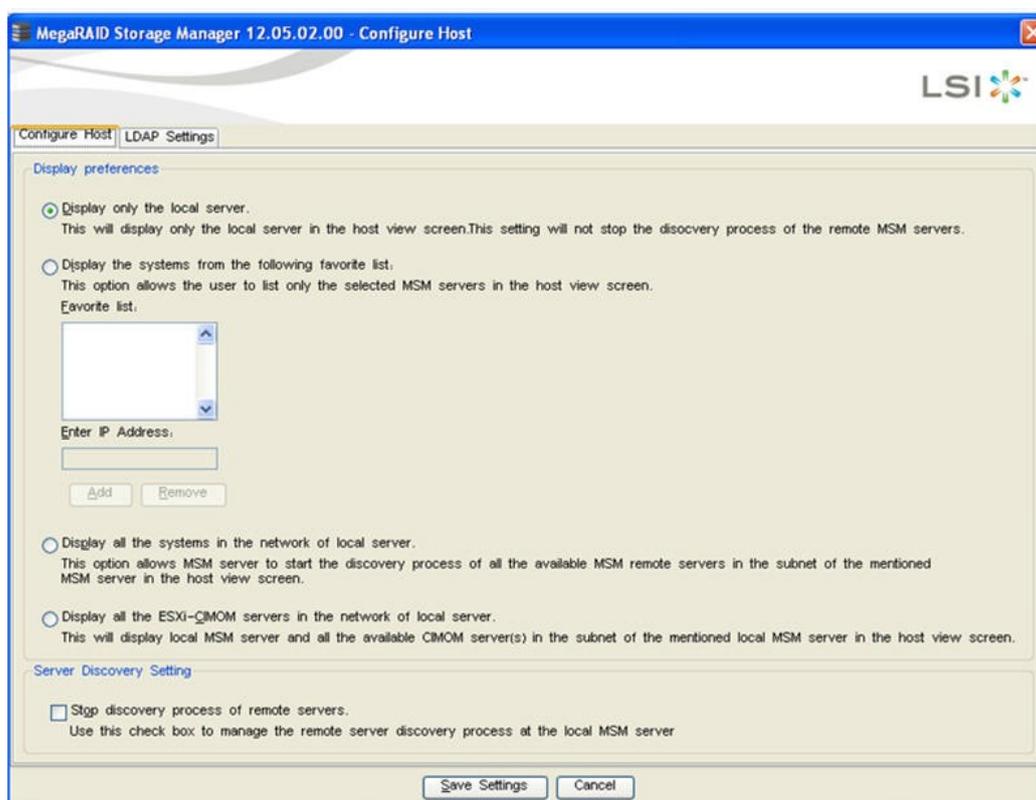
The **Host View** dialog shows an icon for each server on which the MegaRAID Storage Manager software is installed. The servers are color-coded with the following definitions:

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

NOTE Enter a valid MegaRAID Storage Manager server's IP address and select the **Display all the systems in the Network of the local server** option in the following figure.

1. Click **Configure Host** to configure the hosts. The **Configure Host** dialog appears, as shown in the following figure.

Figure 57 Configure Host



The following options are available to configure the host.

- **Display only the local server** – Select this option to display only the Local server or the Server of the IP address entered in the Host View dialog.

- **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 will be able to discover all the ESXi servers.

NOTE If the controller supports High Availability DAS, and you want to view the cluster information in a single pane, select either of the options: **Display the systems from the following favorite list** or **Display all the systems in the Network of the local server**.

2. Click **Save Settings** to save your setting, or on **Cancel** to quit without saving.

If you click **Save Settings**, a confirmation dialog appears asking you to confirm your settings. Click **OK** in the confirmation dialog to start the discovery process.

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

The servers appear in the list of found hosts in the **Host View** dialog.

4. Double-click the icon of the server that you want to access.

The **Server Login** window appears, as shown in the following figure.

Figure 58 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.

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.

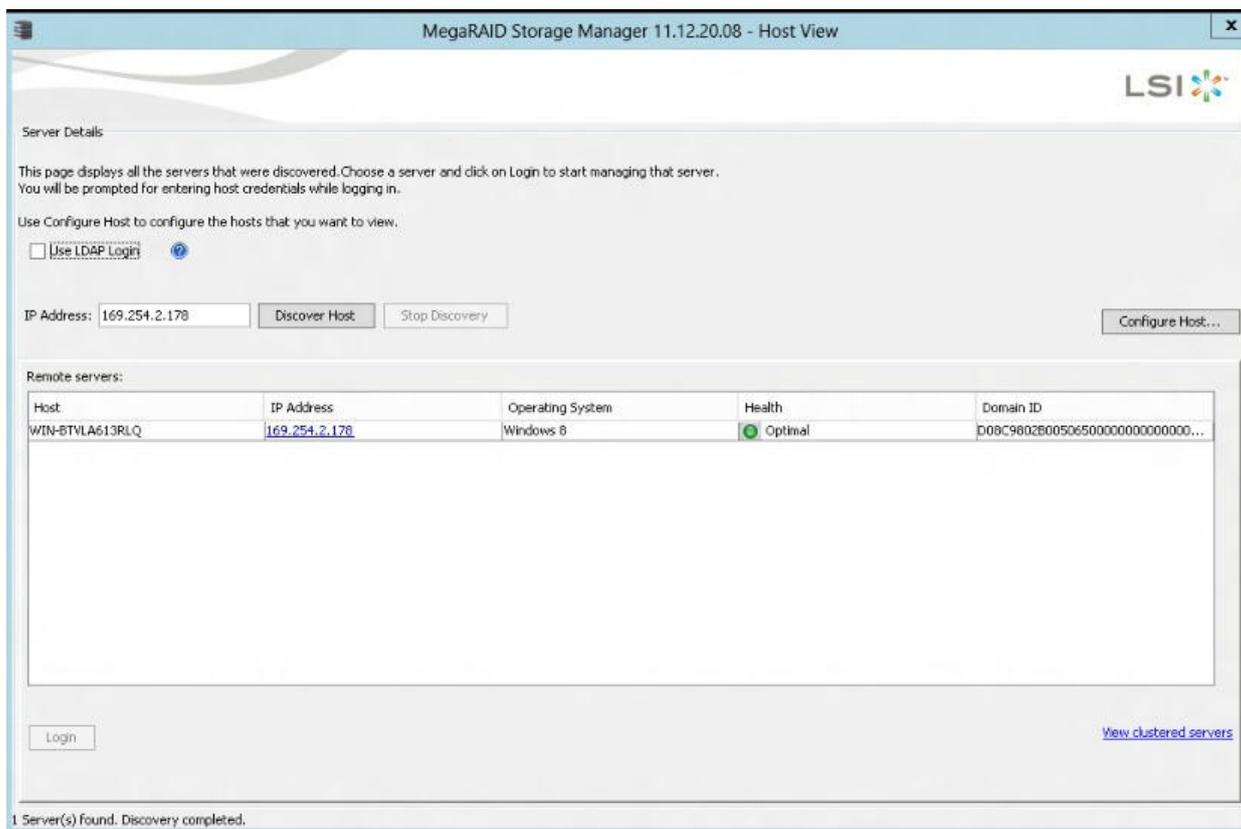
If your user name and password are correct for the Login mode you have chosen, the MegaRAID Storage Manager main menu appears.

4.3 High Availability DAS Support (Not Supported)

NOTE This function is available only when the DAS is attached to the server in use.

If High Availability DAS is supported on the controller when you launch the MegaRAID Storage Manager application, the following dialog appears.

Figure 59 Host View - High Availability DAS



1. Click **View clustered servers** to view all the High Availability cluster servers available.
The **View Clustered Servers** dialog appears, as shown in the following figure.

Figure 60 View Clustered Servers



2. Click on a server link to log into that server.
The **Server Login** window appears.
3. Enter the login details in the **Server Login** window.

4.4 LDAP Support (Not Supported)

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.

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 61 LDAP Login

The screenshot shows the 'LDAP Login' dialog box. The title bar reads 'LDAP Login' and the LSI logo is in the top right. The dialog contains the following fields and controls:

- LDAP Server IP Address:** A text input field.
- User Name:** A text input field with a help icon (question mark) to its right.
- Password:** A text input field.
- Distinguished Name :** A text input field.
- Use Default Port:** A checked checkbox.
- Port:** A text input field containing the value '389'.
- Remember my Login Details:** An unchecked checkbox.
- Login** and **Cancel** buttons at the bottom.

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.

4.5 Configuring LDAP Support Settings (Not Supported)

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 62 Configure Host LDAP

The screenshot shows the 'Configure Host' dialog box with the 'LDAP Settings' tab selected. The dialog contains the following elements:

- Two checkboxes:
 - Use LDAP login as default login mode
 - Do not prompt for credentials when connecting to LDAP
- A 'Server' section with three input fields:
 - IP Address: []
 - Distinguished Name: []
 - Port: []
- A 'Connection' section with two input fields:
 - User Name: []
 - Password: []
- At the bottom, there are two buttons: 'Save Settings' and 'Cancel'.

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.

4.6 MegaRAID Storage Manager Main Menu

This section describes the MegaRAID Storage Manager main menu window:

- [Dashboard / Physical View/ Logical View](#)
- [Properties and Graphical View Tabs](#)
- [Event Log Panel](#)

4.6.1 Dashboard / Physical View/ Logical View

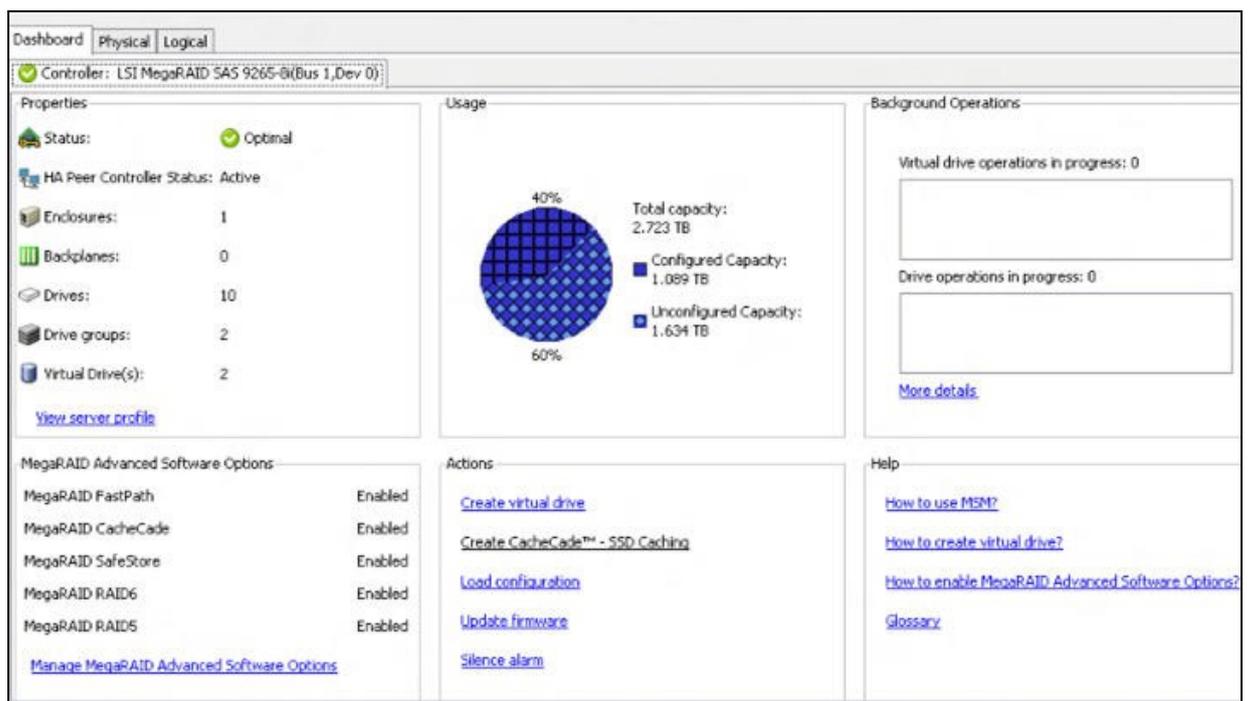
The left panel of the **MegaRAID Storage Manager** window displays the *Dashboard* view, the *Physical* view, or the *Logical* view of the system and the attached devices, depending on which tab is selected.

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 63 MegaRAID Storage Manager Dashboard View



NOTE If the controller supports High Availability DAS, the **HA Peer Controller Status** field appears in the above dialog and displays one of the following values: **Active** (both the servers in the cluster are running), **Inactive** (only one server in the cluster is running), or **Incompatible** (there is incompatibility between the servers).

Physical View

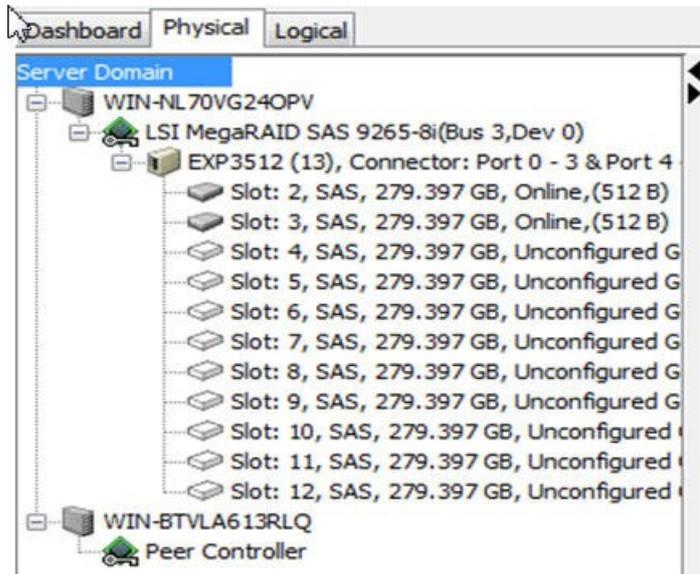
The *Physical* view shows the hierarchy of physical devices in the system. At the top of the hierarchy is the system itself, followed by the controller and the backplane. One or more controllers are installed in the system. The controller label identifies the MegaRAID controller, such as the MegaRAID SAS 9260-8i controller, so that you can easily differentiate between multiple controllers. Each controller has one or more ports. Drives and other devices are attached to the ports. The properties for each item appear in the right panel of the dialog.

Figure 64 MegaRAID Storage Manager Physical View

ID	Error Level	Date / Time	Description
469	[Information, 0]	2010-08-28, 14:00:58	Controller ID: 0 Time established since power on: Time: 2010-08-28, 14:00:58 211533 Seconds
468	[Information, 0]	2010-08-28, 13:00:58	Controller ID: 0 Time established since power on: Time: 2010-08-28, 13:00:58 207939 Seconds
467	[Information, 0]	2010-08-28, 12:00:58	Controller ID: 0 Time established since power on: Time: 2010-08-28, 12:00:58 204333 Seconds
466	[Information, 0]	2010-08-28, 11:00:58	Controller ID: 0 Time established since power on: Time: 2010-08-28, 11:00:58 200732 Seconds
465	[Information, 0]	2010-08-28, 10:00:58	Controller ID: 0 Time established since power on: Time: 2010-08-28, 10:00:58 197132 Seconds
464	[Information, 0]	2010-08-28, 09:00:58	Controller ID: 0 Time established since power on: Time: 2010-08-28, 09:00:58 193532 Seconds
463	[Information, 0]	2010-08-28, 08:00:58	Controller ID: 0 Time established since power on: Time: 2010-08-28, 08:00:58 189932 Seconds
462	[Information, 0]	2010-08-28, 07:00:58	Controller ID: 0 Time established since power on: Time: 2010-08-28, 07:00:58 186332 Seconds
461	[Information, 0]	2010-08-28, 06:00:58	Controller ID: 0 Time established since power on: Time: 2010-08-28, 06:00:58 182732 Seconds
460	[Information, 0]	2010-08-28, 05:00:58	Controller ID: 0 Time established since power on: Time: 2010-08-28, 05:00:58 179132 Seconds
459	[Information, 0]	2010-08-28, 04:00:58	Controller ID: 0 Time established since power on: Time: 2010-08-28, 04:00:58 175532 Seconds
458	[Warning, 1]	2010-08-28, 03:17:19	Controller ID: 0 PD Predictive Failure: Port 0 - 3:1-5
457	[Warning, 1]	2010-08-28, 03:17:19	Controller ID: 0 PD Predictive Failure: Port 0 - 3:1-1
456	[Information, 0]	2010-08-28, 03:00:58	Controller ID: 0 Time established since power on: Time: 2010-08-28, 03:00:58 171932 Seconds
455	[Information, 0]	2010-08-28, 02:48:11	Controller ID: 0 Patrol Read complete
454	[Information, 0]	2010-08-28, 02:00:58	Controller ID: 0 Time established since power on: Time: 2010-08-28, 02:00:58 168332 Seconds
453	[Information, 0]	2010-08-28, 01:21:53	Controller ID: 0 Patrol Read corrected medium error: PD ->4 Location: Qued71e77
452	[Information, 0]	2010-08-28, 01:21:53	Controller ID: 0 Unexpected sense: PD ->4 - Unrecovered read error, CID = 0x2f 0x00 0x0e 0x07 0x1e 0x72 0x00 0x10 0x00 0x00 . Sense = 0x49 0x00 0...

If the controller supports High Availability DAS, an additional parent mode, **Server Domain**, appears on the device tree in the Physical tab, as shown in the following figure.

Figure 65 Physical View for High Availability DAS



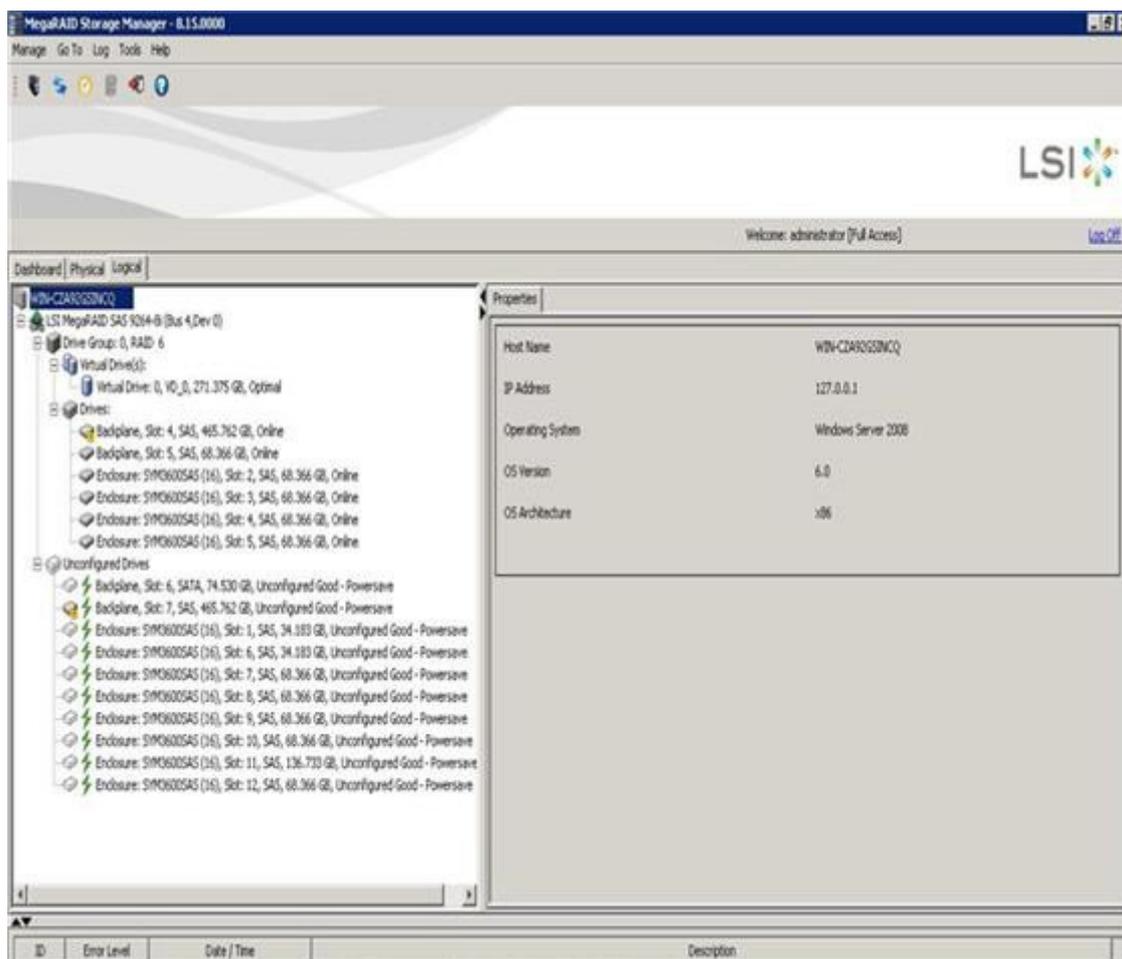
The **Server Domain** is the domain ID of the cluster and shows the two servers that belong to it as child nodes. Information that pertains to the logged-in server in the cluster (such as controller name, enclosures, physical drives) is shown in the Physical tab. For the peer server, no details are shown; the Physical tab just detects that a peer server exists and a controller is attached to it. Right-click **Server Domain** to view the properties of the cluster. A view-only properties dialog appears with two fields; **Domain ID** and **No. of Servers Tagged**.

Logical View

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

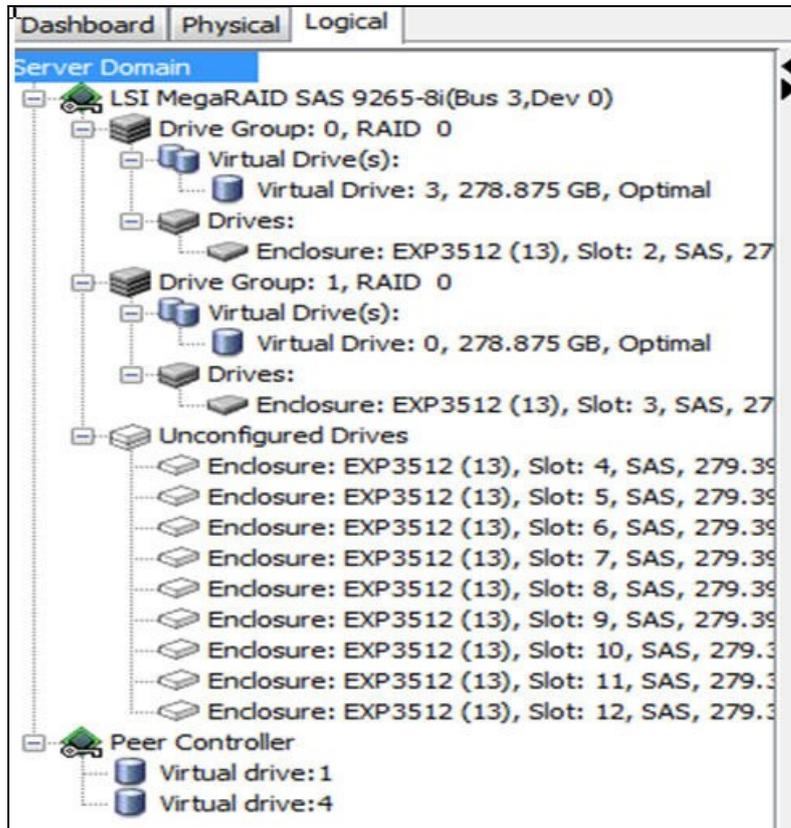
The following figure shows the Logical view.

Figure 66 MegaRAID Storage Manager Logical View



If the controller supports High Availability DAS, an additional parent mode, **Server Domain**, appears on the device tree in the Logical tab, as shown in the following figure.

Figure 67 Logical View for High Availability DAS



The **Server Domain** is the domain ID of the cluster and shows the two servers that belong to it as child nodes. Information that pertains to the logged-in server in the cluster (such as controller name, drive groups, virtual drives) is shown. For the peer server, the Logical tab detects that a peer server exists and a controller is attached to it and shows only the virtual drives created by the peer server. Right-click **Server Domain** to view the properties of the cluster. A view-only properties dialog shows with two fields; **Domain ID** and **No. of Servers Tagged**.

4.6.2 Physical Drive Temperature

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

Figure 68 Physical Drive Temperature

Properties			
General:			
Usable Capacity	927,500 GB	Power Status	Powersave
Raw Capacity	931,513 GB	Revision Level	N506
Logical Sector Size	4 KB	Media Error Count	0
Physical Sector Size	4 KB	Pred Fail Count	0
Certified	No	Enclosure Properties:	
Product ID	ST9100064055	Enclosure ID	252
Vendor ID	SEAGATE	Enclosure Model	Backplane
Serial Number	9XG0192Y	Enclosure Location	External
Device ID	129	Connector	Port 4 - 7
Status	Unconfigured Good	Slot Number	4
Drive Speed	6.0 Gbps	Drive Security Properties:	
Negotiated Link Speed	6.0 Gbps	Full Disk Encryption capable	No
SCSI Device Type	Disk	Secured	No
SAS Address 0	0x5000CS00001A3D4D	Data Protection Properties:	
SAS Address 1	0x0	Data Protection	Incapable
		Shield Counter	0

4.6.3 Shield State

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

Physical devices in MegaRAID firmware transit between different states. If firmware detects a problem or a communication loss for a physical drive, it transitions the physical drive to a bad (FAILED/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).

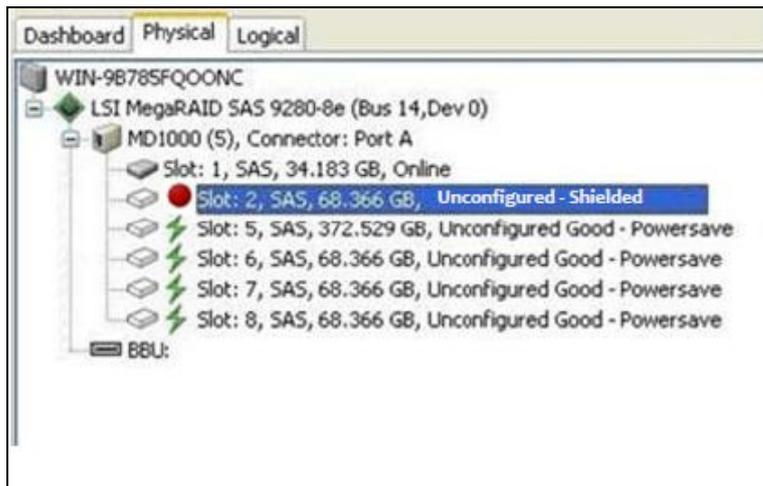
The three possible Shield states are **Unconfigured - Shielded**, **Configured - Shielded**, and **Hotspare - Shielded**.

4.6.4 Shield State Physical View

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

1. Click the **Physical** tab in the device tree. The red dot icon (●) indicates a Shield state.
The Physical View shield state is shown in the following figure.

Figure 69 Physical View Shield State

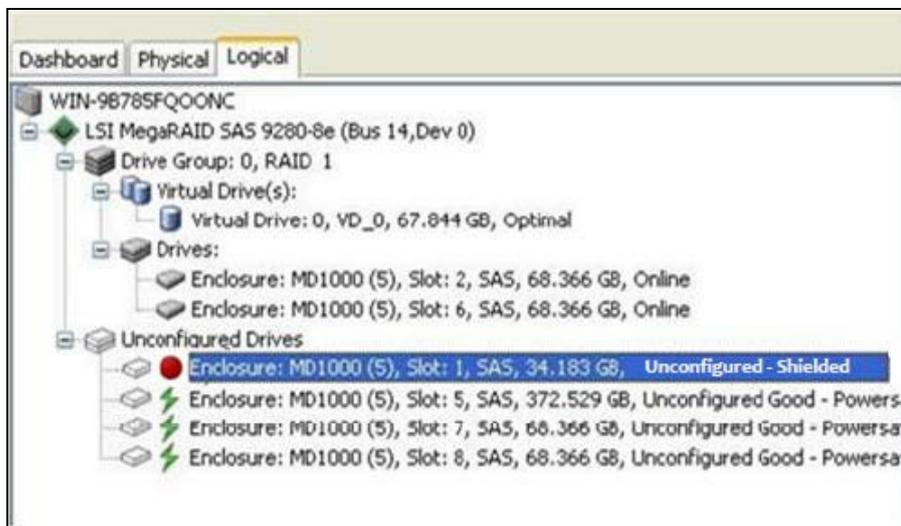


4.6.5 Logical View Shield State

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

1. Click the **Logical** tab in the device tree. The red dot icon (●) indicates a Shield state.
The Logical view Shield state is shown in the following figure.

Figure 70 Logical View Shield State



4.6.6 Viewing the Physical Drive Properties

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

1. Click the **Physical** tab or **Logical** tab in the device tree.

The red dot icon () indicates a Shield state.

2. Click the physical drive in Shield state on Physical view or Logical view of the device tree to view the properties.

The device properties are displayed as shown in the following figure.

Figure 71 Physical Drive Properties of a Drive in Shield State

General:		SAS Address 0	
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	Revision Level	TI35
Certified	No	Media Error Count	0
Product ID	TX43E10100GB0LSI	Pred Fail Count	0
Vendor ID	ATA	Slot Number	4
Serial Number	5L0010ZE	Drive Security Properties:	
Device ID	46	Full Disk Encryption capable	No
Status	Online	Data Protection Properties:	
Drive Speed	6.0 Gbps	Data Protection	Incapable
Negotiated Link Speed	6.0 Gbps	Shield Counter	0
SCSI Device Type	Disk		

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

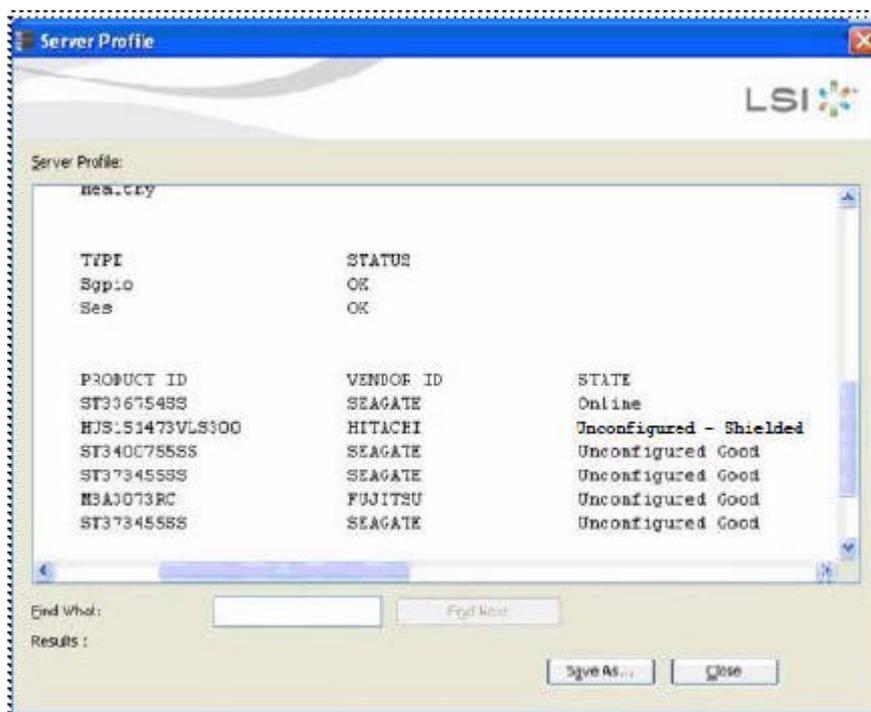
4.6.7 Viewing Server Profile of a Drive in Shield State

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

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

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

Figure 72 View of a Drive in Shield State



4.6.8 Displaying the Virtual Drive Properties

The MegaRAID Storage Manager application displays the following additional virtual drive statistics under controller properties.

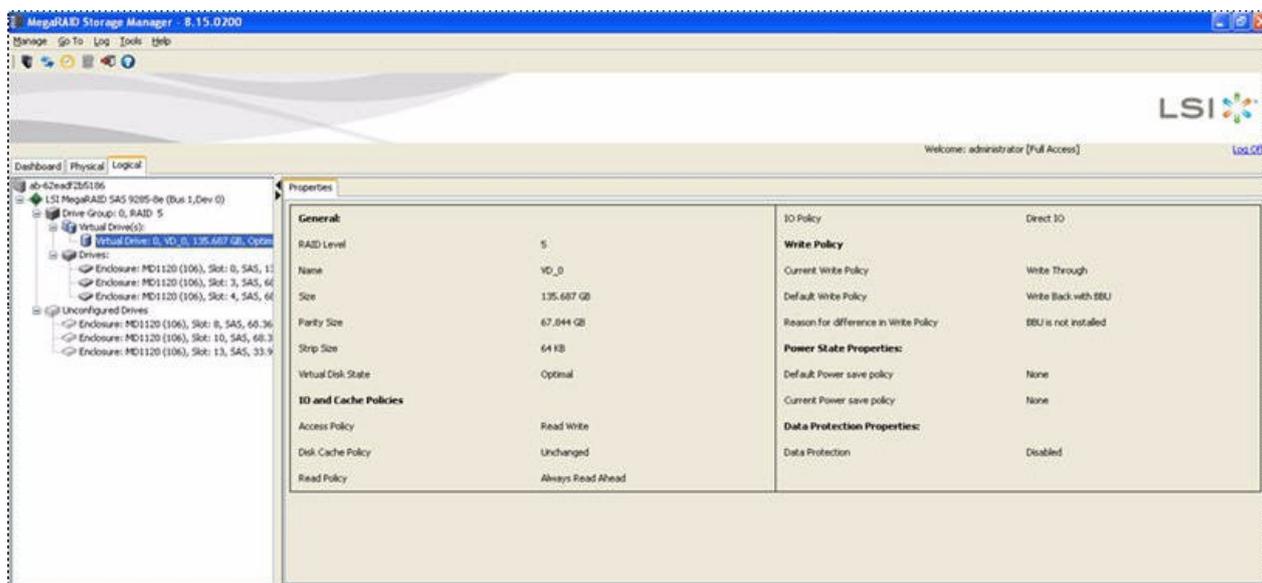
- Parity size
- Mirror date size
- Metadata size

4.6.8.1 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 73 Parity Size



4.6.8.2 Mirror Data Size

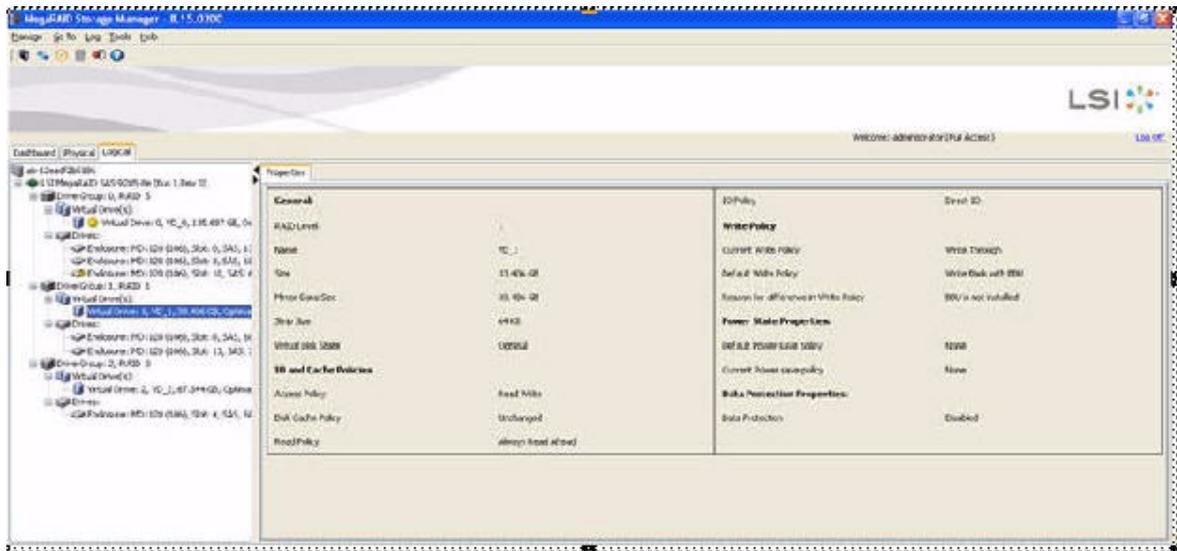
Mirror Data Size is used to determine the size used for storing redundant information on RAID 1 and RAID 10 virtual drives.

Follow these steps to view the Mirror Data Size.

1. In the **Logical** view, click on the Virtual Drive node.

The Mirror data size is displayed for RAID 1 and RAID 10 volumes, as shown in the following figure.

Figure 74 Mirror Data Size



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

4.6.8.3 Metadata Size

The metadata size field displays the total space used for metadata. Follow these steps to view the Metadata Size.

1. In the **Logical** view or the **Physical** view, click the controller node.

The total space used for metadata is displayed in this field, as shown in the following figure.

Figure 75 Metadata Size

Properties			
Alarm Present	Yes	Backend SAS Address 7	0x0
Alarm Enabled	No	Correctable Error Count	0
Cache Flush Interval	4 sec	Memory uncorrectable count	0
Coarson Mode	None	Cluster Enable	No
EDD Present	Yes	Cluster Active	No
NRAM Present	Yes	SSD Guard	Enabled
NRAM Size	32,000 KB	Drive Security Properties:	
BIOS Version	3.18.00_4.09.35.00_G104250000	Drive security capable	No
Native Command Queuing	Enabled	Background Operation Properties:	
Flush Size	8,000 MB	Rebuild Rate	55
Memory Size	512,000 MB	Format Head Rate	38
Metadata Size	500 MB	Reconstruction Rate	38
Power State Properties:		ECR Rate	24
Power savings on unconfigured drives	Enabled	Consistency Check Rate	26
Power savings on hot spares	Enabled	MegaRAID Recovery Properties:	
Power Save Policy for Configured Drives	Auto	MegaRAID Recovery	Enabled
Drive Standby Time	40mins	Firmware Properties:	

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.

4.6.9 Emergency Spare

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

4.6.9.1 Emergency Spare for Physical Drives

The Emergency Spare property determines whether a particular drive is capable of becoming an emergency spare. This property is displayed under the controller properties only if the Global spare for Emergency and the Unconfigured Good for Emergency controller properties are enabled.

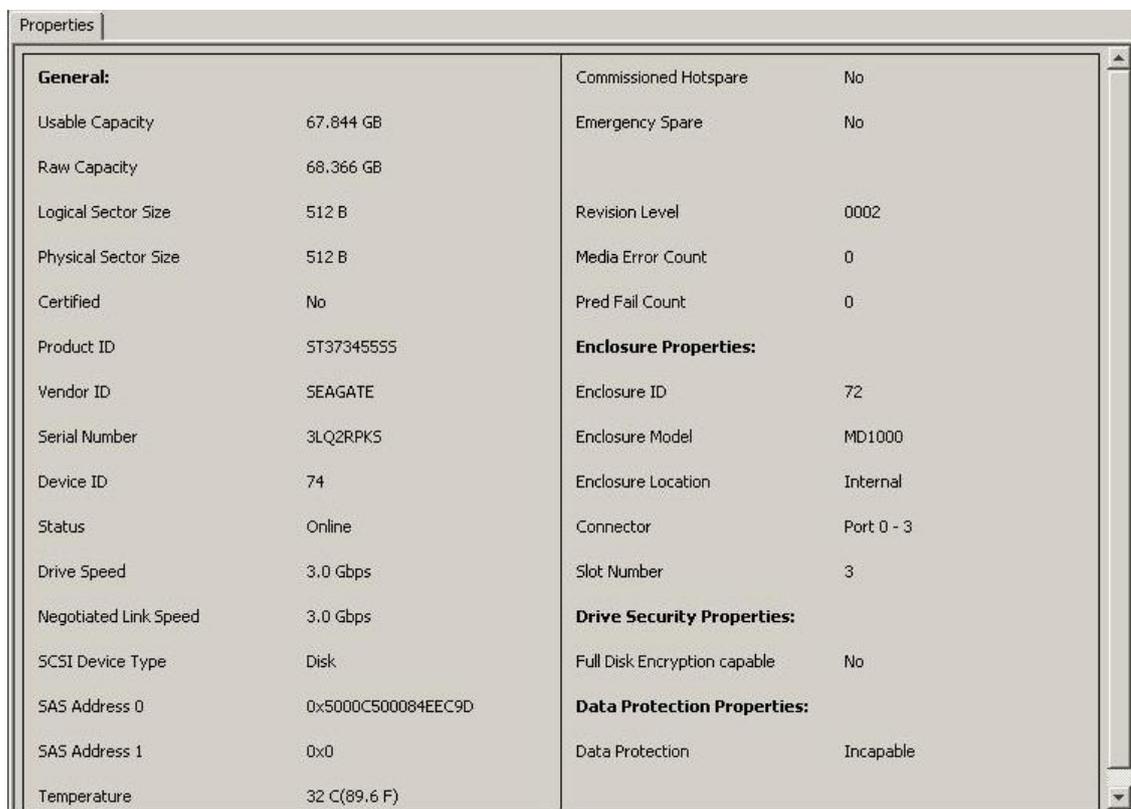
Follow these steps to view the Emergency Spare property.

1. Go to either the **Logical** view or the **Physical** view.
2. Click the drive for which you want to view the spare properties.

The Emergency spare is displayed under general properties. This property denotes whether a particular drive is commissioned as an emergency spare or not an emergency spare.

NOTE This property is displayed only for online physical drives.

Figure 76 Emergency Spare- Physical Drive Properties



General:		Commissioned Hotspare	
Usable Capacity	67.844 GB	Emergency Spare	No
Raw Capacity	68.366 GB	Revision Level	0002
Logical Sector Size	512 B	Media Error Count	0
Physical Sector Size	512 B	Pred Fail Count	0
Certified	No	Enclosure Properties:	
Product ID	ST37345555	Enclosure ID	72
Vendor ID	SEAGATE	Enclosure Model	MD1000
Serial Number	3LQ2RPK5	Enclosure Location	Internal
Device ID	74	Connector	Port 0 - 3
Status	Online	Slot Number	3
Drive Speed	3.0 Gbps	Drive Security Properties:	
Negotiated Link Speed	3.0 Gbps	Full Disk Encryption capable	No
SCSI Device Type	Disk	Data Protection Properties:	
SAS Address 0	0x5000C500084EEC9D	Data Protection	Incapable
SAS Address 1	0x0		
Temperature	32 C(89.6 F)		

4.6.9.2 Emergency Spare Property for Controllers

The Emergency spare properties under the controller properties are configured based on enabling or disabling the following properties:

- Emergency Spare
- Emergency for SMARTer

To view the Emergency spare property for controllers, click the controller node in the device tree. The emergency spare properties are displayed, as shown in the following figure.

Figure 77 Emergency Spare Properties for Controllers

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

4.6.9.3 Commissioned Hotspare

The commissioned hotspare is used to determine whether the online drive has a Commissioned Hotspare. To check if the drive is commissioned with a hotspare, click the online physical drive node in the device tree.

The Commissioned Hotspare property is displayed, as shown in the following figure. This property is displayed only for online physical drives.

Figure 78 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		

4.6.10 SSD Disk Cache Policy

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

The MegaRAID firmware does not allow any user application to modify the write-cache policies of any SSD media. The host applications can modify this property through a new logical device (LD) addition or a LD property change. When SSDs are configured in a mixed disk group with HDDs, the Physical Device Write-Cache Policy setting of all the participating drives are changed to match the SSD cache policy setting.

Follow these steps to view the SSD cache property.

1. Click the controller node in the device tree.

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

Figure 79 Controller Properties – SSD Disk Cache Policy

Properties			
Host Port Count	0	Backend SAS Address 5	0x0
FTU		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
BSI Present	No	VSU Guard	Enabled
NVRAM Present	Yes	SSD Disk Cache Setting	Disabled
NVRAM Size	32,000 KB	Drive Security Properties:	
BIOS Version	3.18.08_4.09.05.01_0x016A000	Drive security enabled	No
Native Command Queuing	Enabled	Drive security method	FDE Only
Flash Size	0.600 PB	Drive security capable	Yes
Memory Size	256,000 MB	EEM 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
Premium Package Version	12.10.6-0015	BGU Rate	30
		Consistency Check Rate	30

4.6.10.1 Virtual Drive Settings

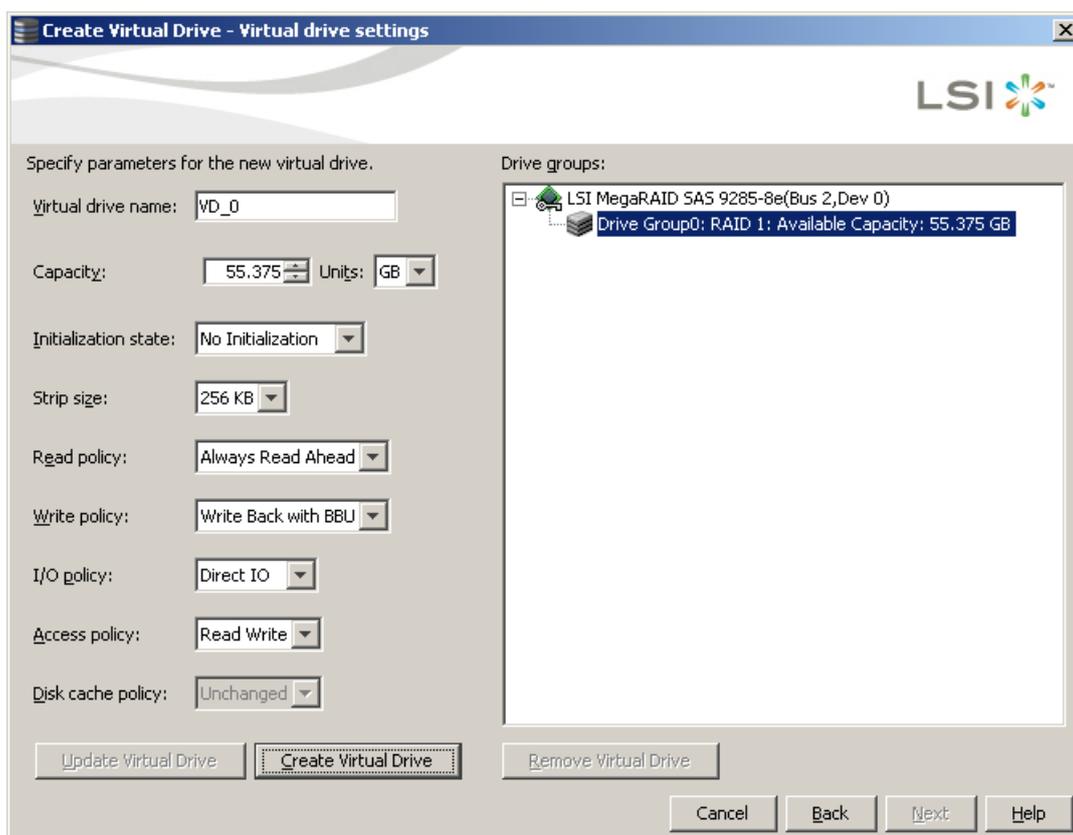
If the SSD cache property is enabled in the controller properties dialog as shown in [Controller Properties – SSD Disk Cache Policy](#), then you cannot select the disk cache policy for the virtual drives having only SSD drives or a mix of SSD drives and HDD drives during virtual drive creation. The value of the disk cache policy is unchanged and the drop-down menu is disabled.

Follow these steps to view the virtual drive settings.

1. Right-click the controller node in the device tree.
2. Select the **Create Virtual Drive** menu option.
3. Select **Advanced Configuration**, and click **Next**.
4. Create **Drive Group**, and click **Next**.

The **Create Virtual Drive – Virtual drive settings** dialog appears, as shown in the following figure.

Figure 80 Virtual Drive Settings



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

4.6.10.2 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.

Figure 81 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:

Ok Cancel

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.

4.6.11 Non-SED Secure Erase Support

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

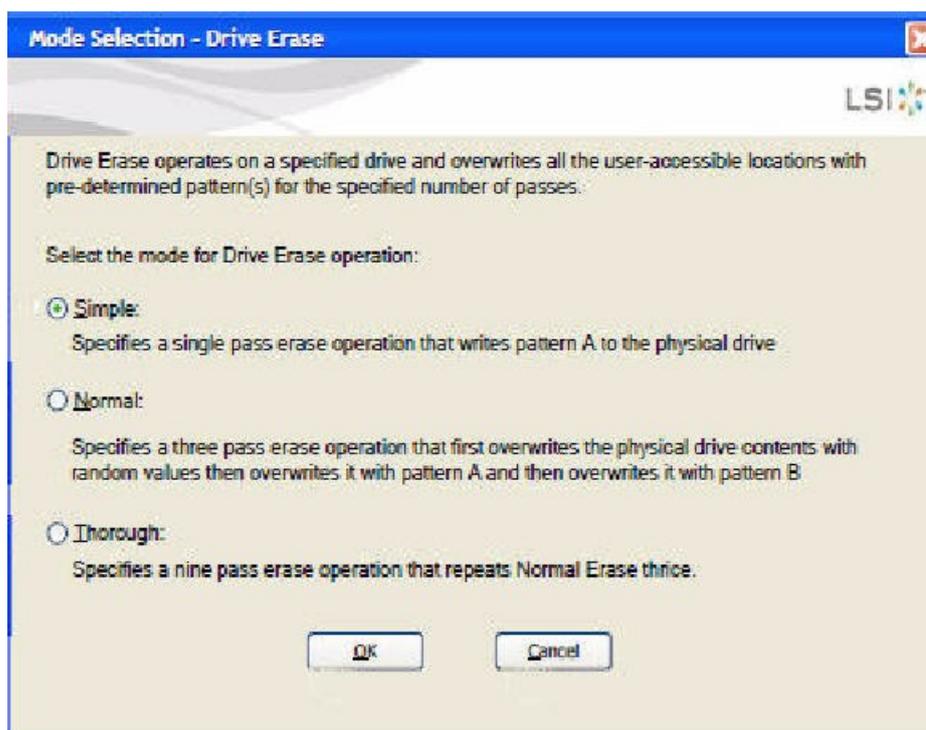
SEDs securely erase their internal encryption keys, effectively destroying all of the data present on the drive. For Non-SED drives, the erase operation consists of a series of write operations to a drive that overwrites every user-accessible sector of the drive with specified patterns. It can be repeated in multiple passes using different data patterns for enhanced security. The sanitization technique is more secure than a simple format operation and is commonly called a “clearing” operation, similar to the existing physical drive clear command.

Follow these steps to set physical drive properties.

1. In the Physical view, right click the **Physical Drive** node.
2. Select the **Drive Erase** option (Alt+E).

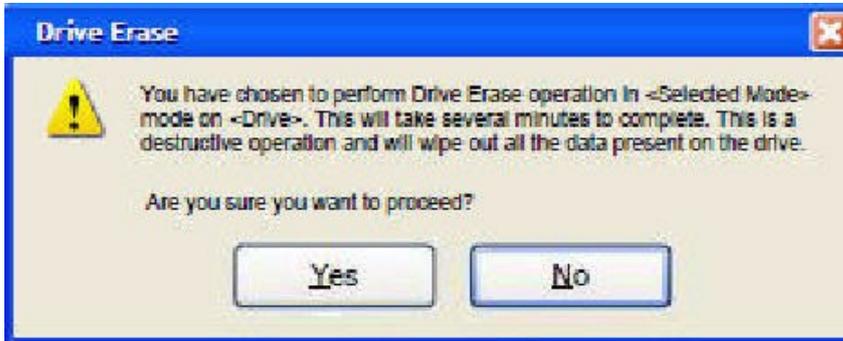
The **Mode Selection - Drive Erase** dialog appears.

Figure 82 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 83 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.

4.6.11.1 Group Show Progress for Drive Erase

Physical drive erase operation is a time-consuming operation and is performed as a background task. It posts events to notify users of the progress.

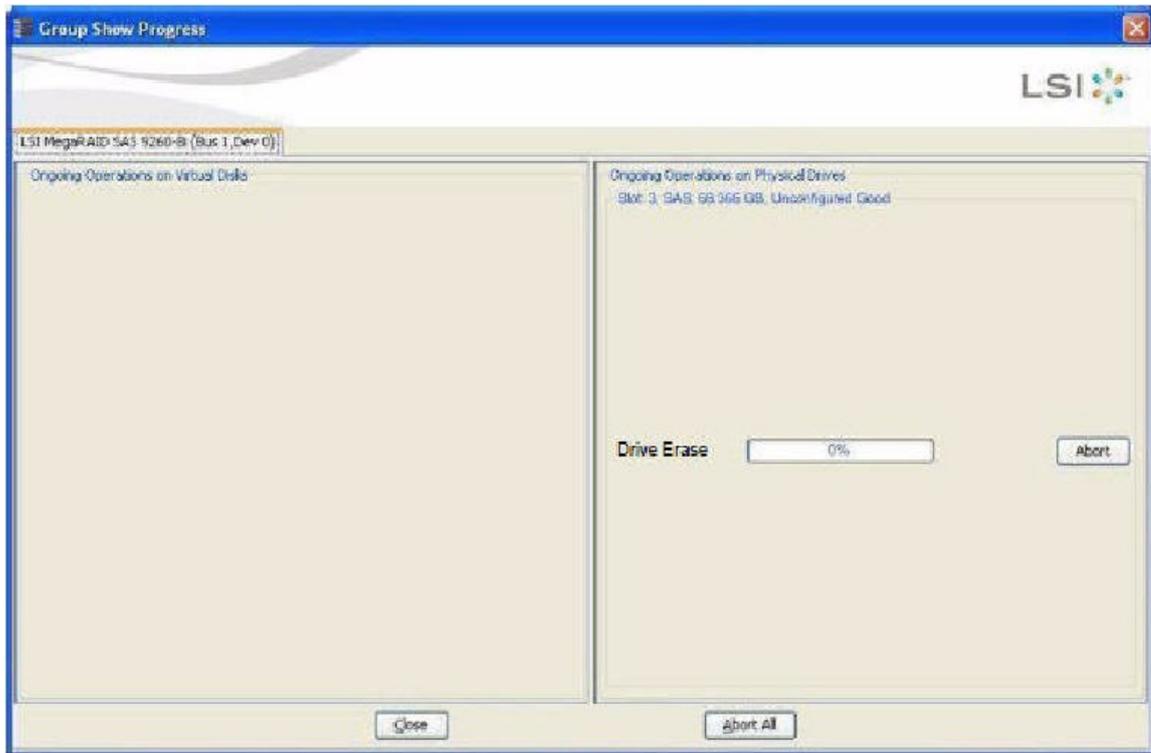
Follow these steps to check the progress of physical drive erase operation.

1. Click the **Show Progress** toolbar icon in the MegaRAID Storage Manager. You can also select **Show Progress** from the dashboard or select **Show Progress** from the Manage menu.

2. Click the **More info** link under the Background Operations portlet.

The progress bar appears.

Figure 84 Group Show Progress



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

4.6.11.2 Virtual Drive Erase

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

Follow these steps to open the Virtual Drive Erase menu.

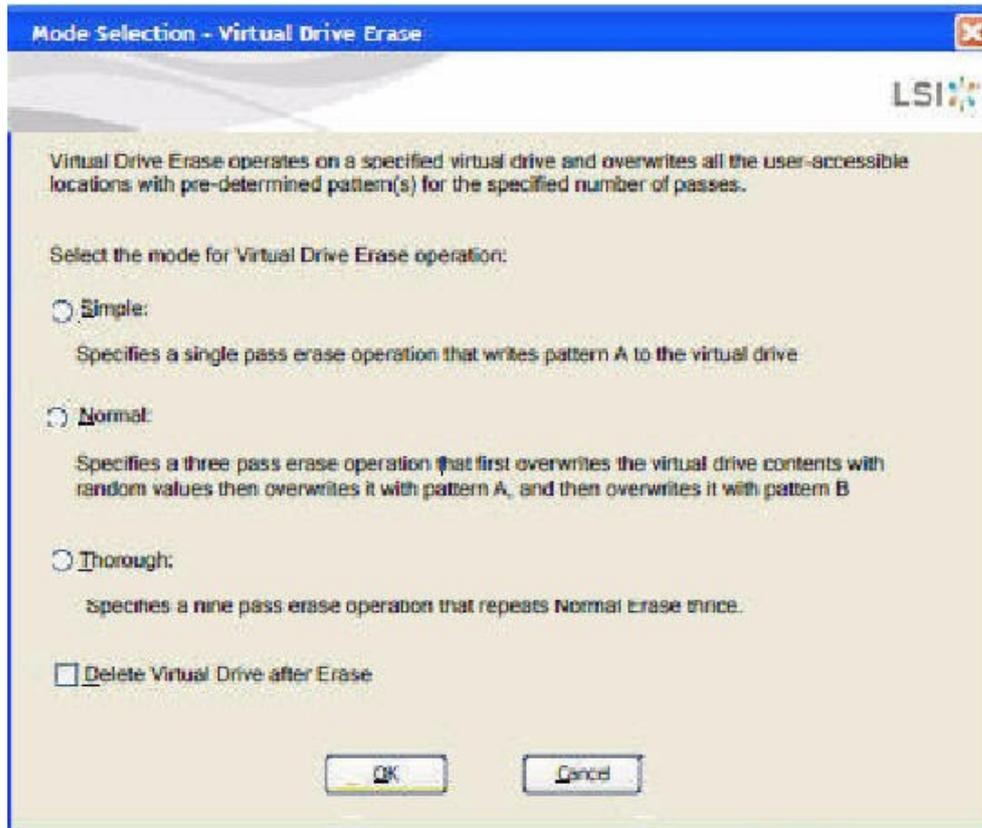
1. In the Logical view, right-click the Virtual Drive node.
2. Click on the Virtual Drive node, select top level navigation and click **Go to**.
3. Select **Virtual Drive** and select **Events & Response**.

The **Logical View - Virtual Drive Erase** menu appears.

4. Select **Virtual Drive Erase**.

The **Virtual Drive Erase Menu** opens, as shown in the following figure.

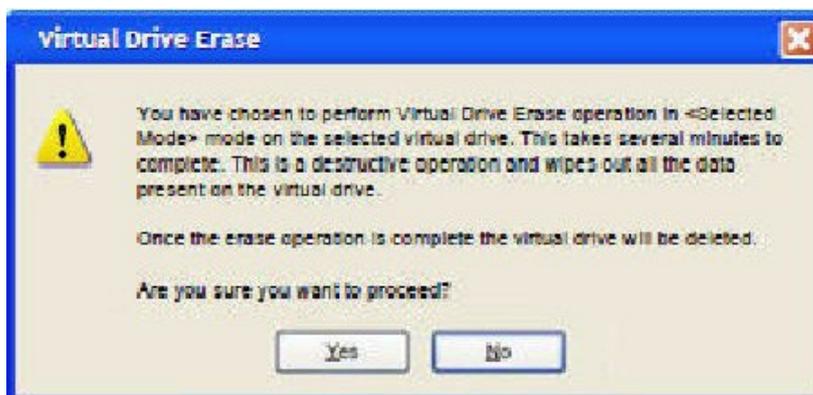
Figure 85 Mode Selection – Virtual Drive Erase Dialog



The menu has the following options.

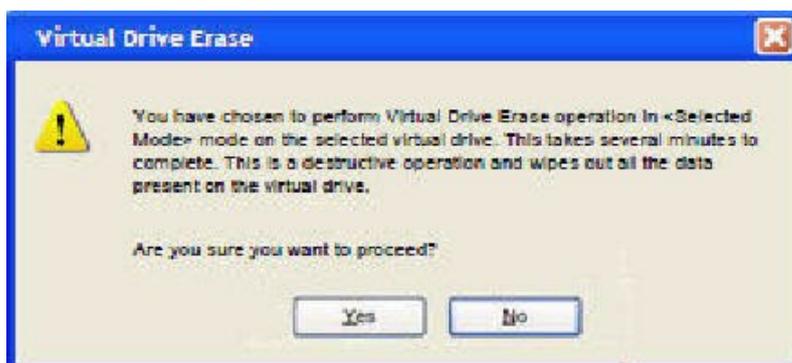
- **Simple** – (Alt + S) – After you select this option and click **OK**, and if **Delete Virtual Drive after Erase** is selected, [Figure 192](#) appears; otherwise, [Figure 193](#) appears.
- **Normal** – (Alt + N) – After you select this option and click **OK**, and if **Delete Virtual Drive after Erase** is selected, [Figure 192](#) appears; otherwise, [Figure 193](#) appears.
- **Thorough** – (Alt + T) – After you select this option and click **OK**, and if **Delete Virtual Drive after Erase** is selected, [Figure 192](#) appears; otherwise, [Figure 193](#) appears.
- **Delete Virtual Drive after Erase** – (Alt + D) – When you select this option, the virtual drive is erased and [Figure 192](#) appears; otherwise, [Figure 193](#) appears.
- **OK** – (Alt + O) – Click **OK** and if **Delete Virtual Drive after Erase** is checked, [Figure 192](#) appears; otherwise, [Figure 193](#) appears.
- **Cancel** – (Alt + C) – When you select this option, the dialog closes, and the MegaRAID Storage Manager navigates back to Physical view.

Figure 86 Warning Message for Virtual Drive Erase



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

Figure 87 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.

4.6.11.3 Group Show Progress for Virtual Drive Erase

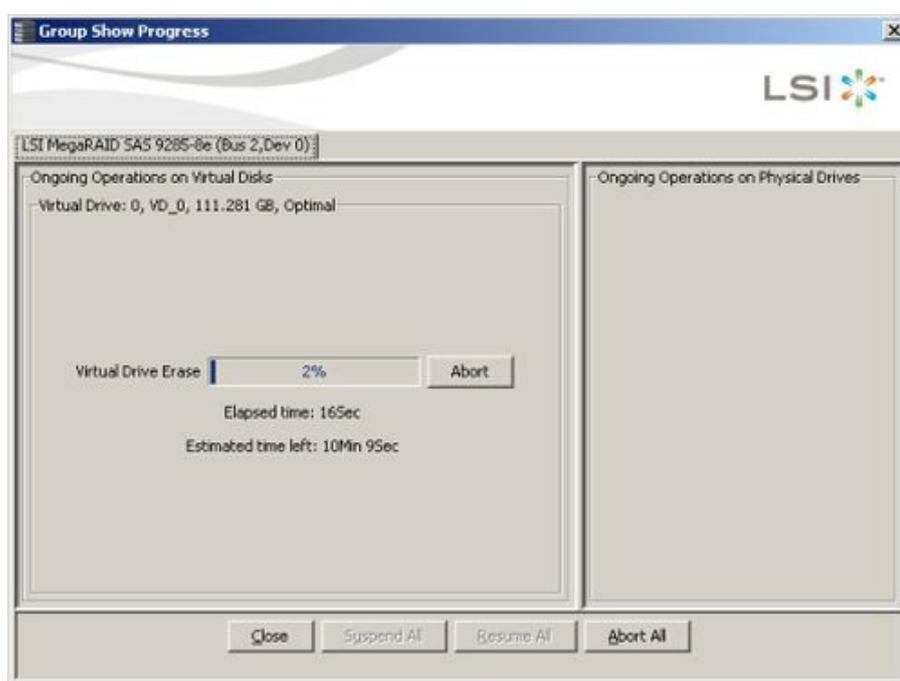
The virtual drive erase operation is a time-consuming operation and is performed as a background task. It posts events to notify users of the progress.

To view the progress of Group Show Progress-Virtual Drive, click the **Show Progress** toolbar icon.

You can also either select **Show Progress** from the Manage menu, or select the **More info** Link under Background Operations portlet on the dashboard.

The Virtual Drive Erase progress bar appears, as shown in the following figure.

Figure 88 Group Show Progress – Virtual Drive



4.6.12 Rebuild Write Cache

MegaRAID firmware supports drive cache properties during a rebuild operation. The MegaRAID solution temporarily enables drive cache for the physical drive that is being rebuilt for the duration of the rebuild operation. Users can enable or disable this feature using the Mega CLI feature.

The MegaRAID software automatically changes the setting for a drive that is being rebuilt. If the PD_CACHE for the rebuilt drive is already set, the firmware does not need to do anything extra.

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

4.6.13 Background Suspend / Resume Support

MegaRAID provides a background Suspend or Resume Support feature that enhances the functionality where in the background operations running on a physical drive or a virtual drive can be suspended for some time, and resumed later using the Resume option.

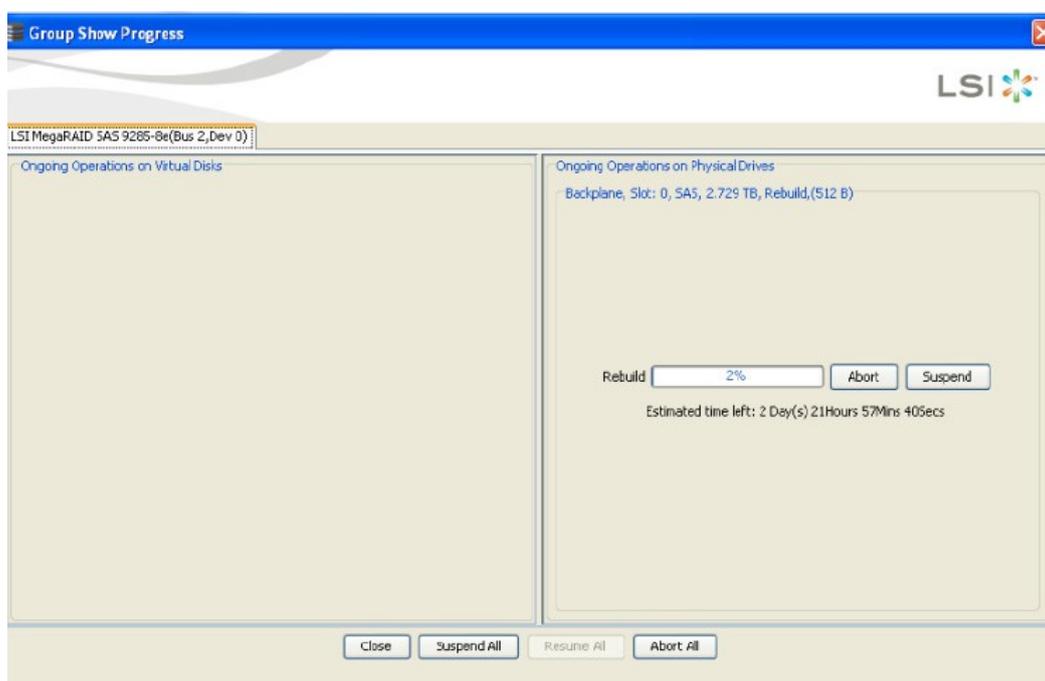
The background operations, including consistency-check, rebuild, replace, and background initialization are supported by an abort operation. If any operation is stopped before completion, it is considered to be aborted. An aborted operation cannot be resumed from the place where it was stopped.

A suspended operation can be resumed later by using the **Resume** option, and the suspended operation resumes from the point where the operation was suspended last.

To perform a suspend and resume operation, go to the **Group Show Progress** dialog, and perform the tasks mentioned below. You also can select **Show Progress** from the **Manage** menu, or select the **More info** link under the **Background Operations** portlet on the dashboard.

The **Group Show Progress** dialog appears, as shown in the following figure. If Patrol Read is running, the **Group Show Progress Patrol Read** dialog appears.

Figure 89 Group Show Progress

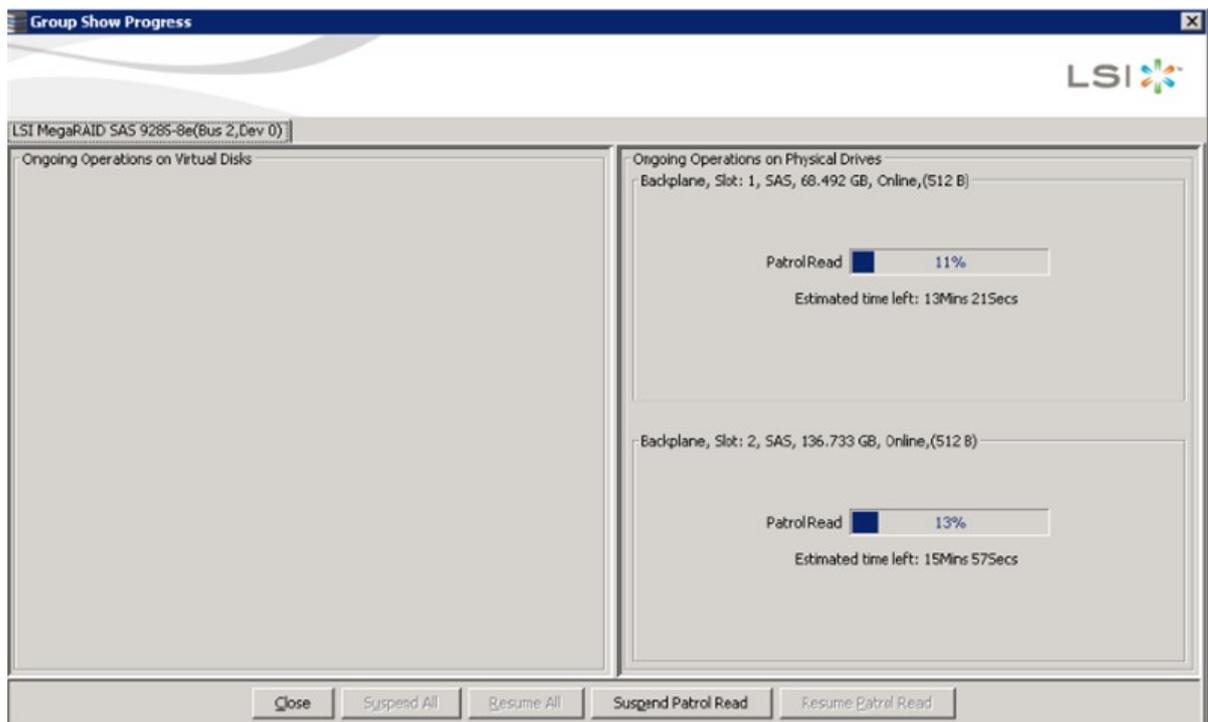


- **Suspend** (Alt + S) – Click the **Suspend** button to suspend the background operation taking place at that particular point of time. When the operations get suspended, the **Resume** button appears instead of the **Suspend** button.

- **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 90 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.

4.6.14 Enclosure Properties

To view the enclosure properties, in the Physical view, click the **Enclosure**  node. The Enclosure Properties are displayed, as shown in the following figure.

Figure 91 Enclosure Properties

Vendor ID	DELL	FRU Number	42R5133
Enclosure ID	5	Part Number	CP-111-006-020
Enclosure Type	SES	Component Properties	
Enclosure Model	MD1000	Number of Temperature Sensors	4
Enclosure Location	External	Number of Fans	4
Firmware Version	A.04	Number of Power Supplies	2
Serial Number	0802V16VTE	Number of Voltage Sensors	0
Connector	Port A		
Number of Slots	15		

4.7 GUI Elements in the MegaRAID Storage Manager Window and Menus

This section describes the graphical user interface (GUI) elements used in the MegaRAID Storage Manager software.

4.7.1 Device Icons

The following icons in the left panel represent the controllers, drives, and other devices.

	Status
	System
	Controller
	Backplane
	Enclosure
	Port
	Drive group
	Virtual drive
	Online drive
	Power save mode
	Dedicated hotspare
	Global hotspare
	Battery backup unit (BBU)
	Tape drive
	CD-ROM
	Foreign drive
	Unconfigured drive
	Locked SED
	Unlocked SED

NOTE The MegaRAID Storage Manager software shows the icons for tape drive devices; however, no tape-related operations are supported by the utility. If these operations are required, use a separate backup application.

A red circle to the right of an icon indicates that the device has failed. For example, this icon indicates that a drive has failed:

A yellow circle to the right of an icon indicates that a device is running in a partially degraded state. For example, this icon indicates that a virtual drive is running in a degraded state because a controller has failed.

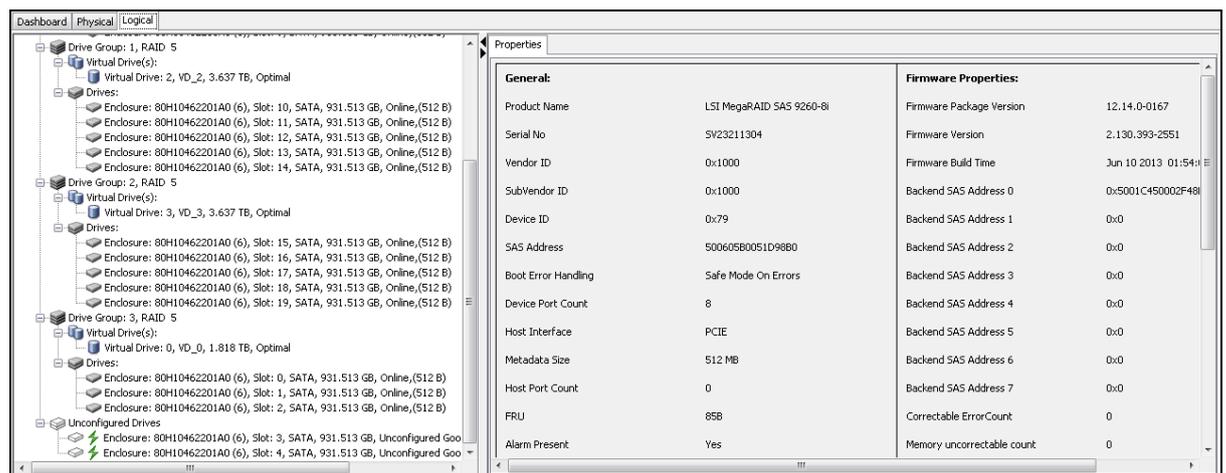
An orange circle to the right of an icon indicates that a device is running in a degraded state.

4.7.2 Properties Tab

The right panel of the MegaRAID Storage Manager window has one tab.

- 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 [Monitoring Controllers](#), [Monitoring Drives](#), and [Monitoring Virtual Drives](#).

Figure 92 Properties Tab and Graphical View Tab



4.7.3 Event Log Panel

The lower part of the **MegaRAID Storage Manager** window displays the system event log entries. New event log entries appear during the session. Each entry has an ID, an error level indicating the severity of the event, the timestamp and date, and a brief description of the event.

For more information about the event log, see **Monitoring Controllers and Their Attached Devices**. For more information about the event log entries, see [Events and Messages](#).

4.7.4 Menu Bar

Here are brief descriptions of the main selections on the MegaRAID Storage Manager menu bar. Specific menu options are described in more detail in the **Configuration** and **Maintaining and Managing Storage Configurations** sections.

Manage Menu

The Manage menu has a **Refresh** option for updating the display in the **MegaRAID Storage Manager** window (refresh is seldom required; the display usually updates automatically) and an **Exit** option to end your session on MegaRAID Storage Manager. The **Server** option shows all the servers that were discovered by a scan. In addition, you can perform a check consistency, initialize multiple virtual groups, and show the progress of group operations on virtual drives.

Go To Menu

The Go To menu is available when you select a controller, drive group, physical drive, virtual drive, or battery backup unit in the main menu dialog. The menu options vary depending on the type of device selected in the left panel of the MegaRAID Storage Manager main menu. The options also vary depending on the current state of the selected device. For example, if you select an offline drive, the **Make Drive Online** option appears in the Physical Drive menu.

Configuration options are also available. This is where you access the Configuration Wizard that you use to configure drive groups and virtual drives. To access the Wizard, select the controller in the left panel, and then select **Go To > Controller > Create Virtual Drive**.

Log Menu

The Log menu includes options for saving and clearing the message log. For more information about the Log menu, see [Events and Messages](#).

Tools Menu

On the Tools menu, you can select **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 [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.

Chapter 5 Configuration

This chapter explains how to use MegaRAID Storage Manager software to create and modify storage configurations on LSI SAS controllers.

The LSI SAS controllers support RAID 0, RAID 1, RAID 5, RAID 6, RAID 00, RAID 10, RAID 50, and RAID 60 storage configurations. The **Configuration** wizard allows you to easily create new storage configurations and modify the configurations. To learn more about RAID and RAID levels, see [Introduction to RAID](#).

NOTE You cannot create or modify a storage configuration unless you are logged on to a server with administrator privileges.

5.1 Creating a New Configuration

You can use the MegaRAID Storage Manager software to create new storage configurations on systems with LSI SAS controllers. You can create the following types of configurations:

- **Simple configuration** specifies a limited number of settings and has the system select drives for you. This option is the easiest way to create a virtual drive.
- **Advanced configuration** lets you choose additional settings and customize virtual drive creation. This option provides greater flexibility when creating virtual drives for your specific requirements.

This section describes the virtual drive parameters and explains how to create simple and advanced storage configurations.

5.1.1 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
- **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.

5.1.2 Optimum Controller Settings for Fast Path

Write Policy: Write Through

IO Policy: Direct IO

Read Policy: No Read Ahead

Stripe Size: 64 KB

5.1.3 Creating a Virtual Drive Using Simple Configuration

Simple configuration is the quickest and easiest way to create a new storage configuration. When you select simple configuration mode, the system creates the best configuration possible using the available drives.

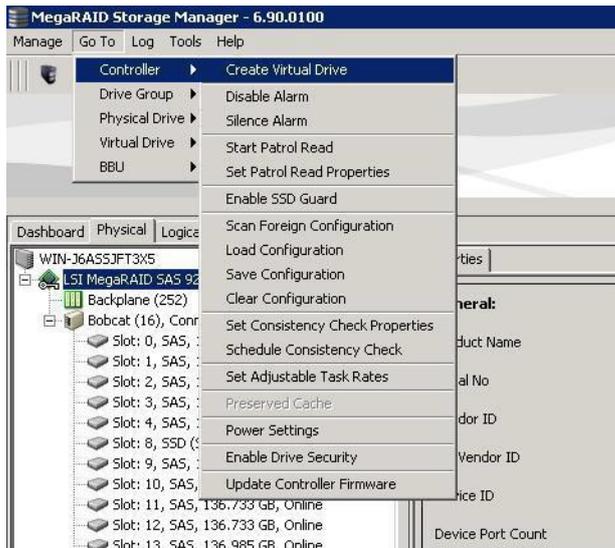
NOTE You cannot create spanned drives using the simple configuration procedure. To create spanned drives, use the advanced configuration procedure described in [Creating a Virtual Drive Using Advanced Configuration](#).

Follow these steps to create a new storage configuration in simple configuration mode.

1. Perform either of the following steps:

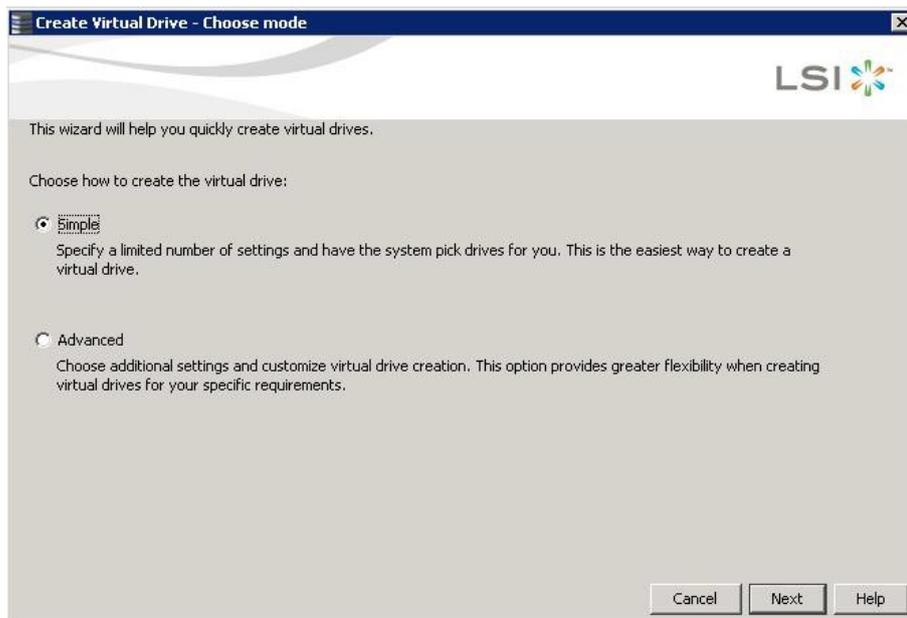
- Right-click the controller node in the device tree in the left frame of the **MegaRAID Storage Manager** window, and select **Create Virtual Drive**.
- Select the controller node, and select **Go To > Controller > Create Virtual Drive** in the menu bar, as shown in the following figure.

Figure 93 Create Virtual Drive Menu Option



The dialog for the configuration mode (simple or advanced) appears, as shown in the following figure.

Figure 94 Create Virtual Drive - Choose mode

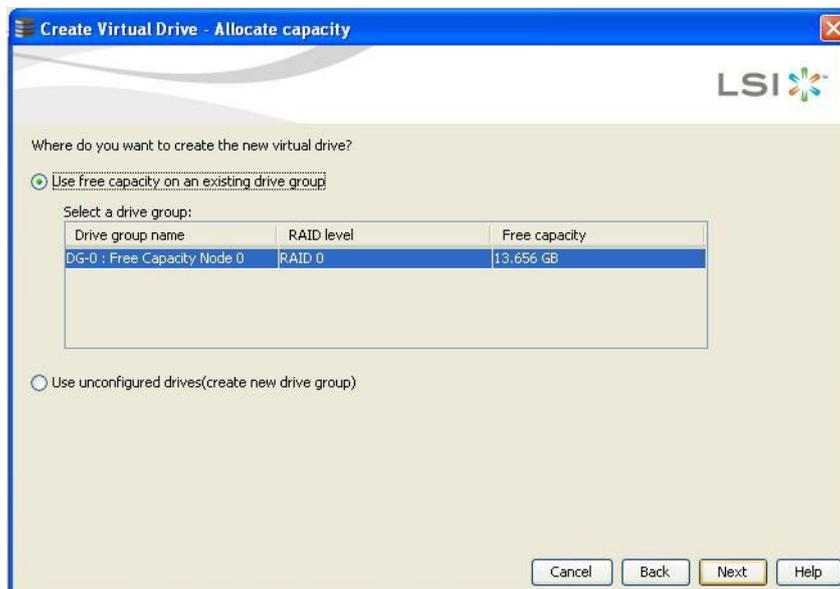


2. Select the **Simple** radio button, and click **Next**.

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

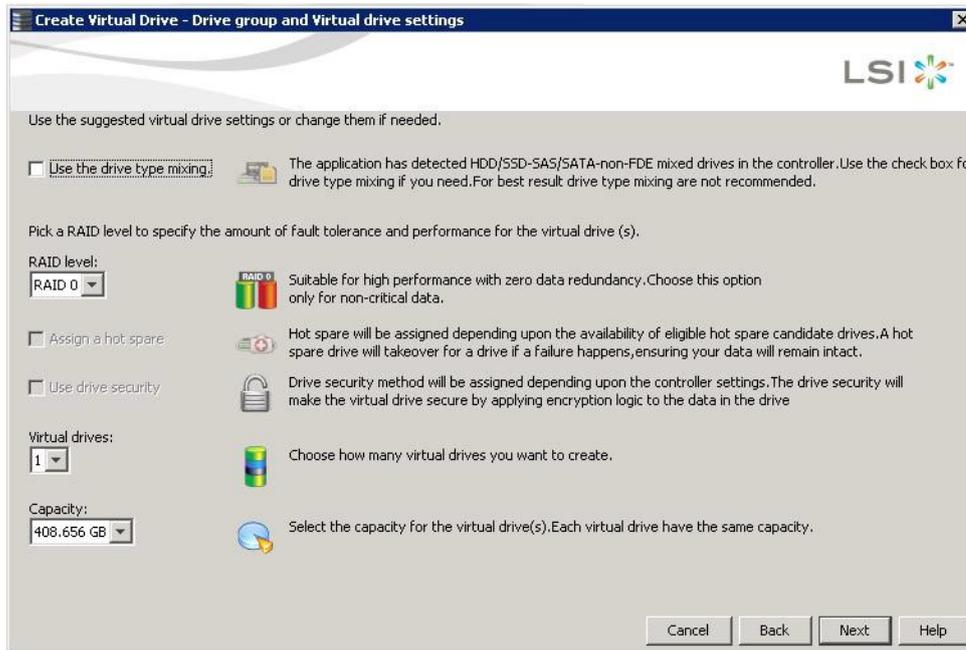
drives are available, the **Create Drive Group Settings** window appears, and you can go to step 4.

Figure 95 Using the Free Capacity of an Existing Drive Group



3. Perform either of the two options:
 - If a drive group exists, select the **Use free capacity on an existing drive group** radio button and click **Next**. Continue with step 4. The **Create Virtual Drive** window appears, as shown in the following figure. If different types of drives are attached to the controller, such as HDD, SDD, SAS, and SATA, an option appears to allow drive type mixing.
 - If unconfigured drives are available, select the radio button to use the unconfigured drives, and click **Next**. Continue with step 10. The Summary window appears as shown in [Figure 203](#).

Figure 96 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.

The LSI SafeStore Data Security Service encrypts data and provides disk-based key management for your data security solution. This solution protects the data in the event of theft or loss of drives. See [LSI MegaRAID SafeStore Encryption Services](#), for more information about the SafeStore feature.

8. Use the drop-down list in the **Virtual drives** field to choose how many virtual drives you want to create.

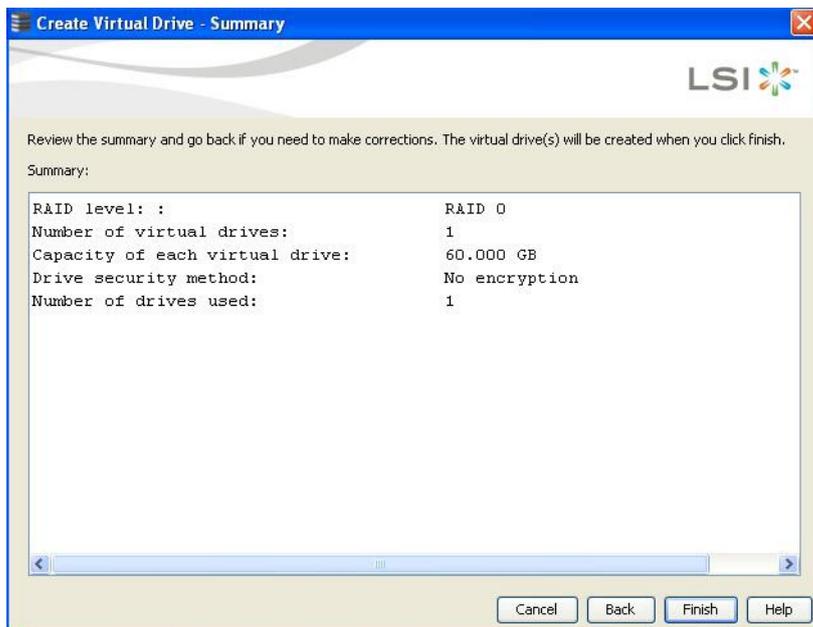
9. Select the capacity of the virtual drives.

Each virtual drive has the same capacity.

10. Click **Next**.

The **Create Virtual Drive - Summary** window appears, as shown in the following figure. This window shows the selections you made for simple configuration.

Figure 97 Create Virtual Drive - Summary Window



NOTE If High Availability DAS is supported on the controller and you are creating a virtual drive using simple configuration, by default, the virtual drive is shared with the other servers in that cluster.

11. Either click **Back** to return to the previous window to change any selections, or click **Finish** to accept and complete the configuration.

The new virtual drive is created after you click **Finish**. After the configuration is completed, a dialog box notifies you that the virtual drives were created successfully.

NOTE If you create a large configuration using drives that are in Power-Save mode, it could take several minutes to spin up the drives. A progress bar appears as the drives spin up. If any of the selected unconfigured drives fail to spin up, a dialog box that identifies these drives appears.

5.1.4 Creating a Virtual Drive Using Advanced Configuration

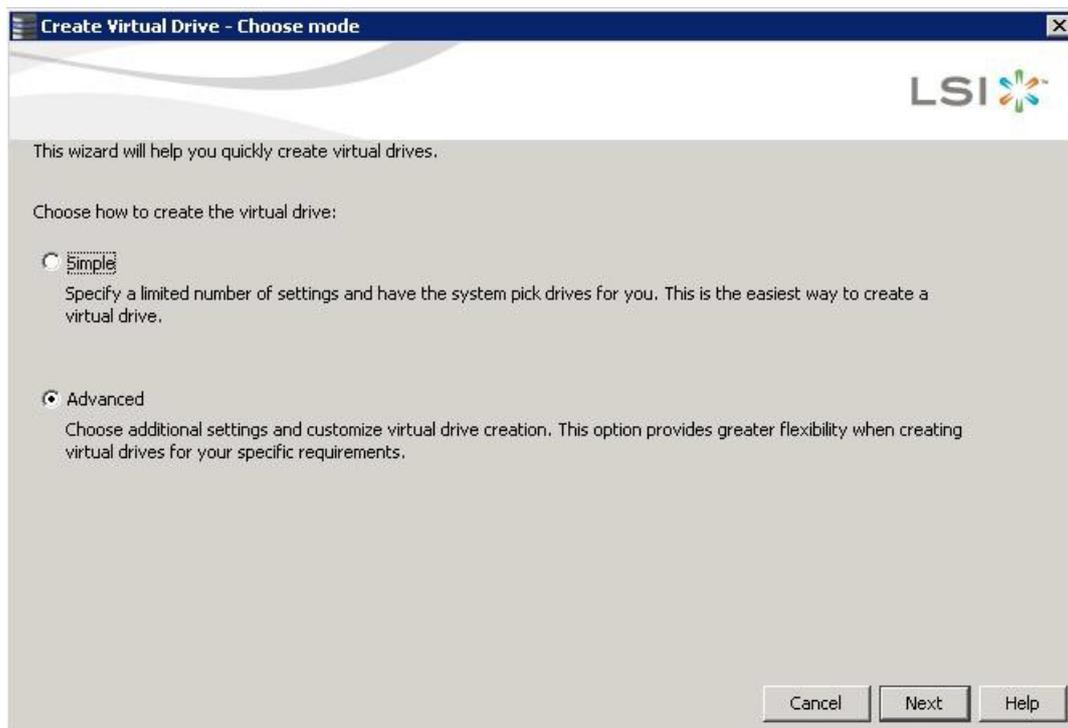
The advanced configuration procedure provides an easy way to create a new storage configuration. Advanced configuration gives you greater flexibility than simple configuration because you can select the drives and the virtual drive parameters when you create a virtual drive. In addition, you can use 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.

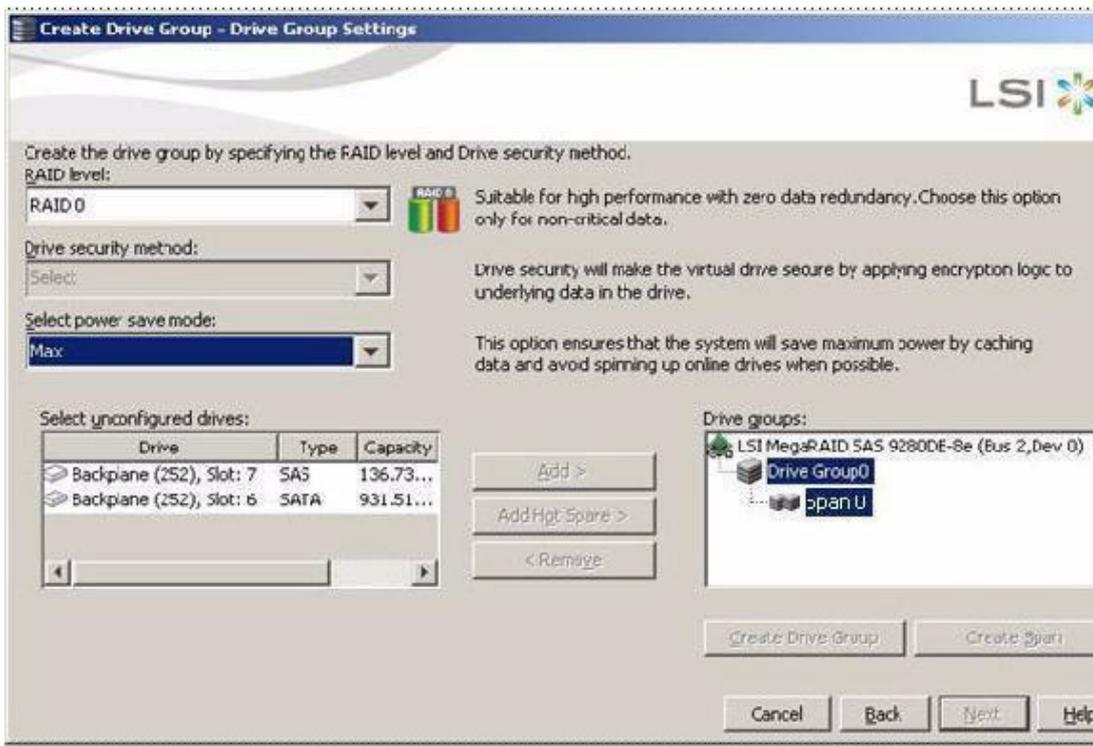
Figure 98 Create Virtual Drive - Choose mode Dialog



2. Select the **Advanced** radio button, and click **Next**.

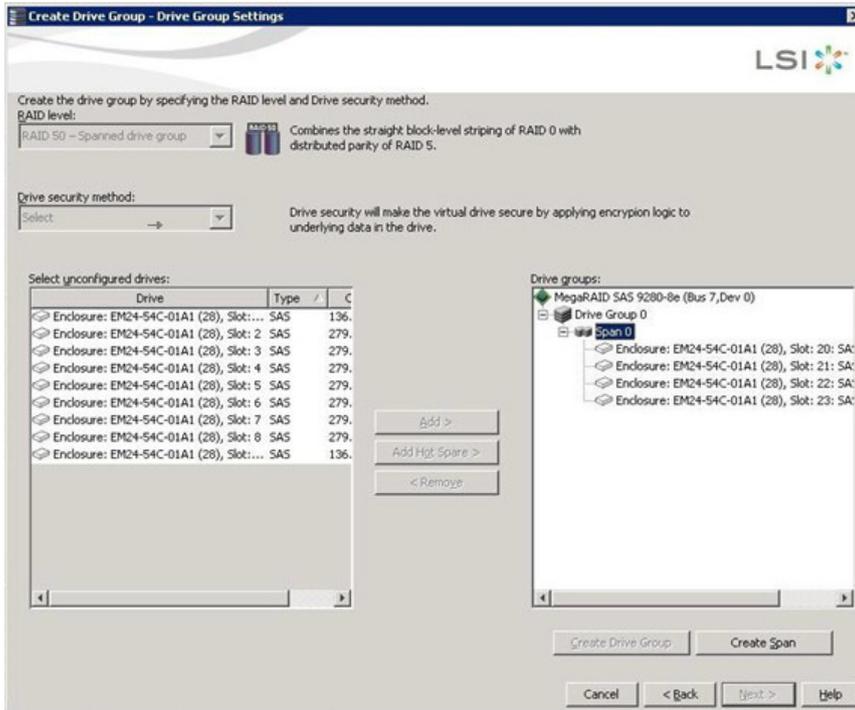
The **Create Drive Group Settings** window appears, as shown in the following figure.

Figure 99 Create Drive Group - Drive Group Settings Window



3. Select the following items on the **Create Drive Group - Drive Group Settings** window:
 - a. Select the RAID level desired for the drive group from the drop-down menu. To make a spanned drive, select **RAID 10**, **RAID 50**, or **RAID 60** in the **RAID level** field.
Drive Group 0 and **Span 0** appear in the **Drive groups** field when you select RAID 10, 50, or 60.
 The RAID controller supports RAID levels 1, 5, 6, 10, 50, and 60. In addition, it supports independent drives (configured as RAID 0 and RAID 00). The dialog text gives a brief description of the RAID level that you select. You can choose the RAID levels depending on the number of available drives.
 - b. Scroll down the menu for the **Drive security method** field if you want to set a drive security method.
 The drive security feature provides the ability to encrypt data and use disk-based key management for your data security solution. This solution provides protection to the data in the event of theft or loss of drives. See [LSI MegaRAID SafeStore Encryption Services](#), for more information about drive security and encryption.
 - c. Select *unconfigured* drives from the list of drives, and click **Add>** to add them to the drive group.
 The selected drives appear under **Span 0** below **Drive Group 0**, as shown in the following figure.

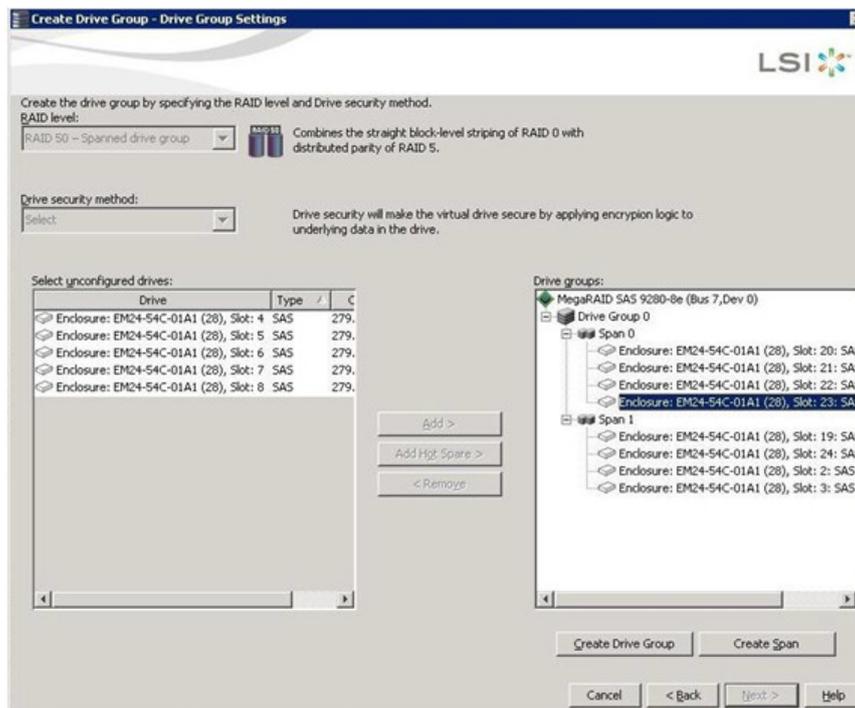
Figure 100 Span 0 of Drive Group 0



- 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 101 Span 0 and Span 1 of Drive Group 0

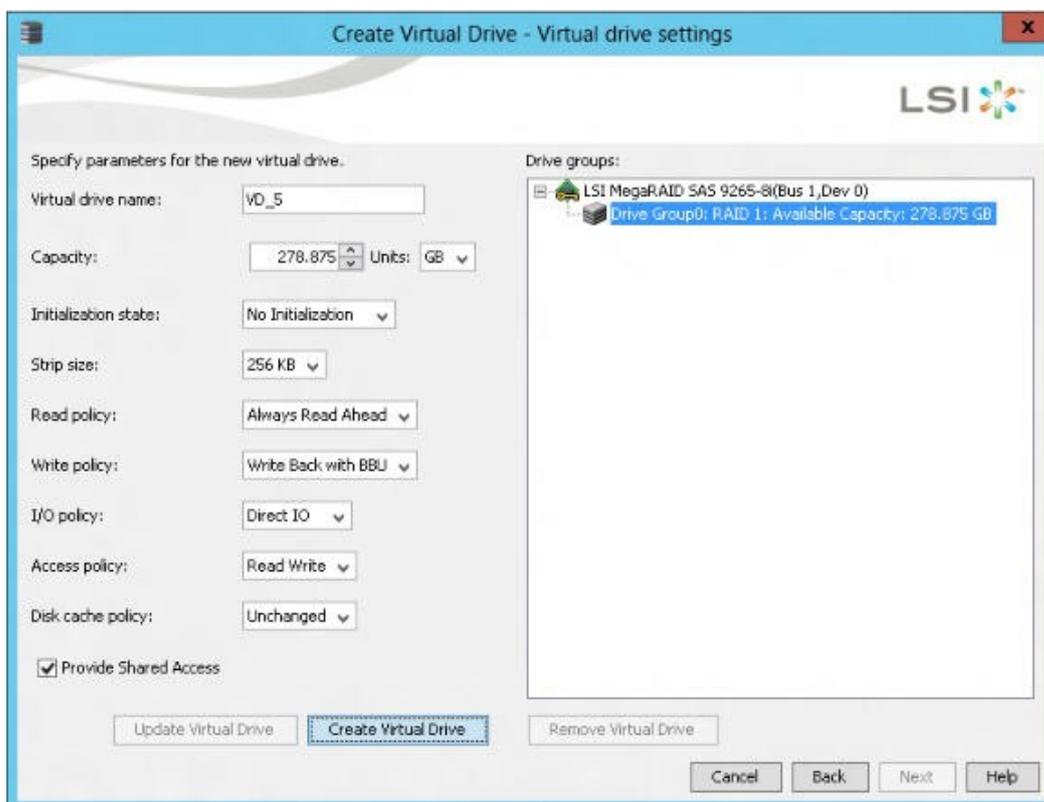


- f. Click **Create Drive Group** to make a drive group with the spans.
- g. Click **Next** to complete this step.

The **Create Virtual Drive - Virtual drive settings** window appears, as shown in the following figure. The drive group and the default virtual drive settings appear. The options to update the virtual drive or remove the virtual drive are grayed out until you create the virtual drive.

NOTE The parameters in the **Create Virtual Drive - Virtual drive settings** window display in Disabled mode (grayed out) for SAS-Integrated RAID (IR) controllers because these parameters do not apply to SAS-IR controllers.

Figure 102 Create Virtual Drive - Virtual Drive Settings Window



NOTE If you select **Write Back with BBU** as the write policy, and no battery exists, the battery is low or failed, or the battery is running through a re-learn cycle, the write policy switches to **Write Through**. This setting eliminates the risk of data loss in case of a power failure. A message window notifies you of this change.

NOTE If the controller supports High Availability DAS, the **Provide Shared Access** option appears in the above dialog. Select this option if you want the virtual drive to be shared between the two servers in a cluster.

4. Change any virtual drive settings, if desired.

See [Selecting Virtual Drive Settings](#), for more information about the virtual drive settings.

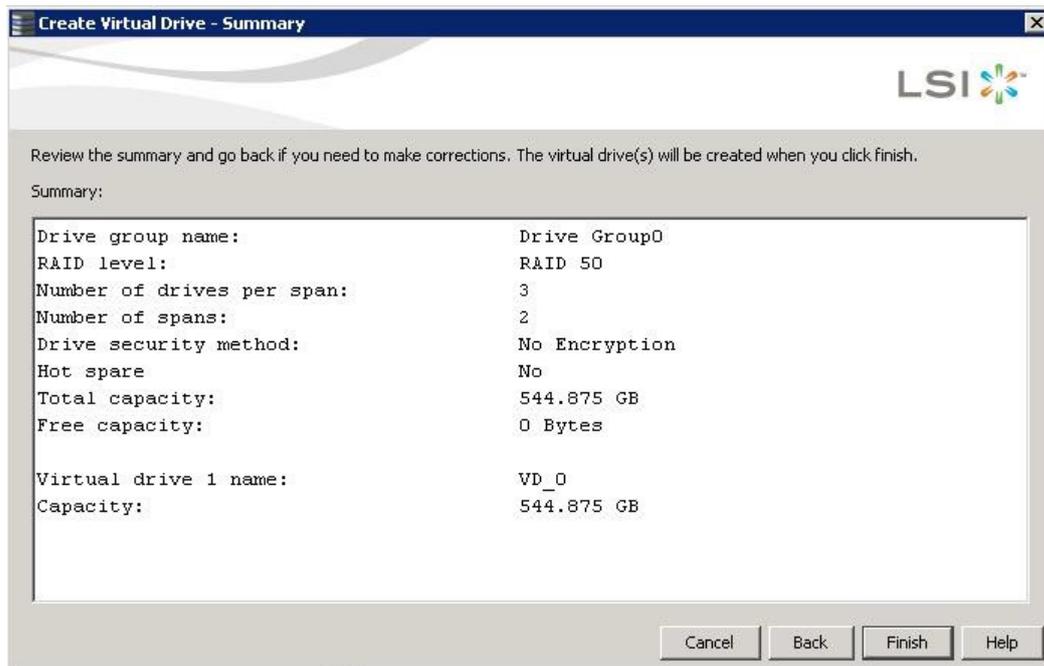
5. Click **Create Virtual Drive**.

The new virtual drive appears under the drive group. The options **Update Virtual Drive** and **Remove Virtual Drive** are available. **Update Virtual Drive** allows you to change the virtual drive settings, and **Remove Virtual Drive** allows you to delete the virtual drive.

6. Click **Next**.

The **Create Virtual Drive - Summary** window appears, as shown in the following figure. This window shows the selections you made for advanced configuration.

Figure 103 Create Virtual Drive - Summary Window



7. Click **Back** to return to the previous window to change any selections, or click **Finish** to accept and complete the configuration.

After you click **Finish**, the new storage configuration is created and initialized according to the selected options.

NOTE If you create a large configuration using drives that are in Power-Save mode, it can take several minutes to spin up the drives. A progress bar appears as the drives spin up. If any of the selected unconfigured drives fail to spin up, a dialog appears that identifies the drives.

After the configuration is completed, a dialog notifies you that the virtual drives were created successfully.

8. Click **OK**.

The **Enable SSD Caching on New Virtual Drives** dialog appears.

The newly created virtual drive is enabled for SSD caching by default.

9. Click **OK** to confirm SSD caching on the virtual drive. Click **No** if you want to disable SSD caching on the virtual drive.

The **All** check box is selected by default. To disable SSD caching on the virtual drives, deselect the **All** check box. If more drive capacity exists, the dialog asks whether you want to create more virtual drives. If no more drive capacity exists, you are prompted to close the configuration session.

10. Select either **Yes** or **No** to indicate whether you want to create additional virtual drives.

If you select **Yes**, the system takes you to the [Create Virtual Drive - Drive group and Virtual drive settings Dialog](#). If you select **No**, the utility asks whether you want to close the wizard.

11. If you selected **No** in the previous step, select either **Yes** or **No** to indicate whether you want to close the Wizard. If you select **Yes**, the **Configuration** wizard closes. If you select **No**, the dialog closes, and you remain on the same page.

5.2 Creating Hot Spare Drives

Hot spares are drives that are available to automatically replace failed drives in a RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, or RAID 60 virtual drive. *Dedicated hot spares* can be used to replace failed drives in a selected drive group only. *Global hot spares* are available to any virtual drive on a specific controller.

To create a dedicated or global hot spare drive, follow these steps:

1. Select the **Physical** tab in the left panel of the MegaRAID Storage Manager main menu, and click the icon of an unused drive.

For each drive, the window displays the port number, enclosure number, slot number, drive state, drive capacity, and drive manufacturer.

2. Either select **Go To > Physical Drive > Assign Global Hot Spare**, or select **Go To > Physical Drive > Assign Dedicated Hot Spare**.
3. Perform one of these actions:

- If you selected **Assign Dedicated Hotspare**, select a drive group from the list that appears. The hot spare is dedicated to the drive group that you select.

NOTE If the controller supports High Availability DAS, dedicated hot spares can be assigned to only one drive group. If you try to assign dedicated hot spares to more than one drive group, an error message appears.

- If you selected **Assign Global Hotspare**, skip this step, and go to the next step. The hot spare is available to any virtual drive on a specific controller.

4. Click **Go** to create the hot spare.

The drive state for the drive changes to dedicated or global hot spare, depending on your selection.

5.3 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

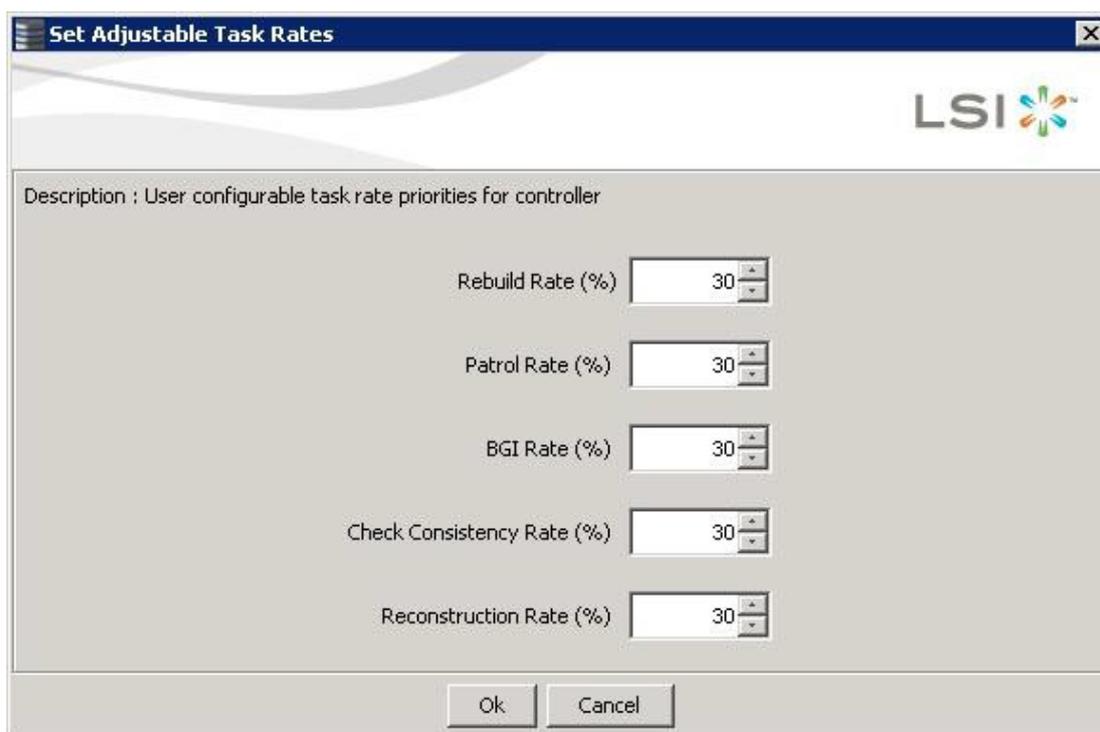
might run very slowly and might not complete within a reasonable time. If you decide to change the values, record the original default values here so you can restore them later, if necessary: **Rebuild rate:** _____, **Background Initialization (BGI) rate:** _____, **Check consistency rate:** _____.

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 104 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.

5.4 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)

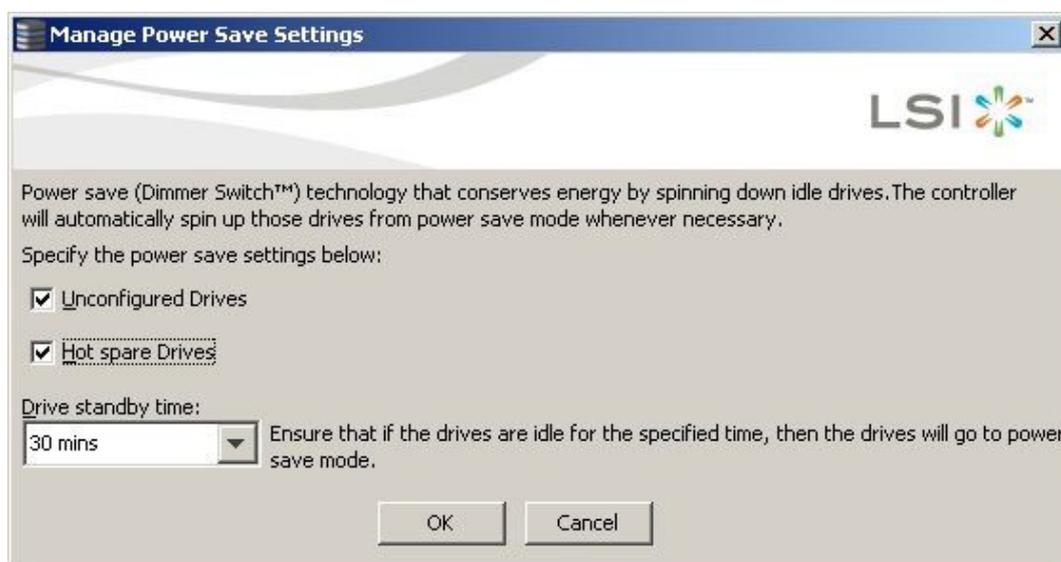
- Update of disk data format (DDF) and other metadata when you make changes to RAID configurations (Commissioned Hotspare drives and unconfigured drives)

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 105 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.

5.5 Recovering and Clearing Punctured Block Entries

You can recover and clear the punctured block area of a virtual drive.

ATTENTION This operation removes any data stored on the physical drives. Back up the good data on the drives before making any changes to the configuration.

When a Patrol Read or a Rebuild operation encounters a media error on the source drive, it punctures a block on the target drive to prevent the use of the data with the invalid parity. Any subsequent read operation to the punctured block completes but with an error. Consequently, the puncturing of a block prevents any invalid parity generation later while using this block.

To recover or clear the punctured block area of a virtual drive, run a Slow (or Full) Initialization to zero out and regenerate new parity causing all bad block entries to be removed from the bad block table.

To run a Slow (or Full) Initialization, see [Selecting Virtual Drive Settings](#).

5.6 Changing Virtual Drive Properties

You can change the read policy, write policy, and other virtual drive properties at any time after a virtual drive is created.

ATTENTION Do not enable drive caching on a mirrored drive group (RAID 1 or RAID 1E). If you do, data can be corrupted or lost in the event of a sudden power loss. A warning appears if you try to enable drive caching for a mirrored drive group.

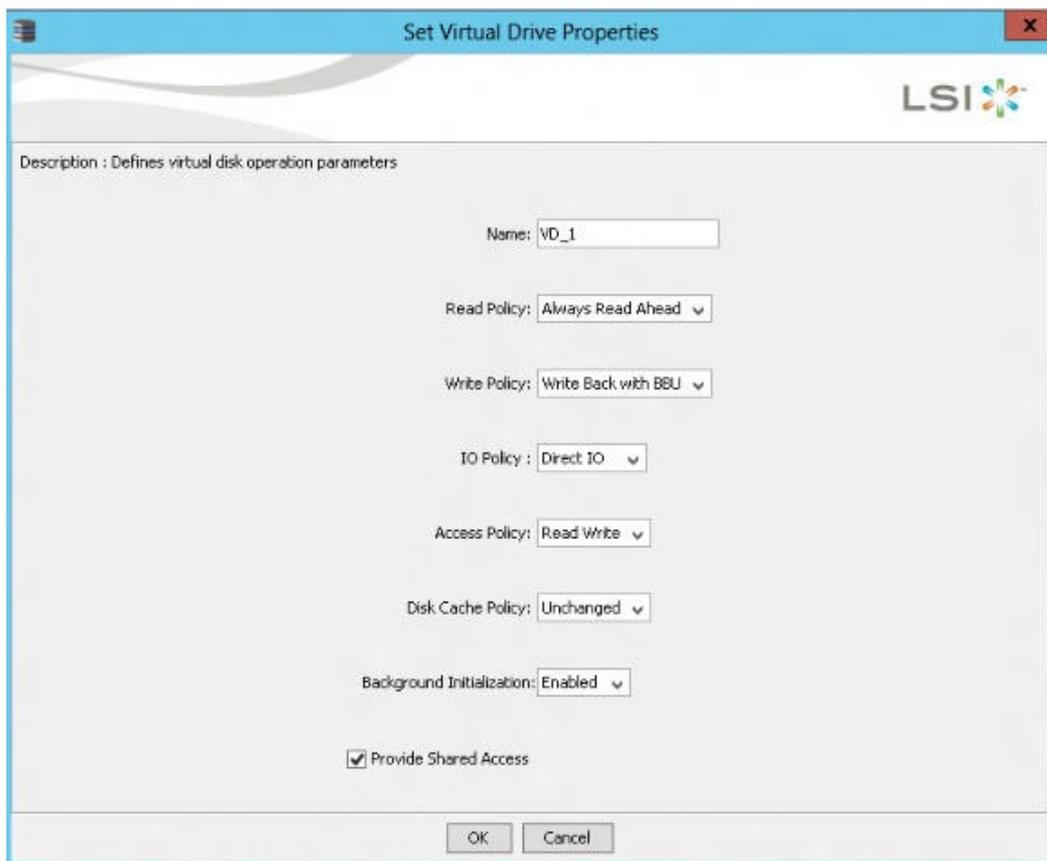
NOTE For virtual drives with SAS drives only, set the drive write cache policy set to **Disabled**, by default. For virtual drives with SATA drives only, set the drive write cache policy to **Enabled**, by default.

To change the virtual drive properties, perform the following steps:

1. Select a virtual drive icon in the **Physical** tab or the **Logical** tab in the left panel of the **MegaRAID Storage Manager** window.
2. Select **Go To > Virtual Drive > Set Virtual Drive Properties** from the menu bar.

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

Figure 106 Set Virtual Drive Properties Dialog



NOTE If the controller supports High Availability DAS, the **Provide Shared Access** check box appears in the above dialog. Select this option if you want the virtual drive to be shared between the two servers in a cluster.

3. Change the virtual drive properties as needed.
For information about these properties, see [Selecting Virtual Drive Settings](#).
4. Click **OK** to accept the changes. The virtual drive settings are updated.

5.7 Changing a Virtual Drive Configuration

You can remove drives from its drive group or replace the drives using configuration in the MegaRAID Storage Manager software.

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.)

5.7.1 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.

5.7.2 Replacing a Drive

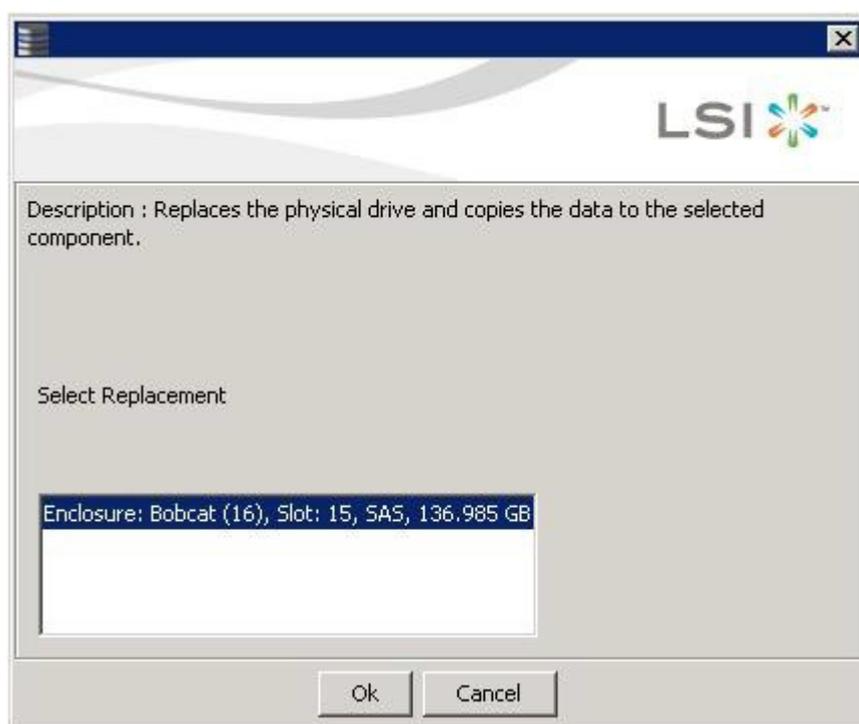
ATTENTION Be sure to back up the data on the virtual drive before you replace a drive.

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 107 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.

5.8 Deleting a Virtual Drive

CAUTION Make sure to back up the data that is on the virtual drive before you delete it. Make sure that the operating system is not installed on this virtual drive.

You can delete virtual drives to rearrange the storage space. To delete a virtual drive, follow these steps.

1. Back up all user data that is on the virtual drive you want to delete.
2. On the **MegaRAID Storage Manager** window, select the **Logical** tab, and click the icon of the virtual drive you want to delete.
3. Select **Go To > Virtual Drive > Delete Virtual Drive**.
4. When the warning messages appear, click **Yes** to confirm that you want to delete the virtual drive.

NOTE You are asked twice if you want to delete a virtual disk to avoid deleting the virtual disk by mistake.

Chapter 6 Monitoring Controllers and Their Attached Devices

This chapter explains how to use the MegaRAID Storage Manager software to monitor the status of drives, virtual drives, and other storage devices.

The MegaRAID Storage Manager software enables you to monitor the activity of all the controllers present in the system and the devices attached to them. It 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.

6.1 Alert Delivery Methods

Based on the severity level (Information, Warning, Critical and Fatal), the default alert delivery methods change. By default, each severity level has one or more alert delivery methods configured for it, as shown in the following table. To modify these alert delivery methods, see [Configuring Alert Notifications](#). The different alert delivery methods are as follows:

- Vivaldi Log/MegaRAID Storage Manager Log
- System Log
- Pop-up Notification
- Email Notification

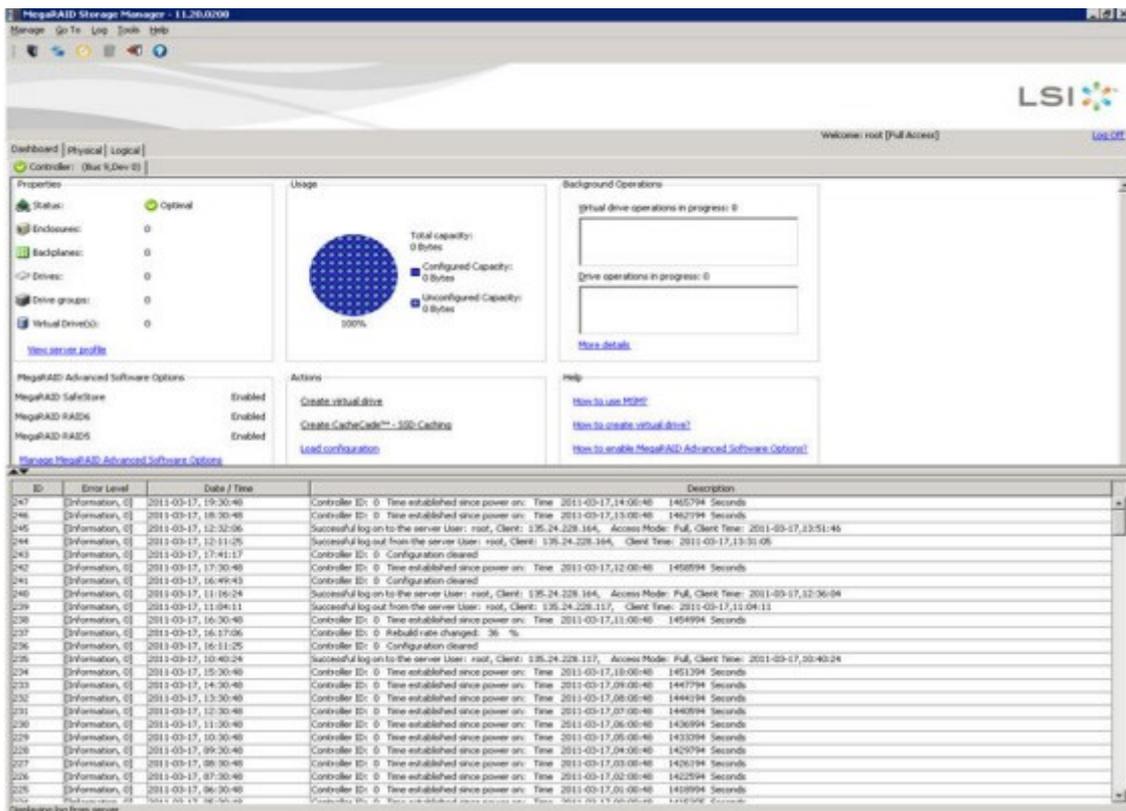
Table 21 Severity Level and Default Alert Delivery Methods

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.

6.1.1 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 108 Vivald 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 [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.

6.1.2 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`.

6.1.3 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.

Figure 109 Pop-up Notification

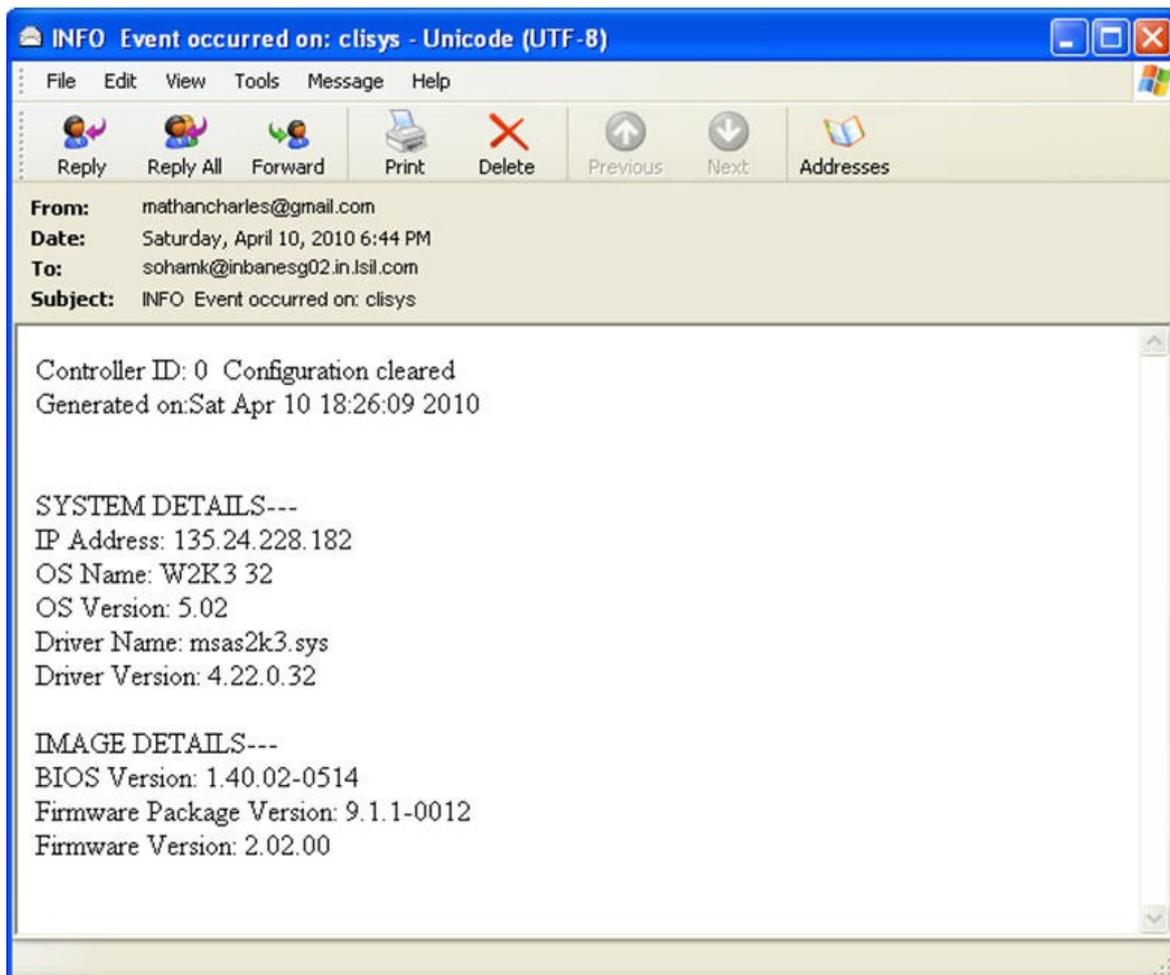


6.1.4 Email Notification

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.

Figure 110 Email Notification



6.2 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.

The **Configure Alerts** window appears, as shown in the following figure. The window contains three tabs: **Alert Settings**, **Mail Server**, and **Email**.

Figure 111 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.

6.3 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 112 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.

6.4 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 113 Change Individual Events



3. Click an event in the list to select it. The current alert delivery methods appear for the selected event in the **Alert Delivery Methods** frame.
4. Select the desired alert delivery methods for the event.
5. Click **OK** to return to the **Configure Alerts** window.
6. You may click **Cancel** to discard your current changes and to go back to the **Configure Alerts** window.
7. In the **Configure Alerts** window, click **OK**.

NOTE

You can click **Restore Defaults** to revert back to the default alert delivery method and the default severity level of an individual event. For more information, see [Roll Back to Default Individual Event Configuration](#).

6.5 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 114 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.

6.6 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, as shown in [Change Individual Events](#). 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.

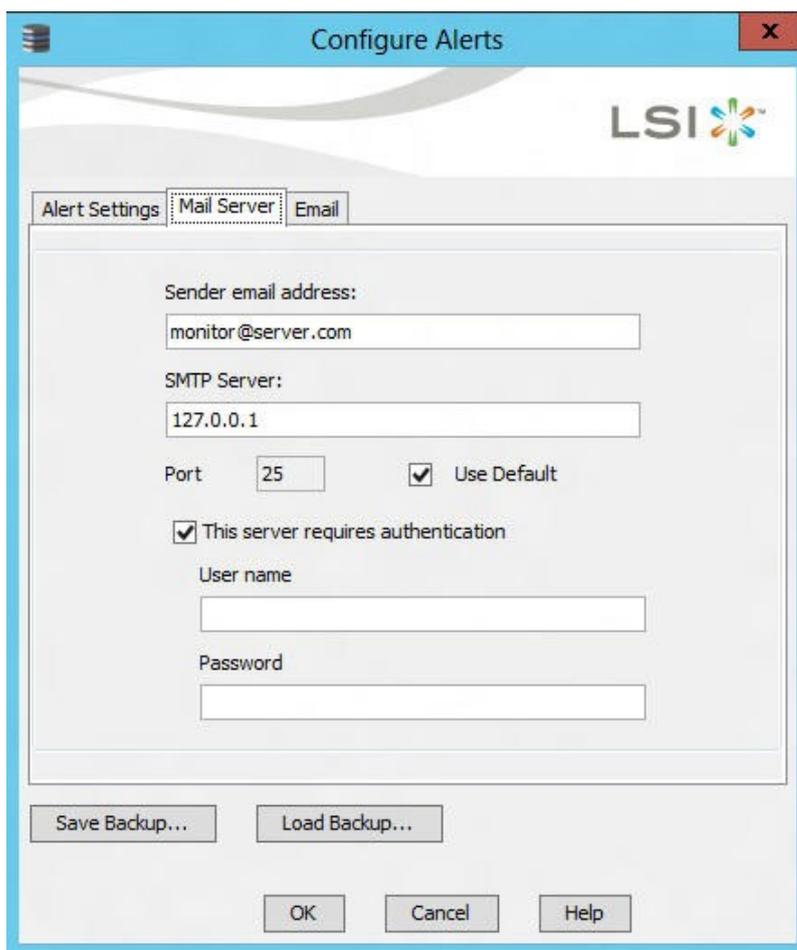
6.7 Entering or Editing the Sender Email Address and SMTP Server

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 115 Mail Server Options



The screenshot shows the 'Configure Alerts' window with the 'Mail Server' tab selected. The 'Sender email address' field contains 'monitor@server.com'. The 'SMTP Server' field contains '127.0.0.1'. The 'Port' field contains '25' and the 'Use Default' checkbox is checked. The 'This server requires authentication' checkbox is also checked. Below these are empty fields for 'User name' and 'Password'. At the bottom of the window are buttons for 'Save Backup...', 'Load Backup...', 'OK', 'Cancel', and 'Help'.

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**.

6.8 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 [Mail Server Options](#).

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**.

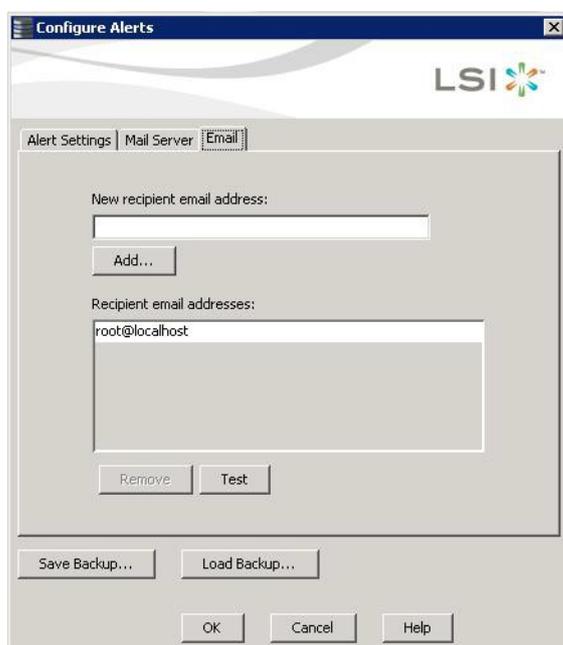
6.9 Adding Email Addresses of Recipients of Alert Notifications

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 116 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.

6.10 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 233](#).
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.

6.11 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 233](#).

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.

6.12 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.

6.13 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.

6.14 Monitoring Server Events

The MegaRAID Storage Manager software enables you to monitor the activity of MegaRAID Storage Manager users in the network.

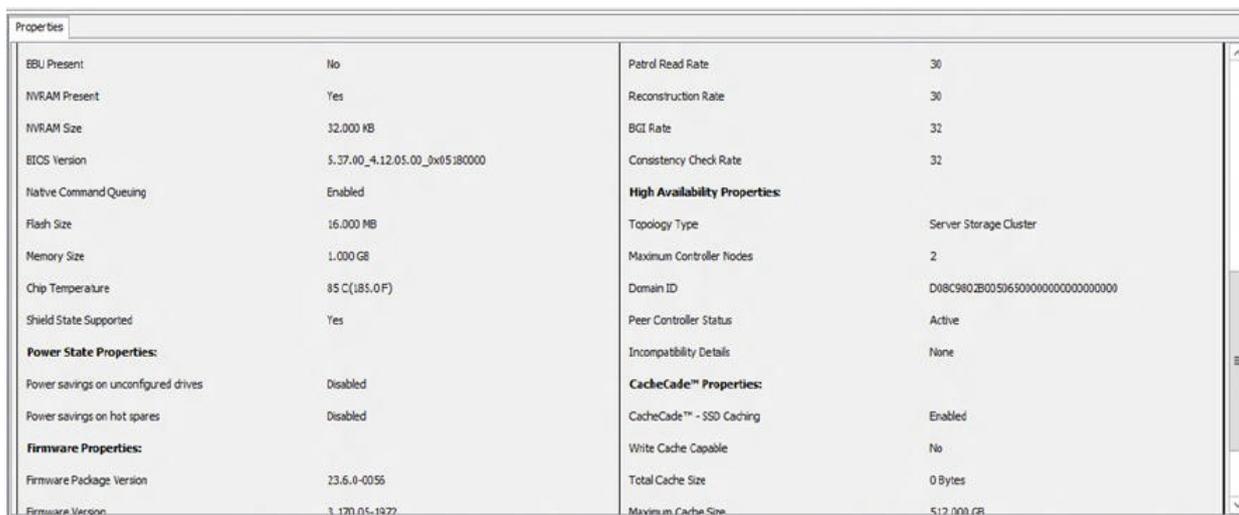
When a user logs on/logs off from the application, the event message appears in the log displayed at the bottom of the MegaRAID Storage Manager screen (the Vivaldi log/MegaRAID Storage Manager Log). These event messages have a severity level, a date and timestamp (User log on / log off time), and a brief description that contains a user name, client IP address, an access mode (full/view only) and a client system time.

6.15 Monitoring Controllers

When the MegaRAID Storage Manager software is running, you can see the status of all the controllers in the left panel. If a controller is operating normally, the controller icon looks like this: . If a controller has failed, a small red circle appears next to the icon.

To display the complete controller information, click on a controller icon in the left panel of the MegaRAID Storage Manager main menu. The controller properties appear in the right panel as shown in the following figure. Most of the information on this tab is self-explanatory.

Figure 117 Controller Properties



EBU Present	No	Patrol Read Rate	30
NVRAM Present	Yes	Reconstruction Rate	30
NVRAM Size	32,000 KB	BGI Rate	32
BIOS Version	5.37.00_4.12.05.00_3x05180000	Consistency Check Rate	32
Native Command Queuing	Enabled	High Availability Properties:	
Flash Size	16,000 MB	Topology Type	Server Storage Cluster
Memory Size	1,000 GB	Maximum Controller Nodes	2
Chip Temperature	85 C (185.0 F)	Domain ID	D08C9802B00506500000000000000000
Shield State Supported	Yes	Peer Controller Status	Active
Power State Properties:		Incompatibility Details	None
Power savings on unconfigured drives	Disabled	CacheCade™ Properties:	
Power savings on hot spares	Disabled	CacheCade™ - SSD Caching	Enabled
Firmware Properties:		Write Cache Capable	No
Firmware Package Version	23.6.0-0055	Total Cache Size	0 Bytes
Firmware Version	1.120.05-1927	Maximum Cache Size	5,12,000 GB

In the above dialog, the following properties appear under the **High Availability Properties** heading if the controller supports High Availability DAS:

- **Topology Type** - Indicates whether clustering is supported or not on the controller. Possible values for this field are **Server Storage Cluster**, or **None**.
- **Maximum Controller Nodes** - Indicates the total number of servers in a cluster.
- **Domain ID** - Shows the domain ID of the two servers in a cluster. The domain ID for both the servers is the same.
- **Peer Controller Status** - Indicates if both the servers in a cluster are running or not. The possible values are **Active**, **Inactive**, or **Incompatible**.
- **Incompatibility Details** - Indicates the reason for the incompatibility between the servers in a cluster. The possible values are **FW Level Mismatch**, **HW Incompatibility**, **Controller Property Mismatch**, **Premium Features Mismatch**, or **None**.

NOTE If the controller does not support High Availability DAS, only the **Topology Type** property appears under the **High Availability Properties** heading, with the value **None**.

The Rebuild rate, Patrol read rate, Reconstruction rate, Consistency check rate, and BGI rate (background initialization) are all user selectable. For more information, see [Changing Adjustable Task Rates](#). The **BBU Present** field indicates whether a battery backup unit is installed.

The **Alarm Enabled** field indicates whether the controller has an alarm to alert the user with an audible tone when there is an error or a problem on the controller. Options are available for disabling or silencing the alarm by right clicking on a controller icon or by selecting **Go To > Controller** menu.

The controller properties are defined in the [Glossary](#).

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**.

6.17 Running a Patrol Read

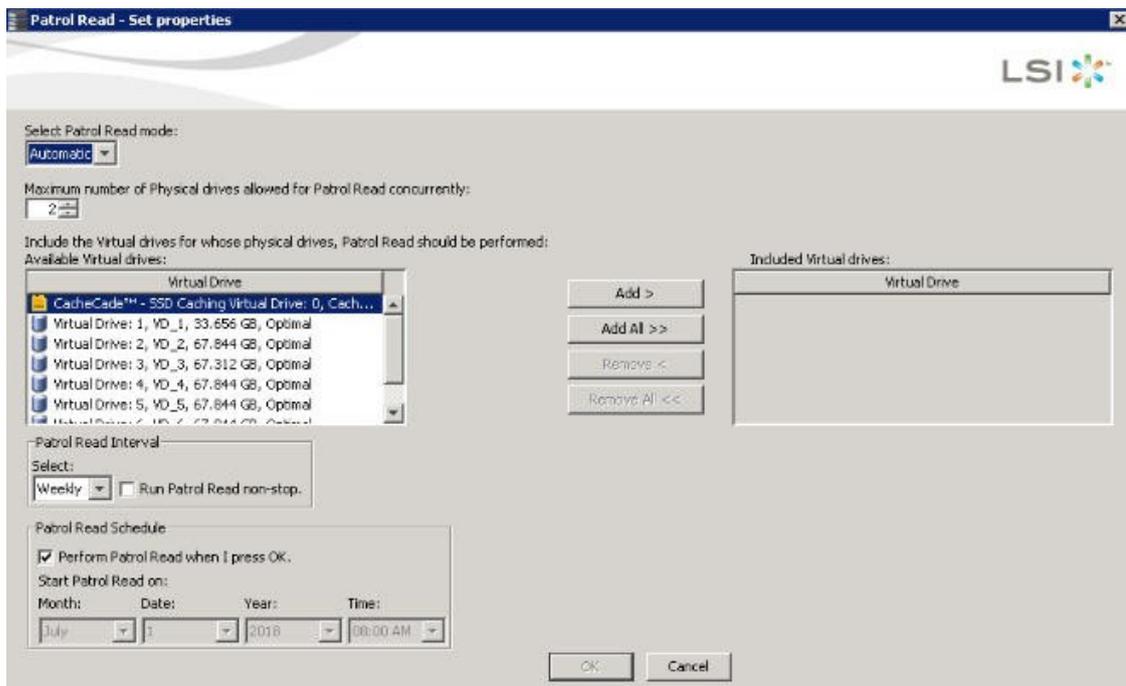
A patrol read periodically verifies all sectors of the drives connected to a controller, including the system reserved area in the RAID configured drives. You can run a patrol read for all RAID levels and for all hot spare drives. A patrol read is initiated only when the controller is idle for a defined 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 119 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.

6.17.1 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.

6.18 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 120 Virtual Drive Properties

Properties	
General:	
RAID Level	1
Name	VD_1
GUID	4C53492020202000105B000010659209316E3D1163984A
Host Access Policy	Shared
Size	278.875 GB
Mirror Data Size	278.875 GB
Strip Size	256 KB
Virtual Disk State	Optimal
IO and Cache Policies:	
Disk Cache Policy	Unchanged
Read Policy	Always Read Ahead
IO Policy	Direct IO
Write Policy:	
Current Write Policy	Write Through
Default Write Policy	Write Back with BBU
Access Policy:	
Current Access Policy	Read Write
Default Access Policy	Read Write
Drive Security Properties:	
Secured	No

If High Availability DAS is supported on the controller, two additional virtual drive properties, **GUID** and **Host Access Policy** appear on the Properties page.

- **GUID** - Indicates a unique ID assigned to this virtual drive by the firmware.
- **Host Access Policy** - Indicates whether or not the virtual drive is shared between the servers in a cluster. The values for this property are **Shared**, **Exclusive**, and **Exclusive to Peer Controller**.

You can change the read policy, write policy, and other virtual drive properties. To change these properties, see [Changing Virtual Drive Properties](#).

NOTE You can change the Read Policy, Write Policy, and other virtual drive properties by selecting the virtual drive icon and then selecting **Go To > Virtual Drive > Set Virtual Drive Properties** in the menu bar.

If the drives in the virtual drive are in a disk enclosure, you can identify them by making their LEDs blink. To identify the drives, follow these steps:

1. Click the virtual drive icon in the left panel.
2. Either select **Go To > Virtual Drive > Start Locating Virtual Drive**, or right-click a virtual drive and select **Start Locating Virtual Drive** from the menu.

The LEDs on the drives in the virtual drive start blinking (except for the hot spare drives).

3. To stop the LEDs from blinking, select **Go To > Virtual Drive > Stop Locating Virtual Drive**, or right-click a virtual drive and select **Stop Locating Virtual Drive** from the menu.

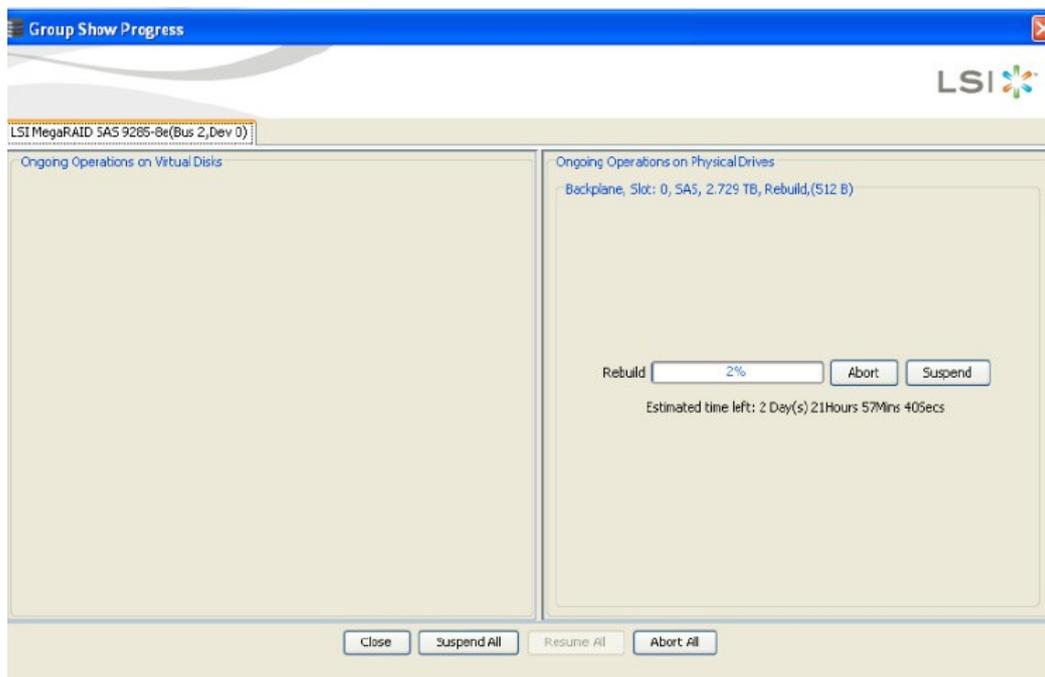
6.19 Monitoring Rebuilds and Other Processes

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 121 Group Show Progress Window



The **Group Show Progress** window displays a percent-complete indicator for drive rebuilds. Rebuilds might take a long time to complete. An up-arrow appears above the drive icon while it is being rebuilt.

Operations on virtual drives appear in the left panel of the window, and operations on drives appear in the right panel. The type of operations that appear in this window are as follows:

- Initialization of a virtual drive
- Rebuild
- Consistency check
- Non FDE Physical Drive Erase
- Virtual Drive Erase
- Patrol Read
- LD Reconstruction
- LD Disassociate
- PD Clear
- Replace
- Background Initialization (BGI)

A Modify Drive Group process cannot be aborted. To abort any other ongoing process, click the **Abort** button next to the status indicator. Click **Abort All** to abort all ongoing processes. Click **Close** to close the window.

Chapter 7 Maintaining and Managing Storage Configurations

This chapter explains how to use the MegaRAID Storage Manager software to maintain and manage storage configurations. Log on to the server in Full Access mode to perform the maintenance and management tasks.

7.1 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.

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

4. Select the **Fast Initialization** check box if you want to use this option.

If you leave the box unselected, the MegaRAID Storage Manager software runs a Full Initialization on the virtual drive. (For more information, see [Selecting Virtual Drive Settings](#).)

5. Click **Start** to begin the initialization.

You can monitor the progress of the initialization. See [Monitoring Rebuilds and Other Processes](#) for more information.

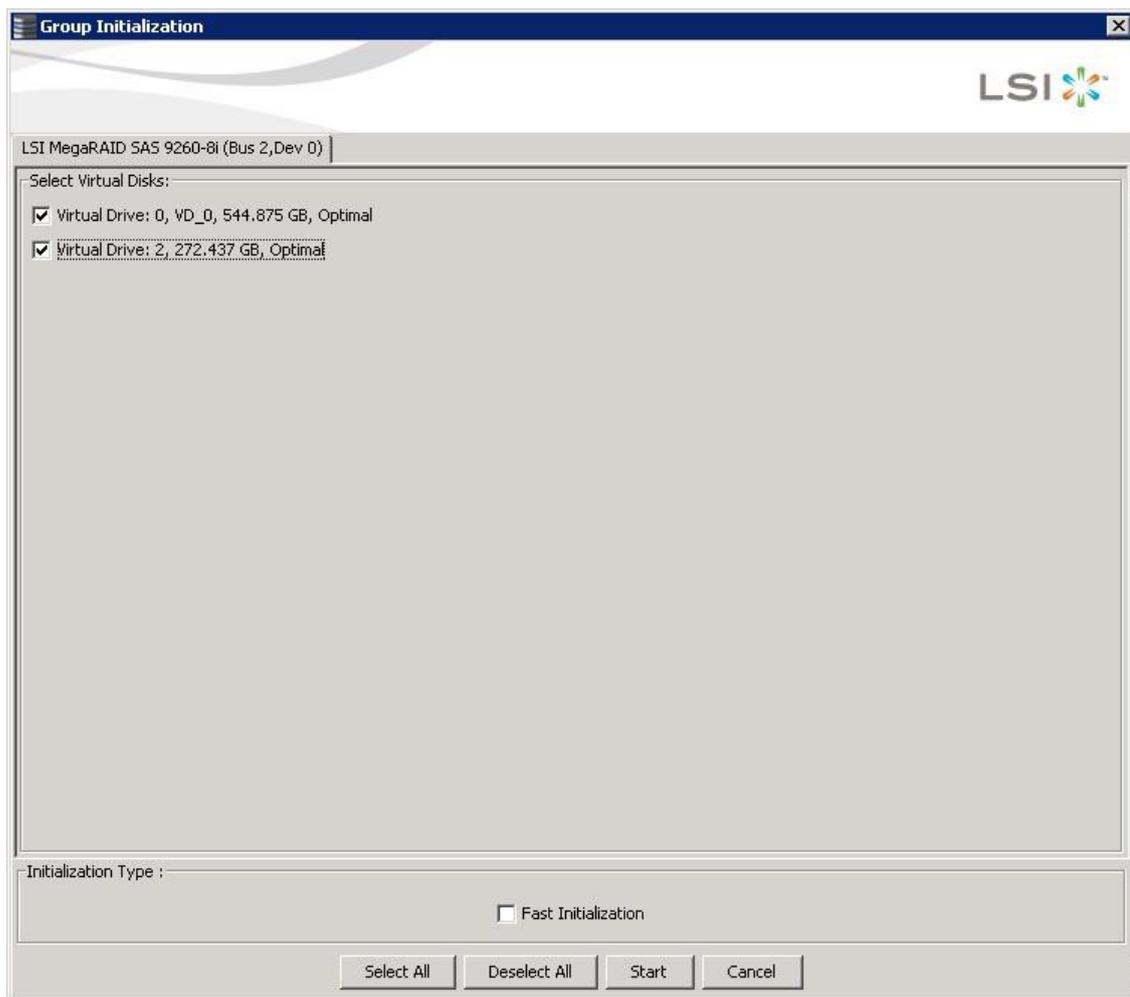
7.1.1 Running a Group Initialization

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 122 Group Initialization Dialog



2. Either check the virtual drives on which to run the initialization, or click **Select All** to select all of the virtual drives.
3. Click **Start**.

You can monitor the progress of the group initialization. See [Monitoring Rebuilds and Other Processes](#) for more information.

7.2 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.

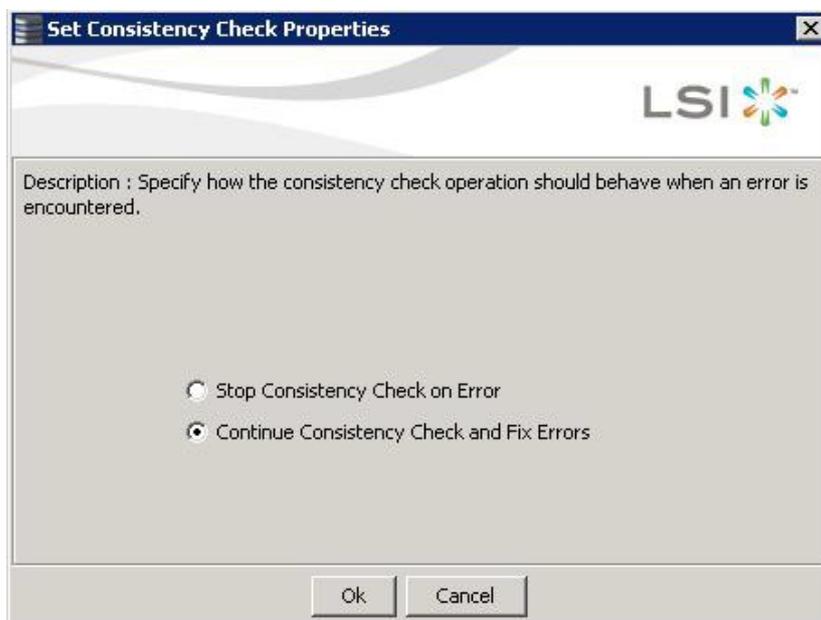
7.2.1 Setting the Consistency Check Settings

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 123 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**.

7.2.2 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 124 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

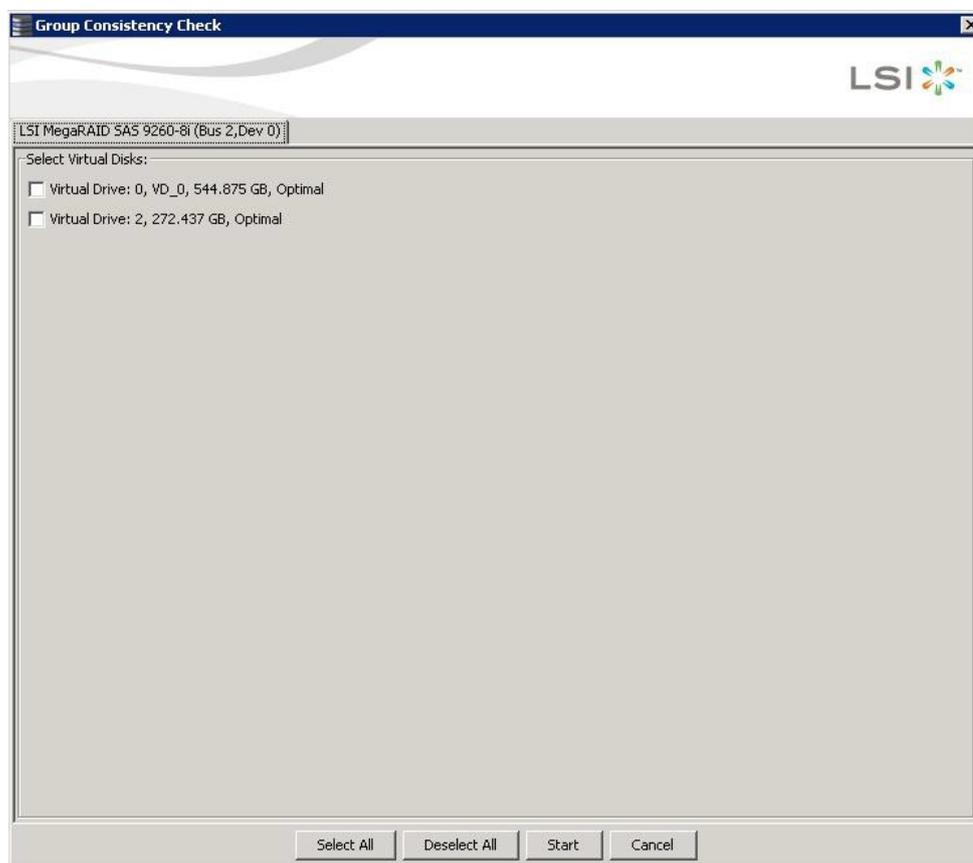
3. Perform the following steps to schedule the consistency check:
 - a. Select how often to run the consistency check from the drop-down list.
You can click **Advanced** for more detailed date options.
 - b. (Optional) Select the **Run consistency check continuously** check box.
 - c. Select the month, day, and year on which to start the consistency check.
 - d. Select the time of day to start the consistency check.
4. Click **Ok**.
You can monitor the progress of the consistency check. See [Monitoring Rebuilds and Other Processes](#) for more information.

7.2.3 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 125 Group Consistency Check Dialog



2. Either check the virtual drives on which to run the consistency check, or click **Select All** to select all of the virtual drives.
3. Click **Start**.

You can monitor the progress of the group consistency check. See [Monitoring Rebuilds and Other Processes](#) for more information.

7.3 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.

7.4 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.

7.5 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.

7.6 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**.

7.7 Upgrading Firmware

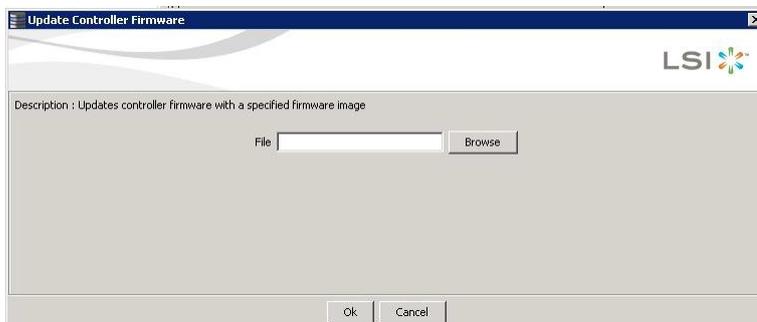
The MegaRAID Storage Manager software enables you to easily upgrade the controller firmware.

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 126 Update Controller Firmware Dialog



4. After you locate the file, click **Open**.
5. When you are prompted to indicate whether you want to upgrade the firmware, click **Yes**.

The MegaRAID Storage Manager software displays the version of the existing firmware.

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.

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 22 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 23 Event Messages

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0000	Information	MegaRAID firmware initialization started (PCI ID %04x/%04x/%04x/%04x)	Logged at firmware initialization.
0x0001	Information	MegaRAID firmware version %s	Logged at firmware initialization to display firmware
0x0002	Fatal	Unable to recover cache data from TBBU	Currently not logged.
0x0003	Information	Cache data recovered from TBBU successfully	Currently not logged.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
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.
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 updation 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.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
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.
0x0024	Information	Patrol Read paused	Logged when patrol read is paused.
0x0025	Information	Patrol Read Rate changed to %d%%	Logged to indicate progress of patrol read in percentage.
0x0026	Information	Patrol Read resumed	Logged when patrol read is resumed.
0x0027	Information	Patrol Read started	Logged when patrol read is started.
0x0028	Information	Reconstruction rate changed to %d%%"	Logged to indicate progress of reconstruction in percentage.
0x0029	Information	Drive group modification rate changed to %d%%	Logged to indicate the change in Drive group modification frequency.
0x002a	Information	Shutdown command received from host	Logged when shutdown command is received from host to controller.
0x002b	Information	Test event: %s	General controller event, with a generic string.
0x002c	Information	Time established as %s; (%d seconds since power on)	Logged when controller time was set from host, also displaying time since power on in seconds.
0x002d	Information	User entered firmware debugger	Logged when user enters controller debug shell.
0x002e	Warning	Background Initialization aborted on %s	Logged to inform about user aborted background initialization on displayed LD number.
0x002f	Warning	Background Initialization corrected medium error (%s at %lx	logged to inform about corrected medium error on displayed LD number, LBALBA number, PD number and PDLBA number in that order.
0x0030	Information	Background Initialization completed on %s	Logged to inform Background Initialization completion on displayed LD.
0x0031	Fatal	Background Initialization completed with uncorrectable errors on %s	Logged to inform Background Initialization completion with error on displayed LD.
0x0032	Fatal	Background Initialization detected uncorrectable double medium errors (%s at %lx on %s)	Logged to inform Background Initialization completion with double medium error on displayed PD, PDLBA and LD in that order.
0x0033	Critical	Background Initialization failed on %s	Logged to inform Background Initialization failure on displayed LD.
0x0034	Progress	Background Initialization progress on %s is %s	Logged to inform Background Initialization progress in percentage of displayed LD.
0x0035	Information	Background Initialization started on %s	Logged to inform Background Initialization started for displayed LD.
0x0036	Information	Policy change on %s from %s to %s	Logged to inform the changed policy for displayed LD with old and new policies.
0x0038	Warning	Consistency Check aborted on %s	Logged to inform aborted Consistency check for displayed LD.
0x0039	Warning	Consistency Check corrected medium error (%s at %lx	Logged when Consistency check corrected medium error
0x003a	Information	Consistency Check done on %s	Logged when Consistency check has completed successfully on the LD.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
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.
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
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 (un recoverable 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.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
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
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.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
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.
0x0070	Information	Drive removed: %s	Logged when a drive is removed from the controller.
0x0071	Warning	Unexpected sense: %s, CDB%s, Sense: %s	Logged when an I/O fails due to unexpected reasons and sense data needs to be logged.
0x0072	Information	State change on %s from %s to %s	Logged when the state of a drive is changed by the firmware or by you.
0x0073	Information	State change by user on %s from %s to %s	Not logged by the firmware.
0x0074	Warning	Redundant path to %s broken	Not logged by the firmware.
0x0075	Information	Redundant path to %s restored	Not logged by the firmware
0x0076	Information	Dedicated Hot Spare Drive %s no longer useful due to deleted drive group	Not logged by the firmware.
0x0077	Critical	SAS topology error: Loop detected	Logged when device discovery fails for a SAS device as a loop was detected.
0x0078	Critical	SAS topology error: Unaddressable device	Logged when device discovery fails for a SAS device as an unaddressable device was found.
0x0079	Critical	SAS topology error: Multiple ports to the same SAS address	Logged when device discovery fails for a SAS device multiple ports with same SAS address were detected.
0x007a	Critical	SAS topology error: Expander error	Not logged by the firmware.
0x007b	Critical	SAS topology error: SMP timeout	Logged when device discovery fails for a SAS device due to SMP timeout.
0x007c	Critical	SAS topology error: Out of route entries	Logged when device discovery fails for a SAS device as expander route table is out of entries.
0x007d	Critical	SAS topology error: Index not found	Logged when device discovery fails for a SAS device as expander route table out of entries.
0x007e	Critical	SAS topology error: SMP function failed	Logged when device discovery fails for a SAS device due to SMP function failure.
0x007f	Critical	SAS topology error: SMP CRC error	Logged when device discovery fails for a SAS device due to SMP CRC error.
0x0080	Critical	SAS topology error: Multiple subtractive	Logged when device discovery fails for a SAS device as a subtractive-to-subtractive link was detected.
0x0081	Critical	SAS topology error: Table to table	Logged when device discovery fails for a SAS device as table-to-table link was detected.
0x0082	Critical	SAS topology error: Multiple paths	Not logged by the firmware.
0x0083	Fatal	Unable to access device %s	Logged when the inserted drive is bad and unusable.
0x0084	Information	Dedicated Hot Spare created on %s (%s)	Logged when a drive is configured as a dedicated spare.
0x0085	Information	Dedicated Hot Spare %s disabled	Logged when a drive is removed as a dedicated spare.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
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 hot spare 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.
0x008b	Information	Deleted %s}	Logged when the firmware removes an LD from its 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
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	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

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
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.
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	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.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
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.
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
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
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.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
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.
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
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
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.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
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
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
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.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
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.
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.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
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.
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.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
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
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	Logged when BBU periodic relearn is pending.
0x0134	Information	Drive security key created	Logged when controller lock key is created.
0x0135	Information	Drive security key backed up	Logged when controller lock key is backed up.
0x0136	Information	Drive security key from escrow, verified	Logged when controller lock key is verified from escrow.
0x0137	Information	Drive security key changed	Logged when controller lock key is re-keyed.
0x0138	Warning	Drive security key, re-key operation failed	Logged when controller lock re-key operation failed.
0x0139	Warning	Drive security key is invalid	Logged when the controller lock is not valid.
0x013a	Information	Drive security key destroyed	Logged when the controller lock key is destroyed.
0x013b	Warning	Drive security key from escrow is invalid	Logged when the controller escrow key is not valid. This escrow key can not unlock any drive.
0x013c	Information	VD %s is now secured	Logged when secure LD is created.
0x013d	Warning	VD %s is partially secured	Logged when all the drives in the array are not secure.
0x013e	Information	PD %s security activated	Logged when PD security key is set.
0x013f	Information	PD %s security disabled	Logged when security key is removed from an FDE drive.
0x0140	Information	PD %s is reprovisioned	Logged when PD security is cleared.
0x0141	Information	PD %s security key changed	Logged when PD lock key is re-keyed.
0x0142	Fatal	Security subsystem problems detected for PD %s	Logged when PD security can not be set.
0x0143	Fatal	Controller cache pinned for missing or offline VD %s	Logged when LD cache is pinned.
0x0144	Fatal	Controller cache pinned for missing or offline VD: %s	Logged when pinned cache is found during OCR.
0x0145	Information	Controller cache discarded by user for VD: %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.

Number	SeverityLevel	Event Text	Generic Conditions when Each Event Occurs
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
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.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
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
0x016f	Information	COD on %s updated as it was stale	Logged when COD in DDF is updated due to various
0x0170	Warning	Power state change failed on %s (from %s to %s)	Logged when power state change failed on a PD.
0x0171	Warning	%s is not available	Logged when a LD was not available.
0x0172	Information	%s is available	Logged when a LD was available.
0x0173	Information	%s is used for CacheCade with capacity 0x%lx logical blocks	Logged when a LD was used for CacheCade with the indicated capacity in logical blocks.
0x0174	Information	%s is using CacheCade %s	Logged when a LD was using CacheCade.
0x0175	Information	%s is no longer using CacheCade %s	Logged when a LD was no longer using CacheCade.
0x0176	Critical	Snapshot deleted due to resource constraints for %s in snapshot	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
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
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.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0184	Warning	%s cannot transition to max power savings	Logged when LD cannot transition to max power savings
0x0185	Information	Host driver is loaded and operational	This event is not reported to the user.
0x0186	Information	%s mirror broken	Logged when the mirror is broken for an LD.
0x0187	Information	%s mirror joined	Logged when joining the LD with its broken mirror.
0x0188	Warning	%s link %d failure in wide port	This event is not reported to the user.
0x0189	Information	%s link %d restored in wide port	This event is not reported to the user.
0x018a	Information	Memory module FRU is %s	This event is not reported to the user.
0x018b	Warning	Cache-vault power pack is sub-optimal. Please replace the pack.	This event is not reported to the user.
0x018c	Warning	Foreign configuration auto-import did not import any drives	Logged when the Foreign configuration auto-import did not import any drives.
0x018d	Warning	Cache-vault microcode update required	Logged when the BMU is not in Normal mode and Cache-vault microcode update required.
0x018e	Warning	CacheCade (%s) capacity exceeds maximum allowed size, extra capacity is not used	Logged when CacheCade capacity exceeds maximum allowed size, extra capacity is not used.
0x018f	Warning	LD (%s) protection information lost	Logged when the protection information is lost for an LD.
0x0190	Information	Diagnostics passed for %s	Logged when the SHIELD Diagnostics passed for a PD.
0x0191	Critical	Diagnostics failed for %s	Logged when the SHIELD Diagnostics failed for a PD.
0x0192	Information	Server Power capability Diagnostic Test Started	Logged when the Server Power capability Diagnostic Test starts.
0x0193	Information	Drive Cache settings enabled during rebuild for %s	Logged when the Drive Cache settings enabled during rebuild for a PD.
0x0194	Information	Drive Cache settings restored after rebuild for %s	Logged when the Drive Cache settings restored after rebuild for a PD.
0x0195	Information	Drive %s commissioned as Emergency spare	Logged when the Drive commissioned as Emergency spare.
0x0196	Warning	Reminder: Potential non-optimal configuration due to drive %s commissioned as emergency spare	Logged when the PD being imported is an Emergency Spare.
0x0197	Information	Consistency Check suspended on %s	Logged when the Consistency Check is suspended on an LD.
0x0198	Information	Consistency Check resumed on %s	Logged when the Consistency Check is resumed on an LD.
0x0199	Information	Background Initialization suspended on %s	Logged when the Background Initialization is suspended on an LD.
0x019a	Information	Background Initialization resumed on %	Logged when the Background Initialization is resumed on an LD.
0x019b	Information	Reconstruction suspended on %s	Logged when the Reconstruction is suspended on an LD.
0x019c	Information	Rebuild suspended on %	Logged when the Rebuild is suspended on a PD.
0x019d	Information	Replace Drive suspended on %s	Logged when the Replace is suspended on a PD.
0x019e	Information	Reminder: Consistency Check suspended on %	Logged as a reminder when the Consistency Check is suspended on an LD.
0x019f	Information	Reminder: Background Initialization suspended on %s	Logged as a reminder when the Background Initialization is suspended on an LD.
0x01a0	Information	Reminder: Reconstruction suspended on %s	Logged as a reminder when the Reconstruction is suspended on an LD.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x01a1	Information	Reminder: Rebuild suspended on %s	Logged as a reminder when the Rebuild is suspended on a PD.
0x01a2	Information	Reminder: Replace Drive suspended on %s	Logged as a reminder when Replace is suspended on a PD
0x01a3	Information	Reminder: Patrol Read suspended	Logged as a reminder when the Patrol Read is suspended.
0x01a4	Information	Erase aborted on %s	Logged when the Erase is aborted on a PD.
0x01a5	Critical	Erase failed on %s (Error %02x)	Logged when the Erase is failed on a PD along with the error.
0x01a6	Progress	Erase progress on %s is %s	Logged to display the Erase progress on a PD along with its current progress.
0x01a7	Information	Erase started on %s	Logged when Erase is started on a PD.
0x01a8	Information	Erase completed on %s	Logged when the Erase is completed on a PD.
0x01a9	Information	Erase aborted on %s	Logged when the Erase is aborted on an LD.
0x01aa	Critical	Erase failed on %s	Logged when the Erase is failed on an LD.
0x01ab	Progress	Erase progress on %s is %s	Logged to display the Erase progress on an LD along with its current progress.
0x01ac	Information	Erase started on %s	Logged when the Erase is started on an LD.
0x01ad	Information	Erase complete on %s	Logged when the Erase is complete on an LD.
0x01ae	Warning	Potential leakage during erase on %s	Logged to inform the Potential leakage during erase on an LD.
0x01af	Warning	Battery charging was suspended due to high battery temperature	Logged when the Battery charging was suspended due to high battery temperature.
0x01b0	Information	NVCache firmware update was successful	This event is not reported to the user.
0x01b1	Warning	NVCache firmware update failed	This event is not reported to the user.
0x01b2	Fatal	%s access blocked as cached data in CacheCade is unavailable	This event is not reported to the user.
0x01b3	Information	CacheCade disassociate started on %s	This event is not reported to the user.
0x01b4	Information	CacheCade disassociate completed on %s	This event is not reported to the user.
0x01b5	Critical	CacheCade disassociate failed on %s	This event is not reported to the user.
0x01b6	Progress	CacheCade disassociate progress on %s is %s	This event is not reported to the user.
0x01b7	Information	CacheCade disassociate aborted by user on %s	This event is not reported to the user.
0x01b8	Information	Link speed changed on SAS port %d and PHY %d	Logged when the Link speed changed on SAS port and PHY.
0x01b9	Warning	Advanced Software Options was deactivated for - %s	This event is not reported to the user.
0x01ba	Information	%s is now accessible	This event is not reported to the user.
0x01bb	Information	%s is using CacheCade	This event is not reported to the user.
0x01bc	Information	%s is no longer using CacheCade	This event is not reported to the user.
0x01bd	Warning	Patrol Read aborted on %s	Logged when the Patrol Read is aborted on a PD.
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.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x01c3	Information	Controller reset requested by host	Logged when the Controller Reset process started on the corresponding controller.
0x01c4	Information	Controller reset requested by host, completed	Logged when the Controller Reset process completed on the corresponding controller.
0x01c7	Warning	Controller booted in headless mode with errors	Logged when the Controller is booted to safe mode due to warning errors.
0x01c8	Critical	Controller booted to safe mode due to critical errors	Logged when the Controller is booted to safe mode due to critical errors.
0x01c9	Warning	Warning Error during boot - %s	Logged when a warning error occurs during booting the controller to safe mode.
0x01ca	Critical	Critical Error during boot - %s	Logged when a critical error occurs during booting the controller to safe mode
0x01cb	Fatal	Fatal Error during boot - %s	Logged when a fatal error occurs during booting the controller to safe mode

Appendix B: Glossary

This appendix provides a glossary for terms used in this document.

A

Absolute state of charge	Predicted remaining battery capacity expressed as a percentage of Design Capacity. Note that the Absolute State of Charge operation can return values greater than 100 percent.
Access policy	A virtual drive property indicating what kind of access is allowed for a particular virtual drive. The possible values are <i>Read/Write</i> , <i>Read Only</i> , or <i>Blocked</i> .
Alarm enabled	A controller property that indicates whether the controller's onboard alarm is enabled.
Alarm present	A controller property that indicates whether the controller has an onboard alarm. If present and enabled, the alarm is sounded for certain error conditions.
Array	See <i>drive group</i> .
Auto learn mode	The controller performs the learn cycle automatically in this mode. This mode offers the following options: <ul style="list-style-type: none"> ■ BBU Auto Learn: Firmware tracks the time since the last learn cycle and performs a learn cycle when due. ■ BBU Auto Learn Disabled: Firmware does not monitor or initiate a learn cycle. You can schedule learn cycles manually. ■ BBU Auto Learn Warn: Firmware warns about a pending learn cycle. You can initiate a learn cycle manually. After the learn cycle is complete, the firmware resets the counter and warns you when the next learn cycle time is reached.
Auto learn period	Time between learn cycles. A learn cycle is a battery calibration operation performed periodically by the controller to determine the condition of the battery.
Average time to empty	One-minute rolling average of the predicted remaining battery life.
Average time to full	Predicted time to charge the battery to a fully charged state based on the one minute rolling average of the charge current.

B

Battery module version	Current revision of the battery pack module.
Battery replacement	Warning issued by firmware that the battery can no longer support the required data retention time.
Battery retention time	Time, in hours, that the battery can maintain the contents of the cache memory.

Battery status	Operating status of the battery. Possible values are Missing, Optimal, Failed, Degraded (need attention), and Unknown.
Battery type	Possible values are intelligent Battery Backup Unit (BBU), intelligent Battery Backup Unit (iBBU), intelligent Transportable Battery Backup Unit (iTBBU), and ZCR Legacy.
BBU present	A controller property that indicates whether the controller has an onboard battery backup unit to provide power in case of a power failure.
BGI rate	A controller property indicating the rate at which the background initialization of virtual drives will be carried out.
BIOS	Basic Input/Output System. The computer BIOS is stored on a flash memory chip. The BIOS controls communications between the microprocessor and peripheral devices, such as the keyboard and the video controller, and miscellaneous functions, such as system messages.

C

Cache	Fast memory that holds recently accessed data. Use of cache memory speeds subsequent access to the same data. When data is read from or written to main memory, a copy is also saved in cache memory with the associated main memory address. The cache memory software monitors the addresses of subsequent reads to see if the required data is already stored in cache memory. If it is already in cache memory (a cache hit), it is read from cache memory immediately and the main memory read is aborted (or not started). If the data is not cached (a cache miss), it is fetched from main memory and saved in cache memory.
Cache flush interval	A controller property that indicates how often the data cache is flushed.
Caching	The process of using a high speed memory buffer to speed up a computer system's overall read/write performance. The cache can be accessed at a higher speed than a drive subsystem. To improve read performance, the cache usually contains the most recently accessed data, as well as data from adjacent drive sectors. To improve write performance, the cache can temporarily store data in accordance with its write back policies.
Capacity	A property that indicates the amount of storage space on a drive or virtual drive.
Coerced capacity	A drive property indicating the capacity to which a drive has been coerced (forced) to make it compatible with other drives that are nominally the same capacity. For example, a 4-GB drive from one manufacturer might be 4,196 MB, and a 4-GB from another manufacturer might be 4,128 MB. These drives could be coerced to a usable capacity of 4,088 MB each for use in a drive group in a storage configuration.
Coercion mode	A controller property indicating the capacity to which drives of nominally

	identical capacity are coerced (forced) to make them usable in a storage configuration.
Consistency check	An operation that verifies that all stripes in a virtual drive with a redundant RAID level are consistent and that automatically fixes any errors. For RAID 1 drive groups, this operation verifies correct mirrored data for each stripe.
Consistency check rate	The rate at which consistency check operations are run on a computer system.
Controller	A chip that controls the transfer of data between the microprocessor and memory or between the microprocessor and a peripheral device such as a drive. RAID controllers perform RAID functions such as striping and mirroring to provide data protection.
Copyback	The procedure used to copy data from a source drive of a virtual drive to a destination drive that is not a part of the virtual drive. The copyback operation is often used to create or restore a specific physical configuration for a drive group (for example, a specific arrangement of drive group members on the device I/O buses). The copyback operation can be run automatically or manually. Typically, a drive fails or is expected to fail, and the data is rebuilt on a hot spare. The failed drive is replaced with a new drive. Then the data is copied from the hot spare to the new drive, and the hot spare reverts from a rebuild drive to its original hot spare status. The copyback operation runs as a background activity, and the virtual drive is still available online to the host.
Current	Measure of the current flowing to (+) or from (-) the battery, reported in milliamperes.
Current write policy	A virtual drive property that indicates whether the virtual drive currently supports Write Back mode or Write Through mode. <ul style="list-style-type: none"> ■ In Write Back mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. ■ In Write Through mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction.
Cycle count	The count is based on the number of times the near fully charged battery has been discharged to a level below the cycle count threshold.

D

Default write policy	A virtual drive property indicating whether the default write policy is Write Through or Write Back. In Write Back mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. In Write Through mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction.
Design capacity	Designed charge capacity of the battery, measured in milliampere-hour units (mAh).
Design charge capacity remaining	Amount of the charge capacity remaining, relative to the battery pack design capacity.
Design voltage	Designed voltage capacity of the battery, measured in millivolts (mV).
Device chemistry	Possible values are NiMH (nickel metal hydride) and LiON (lithium ion).
Device ID	A controller or drive property indicating the manufacturer-assigned device ID.
Device port count	A controller property indicating the number of ports on the controller.
Drive cache policy	A virtual drive property indicating whether the virtual drive cache is enabled, disabled, or unchanged from its previous setting.
Drive group	A group of drives attached to a RAID controller on which one or more virtual drives can be created. All virtual drives in the drive group use all of the drives in the drive group.
Drive state	<p>A drive property indicating the status of the drive. A drive can be in one of the following states:</p> <ul style="list-style-type: none"> ■ Unconfigured Good – A drive accessible to the RAID controller but not configured as a part of a virtual drive or as a hot spare. ■ Hot Spare – A drive that is configured as a hot spare. ■ Online – A drive that can be accessed by the RAID controller and will be part of the virtual drive. ■ Rebuild – A drive to which data is being written to restore full redundancy for a virtual drive. ■ Failed – A drive that was originally configured as Online or Hot Spare, but on which the firmware detects an unrecoverable error. ■ Unconfigured Bad – A drive on which the firmware detects an unrecoverable error; the drive was Unconfigured Good or the drive could not be initialized. ■ Missing – A drive that was Online, but which has been removed from its location. ■ Offline – A drive that is part of a virtual drive but which has invalid data as far as the RAID configuration is concerned. ■ None – A drive with an unsupported flag set. An Unconfigured Good or Offline drive that has completed the prepare for removal operation.

Drive state drive subsystem A collection of drives and the hardware that controls them and connects them to one or more controllers. The hardware can include an intelligent controller, or the drives can attach directly to a system I/O bus controller.

Drive type A drive property indicating the characteristics of the drive.

E

EKM External Key Management

Estimated time to recharge Estimated time necessary to complete recharge of the battery at the current charge rate

Expected margin of error Indicates how accurate the reported battery capacity is in terms of percentage.

F

Fast initialization A mode of initialization that quickly writes zeroes to the first and last sectors of the virtual drive. This allows you to immediately start writing data to the virtual drive while the initialization is running in the background.

Fault tolerance The capability of the drive subsystem to undergo a single drive failure per drive group without compromising data integrity and processing capability. LSI SAS RAID controllers provides fault tolerance through redundant drive groups in RAID levels 1, 5, 6, 10, 50, and 60. They also support hot spare drives and the auto-rebuild feature.

Firmware Software stored in read-only memory (ROM) or programmable ROM (PROM). Firmware is often responsible for the behavior of a system when it is first turned on. A typical example would be a monitor program in a system that loads the full operating system from drive or from a network and then passes control to the operating system.

Foreign configuration A RAID configuration that already exists on a replacement set of drives that you install in a computer system. MegaRAID Storage Manager software allows you to import the existing configuration to the RAID controller, or you can clear the configuration so you can create a new one.

Formatting The process of writing a specific value to all data fields on a drive, to map out unreadable or bad sectors. Because most drives are formatted when manufactured, formatting is usually done only if a drive generates many media errors.

Full charge capacity Amount of charge that can be placed in the battery. This value represents the last measured full discharge of the battery. This value is updated on each learn cycle when the battery undergoes a qualified discharge from nearly full to a low battery level.

G

Gas gauge status Hexadecimal value that represents the status flag bits in the gas gauge status register.

H

Hole In MegaRAID Storage Manager, a *hole* is a block of empty space in a drive group that can be used to define a virtual drive.

Host interface A controller property indicating the type of interface used by the computer host system: for example, *PCIX*.

Host port count A controller property indicating the number of host data ports currently in use.

Host system Any computer system on which the controller is installed. Mainframes, workstations, and standalone desktop systems can all be considered host systems.

Hot spare A standby drive that can automatically replace a failed drive in a virtual drive and prevent data from being lost. A hot spare can be dedicated to a single redundant drive group or it can be part of the global hot spare pool for all drive groups controlled by the controller.

When a drive fails, MegaRAID Storage Manager software automatically uses a hot spare to replace it and then rebuilds the data from the failed drive to the hot spare. Hot spares can be used in RAID 1, 5, 6, 10, 50, and 60 storage configurations.

I

Initialization The process of writing zeros to the data fields of a virtual drive and, in fault-tolerant RAID levels, generating the corresponding parity to put the virtual drive in a Ready state. Initialization erases all previous data on the drives. Drive groups will work without initializing, but they can fail a consistency check because the parity fields have not been generated.

IO policy A virtual drive property indicating whether Cached I/O or Direct I/O is being used. In Cached I/O mode, all reads are buffered in cache memory. In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to cache and the host concurrently. If the same data block is read again, it comes from cache memory. (The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.)

L

Learning cycle A battery calibration operation performed by a RAID controller periodically to determine the condition of the battery. You can start battery learn cycles manually or automatically

Learn delay interval	Length of time between automatic learn cycles. You can delay the start of the learn cycles for up to 168 hours (seven days).
Learn mode	Mode for the battery auto learn cycle. Possible values are Auto, Disabled, and Warning.
Learn state	Indicates that a learn cycle is in progress.
Load-balancing	A method of spreading work between two or more computers, network links, CPUs, drives, or other resources. Load balancing is used to maximize resource use, throughput, or response time.
Low-power storage mode	Storage mode that causes the battery pack to use less power, which save battery power consumption.
LKM	Local Key Management

M

Manufacturing date	Date on which the battery pack assembly was manufactured.
Manufacturing name	Device code that indicates the manufacturer of the components used to make the battery assembly.
Max error	Expected margin of error (percentage) in the state of charge calculation. For example, when Max Error returns 10 percent and Relative State of Charge returns 50 percent, the Relative State of charge is more likely between 50 percent and 60 percent. The gas gauge sets Max Error to 100 percent on a full reset. The gas gauge sets Max Error to 2 percent on completion of a learn cycle, unless the gas gauge limits the learn cycle to the +512/-256-mAh maximum adjustment values. If the learn cycle is limited, the gas gauge sets Max Error to 8 percent unless Max Error was already below 8 percent. In this case Max Error does not change. The gas gauge increments Max Error by 1 percent after four increments of Cycle Count without a learn cycle.
Maximum learn delay from current start time	Maximum length of time between automatic learn cycles. You can delay the start of a learn cycle for a maximum of 168 hours (7 days).
Media error count	A drive property indicating the number of errors that have been detected on the drive media.
Migration	The process of moving virtual drives and hot spare drives from one controller to another by disconnecting the drives from one controller and attaching them to another one. The firmware on the new controller will detect and retain the virtual drive information on the drives.

Mirroring The process of providing complete data redundancy with two drives by maintaining an exact copy of one drive's data on the second drive. If one drive fails, the contents of the other drive can be used to maintain the integrity of the system and to rebuild the failed drive.

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

N

Name A virtual drive property indicating the user-assigned name of the virtual drive.

Next learn time Time at which the next learn cycle starts.

Non-redundant configuration A RAID 0 virtual drive with data striped across two or more drives but without drive mirroring or parity. This provides for high data throughput but offers no protection in case of a drive failure.

NVRAM Acronym for nonvolatile random access memory. A storage system that does not lose the data stored on it when power is removed. NVRAM is used to store firmware and configuration data on the RAID controller.

NVRAM present A controller property indicating whether an NVRAM is present on the

controller. NVRAM size A controller property indicating the capacity of the controller's NVRAM.

O

Offline A drive is offline when it is part of a virtual drive but its data is not accessible to the virtual drive.

P

Patrol read A process that checks the drives in a storage configuration for drive errors that could lead to drive failure and lost data. The patrol read operation can find and sometimes fix any potential problem with drives before host access. This enhances overall system performance because error recovery during a normal I/O operation might not be necessary.

Patrol read rate The user-defined rate at which patrol read operations are run on a computer system.

Predicted battery capacity status (hold 24hr charge)	Indicates whether the battery capacity supports a 24-hour data retention time.
Product info	A drive property indicating the vendor-assigned model number of the drive.
Product name	A controller property indicating the manufacturing name of the controller.

R

RAID	<p>A group of multiple, independent drives that provide high performance by increasing the number of drives used for saving and accessing data.</p> <p>A RAID drive group improves input/output (I/O) performance and data availability. The group of drives appears to the host system as a single storage unit or as multiple virtual drives. Data throughput improves because several drives can be accessed simultaneously. RAID configurations also improve data storage availability and fault tolerance. Redundant RAID levels (RAID levels 1, 5, 6, 10, 50, and 60) provide data protection.</p>
RAID 0	Uses data striping on two or more drives to provide high data throughput, especially for large files in an environment that requires no data redundancy.
RAID 00	Uses data striping on two or more drives in a spanned drive group to provide high data throughput, especially for large files in an environment that requires no data redundancy.
RAID 1	Uses data mirroring on pairs of drives so that data written to one drive is simultaneously written to the other drive. RAID 1 works well for small databases or other small applications that require complete data redundancy.
RAID 5	Uses data striping and parity data across three or more drives (distributed parity) to provide high data throughput and data redundancy, especially for applications that require random access.
RAID 6	Uses data striping and parity data across three or more drives (distributed parity) to provide high data throughput and data redundancy, especially for applications that require random access. RAID 6 can survive the failure of two drives.
RAID 10	A combination of RAID 0 and RAID 1 that uses data striping across two mirrored drive groups. It provides high data throughput and complete data redundancy.
RAID 50	A combination of RAID 0 and RAID 5 that uses data striping across two drive groups with parity data. It provides high data throughput and complete data redundancy.

RAID 60	A combination of RAID 0 and RAID 6 that uses data striping across two drive groups with parity data. It provides high data throughput and complete data redundancy. RAID 60 can survive the failure of two drives in each RAID set in the spanned drive group.
RAID level	A virtual drive property indicating the RAID level of the virtual drive. LSI SAS RAID controllers support RAID levels 0, 1, 5, 6, 10, 50, and 60.
RAID Migration	A feature in RAID subsystems that allows changing a RAID level to another level without powering down the system.
Raw capacity	A drive property indicating the actual full capacity of the drive before any coercion mode is applied to reduce the capacity.
Read policy	A controller attribute indicating the current Read Policy mode. In Always Read Ahead mode, the controller reads sequentially ahead of requested data and stores the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data. In No Read Ahead mode (known as Normal mode in WebBIOS), read ahead capability is disabled.
Rebuild	The regeneration of all data to a replacement drive in a redundant virtual drive after a drive failure. A drive rebuild normally occurs without interrupting normal operations on the affected virtual drive, though some degradation of performance of the drive subsystem can occur.
Rebuild rate	The percentage of central processing unit (CPU) resources devoted to rebuilding data onto a new drive after a drive in a storage configuration has failed.
Reclaim virtual drive	A method of undoing the configuration of a new virtual drive. If you highlight the virtual drive in the Configuration Wizard and click Reclaim, the individual drives are removed from the virtual drive configuration.
Reconstruction rate	The user-defined rate at which a drive group modification operation is carried out.
Redundancy	A property of a storage configuration that prevents data from being lost when one drive fails in the configuration.
Redundant configuration	A virtual drive that has redundant data on drives in the drive group that can be used to rebuild a failed drive. The redundant data can be parity data striped across multiple drives in a drive group, or it can be a complete mirrored copy of the data stored on a second drive.
Relative state of charge	A redundant configuration protects the data in case a drive fails in the configuration. Predicted remaining battery capacity expressed as a percentage of Full Charge Capacity.

Remaining capacity	Amount of remaining charge capacity of the battery as stated in milliamp hours. This value represents the available capacity or energy in the battery at any given time. The gas gauge adjusts this value for charge, self-discharge, and leakage compensation factors.
Revertible hot spare	When you use the Replace Member procedure, after data is copied from a hot spare to a new drive, the hot spare reverts from a rebuild drive to its original hot spare status.
Revision level	A drive property that indicates the revision level of the drive's firmware.
Run time to empty	Predicted remaining battery life at the present rate of discharge in minutes.

S

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

Subvendor ID A controller property that lists additional vendor ID information about the controller.

T

Temperature Temperature of the battery pack, measured in Celsius.

Uncorrectable error count A controller property that lists the number of uncorrectable errors detected on drives connected to the controller. If the error count reaches a certain level, a drive will be marked as failed.

Vendor ID A controller property indicating the vendor-assigned ID number of the

controller. Vendor info A drive property listing the name of the vendor of the drive.

Virtual drive A storage unit created by a RAID controller from one or more drives. Although a virtual drive can be created from several drives, it is seen by the operating system as a single drive. Depending on the RAID level used, the virtual drive can retain redundant data in case of a drive failure.

Virtual drive state A virtual drive property indicating the condition of the virtual drive. Examples include Optimal and Degraded.

W

Write-back In Write-Back Caching mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a drive write transaction. Data is written to the drive subsystem in accordance with policies set up by the controller.

These policies include the amount of dirty/clean cache lines, the number of cache lines available, and elapsed time from the last cache flush.

Write policy See *Default Write Policy*.

Write-through In Write-Through Caching mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data and has completed the write transaction to the drive.