



Embedded MegaRAID[®] Software

User Guide

March 2012

48712-00, Rev. B



Revision History

Version and Date	Description of Changes
48712-00, Rev. B, March 2012	Revised the guide to document changes to the driver installation procedures, configuration utilities, and new event messages.
48712-00, Rev. A, June 2011	Initial release of the document.

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

LSI Corporation reserves the right to make changes to the product(s) or information disclosed herein at any time without notice. LSI Corporation does not assume any responsibility or liability arising out of the application or use of any product or service described herein, except as expressly agreed to in writing by LSI Corporation; nor does the purchase, lease, or use of a product or service from LSI Corporation convey a license under any patent rights, copyrights, trademark rights, or any other of the intellectual property rights of LSI Corporation or of third parties. LSI products are not intended for use in life-support appliances, devices, or systems. Use of any LSI product in such applications without written consent of the appropriate LSI officer is prohibited.

Corporate Headquarters
Milpitas, CA
800-372-2447

Website
www.lsi.com

Document Number: 48712-00, Rev. B
Copyright © 2012 LSI Corporation
All Rights Reserved

Table of Contents

Chapter 1: Overview	7
1.1 Embedded MegaRAID Software Features	7
1.1.1 Device Support	7
1.1.2 RAID Features	7
1.1.3 Error Handling	8
1.1.4 Driver Features	8
1.1.5 BIOS Features	8
1.1.6 RAID Management Utility Features	9
1.2 RAID Overview	9
1.2.1 RAID 0 Description	10
1.2.2 RAID 1 Description	10
1.2.3 RAID 5 Description	11
1.2.4 RAID 10 Description	12
Chapter 2: Driver Installation	13
2.1 Windows 2003 Driver Installation	13
2.1.1 Windows 7, Windows 2008, and Windows Vista Driver Installation	14
2.1.2 Updating the Windows Driver	14
2.1.3 Confirming the Windows Driver Installation	14
2.2 Linux Driver Installation	15
2.2.1 Obtaining the Driver Image File	15
2.2.2 Preparing the Installation Disk(s) for Linux	15
2.2.3 Installing the Red Hat Linux Driver on a New System	16
2.2.4 Updating the Red Hat Linux Driver (Generic)	17
2.2.5 Enabling RAID Mode during Red Hat Linux 5 Driver Installation	18
2.2.6 Known Restrictions for the Driver Installation Process	19
2.2.7 Installing the SuSE Linux Enterprise Server 9, 10, or 11 Driver	20
Chapter 3: LSI Software RAID Configuration Utility	21
3.1 Performing a Quick Configuration	21
3.2 Management Menu	22
3.3 Configuration Menu	23
3.3.1 Configuration Menu Options	23
3.4 Configuring Drive Groups and Virtual Drives	24
3.5 Creating a Storage Configuration	24
3.5.1 Selecting the Configuration Method	24
3.5.2 Using Easy Configuration	26
3.5.3 Using New Configuration	30
3.5.4 Using View/Add Configuration	34
3.6 Clearing a Storage Configuration	37
3.7 Configuring a Bootable Virtual Drive	38
3.8 Initializing Virtual Drives	38
3.8.1 First Initialization Method	39
3.8.2 Second Initialization Method	40
3.9 Rebuilding a Drive	41
3.10 Creating a Global Hotspare Drive	41
3.11 Checking Data Consistency	43

3.12 Displaying and Changing Controller Properties	44
3.12.1 Displaying and Changing Controller Properties	44
3.12.2 Displaying and Changing Drive Properties	47
3.12.3 Viewing or Changing Virtual Drive Properties	48
3.13 Forcing Drives Online or Offline	48
Chapter 4: MegaCLI Command Tool	49
4.1 MegaCLI CT Overview	49
4.2 MegaCLI Commands Not Supported by Embedded MegaRAID Software	50
4.3 Exception Handling	50
4.4 Command Line Abbreviations and Conventions	51
4.4.1 Abbreviations Used in the Command Line	51
4.4.2 Conventions	51
4.5 Adapter Commands	52
4.5.1 Display Adapter Information	52
4.5.2 Enable or Disable Automatic Rebuild	52
4.5.3 Set Adapter Properties	53
4.5.4 Display Specified Adapter Properties	53
4.5.5 Set Time and Date on Controller	54
4.5.6 Display Adapter Time	54
4.5.7 Set Factory Defaults	54
4.6 Event Log Commands	55
4.6.1 Manage the Event Log Entries	55
4.7 Configuration Commands	56
4.7.1 Add RAID 0, RAID 1, or RAID 5 Configuration	56
4.7.2 Configure Each Disk as RAID 0	57
4.7.3 Add RAID 10 Configuration	58
4.7.4 Clear Existing Configuration	58
4.7.5 Display Existing Configuration	58
4.7.6 Save Adapter Configuration	58
4.7.7 Restore Configuration Data from File	59
4.7.8 Delete Virtual Drive(s)	59
4.7.9 Display Free Space	59
4.8 Virtual Drive Commands	60
4.8.1 Display Virtual Drive Information	60
4.8.2 Display Virtual Drive Disk Cache Settings	60
4.8.3 Manage Virtual Drive Initialization	60
4.8.4 Manage Consistency Check	61
4.8.5 View Ongoing Background Initialization	61
4.8.6 Display Virtual Drive and Physical Drive Information	61
4.8.7 Display Number of Virtual Drives	62
4.9 Drive Commands	62
4.9.1 Display Drive Information	62
4.9.2 Set the Drive State to Online	62
4.9.3 Set the Drive State to Offline	62
4.9.4 Change the Drive State to Unconfigured-Good	63
4.9.5 Manage a Drive Initialization	63
4.9.6 Manage Global Hot Spares	63
4.9.7 Rebuild a Drive	64
4.9.8 Locate Physical Disk Drive(s) and Activate LED	64
4.9.9 Replace Configured Disk Drives and Start Automatic Rebuild	64
4.9.10 Prepare Unconfigured Physical Drives for Removal	64
4.9.11 Display Number of Physical Drives	65
4.9.12 Display List of Physical Drives	65
4.9.13 Download Firmware to the Physical Devices	65

4.10 Miscellaneous Commands	66
4.10.1 Display Version Information	66
4.10.2 Display MegaCLI Version	66
4.10.3 Display Help for the MegaCLI Utility	66
4.10.4 Display Summary Information	66
Chapter 5: MegaRAID Storage Manager Overview and Installation	67
5.1 Overview	67
5.1.1 Creating Storage Configurations	67
5.1.2 Monitoring Storage Devices	67
5.1.3 Maintaining Storage Configurations	68
5.2 Hardware and Software Requirements	68
5.3 Installation	69
5.3.1 Installing MegaRAID Storage Manager on Microsoft Windows	69
5.3.2 Installing MegaRAID Storage Manager for Linux	71
Chapter 6: MegaRAID Storage Manager Screen and Menus	72
6.1 Starting MegaRAID Storage Manager	72
6.2 MegaRAID Storage Manager Main Menu Screen	74
6.2.1 Dashboard/Physical View/Logical Views	75
6.2.2 Event Log Panel	78
6.2.3 Menu Bar	79
Chapter 7: Configuration	80
7.1 Creating a New Storage Configuration	80
7.1.1 Selecting Virtual Drive Settings	80
7.1.2 Creating a Virtual Drive Using Simple Configuration	82
7.1.3 Creating a Virtual Drive Using Advanced Configuration	85
7.2 Changing Adjustable Task Rates	89
7.3 Changing Virtual Drive Properties	90
7.4 Deleting a Virtual Drive	92
Chapter 8: Monitoring System Events and Storage Devices	93
8.1 Monitoring System Events	93
8.2 Configuring Alert Notifications	95
8.2.1 Setting Alert Delivery Methods	96
8.2.2 Changing Alert Delivery Methods for Individual Events	97
8.2.3 Changing the Severity Level for Individual Events	98
8.2.4 Multiple Events Displayed in a Single Pop-Up Window	99
8.2.5 Entering or Editing the Sender Email Address and SMTP Server	99
8.2.6 Authenticating a Server	100
8.2.7 Saving Backup Configurations	100
8.2.8 Loading Backup Configurations	100
8.2.9 Adding Email Addresses of Recipients of Alert Notifications	101
8.2.10 Testing Email Addresses of Recipients of Alert Notifications	101
8.2.11 Removing Email Addresses of Recipients of Alert Notifications	102
8.3 Monitoring Controllers	102
8.4 Monitoring Drives	103
8.5 Running a Patrol Read	104
8.5.1 Patrol Read Task Rates	106
8.6 Monitoring Virtual Drives	106
8.7 Monitoring Enclosures	108
8.8 Monitoring Rebuilds and Other Processes	108

Chapter 9: Maintaining and Managing Storage Configurations	110
9.1 Initializing a Virtual Drive	110
9.1.1 Running a Group Initialization	111
9.2 Running a Consistency Check	112
9.2.1 Running a Group Consistency Check	112
9.3 Scanning for New Drives	113
9.4 Rebuilding a Drive	113
9.5 Making a Drive Offline or Missing	114
Appendix A: Events and Messages	115

Chapter 1: Overview

This manual explains the features of the Embedded MegaRAID® Software. It includes instructions for using the LSI® Software RAID configuration utility, the MegaCLI command line utility, and the MegaRAID Storage Manager™ configuration utility.

You can use these three utilities to create RAID storage configurations on drives controlled by Embedded MegaRAID Software. The manual also includes instructions for installing the Embedded MegaRAID Software drivers in Microsoft® Windows® systems and Linux™ systems.

1.1 Embedded MegaRAID Software Features

The Embedded MegaRAID Software supports up to eight ports, depending on the hardware platform. This provides a cost-effective way to achieve higher transfer rates and reliability.

The following sections list the features available for devices, RAID, error handling, drivers, BIOS, Ctrl+M configuration utility, and RAID disk management.

1.1.1 Device Support

The Embedded MegaRAID Software offers the following device support:

- Support for up to 8 physical drives
- Support for SATA 6Gb/s drives
- Support for SAS 3Gb/s drives
- Support for Solid State Drives (SSDs)
- SATA CD/DVD-ROM support¹
- SATA tape device support¹
- Optical device (CD/DVD) hot plug feature used to connect optical devices while the operating system is running
- Hot plug support (online drive insertion and removal)
- Support for drive roaming
- Support for disk coercion (None, 128 MB, and 1 GB)

1.1.2 RAID Features

The Embedded MegaRAID Software supports the following RAID features:

- Support for RAID 0, RAID 1, RAID 5², and RAID 10
- Support for up to 8 virtual drives
- Support for virtual drives larger than 2 TB
- Stripe size of 64 KB only
- Virtual drive availability immediately after creation
- Support for the random deletion of virtual drives
- Support for array cache setting (a RAID 10 volume is considered as a single array, though it might have 2, 3, or 4 spans)

1. AHCI-based chipsets only.
2. RAID 5 is a premium feature.

- Support for migration path from Embedded MegaRAID Software to MegaRAID SAS RAID controllers
- Check consistency for RAID 1, RAID 5, and RAID 10
- Drive group initialization support (fast and full)
- Support for auto or manual rebuild
- Ability to set the rates for the BGI, consistency check, and patrol read
- Automatic resumption of rebuilding, check consistency, full initialization, and BGI processes
- Global hot spare support

1.1.3 Error Handling

The Embedded MegaRAID Software supports the following error handling features:

- Soft Bad Block Management (SBBM) support
- Error/Event logging and notification

1.1.4 Driver Features

The Embedded MegaRAID Software driver supports the following features:

- Error logging and notification
- Support for Microsoft Windows Server® 2003, Microsoft Windows Server 2008, Microsoft Windows Server® 2008R2, Microsoft Windows Vista, and Microsoft Windows 7
- Support for Red Hat® Linux
- Support for SuSE® Linux for 2.4, 2.6, and 3.0 kernels

1.1.5 BIOS Features

The Embedded MegaRAID Software BIOS has the following features:

- Support for Interrupt 13 and Enhanced Disk Drive Specification
- Support for Int19h
- Option ROM size of 64 KB
- Support for BIOS Boot Specification (BBS) (If available in system BIOS, this feature lets you select the controller from which to boot.)
- Support for power-on self test (POST)
- Support for Post Memory Management (PMM): Specification v7, July 2010
- Industry-standard EBDA
- POST and run-time BIOS support for device insertion and removal
- Support for Stop On Error during boot-up

The following feature is supported by the BIOS and the Ctrl+M Configuration Utility:

- Automatic resumption of rebuilding, Check Consistency, and full initialization, and background initialization (BGI; BGI is for RAID 5 configurations only)



NOTE The BIOS program and the BIOS Configuration Utility (Ctrl+M) do not start or resume background initialization (BGI). If BGI is already in progress, you cannot start the Check Consistency operation.

1.1.6 RAID Management Utility Features

The following features are available to manage the virtual drives and the physical drives in the system:

- Configuration information display (in MegaRAID Storage Manager)
- Physical drive properties and virtual drive properties
- Drive group (array) management
- Error logging and notification
- Auto-configuration support of newly added drive
- Ability to save and restore a configuration

1.2 RAID Overview

This section provides a brief overview of the types of RAID configurations that Embedded MegaRAID Software supports.

The first step in creating a RAID storage configuration is to configure drives in drive groups (also known as arrays). As defined for Embedded MegaRAID Software, a drive group is a group of one to eight physical disks that is seen by the host computer system as one large disk drive, or virtual drive. Only one RAID level can be assigned to a drive group.

- A RAID 0 drive group consists of one to eight drives.
- A RAID 1 drive group consists of two drives.
- A RAID 5 drive group consists of three to eight drives.
- A RAID 10 drive group consists of four, six, or eight drives.



NOTE Some hardware configurations do not support eight drives. Depending on the hardware, the actual maximum number of drives for RAID 0, RAID 5 and RAID 10 drive groups can be fewer than eight.

You can use any of these three strategies when creating RAID drive groups and virtual drives:

- **Maximize Fault Tolerance** – You can maximize fault tolerance to protect against loss of data by creating a RAID 1 drive group with mirroring. All data is written to the primary drive in the drive group and is also written (mirrored) to a second drive.
- **Maximize Virtual Drive Performance** – You can maximize virtual drive performance by creating a RAID 0 drive group with striping. Data is broken into segments and can be simultaneously written to or read from several different stripes on several different drives in the drive group.
RAID 10 drive groups combine both striping and mirroring to provide high data transfer rates and data redundancy.
- **Maximize Storage Capacity** – You can maximize storage capacity when selecting a RAID level. Striping alone (RAID 0) requires less storage space than mirrored data (RAID 1) or distributed parity (RAID 5). RAID 5, which provides redundancy for one drive failure without duplicating the contents of entire drives, requires less space than RAID 1.

1.2.1 RAID 0 Description

RAID 0 provides disk striping across all drives in the drive group. RAID 0 does not provide any data redundancy, but does offer the best performance of any RAID level. RAID 0 breaks up data into smaller segments called strips, and then stripes the data segments across each drive in the drive group. The size of each data segment is determined by the strip size, which is 64 KB.



NOTE It is possible to create each disk as a single-drive RAID 0 drive group. However, spanning across single drive RAID 0 drive groups is not supported.

By breaking up a large file into smaller segments, and writing or reading from several drives at once, the Embedded MegaRAID Software utility can read or write the file faster. This feature makes RAID 0 ideal for applications that require high bandwidth but do not require fault tolerance.

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; all data lost if any drive fails

Drives One to eight

The following figure shows a RAID 0 drive group with two drives.

Figure 1 RAID 0 Drive Group Example with Two Drives



1.2.2 RAID 1 Description

RAID 1 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.

Uses Databases or any other mission critical environment that requires fault tolerance


Strong Points Provides complete data redundancy; RAID 1 is ideal for any application that requires fault tolerance

Weak Points Requires twice as many drives; performance is impaired during drive rebuilds

Drives Two

The following figure shows a RAID 1 drive group.

Figure 2 RAID 1 Drive Group



1.2.3 RAID 5 Description

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 distributed to all drives. RAID 5 is best suited for networks that perform a lot of small input/output (I/O) transactions simultaneously.




NOTE RAID 5 is a premium feature. You might need to install a software key to enable RAID 5. The key you need depends on your supplier. Contact your supplier for more information.

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.

Uses	Provides high data throughput. Use RAID 5 for transaction processing applications because each drive can read and write independently. If a drive fails, the RAID controller uses the parity drive to recreate all missing information. Use also for office automation and online customer service that requires fault tolerance. Use for any application that has high read request rates but low write request rates.
Strong Points	Provides data redundancy, high read rates, and good performance in most environments. Provides redundancy with lowest loss of capacity.
Weak Points	Not well suited to tasks requiring lot of small writes. Drive performance will be reduced if a drive is being rebuilt or a background initialization is in progress. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.
Drives	Three to eight.

The following figure shows a RAID 5 drive group with six drives.

Figure 3 RAID 5 Drive Group



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

1.2.4 RAID 10 Description

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




NOTE On a RAID 10 drive group, you can create only one virtual drive, and that virtual drive must occupy the entire space of the RAID 10 drive group.

- Uses** 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); 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
- Drives** Four, six, or eight

The following figure shows a RAID 10 drive group with four drives.

Figure 4 RAID 10 Drive Group



Chapter 2: Driver Installation

This chapter explains how to install the Embedded MegaRAID Software drivers for the following operating systems:

- Microsoft Windows Server 2003, 2008, and 2008R2
- Microsoft Windows 7 Workstation
- Microsoft Windows Vista Workstation
- Red Hat Enterprise Linux™ (RHEL) 5 and 6
- SuSE Linux Enterprise (SLES) 10 and 11

2.1 Windows 2003 Driver Installation

Perform the following steps to install the MegaRAID device driver in a new Windows 2003 operating system. The Windows 2003 operating system automatically adds the driver to the registry and copies the driver to the appropriate directory.

1. Start the Windows 2003 installation by booting from the appropriate Windows CD-ROM.
The system BIOS must support booting from a CD-ROM. BIOS settings might require changes to allow CD-ROM booting. See your system documentation.
2. Press **F6** when the following displays at the bottom of the screen, unless you are installing Windows Vista:
"Press F6 if you need..."



NOTE For the system to recognize the new driver for Windows Server 2003, you must press **F6**. If you are installing Windows Vista, you do not need to press F6. For Windows Vista, after the first installation screen, an option displays to allow you to load the driver from either a floppy diskette or a USB key. Otherwise, it will load the default driver from the DVD.

3. Select **S** to specify an additional device when the screen displays:
Setup could not determine the type of one or more mass storage devices...
The system prompts for the manufacturer-supplied hardware support disk.



NOTE If the screen does not display this message after you press **F6**, then the setup program did not recognize the F6 command. Reboot the system, and return to step 2.

4. Insert the driver diskette containing the Windows device driver and press **Enter**.
5. Select the appropriate MegaRAID adapter from the menu by using the arrow key to highlight it, and then press **Enter** to proceed.
6. Press **Enter** again to proceed.
7. Return to the Windows Setup screen.
Windows displays a **Welcome to Setup** window.
8. Press **Enter** to continue.
9. Press **C** to continue the Microsoft Windows installation procedure.
10. Follow the Windows installation procedure.
11. Repeat this process for all the adapters on your system.

2.1.1 Windows 7, Windows 2008, and Windows Vista Driver Installation

Perform the following steps to install the MegaRAID device driver in a new Windows 7, 2008, or Vista operating system.

The Windows 2003 operating system automatically adds the driver to the registry and copies the driver to the appropriate directory.

1. Start the Windows 7, 2008, or Vista installation by booting from the appropriate Windows DVD.
The system BIOS must support booting from a DVD. BIOS settings might require changes to allow DVD booting. See your system documentation.
Windows loads the file and the first installation screen appears.
2. Select your language and other settings based upon your location and preference.
3. Press **Next**.
The Windows Install screen appears.
4. Press **Install Now** to start the installation wizard.
The software license screen appears.
5. Click the checkbox to accept the software license and click **Next**.
The next installation screen appears.
6. Select the type of installation you want.
7. Follow the prompts to select the location where you want to install Windows and click **Next**.
The program begins installing the files. Your system will restart several time during the installation process.

2.1.2 Updating the Windows Driver

Perform the following steps to update the Embedded MegaRAID Software driver for Windows or to install this driver on an existing system booted from a standard IDE drive.

1. Click **Start**, point to **Settings**, and then click **Control Panel**.
2. Double-click **System**, click the **Hardware** tab, and then click **Device Manager**.
Device Manager starts.
3. In Device Manager, double-click **SCSI and RAID Controllers**, right-click the device for which you are installing the driver, and then click **Properties**.
4. On the Driver tab, click **Update Driver** to open the Update Device Driver wizard, and then follow the wizard instructions to update the driver.

2.1.3 Confirming the Windows Driver Installation

Perform the following steps to confirm that the Embedded MegaRAID Software driver for Windows is installed correctly.

1. Click **Start**, point to **Settings**, and then click **Control Panel**.
2. Double-click **System**, click the **Hardware** tab, and then click **Device Manager**.
Device Manager starts.
3. In Device Manager, double-click **SCSI and RAID Controllers**, right-click the device for which you are installing the driver, and then click **Properties**.
4. On the Driver tab, click **Driver Details** and verify that the driver information is correct.

2.2 Linux Driver Installation

This section explains the steps to install the Embedded MegaRAID device driver in a Red Hat Enterprise Linux installation or a SuSE Linux Enterprise Server installation.

2.2.1 Obtaining the Driver Image File

The Linux driver is offered in the form of a driver update disk. The required file is `dud-[driver version].img`, which is the driver update disk for the Embedded MegaRAID Software stack.

You can obtain the latest driver files from the Download Center on the LSI website.

2.2.2 Preparing the Installation Disk(s) for Linux

This section describes how to prepare the Linux installation disk(s) from the driver image files, using either the Windows operating system or the Linux operating system.

2.2.2.1 Preparing Installation Disks with the Windows Operating System

Under Windows, you can use the RaWrite floppy image writer utility to create disk images from image files. The image writer can be downloaded from the Internet. Perform the following steps to build installation diskettes.

1. Copy the driver update disk image `dud-[driver version].img` and the file `raw write.exe` to a directory.
2. Confirm that the files are in the selected directory.
3. If necessary, use this command to change the filename of the driver update disk to a name with fewer than eight characters:

```
copy dud-[driver version].img dud.img
```
4. Open the DOS Command Prompt and navigate to the directory where `raw write.exe` is located.
5. Type the following command to create the installation diskette:

```
raw write
```
6. Press **Enter**.
You are prompted to enter the name of the boot image file.
7. Type the following:

```
dud.img
```
8. Press **Enter**.
You are prompted for the target drive diskette.
9. Insert a floppy diskette into the floppy drive and type:

```
A:
```
10. Press **Enter**.
11. Press **Enter** again to start copying the file to the diskette.
12. After the command prompt returns and the floppy disk drive LED goes out, remove the diskette.
13. Label the diskette with the image name.

2.2.2.2 Preparing Installation Disks with the Linux Operating System

Under Red Hat Linux and SuSE Linux, you can use a driver diskette utility to create disk images from image files. Perform the following steps to create the driver update disk:

1. Copy the driver update disk image `dud-[driver version].img` to a Linux system.
2. Insert a blank floppy diskette into the floppy drive.
3. Confirm that the files are in the selected directory.
4. Create the driver update diskette using the following command:

```
dd if=dud-[driver version].img of=/dev/fd0
```
5. After the command prompt returns and the floppy disk drive LED goes out, remove the diskette.
6. Label the diskette with the image name.

2.2.3 Installing the Red Hat Linux Driver on a New System

This section describes the fresh installation of the Red Hat Enterprise Linux 4, 5, or 6 device driver on systems with the Embedded MegaRAID Software stack. After you prepare the installation disks with the driver image, perform the following steps to install the driver on a new system.

1. Boot to CD-ROM (Disk 1).
The Red Hat introductory screen appears.
2. Type the following at the boot prompt:

```
linux dd noprobe
```

 (it depends on the number of drives)
For example, to install Red Hat Linux on a RAID 5 configuration with three drives, enter:

```
linux dd noprobe=ata1 noprobe=ata2 noprobe=ata3
```
3. Press **Enter**.
The prompt asks whether you have a driver disk.
4. Use the arrow key to select **Yes**, and then press **Enter**.
5. Select **fd0** to indicate you have a floppy diskette with the driver on it.
6. Insert the floppy diskette in the A:/ drive and press **Enter**.
The installer locates and loads the driver for your device. The following message appears:

```
Loading megasr driver...
```


The prompt at the next screen asks whether you have another driver.
7. Follow the Red Hat Linux installation procedure to complete the installation.
8. Before you reboot, go to text console and follow these steps:
 - a. Press **Ctrl+Alt+F2** goes to the text console.
 - b. Enter the command `cat /proc/partitions` to get the major and minor number of floppy drive.
 - c. Execute `mknod /dev/sd(x) b major minor`.
 - d. Create a directory, such as `mkdir swr`.
 - e. Mount the floppy drive to that directory with the `mount /dev/sd(x) swr` command.
 - f. Run the script `./replaceachi.sh`.
9. Reboot the system.

2.2.4 Updating the Red Hat Linux Driver (Generic)

Perform the following steps to update the Red Hat Linux driver or to install the Red Hat Linux driver in an existing system booted from a standard SATA drive or systems with the Embedded Software RAID stack.

1. Boot the system with the Red Hat Linux Installation CD from the primary controller or disk.
The Red Hat introductory screen appears.
2. Mount the driver update diskette (DUD) using the following command:

```
#mount /dev/fd0 /mnt/floppy
```
3. Unzip the `modules.cgz` file that is on the DUD to get driver images for different Red Hat operating systems:

```
#mkdir -p /home/megasr  
#cd /home/megasr  
#cp /mnt/floppy/modules.cgz  
#gunzip -S .cgz modules.cgz
```

This action generates a new file named `modules`:

```
#cpio -ivd < modules
```

This action provides the following driver images:

```
{<kernel version>,<kernel version>smp,  
<kernel version>BOOT }/megasr.o
```
4. Update the Megasr driver module for the required kernels using the following commands:

```
#cd /home/megasr
```

If the `/lib/modules/<kernel version>/update/` directory is present, use the following command:

```
# cp <kernel version>/megasr.[o/ko]  
/lib/modules/<kernel version>/update/megasr.[o/ko]
```

If the `/lib/modules/<kernel version>/update/` directory is not present, use the following command:

```
# cp <kernel version>/megasr.[o/ko]  
/lib/modules/<kernel version>  
/kernel/drivers/scsi/megasr.[o/ko]
```
5. Create a Megasr driver entry in the configuration file. The Red Hat configuration file is `/etc/modules.conf`.
If the Megasr entry is not present in `/etc/modules.conf`, add the following line:

```
alias scsi_hostadapter megasr
```

If the `ahci` SCSI driver entry (located on the following paragraph) is present in `/etc/modules.conf`, remove it. It must be removed; otherwise, the `ahci` driver would take control of the RAID controller without checking the subsystem device or Vendor ID. The `ahci` SCSI driver entry is the following:

```
alias scsi_hostadapter ahci
```
6. Create a new `initrd` image for the required kernel.
Red Hat installation uses the `mk_initrd` command to create an `initrd` image. The following command creates an `initrd` image for the `<kernel version>smp` kernel in the boot directory. Refer to the `mk_initrd` man page for more information. The command is:

```
#mkinitrd /boot/initrd<kernel version>smp.img.new  
<kernel version>smp
```
7. Modify the `lilo.conf/grub.conf` file by adding newly created `initrds` as new entries in the `/etc/lilo.conf` file.
The suggested method is to copy an existing `lilo` entry in the file and paste it as a new one. Then modify its kernel image name, `initrd` image name, and label name.

Sample Lilo Entry

```
image=/boot/vmlinuz-<kernel version>smp label=linuxnew  
initrd=/boot/initrd-<kernel version>smp.img.new  
read-only appended=root=LABEL=/ "
```

Sample Grub Entry

```
title Red Hat Linux (<kernel version> with Megasr driver)  
root (hd0,0)  
kernel /vmlinuz-<kernel version> ro root=LABEL=/  
initrd /initrd-<kernel version>.img.new
```

8. Update the boot loader. If the boot loader is Lilo, run the lilo command to update the boot loader:
#lilo
9. Reboot the system to the new boot loader entry.

2.2.5 Enabling RAID Mode during Red Hat Linux 5 Driver Installation

This section documents how to load the RHEL5 operating system drivers so the drivers recognize RAID mode and support RAID 0, 1, 5, and 10 functionality.

On servers with on-board, AHCI-based SATA controllers, the controllers can be set to HBA mode or RAID mode using the system BIOS. The system BIOS changes the PCI sub-system IDs based on the selected mode. The MegaRAID software RAID (megasr) driver can be set to RAID mode for the AHCI-based SATA controllers and support RAID levels 0, 1, 5, and 10. The megasr driver considers all four PCI IDs to decide whether the driver supports the controller.

In Red Hat Enterprise Linux 5 (RHEL 5), however, the native AHCI driver Linux kernel looks only at the vendor ID and the device ID, and ignores the subsystem IDs. As a result, RHEL 5 ignores the RAID mode, and does not support RAID mode. The RHEL 5 installation process loads the native AHCI driver even when the megasr DUD is present. Also, the process includes AHCI in the initrd (during which a temporary file system is loaded into memory in the Linux kernel boot process).

Perform the following steps to load the RHEL5 operating system drivers so they recognize the RAID mode.

1. Use the `noprobe` option at the installation kernel boot prompt, in addition to the `dd` option.
This step prevents AHCI from loading ahead of megasr, allowing you to install RHEL 5 on the megasr virtual drive. The following string is an example of the `noprobe` option:

```
boot: linux dd noprobe=ata1 noprobe=ata2 noprobe=ata3 noprobe=ata4
```


At the last installation step, after all of the packages are installed, RHEL5 prompts you to reboot. At this point, the `initrd` is built with the AHCI driver. You must add `megasr` and delete `ahci` from `initrd` before you reboot.
2. Press **CTRL+ALT+F2**.
This step takes you to a text console prompt.
3. Before you reboot, go to text console and follow these steps:
 - a. Press **Ctrl+Alt+F2** goes to the text console.
 - b. Enter the command `cat /proc/partitions` to get the major and minor number of floppy drive.
 - c. Execute `mknod /dev/sd(x) b major minor`.
 - d. Create a directory, such as `mkdir swr`.
 - e. Mount the floppy drive to that directory with the `mount /dev/sd(x) swr` command.
 - f. Run the script `./replaceachi.sh`.

If the script is in the floppy drive, you have to mount the device first. In the normal system, you use `/dev/sd<x>` device names to mount a device. In the pre-installation environment, these names might not exist. You can find the major and minor numbers of your device by reading the `/proc/partitions` file.

4. Use the major and minor numbers of your device in the `mknod` command to create your own device name.

2.2.6 Known Restrictions for the Driver Installation Process

This section documents known restrictions that you must follow when you install the operating system drivers for various operating systems.

2.2.6.1 Operating System Using DUD Images

Follow these steps when you install the operating system using DUD images:

- For SLES (32-bit and 64-bit) platforms:
 - Enter `brokenmodules=ahci` while installing the driver.
- For SLES 11 SP2 (32-bit and 64-bit) platforms:
 - Enter `brokenmodules=ahci brokenmodules=isci` while installing the driver.
- For RHEL 6 GA (32-bit and 64-bit) platform operating system installation using the SWR DUD images:
 - While booting from DVD, press **ESC** so you can install the third party driver.
 - Enter the following command to install the driver: `Linux dd blacklist=ahci`.
- For RHEL 6.1 GA and RHEL 6.2 GA (32-bit and 64-bit) platform operating system installation using the SWR DUD images, follow these steps:
 - While booting from DVD, press **ESC** so you can install the third party driver.
 - Enter the following command to install the driver: `Linux dd blacklist=isci blacklist=ahci nodmraid`.
- For RHEL 5.7 GA (32-bit and 64-bit) platform operating system installation using the SWR DUD images, follow these steps:
 - While booting from DVD, press **ESC** so you can install the third party driver.
 - Enter the following command to install the driver: `Linux dd blacklist=isci blacklist=ahci`.
- For all SLES platforms (32-bit and 64-bit), platform operating system installation using the SWR DUD images, follow the steps in this example for SLES 11 SP1 (32-bit and 64-bit) platforms, in which the dud image size exceeds the floppy disk size:
 - Copy the image file to any linux system. The image files are:
 - `megasr-14.00.0722.2010-1-sles11-sp1-x86.img` for 32-bit SLES11 SP1
 - `megasr-14.00.0722.2010-1-sles11-sp1-x86_64.img` for 64-bit SLES11 SP1
 - Create a directory, for example, `mkdir image`.
 - Mount the dud image on the image directory using the following command:
 - `mount -o loop megasr-14.00.0722.2010-1-sles11-sp1-<arch>.img image`.
 - Change the directory to image (for example, `cd image`).
 - Copy the contents of the image directory to a USB drive. (Contents of the image start with the 01 directory, so copy the 01 directory to USB drive).
 - Use the USB drive to provide the third-party driver during the operating system installation.
 - Enter the following command while installing the driver: `brokenmodules=ahci`.
 - Make sure the USB drive is formatted before using it for operating system installation using DUD images.
- This package has limitations on RHEL4.x (64-bit) platforms. For RHEL4.x (64-bit) platform operating system installation using the SWR DUD images, there are two images.
 - Use the default DUD image for 2.6.x-xx.EL and 2.6.x-xx.ELsmp kernels, and use ext DUD image for 2.6.x-xx.EL and 2.6.x-xx.ELlargesmp kernels.

2.2.7 Installing the SuSE Linux Enterprise Server 9, 10, or 11 Driver

This section describes the fresh installation of the SuSE Linux Enterprise Server 9, 10, or 11 driver on a system with the Embedded MegaRAID Software stack. Prepare installation disks with the driver image, and then perform the following steps to install the driver.

1. Create a RAID drive group using one of the configuration utilities.
2. Boot the system using the SLES Disk.
3. When the first screen appears, select **Installation** on the menu.
4. Type the following in the Boot Options field:
`brokenmodules=achi`
5. Press **F6** for the driver and select **Yes**.
6. Insert the driver update diskette in the **A :** / drive and press **Enter**.
"Yes" appears under the F6 Driver heading.
7. Press **OK**.
The following message appears:
`LSI Soft RAID Driver Updates added.`
8. At the menu, select the driver update medium and press the **Back** button.
9. Continue and complete the installation process.

Chapter 3: LSI Software RAID Configuration Utility

Use the LSI Software RAID Configuration Utility (CU) to configure disk drive groups and virtual drives, and to perform other configuration tasks in a pre-boot environment.

You can perform the following functions using the LSI Software RAID CU:

- Select a configuration method for drive groups and virtuals drives
- Create drive groups
- Define virtual drives
- Initialize virtual drives
- Access controllers, virtual drives, and drive groups to view their properties
- Create hot spare drives
- Verify that the redundancy data in virtual drives using RAID level 1, 5, and 10 is correct
- Rebuild failed drives
- Reconstruct virtual drives after changing RAID levels or adding a drive to a drive group
- Select a MegaRAID host adapter



NOTE If the configuration utility does not display, go into BIOS setup and disable Quick boot, Fast boot, Silent boot, Intel® Rapid boot, and Quick POST, then reboot. If still unable to access the configuration utility, check for a system BIOS upgrade.

3.1 Performing a Quick Configuration

This section provides high-level instructions for quickly configuring drive groups and virtual drives with the LSI Software RAID CU. These instructions are intended for users that are familiar with configuration utilities and tools. See Section 3.4, [Configuring Drive Groups and Virtual Drives](#), on page 24, for detailed configuration instructions. To ensure the best performance, select the optimal RAID level for the virtual drive you create. For an explanation of RAID levels, see Section 1.2, [RAID Overview](#), on page 9.


Perform the following steps to configure drive groups and virtual drives using the LSI Software RAID CU:

1. Boot the system.
2. Press **Ctrl+M** to start the LSI Software RAID CU.
3. Select **Configure** on the Management Menu screen.
4. Select a configuration method from the Configuration menu (**Easy Configuration**, **New Configuration**, or **View/Add Configuration**).
5. Create drive groups using the available drives.
6. Designate hot spare disks (optional).
7. Define the virtual drive(s) using the space in the drive groups.
8. Initialize the new virtual drive(s).

3.2 Management Menu

The Management Menu screen appears when you start the LSI Software RAID CU:

Figure 5 LSI Software RAID Configuration Utility Management Menu Screen



NOTE The minimum screen resolution for the LSI Software RAID CU is 640x480.

3.3 Configuration Menu

Use the Configuration Menu screen to configure drive groups and virtual drives. This section describes the configuration options.

3.3.1 Configuration Menu Options

The Configuration Menu screen provides four methods to modify and/or create a virtual drive configuration: Easy Configuration, New Configuration, View/Add Configuration, and Clear Configuration, as shown in the following figure:

Figure 6 LSI Software RAID Configuration Menu

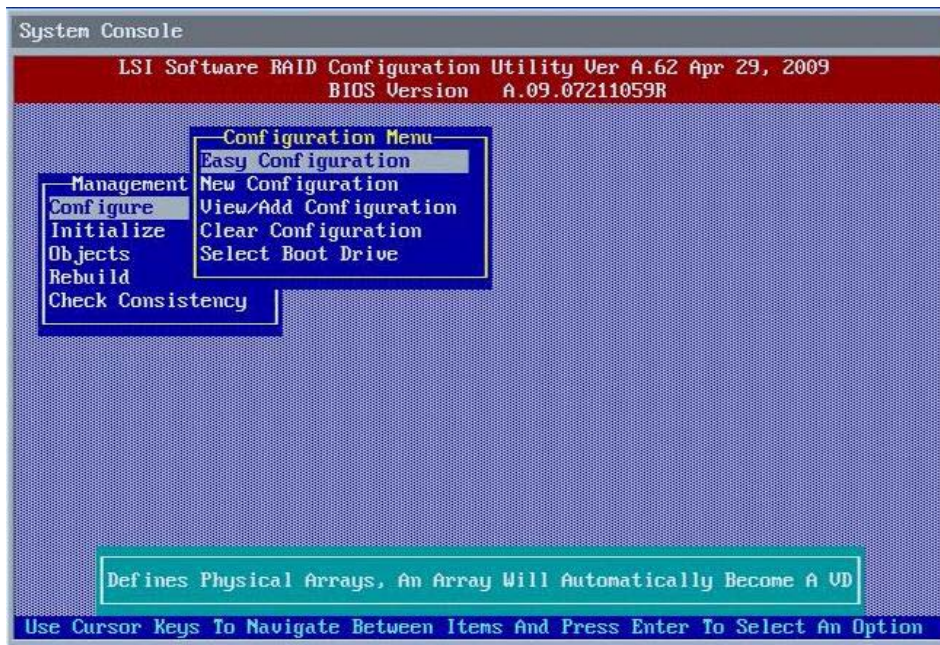


Table 1 describes these methods. The Configuration Menu has an Advance submenu that enables you to set specific options. The available options depend upon the configuration method you use.

Table 1 LSI Software RAID Configuration Utility – Configuration Menu

Option	Description
Easy Configuration	Easy Configuration automatically associates every drive group with one virtual drive. Through the Advance Menu, Easy Configuration allows you to modify the RAID level and stripe size. Section 3.5.2, Using Easy Configuration , on page 26, provides detailed instructions.
New Configuration	New Configuration allows you to modify the RAID level, stripe size, virtual drive size, and disk spanning (associating virtual drives with multiple drive groups). If you select New Configuration, the CU deletes the existing configuration information on the selected controller when saving the new configuration. Section 3.5.3, Using New Configuration , on page 30, provides detailed instructions.
View/Add Configuration	View/Add Configuration allows you to control the same virtual drive parameters as New Configuration without disturbing the existing configuration information. The View/Add configuration also allows you to enable the Configuration on Disk feature. Section 3.5.4, Using View/Add Configuration , on page 34, provides detailed instructions.
Clear Configuration	Clear Configuration erases the current configuration information. Section 3.6, Clearing a Storage Configuration , on page 37, provides detailed instructions.
Specify Boot Drive	Specify Boot Drive enables you to specify a virtual drive as the boot drive on the controller, if you have created virtual drives. Section 3.7, Configuring a Bootable Virtual Drive , on page 38, provides detailed instructions.

3.4 Configuring Drive Groups and Virtual Drives

The following sections provide detailed instructions for configuring drive groups and virtual drives with the LSI Software RAID CU.



NOTE Use drives with the same capacity when you create a storage configuration. If you use drives with different capacities in one drive group, the CU limits each drive to the capacity of the smallest drive.

The number of physical drives in a specific drive group determines the possible RAID levels that you can implement with the drive group.

- RAID 0 requires from one to eight physical drives.
- RAID 1 requires two physical drives.
- RAID 5 required three to eight physical drives.
- RAID 10 requires four, six, or eight physical drives.

3.5 Creating a Storage Configuration

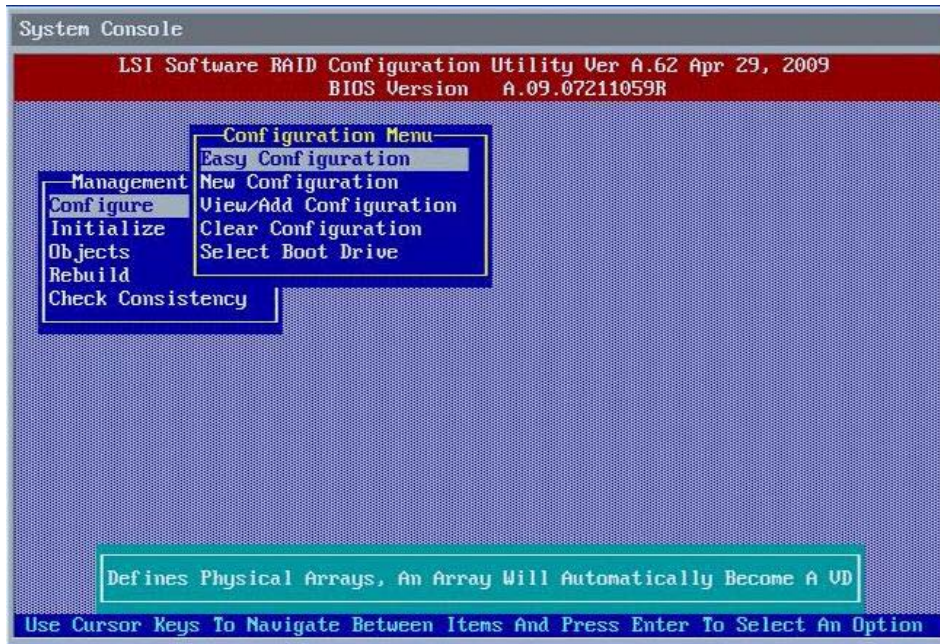
This section explains how to use the LSI Software RAID Configuration Utility to configure RAID drive groups and virtual drives to create storage configurations:

3.5.1 Selecting the Configuration Method

Easy Configuration automatically associates each drive group with one virtual drive. Follow these steps to open the Configuration Menu screen, and select a configuration method:

1. On the Management Menu screen, highlight **Configure** with your cursor, and press **Enter**.
The Configuration Menu screen appears, as shown in the following figure.

Figure 7 Configuration Menu Screen



2. Select a configuration option.
 - **Easy Configuration** – Automatically associates every drive group with one virtual drive.
 - **New Configuration** – Clears the existing configuration and lets you create a new configuration.
 - **View/Add Configuration** – Retains the existing storage configuration and adds new drives to it (this method does not cause any data loss).
 - **Clear Configuration** – Clears the existing configuration.



CAUTION If you choose to clear the existing configuration or to create a new configuration, the system deletes all of the existing data in the configuration. Make a backup of any data that you want to keep before you choose an option.

3. Press **Enter**.

The configuration screen appear for the configuration option that you selected.

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

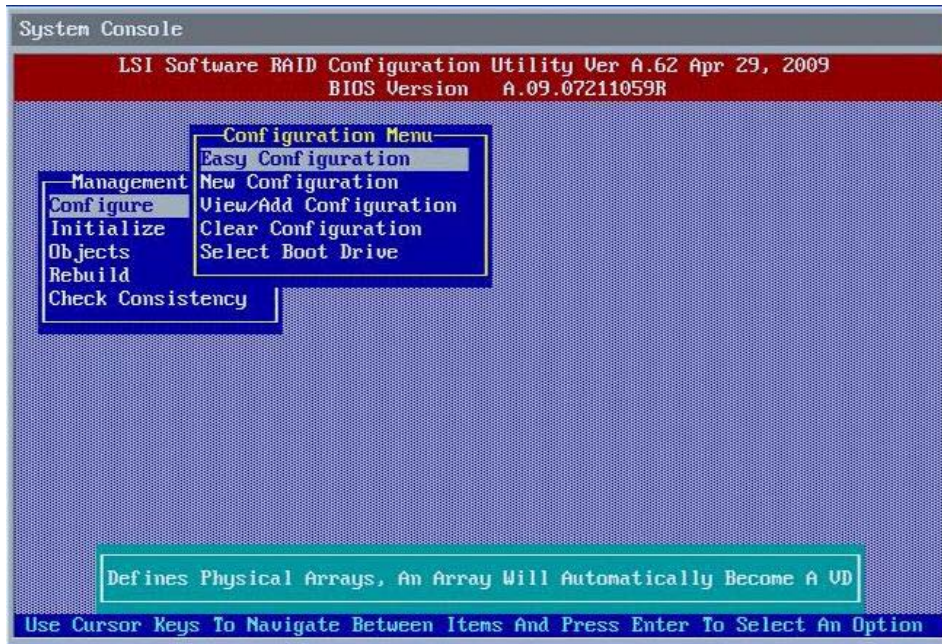
The following sections describe the configuration steps for each configuration method.

3.5.2 Using Easy Configuration

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

1. On the Management Menu screen, highlight **Configure** with your cursor, and press **Enter**.
The Configuration Menu screen appears, as shown in the following figure.

Figure 8 Configuration Menu Screen



2. Use your cursor to highlight **Easy Configuration** and press **Enter**.
The following screen appears.

Figure 9 Easy Configuration – Array Selection Menu



3. Press the space bar to select a drive and add it to the drive group.
4. Repeat step 3 to add additional drives to the drive group.
5. Press **F10** to continue configuration.

The Select Configurable Array(s) dialog box appears, as shown in the following figure.

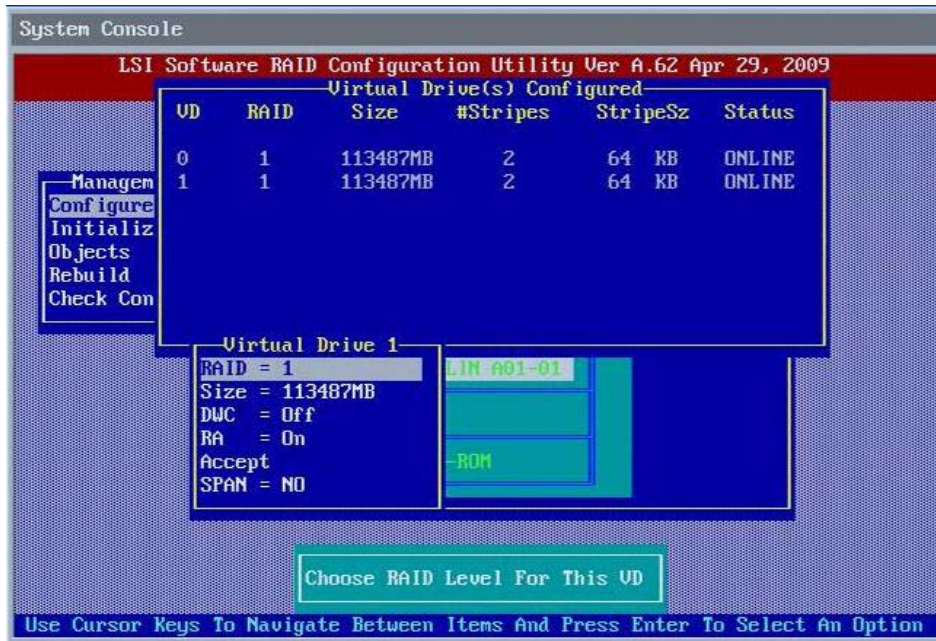
Figure 10 Drive Group Selection



6. Press the space bar to select the drive group.
7. Press **F10** to continue configuration.

The Virtual Drive(s) Configured dialog box appears, as shown in the following screen. This screen shows the virtual drive number, RAID level, drive group size, number of stripes, stripe size and drive status.

Figure 11 Virtual Drive Parameters Dialog Box



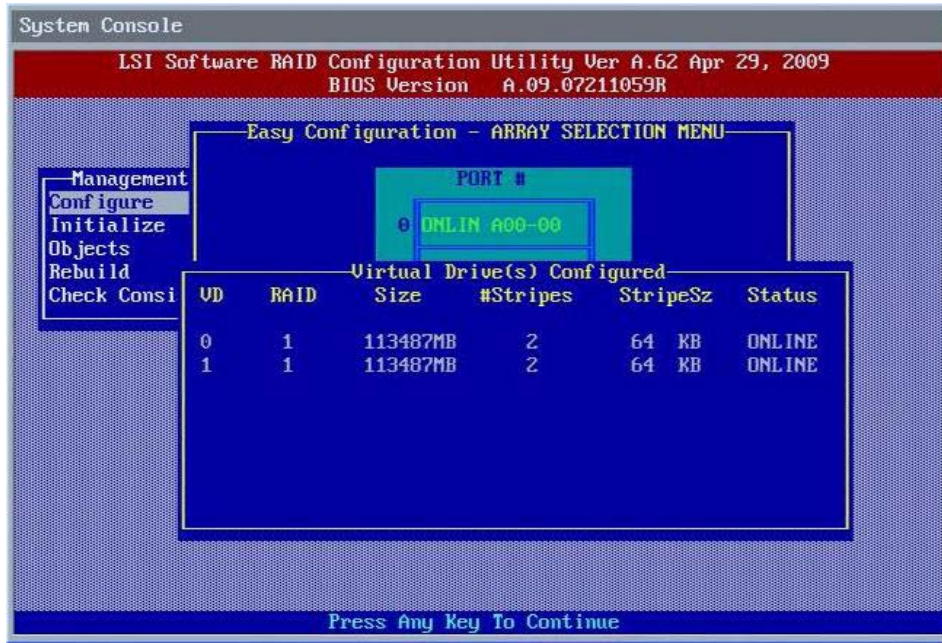
8. Change the virtual drive options from the defaults listed on the screen as needed. Use the cursor keys to navigate between the virtual drive parameters and press **Enter** to select an option.

Here are brief explanations of the virtual drive options:

- **RAID Level** – The drop-down menu lists the possible RAID levels for the virtual drive. In some cases, only one RAID level is available, depending on the number of drives in the drive group.
- **Size** – This setting specifies the capacity of the virtual drive.
- **Disk Write Cache Policy** – When the disk Write Cache is **On**, a write transaction is considered to be complete when all the data has been written to the disk cache. When disk Write Cache is **Off**, the write transaction is complete only when the data has been written to the disk.
- **Read Ahead Policy** – When disk Read Ahead is **On**, extra data is read sequentially ahead of the data that is actually requested, and this extra data is stored in cache memory. If the additional read-ahead data is then requested, it can be read faster from the cache than from the disk directly. This setting speeds up reads for sequential data, but there is little improvement when accessing random data.
- **Accept** – Select this option to accept the virtual drive parameters.
- **SPAN** – Choose whether to span drive groups. This setting is available only for RAID 10 drive groups.

9. Highlight **Accept** and press **Enter** after you select your virtual drive parameters.
The virtual drive configuration appears, as shown in the following screen.

Figure 12 Virtual Drive Configuration Parameters



10. Press any key to continue.
11. Click **Yes** at the prompt to save the configuration.

The LSI Software RAID Configuration Utility begins a background initialization of the virtual drives.

New RAID 5 virtual drives require a minimum number of drives for a background initialization to start. If there are fewer drives than the minimum, the background initialization will not start.

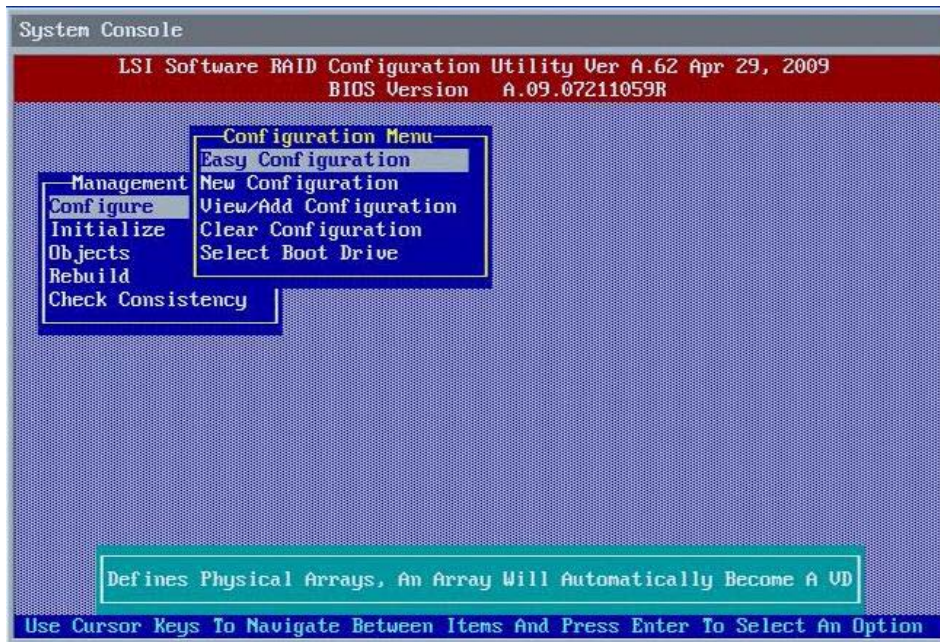
3.5.3 Using New Configuration

If you select New Configuration, the CU deletes the existing configuration information on the selected controller when it saves the new configuration.

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

1. On the Management Menu screen, highlight **Configure** with your cursor, and press **Enter**.
The Configuration Menu screen appears, as shown in the following figure.

Figure 13 Configuration Menu Screen



2. Highlight **New Configuration** with your cursor, and press **Enter**.
The following screen appears. This screen displays the drives.

Figure 14 New Configuration – Array Selection Menu



3. Press the space bar to select a drive and add it to the drive group.
4. Repeat step 3 to add additional drives to the drive group.
5. Press **F10** to continue the configuration.
The Select Configurable Array(s) dialog box appears, as shown in the following figure.

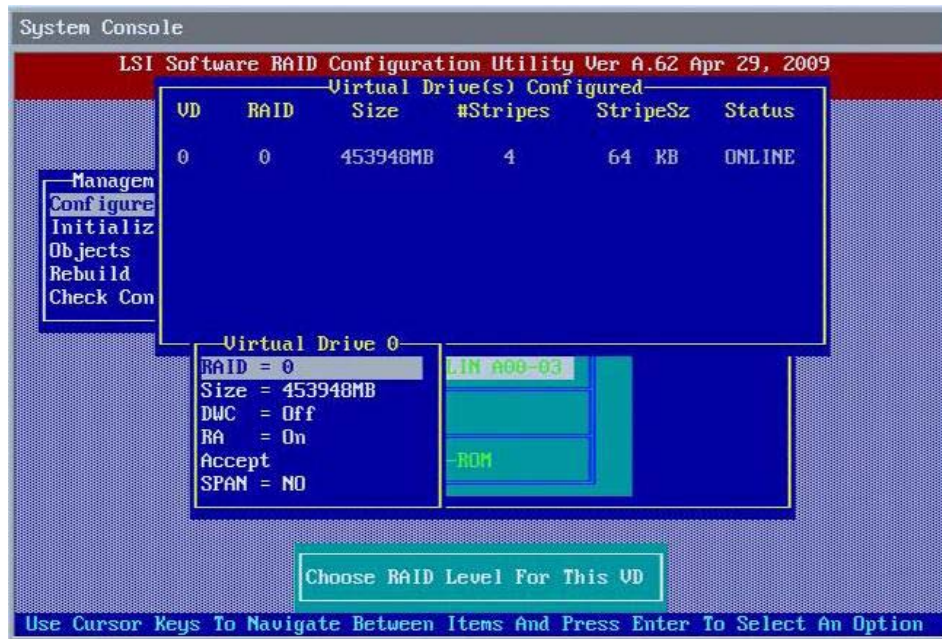
Figure 15 Drive Group Selection



6. Press the space bar to select the drive group.
7. Press **F10** to continue the configuration.

The Virtual Drive(s) Configured dialog box appears, as shown in the following screen. This screen shows the virtual drive number, RAID level, drive group size, number of stripes, stripe size and drive status.

Figure 16 Virtual Drive Parameters Dialog Box



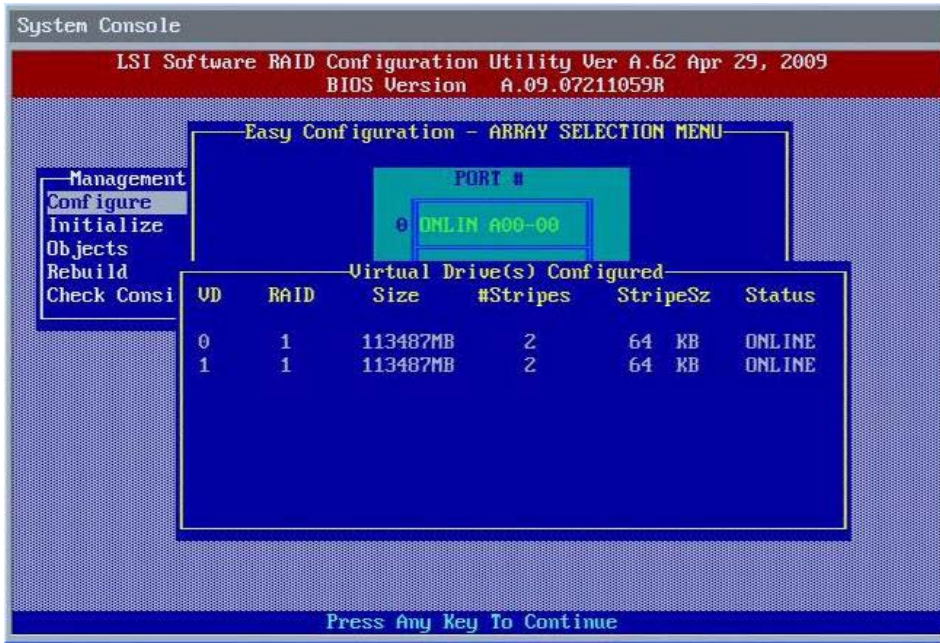
8. Change the virtual drive options from the defaults listed on the screen as needed. Use the cursor keys to navigate between the virtual drive parameters and press **Enter** to select an option.

Here are brief explanations of the virtual drive options:

- **RAID Level** – The drop-down menu lists the possible RAID levels for the virtual drive. In some cases, only one RAID level is available, depending on the number of drives in the drive group.
- **Size** – This setting specifies the capacity of the virtual drive.
- **Disk Write Cache Policy** – You can disable the write cache option when you create a virtual drive, but you can enable this option later using the configuration utilities. When the disk Write Cache is **On**, a write transaction is considered to be complete when all the data has been written to the disk cache. When disk Write Cache is **Off**, the write transaction is complete only when the data has been written to the disk.
- **Read Ahead Policy** – When disk Read Ahead is **On**, extra data is read sequentially ahead of the data that is actually requested, and this extra data is stored in cache memory. If the additional read-ahead data is then requested, it can be read faster from the cache than from the disk directly. This setting speeds up reads for sequential data, but there is little improvement when accessing random data.
- **Accept** – Select this option to accept the virtual drive parameters.
- **SPAN** – Choose whether to span drive groups. This setting is available only for RAID 10 drive groups.

9. Highlight **Accept** with your cursor, and press **Enter** after you select your virtual drive parameters.
The virtual drive configuration displays, as shown in the following screen.

Figure 17 Virtual Drive Configuration Parameters



10. Press any key to continue.
11. Click **Yes** at the prompt to save the configuration.

The LSI Software RAID Configuration Utility begins a background initialization of the virtual drives.

New RAID 5 virtual drives require a minimum number of drives for a background initialization to start. If there are fewer drives than the minimum, the background initialization will not start.

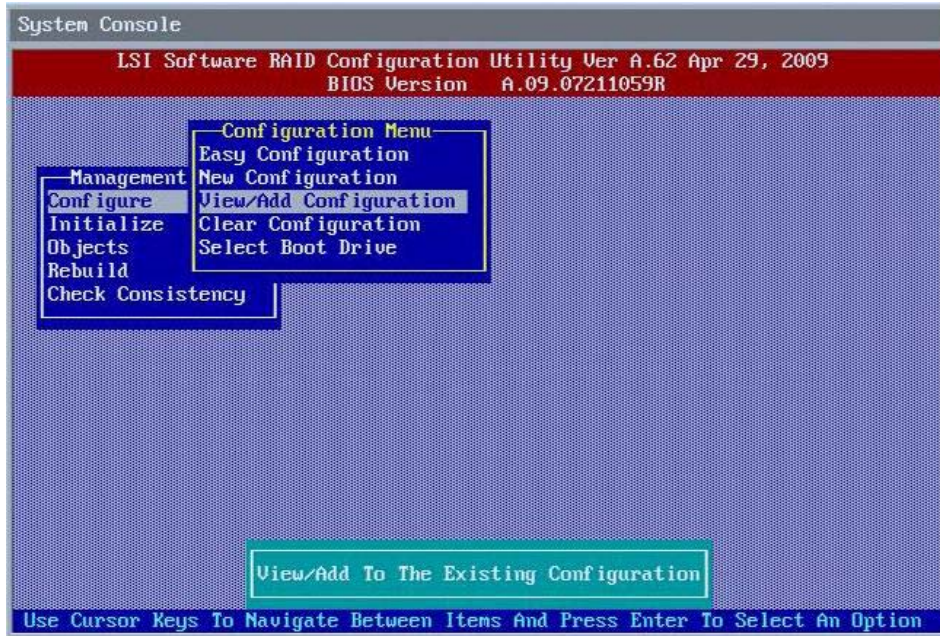
3.5.4 Using View/Add Configuration

View/Add Configuration allows you to control the same virtual drive parameters as New Configuration without disturbing the existing configuration information.

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

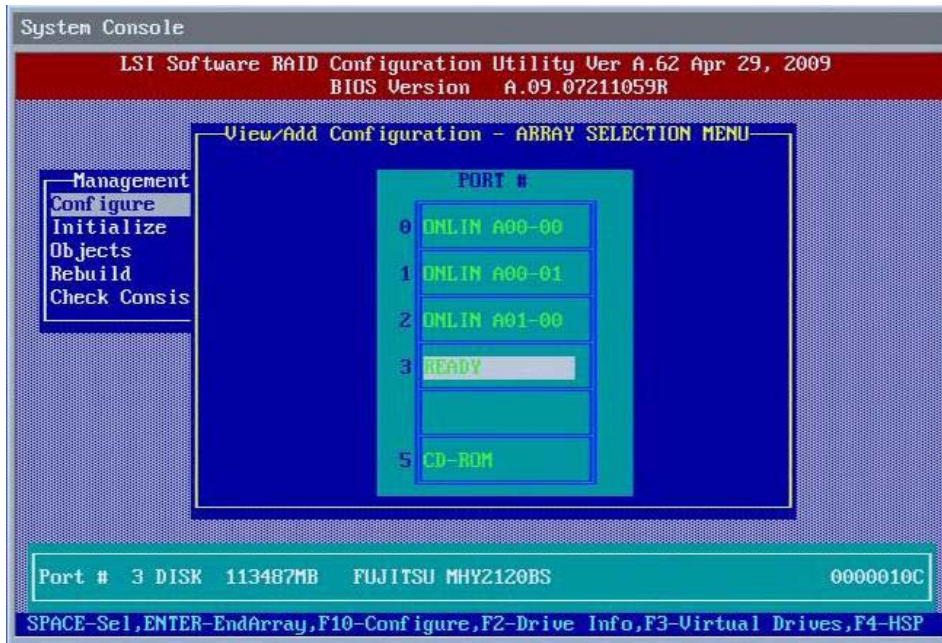
1. On the Management Menu screen, highlight **Configure** with your cursor, and press **Enter**.
The Configuration Menu screen appears, as shown in the following figure.

Figure 18 Configuration Menu Screen



2. Highlight **View/Add Configuration** with your cursor, and press **Enter**.
The following screen appears.

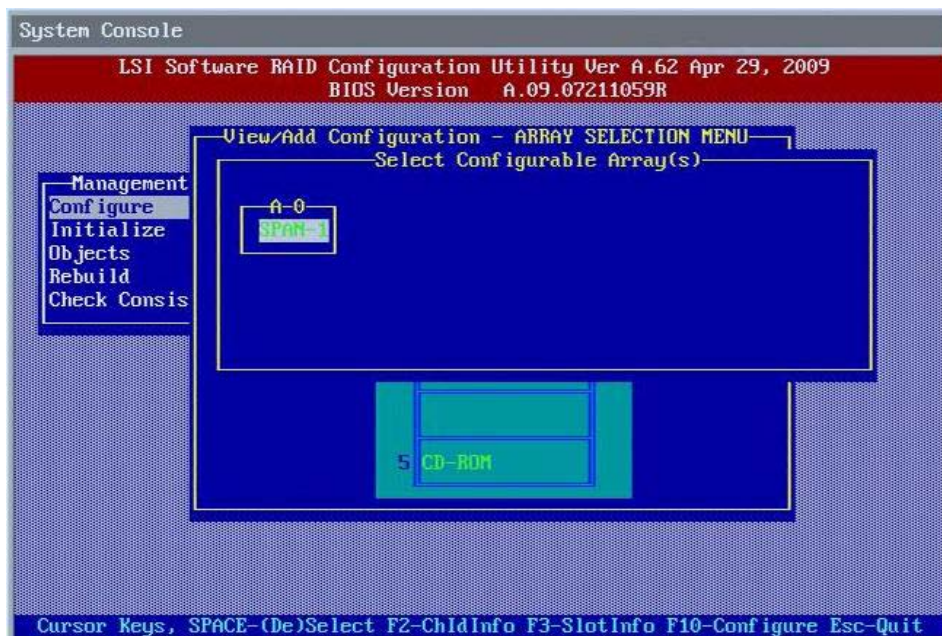
Figure 19 View/Add Configuration – Array Selection Menu



3. Press the space bar to select a drive and add it to the drive group.
4. Repeat step 3 to add additional drives to the drive group.
5. Press **F10** to continue the configuration.

The Select Configurable Array(s) dialog box appears, as shown in the following figure.

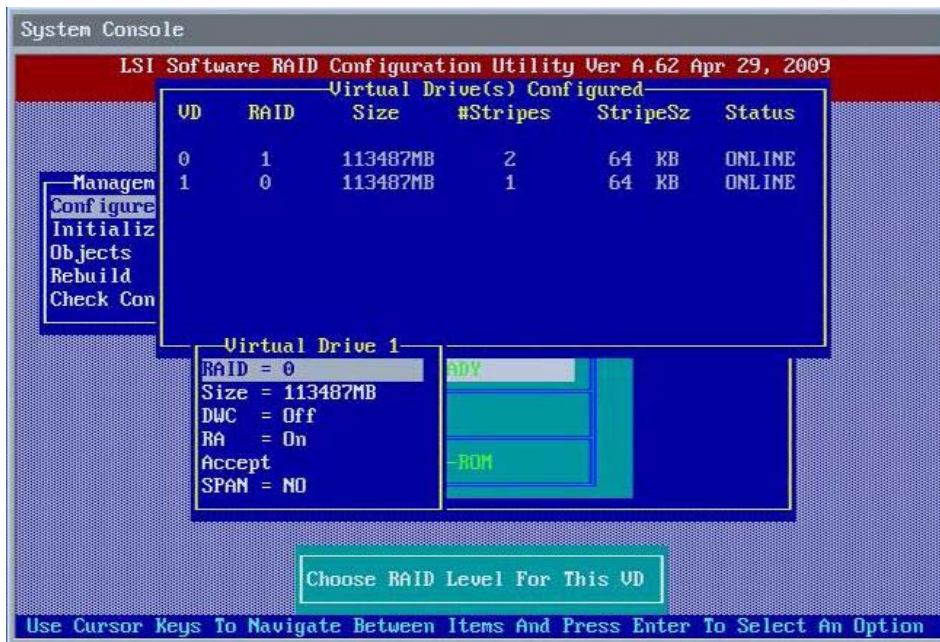
Figure 20 Drive Group Selection



6. Press the space bar to select the drive group.
7. Press **F10** to continue the configuration.

The Virtual Drive(s) Configured dialog box appears, as shown in the following screen. This screen shows the virtual drive number, RAID level, drive group size, number of stripes, stripe size and drive status.

Figure 21 Virtual Drive Parameters Dialog Box



8. Change the virtual drive options from the defaults listed on the screen as needed. Use the cursor keys to navigate between the virtual drive parameters and press **Enter** to select an option.

Here are brief explanations of the virtual drive options:

- **RAID Level** – The drop-down menu lists the possible RAID levels for the virtual drive. In some cases, only one RAID level is available, depending on the number of drives in the drive group.
- **Size** – This setting specifies the capacity of the virtual drive.
- **Disk Write Cache Policy** – You can disable the write cache option when you create a virtual drive, but you can enable this option later using the configuration utilities. When the disk Write Cache is **On**, a write transaction is considered to be complete when all the data has been written to the disk cache. When disk Write Cache is **Off**, the write transaction is complete only when the data has been written to the disk.
- **Read Ahead Policy** – When disk Read Ahead is **On**, extra data is read sequentially ahead of the data that is actually requested, and this extra data is stored in cache memory. If the additional read-ahead data is then requested, it can be read faster from the cache than from the disk directly. This setting speeds up reads for sequential data, but there is little improvement when accessing random data.
- **Accept** – Select this option to accept the virtual drive parameters.
- **SPAN** – Choose whether to span drive groups. This setting is available only for RAID 10 drive groups.

9. Highlight **Accept** with your cursor, and press **Enter**.

The virtual drive configuration displays.

10. Press any key to continue.

11. Click **Yes** at the prompt to save the configuration.

The LSI Software RAID Configuration Utility begins a background initialization of the virtual drives.

New RAID 5 virtual drives require a minimum number of drives for a background initialization to start. If there are fewer drives than the minimum, the background initialization will not start.

3.6 Clearing a Storage Configuration

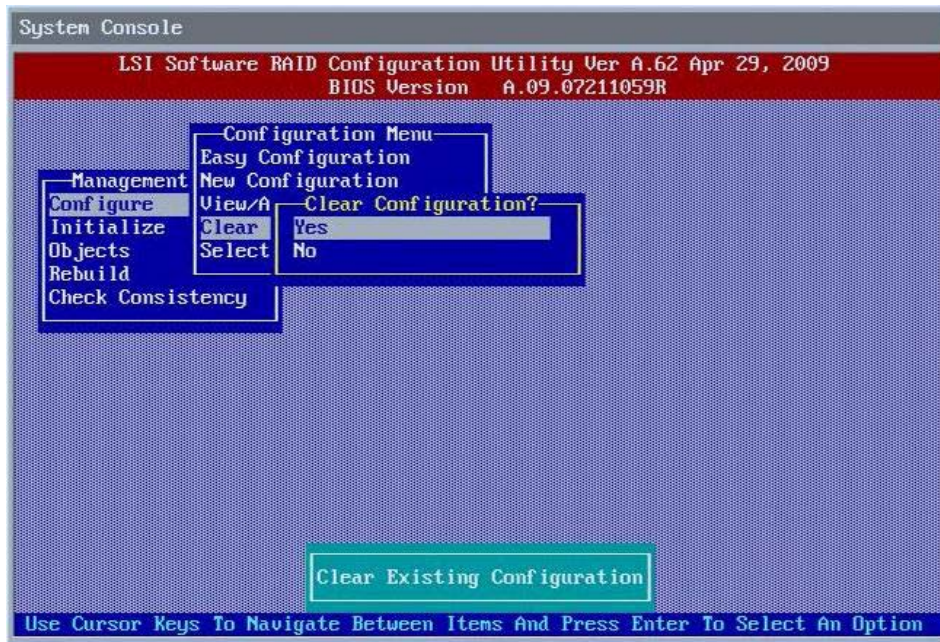


CAUTION Before you clear a storage configuration, be sure to back up all the data you want to keep.

To clear a storage configuration, follow these steps:

1. On the Management Menu screen, select **Configure > Clear Configuration**.
The following screen appears.

Figure 22 Clear Configuration Option



2. At the prompt, select **Yes** to confirm and press **Enter**.
The virtual drive is deleted from the configuration.

3.7 Configuring a Bootable Virtual Drive

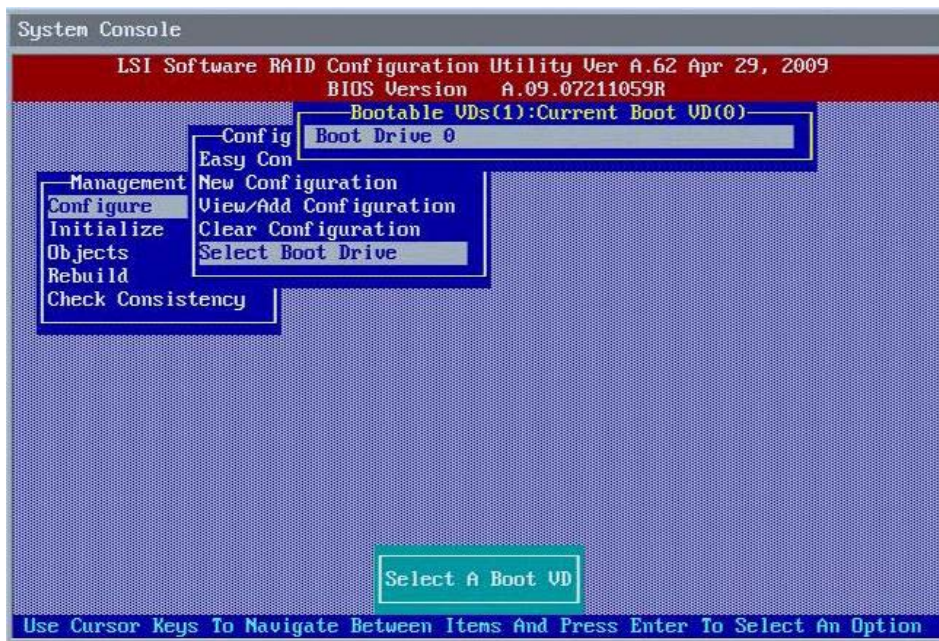
The default boot virtual drive is LD 0. If you change the boot drive to another virtual drive, the BIOS and the CU preserve this change. However, if you delete the new boot virtual drive, you must be sure to configure another virtual drive for booting. The CU does not automatically select a different boot virtual drive.

Follow these steps to configure a bootable virtual drive:

1. On the Management Menu screen, select **Configure > Select Boot Drive**.

The following screen appears.

Figure 23 Select Boot Drive Option



2. Select a virtual drive from the list to be the designated boot drive.
The virtual drive is configured to be the boot drive.

3.8 Initializing Virtual Drives

This section explains the two methods used to initialize a virtual drive with the LSI Software RAID Configuration Utility.

If the *Fast Init* property is enabled, fast initialization is used. In fast initialization, the LSI Software RAID CU quickly writes zeroes to the first and last 8-MB regions of the new virtual drive.

If the *Fast Init* property is not enabled, the utility performs a complete initialization on the virtual drive. This process can take a long time if the physical disk drives are large.



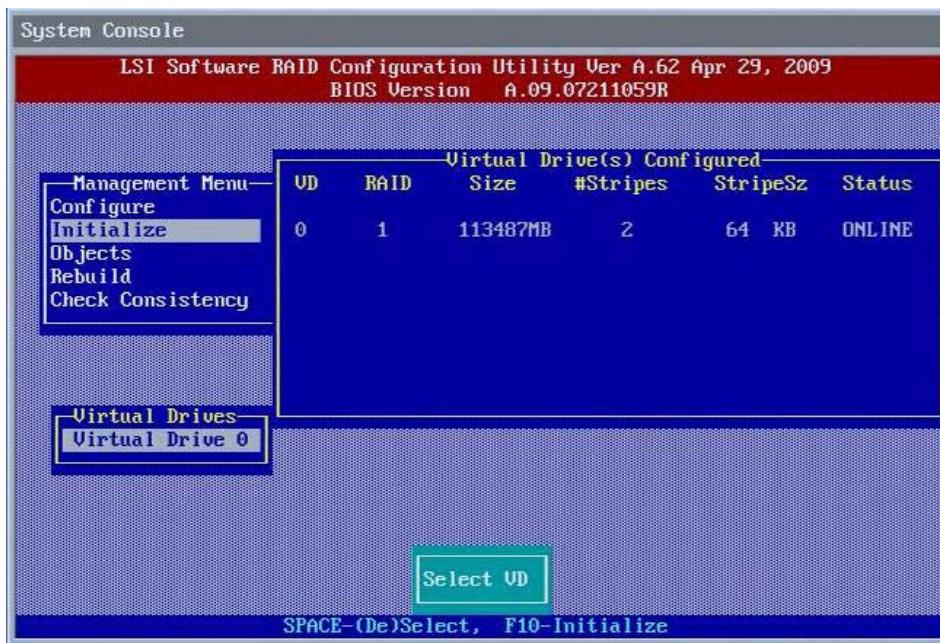
CAUTION When you initialize a virtual drive, all of the existing data on the virtual drive is erased.

3.8.1 First Initialization Method

Follow these steps to initialize a virtual drive using the Initialize menu.

1. On the Management Menu screen, highlight **Initialize** with your cursor, and press **Enter**.
The list of virtual drives appears.
2. Use the cursor to highlight the virtual drive you want to initialize (if there is more than one virtual drive).
3. Press **Enter**.
The following screen appears.

Figure 24 Initializing a Virtual Drive – First Method



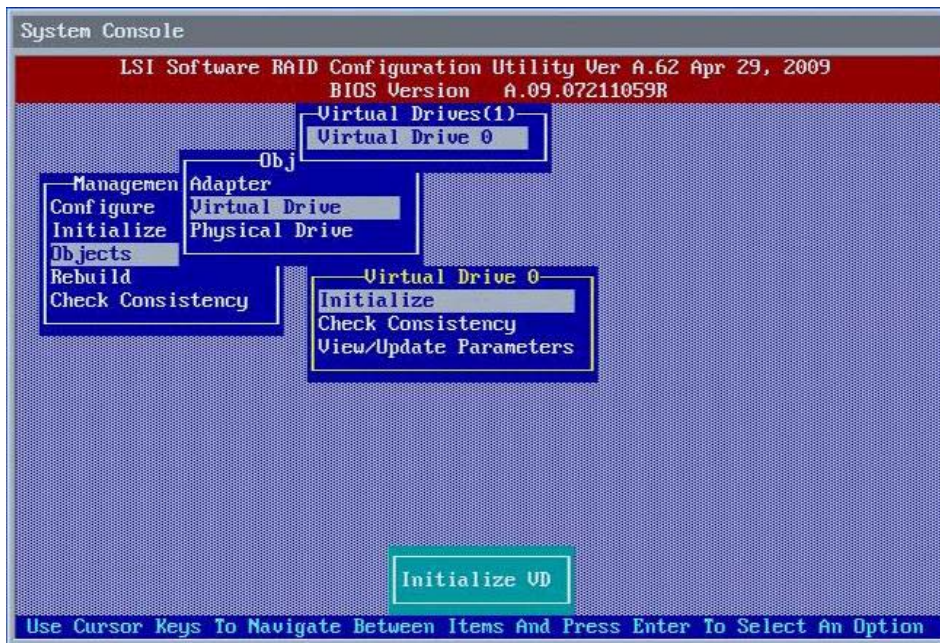
4. Press **F10**.
5. Select **Yes** at the prompt and press **Enter** to begin the initialization.
A graph shows the progress of the initialization until it is complete.
6. After the initialization is complete, press **Esc** to return to previous menus.
If you press **Esc** while initialization is in progress, the following options appear:
 - **Stop** – (Available only if *Auto Resume* is enabled on the controller: **Management Menu > Objects > Adapter > AutoResume**.) The initialization is stopped, and the CU stores the percentage of the initialization already completed. If *Auto Resume* is enabled, and if *Fast Init* is *not* enabled, the initialization resumes where it left off when you restart it, instead of starting over from 0 percent.
 - **Continue** – The initialization continues normally.
 - **Abort** – The initialization is completely aborted. If you restart initialization, it begins at 0 percent.

3.8.2 Second Initialization Method

Follow these steps to initialize a virtual drive using the Objects menu.

1. On the Management Menu screen, select **Objects > Virtual Drive**.
A list of configured virtual drives appears.
2. Use the cursor to highlight the virtual drive you want to initialize (if there is more than one virtual drive), and press **Enter**.
The following screen appears.

Figure 25 Initializing a Virtual Drive – Second Method



3. Highlight **Initialize** with your cursor on the virtual drive submenu, and press **Enter**.
4. Select **Yes** at the prompt, and press **Enter**.
The utility displays a bar graph showing the initialization progress.
5. When initialization completes, press **Esc** to return to the previous menu.

If you press **Esc** while initialization is in progress, the **Stop**, **Continue**, and **Abort** options are available, as explained in Section 3.8.1, [First Initialization Method](#), on page 39.

3.9 Rebuilding a Drive

The LSI Software RAID Configuration Utility enables you to rebuild a drive of a redundant drive group if the drive group has a failed drive. If the failed drive is still good (that is, if the drive is physically present and its capacity is greater than or equal to the defined capacity of the drive group), it will be rebuilt. If the drive is too small, an error message appears and the CU does not allow the drive to be rebuilt.



NOTE You cannot rebuild a failed drive if the drives capacity is even 1 byte smaller than the defined capacity of the drive group.

Follow these steps to rebuild a drive:

1. On the Management Menu screen, highlight **Rebuild** and press **Enter**.
2. When the list of drives appears, highlight the failed (*FAIL*) drive that you want to rebuild and press the spacebar to select it.
3. After you select the drive, press **F10** to start the rebuild, and then select **Yes** at the confirmation prompt.

The rebuild process begins, and a graph shows the progress of the rebuild until it is complete. Although the CU changes the disk drive state to *Rebuild* at this point, the change does not appear on the screen while the rebuild is in progress.

If the CU detects a media error on the source drive during rebuild, it initiates a sector read for that block. If the sector read fails, the CU adds entries to the Soft Bad Block Management (SBBM) table, writes this table to the target drive, and displays an error message.

Additional error messages appear if the SBBM table is 80 percent full or 100 percent full. If the SBBM table is completely full, the rebuild operation is aborted, and the drive is marked as *FAIL*.

4. When the rebuild is complete, the CU displays the message that the rebuild is successful.
5. Press **Esc** to display the Management Menu screen.

The state of the rebuilt disk drive changes from *FAIL* to *ONLIN*.

If you press **Esc** while the rebuild is running, the following options display:

- **Stop** – (Available only if AutoResume is enabled on the adapter: **Management Menu > Objects > Adapter > AutoResume**.)
The rebuild is stopped, and the CU stores the percentage of the rebuild already completed. If AutoResume is enabled, the rebuild resumes where it left off when you restart it, instead of starting over from 0 percent.
- **Continue** – The rebuild continues normally.
- **Abort** – The rebuild is completely aborted and the disk drive remains in the *FAIL* state. If you restart the rebuild, it begins at 0 percent.

3.10 Creating a Global Hotspare Drive

The LSI Software RAID Configuration Utility enables you to create global hot spare drives to protect against data loss. A hot spare is an unused drive that you can use to rebuild the data from a failed drive and re-establish redundancy, in case of a disk failure in a redundant RAID drive group (RAID 1, RAID 5, or RAID 10).



NOTE Dedicated hot spare drives are not supported by the LSI Software RAID Configuration Utility.



NOTE When you select a drive to change into a global hot spare, be sure it is the same type of drive as the drives in the drive group that it will protect.

You can create a hotspare when you are configuring a new storage configuration, as described in the previous sections. To add a hotspare drive to an existing redundant storage configuration, follow these steps:

1. On the Management menu, select **Configure > View/Add Configuration**.
2. Select **Physical Drive**.
A list of physical drives appears.
3. Highlight an unconfigured drive or a *Ready* drive with your cursor, and press **Enter**.
The following screen appears.



4. In the HotSpare dialog box, select **Yes** and press **Enter**.
5. Select **Yes** from the pop-up menu to create the hot spare drive.
This procedure creates a hot spare for the drive group.

3.11 Checking Data Consistency

The Check Consistency feature verifies the consistency of the data on the physical drives that are part of RAID 1, 5, or 10 virtual drives. The LSI Software RAID Configuration Utility automatically corrects any differences found in the data when a consistency check is run.

Follow these steps to check consistency:

1. On the Management Menu screen, highlight **Check Consistency** with your cursor, and press **Enter**.
A list of configured virtual drives appears.
2. Use the cursor to highlight a virtual drive (if there is more than one virtual drive).
3. Press the spacebar to select the virtual drive.



NOTE If you select a RAID 0 virtual drive, a message appears stating that a Check Consistency cannot be performed. To continue, deselect the virtual drive, highlight a redundant virtual drive, and press the spacebar again.

4. Press **F10**.
5. At the prompt, select **Yes** to start the Check Consistency process, and press **Enter**.

A graph shows the progress of the Check Consistency operation until it is complete.

If the LSI Software RAID CU finds any data inconsistencies while comparing the source drive and the target drive, the utility fixes the inconsistency by writing the source data to the target drive. When this happens, a message notifies you that inconsistent data exists on the drives and that the repair has been performed.

If you press **Y**, the program skips the bad block and continues. If you press **N**, the program aborts the consistency check. The same message appears if the program finds a hard media error on the target drive.

If you press **Esc** while a Check Consistency is running, the following options appear:

- **Stop** – (Available only if *AutoResume* is enabled on the adapter: **Management Menu > Objects > Adapter > AutoResume**.) The Check Consistency is stopped, and the CU stores the percentage of the task already completed. If *AutoResume* is enabled, the Check Consistency resumes where it left off when you restart it, instead of starting over from 0 percent.
- **Continue** – The Check Consistency continues normally.
- **Abort** – The Check Consistency is completely aborted. If you restart it, it begins at 0 percent.

3.12 Displaying and Changing Controller Properties

This section explains how you can use the LSI Software RAID Configuration Utility to display and change the properties for the controllers.

3.12.1 Displaying and Changing Controller Properties

Follow these steps to display the properties of a controller.

1. On the Management Menu screen, select **Objects > Adapter**.
The list of controllers appear in a dialog box.
2. Highlight a controller with your cursor, and press **Enter**.
The following figure appears.

Figure 26 Controller Properties Screen



Table 2 describes the entries/options listed on the controller properties screen. LSI recommends that you leave these options at their default settings to achieve the best performance, unless you have a specific reason for changing them.

Table 2 Controller Properties and Values

Property	Options	Default
Rebuild Rate	0 to 100 (percentage of system resources) Use this option to select the rebuild rate for drives connected to the selected controller. The rebuild rate is the percentage of system resources dedicated to rebuilding a failed drive. The higher the number, the more system resources devoted to a rebuild.	30
Chk Const Rate	0 to 100 (percentage of system resources) Use this option to select the amount of system resources dedicated to consistency checks of virtual drives connected to the selected controller.	30
FGI Rate (Foreground Initialization Rate)	0 to 100 (percentage of system resources) Use this option to select the amount of system resources dedicated to foreground initialization of virtual drives connected to the selected controller.	30
BGI Rate (Background Initialization Rate)	0 to 100 (percentage of system resources) Use this option to select the amount of system resources dedicated to background initialization of virtual drives connected to the selected controller.	30
Disk WC (Disk Write Cache)	Off, On You can disable the disk write cache option when you create a virtual drive, but you can enable this option later using the configuration utilities. When the disk Write Cache is On, a write transaction is considered to be complete when all the data has been written to the disk cache. When disk Write Cache is Off, the write transaction is complete only when the data has been written to the disk.	Off (Write Through Enabled)
Read Ahead (RA)	On, Off When disk Read Ahead is On, extra data is read sequentially ahead of the data that is actually requested, and this extra data is stored in cache memory. If the additional read-ahead data is then requested, it can be read faster from the cache than from the disk directly. This setting speeds up reads for sequential data, but there is little improvement when accessing random data.	On
BIOS State	Enable, Disable Use this option to enable or disable the BIOS for the selected controller. If the boot device is on the selected controller, the BIOS must be enabled; otherwise, the BIOS should be disabled or it might not be possible to use a boot device elsewhere.	Enable
Cont On Error	No, Yes Enable this option if you want the boot process to continue when the controller BIOS encounters an error during boot-up.	Yes
Fast Init	Enable, Disable 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. It is seldom necessary to use this option, because the virtual drive was already initialized when you created it.	Enable
Auto Rebuild	On, Off 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.	On

Table 2 Controller Properties and Values (Continued)

Property	Options	Default
Auto Resume	Enable, Disable When Enabled, you can stop a consistency check, rebuild, or initialization, and resume it later where it left off, instead of aborting it and starting over.	Enable
Disk Coercion NOTE The Disk Coercion property can be accessed only when no configuration is present for the controller. Otherwise, an error message appears.	None, 128MB, 1GB Drive coercion is a tool for forcing drives of varying capacities to the same capacity so they can be used in a drive group. The coercion mode options are None, 128MB-way, and 1GB-way. The number you choose depends on how much the drives from various vendors vary in their actual size.	1GB
Factory Default	Returns the settings to the factory default settings.	N/A

3. To change the value of a controller property, use the cursor to highlight the property, and then press **Enter**.
4. Select or type a different value for the property, and press **Enter**.
When you are finished, press **Esc** until you return to the Management Menu screen.

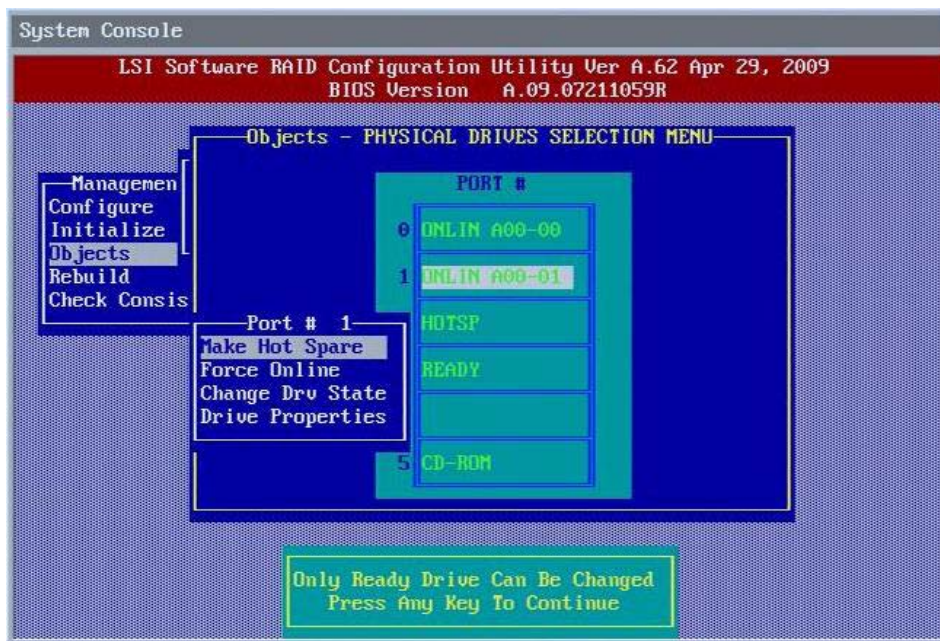
3.12.2 Displaying and Changing Drive Properties

The LSI Software RAID Configuration Utility displays properties, policies, and operations for virtual drives.

Follow these steps to display the virtual drive information.

1. On the Management Menu screen, select **Objects > Adapter**.
The list of controllers appear in a dialog box.
2. Highlight a controller with your cursor, and press **Enter**.
The Physical Drive Selection Menu screen appears.
3. Highlight a drive with your cursor, and press **Enter**.
The following screen appears. This screen lists the actions you can perform.

Figure 27 Virtual Drive Parameters Screen



4. Select the action you want to perform and press **Enter**.
The options are:
 - **Make Hot Spare:** Only a drive in *Ready* state can be changed to a hot spare.
 - **Force Online:** Only a drive in *Failed* state can be forced online.
 - **Change Drv State:** Change an online (*Onlin*) drive to *Failed* state or a hot spare to *Ready* state.
 - **Drive Properties:** Display the device type, capacity, manufacturer name and model, and revision number.

3.12.3 Viewing or Changing Virtual Drive Properties

You can disable the write cache option when you create a virtual drive, but you can enable this option later using the configuration utilities. When the Disk Write Cache is **On**, a write transaction is considered to be complete when all the data has been written to the disk cache. When Disk Write Cache is **Off**, the write transaction is complete only when the data has been written to the disk.

When disk Read Ahead is **On**, extra data is read sequentially ahead of the data that is actually requested, and this extra data is stored in a cache. If the additional read-ahead data is then requested, it can be read faster from the cache than from the disk directly.



NOTE When the Disk Write Cache is **On**, there is a danger that data could be lost if the power fails before the cached data is written to disk.

Follow these steps to view the Disk Write Cache or Read Ahead settings, or enable the Disk Write Cache setting:

1. On the Management Menu screen, select **Objects > Virtual Drive**.
The list of virtual drives appears.
2. Highlight a virtual drive with your cursor, and press **Enter**.
The list of virtual drive parameters appears.
3. Highlight **Disk WC** with your cursor, and press **Enter**.
The **On/Off** dialog box opens.
4. Change the setting to **On** to enable the Disk Write Cache setting.
The Disk Write Cache option is enabled.

3.13 Forcing Drives Online or Offline

To force a drive online or offline, follow these steps:

1. On the Management Menu, select **Objects > Physical Drive**.
2. Highlight a physical drive that is a member of an drive group, and press **Enter**.
3. Select **Force Offline** or **Force Online** from the menu.
If the drive was online, its status changes to *FAIL*. If it was offline, its status changes to *ONLINE*.

Chapter 4: MegaCLI Command Tool

This chapter describes the commands for the MegaCLI Command Tool and how to use them.

The MegaCLI Command Tool (CT) is a command line interface application that you can use to configure and maintain storage configurations created with Embedded MegaRAID Software.



NOTE The MegaCLI CT utility runs in the Microsoft® Windows® operating environment and the Linux™ operating environment. For DOS, the utility is called MegaDCLI CT and it supports a subset of the full command set.

4.1 MegaCLI CT Overview

MegaCLI CT and MegaDCLI CT are command line interface applications you can use to configure and manage storage configurations under Embedded MegaRAID Software. You can use these command tools to perform the following tasks:

- Configure virtual drives and create configurations on the adapter
- Display the configuration on the adapter
- Display and change virtual drive's properties on the adapter
- Display and change physical drive's properties on the adapter
- Display and change adapter properties
- Load a configuration to the adapter from a file
- Save an adapter configuration to a file
- Start or stop rebuild, consistency check, and initialization
- Suspend and display an ongoing background initialization
- Display relevant user messages on the console or write them to the log file
- Work in silent mode, if selected (no messages are displayed on the console)
- Display adapter inventory data in a single command
- Customize output strings
- Exit with predefined success or failure exit codes
- Set some predefined environment variables, such as number of adapters and number of virtual drives after the execution of CT
- Display help on how to use the command line options

4.2 MegaCLI Commands Not Supported by Embedded MegaRAID Software

Some MegaCLI CT commands that are supported in hardware RAID configurations are *not* supported under Embedded MegaRAID Software. These include the following commands:

- Adapter Cache Flush (-AdpCacheFlush)
- Cluster Enable (-AdpGetProperties, -AdpSetProperties)
- Diagnostics (-AdpDiag, -AdpBatTest, -AdpNVRAM)
- Display Enclosure Information (-EncInfo)
- Display List of Missing Physical Disk Drives (-PDGetMissing)
- Foreign Configuration (-CfgForeign)
- Mark Configured Physical Disk Drive as Missing (-PDMarkMissing)
- Set Virtual Drive Properties (-LDSetProp)
- Virtual Drive Reconstruction (-LDRecon)



NOTE The Embedded MegaRAID Software utility does not support Write Policy, Read Policy, or I/O Policy. Software RAID supports only enabling and disabling the drive's write cache and read-ahead functions.



NOTE MegaCLI CT for Embedded MegaRAID Software does not support the concept of disk enclosures, except for a default *Enclosure 0* (E0). Some of the commands support an E0:Sn syntax that enables you to specify a drive in a particular "slot" in this default enclosure.

4.3 Exception Handling

MegaCLI CT exits with exit code 0 for all successful operations. In case of failure, it exits with exit code 1 to 255, depending on the failure conditions. For example, assume that a rebuild is started on three physical drives. MegaCLI CT successfully starts rebuilding the first drive but fails to start rebuilding the second drive. If this happens, MegaCLI CT does not attempt to start rebuilding the third drive; instead, it exits with an error exit code. In this case, the error code is EXIT_ERR_START_RBLD. So even if the command was partially successful, an error code is still generated.

Some operations, such as `GetNumberOfAdapters` or `GetNumberOfLogicaldrives` return with the actual number of controllers or virtual drives. These return values are special cases and do not reflect any error conditions based on the return code, which in these cases contains meaningful values.

4.4 Command Line Abbreviations and Conventions

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

4.4.1 Abbreviations Used in the Command Line

Table 3 lists the abbreviations for the virtual drive parameters used in the following sections.

Table 3 Command Line Abbreviations

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

4.4.2 Conventions

You can specify multiple values for some options. You can enter commands for a single adapter (-aN), multiple selected adapters (-a0, 1, 2) or all adapters (-aALL). These options are denoted as -aN | -a0, 1, 2 | -aALL in the command line syntax used in this chapter. Table 4 lists all of the conventions used in the command line options and specifies that you can enter commands for one controller, multiple controllers, or all controllers.



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

Table 4 Conventions

Convention	Description
MegaCLI	Specifies the command line interface used. Type MegaCLI under Microsoft Windows and Linux.
	Specifies "or," meaning you can choose between options.
-aN	N specifies the adapter number for the command.
-a0, 1, 2	Specifies the command is for adapters 0, 1, and 2. You can select two or more controllers in this manner.
-aALL	Specifies the command is for all adapters.
-Lx	x specifies the virtual drive number for the command.
-L0, 1, 2	Specifies the command is for virtual drives 0, 1, and 2. You can select two or more virtual drives in this manner.
-Lall	Specifies the command is for all virtual drives.
[E0 : S0, E1, S1, ...]	Specifies when one or more physical devices need(s) to be specified in the command line. Each [E:S] pair specifies one physical device where E means device ID of the enclosure in which a drive resides, and S means the slot number of the enclosure.
[]	Indicates that the parameter is optional.
{ }	Indicates that the parameters are grouped and that they must be given at the same time.
-Force	Specifies that the MegaCLI utility does not ask you for confirmation before it performs this command. You might lose data using this option with some commands.

You can specify the `-Silent` command line option for all possible functions of the MegaCLI CT. If you enter this option at the command line, no messages appear on the screen.

4.5 Adapter Commands

You can use the commands in this section to set or display properties related to the adapter(s).



NOTE The Embedded MegaRAID Software drivers function as virtual “adapters” or “controllers.” Because the drivers are not actual hardware components, some of the adapter parameters do not apply to them.

4.5.1 Display Adapter Information

Use the command in [Table 5](#) to display information on adapter parameters such as the number of virtual drives and initiator ID.

Table 5 Display Adapter Information

Syntax	MegaCLI <code>-AdpAllInfo -aN -a0,1,2 -aALL</code>
Description	Displays parameters on the selected adapter(s). Displayed information includes initiator ID, current status of auto rebuild, alarm, number of virtual drives, rebuild rate, bus number/device number, present RAM, settings, serial number of the board, and address.

4.5.2 Enable or Disable Automatic Rebuild

Use the command in [Table 6](#) to turn automatic rebuild on or off for the selected adapter(s). If you have configured hot spares and enabled automatic rebuild, the RAID adapter automatically tries to use them to rebuild failed disks. Automatic rebuild also controls whether a rebuild starts when a drive that was part of the array is reinserted.

Table 6 Enable or Disable Automatic Rebuild

Syntax	MegaCLI <code>-AdpAutoRbld -Enbl -Dsbl -Dsply -aN -a0,1,2 -aALL</code>
Description	Enables or disables automatic rebuild on the selected adapter(s). The <code>-Dsply</code> option shows the status of the automatic rebuild state.

4.5.3 Set Adapter Properties

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



NOTE The default rebuild rate of 30 percent is recommended.

Use the command in [Table 7](#) to display or set adapter properties.

Table 7 Set Adapter Properties

Syntax	MegaCLI -AdpSetProp {RebuildRate -val} {PatrolReadRate -val} {BgiRate -val} {CCRate -val} {ReconRate -val} {CoercionMode -val} {EnblSSDPatrolRead -val} {-PrCorrectUncfgdAreas -val} -aN -a0,1,2 -aALL
Description	<p>Sets the properties on the selected controller(s). The possible settings are:</p> <ul style="list-style-type: none"> ■ RebuildRate – Rebuild rate. Values: 0 to 100. ■ PatrolReadRate – Patrol read rate. Values: 0 to 100. ■ BgiRate – Background initialization rate. Values: 0 to 100. ■ CCRate – Consistency check rate. Values: 0 to 100. ■ ReconRate – Reconstruction rate. Values: 0 to 100. ■ CoercionMode – Drive capacity Coercion mode. Values: 0 – None, 1 – 128 MB, 2 – 1 GB. <hr/> <p>NOTE The coercion type can be set only when there is no configuration present on the adapter.</p> <ul style="list-style-type: none"> ■ EnblSSDPatrolRead – Enable the patrol read operation (media scan) on a SSD. ■ PrCorrectUncfgdAreas –

4.5.4 Display Specified Adapter Properties

Use the command in [Table 8](#) to display specified properties on the selected adapter(s).

Table 8 Display Specified Adapter Properties

Syntax	MegaCLI -AdpGetProp RebuildRate PatrolReadRate BgiRate CCRate ReconRate CoercionMode -aN PrCorrectUncfgdAreas EnblSSDPatrolRead -a0,1,2 -aALL
Description	<p>Displays the properties on the selected controller(s).</p> <ul style="list-style-type: none"> ■ RebuildRate – Rebuild rate. Values: 0 to 100. ■ PatrolReadRate – Patrol read rate. Values: 0 to 100. ■ BgiRate – Background initialization rate. Values: 0 to 100. ■ CCRate – Consistency check rate. Values: 0 to 100. ■ ReconRate – Reconstruction rate. Values: 0 to 100. ■ CoercionMode – Drive capacity Coercion mode. Values: 0 – None, 1 – 128 MB, 2 – 1 GB. ■ PrCorrectUncfgdAreas – ■ EnblSSDPatrolRead – Enable the patrol read operation (media scan) on a SSD.

4.5.5 Set Time and Date on Controller

Use the command in [Table 9](#) to set the time and date on the selected controller(s).

Table 9 Set Time and Date on Controller

Convention	MegaCLI -AdpSetTime yyyyymmdd HH:mm:ss -aN -a0,1,2 -aALL
Description	Sets the time and date on the controller. This command uses a 24-hour format. For example, 7 p.m. displays as 19:00:00. The order of date and time is reversible.

4.5.6 Display Adapter Time

Use the command in [Table 10](#) to display the current time and date of the selected adapter.

Table 10 Display Adapter Time

Syntax	MegaCLI -AdpGetTime -aN
Description	Displays the time and date on the controller. This command uses a 24-hour format. For example, 7 p.m. would display as 19:00:00. The order of date and time is reversible.

4.5.7 Set Factory Defaults

Use the command in [Table 11](#) to set the factory defaults on the selected adapter(s).

Table 11 Set Factory Defaults

Syntax	MegaCLI -AdpFacDefSet -aN -a0,1,2 -aALL
Description	Sets the factory defaults on the selected adapter(s). You cannot set the factory defaults if the adapter already has a configuration defined on it.

4.6 Event Log Commands

4.6.1 Manage the Event Log Entries

Use the command in [Table 12](#) to manage the event entries in the event log for the selected adapter(s).

Table 12 Event Log Management

Syntax	MegaCLI -AdpEventLog -GetEventlogInfo {-GetEvents GetSinceShutdown GetSinceReboot IncludeDeleted {GetLatest <number>} -f <filename>} Clear -aN -a0,1,2 -aALL
Description	<p>The RAID driver maintains a volatile circular list of 100 events, which is deleted at reboot. The following command options are available:</p> <ul style="list-style-type: none"> ■ -GetEventlogInfo - Displays overall event information such as total number of events, newest sequence number, oldest sequence number, shutdown sequence number, reboot sequence number, and clear sequence number. ■ -GetEvents - Gets event log entry details. The information shown consists of the total number of entries and the details of each error log entry. <i>Start_entry</i> specifies the initial event log entry when displaying the log. ■ -GetSinceShutdown - Displays all the events since last adapter shutdown. ■ -GetSinceReboot - Displays all the events since last adapter reboot. ■ -IncludeDeleted - Displays all events, including deleted events. ■ -GetLatest - Displays the latest number of events, if it exists. The event data is written to the file in reverse order. ■ -Clear - Clears the event log for the selected adapter(s).

4.7 Configuration Commands

You can use the commands in this section to create storage configurations.

4.7.1 Add RAID 0, RAID 1, or RAID 5 Configuration

Use the command in [Table 13](#) to add a RAID level 0, 1, or 5 configuration to the existing configuration on the selected adapter. For RAID level 10, see [Section 4.7.3, Add RAID 10 Configuration](#).

Table 13 Add RAID 0, 1, or 5 Configuration

Syntax	MegaCLI -CfgLDAdd -R0 -R1 -R5[[E0]:Sn] [-szXXXXXXXX [-szYYYYYYY [...]]] [-strpszM] [-Hsp[E5:S5,...]] [-afterLdX] -aN
Description	<p>Adds a RAID level 0, 1 or 5 configuration to a specified adapter. Even if no configuration is present, you have the option to write the configuration to the adapter.</p> <ul style="list-style-type: none"> ■ -Rx[[E0]:Sn] – Used to specify the RAID level and the physical drive enclosure/slot numbers for a drive group. <hr/> <p>NOTE No enclosure number is needed for simple swap configurations, in which a drive is not part of an enclosure. For example, in a RAID 1 configuration with a drive in slot 2 that is not in an enclosure, and a drive in an enclosure in slot 3, the syntax is -R1[:2,1:3].</p> <hr/> <ul style="list-style-type: none"> ■ -strpszM – Used (optionally) to specify a stripe size. You can set the stripe size to 64 KB. ■ -Hsp[Ex:Sx,...] – Used to create a global hot spare, with physical drive enclosure/slot numbers specified. ■ -szXXXX – Used to specify the size of a virtual drive, where XXXX is a decimal number of MB. However, the actual size of the virtual drive might be smaller, because the driver requires the number of blocks from physical drives in each virtual drive to be aligned to the stripe size. This option can also be used to create a configuration on the free space available in the array. ■ -AfterLdX – Used to specify which free slot should be used (optional). By default, MegaCLI CT uses the first free slot available in the array. This option has no meaning if the array is not already used for configuration.

4.7.2 Configure Each Disk as RAID 0

Use the command in [Table 14](#) to configure each physical disk in Unconfigured-Good state as RAID 0.



NOTE The MegaCLI CT does *not* support spanning across these single-drive RAID 0 configurations.

Table 14 Configure Each Disk as RAID 0

Syntax	MegaCLI -CfgEachDskRAID0 [{WT WB}] [{NORA RA ADRA}] [{Direct Cached}] [{-strpszM} -aN -a0,1,2 -aALL
Description	<p>Configures each physical disk in Unconfigured-Good state as RAID 0 on this controller. The options {WT WB} {NORA RA ADRA} {Direct Cached} must be entered in the sequence that is shown.</p> <ul style="list-style-type: none"> ■ WT – Used to select Write-through caching, in which a write transaction is considered to be complete when all the data has been written to the disk cache. ■ WB – Used to select Write-back caching, in which the write transaction is complete only when the data has been written to the disk. ■ NORA – Used to select Normal Read Ahead caching, which specifies that the controller reads only the requested data and does not read ahead for the current virtual drive. ■ RA – Used to select Read Ahead caching, which specifies that data is read sequentially ahead of the data that is actually requested and is stored in a cache. If the additional read-ahead data is then requested, it can be read faster from the cache than from the disk directly. Read-Ahead supplies sequential data faster, but is not as effective when accessing random data. ■ ADRA – Used to select Adaptive Read Ahead, which specifies that the controller begins using Read Ahead caching if the two most recent disk accesses occurred in sequential sectors. If all read requests are random, the algorithm reverts to No Read Ahead; however, all requests are still evaluated for possible sequential operation. <p>If you select Read Ahead, there is a danger that data could be lost if the power fails before the cached data is written to disk.</p> <ul style="list-style-type: none"> ■ Direct – Used to specify that the controller does not buffer reads 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. ■ Cached – Used to specify that the controller buffers all reads in cache memory. ■ {-strpszM} – Used to specify the size of the segments written to each drive in the configuration. You can set the stripe size to 64 KB.

4.7.3 Add RAID 10 Configuration

Use the command in [Table 15](#) to add a RAID 10 configuration to the existing configuration on the selected adapter. For RAID levels 0, 1, or 5, see [Section 4.7.1, Add RAID 0, RAID 1, or RAID 5 Configuration](#), on page 56.



NOTE On a RAID 10 array, you can create only one virtual drive, and that virtual drive must occupy the entire space of the RAID 10 array.

Table 15 Add RAID 10 Configuration

Syntax	MegaCLI -CfgSpanAdd -R10 -Array0[E0:Sn] -Array1[E0:Sn] [...] [{WT WB}] [{NORA RA ADRA}] [{Direct Cached}] [{-strpszM}] -aN
Description	Creates a RAID level 10 (spanned) configuration from the specified arrays. Even if no configuration is present, you must use this option to write the configuration to the adapter. Multiple arrays are specified using the -ArrayX[E0:Sn, ...] option. (Note that X starts from 0, not 1.) All the arrays must have the same number of physical drives. At least two arrays must be provided. The options {WT WB} {NORA RA ADRA} {Direct Cached} must be entered in the sequence that is shown.

4.7.4 Clear Existing Configuration

Use the command in [Table 16](#) to clear the existing storage configuration on the selected adapter(s).

Table 16 Clear Existing Configuration

Syntax	MegaCLI -CfgClr -aN -a0,1,2 -aALL
Description	Clears the existing storage configuration.

4.7.5 Display Existing Configuration

Use the command in [Table 17](#) to display the virtual drive and physical disk drive information for the configuration on the selected adapter(s). This command also provides information about the remaining unconfigured space.

Table 17 Display Existing Configuration

Syntax	MegaCLI -CfgDsply -aN -a0,1,2 -aALL
Description	Displays the existing configuration on the selected adapter(s), which includes the virtual drive and component physical drive related details.

4.7.6 Save Adapter Configuration

Use the command in [Table 18](#) to save the configuration for the selected adapter(s) to the given filename.

Table 18 Save Adapter Configuration

Syntax	MegaCLI -CfgSave -fFileName -aN
Description	Saves the configuration for the selected adapter(s) to the given filename, in binary format. The command also stores the controller properties structure in the file.

4.7.7 Restore Configuration Data from File

Use the command in [Table 19](#) to read the configuration from the file and load it on the selected controller(s). You can restore the read/write properties and RAID configuration using hot spares.

Table 19 Restore Configuration Data from File

Syntax	MegaCLI -CfgRestore -fFileName -aN
Description	<p>Reads the configuration from the file and loads it on the adapter. MegaCLI can store or restore all read and write adapter properties, all read and write properties for virtual drives, and the RAID configuration including hot spares.</p> <hr/> <p>NOTE The -CfgSave option stores the configuration data and adapter properties in the file. Configuration data has only the device ID and sequence number information of the physical drives used in the configuration. The CfgRestore option will fail if the same device IDs of the physical drives are not present. The utility does not validate the setup before restoring the configuration.</p>

4.7.8 Delete Virtual Drive(s)

Use the command in [Table 20](#) to delete one or more virtual drives on the selected adapter(s).

Table 20 Delete Virtual Drives

Syntax	MegaCLI -CfgLDDe1-Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL
Description	Deletes one virtual drive, multiple virtual drives, or all the selected virtual drives on selected adapter(s).

4.7.9 Display Free Space

Use the command in [Table 21](#) to display the free space that is available to use for configuration on the selected adapter(s).

Table 21 Display Free Space

Syntax	MegaCLI -CfgFreeSpaceInfo -aN -a0,1,2 -aALL
Description	Displays all the free space available for configuration on the selected adapter(s). The information includes the number of disk groups, the number of spans in each disk group, the number of free space slots in each disk group, the start block, and the size (in both blocks and MB) of each free space slot.

4.8 Virtual Drive Commands

You can use the commands in this section to select settings for the virtual drives and to perform actions on them.

4.8.1 Display Virtual Drive Information

Use the command in [Table 22](#) to display information about virtual drives on the selected adapter(s).

Table 22 Display Virtual Drive Information

Syntax	MegaCLI -LDInfo -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL
Description	Displays information about the virtual drive(s) on the selected adapter(s). This information includes the name, RAID level, RAID level qualifier, size in MB, state, stripe size, number of drives, and span depth. It also displays activity progress, if any, including initialization, background initialization, and consistency check.

4.8.2 Display Virtual Drive Disk Cache Settings

Use the command in [Table 23](#) to display the disk cache settings for the virtual drive(s) on the selected adapter(s).

Table 23 Display Virtual Drive Cache Settings

Syntax	MegaCLI -LDGetProp -DskCache -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL
Description	<p>Displays the disk cache settings of the virtual drive(s).</p> <ul style="list-style-type: none"> ■ -DskCache – Displays physical drive cache policy. <hr/> <p>NOTE Embedded MegaRAID Software does not support setting virtual drive parameters.</p>

4.8.3 Manage Virtual Drive Initialization

Use the command in [Table 24](#) to manage initialization of the virtual drive(s) on the selected adapter(s).



NOTE This command is not supported in DOS.

Table 24 Manage Virtual Drive Initialization

Syntax	MegaCLI -LDInit -Start [Fast Full] -Abort -ShowProg -ProgDsply -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL
Description	<p>Allows you to select the following actions for virtual drive initialization:</p> <ul style="list-style-type: none"> ■ -Start – Starts the initialization (writing 0s) on the virtual drive(s) and displays the progress (optional). The fast initialization option initializes the first 100 MB on the virtual drive. The full option allows you to initialize the entire virtual drive. ■ -Abort – Aborts the ongoing initialization on the LD(s). ■ -ShowProg – Displays the snapshot of the ongoing initialization, if any. ■ -ProgDsply – Displays the progress of the ongoing initialization until at least one initialization is completed or a key is pressed.

4.8.4 Manage Consistency Check

Use the command in [Table 25](#) to manage a data consistency check (CC) on the virtual drives for the selected adapter(s).



NOTE This command is not supported in DOS.

Table 25 Manage Consistency Check

Syntax	MegaCLI -LDCC -Start -Abort -ShowProg -ProgDsply -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL
Description	<p>Allows you to select the following actions for a data consistency check:</p> <ul style="list-style-type: none"> ■ -Start - Starts a CC on the virtual drive(s), and then displays the progress (optional) and time remaining. ■ -Abort - Aborts an ongoing CC on the virtual drive(s). ■ -ShowProg - Displays a snapshot of an ongoing CC. ■ -ProgDsply - Displays ongoing CC progress until at least one CC is completed or a key is pressed.

4.8.5 View Ongoing Background Initialization

Use the command in [Table 26](#) to view ongoing background initialization of the selected virtual drives, after the -LDInit command has been issued to start the initialization. This function completes only when all background initialization processes complete or the user presses a key to exit.



NOTE This command is not supported in DOS.

Table 26 View Ongoing Background Initialization

Syntax	MegaCLI -LDBI -Enbl -Dsbl -GetSetting -ShowProg -ProgDsply -Lx -L0,1,2 -Lall -aN -a0,1,2 -aALL
Description	<p>Displays ongoing background initialization of the selected virtual drives. This function completes only when all background initialization processes complete or the user presses a key to exit.</p> <ul style="list-style-type: none"> ■ -Enbl, -Dsbl - Enables or disables the background initialization on the given adapter(s). ■ -ProgDsply - Allows the user to view ongoing background initialization until all background initialization processes complete or the user presses a key to exit. ■ -ShowProg - Displays current progress value. ■ -GetSetting - Displays current background initialization setting (enabled or disabled).

4.8.6 Display Virtual Drive and Physical Drive Information

Use the command in [Table 27](#) to display information about the virtual drives and physical disk drives for the selected adapter(s), such as the number of virtual drives, RAID level, and physical disk drive size.

Table 27 Display Virtual Drive and Physical Disk Drive Information

Syntax	MegaCLI -LDPDInfo -aN -a0,1,2 -aALL
Description	Displays information about the virtual drive(s) and physical disk drive(s) on the selected adapter(s). Displays information including the number of virtual drives, the RAID level of the virtual drives, and physical drive size information, which includes the raw size, coerced size, uncoerced size, and address.

4.8.7 Display Number of Virtual Drives

Use the command in [Table 28](#) to display the number of virtual drives attached to the adapter.

Table 28 Display Number of Virtual Drives

Syntax	MegaCLI -LDGetNum -aN -a0,1,2 -aALL
Description	Displays the number of virtual drives attached to the adapter. The return value is the number of virtual drives.

4.9 Drive Commands

You can use the commands in this section to select settings for the drives and perform actions on them.

4.9.1 Display Drive Information

Use the command in [Table 29](#) to display information about the drives on the selected adapter(s).

Table 29 Display Physical Disk Drive Information

Syntax	MegaCLI -PDInfo -aN -a0,1,2 -aALL -PhysDrv[E0:Sn,...]
Description	Provides information about the drives connected to the enclosure and controller slot. This includes information such as the enclosure number, slot number, device ID, sequence number, drive type, capacity (if a drive), foreign state, and inquiry data. <ul style="list-style-type: none"> ■ -Physdrv[E0:S0,...]: Specifies the physical drive enclosure and the slots for the drives to provide information about.

4.9.2 Set the Drive State to Online

Use the command in [Table 30](#) to set the state of a drive to Online. In an online state, the drive is working normally and is a part of a configured virtual drive.

Table 30 Set Drive State to Online

Syntax	MegaCLI -PDOnline -PhysDrv[E0:Sn....] -aN -a0,1,2 -aALL
Description	Changes the drive state to <i>Online</i> . <ul style="list-style-type: none"> ■ -Physdrv[E0:S0,...]: Specifies the physical drive enclosure and the slots for the drives.

4.9.3 Set the Drive State to Offline

Use the command in [Table 31](#) to set the state of a drive to Offline. In the offline state, the virtual drive is not available to the controller.

Table 31 Set Physical Disk Drive State to Offline

Syntax	MegaCLI -PDOffline -PhysDrv[E0:S....] -aN -a0,1,2 -aALL
Description	Changes the drive state to <i>Offline</i> . <ul style="list-style-type: none"> ■ -Physdrv[E0:S0,...]: Specifies the physical drive enclosure and the slots for the drives.

4.9.4 Change the Drive State to Unconfigured-Good

Use the command in [Table 32](#) to change the state of a drive from Unconfigured-Bad to Unconfigured-Good.

Table 32 Set Physical Disk Drive State to Unconfigured-Good

Syntax	MegaCLI -PDMakeGood -PhysDrv[E0:Sn...] -aN -a0,1,2 -aALL
Description	<p>Changes the drive state to <i>Unconfigured Good</i>.</p> <ul style="list-style-type: none"> ■ -Physdrv[E0:S0,...]: Specifies the physical drive enclosure and the slots for the drives. ■ Force: Force the drive to the <i>Unconfigured Good</i> state.

4.9.5 Manage a Drive Initialization

Use the command in [Table 33](#) to manage a drive initialization on the selected controller(s).

Table 33 Drive Initialization

Convention	MegaCLI -PDClear -Start -Stop -ShowProg -ProgDsply -PhysDrv[E0:S0,E1:S1...] -aN -a0,1,2 -aALL
Description	<p>Manages initialization or displays initialization progress on a single controller, multiple controllers, or all controllers:</p> <ul style="list-style-type: none"> ■ -Start: Starts initialization on the selected drive(s). ■ -Stop: Stops an ongoing initialization on the selected drive(s). ■ -ShowProg: Displays the current progress percentage and time remaining for the initialization. This option is useful for running the application through scripts. ■ -ProgDsply: Displays the ongoing clear progress. The routine continues to display the initialization progress until at least one initialization is completed or a key is pressed.

4.9.6 Manage Global Hot Spares

Use the command in [Table 34](#) to manage the configuration and assignment of global hot spares. Make sure the capacity of the hot spare drive is equal to or larger than the capacity of the disks in the drive group and that it is the same type of drive.



NOTE Dedicated hot spare drives are not supported by the MegaCLI Command Tool.



NOTE The hot spare drive rebuilds a failed drive even if the hot spare drive and the drives in the drive group are different types of drives. When the rebuild is completed, LSI recommends that you replace the new drive group member with a drive of the same type.

Table 34 Manage Hot Spares

Syntax	MegaCLI -PDHSP {-Set -Rmv} -PhysDrv[E0:Sn...] -aN -a0,1,2 -aALL
Description	<ul style="list-style-type: none"> ■ -Set - Changes the physical disk drive state to hot spare for the enclosure. ■ -Rmv - Changes the physical drive state to ready (removes the hot spare).

4.9.7 Rebuild a Drive

Use the command in [Table 35](#) to start or stop a rebuild on a physical disk drive and display the rebuild progress. When a physical disk in an array fails, you can rebuild the physical disk by recreating the data that was stored on the physical disk before it failed.

Table 35 Rebuild Physical Disk Drive

Syntax	MegaCLI -PDRbld -Start -Stop -ShowProg -ProgDsply -PhysDrv [E0:Sn...] -aN -a0,1,2 -aALL
Description	<p>Manages a drive rebuild or displays the rebuild progress on a single controller, multiple controllers, or all controllers. Note that the drive must meet the capacity requirements before it can be rebuilt, and it must be part of a drive group:</p> <ul style="list-style-type: none"> ■ -Start: Starts a rebuild on the selected drive(s) and displays the rebuild progress (optional). ■ -Stop: Stops an ongoing rebuild on the selected drive(s). ■ -ShowProg: Displays the current progress percentage and time remaining for the rebuild. This option is useful for running the application through scripts. ■ -ProgDsply: Displays the ongoing rebuild progress. This routine displays the rebuild progress until at least one initialization is completed or a key is pressed. ■ -Physdrv [E0:S0, . . .]: Specifies the physical drive enclosure and the slots for the drives.

4.9.8 Locate Physical Disk Drive(s) and Activate LED

Use the command in [Table 36](#) to locate physical disk drive(s) by flashing the drive activity LED.

Table 36 Locate Drive and Activate the LED

Syntax	MegaCLI -PDLocate -Start -Stop -PhysDrv[E0:Sn...] -aN -a0,1,2 -aALL
Description	<p>Locates the drive(s) for the selected controller(s) and activates the drive activity LED.</p> <ul style="list-style-type: none"> ■ -Physdrv [E0:S0, . . .]: Specifies the physical drive enclosure and the slots for the drives.

4.9.9 Replace Configured Disk Drives and Start Automatic Rebuild

Use the command in [Table 37](#) to replace a configured physical disk drive and start an automatic rebuild of the drive.

Table 37 Replace Configured Disk Drives and Start Automatic Rebuild

Syntax	MegaCLI -PDReplaceMissing -PhysDrv[E0:Sn] -ArrayX -RowY -aN
Description	Replaces the configured physical drives, and then starts an automatic rebuild. The specified array Index and row must be a missing drive.

4.9.10 Prepare Unconfigured Physical Drives for Removal

Use the command in [Table 38](#) to prepare an unconfigured physical disk drive(s) for removal from the selected adapter(s).

Table 38 Prepare Unconfigured Physical Disk Drives for Removal

Syntax	MegaCLI -PDPrpRmv [-Undo] -PhysDrv[E0:Sn...] -aN -a0,1,2 -aALL
Description	<p>Prepares unconfigured physical drive(s) for removal. The drive is spun down, and the drive state is set to unaffiliated, which marks it as offline even though it is not a part of configuration.</p> <p>The -Undo option undoes this operation, and the physical disk is marked as Unconfigured-Good.</p>

4.9.11 Display Number of Physical Drives

Use the command in [Table 39](#) to display the total number of physical disk drives attached to an adapter.

Table 39 Display Number of Physical Disk Drives

Syntax	MegaCLI -PDGetNum -aN -a0,1,2 -aALL
Description	Displays the total number of physical disk drives attached to an adapter. The return value is the number of physical disk drives.

4.9.12 Display List of Physical Drives

Use the command in [Table 40](#) to display a list of the physical drives connected to the selected adapter(s).

Table 40 Display List of Physical Drives

Syntax	MegaCLI -PDList -aN -a0,1... -aAll
Description	Displays information about all physical disk drives connected to the selected adapter(s). This includes information, such as the drive type, size, and serial number.

4.9.13 Download Firmware to the Physical Devices

Use the command in [Table 41](#) to download firmware to the physical devices connected to the selected adapter(s).

Table 41 Download Firmware to the Physical Devices

Syntax	MegaDCLI -PdFwDownload [offline] {[-SataBridge] -PhysDrv[E0:Sn...]} {EncdevId[devId1,devId2,...]} -f <filename> -aN -a0,1,2 -aALL
Description	Flashes the firmware with the file specified at the command line. The firmware files used to flash a physical device can be of any format. The CLI utility assumes that you provide a valid firmware image, and it flashes the same. The physical device needs to do error checking. <ul style="list-style-type: none"> ■ -SataBridge: Allows you to download the SATA bridge firmware in online mode. ■ -Physdrv[E0:S0,...]: Specifies the physical drive enclosure and the slots for the drives. Flashes the firmware to only one physical drive. ■ -EncdevId[devId]: Specifies the enclosure device ID. See Section 5.16.1, Display Enclosure Information for more enclosure information.

4.10 Miscellaneous Commands

The commands in this section are used to display various information about the MegaCLI utility.

4.10.1 Display Version Information

Use the command in [Table 42](#) to display the version number of the MegaCLI utility, the version of the device driver, the firmware versions for the attached physical device, and the enclosure.

Table 42 Display MegaCLI Version

Syntax	MegaCLI -v -Cli Ctrl Driver Pd -aN
Description	Displays the firmware versions and the other code levels installed on the controller, the MegaCLI version, the version of the device driver, the firmware versions for the attached physical device, and the enclosure in a list as location information, model string, and firmware version.

4.10.2 Display MegaCLI Version

Use the command in [Table 43](#) to display the version number of the MegaCLI utility.

Table 43 Display MegaCLI Version

Syntax	MegaCLI -v
Description	Displays the version number of the MegaCLI utility.

4.10.3 Display Help for the MegaCLI Utility

Use the command in [Table 44](#) to display help information for the MegaCLI utility.

Table 44 Display MegaCLI Help

Syntax	MegaCLI -h -Help ?
Description	Displays help for the MegaCLI utility.

4.10.4 Display Summary Information

Use the command in [Table 45](#) to display help information for the MegaCLI utility.

Table 45 Display MegaCLI Help

Syntax	MegaCLI -ShowSummary [-f filename] -aN
Description	Displays a summary of the system information, the controller information, the drive information, the virtual drive information, and the enclosure information.

Chapter 5: MegaRAID Storage Manager Overview and Installation

MegaRAID Storage Manager (MSM) is a configuration and monitoring utility used with Embedded MegaRAID Software. This chapter provides a brief overview of the MegaRAID Storage Manager and explains how to install it on the supported operating systems.

5.1 Overview

MegaRAID Storage Manager can be installed either on the local system which contains the MegaRAID controllers that you wish to monitor, or on a remote system such as a workstation or a laptop with network connectivity to the system you wish to monitor. Because of this, the system hardware and software requirements for installing MegaRAID Storage Manager might differ from the system requirements for the MegaRAID controller(s) you are using.

MegaRAID Storage Manager enables you to configure, monitor, and maintain storage configurations created under Embedded MegaRAID Software. The MegaRAID Storage Manager graphical user interface (GUI) makes it easy for you to create and manage storage configurations.



NOTE MegaRAID Storage Manager can be used to manage a wide range of MegaRAID controllers. Some MegaRAID Storage Manager features are not applicable for Embedded MegaRAID Software.

5.1.1 Creating Storage Configurations

MegaRAID Storage Manager enables you to easily configure the controllers, disk drives, and virtual disks on your workstation or server. The Simple Configuration mode and the Advanced Configuration mode greatly simplify the process of creating drive groups and virtual drives. You can use these configuration modes to create the best possible configuration with the available hardware.

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.

See Section 7.1, [Creating a New Storage Configuration](#), on page 80, for the procedures used to create storage configurations.

5.1.2 Monitoring Storage Devices

MegaRAID Storage Manager displays the status of virtual disks, physical disks, and other storage devices on the workstation or server that you are monitoring. System errors and events are recorded in an event log file and are displayed on the screen. Special device icons appear on the screen to notify you of disk failures and other events that require immediate attention.

5.1.3 Maintaining Storage Configurations

You can use MegaRAID Storage Manager to perform system maintenance tasks such as running consistency checks on arrays that support redundancy.

5.2 Hardware and Software Requirements

MegaRAID Storage Manager has the following hardware requirements:

- PC-compatible computer with an IA-32 (32-bit) Intel Architecture processor or an EM64T (64-bit) processor and at least 128 MB of system memory



NOTE LSI recommends that you use at least 1 GB of system memory.

-
- Physical drive with at least 50 MB available free space

MegaRAID Storage Manager supports these operating systems:

- Microsoft Windows Server 2003, 2008, and 2008R2
- Microsoft Windows Workstation 7
- Microsoft Windows Vista Workstation
- Red Hat Enterprise Linux (RHEL) 5 and 6
- SuSE Linux Enterprise Server (SLES) 10 and 11, with latest updates and service packs



NOTE MegaRAID Storage Manager supports the RHEL6 operating system (OS) with a few prerequisites. By nature, the RHEL6 OS default installation deploys only limited inbox libraries though it is shipped in the CD image for the full installation. This design from the RHEL6 OS adds an overhead for the Java® application programming users, and a few required libraries are not deployed during the default RHEL6 operating system installation. Because of this, users need to port the application in to the native 64-bit version.

Refer to your server documentation and to the operating system documentation for more information on hardware requirements and operating system requirements.

5.3 Installation

This section explains how to install (or reinstall) MegaRAID Storage Manager on your workstation or server.

5.3.1 Installing MegaRAID Storage Manager on Microsoft Windows

Follow these steps to install MegaRAID Storage Manager on a system running a Microsoft Windows operating system:

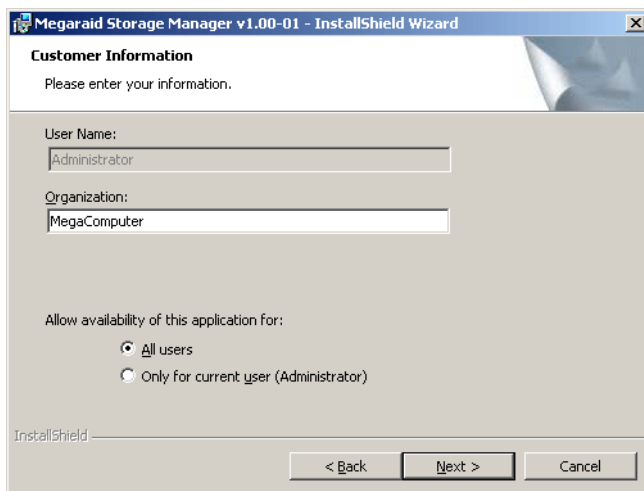
1. Unzip the zip file containing MegaRAID Storage Manager in Windows Explorer.
2. Double-click the `setup.exe` file to start the installation program.
3. When the Welcome screen appears, click **Next**.

If MegaRAID Storage Manager is already installed on this system, the Program Maintenance screen appears. Read the screen text and select **Modify, Repair, or Remove**.

4. When the next screen appears, read and accept the user license, and click **Next**.

The Customer Information screen appears, as shown in the following figure.

Figure 28 Customer Information Screen

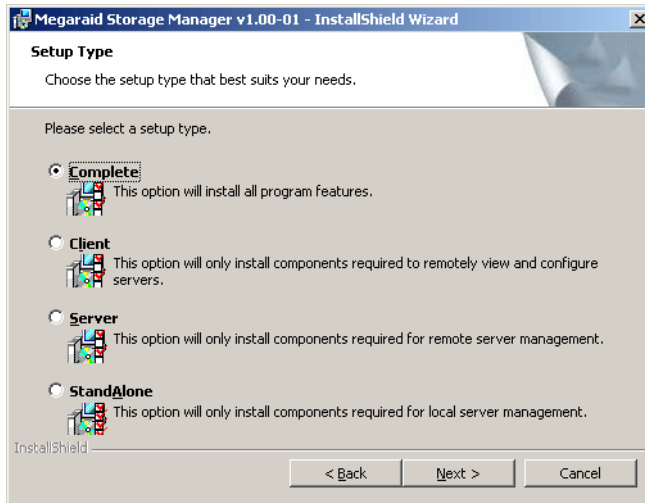


5. Enter your user name and organization name. In the bottom part of the screen, select an installation option:
 - If you select **All users**, any user with administrative privileges can use this version of MegaRAID Storage Manager to view or change storage configurations.
 - If you select **Only for current user**, the MegaRAID Storage Manager shortcuts and associated icons will be available only to the user with this user name.
6. Click **Next** to continue.

7. On the next screen, accept the default Destination Folder, or click **Change** to select a different destination folder. Click **Next** to continue.

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

Figure 29 Setup Type Screen



8. Select one of the Setup options. The screen text explains the options.
 - Normally, you would select **Complete** if you are installing MegaRAID Storage Manager on a server. This option installs the complete MSM.
 - Select **Client** if you are installing MegaRAID Storage Manager on a PC that will be used to view and configure servers over a network.

Master GUI, Monitor Configurator, Help files for both MSM and Monitor Configurator, and DebugLog are available. The following are not available: Popup, SNMP, Monitor, Framework, Storelib and Storelib-IR, StorelibJNI, and StorelibIRJNI.
 - Select **Server** to install only those components required for remote server management.

Popup, SNMP, Monitor, Framework, Storelib and Storelib-IR, StorelibJNI, and StorelibIRJNI are available. Master GUI, Monitor Configurator, Help files for both MSM and Monitor Configurator, and DebugLog are not available.
 - Select **StandAlone** if you will use MegaRAID Storage Manager to create and manage storage configurations on a standalone workstation.

All of the components are available except for the Network Capability Plugin.
9. Click **Next** to proceed.
10. Click **Install** to install the program.
11. Click **Finish** when the final Configuration Wizard screen appears.

If you select **Client** installation for a PC used to monitor servers, and if there are no available servers with a registered framework on the local subnet (that is, servers with a complete installation of MegaRAID Storage Manager), you cannot connect to a remote server unless you first edit the `startupui.bat` file. Specifically, you must add the IP address of the remote server to the end of the `startupui.bat` file.

For example, to connect to a remote framework on server 192.168.0.10, add the IP address to the end of `startupui.bat` as shown in this example:

```
start JRE\bin\javaw -classpath .;GUI.jar GUI.VivaldiStartupDialog  
ajsgyqkj=71244 192.168.0.10
```

Be sure to include a space in front of the IP address, as shown in the example.

5.3.2 Installing MegaRAID Storage Manager for Linux

Follow these steps to install MegaRAID Storage Manager on a system running Red Hat Linux or SUSE Linux:

1. Copy the `MSM_linux_installer...tar.gz` file to a temporary folder.
2. Untar the `MSM_linux_installer...tar.gz` file using the following command:

```
tar -zxvf MSM_linux_installer...tar.gz
```

A new disk directory is created.

3. Go to the new `disk` directory.
4. In the `disk` directory, find and read the `readme.txt` file.
5. To start the installation, enter the following command:

```
./install.sh
```

If you select **Client** installation for a PC used to monitor servers, and if there are no available servers with a registered framework on the local subnet (that is, servers with a complete installation of MegaRAID Storage Manager), you cannot connect to a remote server unless you first edit the `startupui.sh` file. Specifically, you must add the IP address of the remote server to the end of the `startupui.sh` file.

For example, to connect to a remote framework on server 192.168.0.10, add the IP address to `startupui.sh` as shown in this example:

```
start JRE\bin\javaw -classpath .;GUI.jar GUI.VivaldiStartupDialog  
ajsgyqkj=71244 192.168.0.10
```

Be sure to include a space in front of the IP address, as shown in the example.

5.3.2.1 Linux Installation Messages

One or more of the following messages can appear while you are installing MegaRAID Storage Manager on a Linux system:

- More than one copy of MegaRAID Storage Manager has been installed.
This message indicates that the user has installed more than one copy of MegaRAID Storage Manager. (This can be done by using the `rpm-force` command to install the `rpm` file directly, which is not recommended, instead of using the `install.sh` file.) In such cases, the user must uninstall all the `rpm` files manually before installing MegaRAID Storage Manager with the procedure listed previously.
- The version is already installed.
This message indicates that the version of MegaRAID Storage Manager you are trying to install is already installed on the system.
- The installed version is newer.
This message indicates that a version of MegaRAID Storage Manager is already installed on the system, and it is a newer version than the version you are trying to install.
- Exiting installation.
This message appears when the installation is complete.
- RPM installation failed.
This message indicates that the installation failed for some reason. Additional message text explains the cause of the failure.

Chapter 6: MegaRAID Storage Manager Screen and Menus

This chapter explains how to start LSI Software RAID Configuration Utility and describes the LSI Software RAID Configuration Utility main menu screen and menus.

6.1 Starting MegaRAID Storage Manager

Follow these steps to start MegaRAID Storage Manager and view the main menu screen:

1. Start the program using the method required for your operating system environment:
 - To start MegaRAID Storage Manager on a Microsoft Windows system, select **Start > Programs > MegaRAID Storage Manager > StartupUI**, or double-click the MegaRAID Storage Manager shortcut on the desktop.

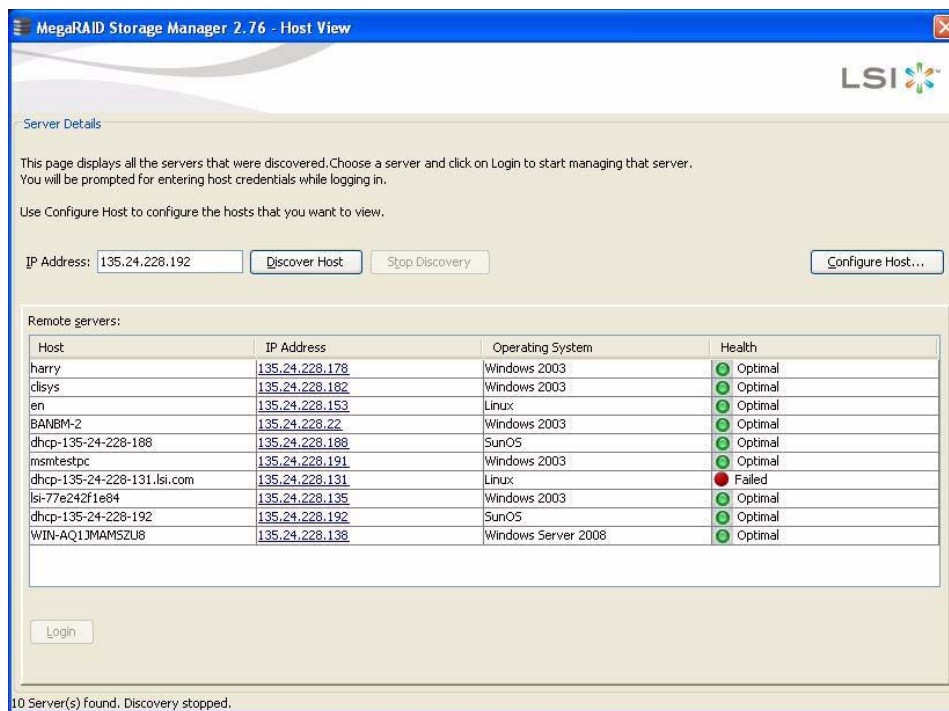


NOTE If a warning appears stating that Windows Firewall has blocked some features of the program, click **Unblock** to allow MegaRAID Storage Manager to start. (The Windows Firewall sometimes blocks the operation of programs that use Java.)

- To start MegaRAID Storage Manager on a Red Hat Linux system, select **Applications > System Tools > MegaRAID Storage Manager StartupUI**.
- To start MegaRAID Storage Manager on a SUSE Linux/SLES system, select **Start > System > More Programs > MegaRAID Storage Manager**.

When the program starts, the Select Server window appears, as shown in the following figure. The remote servers display, along with their IP address, operating system, and health status.

Figure 30 Select Server Window



If the circle in the server icon is orange instead of green, it means that the server is running in a degraded state, for example, because a drive used in a virtual drive has failed. If the circle is red, the storage configuration in the server has failed.



NOTE To access servers on a different subnet, type in the box at the bottom of the screen the IP address of a server in the desired subnet where MegaRAID Storage Manager is running, and click **Update**. If you check the **Connect to remote server at: IP address** box, you can also access a standalone (remote) installation of MegaRAID Storage Manager, if it has a network connection.

2. Double-click the icon of the server that you want to access.
The Server Login window appears, as shown in the following figure.

Figure 31 Server Login Window



3. Enter your user name and password.
The question mark icon opens a dialog box that explains what you need for full access to the server and for view-only access to the server.
4. Select an access mode from the drop-down menu for **Login Mode**, and click **Login**.
 - Select **Full Access** if you need to both view and change the current configuration.
 - Select **View Only** if you need to only view and monitor the current configuration.



NOTE If the computer is connected to a network, this procedure is the login to the computer itself, not the network login.

5. Enter the root/administrator user name and password to use the Full Access mode.



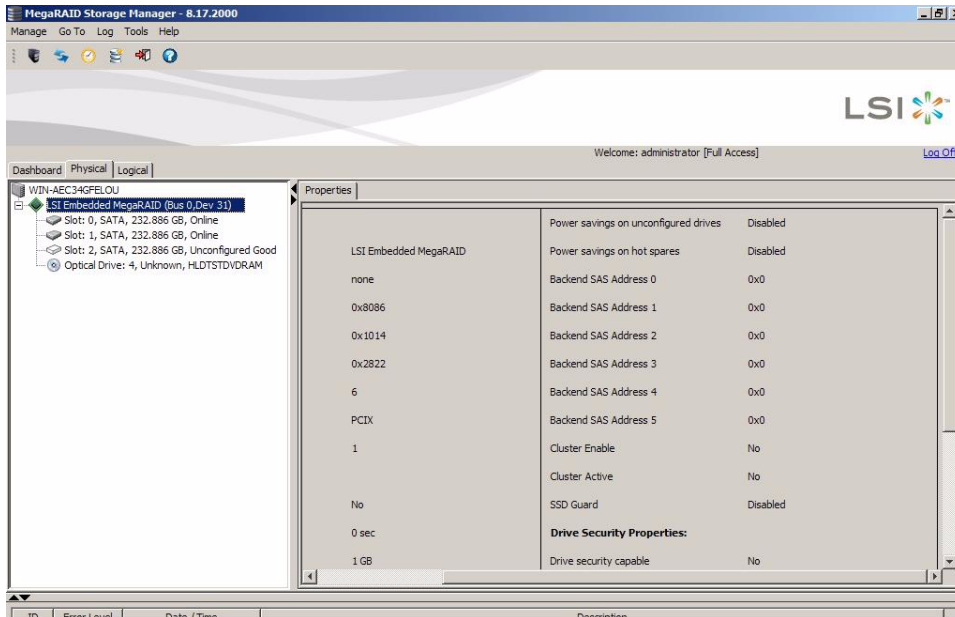
NOTE In Linux, users belonging to the root group can log in. You do not have to be the user "root".

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

6.2 MegaRAID Storage Manager Main Menu Screen

This section describes the MegaRAID Storage Manager main menu screen, which is shown in the following figure.

Figure 32 MegaRAID Storage Manager Main Menu Screen



The following sections describe the panels and the menu options that appear on this screen.

6.2.1 Dashboard/Physical View/Logical Views

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

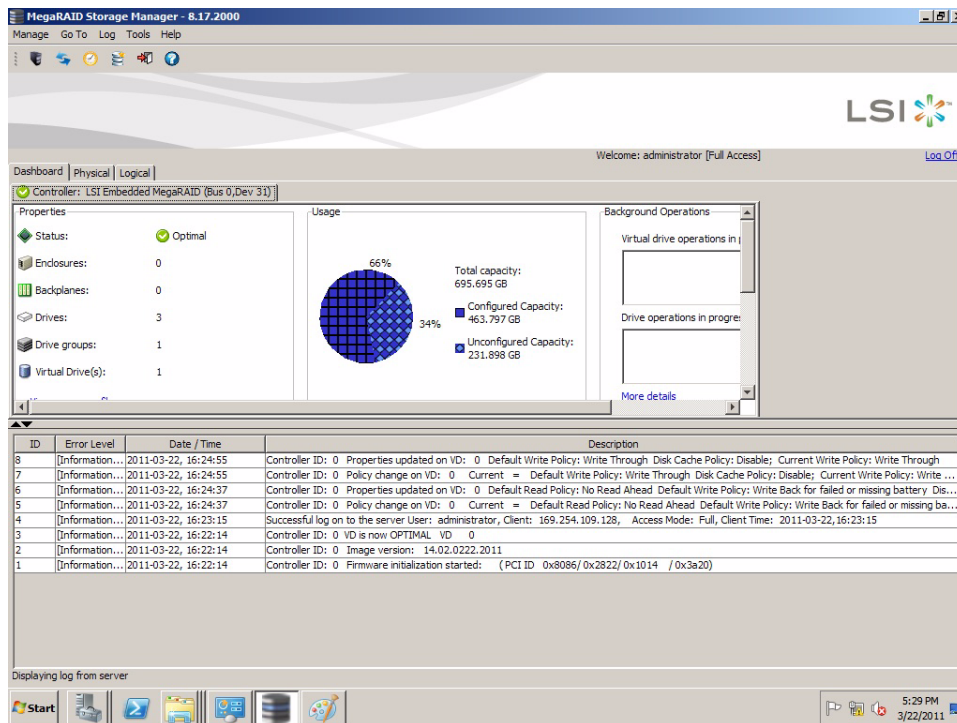
6.2.1.1 Dashboard View

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

- Properties of the virtual drives and the physical drives
- Total capacity, configured capacity, and unconfigured capacity
- Background operations in progress
- MSM features and their status (enabled or disabled)
- Actions you can perform
- Links to Online Help

The following figure shows the Dashboard view.

Figure 33 MSM Dashboard View



6.2.1.2 Physical View

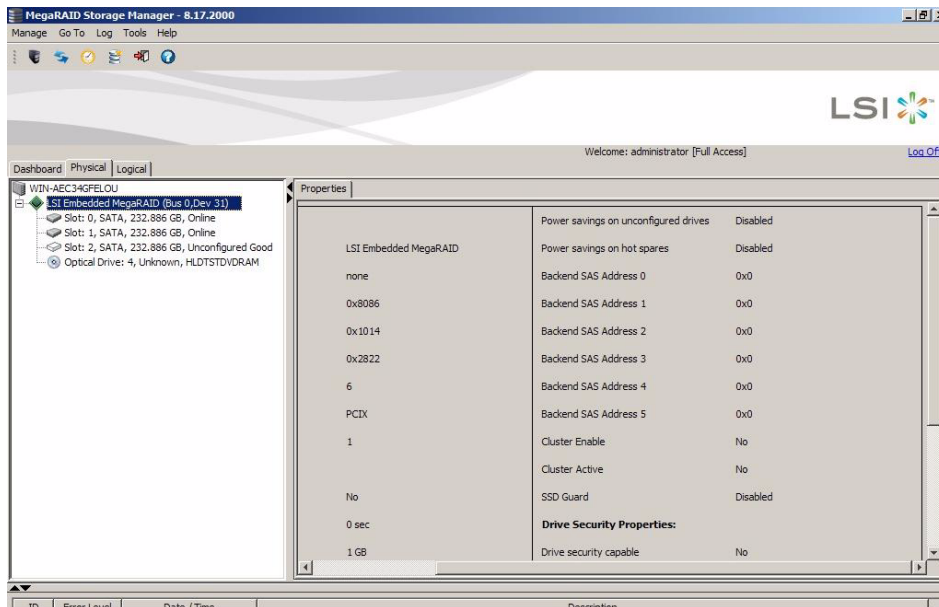
The *Physical* view shows the hierarchy of physical devices in the system. At the top of the hierarchy is the system itself, followed by the controller and the backplane.

One or more controllers are installed in the system. The controller label identifies the 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 device appear in the right panel of the screen under the Properties tab.

The following figure shows the Physical view and the Properties tab.

Figure 34 MSM Physical View

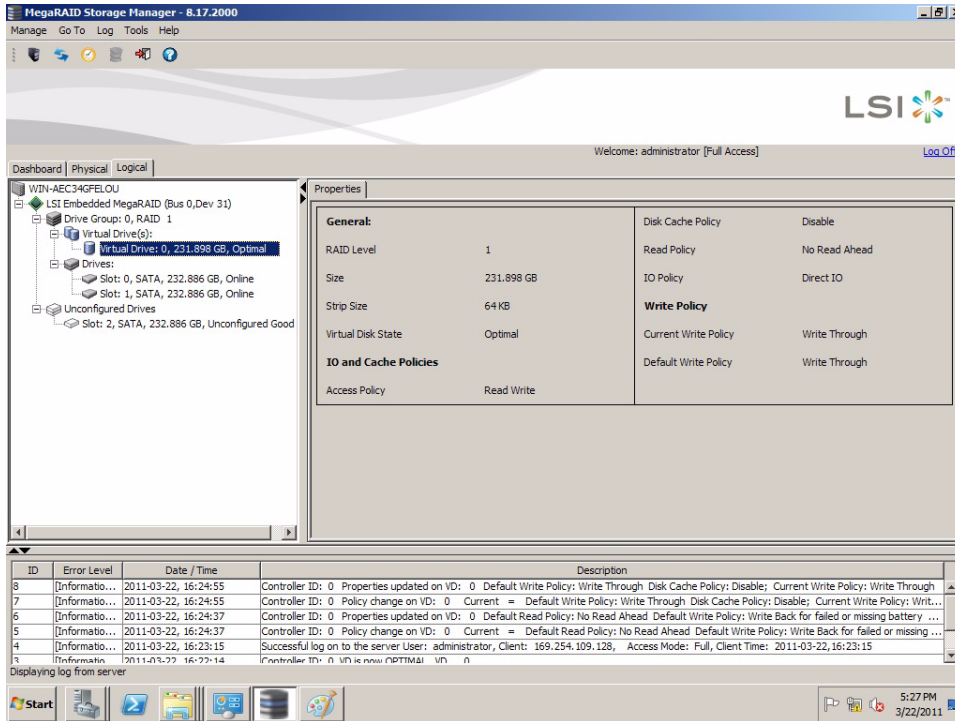


6.2.1.3 Logical View

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

The following figure shows the Logical view.

Figure 35 MSM Logical View



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

Table 46 Icons and Descriptions

Icon	Description
	Status
	System
	Controller
	Backplane
	Port
	Drive group
	Virtual drive
	Online drive
	Global hot spare
	Tape drive
	CD-ROM
	Foreign Drive
	Unconfigured Foreign Drive
	Unconfigured Drive



NOTE MegaRAID Storage Manager shows the icons for tape drive devices, but the utility does not support tape-related operations. If these operations are required, use a separate backup application.

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

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

6.2.2 Event Log Panel

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

For more information about the event log, see Section 8.1, [Monitoring System Events](#), on page 93. For more information about the event log entries, see Appendix A, [Events and Messages](#), on page 115.

6.2.3 Menu Bar

Here are brief descriptions of the main selections on the MegaRAID Storage Manager menu bar. Specific menu options are described in more detail in [Chapter 8, Monitoring System Events](#), and [Chapter 6, Maintaining and Managing Storage Configurations](#), of this manual.

6.2.3.1 Manage Menu

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

6.2.3.2 Go To Menu

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

Configuration options are also available, in which you access the Configuration Wizard 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**.

6.2.3.3 Log Menu

The Log menu includes options for saving and clearing the message log. For more information about the Log menu, see [Appendix A, Events and Messages](#), on page 115.

6.2.3.4 Tools Menu

On the Tools menu you can select **Tools > Configure Alerts** to access the Configure Alerts screen, which you can use to set the alert delivery rules, event severity levels, exceptions, and email settings. For more information, see [Section 8.2, Configuring Alert Notifications](#), on page 95.

6.2.3.5 Help Menu

On the Help menu you can select **Help > Contents** to view the MegaRAID Storage Manager online help file. You can select **Help > About MegaRAID Storage Manager** to view version information for MegaRAID Storage Manager.



NOTE When you use the MegaRAID Storage Manager Online Help, you might see a warning message that Internet Explorer has restricted the file from showing active content. If this warning appears, click on the active content warning bar and enable the active content.



NOTE If you are using the Linux operating system, you must install the Firefox[®] browser or the Mozilla[®] browser for the MegaRAID Storage Manager Online Help to display.

Chapter 7: Configuration

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

The LSI RAID controllers support RAID 0, RAID 1, RAID 5, and RAID 10 storage configurations. The Configuration Wizard allows you to create new storage configurations and modify the configurations. To learn more about RAID and RAID levels, see Section 1.2, [RAID Overview](#), on page 9.



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

7.1 Creating a New Storage Configuration

You can use the MegaRAID Storage Manager to create new storage configurations on systems with LSI RAID 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 storage configurations and advanced storage configurations.

7.1.1 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. 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:* This option allows you to start writing data to the virtual drive immediately.
 - *Full Initialization:* A complete initialization is done on the new configuration. You cannot write data to the new virtual drive until the initialization is complete. This can take a long time if the drives are large.



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

- **Stripe size:** The setting is 64 KB only.
- **Read policy:** Specify the read policy for this virtual drive:
 - *No read ahead:* (the default) Disables the read ahead capability. 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. Read ahead capability 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:
 - *Write Through:* In this mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This option eliminates the risk of losing cached data in case of power failure.
 - *Always Write Back:* In this mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction.
 - **I/O policy:** The I/O policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.
 - *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.
 - *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.

7.1.2 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.



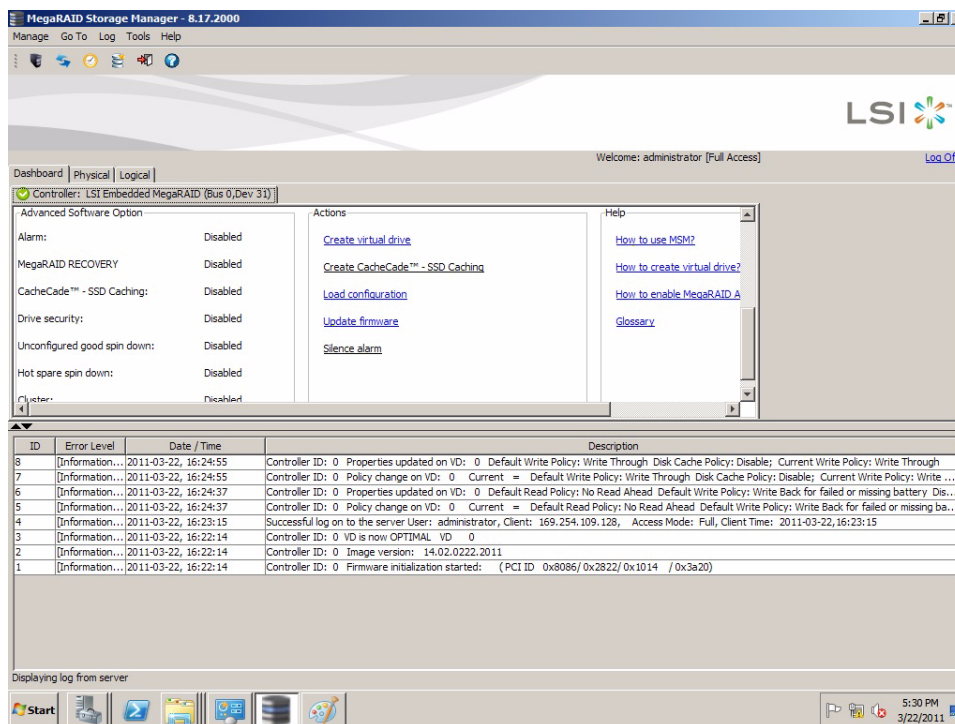
NOTE You cannot create RAID 10 spanned drives using the simple configuration procedure. To create RAID 10 spanned drives, use the advanced configuration procedure described in Section 7.1.3, [Creating a Virtual Drive Using Advanced Configuration](#), on page 85.

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

1. Click the **Dashboard** tab on the main menu screen.

The Create Virtual Drive option appears in the Actions section of the Dashboard, as shown in the following figure.

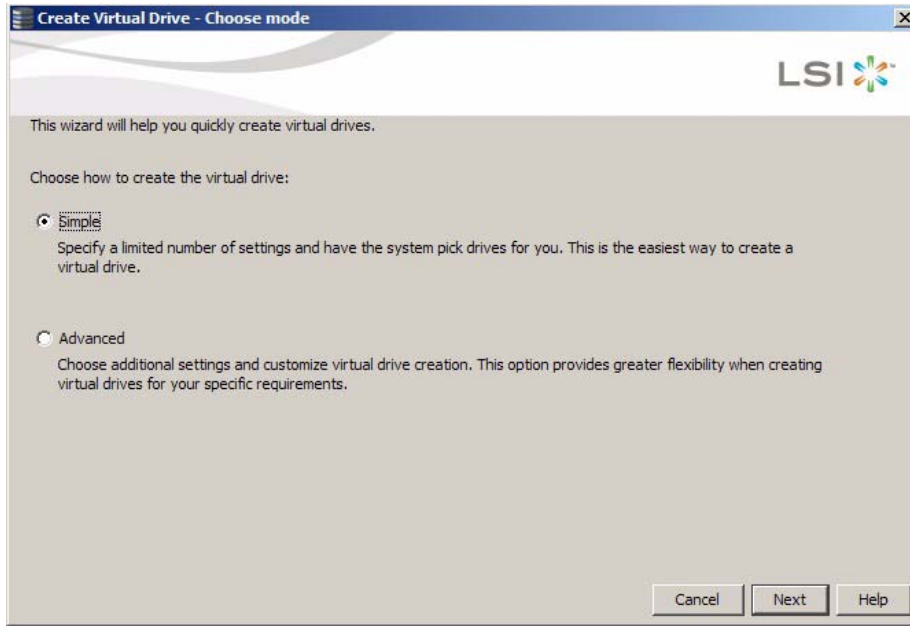
Figure 36 Virtual Drive Creation Menu



2. Click **Create Virtual Drive**.

The dialog box for the configuration mode (Simple or Advanced) appears, as shown in the following figure.

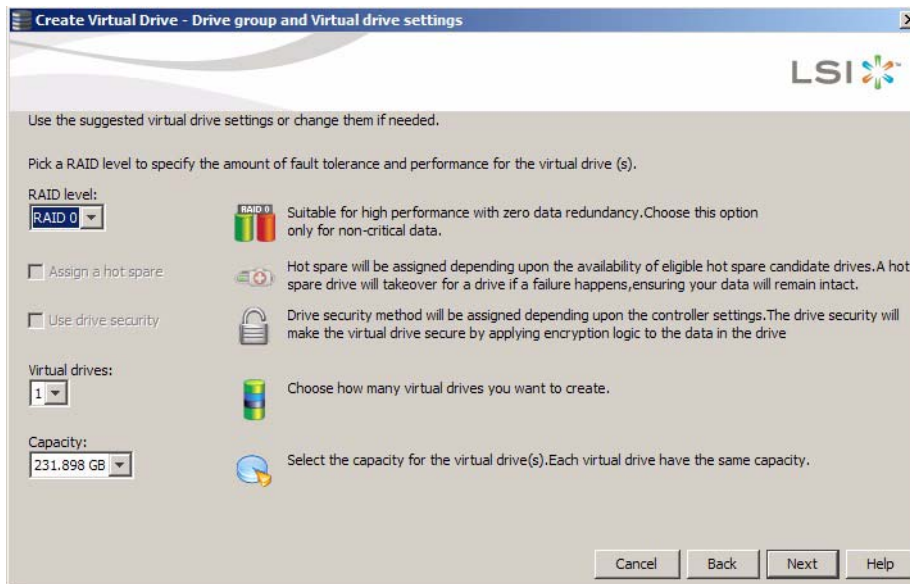
Figure 37 Virtual Drive Simple Configuration Mode



3. Click **Simple** and press **Next**.

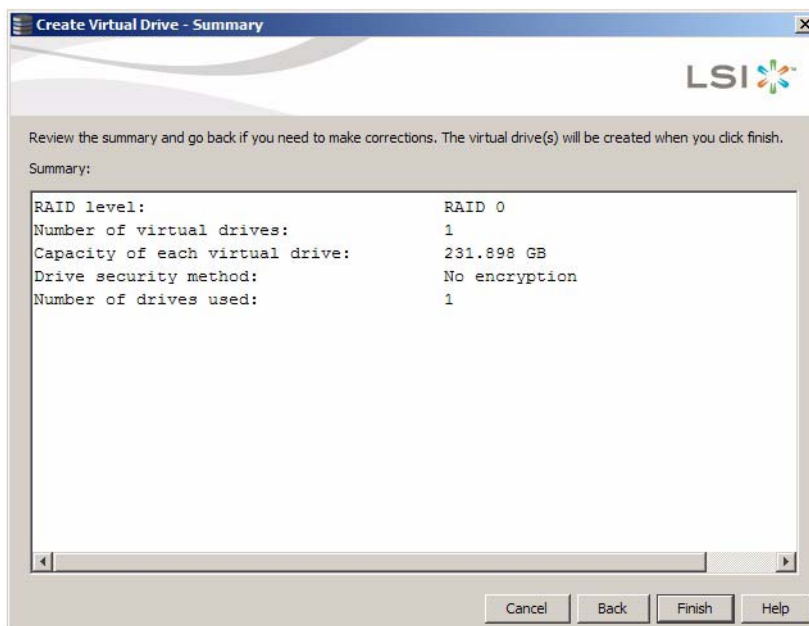
The Create Virtual Drive screen appears, as shown in the following figure.

Figure 38 Create Virtual Drive Screen



4. Select the RAID level desired for the virtual drive.
When you use simple configuration, the RAID controller supports RAID levels 1 and 5. In addition, it supports independent non-redundant drives (configured as RAID 0). The screen text gives a brief description of the RAID level that you select. The RAID levels that you can choose depend on the number of drives available. To learn more about RAID levels, see Section 1.2, [RAID Overview](#), on page 9.
5. Use the drop-down menu in the **Virtual drives** field to choose how many virtual drives you want to create.
6. Select the capacity of the virtual drive(s).
Each virtual drive has the same capacity.
7. 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 39 Create Virtual Drive - Summary Window



8. Click **Back** to return to the previous screen to change any selections or click **Finish** to accept and complete the configuration.
The new virtual drive is created after you click **Finish**. After the configuration is completed, a dialog box notifies you that the virtual drives were created successfully.

7.1.3 Creating a Virtual Drive Using Advanced Configuration

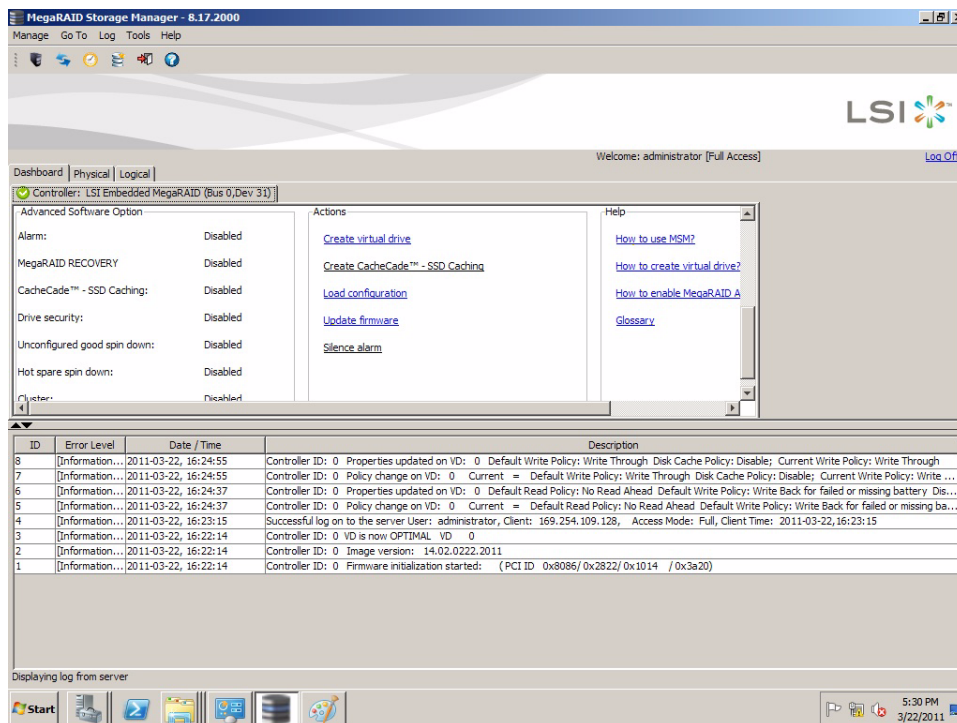
The advanced configuration procedure provides an easy way to create a new storage configuration. Advanced configuration gives you greater flexibility than simple configuration because you can select the drives and the virtual drive parameters when you create a virtual drive. In addition, you can use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

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

1. Click the **Dashboard** tab on the MSM main menu screen.

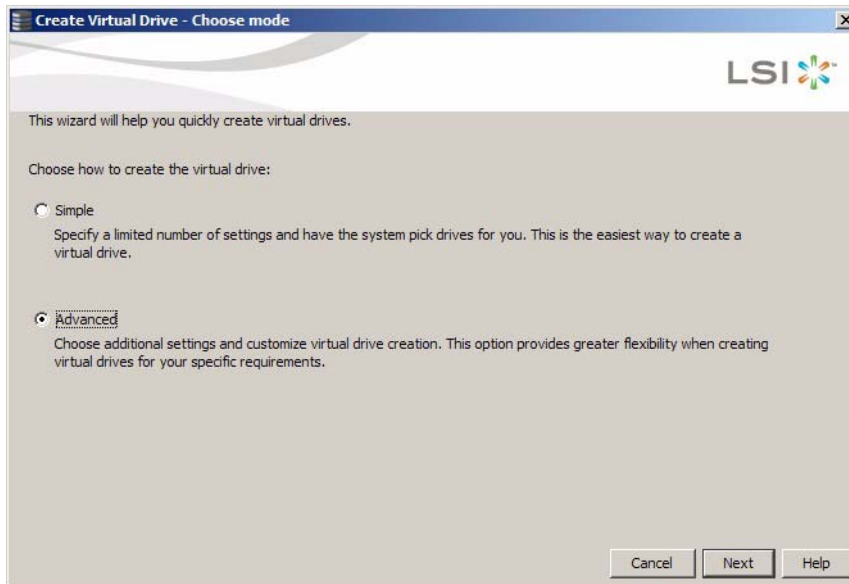
The Create Virtual Drive option appears in the Actions section of the Dashboard, as shown in the following figure.

Figure 40 Virtual Drive Creation Menu



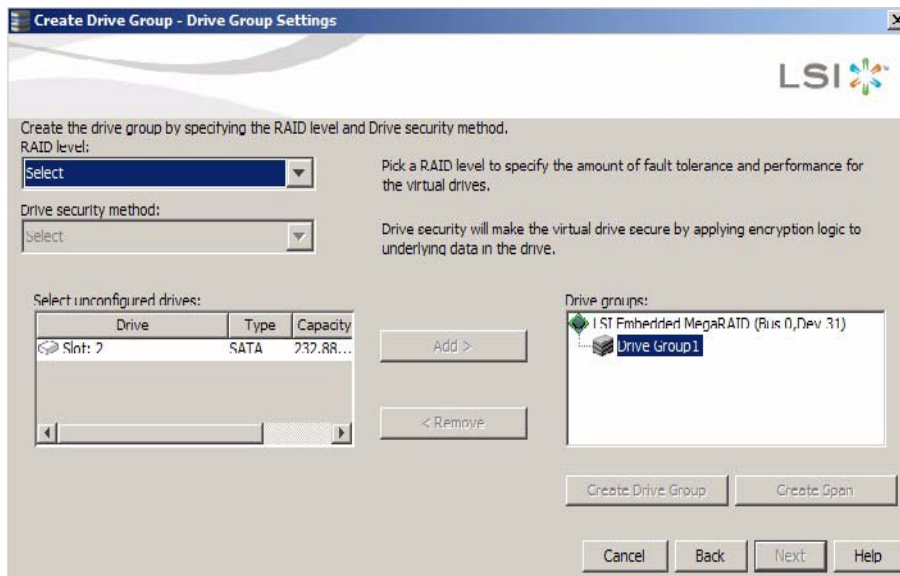
2. Click **Create Virtual Drive** in the Actions section of the screen.
The dialog box for the configuration mode (Simple or Advanced) appears, as shown in the following figure.

Figure 41 Virtual Drive Advanced Configuration Mode



3. Click **Advanced** and then click **Next**.
The Create Drive Group Settings dialog box appears, as shown in the following figure.

Figure 42 Create Drive Group Settings Dialog Box



4. Perform the following actions in the Create Drive Group Settings dialog box:
 - a. Select the RAID level desired for the drive group from the drop-down menu.
The drive group appears in the **Drive groups** field.
The RAID controller supports RAID levels 1, 5, and 10. In addition, it supports independent drives configured as RAID 0. The screen text gives a brief description of the RAID level you select. Which RAID levels you can choose depend on the number of drives available.
 - b. Select *unconfigured* drives from the list of drives and click **Add>** to add them to the drive group.
The selected drives appear under **Drive Group {number}**.
 - c. If you are creating a RAID 10 drive group, perform the following steps. If not, go to step **d**.

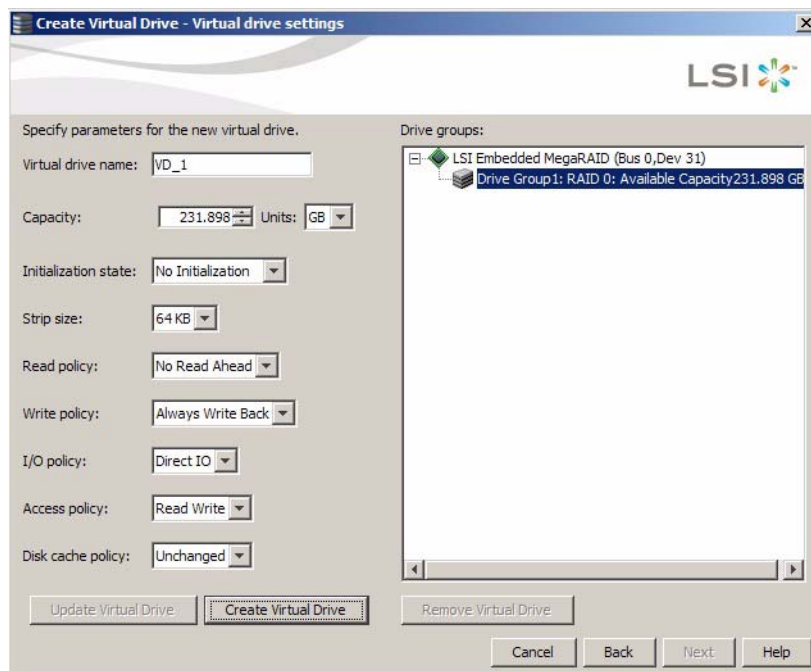


NOTE On a RAID 10 array, you can create only one virtual drive, and that virtual drive must occupy the entire space of the RAID 10 array.

- Click **Create Span** to create a second span in the drive group.
- Select *unconfigured* drives from the list of drives and click **Add>** to add them to the second drive group.
- The selected drives appear under the second span, **Span {number}**, below the second drive group, **Drive Group {number}**.
- Click **Create Drive Group** to make a drive group with the spans.
- Click **Next** to complete the steps for a RAID 10 configuration.
- d. Click **Create Drive Group** to make a drive group.
- e. Click **Next** to complete this step.

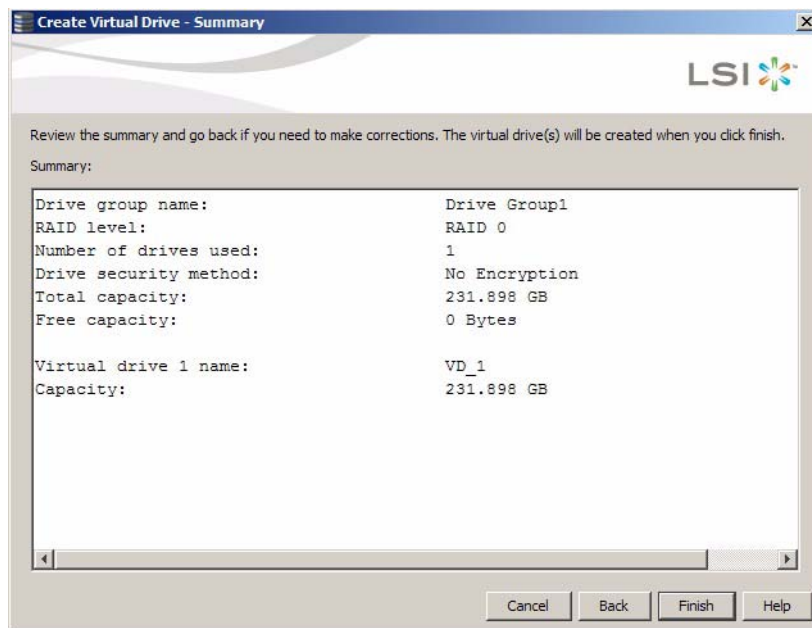
The Virtual drive settings window appears, as shown in the following figure. The drive group and the default virtual drive settings appear. The options to update the virtual drive or remove the virtual drive are grayed out until you create the virtual drive.

Figure 43 Virtual Drive Settings Dialog Box



5. Select the virtual drive settings to fit your environment.
See Section 7.1.1, [Selecting Virtual Drive Settings](#), on page 80, for more information about the virtual drive settings.
6. 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.
7. Click **Next**.
The Create Virtual Drive - Summary window appears, as shown in the following figure. This window shows the selections you made for the advanced configuration.

Figure 44 Virtual Drive Summary Window



8. Click **Back** to return to the previous screen to change any selections or click **Finish** to accept and complete the configuration.
After you click **Finish**, the new storage configuration is created and initialized.
After the configuration is completed, a dialog box notifies you that the virtual drives have been successfully created. If more drive capacity exists, the dialog box asks whether you want to create more virtual drives. If no more drive capacity exists, you are prompted to close the configuration session.
9. Select **Yes** or **No** to indicate whether you want to create additional virtual drives.
If you select **Yes**, the system takes you to the Create Virtual Drive screen, as shown in [Figure 38](#). If you select **No**, the utility asks whether you want to close the wizard.
10. If you selected **No** in the previous step, select **Yes** or **No** to indicate whether you want to close the wizard.
If you select **Yes**, the configuration procedure closes. If you select **No**, the dialog box closes and you remain on the same page.

7.2 Changing Adjustable Task Rates

Follow these steps if you need to change the adjustable rates for rebuilds, and other system tasks that run in the background:



NOTE Leave the adjustable task rates at their default settings to achieve the best system performance. If you raise the task rates above the defaults, foreground tasks will run more slowly and it might seem that the system is not responding. If you lower the task rates below the defaults, rebuilds and other background tasks might run very slowly and might not complete within a reasonable time. If you decide to change the values, record the original default value here so you can restore them later, if necessary:

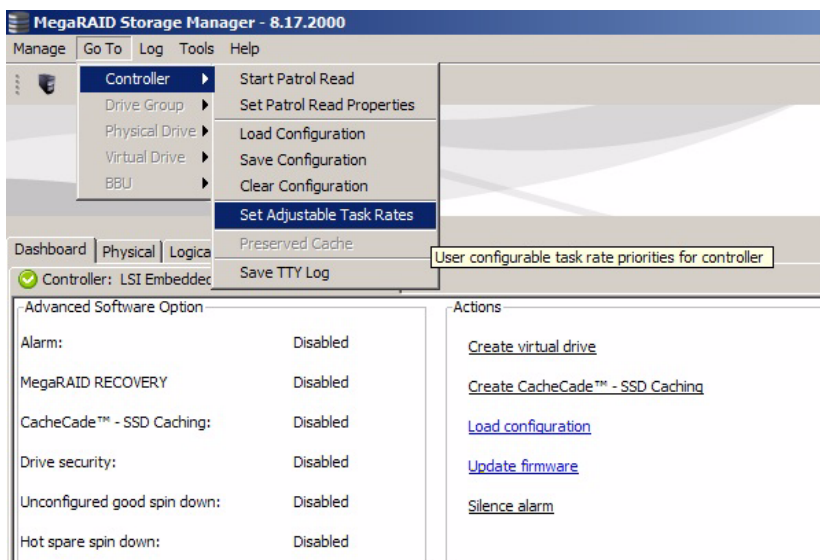
Rebuild Rate: _____

Background Initialization (BGI) Rate: _____

Check Consistency Rate: _____

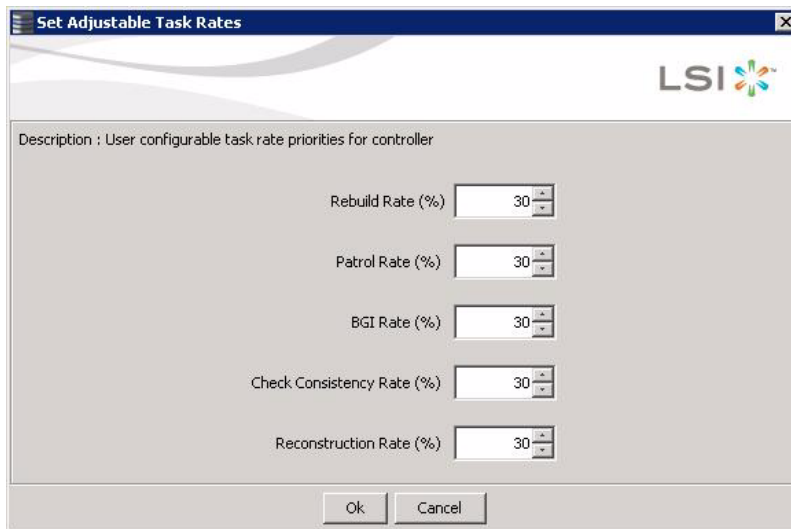
1. Select a controller icon in the **Physical** tab or the **Logical** tab in the left panel of the MegaRAID Storage Manager main menu screen.
2. Select **Go To > Controller > Set Adjustable Task Rates** from the menu bar, as shown in the following figure.

Figure 45 Set Adjustable Task Rates Menu



The Set Adjustable Task Rates dialog box appears, as shown in the following figure.

Figure 46 Set Adjustable Task Rates Dialog Box



3. Enter changes, as needed, to the following task rates:

- Rebuild Rate
- Patrol Read
- Background Initialization (BGI) (for fast initialization)
- Check Consistency (for consistency checks).
- Reconstruction

Each task rate can be set from 0 to 100 percent. The higher the number, the faster the activity runs in the background, possibly impacting other system tasks.

4. Click **OK** to accept the new task rates.
5. When the warning message appears, click **OK** to confirm that you want to change the task rates.

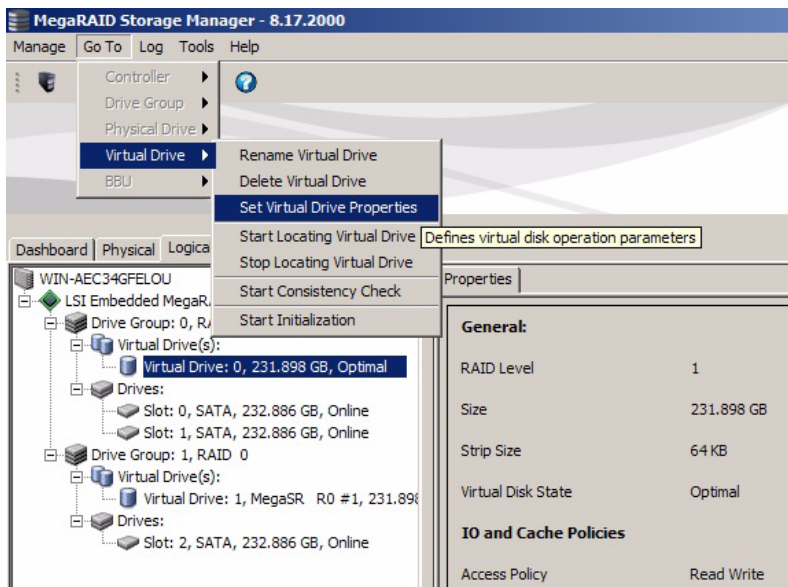
7.3 Changing Virtual Drive Properties

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

1. Select a virtual drive icon in the **Physical** tab or the **Logical** tab in the left panel of the MegaRAID Storage Manager main menu screen.

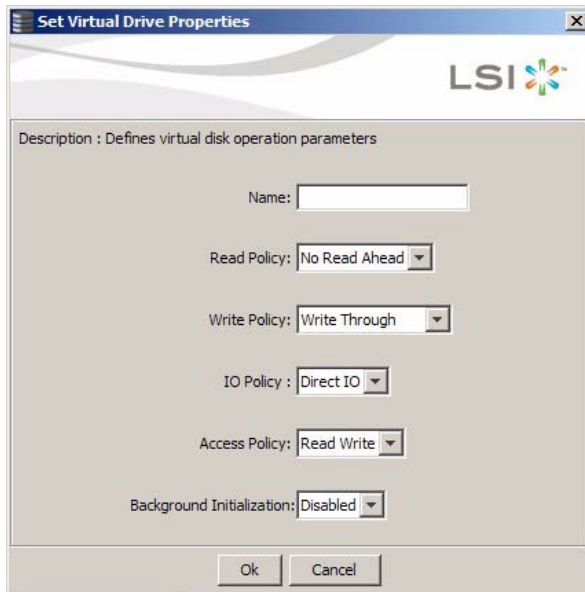
2. Select **Go To > Virtual Drive > Set Virtual Drive Properties** from the menu bar, as shown in the following figure.

Figure 47 Set Virtual Drive Properties Menu



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

Figure 48 Set Virtual Drive Properties Screen



3. Change the virtual drive properties as needed.
For information about these properties, see Section 7.1.1, [Selecting Virtual Drive Settings](#), on page 80.
4. Click **Ok** to accept the changes.
The virtual drive settings are updated.

7.4 Deleting a Virtual Drive



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

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

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



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

Chapter 8: Monitoring System Events and Storage Devices

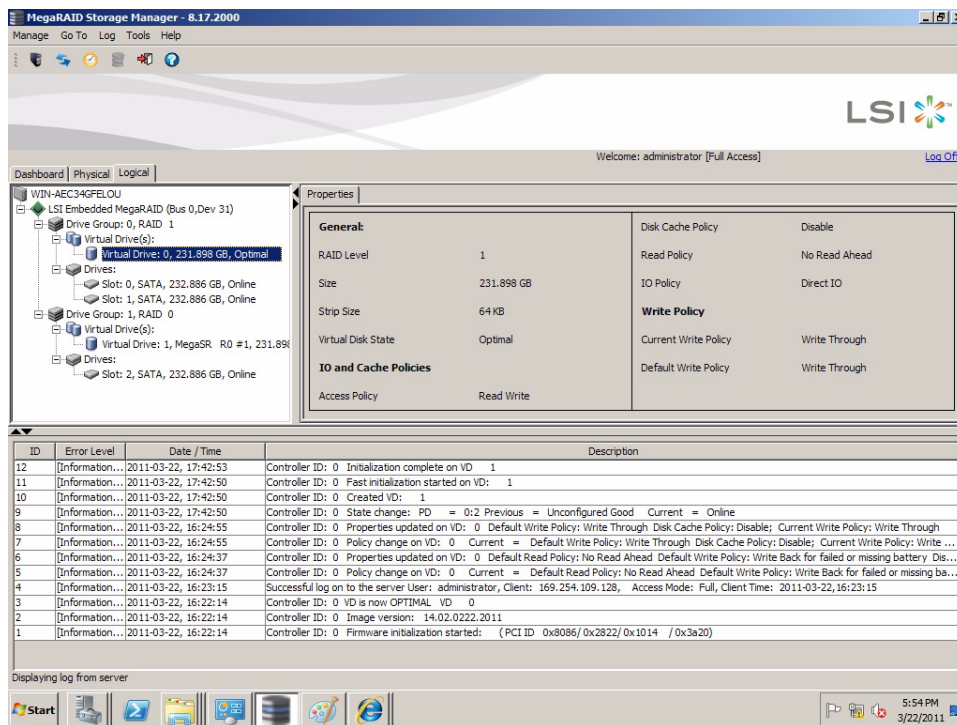
This chapter explains how to use MegaRAID Storage Manager to monitor the status of drives, virtual drives, and other storage devices.

8.1 Monitoring System Events

The MegaRAID Storage Manager utility monitors the activity and performance of all controllers in the system and the storage devices connected to them. When an event occurs (such as the creation of a new virtual drive or the removal of a drive) an event message appears in the log at the bottom of the MegaRAID Storage Manager main menu screen, as shown in the following figure.

You can use MegaRAID Storage Manager to alert you about events. There are settings for the delivery of alerts, the severity level of events, exceptions, and email settings.

Figure 49 Event Information Window



Each message that appears in the event log has a severity level that indicates the importance of the event, as shown in [Table 47](#), a date and timestamp, and a brief description. You can click an event to display the same information in a window. (For a list of all events, see [Appendix A, Events and Messages](#), on page 115).

Table 47 Event Severity Levels

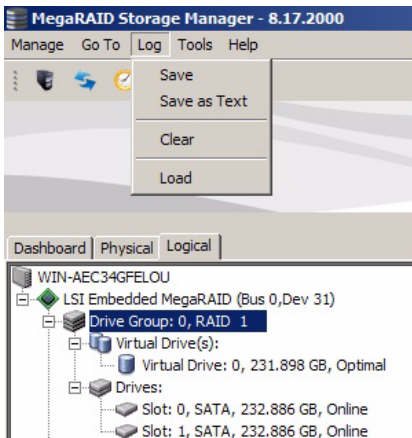
Severity Level	Meaning
Information	Informational message. No user action is necessary.
Warning	Some component might be close to a failure point.
Critical	A component has failed, but the system has not lost data.
Fatal	A component has failed, and data loss has occurred or will occur.

The Log menu has four options:

- **Save** – Saves the current log to a .log file.
- **Save as Text** – Saves the current log in .txt format.
- **Clear** – Clears the current log information. You have the option of saving the log first.
- **Load** – Enables you to load a local .log file.

The following figure shows the log menu.

Figure 50 Log Menu



8.2 Configuring Alert Notifications

The Alert Notification Configuration feature allows you to control and configure the alerts that MegaRAID Storage Manager sends when various system events occur.

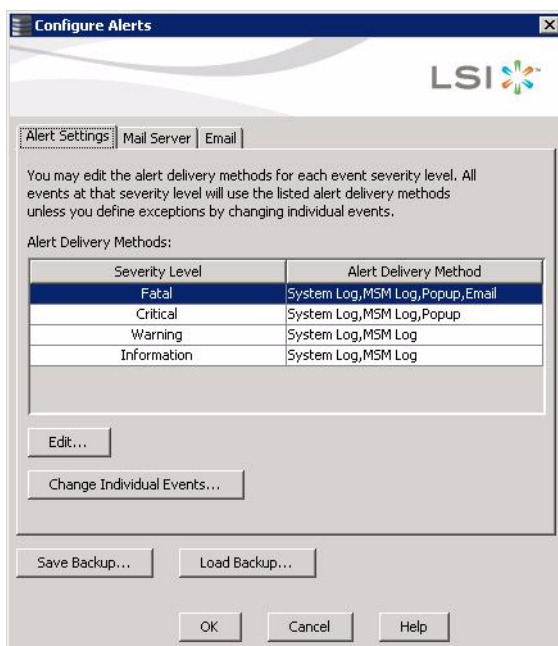
To access this screen, select **Tools > Configure Alerts** on the main menu screen, as shown in the following figure.

Figure 51 Alert Notification Configuration Menu



The Alerts Notification Configuration screen appears, as shown in the following figure. The screen contains three tabs: **Alert Settings**, **Mail Server**, and **Email**. You can use each tab to perform tasks for that topic.

Figure 52 Alerts Notification Configuration Screen



You can select the **Alert Settings** tab to perform the following actions:

- Select the methods for the delivery of alerts.
- Change the severity level of events.
- Save an .xml backup file of the entire alert configuration.

- Load all of the values from a previously saved backup into the dialog to edit or send to the monitor.



NOTE When you load a saved backup file, all unsaved changes made in the current session are lost.

Select the **Mail Server** tab to perform the following actions:

- Enter or edit the sender email address.
- Enter the name of the SMTP server.
- Require authentication of the email server.
- Save an `.xml` backup file of the entire alert configuration.
- Load all of the values from a previously saved backup into the dialog to edit or send to the monitor.



NOTE When you load a saved backup file, all unsaved changes made in the current session will be lost.

Select the **Email** tab to perform the following actions:

- Add new email addresses for recipients of alert notifications.
- Send test messages to the recipient email addresses.
- Remove email addresses of recipients of alert notifications.
- Save an `.xml` backup file of the entire alert configuration.
- Load all of the values from a previously saved backup into the dialog to edit or send to the monitor.



NOTE When you load a saved backup file, all unsaved changes made in the current session will be lost.

8.2.1 Setting Alert Delivery Methods


You can select the methods used to send alert deliveries, including by popup, email, system log, or MSM log. You can select the alert delivery methods for each event severity level (Information, Warning, Critical and Fatal).

Perform the following steps to select the alert delivery methods:

1. On the Alerts Notification Configuration screen, click the **Alert Settings** tab.
2. Under the **Alerts Delivery Methods** heading, select one of the severity levels.
3. Click **Edit**.

The Alert Notification Delivery Methods dialog box appears, as shown in the following figure.

Figure 53 Alert Notification Delivery Methods Dialog Box



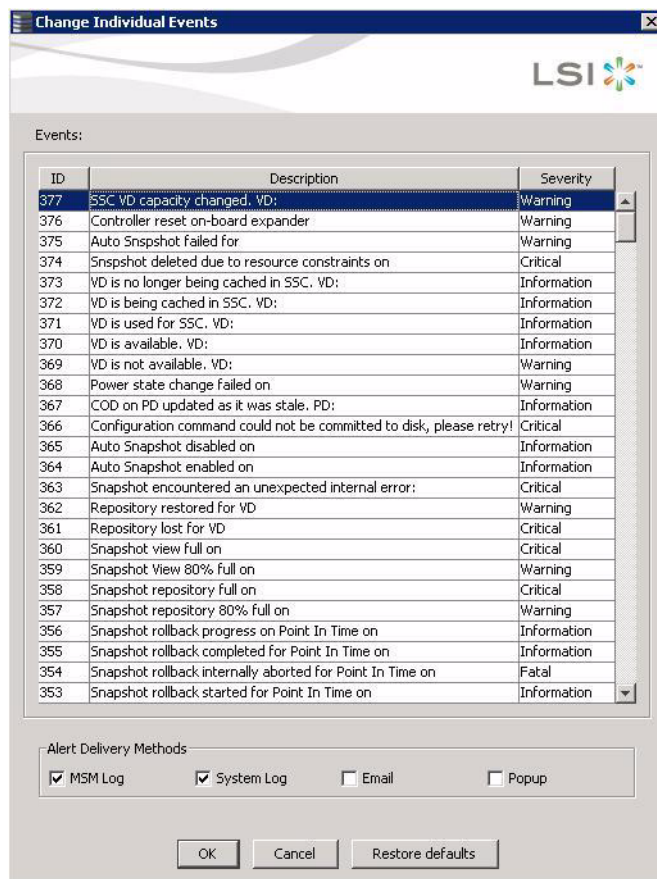
4. Select the desired alert delivery methods for alert notifications at the event severity level.
5. Click **OK** to set the delivery methods used for the severity level that you selected.

8.2.2 Changing Alert Delivery Methods for Individual Events

You can change the alert delivery options for an event without changing the severity level.

1. On the Alerts Notification Configuration screen, click the **Alert Settings** tab.
The **Alerts Setting** portion of the screen appears, as shown in [Figure 52](#).
2. Click **Change Individual Events**.
The Change Individual Events dialog box appears, as shown in the following figure.
The dialog box shows the events by their ID number, description, and severity level.

Figure 54 Change Individual Events Dialog Box



3. Click an event in the list to select it.
The current alert delivery methods appear for the selected event under the **Alert Delivery Methods** heading.
4. Select the desired alert delivery methods for the event.
5. Press **Esc** to return to the Alerts Notification Configuration screen.
6. Click **OK**.
This saves all of the changes made to the event.

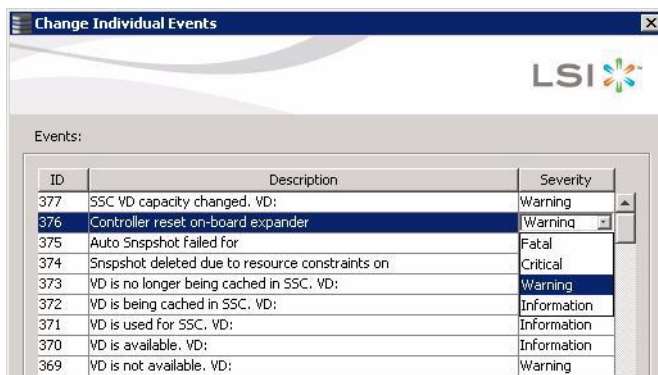
8.2.3 Changing the Severity Level for Individual Events

See [Table 47](#) for details about the severity levels.

To change the event severity level for a specific event, perform the following steps.

1. On the Alerts Notification Configuration screen, click the **Alert Settings** tab.
The **Alerts Setting** portion of the screen appears.
2. Click **Change Individual Events**.
The Change Individual Events dialog box appears, as shown in [Figure 54](#). The dialog box shows the events by their ID number, description, and severity level.
3. Click an event in the list to select it.
The current alert delivery methods appear for the selected event.
4. Click the **Severity** cell for the event.
The Event Severity drop-down menu appears for that event, as shown in the following figure.

Figure 55 Change Individual Events Severity Level Menu

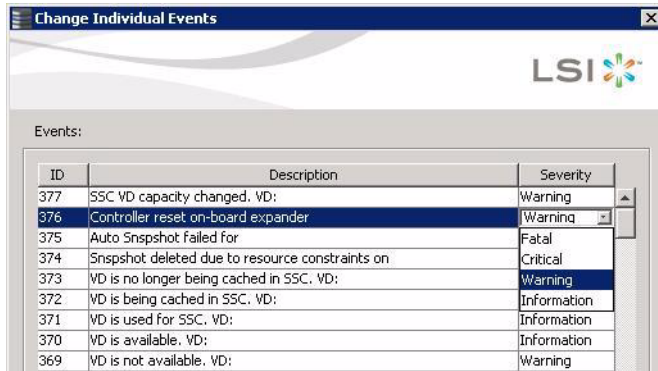


5. Select a different severity level for the event from the menu.
6. Press Esc to return to the Alerts Notification Configuration screen.
7. Click **OK** to save all of the changes made to the events.

8.2.4 Multiple Events Displayed in a Single Pop-Up Window

You can view multiple events in a single pop-up window, as shown in the following figure.

Figure 56 Pop-Up for Multiple Events



8.2.5 Entering or Editing the Sender Email Address and SMTP Server

You can use the Alerts Notification Configuration screen to enter or edit the sender e-mail address and the SMTP server.

1. On the Alerts Notification Configuration screen, click the **Mail Server** tab.
The **Mail Server** options appear, as shown in the following figure.

Figure 57 Mail Server Options



2. Enter a new sender email address in the **Sender email address** field or edit the existing sender email address.
3. Click **OK**.

8.2.6 Authenticating a Server

You can use the Alerts Notification Configuration screen to authenticate the SMTP server, providing an extra level of security. The authentication check box enables the **User name** and **Password** fields when selected by default. Clearing the check box disables these fields.

Perform the following steps to enter or edit the address:

1. On the Alerts Notification Configuration screen, click the **Mail Server** tab.
The **Mail Server** options appears, as shown in [Figure 57](#). The authentication check box is selected by default.
2. Enter a user name in the **User name** field.
3. Enter the password in the **Password** field.
4. Click **OK**.

8.2.7 Saving Backup Configurations

You can save an `.xml` backup file of the entire alert configuration. This includes all the settings on the three tabs.

1. On the Alerts Notification Configuration screen, click the **Alert Settings** tab, **Mail Server** tab, or **Email** tab.
2. Click **Save Backup**.
The drive directory appears.
3. Enter a file name with an `.xml` extension for the backup configuration (in the format `filename.xml`).
4. Click **Save**.
The drive directory disappears.
5. Click **OK**.
The backup configuration is saved, and the Alert Notification Configuration screen closes.

8.2.8 Loading Backup Configurations

You can load all of the values from a previously saved backup into the dialog (all tabs) to edit or send to the monitor.



CAUTION If you choose to load a backup configuration and the Configure Alerts dialog currently contains changes that have not yet been sent to the monitor, the changes will be lost. You are prompted to confirm your choice.

-
1. On the Alerts Notification Configuration screen, click the **Alert Settings** tab, **Mail Server** tab, or **Email** tab.
 2. Click **Load Backup**.
A message warns that when you load a saved backup file, all unsaved changes made in the current session will be lost.
 3. Click **Yes**.
The drive directory appears, from which you can select a backup configuration to load.
 4. Select the backup configuration file (it should be in `.xml` format).
 5. Click **Open**.
The drive directory disappears.
 6. Click **OK**.
The backup configuration is loaded and the Alerts Notification Configuration screen closes.

8.2.9 Adding Email Addresses of Recipients of Alert Notifications

The **Email** tab portion of the Alerts Notification Configuration screen shows the email addresses of recipients of the alert notifications. MegaRAID Storage Manager sends alert notifications to those email addresses. Use the screen to add or remove email addresses of recipients, and to send test messages to recipients that you add.

To add email addresses of recipients of the alert notifications, perform the following steps:

1. Click the **Email** tab on the Event Notification Configuration screen.
The **Email** section of the screen appears, as shown in the following figure.

Figure 58 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.

8.2.10 Testing Email Addresses of Recipients of Alert Notifications

Use the **Email** tab portion of the Alerts Notification Configuration screen to send test messages to the email addresses that you added for the recipients of alert notifications.

1. Click the **Email** tab on the Event Notification Configuration screen.
The **Email** section of the screen appears, as shown in the following figure.
2. Click an email address in the **Recipient email addresses** field.
3. Click **Test**.
4. Confirm whether the test message was sent to the email address.
If MegaRAID Storage Manager cannot send an email message to the email address, an error message appears.

8.2.11 Removing Email Addresses of Recipients of Alert Notifications


Use the **Email** tab portion of the Alerts Notification Configuration screen to remove email addresses of the recipients of alert notifications.

1. Click the **Email** tab on the Event Notification Configuration screen.
The **Email** section of the screen appears, as shown in [Figure 58](#).
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.

8.3 Monitoring Controllers

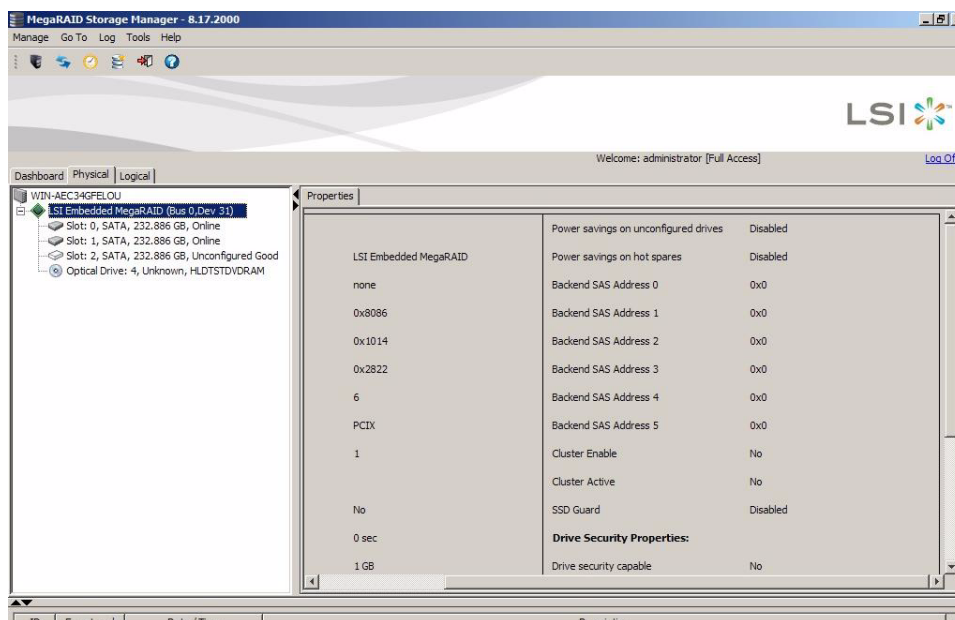


NOTE The Embedded MegaRAID Software drivers act as virtual “controllers.” Because these are not actual hardware storage controllers installed in the computer system, some of the controller properties shown in the following screen do not apply to them.

When MegaRAID Storage Manager is running, you can see the status of all controllers in the left panel of the MegaRAID Storage Manager main menu screen. If the controller is operating normally, the controller icon looks like this: . If the controller has failed, a small red circle appears to the right of the icon. (See [Section 6.2.1, Dashboard/Physical View/Logical Views](#), on page 75 for a complete list of device icons.)

To display complete controller information, click a controller icon in the left panel of the MegaRAID Storage Manager main menu screen. The controller properties display in the right panel, as shown in the following figure.


Figure 59 Controller Properties




Most of the information on this screen is self-explanatory. Note that the *Rebuild Rate*, *Patrol Read Rate*, *Reconstruction Rate*, *Consistency Check Rate*, and *BGI Rate* (background initialization) are all user selectable. For more information about these rates, see Section 7.2, [Changing Adjustable Task Rates](#), on page 89.

The controller properties are defined in Appendix A, [Events and Messages](#), on page 115.

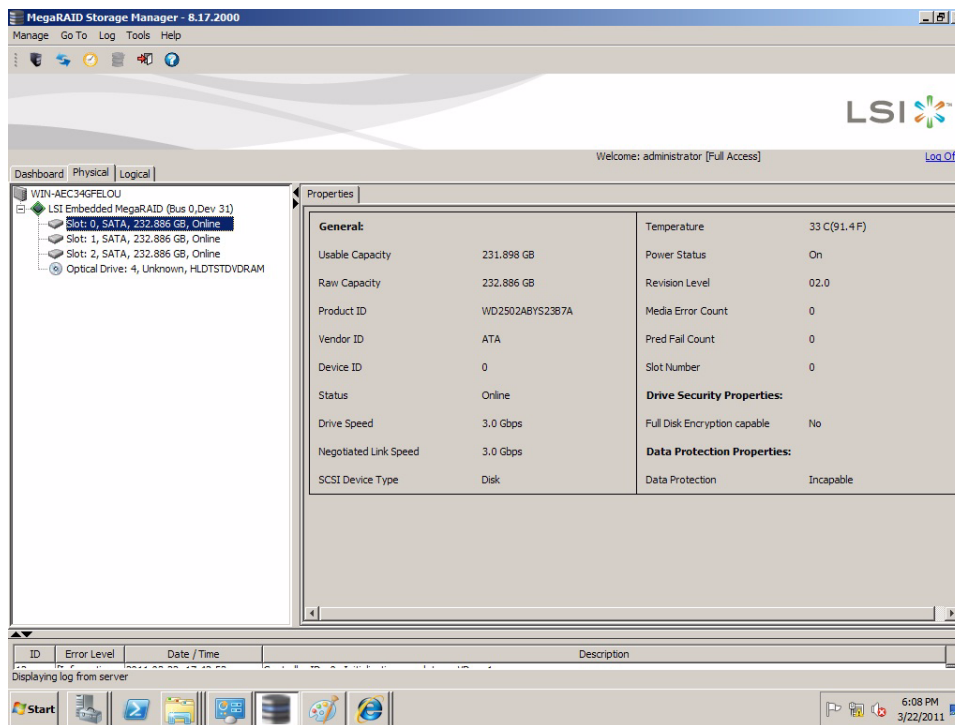
8.4 Monitoring Drives

When MegaRAID Storage Manager is running, you can see the status of all drives in the left panel of the MegaRAID Storage Manager main menu screen. If the drive is operating normally, its icon looks like this: .

If the drive has failed, a small red circle appears to the right of the icon, like this: . (See Section 6.2.1, [Dashboard/Physical View/Logical Views](#), on page 75 for a complete list of device icons.)

To display complete drive information, click a drive icon in the left panel of the MegaRAID Storage Manager main menu screen. The drive properties appear in the right panel, as shown in the following figure.

Figure 60 Drive Information



The information on this panel is self-explanatory. There are no user-selectable properties for physical devices. Icons for other storage devices such as CD-ROM drives and DAT drives can also appear in the left panel.

If the drives are in a drive enclosure, you can identify which drive is represented by each drive LED on the enclosure. Follow these steps to locate the drive:

1. Click the drive icon in the left panel.
2. Click **Go To > Physical Drive > Start Locating Drive**.

The LED on the drive in the enclosure starts blinking to show its location.



NOTE LEDs on drives that are global hot spares do not blink.

3. To stop the drive LED on the enclosure from blinking, select **Go To > Physical Drive > Stop Locating Drive**.

To display a graphical view of a drive, click a drive icon in the left panel of the MegaRAID Storage Manager main menu screen, and click the **Graphical View** tab. In Graphical View, the drive's storage capacity is color coded according to the legend shown on the screen:

- Configured space is blue
- Available space is white
- Reserved space is red.

When you select a virtual drive from the drop-down menu, the drive space used by that virtual drive appears in green.

8.5 Running a Patrol Read

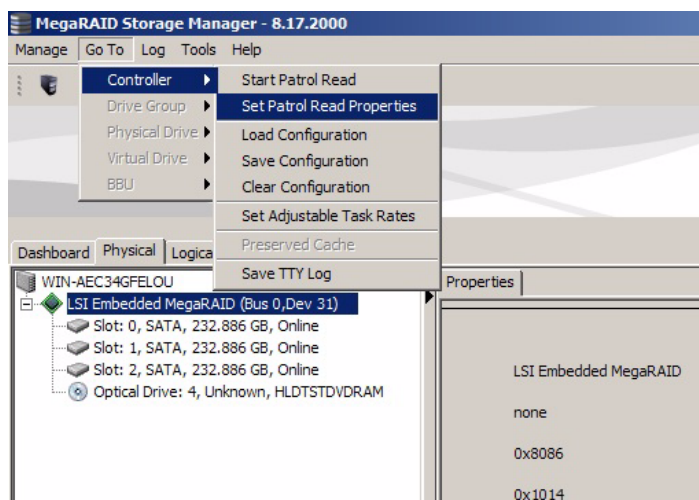
A patrol read periodically verifies all sectors of drives connected to a controller, including the system reserved area in the RAID configured drives. A patrol read can be used for all RAID levels and for all hot spare drives. This operation is initiated only when the controller is idle for a defined time period and has no other background activities.

You can set the patrol read properties and start the patrol read operation, or you can start the patrol read without changing the properties.

To set the patrol read properties and then start a patrol read, follow these steps:

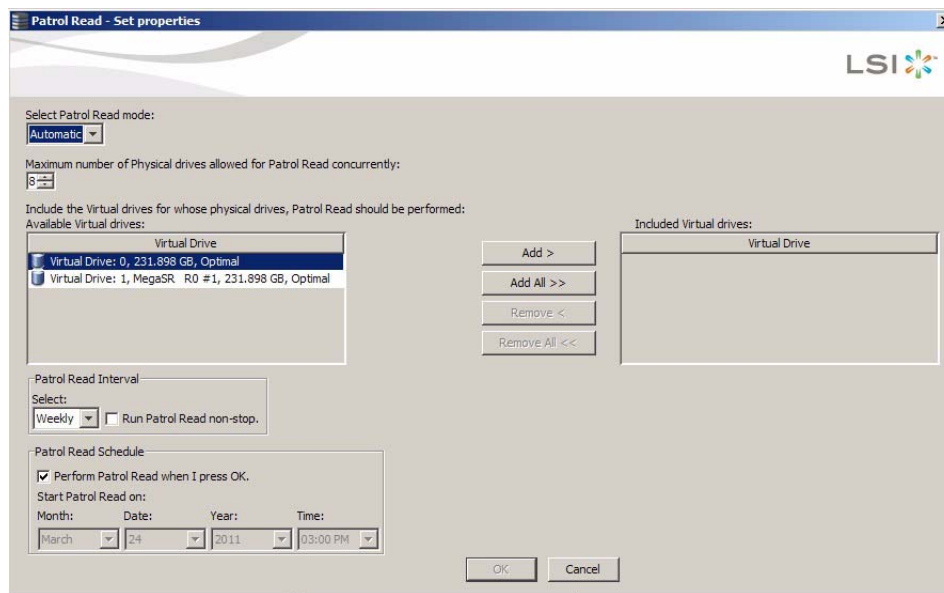
1. Click a controller icon in the left panel of the MegaRAID Storage Manager main menu screen.
2. Select **Go To > Controller > Set Patrol Read Properties**.

Figure 61 Start Patrol Read Menu



The Patrol Read - Set properties screen displays, as shown in the following figure.

Figure 62 Patrol Read Configuration



3. Select a mode for a patrol read. The options are:
 - **Automatic:** Patrol read runs automatically at the time interval you specify on this screen.
 - **Manual:** Patrol read runs only when you manually start it by selecting **Start Patrol Read** from the controller Options panel.
 - **Disabled:** Patrol read does not run.
4. Specify the maximum number of drives to include in the patrol read.
Eight drives is the maximum supported by Embedded MegaRAID Software.
5. Click virtual drives in the list under the heading **Virtual Drives** to include in the patrol read and click **Add >** or click **Add All >>** to include all of the virtual drives.
6. (Optional) Change the frequency at which the patrol read will run.
The default frequency is weekly (168 hours), which is suitable for most configurations. The other options are hourly, daily, and monthly.



NOTE LSI recommends that you leave the patrol read frequency and other patrol read settings at the default values to achieve the best system performance. If you decide to change the values, record the original default value here so you can restore them later, if necessary:

Patrol Read Frequency: _____

Continuous Patrolling: Enabled/Disabled

Patrol Read Task Rate: _____

7. (Optional) Set Patrol Read to run at a specific time.
The default is for the patrol read to start when you click **OK** on this screen. To change the default so that the patrol read starts at a specific time, follow these steps (otherwise, skip this step and proceed to the next step):
 - a. Uncheck the box **Perform Patrol Read when I click OK.**
 - b. Select the month, year, day, and time to start patrol read.

8. Click **OK** to enable your patrol read selections.



NOTE Patrol read does not report on its progress while it is running. The patrol read status is reported in the event log only.

To start a patrol read without changing the patrol read properties, follow these steps:



1. Click a controller icon in the left panel of the MegaRAID Storage Manager main menu screen.
2. Select **Go To > Controller > Start Patrol Read** in the menu bar.
3. When prompted, click **Yes** to confirm that you want to start a patrol read.

8.5.1 Patrol Read Task Rates

You can change the patrol read *task rate*. The task rate determines the amount of system resources that are dedicated to a patrol read when it is running. LSI recommends, however, that you leave the patrol read task rate at its default setting.

If you raise the task rate above the default, foreground tasks will run more slowly and it might seem that the system is not responding. If you lower the task rate below the default, rebuilds and other background tasks might run very slowly and might not complete within a reasonable time. For more information, about the patrol read task rate, see [Section 7.2, Changing Adjustable Task Rates, on page 89](#).

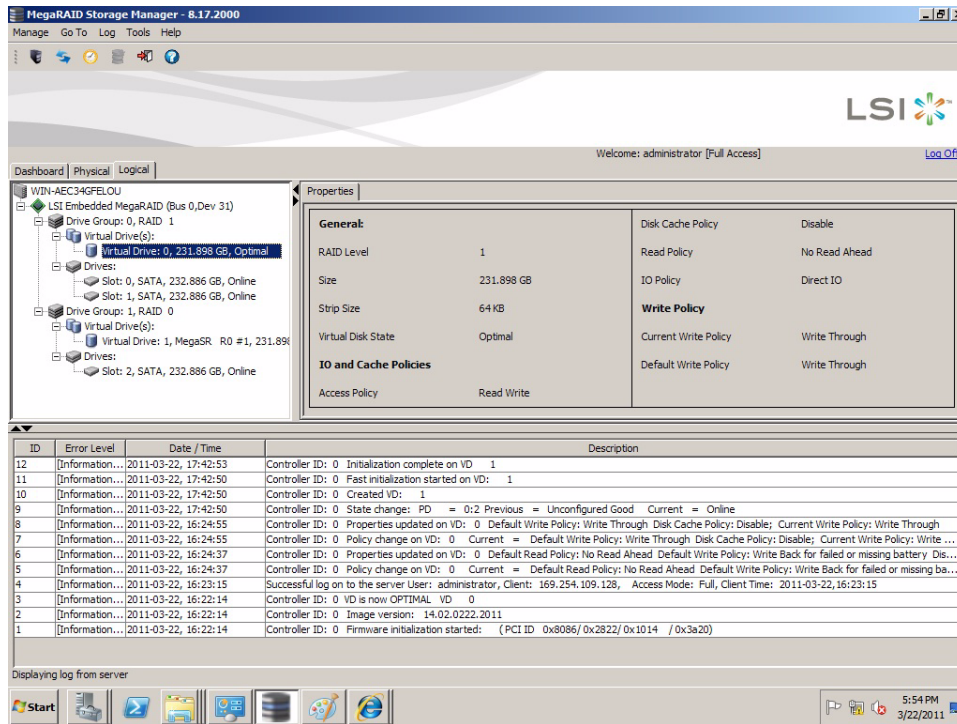
8.6 Monitoring Virtual Drives

When MegaRAID Storage Manager is running, you can see the status of all virtual drives. If a virtual drive is operating normally, the icon looks like this: . If the virtual drive is running in Degraded mode (for example, if a drive has failed), a small yellow circle appears to the right of the icon: . A red circle indicates that the virtual drive has failed and data has been lost.

When the Logical tab is selected, the left panel of the MegaRAID Storage Manager main menu screen shows which drives are used by each virtual drive. The same drive can be used by multiple virtual drives.

To display complete virtual drive information, click the **Logical** tab in the left panel and click a virtual drive icon in the left panel. The properties appear in the right panel. The following figure shows the Properties panel for a virtual drive.

Figure 63 Virtual Drive Properties



The RAID level, stripe size, and access policy of the virtual drive are set when the virtual drive is configured.




NOTE You can change the read policy, write policy, and other virtual drive properties. See Section 7.3, [Changing Adjustable Task Rates](#), on page 90, for the procedure you can use to change these properties.

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

1. Click the virtual drive icon in the left panel.
2. Click **Go To > Virtual Drive > Start Locating Virtual Drive** or right-click a virtual drive and select **Start Locating Virtual Drive** from the menu.
The LEDs on the drives in the virtual drive start blinking (except for the LEDs for hot spare drives).
3. To stop the LEDs from blinking, click **Go To > Virtual Drive > Stop Locating Virtual Drive**.

8.7 Monitoring Enclosures

When MegaRAID Storage Manager is running, you can see the status of all enclosures connected to the server by selecting the **Physical** tab in the left panel. If an enclosure is operating normally, the icon looks like this: . If the enclosure is not functioning normally—for example, if a fan has failed—a small yellow or red circle appears to the right of the icon.

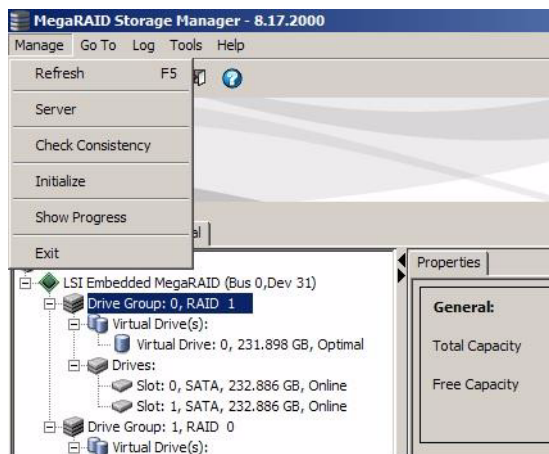
8.8 Monitoring Rebuilds and Other Processes

MegaRAID Storage Manager allows you to monitor the progress of rebuilds and other lengthy operations in the Group Show Progress window.

Follow these steps to monitor the progress of these operations.

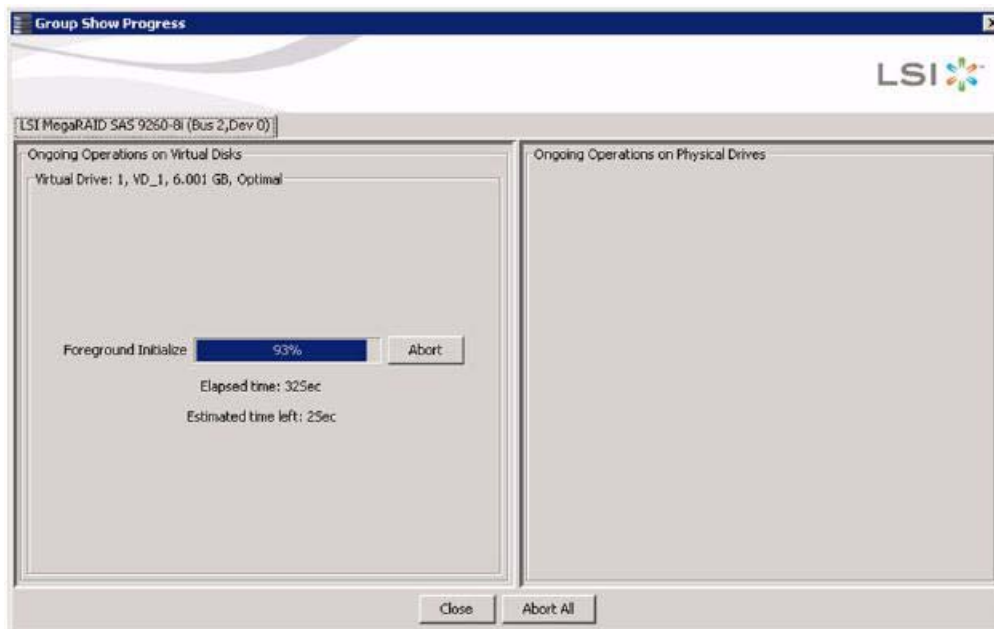
1. Select **Manage > Show Progress** on the menu bar, as shown in the following figure.

Figure 64 Group Show Progress Menu



The Group Show Progress window appears, as shown in the following screen.

Figure 65 Group Show Progress Window



Operations on virtual drives appear in the left panel of the Group Show Progress window, and operations on drives appear in the right panel. The following operations appear in this window:

- Background or foreground initialization of a virtual drive (see Section 9.1, [Initializing a Virtual Drive](#), on page 110)
 - Rebuild (see Section 9.4, [Rebuilding a Drive](#), on page 113)
 - Check Consistency (see Section 9.2, [Running a Consistency Check](#), on page 112)
2. (Optional) Click **Abort All** to abort all ongoing processes.
 3. Click **Close** to close the window.

Chapter 9: Maintaining and Managing Storage Configurations

This chapter explains how to use Embedded MegaRAID Software to maintain and manage storage configurations.

9.1 Initializing a Virtual Drive

To initialize a virtual drive after completing the configuration process, follow these steps:

1. Select the **Logical** tab in the left panel of the MegaRAID Storage Manager main menu screen, and click the icon of the virtual drive that you want to initialize.
2. Select **Go To > Virtual Drive > Start Initialization**.
The initialize dialog box appears.
3. Select the virtual drive(s) to initialize.



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

4. Select the **Fast Initialization** check box if you want to use this option.
If you leave the box unchecked, MegaRAID Storage Manager runs a Full Initialization on the virtual drive. For more information, see Section 9.1.1, [Running a Group Initialization](#), on page 111.
5. Click **Start** to begin the initialization.
You can monitor the progress of the initialization. For more information, see Section 8.8, [Monitoring Rebuilds and Other Processes](#), on page 108.

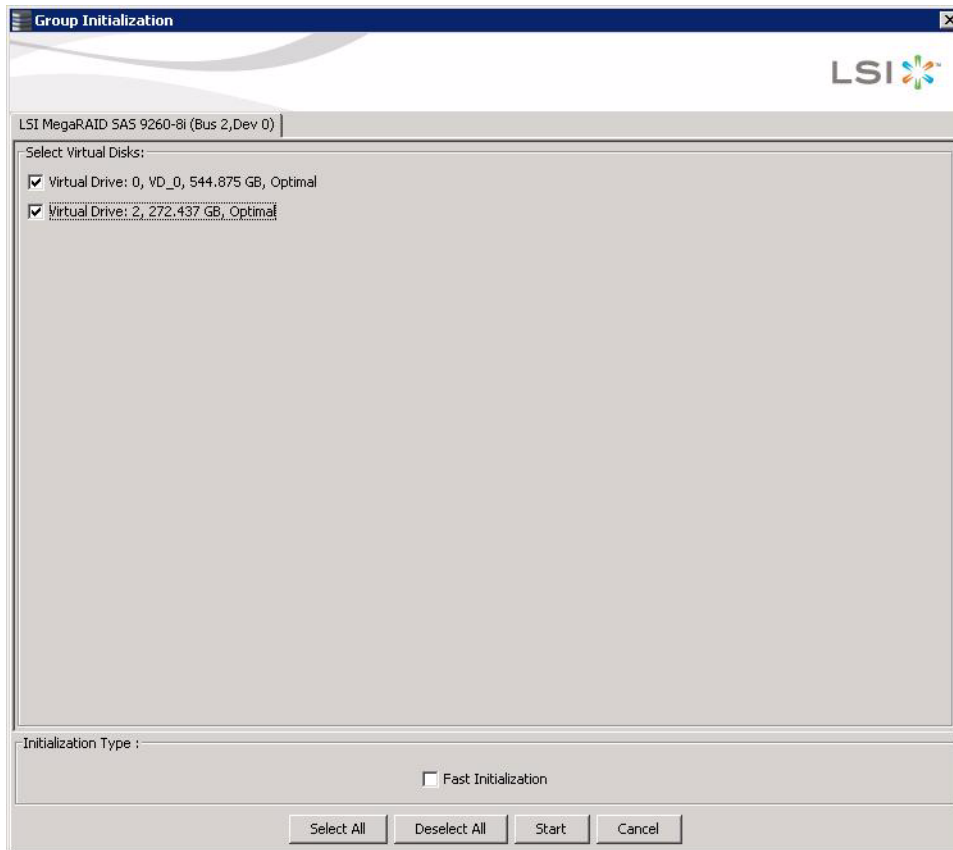
9.1.1 Running a Group Initialization

Initialization prepares the storage medium for use. You can run an initialization on multiple drives at one time. Follow these steps to run a group initialization.

1. Click **Manage > Initialize**.

The Group Initialization appears, as shown in the following figure.

Figure 66 Group Initialization Dialog Box



2. Check the virtual drives to run the initialization on or click **Select All** to select all of the virtual drives.
3. Click **Start**.
You can monitor the progress of the group initialization. See Section 8.8, [Monitoring Rebuilds and Other Processes](#), on page 108, for more information.

9.2 Running a Consistency Check

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

You should run a consistency check on fault-tolerant virtual drives periodically. You must run the consistency check if you suspect that the virtual drive data might be corrupted. Be sure to back up the data before running a consistency check if you think the data might be corrupted.

To run a consistency check, first set the consistency check properties and then schedule the consistency check. This section explains how to set the properties, schedule the check, and run the consistency check.


9.2.1 Running a Group Consistency Check

You can run a consistency check on multiple drives at one time. Follow these steps to run a group consistency check.

1. Click **Manage > Check Consistency**.

The Group Consistency Check appears, as shown in the following figure.

Figure 67 Group Consistency Check Dialog Box



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

You can monitor the progress of the group consistency check. See Section 8.8, [Monitoring Rebuilds and Other Processes](#), on page 108, for more information.

9.3 Scanning for New Drives

You can use the Scan for Foreign Configuration option to find drives with foreign configurations. A foreign configuration is a RAID configuration that already exists on a replacement set of physical disks that you install in a computer system. In addition, if one or more drives are removed from a configuration, by a cable pull or drive removal, for example, the configuration on those drives is considered a foreign configuration by the RAID controller. Drives that are foreign are listed on the physical drives list with a special symbol in MegaRAID Storage Manager.

The utility allows you to import the existing configuration to the RAID controller or clear the configuration so you can create a new configuration using these drives. You can preview the foreign configuration before you decide whether to import it.



MegaRAID Storage Manager normally detects newly installed drives and displays icons for them in the MegaRAID Storage Manager main menu screen. If for some reason MegaRAID Storage Manager 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 main menu screen.
2. Select **Go To > Controller > Scan for Foreign Configuration**.
If MegaRAID Storage Manager detects any new drives, it displays a list of them on the screen. If not, it notifies you that no foreign configuration is found.
3. Follow the instructions on the screen to complete the drive detection.

9.4 Rebuilding a Drive

If a single drive in a RAID 1, RAID 5, or RAID 10 virtual drive fails, the system is protected from data loss. If hot spare disks are available, a failed drive is rebuilt automatically without any user intervention. A failed drive must be replaced, and the data on the drive must be rebuilt on a new drive to restore the system to fault tolerance. (You can choose to rebuild the data on the failed drive if the drive is still operational.) If hot spare drives are available, the failed drive is rebuilt automatically without any user intervention.

If a drive has failed, a red circle appears to the right of the drive icon: . A small yellow circle appears to the right of the icon of the virtual drive that uses this drive: . This indicates that the virtual drive is in a degraded state; the data is still safe, but data could be lost if another drive fails.

Follow these steps if you need to rebuild a drive:

1. Right-click the icon of the failed drive, and select **Rebuild**.
2. Click **Yes** when the warning message appears.
If the drive is still good, a rebuild 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 MegaRAID Storage Manager.
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**.

9.5 Making a Drive Offline or Missing

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



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

To remove the drive from the configuration without harming the data on the virtual drive, follow these steps:

1. In the MegaRAID Storage Manager main menu, click **Go To > Physical Drive > Make Drive (O)ffline**.
The drive status changes to Offline.
2. Click **Go To > Physical Drive > (M)ark Drive as Missing**.
The drive status changes to Unconfigured Good.



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

3. If necessary, create a hot spare drive for the virtual drive from which you have removed the drive.
When a hot spare is available, the data on the virtual drive is then rebuilt. You can now use the removed drive for another configuration.



CAUTION If MegaRAID Storage Manager detects that a drive in a virtual drive has failed, it makes the drive offline. If this happens, you must remove the drive and replace it. You can make the drive so that another configuration cannot use it by using the **Mark physical disk as missing** command and the **Rescan** command.

Appendix A: Events and Messages

MegaRAID Storage Manager 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 screen.

This appendix lists the MegaRAID Storage Manager events that can appear in the event log.



NOTE MegaRAID Storage Manager can be used to manage a wide range of MegaRAID controllers. Some of the events and messages listed in this appendix are not applicable to Embedded MegaRAID Software.

Each message that appears in the event log has an error level that indicates the severity of the event, as shown in [Table 48](#).

Table 48 Event Error Levels

Error Level	Meaning
Information	Informational message; no user action is necessary.
Warning	Some component may be close to a failure point.
Caution	A component has failed, but the system has not lost data.
Fatal	A component has failed, and data loss has occurred or will occur.
Dead	A catastrophic error has occurred, and the controller has died. This event is seen only after the controller has been restarted.

[Table 49](#) 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. Some of the error messages are relevant only for hardware RAID.

Table 49 Event Messages

Number (Hex)	Number (Decimal)	Type	Event Text
0x0000	0	Information	Firmware initialization started (PCI ID %04x/%04x/%04x/%04x)
0x0001	1	Information	Firmware version %s
0x0002	2	Fatal	Unable to recover cache data from TBBU
0x0003	3	Information	Cache data recovered from TBBU successfully
0x0004	4	Information	Configuration cleared
0x0005	5	Warning	Cluster down; communication with peer lost
0x0006	6	Information	%s ownership changed from %02x to %02x
0x0007	7	Information	Alarm disabled by user
0x0008	8	Information	Alarm enabled by user
0x0009	9	Information	Background initialization rate changed to %d%%
0x000a	10	Fatal	Controller cache discarded due to memory/battery problems
0x000b	11	Fatal	Unable to recover cache data due to configuration mismatch
0x000c	12	Information	Cache data recovered successfully
0x000d	13	Fatal	Controller cache discarded due to firmware version incompatibility
0x000e	14	Information	Consistency Check rate changed to %d%%
0x000f	15	Dead	Fatal firmware error: %s

Table 49 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x0010	16	Information	Factory defaults restored
0x0011	17	Information	Flash downloaded image corrupt
0x0012	18	Caution	Flash erase error
0x0013	19	Caution	Flash timeout during erase
0x0014	20	Caution	Flash error
0x0015	21	Information	Flashing image: %s
0x0016	22	Information	Flash of new firmware image(s) complete
0x0017	23	Caution	Flash programming error
0x0018	24	Caution	Flash timeout during programming
0x0019	25	Caution	Flash chip type unknown
0x001a	26	Caution	Flash command set unknown
0x001b	27	Caution	Flash verify failure
0x001c	28	Information	Flush rate changed to %d seconds
0x001d	29	Information	Hibernate command received from host
0x001e	30	Information	Event log cleared
0x001f	31	Information	Event log wrapped
0x0020	32	Dead	Multi-bit ECC error: ECAR=%x, ELOG=%x, (%s)
0x0021	33	Warning	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s)
0x0022	34	Dead	Not enough controller memory
0x0023	35	Information	Patrol Read complete
0x0024	36	Information	Patrol Read paused
0x0025	37	Information	Patrol Read Rate changed to %d%%
0x0026	38	Information	Patrol Read resumed
0x0027	39	Information	Patrol Read started
0x0028	40	Information	Rebuild rate changed to %d%%
0x0029	41	Information	Reconstruction rate changed to %d%%
0x002a	42	Information	Shutdown command received from host
0x002b	43	Information	Test event: %s
0x002c	44	Information	Time established as %s; (%d seconds since power on)
0x002d	45	Information	User entered firmware debugger
0x002e	46	Warning	Background Initialization aborted on %s
0x002f	47	Warning	Background Initialization corrected medium error (%s at %lx
0x0030	48	Information	Background Initialization completed on %s
0x0031	49	Fatal	Background Initialization corrected medium error (%s at %lx, %s at %lx)
0x0032	50	Fatal	Background Initialization detected uncorrectable double medium errors (%s at %lx on %s)
0x0033	51	Caution	Background Initialization failed on %s
0x0034	52	Progress	Background Initialization progress on %s is %s
0x0035	53	Information	Background Initialization started on %s
0x0036	54	Information	Policy change on %s from %s to %s
0x0038	56	Warning	Consistency Check aborted on %s

Table 49 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x0039	57	Warning	Consistency Check corrected medium error (%s at %lx, %s at %lx)
0x003a	58	Information	Consistency Check done on %s
0x003b	59	Information	Consistency Check done with corrections on %s, (corrections=%d)
0x003c	60	Fatal	Consistency Check detected uncorrectable double medium errors (%s at %lx on %s)
0x003d	61	Caution	Consistency Check failed on %s
0x003e	62	Fatal	Consistency Check failed with uncorrectable data on %s
0x003f	63	Warning	Consistency Check found inconsistent parity on %s at strip %lx
0x0040	64	Warning	Consistency Check inconsistency logging disabled on %s (too many inconsistencies)
0x0041	65	Progress	Consistency Check progress on %s is %s
0x0042	66	Information	Consistency Check started on %s
0x0043	67	Warning	Initialization aborted on %s
0x0044	68	Caution	Initialization failed on %s
0x0045	69	Progress	Initialization progress on %s is %s
0x0046	70	Information	Fast initialization started on %s
0x0047	71	Information	Full initialization started on %s
0x0048	72	Information	Initialization complete on %s
0x0049	73	Information	Properties updated to %s (from %s)
0x004a	74	Information	Reconstruction complete on %s
0x004b	75	Fatal	Reconstruction of %s stopped due to unrecoverable errors
0x004c	76	Fatal	Reconstruct detected uncorrectable double medium errors (%s at %lx on %s at %lx)
0x004d	77	Progress	Reconstruction progress on %s is %s
0x004e	78	Information	Reconstruction resumed on %s
0x004f	79	Fatal	Reconstruction resume of %s failed due to configuration mismatch
0x0050	80	Information	Reconstructing started on %s
0x0051	81	Information	State change on %s from %s to %s
0x0052	82	Information	Clear aborted on %s
0x0053	83	Caution	Clear failed on %s (Error %02x)
0x0054	84	Progress	Clear progress on %s is %s
0x0055	85	Information	Clear started on %s
0x0056	86	Information	Clear completed on %s
0x0057	87	Warning	Error on %s (Error %02x)
0x0058	88	Information	Format complete on %s
0x0059	89	Information	Format started on %s
0x005a	90	Caution	Hot Spare SMART polling failed on %s (Error %02x)
0x005b	91	Information	Inserted: %s
0x005c	92	Warning	%s is not supported
0x005d	93	Warning	Patrol Read corrected medium error on %s at %lx
0x005e	94	Progress	Patrol Read progress on %s is %s
0x005f	95	Fatal	Patrol Read found an uncorrectable medium error on %s at %lx
0x0060	96	Caution	Predictive failure: %s

Table 49 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x0061	97	Fatal	Puncturing bad block on %s at %lx
0x0062	98	Information	Rebuild aborted by user on %s
0x0063	99	Information	Rebuild complete on %s
0x0064	100	Information	Rebuild complete on %s
0x0065	101	Caution	Rebuild failed on %s due to source drive error
0x0066	102	Caution	Rebuild failed on %s due to target drive error
0x0067	103	Progress	Rebuild progress on %s is %s
0x0068	104	Information	Rebuild resumed on %s
0x0069	105	Information	Rebuild started on %s
0x006a	106	Information	Rebuild automatically started on %s
0x006b	107	Caution	Rebuild stopped on %s due to loss of cluster ownership
0x006c	108	Fatal	Reassign write operation failed on %s at %lx
0x006d	109	Fatal	Unrecoverable medium error during rebuild on %s at %lx
0x006e	110	Information	Corrected medium error during recovery on %s at %lx
0x006f	111	Fatal	Unrecoverable medium error during recovery on %s at %lx
0x0070	112	Information	Removed: %s
0x0071	113	Warning	Unexpected sense: %s, CDB%s, Sense: %s
0x0072	114	Information	State change on %s from %s to %s
0x0073	115	Information	State change by user on %s from %s to %s
0x0074	116	Warning	Redundant path to %s broken
0x0075	117	Information	Redundant path to %s restored
0x0076	118	Information	Dedicated Hot Spare PD %s no longer useful due to deleted array
0x0077	119	Caution	SAS topology error: Loop detected
0x0078	120	Caution	SAS topology error: Unaddressable device
0x0079	121	Caution	SAS topology error: Multiple ports to the same SAS address
0x007a	122	Caution	SAS topology error: Expander error
0x007b	123	Caution	SAS topology error: SMP timeout
0x007c	124	Caution	SAS topology error: Out of route entries
0x007d	125	Caution	SAS topology error: Index not found
0x007e	126	Caution	SAS topology error: SMP function failed
0x007f	127	Caution	SAS topology error: SMP CRC error
0x0080	128	Caution	SAS topology error: Multiple subtractive
0x0081	129	Caution	SAS topology error: Table to table
0x0082	130	Caution	SAS topology error: Multiple paths
0x0083	131	Fatal	Unable to access device %s
0x0084	132	Information	Dedicated Hot Spare created on %s (%s)
0x0085	133	Information	Dedicated Hot Spare %s (%s) disabled
0x0086	134	Caution	Dedicated Hot Spare %s no longer useful for all arrays
0x0087	135	Information	Spare created on %s (%s)
0x0088	136	Information	Spare %s (%s) disabled

Table 49 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x0089	137	Caution	Spare %s does not cover all arrays
0x008a	138	Information	Created %s
0x008b	139	Information	Deleted %s
0x008c	140	Information	Marking %s inconsistent due to active writes at shutdown
0x008d	141	Information	Battery Present
0x008e	142	Warning	Battery Not Present
0x008f	143	Information	New Battery Detected
0x0090	144	Information	Battery has been replaced
0x0091	145	Caution	Battery temperature is high
0x0092	146	Warning	Battery voltage low
0x0093	147	Information	Battery started charging
0x0094	148	Information	Battery is discharging
0x0095	149	Information	Battery temperature is normal
0x0096	150	Fatal	Battery needs replacement - SOH Bad
0x0097	151	Information	Battery relearn started
0x0098	152	Information	Battery relearn in progress
0x0099	153	Information	Battery relearn completed
0x009a	154	Caution	Battery relearn timed out
0x009b	155	Information	Battery relearn pending: Battery is under charge
0x009c	156	Information	Battery relearn postponed
0x009d	157	Information	Battery relearn will start in 4 days
0x009e	158	Information	Battery relearn will start in 2 day
0x009f	159	Information	Battery relearn will start in 1 day
0x00a0	160	Information	Battery relearn will start in 5 hours
0x00a1	161	Information	Battery removed
0x00a2	162	Information	Current capacity of the battery is below threshold
0x00a3	163	Information	Current capacity of the battery is above threshold
0x00a4	164	Information	Enclosure (SES) discovered on %s
0x00a5	165	Information	Enclosure (SAFTE) discovered on %s
0x00a6	166	Caution	Enclosure %s communication lost
0x00a7	167	Information	Enclosure %s communication restored
0x00a8	168	Caution	Enclosure %s fan %d failed
0x00a9	169	Information	Enclosure %s fan %d inserted
0x00aa	170	Caution	Enclosure %s fan %d removed
0x00ab	171	Caution	Enclosure %s power supply %d failed
0x00ac	172	Information	Enclosure %s power supply %d inserted
0x00ad	173	Caution	Enclosure %s power supply %d removed
0x00ae	174	Caution	Enclosure %s EMM %d failed
0x00af	175	Information	Enclosure %s EMM %d inserted
0x00b0	176	Caution	Enclosure %s EMM %d removed

Table 49 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x00b1	177	Warning	Enclosure %s temperature sensor %d below warning threshold
0x00b2	178	Caution	Enclosure %s temperature sensor %d below error threshold
0x00b3	179	Warning	Enclosure %s temperature sensor %d above warning threshold
0x00b4	180	Caution	Enclosure %s temperature sensor %d above error threshold
0x00b5	181	Caution	Enclosure %s shutdown
0x00b6	182	Warning	Enclosure %s not supported; too many enclosures connected to port
0x00b7	183	Caution	Enclosure %s firmware mismatch (EMM %d)
0x00b8	184	Warning	Enclosure %s sensor %d bad
0x00b9	185	Caution	Enclosure %s phy bad for slot %d
0x00ba	186	Caution	Enclosure %s is unstable
0x00bb	187	Caution	Enclosure %s hardware error
0x00bc	188	Caution	Enclosure %s not responding
0x00bd	189	Information	SAS/SATA mixing not supported in enclosure; %s disabled
0x00be	190	Information	Enclosure (SES) hotplug on %s was detected, but is not supported
0x00bf	191	Information	Clustering enabled
0x00c0	192	Information	Clustering disabled
0x00c1	193	Information	PD too small to be used for auto-rebuild on %s
0x00c2	194	Information	BBU enabled; changing WT virtual disks to WB
0x00c3	195	Warning	BBU disabled; changing WB virtual disks to WT
0x00c4	196	Warning	Bad block table on %s is 80% full
0x00c5	197	Fatal	Bad block table on %s is full; unable to log block %lx
0x00c6	198	Information	Consistency Check Aborted Due to Ownership Loss on %s
0x00c7	199	Information	Background Initialization (BGI) Aborted Due to Ownership Loss on %s
0x00c8	200	Caution	Battery/charger problems detected; SOH Bad
0x00c9	201	Warning	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); warning threshold exceeded
0x00ca	202	Caution	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); critical threshold exceeded
0x00cb	203	Caution	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); further reporting disabled
0x00cc	204	Caution	Enclosure %s Power supply %d switched off
0x00cd	205	Information	Enclosure %s Power supply %d switched on
0x00ce	206	Caution	Enclosure %s Power supply %d cable removed
0x00cf	207	Information	Enclosure %s Power supply %d cable inserted
0x00d0	208	Information	Enclosure %s Fan %d returned to normal
0x00d1	209	Information	BBU Retention test was initiated on previous boot
0x00d2	210	Information	BBU Retention test passed
0x00d3	211	Caution	BBU Retention test failed!
0x00d4	212	Information	NVRAM Retention test was initiated on previous boot
0x00d5	213	Information	NVRAM Retention test passed
0x00d6	214	Caution	NVRAM Retention test failed!
0x00d7	215	Information	%s test completed %d passes successfully
0x00d8	216	Caution	%s test FAILED on %d pass. Fail data: errorOffset=%x goodData=%x badData=%x

Table 49 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x00d9	217	Information	Self check diagnostics completed
0x00da	218	Information	Foreign Configuration Detected
0x00db	219	Information	Foreign Configuration Imported
0x00dc	220	Information	Foreign Configuration Cleared
0x00dd	221	Warning	NVRAM is corrupt; reinitializing
0x00de	222	Warning	NVRAM mismatch occurred
0x00df	223	Warning	SAS wide port %d lost link on PHY %d
0x00e0	224	Information	SAS wide port %d restored link on PHY %d
0x00e1	225	Warning	SAS port %d, PHY %d has exceeded the allowed error rate
0x00e2	226	Warning	Bad block reassigned on %s at %lx to %lx
0x00e3	227	Information	Controller Hot Plug detected
0x00e4	228	Warning	Enclosure %s temperature sensor %d differential detected
0x00e5	229	Information	Disk test cannot start. No qualifying disks found
0x00e6	230	Information	Time duration provided by host is not sufficient for self check
0x00e7	231	Information	Marked Missing for %s on array %d row %d
0x00e8	232	Information	Replaced Missing as %s on array %d row %d
0x00e9	233	Information	Enclosure %s Temperature %d returned to normal
0x00ea	234	Information	Enclosure %s Firmware download in progress
0x00eb	235	Warning	Enclosure %s Firmware download failed
0x00ec	236	Warning	%s is not a certified drive
0x00ed	237	Information	Dirty cache data discarded by user
0x00ee	238	Information	PDs missing from configuration at boot
0x00ef	239	Information	VDs missing drives and will go offline at boot: %s
0x00f0	240	Information	VDs missing at boot: %s
0x00f1	241	Information	Previous configuration completely missing at boot
0x00f2	242	Information	Battery charge complete
0x00f3	243	Information	Enclosure %s fan %d speed changed
0x00f4	244	Information	Dedicated spare %s imported as global due to missing arrays
0x00f5	245	Information	%s rebuild not possible as SAS/SATA is not supported in an array
0x00f6	246	Information	SEP %s has been rebooted as a part of enclosure firmware download. SEP will be unavailable until this process completes.
0x00f7	247	Information	Inserted: %s Info: %s
0x00f8	248	Information	Removed: %s Info: %s
0x00f9	249	Information	%s is now OPTIMAL
0x00fa	250	Warning	%s is now PARTIALLY DEGRADED
0x00fb	251	Caution	%s is now DEGRADED
0x00fc	252	Fatal	%s is now OFFLINE
0x00fd	253	Warning	Battery requires reconditioning; please initiate a LEARN cycle
0x00fe	254	Warning	VD %s disabled because RAID-5 is not supported by this RAID key
0x00ff	255	Warning	VD %s disabled because RAID-6 is not supported by this controller

Table 49 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x0100	256	Warning	VD %s disabled because SAS drives are not supported by this RAID key
0x0101	257	Warning	PD missing: %s
0x0102	258	Warning	Puncturing of LBAs enabled
0x0103	259	Warning	Puncturing of LBAs disabled
0x0104	260	Critical	Enclosure %s EMM %d not installed
0x0105	261	Information	Package version %s
0x0106	262	Warning	Global affinity Hot Spare %s commissioned in a different enclosure
0x0107	263	Warning	Foreign configuration table overflow
0x0108	264	Warning	Partial foreign configuration imported, PDs not imported:%s
0x0109	265	Information	Connector %s is active
0x010a	266	Information	Board Revision %s
0x010b	267	Warning	Command timeout on PD %s, CDB:%s
0x010c	268	Warning	PD %s reset (Type %02x)
0x010d	269	Warning	VD bad block table on %s is 80% full
0x010e	270	Fatal	VD bad block table on %s is full; unable to log block %lx (on %s at %lx)
0x010f	271	Fatal	Uncorrectable medium error logged for %s at %lx (on %s at %lx)
0x0110	272	Information	VD medium error corrected on %s at %lx
0x0111	273	Warning	Bad block table on PD %s is 100% full
0x0112	274	Warning	VD bad block table on PD %s is 100% full
0x0113	275	Fatal	Controller needs replacement, IOP is faulty
0x0114	276	Information	CopyBack started on PD %s from PD %s
0x0115	277	Information	CopyBack aborted on PD %s and src is PD %s
0x0116	278	Information	CopyBack complete on PD %s from PD %s
0x0117	279	Progress	CopyBack progress on PD %s is %s
0x0118	280	Information	CopyBack resumed on PD %s from %s
0x0119	281	Information	CopyBack automatically started on PD %s from %s
0x011a	282	Critical	CopyBack failed on PD %s due to source %s error
0x011b	283	Warning	Early Power off warning was unsuccessful
0x011c	284	Information	BBU FRU is %s
0x011d	285	Information	%s FRU is %s
0x011e	286	Information	Controller hardware revision ID %s
0x011f	287	Warning	Foreign import shall result in a backward incompatible upgrade of configuration metadata
0x0120	288	Information	Redundant path restored for PD %s
0x0121	289	Warning	Redundant path broken for PD %s
0x0122	290	Information	Redundant enclosure EMM %s inserted for EMM %s
0x0123	291	Information	Redundant enclosure EMM %s removed for EMM %s
0x0124	292	Warning	Patrol Read can't be started, as PDs are either not ONLINE, or are in a VD with an active process, or are in an excluded VD
0x0125	293	Information	Copyback aborted by user on PD %s and src is PD %s
0x0126	294	Critical	Copyback aborted on hot spare %s from %s, as hot spare needed for rebuild

Table 49 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x0127	295	Warning	Copyback aborted on PD %s from PD %s, as rebuild required in the array
0x0128	296	Fatal	Controller cache discarded for missing or offline VD %s When a VD with cached data goes offline or missing during runtime, the cache for the VD is discarded. Because the VD is offline, the cache cannot be saved.
0x0129	297	Information	Copyback cannot be started as PD %s is too small for src PD %s
0x012a	298	Information	Copyback cannot be started on PD %s from PD %s, as SAS/SATA is not supported in an array
0x012b	299	Information	Microcode update started on PD %s
0x012c	300	Information	Microcode update completed on PD %s
0x012d	301	Warning	Microcode update timeout on PD %s
0x012e	302	Warning	Microcode update failed on PD %s
0x012f	303	Information	Controller properties changed
0x0130	304	Information	Patrol Read properties changed
0x0131	305	Information	CC Schedule properties changed
0x0132	306	Information	Battery properties changed
0x0133	307	Warning	Periodic Battery Relearn is pending. Please initiate manual learn cycle as Automatic learn is not enabled
0x0134	308	Information	Drive security key created
0x0135	309	Information	Drive security key backed up
0x0136	310	Information	Drive security key from escrow, verified
0x0137	311	Information	Drive security key changed
0x0138	312	Warning	Drive security key, re-key operation failed
0x0139	313	Warning	Drive security key is invalid
0x013a	314	Information	Drive security key destroyed
0x013b	315	Warning	Drive security key from escrow is invalid
0x013c	316	Information	VD %s is now secured
0x013d	317	Warning	VD %s is partially secured
0x013e	318	Information	PD %s security activated
0x013f	319	Information	PD %s security disabled
0x0140	320	Information	PD %s is reprovisioned
0x0141	321	Information	PD %s security key changed
0x0142	322	Fatal	Security subsystem problems detected for PD %s
0x0143	323	Fatal	Controller cache pinned for missing or offline VD %s
0x0144	324	Fatal	Controller cache pinned for missing or offline VD: %s
0x0145	325	Information	Controller cache discarded by user for VD: %s
0x0146	326	Information	Controller cache destaged for VD %s
0x0147	327	Warning	Consistency Check started on an inconsistent VD %s
0x0148	328	Warning	Drive security key failure, cannot access secured configuration
0x0149	329	Warning	Drive security password from user is invalid
0x014a	330	Warning	Detected error with the remote battery connector cable
0x014b	331	Information	Power state change on PD %s from %s to %s
0x014c	332	Information	Enclosure %s element (SES code 0x%x) status changed

Table 49 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x014d	333	Information	PD %s rebuild not possible as HDD/CacheCade software mix is not supported in a drive group
0x014e	334	Information	Copyback cannot be started on PD %s from %s, as HDD/CacheCade software mix is not supported in a drive group
0x014f	335	Information	VD bad block table on %s is cleared
0x0150	336	Caution	SAS topology error: 0x%lx
0x0151	337	Information	VD cluster of medium errors corrected for %s at %lx (on %s at %lx)
0x0152	338	Information	Controller requests a host bus rescan
0x0153	339	Information	Controller repurposed and factory defaults restored
0x0154	340	Information	Drive security key binding updated
0x0155	341	Information	Drive security is in EKM mode
0x0156	342	Warning	Drive security failed to communicate with EKMS
0x0157	343	Information	%s needs key to be %s %s
0x0158	344	Warning	%s secure failed
0x0159	345	Critical	Controller encountered a fatal error and was reset
0x015a	346	Information	Snapshots enabled on %s (Repository %s)'
0x015b	347	Information	Snapshots disabled on %s (Repository %s) by the user
0x015c	348	Critical	Snapshots disabled on %s (Repository %s), due to a fatal error
0x015d	349	Information	Snapshot created on %s at %s
0x015e	350	Information	Snapshot deleted on %s at %s
0x015f	351	Information	View created at %s to a snapshot at %s for %s
0x0160	352	Information	View at %s is deleted, to snapshot at %s for %s
0x0161	353	Information	Snapshot rollback started on %s from snapshot at %s
0x0162	354	Fatal	Snapshot rollback on %s internally aborted for snapshot at %s'
0x0163	355	Information	Snapshot rollback on %s completed for snapshot at %s
0x0164	356	Information	Snapshot rollback progress for snapshot at %s, on %s is %s
0x0165	357	Warning	Snapshot space for %s in snapshot repository %s, is 80%% full
0x0166	358	Critical	Snapshot space for %s in snapshot repository %s, is full
0x0167	359	Warning	View at %s to snapshot at %s, is 80%% full on snapshot repository %s
0x0168	360	Critical	View at %s to snapshot at %s, is full on snapshot repository %s
0x0169	361	Critical	Snapshot repository lost for %s
0x016a	362	Warning	Snapshot repository restored for %s
0x016b	363	Critical	Snapshot encountered an unexpected internal error: 0x%lx
0x016c	364	Information	Auto Snapshot enabled on %s (snapshot repository %s)
0x016d	365	Information	Auto Snapshot disabled on %s (snapshot repository %s)
0x016e	366	Critical	Configuration command could not be committed to disk, please retry
0x016f	367	Information	COD on %s updated as it was stale
0x0170	368	Warning	Power state change failed on %s (from %s to %s)
0x0171	369	Warning	%s is not available
0x0172	370	Information	%s is available

Table 49 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x0173	371	Information	%s is used for CacheCade with capacity 0x%lx logical blocks
0x0174	372	Information	%s is using CacheCade %s
0x0175	373	Information	%s is no longer using CacheCade %s
0x0176	374	Critical	Snapshot deleted due to resource constraints for %s in snapshot repository %s
0x0177	375	Warning	Auto Snapshot failed for %s in snapshot repository %s
0x0178	376	Warning	Controller reset on-board expander
0x0179	377	Warning	CacheCade (%s) capacity changed and is now 0x%lx logical blocks
0x017a	378	Warning	Battery cannot initiate transparent learn cycles
0x017b	379	Information	Premium feature %s key was applied for - %s
0x017c	380	Information	Snapshot schedule properties changed on %s
0x017d	381	Information	Snapshot scheduled action is due on %s
0x017e	382	Information	Performance Metrics: collection command 0x%lx
0x017f	383	Information	Premium feature %s key was transferred - %s
0x0180	384	Information	Premium feature serial number %s
0x0181	385	Warning	Premium feature serial number mismatched. Key-vault serial num - %s
0x0182	386	Warning	Battery cannot support data retention for more than %d hours. Please replace the battery
0x0183	387	Information	%s power policy changed to %s (from %s)
0x0184	388	Warning	%s cannot transition to max power savings
0x0185	389	Information	Host driver is loaded and operational
0x0186	390	Information	%s mirror broken
0x0187	391	Information	%s mirror joined
0x0188	392	Warning	%s link %d failure in wide port
0x0189	393	Information	%s link %d restored in wide port
0x018a	394	Information	Memory module FRU is %s
0x018b	395	Warning	Cache-vault power pack is sub-optimal. Please replace the pack
0x018c	396	Warning	Foreign configuration auto-import did not import any drives
0x018d	398	Warning	Cache-vault microcode update required
0x018e	399	Warning	CacheCade (%s) capacity exceeds maximum allowed size, extra capacity is not used
0x018f	399	Warning	LD (%s) protection information lost
0x0190	400	Information	Diagnostics passed for %s
0x0191	401	Critical	Diagnostics failed for %s
0x0192	402	Information	Server Power capability Diagnostic Test Started
0x0193	403	Information	Drive Cache settings enabled during rebuild for %s
0x0194	404	Information	Drive Cache settings restored after rebuild for %s
0x0195	405	Information	Drive %s commissioned as Emergency spare
0x0196	406	Warning	Reminder: Potential non-optimal configuration due to drive %s commissioned as emergency spare
0x0197	407	Information	Consistency Check suspended on %s
0x0198	408	Information	Consistency Check resumed on %s
0x0199	409	Information	Background Initialization suspended on %s

Table 49 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x019a	410	Information	Background Initialization resumed on %
0x019b	411	Information	Reconstruction suspended on %s
0x019c	412	Information	Rebuild suspended on %
0x019d	413	Information	Copyback suspended on %s
0x019e	414	Information	Reminder: Consistency Check suspended on %
0x019f	415	Information	Reminder: Background Initialization suspended on %s
0x01a0	416	Information	Reminder: Reconstruction suspended on %s
0x01a1	417	Information	Reminder: Rebuild suspended on %s
0x01a2	418	Information	Reminder: Copyback suspended on %s
0x01a3	419	Information	Reminder: Patrol Read suspended
0x01a4	420	Information	Erase aborted on %s
0x01a5	421	Critical	Erase failed on %s (Error %02x)
0x01a6	422	Progress	Erase progress on %s is %s
0x01a7	423	Information	Erase started on %s
0x01a8	424	Information	Erase completed on %s
0x01a9	425	Information	Erase aborted on %s
0x01aa	426	Critical	Erase failed on %s
0x01ab	427	Progress	Erase progress on %s is %s
0x01ac	428	Information	Erase started on %s
0x01ad	429	Information	Erase complete on %s
0x01ae	430	Warning	Potential leakage during erase on %s
0x01af	431	Warning	Battery charging was suspended due to high battery temperature
0x01b0	432	Information	NVCache firmware update was successful
0x01b1	433	Warning	NVCache firmware update failed
0x01b2	434	Fatal	%s access blocked as cached data in CacheCade is unavailable
0x01b3	435	Information	CacheCade disassociate started on %s
0x01b4	436	Information	CacheCade disassociate completed on %s
0x01b5	437	Critical	CacheCade disassociate failed on %s
0x01b6	438	Progress	CacheCade disassociate progress on %s is %s
0x01b7	439	Information	CacheCade disassociate aborted by user on %s
0x01b8	440	Information	Link speed changed on SAS port %d and PHY %d
0x01b9	441	Warning	Advanced Software Options was deactivated for - %s
0x01ba	442	Information	%s is now accessible
0x01bb	443	Information	%s is using CacheCade
0x01bc	444	Information	%s is no longer using CacheCade
0x01bd	445	Information	Patrol Read aborted on %s



48712-00B

