lenovo

ServeRAID® C100 and ServeRAID C105

User's Guide

Fourth Edition
July 2015

Fourth Edition (July 2015)

Revision History

Version and Date	Description of Changes
Fourth Edition, July 2015	Updated the guide to Lenovo Brand and corrected the operating system and system support.
Third Edition, April 2014	Updated section1.4.1, Driver Features. Added a note in Appendix A, MDRAID Migration. Updated Table 2, Controller Properties.
Second Edition, September 2013	Added Chapter 3, Human Interface Infrastructure Configuration Utility. Added Appendix A, MDRAID Migration.
First Edition, June 2012	Initial release of the document.

1 Safety

Before installing this product, read the Safety Information.

قبل تركيب هذا المنتج، يجب قراءة الملاحظات الأمنية

Antes de instalar este produto, leia as Informações de Segurança.

在安装本产品之前，请仔细阅读 Safety Information
(安全信息)。

安裝本產品之前，請先閱讀「安全資訊」。

Prije instalacije ovog produkta obavezno pročitajte Sigurnosne Upute.

Před instalací tohoto produktu si přečtěte příručku bezpečnostních instrukcí.

Læs sikkerhedsforskrifterne, før du installerer dette produkt.

Lees voordat u dit product installeert eerst de veiligheidsvoorschriften.

Ennen kuin asennat tämän tuotteen, lue turvaohjeet kohdasta Safety Information.

Avant d'installer ce produit, lisez les consignes de sécurité.

Vor der Installation dieses Produkts die Sicherheitshinweise lesen.

Πριν εγκαταστήσετε το προϊόν αυτό, διαβάστε τις πληροφορίες ασφάλειας
(safety information).

לפני שתתקינו מוצר זה, קראו את הוראות הבטיחות.

A termék telepítése előtt olvassa el a Biztonsági előírásokat!

Prima di installare questo prodotto, leggere le Informazioni sulla Sicurezza.

製品の設置の前に、安全情報をお読みください。

본 제품을 설치하기 전에 안전 정보를 읽으십시오.

Пред да се инсталира овој продукт, прочитајте информацијата за безбедност.

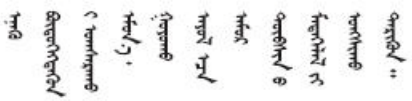

Les sikkerhetsinformasjonen (Safety Information) før du installerer dette produktet.

Przed zainstalowaniem tego produktu, należy zapoznać się
z książką "Informacje dotyczące bezpieczeństwa" (Safety Information).

Antes de instalar este produto, leia as Informações sobre Segurança.

Перед установкой продукта прочтите инструкции по
технике безопасности.

Pred inštaláciou tohto zariadenia si pečítaje Bezpečnostné predpisy.

Pred namestitvijo tega proizvoda preberite Varnostne informacije.

Antes de instalar este producto, lea la información de seguridad.

Läs säkerhetsinformationen innan du installerar den här produkten.

ཐོན་རྫས་འདི་བདེ་སྒྲིད་མ་བྱས་གོང་། ་སྐོར་གྱི་ཡིད་གཟབ་
བྱ་འདྲ་མིན་ཡོད་པའི་འོད་སྣེར་བལྟ་དགོས།

Bu ürünü kurmadan önce güvenlik bilgilerini okuyun.

ھەزكۇر مەھسۇلاتنى ئورنىتىشتىن بۇرۇن بىخەتەرلىك ئۇچۇرلىرىنى ئوقۇپ چىقىڭ.

Youq mwngz yungh canjbinj neix gaxgonq, itdingh aeu doeg aen
canjbinj soengq cungj vahgangj ancien siusik.

2 Safety statements

These statements provide the caution and danger information that is used in this documentation.

Important:

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

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

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

Statement 1:

DANGER

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

To avoid a shock hazard:

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

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

Statement 3:

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

- **Do not remove the covers. Removing the covers of the laser product could result in exposure to hazardous laser radiation. There are no serviceable parts inside the device.**
- **Use of controls or adjustments or performance of procedures other than those specified herein might result in hazardous radiation exposure.**

DANGER

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

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

Statement 8:

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

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

Statement 28:

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

Table of Contents

Chapter 1: Overview

This guide documents the ServeRAID® C100 and ServeRAID C105 features, and includes instructions for using the Human Interface Infrastructure Configuration Utility, the StorCLI command line utility, and MegaRAID Storage Manager™. You can use these utilities to create storage configurations on drives controlled by ServeRAID C105. The manual also includes instructions for installing the ServeRAID C100 and ServeRAID C105 drivers in the Microsoft® Windows® operating systems and the Linux™ operating system.

1.1 Notices and statements

The caution and danger statements in this document are also in the multilingual *Safety Information* document, which is on the *Documentation* CD. Each statement is numbered for reference to the corresponding statement in the *Safety Information* document.

Notices and statements in this document

The following notices and statements are used in this document:

- **Note:** These notices provide important tips, guidance, or advice.
- **Important:** These notices provide information or advice that might help you avoid inconvenient or problem situations.
- **Attention:** These notices indicate possible damage to programs, devices, or data. An attention notice is placed just before the instruction or situation in which damage might occur.
- **Caution:** These statements indicate situations that can be potentially hazardous to you. A caution statement is placed just before the description of a potentially hazardous procedure step or situation.
- **Danger:** These statements indicate situations that can be potentially lethal or hazardous to you. A danger statement is placed just before the description of a potentially lethal or hazardous procedure step or situation.

1.2 Parts listing

Replaceable components consist of consumable parts and field replaceable units (FRUs):

- **Consumable part:** Purchase and replacement of consumable parts (components, such as batteries and printer cartridges, that have depletable life) is your responsibility. If Lenovo acquires or installs a consumable part at your request, you will be charged for the service.
- **Field replaceable unit (FRU):** FRUs must be replaced only by a trained service technician, unless they are classified as customer replaceable units (CRUs):
  - **Tier 1 customer replaceable unit (CRU):** Replacement of Tier 1 CRUs is your responsibility. If Lenovo installs a Tier 1 CRU at your request without a service contract, you will be charged for the installation.
  - **Tier 2 customer replaceable unit:** You may install a Tier 2 CRU yourself or request Lenovo to install it, at no additional charge, under the type of warranty service that is designated for your product.

For information about the terms of the warranty and getting service and assistance, see the *Lenovo Safety, Support, and Warranty Information* document.

1.3 Conventions

The following is a list of notational conventions used throughout this guide.

Table 1 Conventions

Notation	Example	Meaning and Use
Courier typeface	`.nwk` file	Names of commands, files, and directories are shown in Courier typeface.
Bold typeface	**fd1sp**	In a command line, keywords are shown in bold, non-italic typeface. Enter them exactly as shown.
Italics	*module*	In command lines and names italics indicate user variables. Italicized text must be replaced with appropriate user-specified items. Enter items of the type called for, using lowercase.
Italic underscore	*full_pathname*	When an underscore appears in an italicized string, enter a user-supplied item of the type called for with no spaces.
Initial capital letters	Undo Edit Apply	Names of menu commands, options, check buttons, text buttons, options buttons, text boxes, list boxes, and so on, are shown in text with initial capital lettering to avoid misreading. These elements may appear on your screen in all lowercase.
Brackets	[*version*]	You may, but need not, select one item enclosed within brackets. Do not enter the brackets.
Bar	les | les.out2	You may select one (but not more than one) item from a list separated by bars. Do not enter the bar.

Table 1 Conventions

Notation	Example	Meaning and Use
Braces	{property | -all}	You must select one (but not more than one) item enclosed within braces. Do not enter the braces.
Ellipses	option...	In command formats, elements preceding ellipses may be repeated any number of times. Do not enter the ellipses. In menu items, if an ellipsis appears in an item, clicking that item brings up a dialog box.
Semicolon and other punctuation		Use as shown in the text.

NOTE Notes contain supplementary information that can affect system performance.

CAUTION Attentions are notifications that an action has the potential to adversely affect equipment operation, system performance, or data integrity.

1.4 ServeRAID C100 and ServeRAID C105 Features

ServeRAID C100 and ServeRAID C105 offer software RAID functionality and provide the ability to create and manage storage configurations on drives controlled by ServeRAID C100 and ServeRAID C105. ServeRAID C100 and ServeRAID C105 provide a cost-effective way to achieve high transfer rates and reliability.

The following sections list the driver features, BIOS features, disk management features, and UEFI support features for ServeRAID C100 and ServeRAID C105. In addition, the following section documents the differences between ServeRAID C100 and ServeRAID C105.

1.4.1 Driver Features

The driver for ServeRAID C100 and ServeRAID C105 supports the following features:

- Support for 48-bit Logical Block Addressing (LBA)
- Support for drive roaming
- Support for virtual drives larger than 2 terabytes
- Support for migration path from ServeRAID C100 and ServeRAID C105 to ServeRAID SATA hardware (this feature requires support from hardware RAID)
- Support for Patrol Read
- Automatic resumption of rebuilding, Check Consistency, full initialization, and background initialization
- Online mirror rebuilding
- Support for auto rebuild
- Support for SATA 6Gb/s drives

NOTE ServeRAID C105 supports only SATA HDDs; it does not support SSD or SAS.

- Check Consistency for RAID 1 and RAID 10

- Global hotspare support

NOTE ServeRAID C105 supports global hotspares, but it does not support dedicated hotspares.

- Soft Bad Block Management (SBBM) support
- Support for up to 8 physical drives and eight virtual drives
- Stripe size of 64 Kbytes only
- Support for Disk Coercion (1Gbyte)
- Hot Plug support (drive insertion and removal) in the Microsoft® Windows® operating system and the Linux™ operating system
- Support for random deletion of virtual drives
- Error logging and notification
- Support for Microsoft Windows Server 2008, Microsoft Windows Server 2008 R2, Microsoft Windows Server 2012, and Microsoft Windows Server 2012 R2 operating systems
- Support for Red Hat Enterprise Linux (RHEL) and SuSE Linux Enterprise Server (SLES) using kernels 2.6 and above

1.4.2 BIOS Features

The ServeRAID C100 and ServeRAID C105 BIOS support has the following features:

- Support for Interrupt 13 and Enhanced Disk Drive Specification
- Support for Int19h
- Support for BIOS Boot Specification (BBS) (If available in system BIOS, this feature allows the user to 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 features are supported by the BIOS:

- Automatic resumption of rebuilding, Check Consistency, and full initialization, and Background Initialization (BGI)

NOTE The BIOS does not start or resume background initialization. If BGI is already in progress, you cannot start Check Consistency.

- Global hotspare support
- Soft Bad Block Management (SBBM) support
- Support for RAID levels 0, 1, and 10
- Support for auto rebuild
- Support for up to eight physical drives and eight virtual drives
- Stripe size of 64 Kbytes only
- Support for Disk Coercion, with options *None*, *128 Mbytes*, and *1 Gbyte*

1.4.3 HII Configuration Features

The HII Configuration Utility supports the following features:

- Ability to configure controllers, drive groups, and virtual drives in a pre-boot environment
- Ability to perform other configuration tasks in a pre-boot environment
- Ability to select and change the settings for some virtual drive parameters
- Ability to select a virtual drive as boot device (by default, virtual drive 0 is the boot drive)
- Support for RAID levels 0, 1, and 10
- Support for running a consistency check
- Support for running a patrol read
- Ability to set the rates for the BGI, consistency check, and patrol read
- Ability to place drives online or offline

1.4.4 Manageability/Disk Console Features

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

- Configuration information display (in HII Configuration Utility and MegaRAID Storage Manager)
- Support for RAID levels 0, 1, and 10
- Rebuilding mirror while system is operational
- Consistency check while system is operational
- Array management software
- Error logging and notification
- Support for hot device insertion and removal
- Automatic resume of rebuilding on restart
- Support for manual rebuild
- Ability to create up to eight virtual drives per configuration
- Auto-configuration support of newly added drive
- Support for global hotspares
- Support for disk coercion
- Drive group initialization support (fast and normal)
- Virtual drive availability immediately after creation
- Supported stripe size of 64 Kbytes only

1.4.5 UEFI Support

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

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

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

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

1.4.6 Differences Between ServeRAID C100 and ServeRAID C105

The following table identifies the differences between ServeRAID C100 and ServeRAID C105.

Table 2 Controller Properties

Property	ServeRAID C100	ServeRAID C105
Number of Ports	■ 4 AHCI ports	■ 4 standard SCU ports ■ 4 SCU ports with Feature-on-Demand (FoD) key
Drive Support	SATA HDD	SATA HDD
Drive Swap Support	Simple swap and hotswap based on system support	Simple swap and hotswap
Optical Device Support	■ Supported ■ Needs to be enabled in F1 Setup	■ Not supported
RAID Support	Needs to be enabled in F1 Setup	Always enabled
RAID Level Support	RAID 0, 1, 10	RAID 0, 1, 10
Server Support	x3100 M4, 3100 M5, x3250 M4, x3250 M5, nx360 M4, x222	x3630 M4, x3530 M4, x220, x3300 M4

1.5 RAID Overview

This section provides a brief overview of the types of RAID configurations that ServeRAID C100 and ServeRAID C105 support.

The first step in creating a RAID storage configuration is to configure drives in *drive groups* (also known as *arrays*). As defined for ServeRAID C105, a *drive group* is a group of one to eight drives that is seen by the host computer system as one large disk drive, or *virtual drive*. Only one RAID level can be assigned to each array.

- A RAID 0 drive group consists of one to eight drives.
- A RAID 1 drive group consists of two 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 and RAID 10 arrays might 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 array and is also written (mirrored) to a second drive.
- **Maximize Virtual Drive Performance**: You can maximize virtual drive performance by creating a RAID 0 array 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 array.

  RAID 10 arrays 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).

1.5.1 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 array. The size of each data segment is determined by the strip size, which is 64 Kbytes.

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

By breaking up a large file into smaller segments, and writing or reading from several drives at once, ServeRAID C105 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 array with two drives.

Figure 1 RAID 0 Array Example with Two Drives

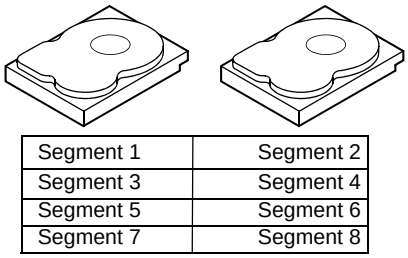

1.5.2 RAID 1 Description

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 Array

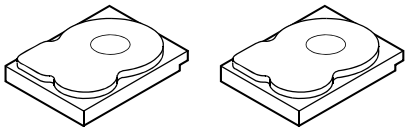

Segment 1	Segment 1 Duplicated
Segment 2	Segment 2 Duplicated
Segment 3	Segment 3 Duplicated
Segment 4	Segment 4 Duplicated

1.5.3 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 Kbytes. RAID 10 can sustain one drive failure in each drive group while maintaining data integrity.

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.

Uses	Works best for data storage that must have 100% 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 array with four drives.

Figure 3 RAID 10 Array

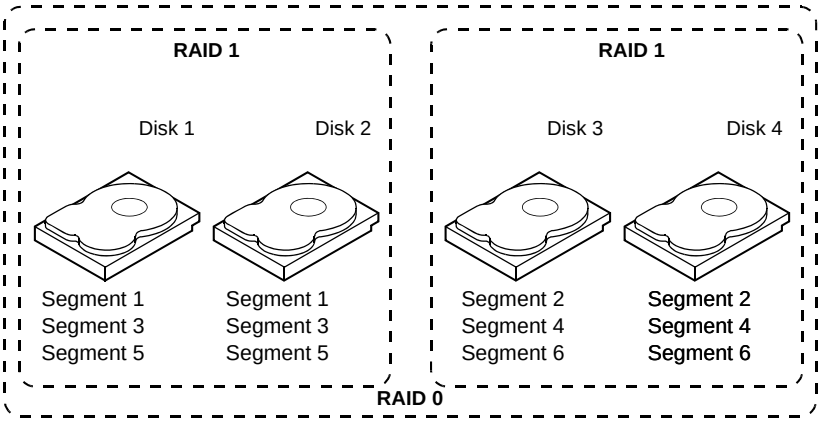

Chapter 2: Driver Installation

Device drivers for the ServeRAID C100 and ServeRAID C105 can be obtained from the **Support & downloads** section of the IBM web site (https://www-947.ibm.com/support/entry/portal/docdisplay?lndocid=SERV-RAID). Usage instructions and step-by-step installation instructions for your specific device driver are available in the README file that accompanies that driver on the download site.

Chapter 3: Human Interface Infrastructure Configuration Utility

The Human Interface Infrastructure (HII) Configuration Utility (CU) is used to configure controllers, drive groups, and virtual drives, and to perform other configuration tasks in a pre-boot environment.

This chapter describes how to configure controllers, drive groups, and virtual drives with the HII Configuration Utility. To ensure the best performance, select the optimal RAID level for the virtual drive you create. For an explanation of RAID levels, see Section 1.5, RAID Overview.

3.1 Accessing the Configuration Options Screen

You can use the Configuration Options screen to manage the controller, virtual drives, and physical drives.

Perform the following steps to access the Configuration Options screen.

1. Boot the system.
2. Press **F1** during bootup to enter setup.
   The System Configuration and Boot Management screen appears, as shown in the following figure.

Figure 1 System Configuration and Boot Management Screen

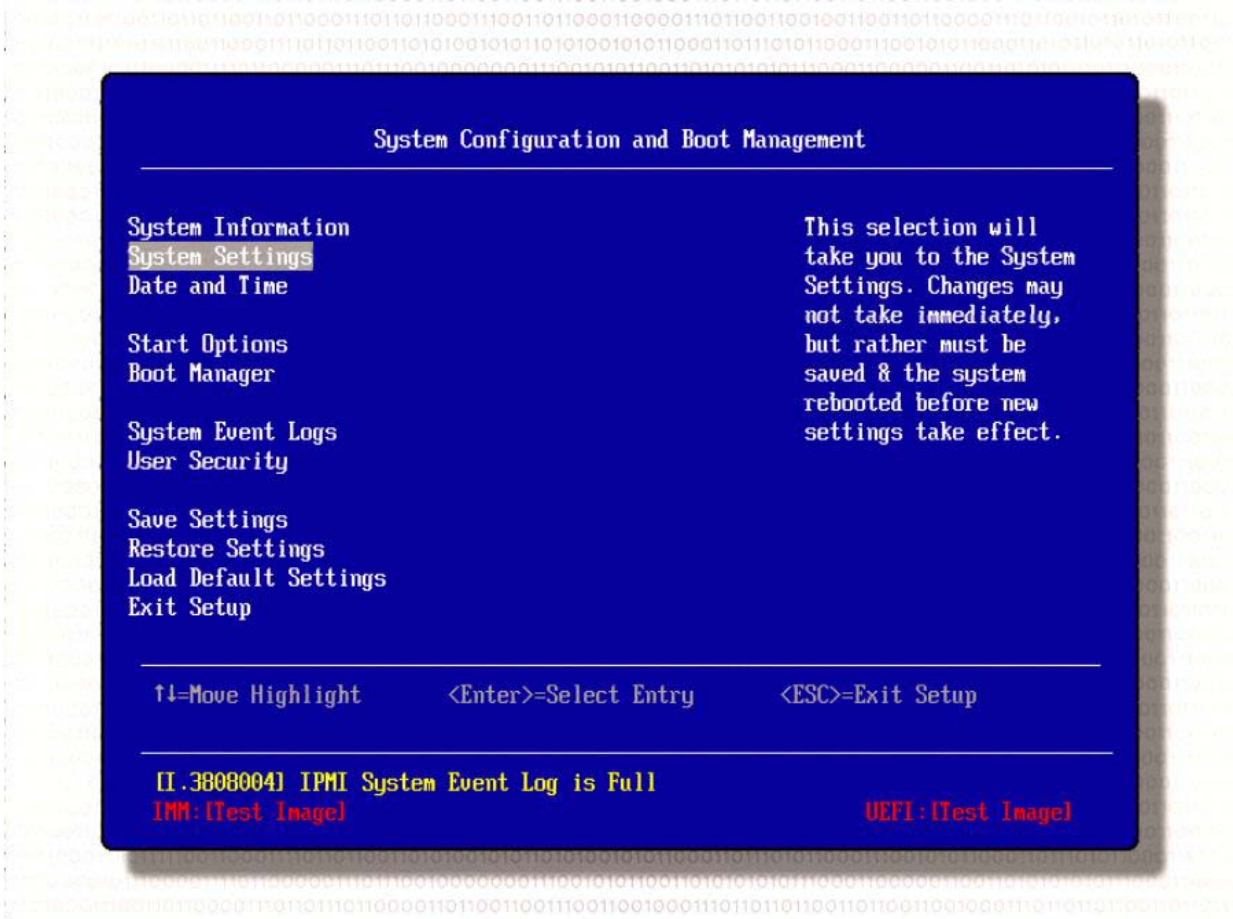

Select **System Settings** and press **Enter**, the System Settings Screen appears as shown in the following figure:

3. Highlight **System Settings** and press **Enter**.
   The System Settings screen appears, as shown in the following figure.

Figure 2 System Settings Screen

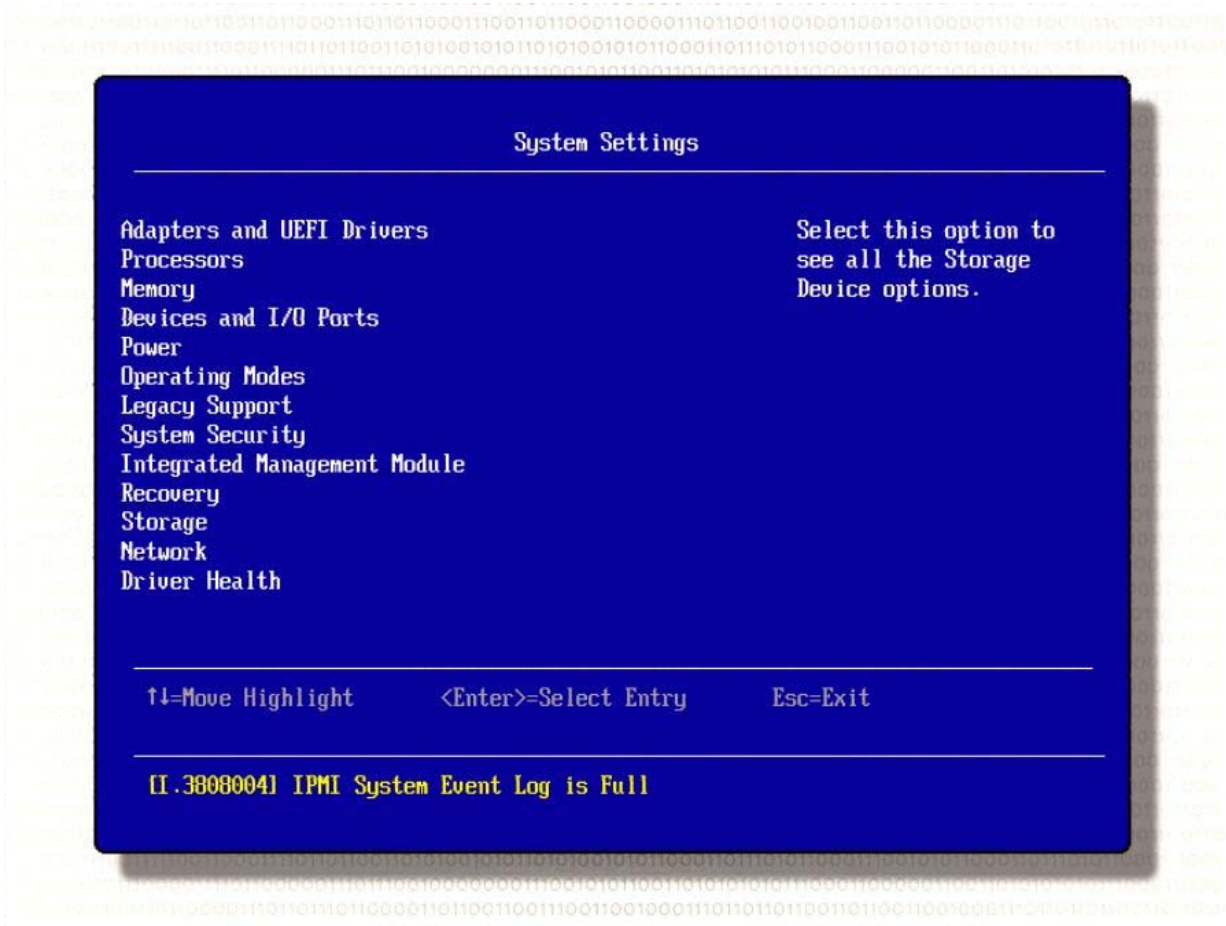

NOTE Some Lenovo systems have the ability to enable or disable the RAID C100/C105 controller and some do not. If your system has this ability, follow steps 4 to 6 to ensure the RAID controller is enabled. If your system does not allow the RAID controller to be disabled, you may skip steps 4 to 6

4. Highlight **Devices and I/O Ports** and press **Enter**.
   The Devices and I/O Ports screen appears.
5. Verify that the **Configure SATA as** field is set to **RAID** on the Devices and I/O Ports screen.
6. If this field is set to RAID, press **Esc** to return to the **System Settings** screen.
   If this field is not set to RAID, perform the following steps:
   a. Highlight the **Configure SATA as** options and press **Enter**.
      IDE, **AHCI**, and **RAID** appear as the options.
   b. Highlight **RAID** and press **Enter**.
   c. Press **Esc** until you return to the System Configuration and Boot Management screen.

d. Highlight **Save Settings** and press **Enter**.
e. Reboot your system.
f. Press **F1** during bootup to enter setup.
The System Configuration and Boot Management screen appears, as shown in the following figure.
g. Highlight **System Settings** and press **Enter**.

The System Settings screen appears. The Storage menu option is now available on this screen.

7. Highlight **Storage** on the System Settings screen and press **Enter**.
The Storage screen appears, as shown in the following figure.

Figure 3 Storage Screen

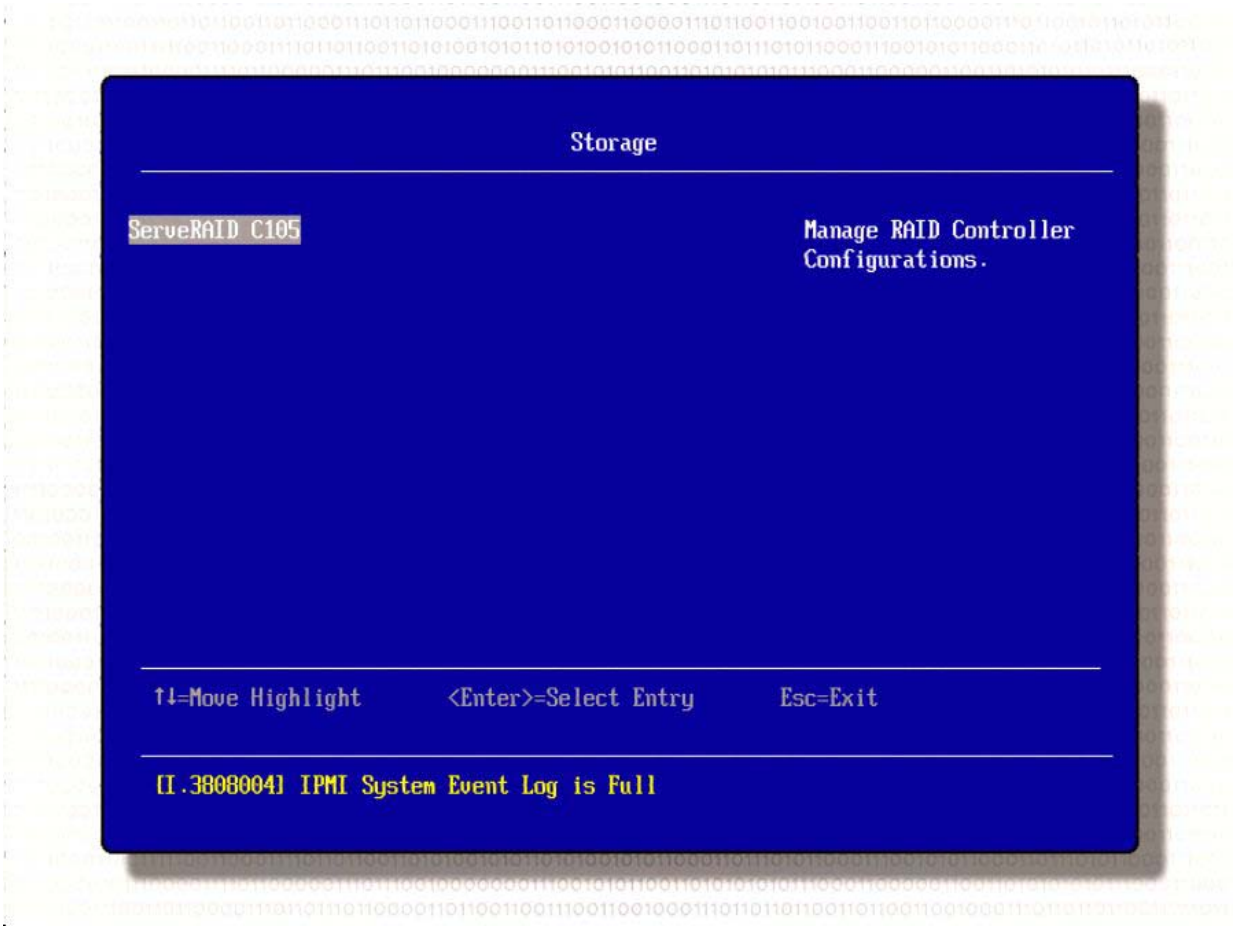

8. Highlight your controller and press **Enter**.
The Configuration Options screen appears, as shown in the following figure. This screen lists the following management options:

— **Controller Management**: Select this option to view and/or change the controller properties, save configurations, and perform other tasks. You can view information about the controller and the devices connected to it. For more information, see Section 3.2, Managing Controllers.
— **Virtual Drive Management**: Select this option to access the Virtual Drive Management screen, where you can create virtual drive configurations, change the virtual drive name, delete virtual drives, initialize drives, and perform other tasks. For more information, see Section 3.3, Managing Virtual Drives.

— **Drive Management**: Select this option to access the Drive Management screen, where you can view drive properties and perform tasks. For more information, see Section 3.4, Managing Drives.

Figure 4 Configuration Options Screen

3.2 Managing Controllers

This section explains how you can use the HII configuration utility to view and change the properties for controllers, and clear a configuration.

3.2.1 Viewing Controller Properties

The HII configuration utility displays information for one controller at a time.

Perform the following steps to view the controller properties.

1. Highlight **Controller Management** on the Configuration Options screen (Figure 4), and press **Enter**. The Controller Management screen appears, as shown in the following figure.

Figure 5 Controller Management Screen

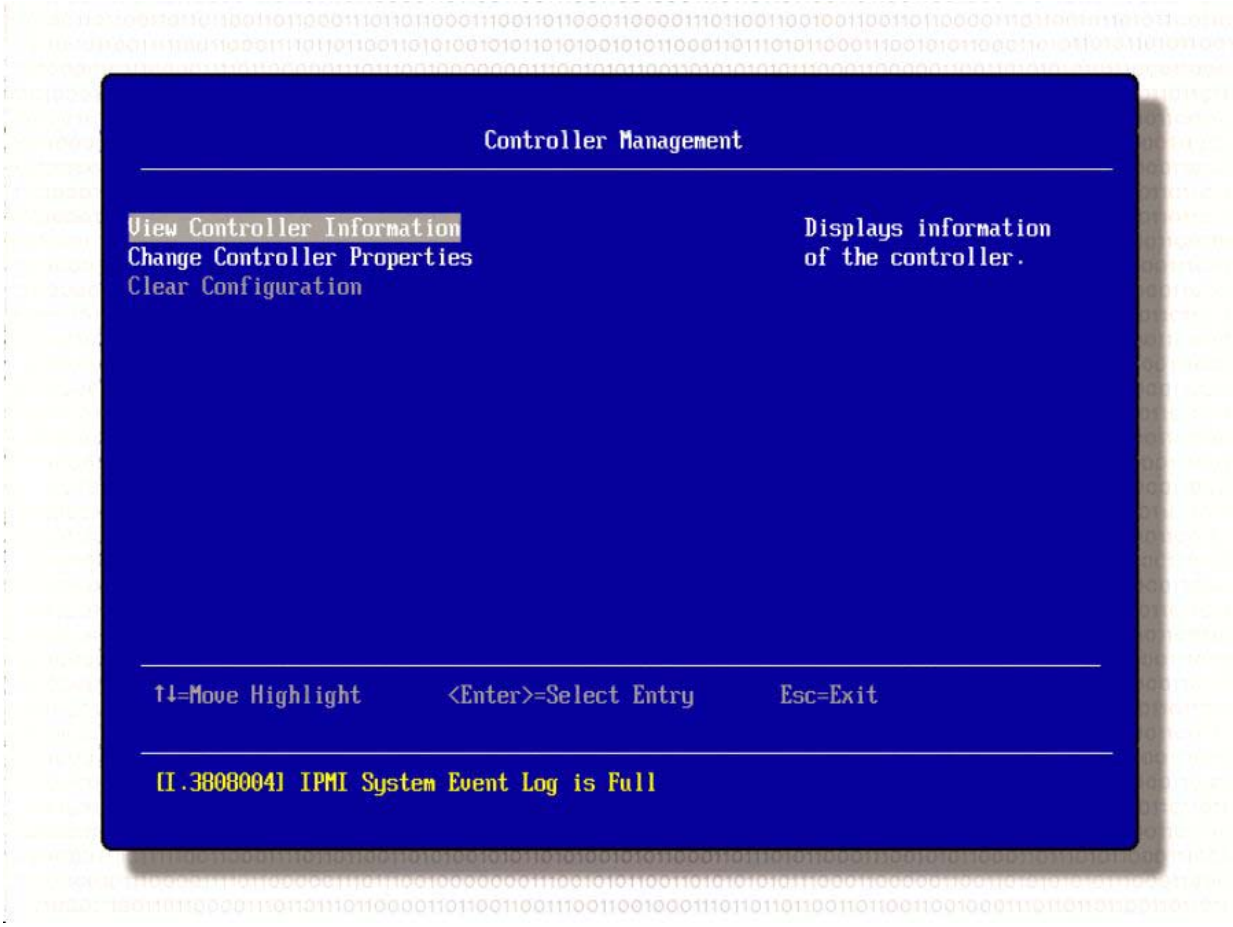

2. Highlight **View Controller Information** and press **Enter**.

   The Controller Management >> View Controller Information screen appears, as shown in the following figure.

Figure 6 Controller Management >> View Controller Information Screen

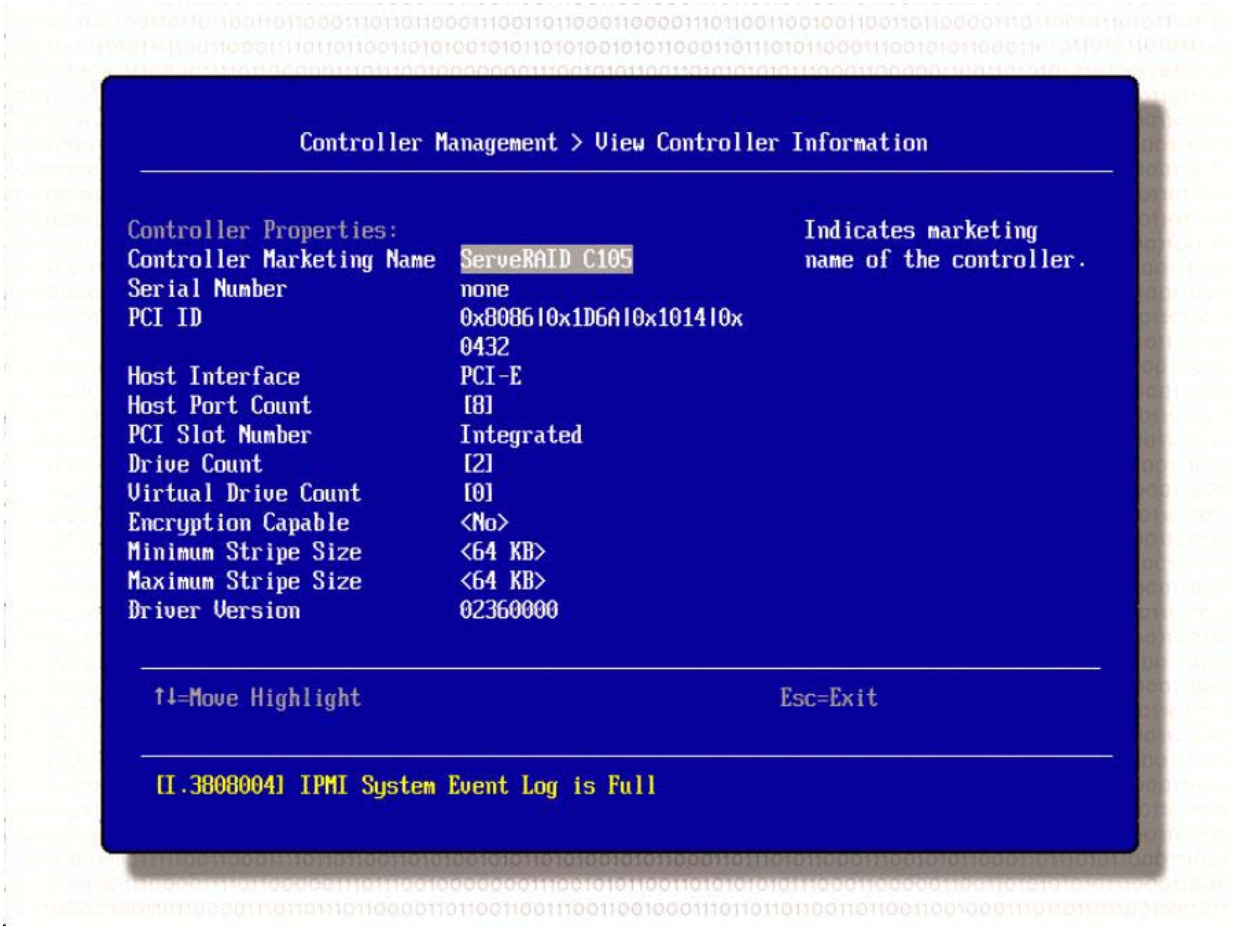

The information on this screen is read-only; it cannot be modified directly. Most of this information is self-explanatory, such as the number of virtual drives that are defined on this controller, and the number of drives connected to the controller.

The following table defines the controller properties.

Table 1 Controller Properties

Property	Description
Controller Marketing Name	ServeRAID C105
Serial Number	Indicates the manufacturer-assigned serial number.
PCI ID	The ID number for the Peripheral Component Interconnect local bus.
Host Interface	Indicates the type of interface used by the computer host system, such as PCI-E.
Host Port Count	Indicates the maximum number of ports supported by the software RAID controller in which devices (such as CD-ROM and disks) can be connected.
PCI Slot Number	N/A for software RAID
Drive Count	The number of drives connected to the selected controller.
Virtual Drive Count	The number of virtual drives supported by the selected controller.

Table 1 Controller Properties

Property	Description
Encryption Capable	N/A for software RAID
Minimum Stripe Size	64 KB is the only stripe size available.
Maximum Stripe Size	64 KB is the only stripe size available.
Driver Version	The version of the device driver.

3.2.2 Changing Controller Properties

You can use the HII configuration utility to change the properties for a controller. Perform the following steps to change information for a controller.

1. Highlight **Controller Management** on the Configuration Options screen (Figure 4) and press **Enter**.
   The Controller Management screen appears, as shown in the following figure.

Figure 7 Controller Management Screen

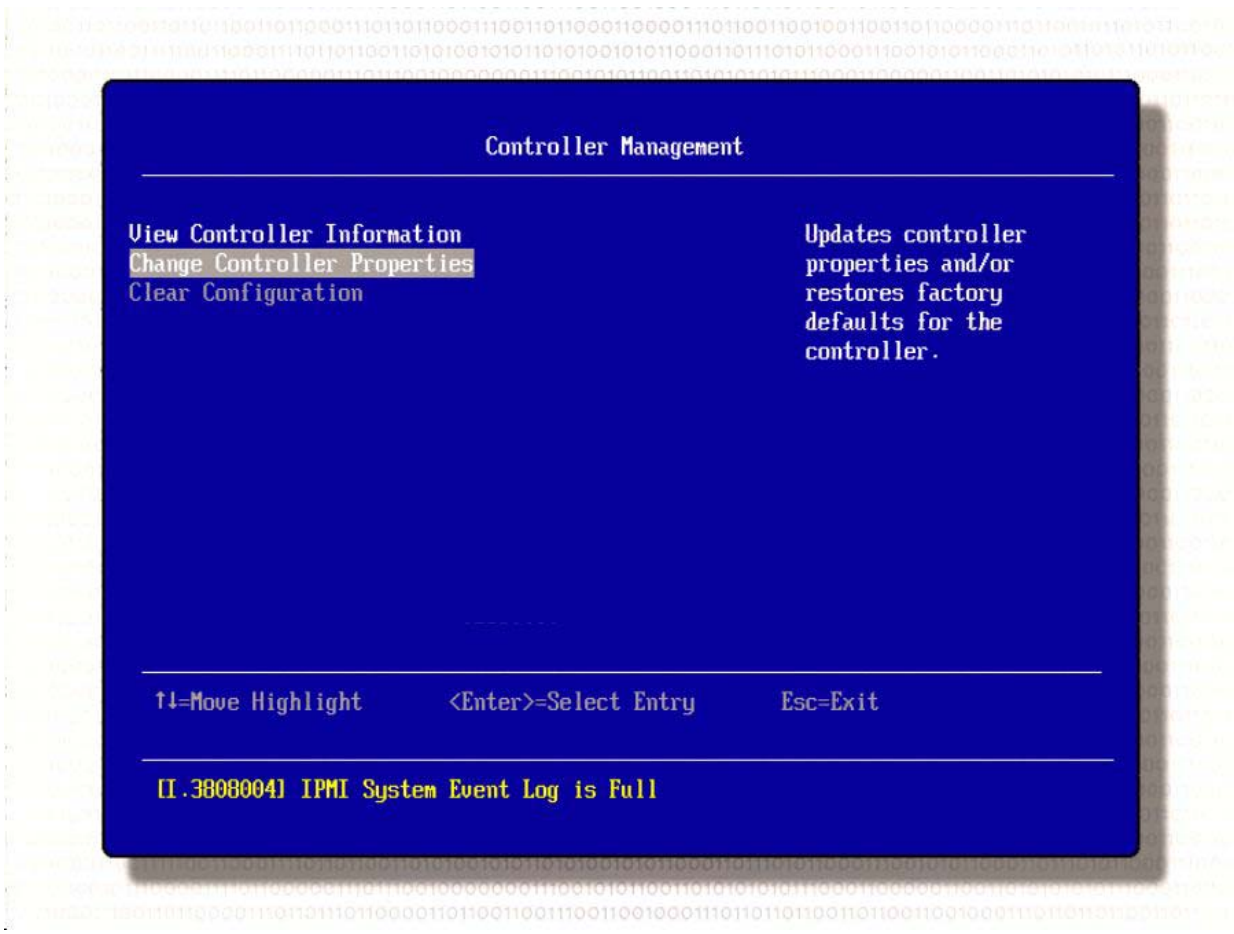

2. Highlight **Change Controller Properties** and press **Enter**.
   The Controller Management >> Change Controller Properties screen appears, as shown in the following figure.

Figure 8 Controller Management >> Change Controller Properties Screen

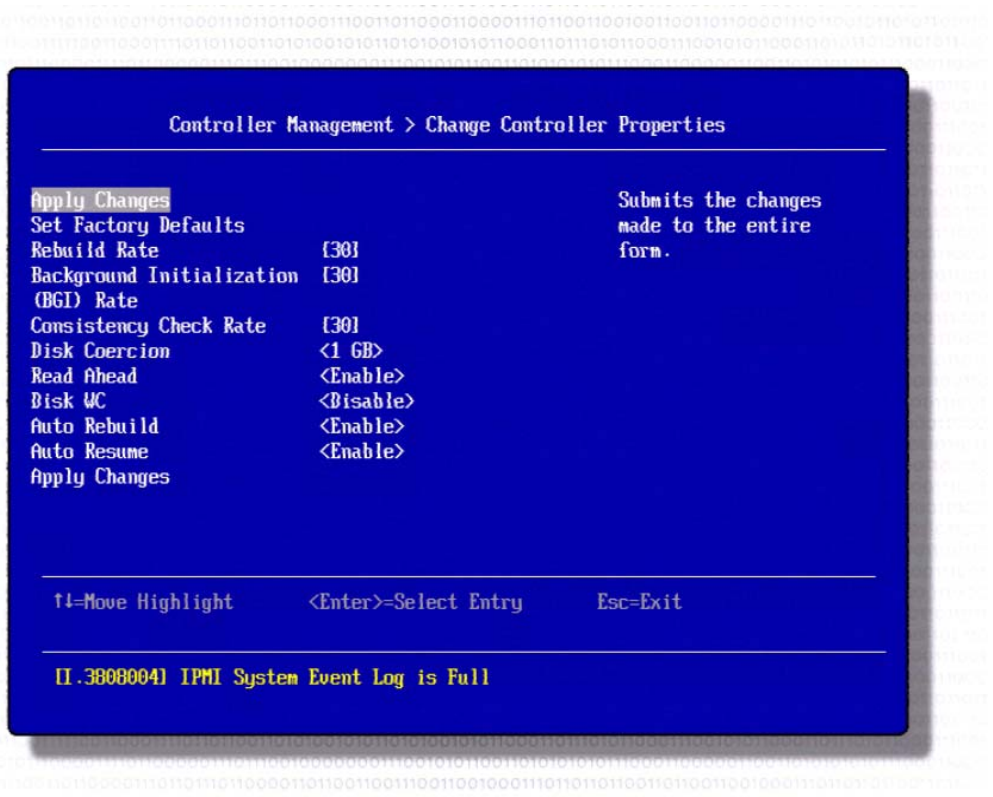

3. Change the following controller properties as desired.
   a. **Set Factory Defaults**: Use the arrow keys to move the cursor to this property and press **Enter**. On the Confirm screen, select **Yes** to confirm your selection.
   b. **Rebuild Rate**: Use the arrow keys to move the cursor to this property. Press the plus (+) key to increase the rate or the minus (-) key to decrease the rate.
   c. **Background Initialization (BGI) Rate**: Use the arrow keys to move the cursor to this property. Press the plus (+) key to increase the rate or the minus (-) key to decrease the rate.
   d. **Consistency Check Rate**: Use the arrow keys to move the cursor to this property. Press the plus (+) key to increase the rate or the minus (-) key to decrease the rate.
   e. **Disk Coercion**: Use the arrow keys to move the cursor to this property. Press the plus key (+) to set the value
   f. to **128 MB**, **1 GB** or **None**.
   g. **Read Ahead**: Use the arrow keys to **Enable** or **Disable**.
   h. **Disk WC**: Use the arrow keys to **Enable** or **Disable**.
   i. **Auto Rebuild**: Use the arrow keys to **Enable** or **Disable**.
   j. **Auto Resume**: Use the arrow keys to **Enable** or **Disable**.

The following table defines these controller properties.

Table 2 Changing Controller Properties

Property	Description
Set Factory Defaults	Resets factory default values for all of the controller properties.
Set Boot Devices	Selects the virtual drive to use as the boot device.
Rebuild Rate	The percentage of central processing unit (CPU) resources devoted to rebuilding data onto a new drive after a drive in a storage configuration has failed. The default value is 30 percent.
Background Initialization (BGI) Rate	Background initialization is a check for media errors on the drives when you create a virtual drive. It is an automatic operation that starts five minutes after you create the virtual drive. This check ensures that striped data segments are the same on all of the drives in the drive group. The default value is 30 percent.
Consistency Check Rate	A consistency check is an operation that verifies that all stripes in a virtual drive with a redundant RAID level are consistent and that automatically fixes any errors. The consistency check rate is the rate at which consistency check operations are run on a computer system. The default value is 30 percent.
Disk Coercion	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.
Read Ahead	When disk Read Ahead is enabled, 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.
Disk WC	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 enabled, a write transaction is considered to be complete when all the data has been written to the disk cache. When disk Write Cache is disabled, the write transaction is complete only when the data has been written to the disk.
Auto Rebuild	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.
Auto Resume	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.

4. Highlight **Apply Changes** and press **Enter** to register your changes.
   Your changes are registered and the controller properties are changed.

3.2.3 Clearing Configurations

Perform the following steps to clear all of the existing configurations on the controller.

1. Highlight **Controller Management** on the Configuration Options screen (Figure 4) and press **Enter**.
2. Highlight **Clear Configuration** on the Controller Management screen (Figure 5) and press **Enter**.
3. Highlight **Yes** to confirm your selection on the Confirm screen.
   This action clears the existing configurations.

3.3 Managing Virtual Drives

You can use the Virtual Drive Management screen to create virtual drive configurations, view and change the virtual drive properties, delete virtual drives, initialize drives, and perform other tasks related to the virtual drives.

3.3.1 Configuring Virtual Drives

This section provides detailed instructions for configuring drive groups and virtual drives with the HII configuration utility.

It is recommended that you use drives with the same capacity when you create a storage configuration. If you use drives with different capacities in one array, the configuration utility limits each drive to the capacity of the smallest drive.

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

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

NOTE The stripe size is read-only. You cannot change the settings for this property.

Follow these steps to access the Virtual Drive Management screen and create a virtual drive configuration.

1. Highlight **Virtual Drive Management** on the Configuration Options screen (Figure 4) and press **Enter**. The Virtual Drive Management screen appears, as shown in the following figure.

Figure 9 Virtual Drive Management Screen

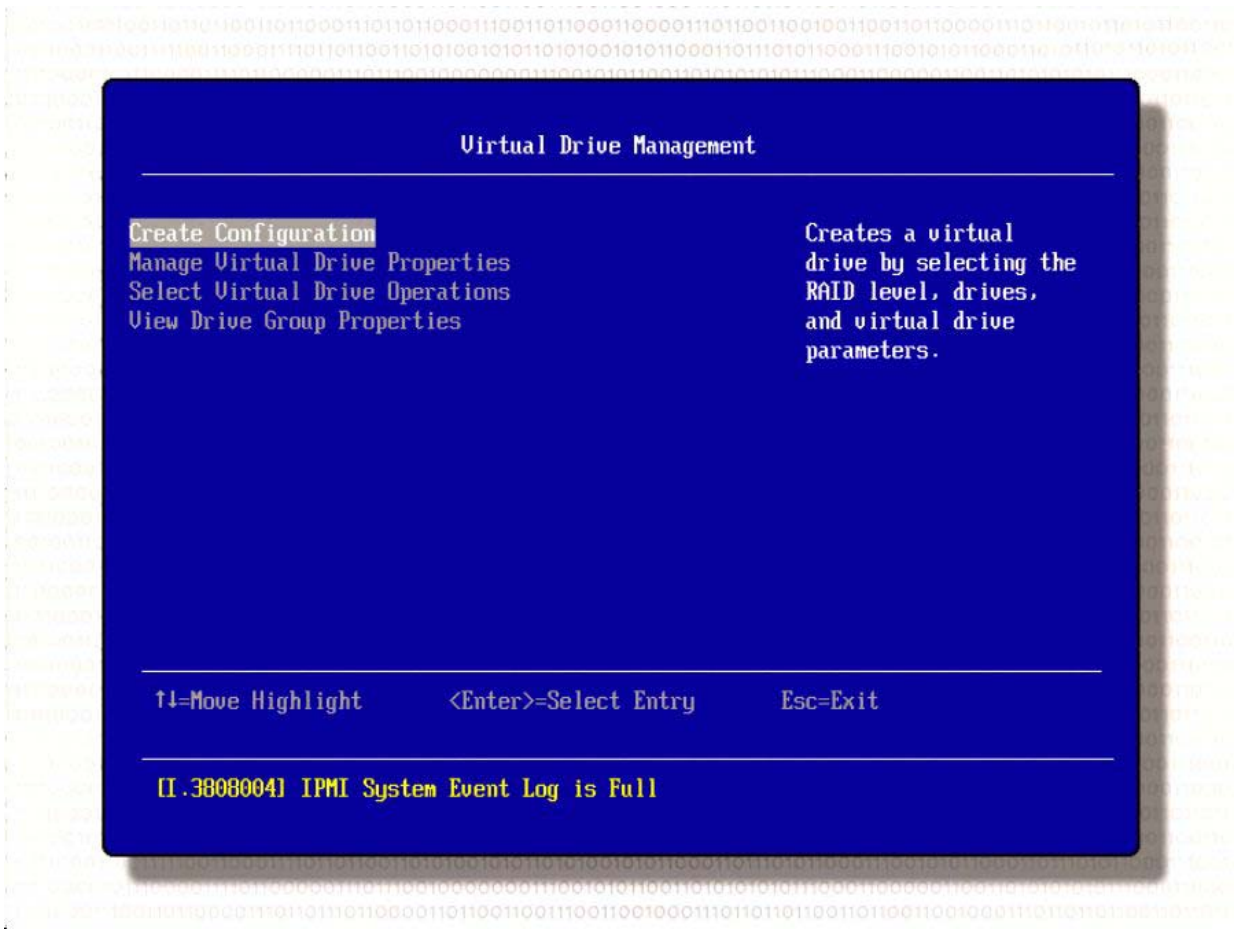

2. Highlight **Create Configuration** and press **Enter**.

   The Virtual Drive Management >> Create Configuration screen appears, as shown in the following figure. The default settings display in the fields.

Figure 10 Virtual Drive Management >> Create Configuration Screen

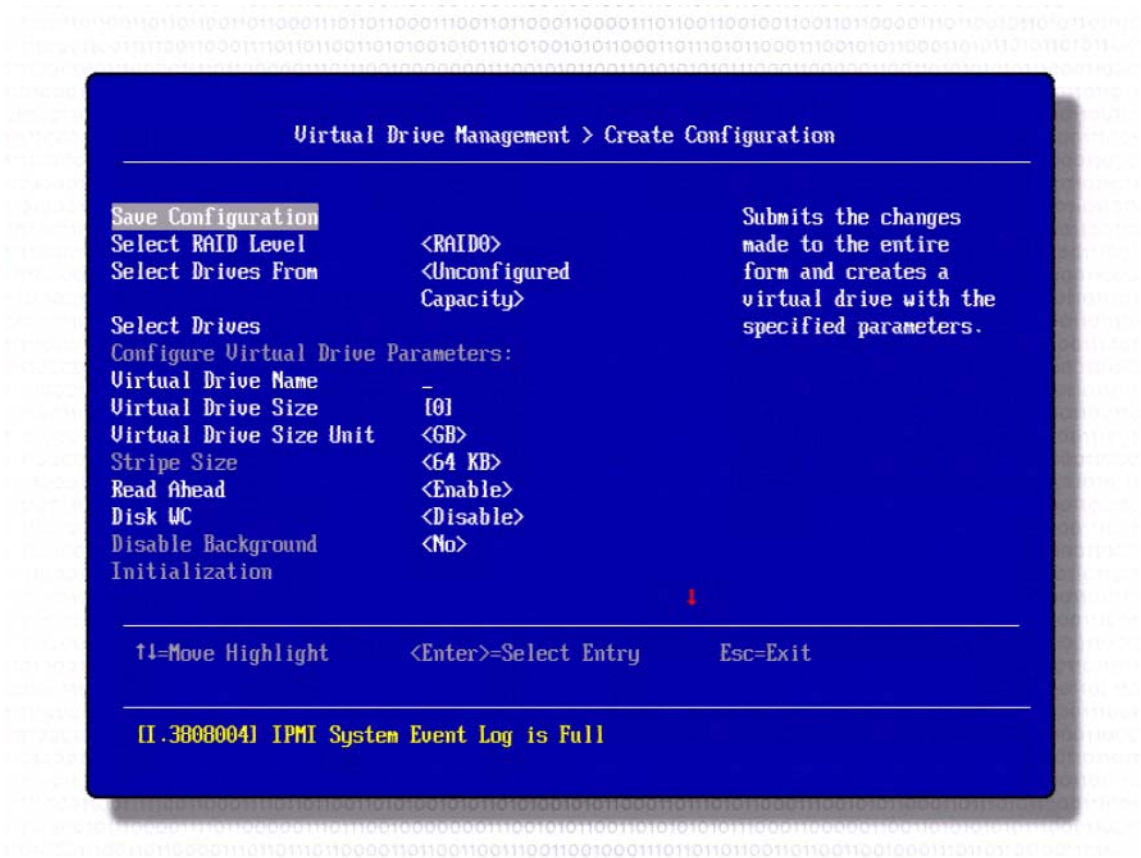

3. Use the arrow keys to select any highlighted fields (one at a time) that you want to change the setting for and press **Enter**.
4. Select the setting for each property that you want to change from the default.

   You can change the settings for the following fields on this screen:

   — **Select RAID Level**: The menu lists the possible RAID levels for the virtual drive. See Section 1.5, RAID Overview, for more information about RAID levels.

   — **Select Drives From**: The menu lists the sources that you can use to select drives for the virtual drive. The options are **Unconfigured Capacity** and **Free Capacity**.

   — **Select Drives**: Press this button and a screen appears that lists the Unconfigured Good drives or the free capacity, depending on the value you selected in the Select Drive From field.

NOTE If you exit and re-enter this screen, the previous selections will be reset and you will need to select the drives again.

   — **Virtual Drive Name**: Enter the name of the virtual drive.

   — **Virtual Drive Size**: Enter the capacity of the virtual drive. Normally, this value is the full capacity of the drives in the virtual drive.

— **Virtual Drive Size Unit**: Enter the unit of capacity you want to use for the virtual drive. The options are **MB**, **GB**, and **TB**.

— **Read Ahead –** When disk Read Ahead is enabled, 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.

— **Disk WC –** 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 enabled, a write transaction is considered to be complete when all the data has been written to the disk cache. When disk Write Cache is disabled, the write transaction is complete only when the data has been written to the disk.

— **Disable Background:** Use this option to select the amount of system resources dedicated to background initialization of virtual drives connected to the selected controller.

The following setting is read-only and cannot be changed.

— **Stripe Size:** A stripe consists of the data segments that the RAID controller writes across multiple drives, not including parity drives. The default is 64 KB.

3.3.2 Managing Virtual Drive Properties

After you create a virtual drive, you can use the Virtual Drive Management screen to change the name of the virtual drive.

NOTE You can change only the virtual drive name for the virtual drive that is selected. The other virtual drive properties are read-only.

Perform the following steps to access the virtual drive properties screen and change the name of the virtual drive.

1. Highlight **Virtual Drive Management** on the Configuration Options screen (Figure 4) and press **Enter**. The Virtual Drive Management screen appears, as shown in the following figure.

Figure 11 Virtual Drive Management Screen

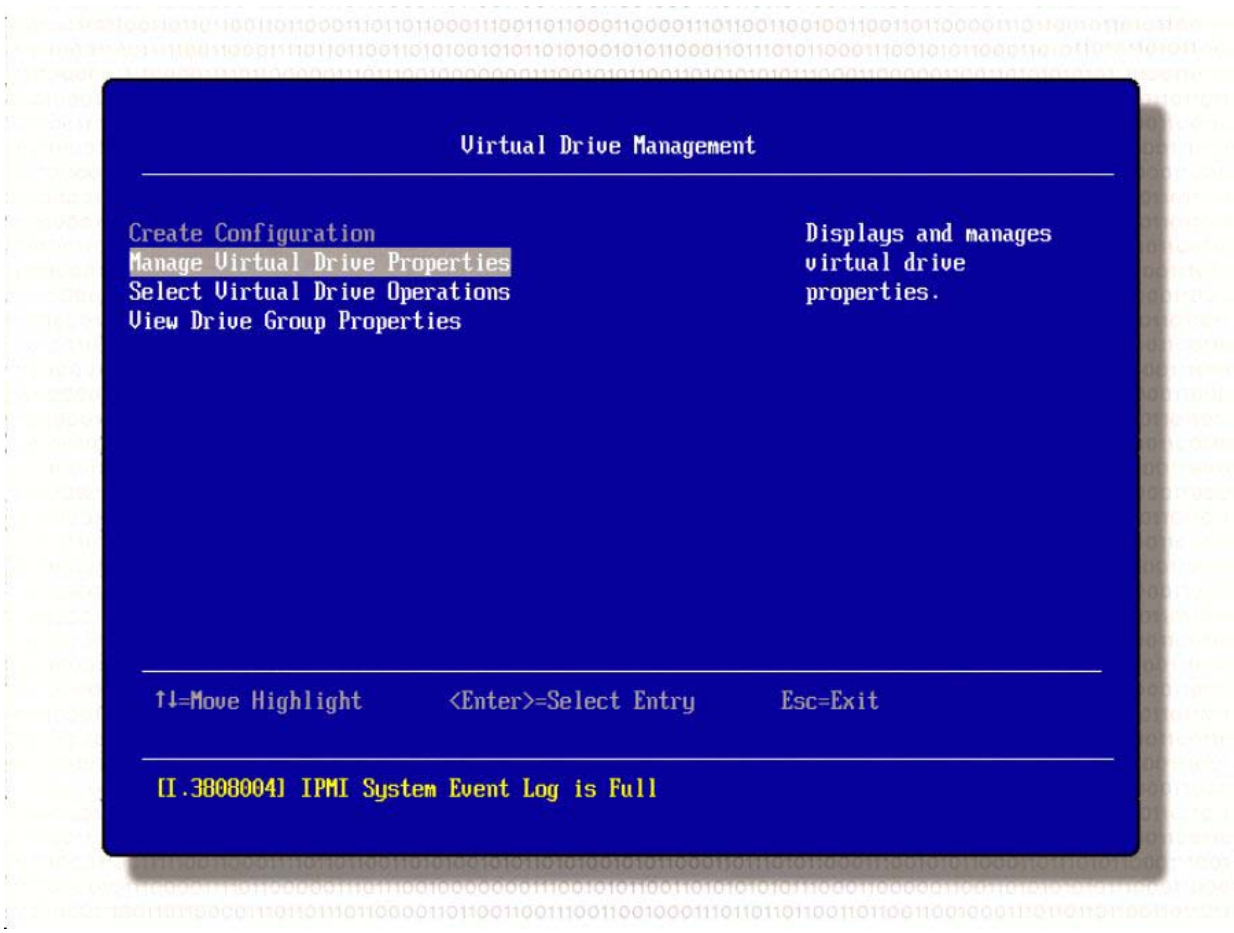

2. Highlight **Manage Virtual Drive Properties** and press **Enter**.
   The Virtual Drive Management >> Manage Virtual Drive Properties screen appears, as shown in the following figure.

Figure 12 Virtual Drive Management >> Manage Virtual Drive Properties Screen

3. Highlight the **Virtual Drive Name** field and press **Enter**.
4. Enter a different name for the virtual drive and press **Enter**.
5. Highlight **Apply Changes** and press **Enter** to make the selected change.
   The name of the virtual drive is changed.

3.3.2.1 Viewing Associated Drives

Perform the following steps to view associated drives for a virtual drive properties:

1. Highlight the **View Associated Drives** in the **Manage Virtual Drive Properties** menu and Press **Enter**.
   The drives associated with the current virtual drive appears, as shown in the following figure.

Figure 13 Associated Drives for a Virtual Drive

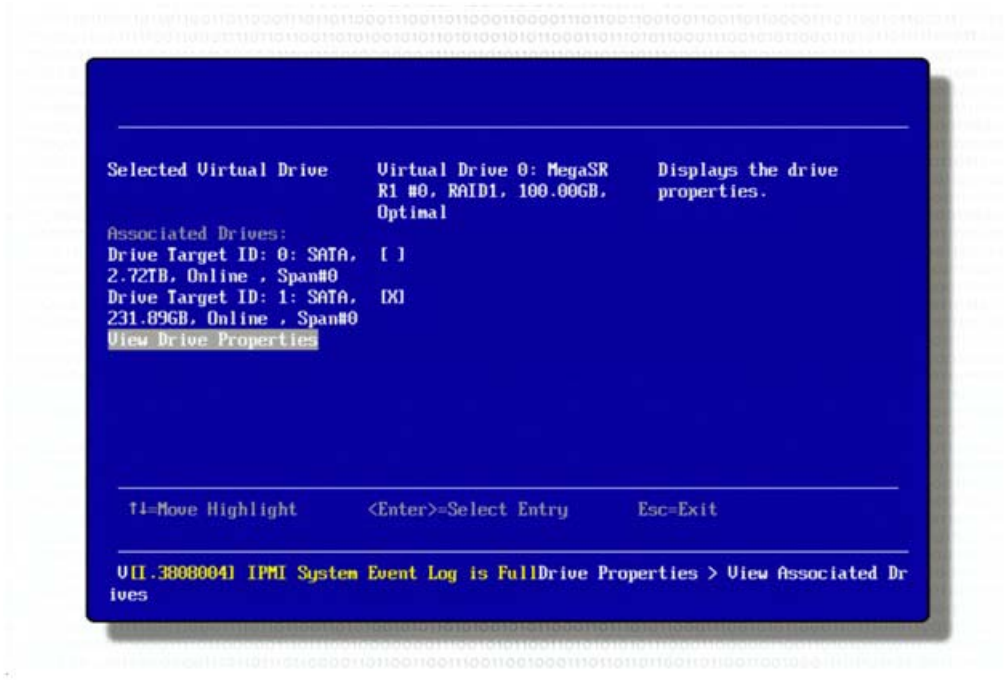

2. Highlight **View Drive Properties** and Press **Enter**.
   The Drive properties page for the physical drive appears.

3.3.3 Selecting Virtual Drive Operations

You can use the Virtual Drive Management screen to select and launch operations related to the virtual drive properties.

Perform the following steps to access the virtual drive operations screen and carry out operations for the virtual drive properties.

1. Highlight **Virtual Drive Management** on the Configuration Options screen (Figure 4) and press **Enter**.

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

Figure 14 Virtual Drive Management Screen

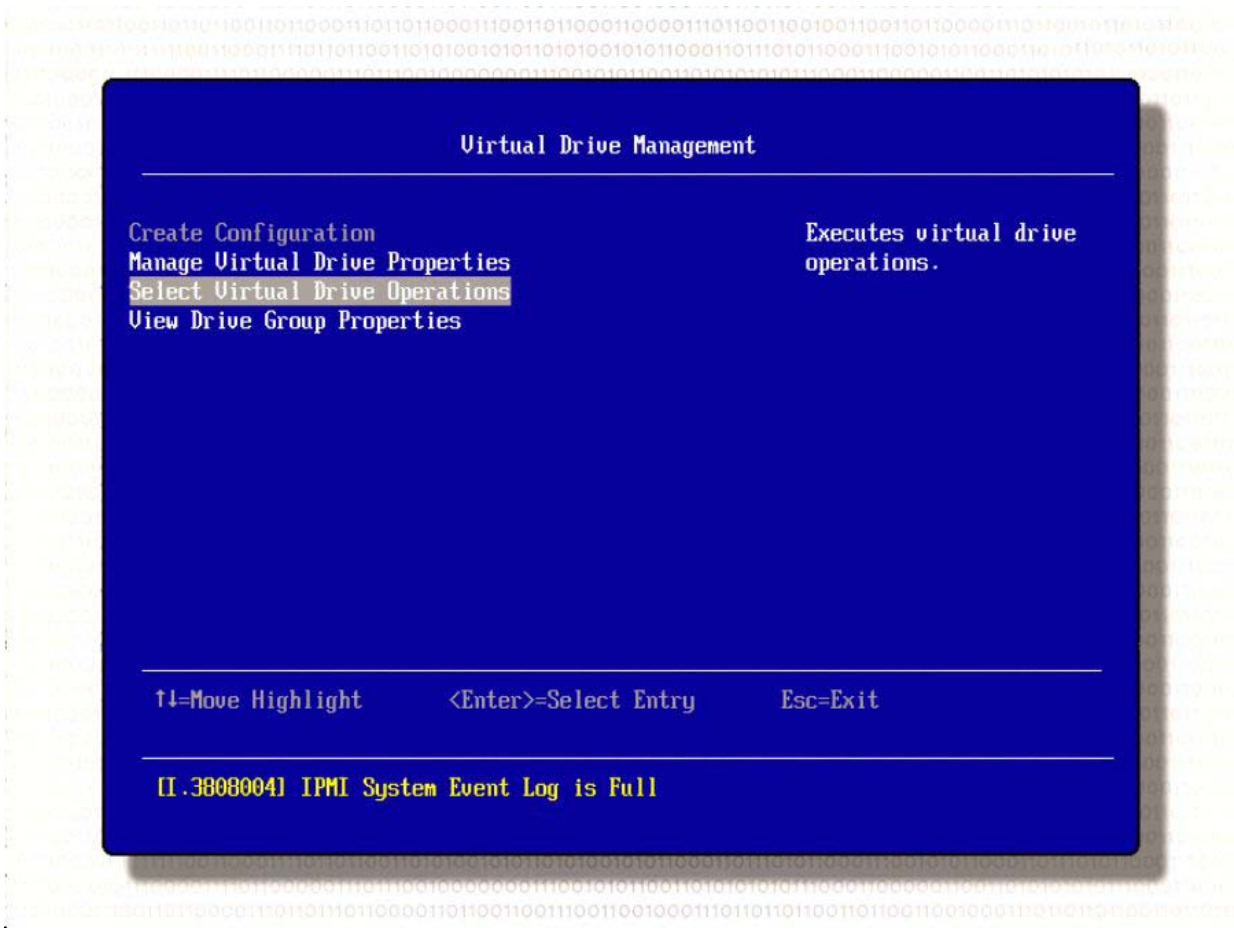

2. Highlight **Select Virtual Drive Operations** and press **Enter**.

   The Virtual Drive Management >> Select Virtual Drive Operations screen appears, as shown in the following figures.

Figure 15 Select Virtual Drive Operations Screen

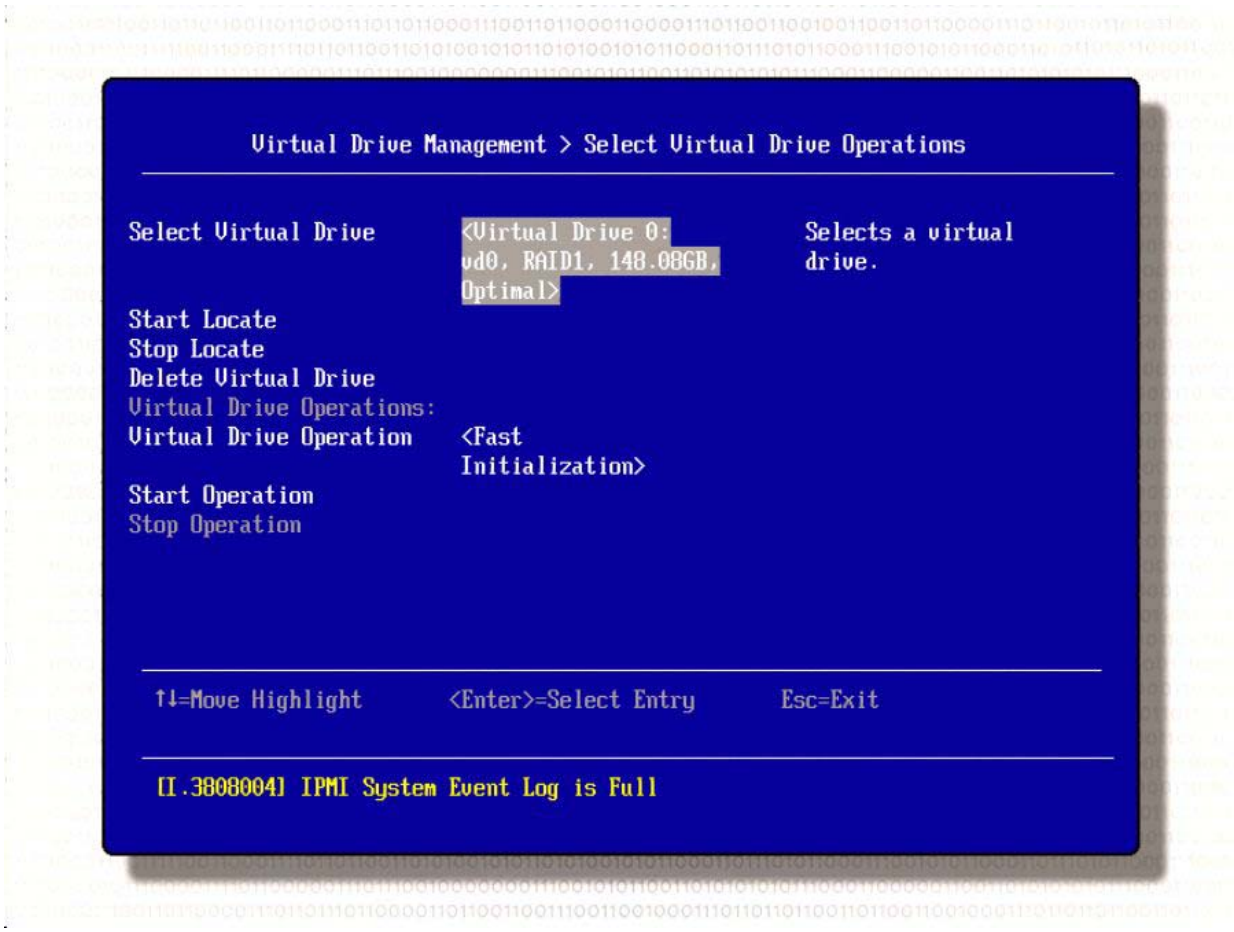

3. Highlight each operation you want to carry out and press **Enter**.

   The options for the virtual drive operations appear. Here are brief explanations of the virtual drive operations:

   — **Select Virtual Drive**: Select the virtual drive you want to carry out operations on. The screen displays the virtual drive number, virtual drive name, RAID level, virtual drive capacity, and virtual drive status.

   — **Start Locate**: Flash the LEDs on all the drives associated with the virtual drive.

   — **Stop Locate**: Stop flashing the LEDs on the drives associated to the virtual drive.

   — **Virtual Drive Operation**: Select **Fast Initialization** or **Full Initialization** to initialize this virtual drive. A fast initialization quickly writes zeroes to the first and last 10-MB regions of the new virtual drive and then completes the initialization in the background. A full initialization is not complete until the entire virtual drive has been initialized with zeroes. It is seldom necessary to use this option, because the virtual drive was initialized right after you created it.

   — **Consistency Checks**: If you have redundant virtual drives (i.e., virtual drives that use RAID 1 or 10), you have the option to perform the **Check Consistency** operation. The consistency check operation corrects inconsistencies between a protected volume and its replica

4. Highlight **Start Operation** and press **Enter**.
   The confirmation screen appears.

> **NOTE** The Start Operation option and the Stop Operation option toggle based on the current status of the operation. For example, after you start an operation, the Start Operation field is no longer highlighted, and the Stop Operation field is highlighted.

5. Highlight **Yes** and press **Enter** to confirm that you want to carry out the operations.

> **NOTE** After you choose an operation and select **Start Operation**, the progress bar appears, showing zero percent. To see the change in the percentage of progress, return to the previous page and then access this page again.

3.3.4 Viewing Drive Group Properties

You can use the Virtual Drive Management screen to access the drive group properties screen and view the properties. The drive group contains the drives that are used to create the virtual drive.

Perform the following steps to access the drive group properties screen and view the drive group properties.

1. Highlight **Virtual Drive Management** on the Configuration Options screen (Figure 4) and press **Enter**.
   The Virtual Drive Management screen appears, as shown in the following figure.

Figure 16 Virtual Drive Management Screen

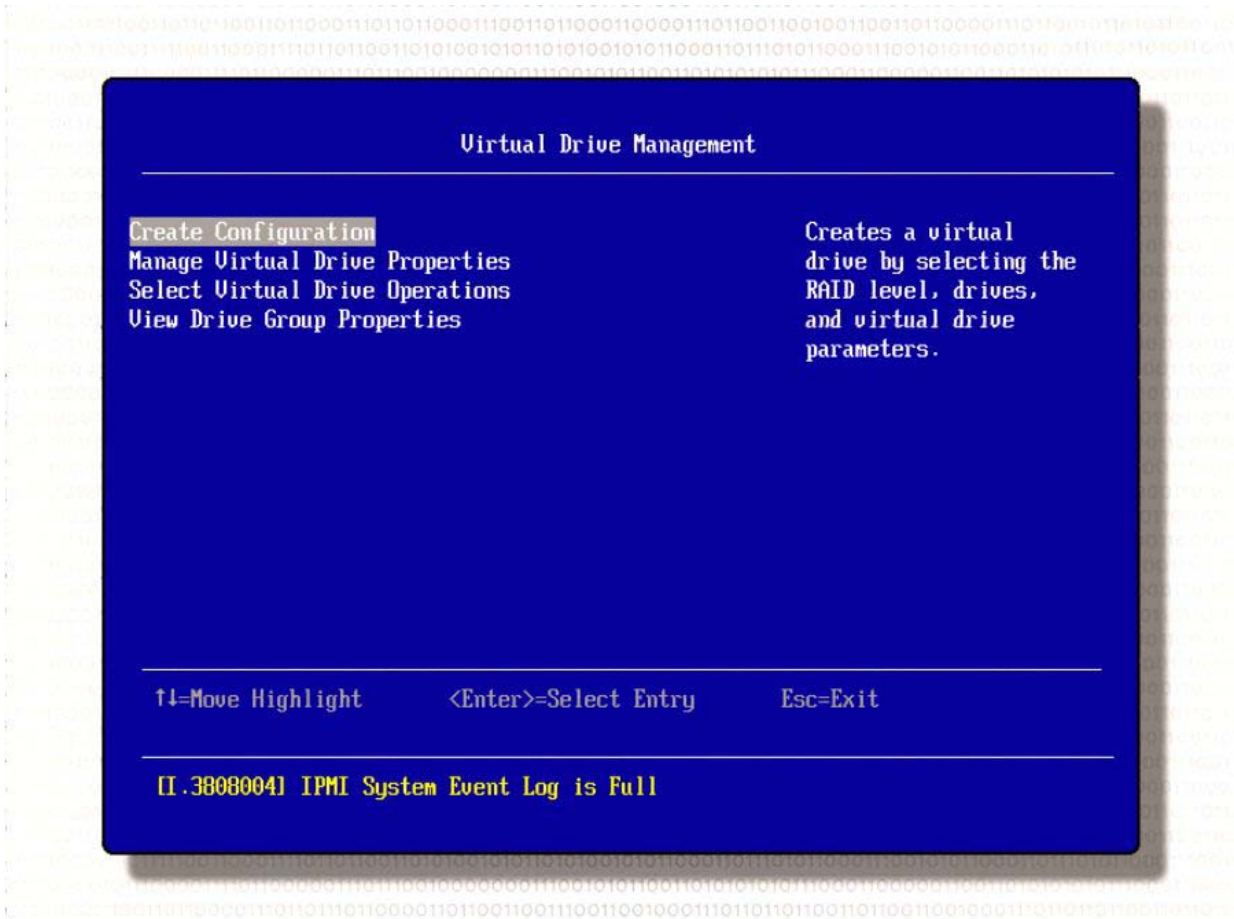

2. Highlight **View Drive Group Properties** and press **Enter**.
   The Virtual Drive Management >> View Drive Group Properties screen appears, as shown in the following figure.

Figure 17 View Drive Group Properties Screen

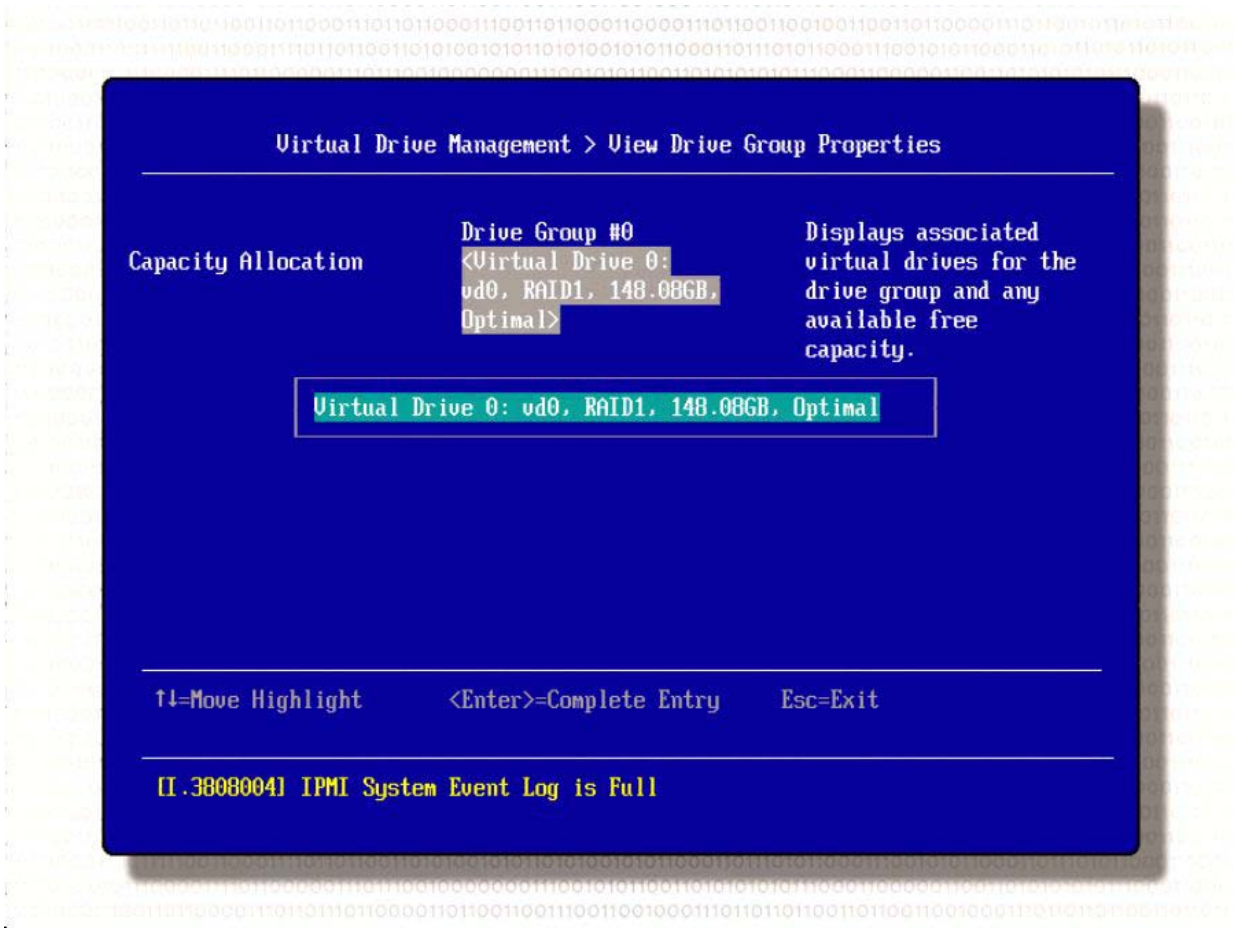

The screen displays the drive group number, drive group name, RAID level, virtual drive capacity, and virtual drive status.

3. Highlight the drive group and press **Enter**.
   The drive group properties and any free capacity appear in the list, as shown in the following figure.

Figure 18 Drive Group Properties and Free Space

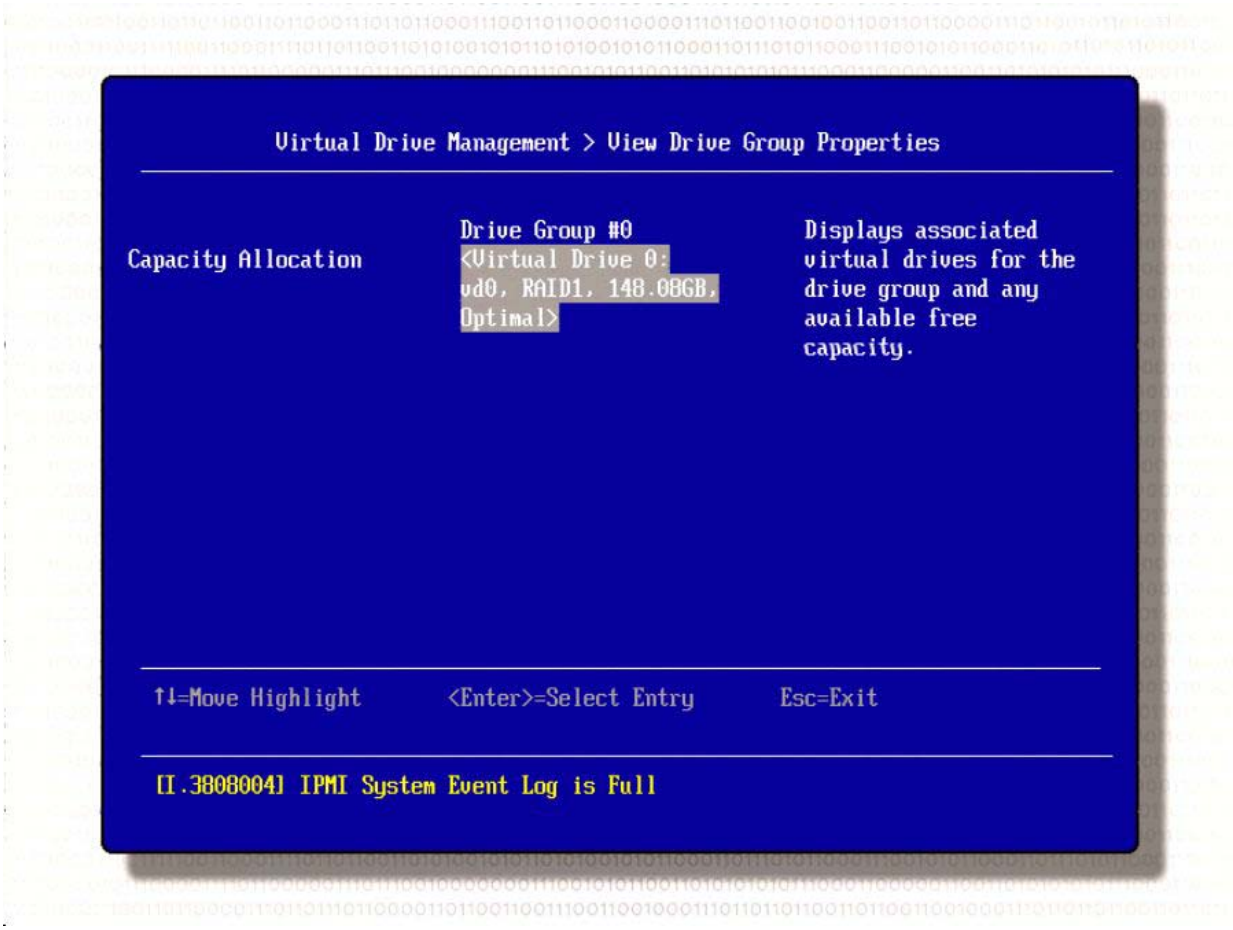

3.4 Managing Drives

You can use the Drive Management screen to view and change the physical drive properties of the drives in a drive group.

3.4.1 Viewing Drive Properties

Perform the following steps to access the drive properties screen and view the values for the properties.

1. Highlight **Drive Management** on the Configuration Options screen (Figure 4) and press **Enter**.

The Drive Management screen appears, as shown in Figure 19.

Figure 19 Drive Management Screen

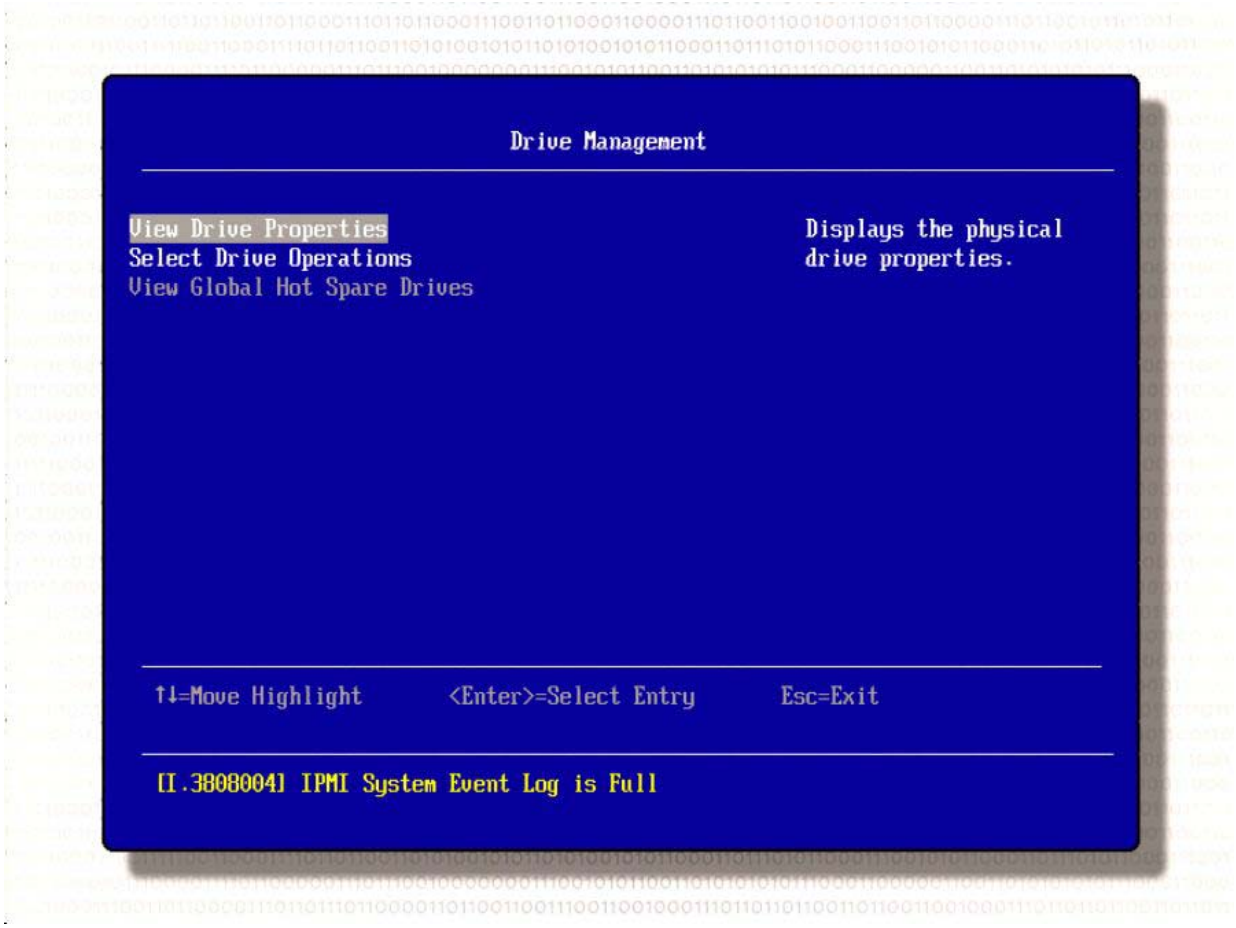

2. Highlight **View Drive Properties** and press **Enter**.
   The Drive Management >> View Drive Properties screen appears, as shown in the following figures.

NOTE The following two figures show the drive properties. Scroll down on Figure 21 to see the additional properties on the View Drive Properties screen.

Figure 20 Drive Management >> View Drive Properties Screen

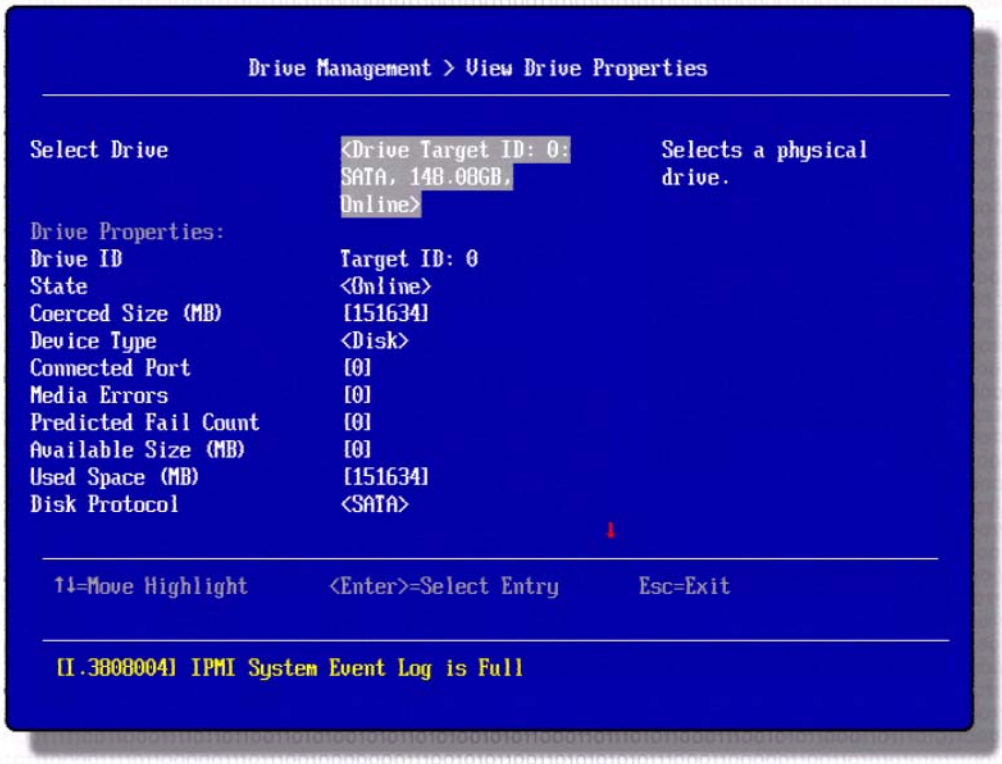

Figure 21 Drive Management >> View Drive Properties Screen (Continued)

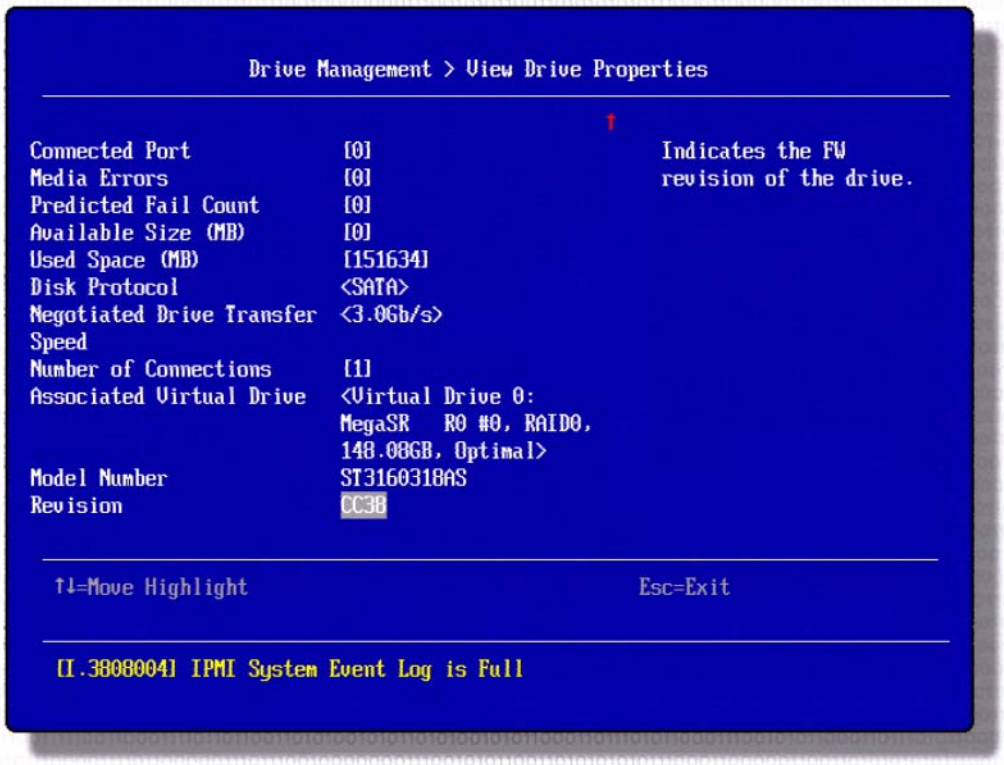

You can highlight the Associated Virtual Drive field and press **Enter** to view the properties. However, you cannot change the values in this field.

Here are brief explanations of the drive properties and their values:

- **Select Drive**: The drive number, drive slot, drive type, drive capacity, and drive status of the selected drive.
- **Drive ID**: Indicates the ID and drive slot of the selected drive.
- **State**: Indicates the state of the selected drive, such as Online, Offline, or Unconfigured Good.
- **Coerced Size (MB)**: Indicates the capacity to which the selected drive has been coerced to make it compatible with other drives that are nominally the same capacity.
- **Device Type**: Indicates the type of device selected, such as a drive, tape, or CD/DVD-ROM.
- **Connected Por**t: Indicates the port that the selected drive is connected to.
- **Media Errors**: Indicates the number of media errors on the drive. Media errors are physical defects on the drive.
- **Predicted Fail Count**: Indicates the predicted number of drive failures.
- **Available Size (MB)**: Indicates the free capacity of the selected drive.
- **Used Space (MB)**: Indicates the configured space of the selected drive in MB.
- **Disk Protocol**: Indicates the type of drive selected, such as SATA.
- **Negotiated Drive Transfer Speed**: Indicates the negotiated link speed for the data transfer to or from the selected drive.
- **Number of Connections**: Indicates the number of devices connected.
- **Associated Virtual Drive**: Indicates the virtual drive number, virtual drive name, RAID level, virtual drive capacity, and virtual drive status.

— **Model Number**: Indicates the Model number of the selected drive.
— **Revision:** The firmware revision of the selected drive.

3.4.2 Selecting Drive Operations

Perform the following steps to access the drive operations screen and carry out the operations.

1. Highlight **Drive Management** on the Configuration Options screen (Figure 4) and press **Enter**. The Drive Management >> Select screen appears, as shown in the following figure.

Figure 22 Drive Management Screen

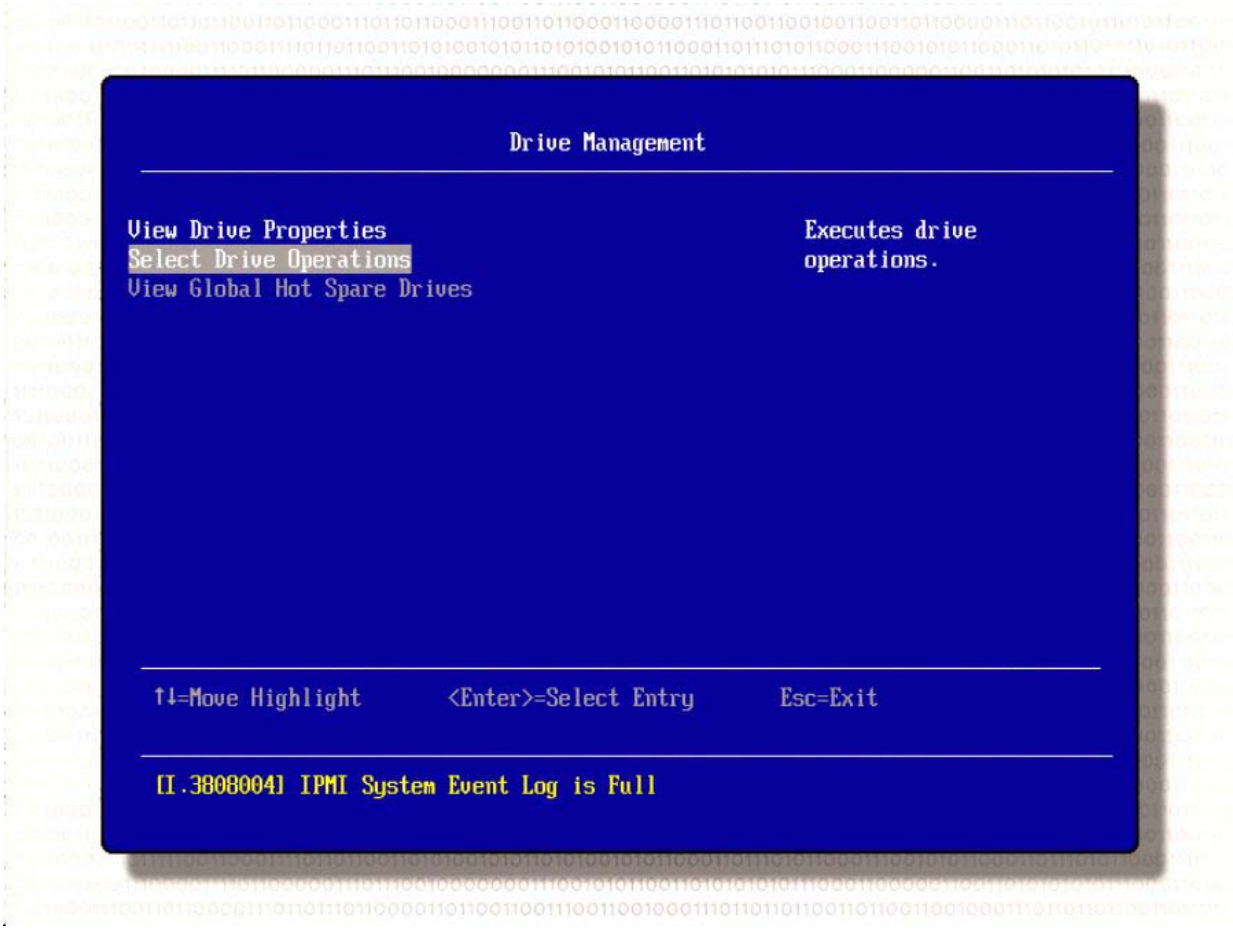

2. Highlight **Select Drive Operations** and press **Enter**.

The Drive Management >> Select Drive Operations screen appears, as shown in the following figure.

Figure 23 Drive Management >> Select Drive Operations Screen

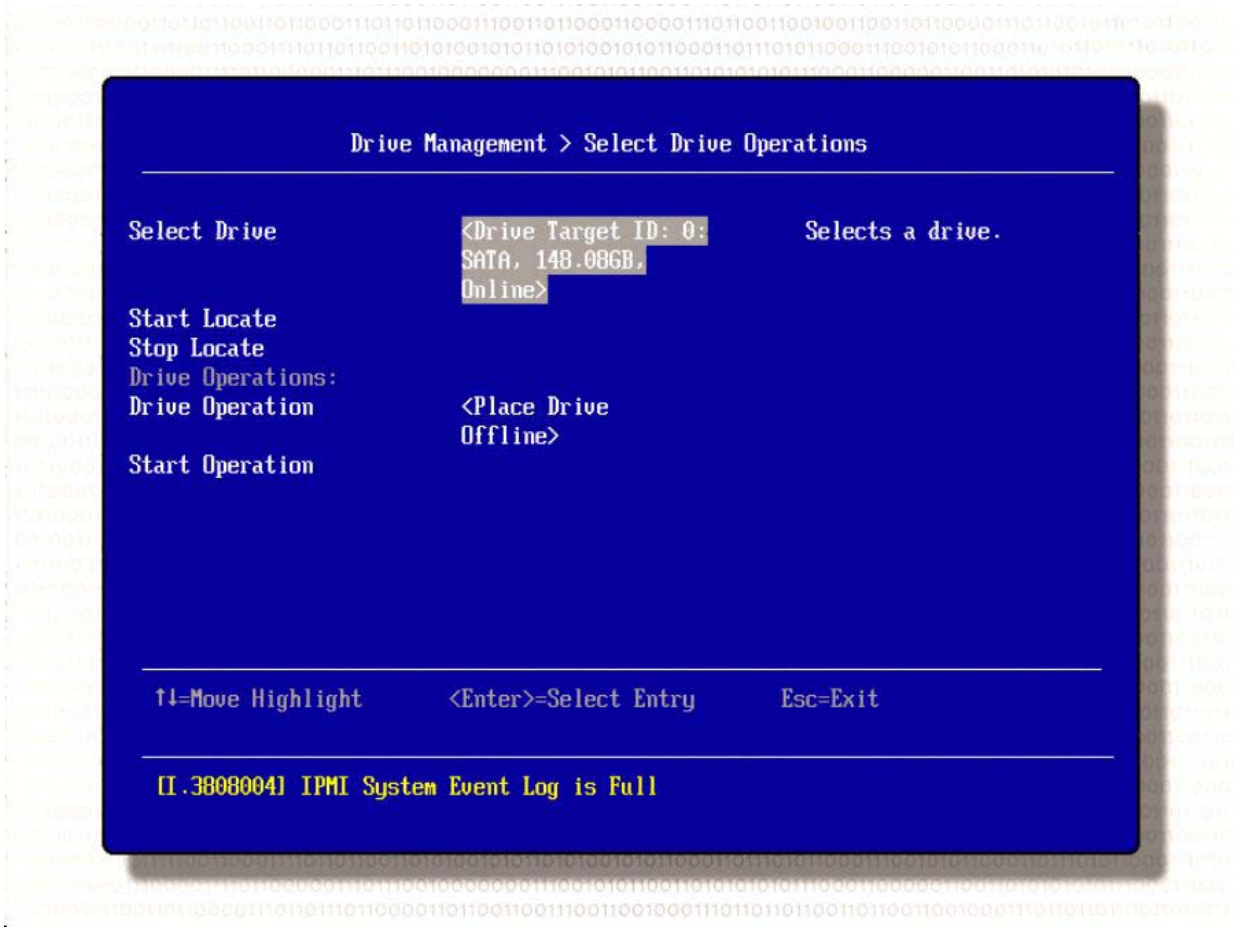

3. Highlight an operation and press **Enter**.

   The following items describe the drive operations:

   — **Select Drive**: Select the drive you want to carry out operations on. The screen displays the drive number, slot number, drive type, drive capacity, and drive status.
   — **Start Locate**: Flash the LED on the selected drive.
   — **Stop Locate**: Stop flashing the LED on the selected drive.
   — **Drive Operation**: Make a drive offline or online. The options are **Place Drive Offline** and **Place Drive Online**.

   NOTE If a good drive is part of a redundant drive group that includes a hotspare, and you force the drive offline, the data on the drive is rebuilt on the hotspare drive. The good drive that you forced offline goes into the *Unconfigured Bad* state.

4. Highlight **Start Operation** and press **Enter**.

   The confirmation screen appears.

5. Highlight **Yes** and press **Enter**.

   The controller performs the selected operations.

3.4.3 Viewing Global Hotspare Drives

Highlight **View Global Hot Spare** and press **Enter** to display all the assigned global hot spare on the RAID controller. A hot spare disk is a disk that replaces a failing or failed disk in a RAID configuration. The replacement is automatic or manually, depending in the hot spare policy,

Chapter 4: StorCLI Command Line Interface

The Storage Command Line Interface (StorCLI) utility is a command line interface application that you can use to configure and maintain storage configurations created with ServeRAID C100 and ServeRAID C105.

NOTE The StorCLI utility runs in the Microsoft® Windows® operating system environment and the Linux® operating system environment.

NOTE The StorCLI utility supports ServeRAID products. However, some of the commands might not be supported by ServeRAID C100 and ServeRAID C105.

4.1 Support for MegaCLI Commands

The MegaCLI commands can be executed on the StorCLI tool. A single binary is output for the StorCLI commands and its equivalent MegaCLI commands. See Appendix A: MegaCLI Commands to StorCLI Command Conversion for the information for conversion from MegaCLI commands to StorCLI commands.

4.2 StorCLI Overview

You can use the StorCLI utility to perform the following tasks:

- Configure virtual drives and create configurations on the controller
- Display the configuration on the controller
- Display and change virtual drive properties on the controller
- Display and change physical drive properties on the controller
- Set the status of physical drives
- Display and change controller properties
- Initialize drives
- Save an controller configuration to a file
- View, import, or delete foreign configurations
- Manage the security key for the controller
- 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
- Display controller inventory data in a single command
- Exit with predefined success or failure exit codes
- Set some predefined environment variables, such as number of controllers and the number of virtual drives
- Display help on how to use the command line options

4.3 StorCLI Command Syntax

This chapter describes the StorCLI command syntax and the valid values for each parameter in the general command syntax.

NOTE To get the output in JSON format, add `J` at the end of the command syntax.
Example: `storcli /cx show <property1>|<property2> J`

NOTEJ SON format output is not supported in the EFI operating system. The EFI platform ignores the `J` when it is added at the end of the command syntax.

NOTE Background operations are blocked in the EFI and HII environments and these operations are resumed in the operating system environments.

The StorCLI syntax uses the following general format:

`<[object identifier]> <verb> <[adverb | attributes | properties] > <[key=value]>`

The StorCLI tool supports the object identifiers listed in the following table.

Table 1 Object Identifiers in the StorCli Command Syntax

Object Identifier	Description
No object identifier specified	If there is no object identifier, the command is a system command.
/c*x*	This object identifier is for controller *x*.
/c*x*/v*x*	This object identifier is for a virtual drive *x* on controller *x*.
/c*x*/vall	This object identifier is for all virtual drives on controller *x*.
/c*x*/e*x*	This object identifier is for an enclosure x on controller *x*.
/c*x*/eall	This object identifier is for all enclosures on controller *x*.
/c*x*/f*x*	This object identifier is for a foreign configuration x on controller *x*.
/c*x*/fall	This object identifier is for all foreign configurations on controller *x*.
/c*x*/e*x*/s*x*	This object identifier is for the drive is slot *x* on enclosure *x* on controller *x*.
/c*x*/s*x*	This object identifier represents the drives that are directly attached to controller *x*.
/c*x*/e*x*/sall	This object identifier is for all the drives on enclosure *x* on controller *x*.
/c*x*/d*x*	This object identifier is for the drive group *x* on enclosure *x* on controller *x*.
/c*x*/dall	This object identifier is for the all drive groups on enclosure *x* on controller *x*.
/c*x*/p*x*	This object identifier is for a PHY operation *x* on controller *x*
/c*x*/pall	This object identifier is for all PHY operations on controller *x*
/c*x*/bbu	This object identifier is for a BBU *x* on controller *x*
/c*x*/cv	This object identifier is for a cache vault *x* on controller *x*

NOTE If enclosures are not used to connect physical drives to the controller, you do not specify the enclosure ID in the command.

The StorCLI tool supports the following verbs.

Table 2 Verbs in the StorCli Command Syntax

Verbs	Description
add	This verb adds virtual drives, JBODs, and so on to the object identifier.
del	This verb deletes a drive, value, or property of the object identifier.
set	This verb sets a value of the object identifier.
show	This verb shows the value and properties of the object identifier.
pause	This verb pauses an ongoing operation.
resume	This verb resumes paused operation.
compare	This verb compares an input value with a system value.
download	This verb downloads and flashes a file to the target.
start	This verb starts an operation.
flush	This verb flushes a controller cache or a drive cache.
stop	This verb stops an operation that is in progress. A stopped process cannot be resumed.
import	This verb imports the foreign configuration into the drive.
expand	This verb expands the size of the virtual drive.
insert	This verb replaces the configured drive that is identified as missing, and starts an automatic rebuild.
flasherase	This verb erases the flash memory on the controller.
transform	This verb downgrades the firmware memory on the controller.
restart	This verb restarts the controller without a system reboot.
apply	This verb applies the activation Key to a WarpDrive® card.

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

4.4 Working with the Storage Command Line Tool

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

NOTE The Storage Command Line Tool is not case sensitive.

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

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

4.4.1 System Commands

4.4.1.1 System Show Commands

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

```
storcli show
storcli show all
```

```
storcli show ctrlcount
storcli show help
storcli -v
```

The detailed description for each command follows.

storcli show

This command shows a summary of controller and controller-associated information for the system. The summary includes the number of controllers, the host name, the operating system information, and the overview of the existing configuration.

storcli show all

This command shows the list of controllers and controller-associated information, information about the drives that need attention, and advanced software options.

storcli show ctrlcount

This command shows the number of controllers detected in the server.

storcli show help

This command shows help for all commands at the server level.

storcli -v

This command shows the version of the Storage Command Line Tool.

4.4.2 Controller Commands

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

4.4.2.1 Show and Set Controller Properties Commands

Table 3 Controller Commands Quick Reference Table

Commands	Value Range	Description
`show <properties>`	See Table 4	Shows specific controller properties.
`set <properties>`	See Table 4	Sets controller properties.
`show`	`all`: Shows all properties of the virtual drive. `freespace`: Shows the freespace in the controller. See Controller Show Commands.	Shows physical drive information.

This section provides command information to show and set controller properties.

NOTE You cannot set multiple properties with a single command.

storcli /c*x* show <*property*>

This command shows the current value of the specified property on the specified controller.

General example output:

```
Status Code = 0
Status = Success
Description = None
Controller: 0
```

```
Property_name = Property_value
```

You can show the following properties using the `storcli /cx show <property1>|<property2>` command.

```
storcli /cx show abortcconerror
storcli /cx show activityforlocate
storcli /cx show alarm
storcli /cx show backplane
storcli /cx show batterywarning
storcli /cx show bgirate
storcli /cx show bootwithpinnedcache
storcli /cx show cachebypass
storcli /cx show cacheflushint
storcli /cx show ccrate
storcli /cx show coercion
storcli /cx show consistencycheck|cc
storcli /cx show copyback
storcli /cx show directpdmapping
storcli /cx show dimmerswitch|ds
storcli /cx show eccbucketleakrate
storcli /cx show eccbucketsize
storcli /cx show eghs
storcli /cx show jbod
storcli /cx show loadbalancemode
storcli /cx show maintainpdfailhistory
storcli /cx show migraterate
storcli /cx show ncq
storcli /cx show patrolread|pr
storcli /cx show perfmode
storcli /cx show pi
storcli /cx show prcorrectunconfiguredareas
storcli /cx show prrate
storcli /cx show rebuildrate
storcli /cx show rehostinfo
storcli /cx show restorehotspare
storcli /cx show safeid
storcli /cx show smartpollinterval
storcli /cx show spinupdelay
storcli /cx show spinupdrivecount
storcli /cx show time
storcli /cx show usefdeonlyencrypt
storcli /cx show badblocks
storcli.exe /cx show wbsupport
storcli.exe /cx show DPM
storcli.exe /cx show SGPIOforce
```

storcli /*cx* set <*property*> = <*value*>

General example output:

```
Status Code = 0
Status = Success
Description = None

Controller 0, new Property_name = Property_value
```

The following commands are examples of the properties that can be set using the `storcli /cx set <property>=<value>` command:

```
storcli /cx set  abortcconerror=<on|off>
storcli /cx set  termlog[=on|off|offthisboot]
storcli /cx set  activityforlocate=<on|off>
storcli /cx set  alarm=<on|off|silence>
storcli /cx set  batterywarning=<on|off>
storcli /cx set  bgirate=<value>
storcli /cx set  bootwithpinnedcache=<on|off>
storcli /cx set  cachebypass=<on|off>
storcli /cx set  cacheflushinterval=<value>
storcli /cx set  ccrate=<value>
storcli /cx set  coercion=<value>
storcli /cx set  consistencycheck|cc=[off|seq|conc][delay=value]
[starttime=yyyy/mm/dd hh] [excludevd=x-y,z]
storcli /cx set  copyback=<on|off> type=<smartssd|smarthdd|all>
storcli /cx set  directpdmapping=<on|off>
storcli /cx set  eccbucketleakrate=<value>
storcli /cx set  eccbucketsize=<value>
storcli /cx set  eghs [state=<on|off>][smarter=<on|off>][eug=<on|off>]
storcli /cx set  backplane [mode=<0-3>][expose=<on|off>]
storcli /cx set  dimmerswitch|ds=<on|off type=1|2|3|4>
storcli /cx set  foreignautoimport=<on|off>
storcli /cx set  jbod=<on|off>
storcli /cx set  loadbalancemode=<value>
storcli /cx set  maintainpdfailhistory=<on|off>
storcli /cx set  migraterate=<value>
storcli /cx set  ncq=<on|off>
storcli /cx set  patrolread|pr {=on mode=<auto|manual>}|{off}
storcli /cxvset  perfmode=<value>
storcli /cx set  pi [state=<on|off>][import=<on|off>]
storcli /cx set  prcorrectunconfiguredareas=<on|off>
storcli /cx set  prrate=<value>
storcli /cx set  rebuildrate=<value>
storcli /cx set  restorehotspare=<on|off>
storcli /cx set  smartpollinterval=<value>
storcli /cx set  spinupdelay=<value>
storcli /cx set  spinupdrivecount=<value>
storcli /cx set  stoponerror=<on|off>
storcli /cx set  usefdeonlyencrypt=<on|off>
storcli /cx set  time=yyyymmdd hh:mm:ss|systemtime
storcli /cx set  usefdeonlyencrypt=<on|off>
storcli.exe /cx set  DPM=<on|off>
storcli.exe /cx set  supportssdpatrolread=<on|off>
storcli.exe /cx set  SGPIOforce=<on|off>
storcli /cx set  immediateio=<on|off>
storcli /cx set  driveactivityled=<on|off>
storcli /cx set  sesmonitoring=<on|off>
```

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

Table 4 Properties for Show and Set Commands

Property Name	Set Command Range	Description
`abortcconerror`	`on|off`	Aborts consistency check when it detects an inconsistency.
`activityforlocate`	`on|off`	Enables/disables drive activity, drive activity locates function for systems without SGPIO/SES capabilities.
`alarm`	`on|off|silence` `silence`: Silences the alarm.	Enables/disables alarm on critical errors.
`batterywarning`	`on|off`	Enables/disables battery warnings.
`bgirate`	0 to 100	Sets background initialization rate in percentage.
`cacheflushint`	0 to 255, default value `4`	Sets cache flush interval in seconds.
`ccrate`	0 to 100	Sets consistency check rate in percentage.
`coercion`	`0`: No coercion `1`: 128 MB `2`: 1 GB	Sets drive capacity in coercion mode.
`consistencycheck`	See Consistency Check.	See Consistency Check.

`copyback`	`on|off`
	`type = smartssd|smarthdd|all`
	`smartssd`: Copy back enabled for SSD drives.
	`smarthdd`: Copy back enabled for HDD drives.
	`all`: Copy back enabled for both ssd drives and HDD drives.
	Example:
	`storcli /cx set copyback=on type=all`
		Enables/disables copy back for drive types.

`directpdmapping`	`on|off`	Enables/disables direct physical drive mapping. When enclosures are used, this feature is disabled; otherwise it should be enabled.
`eccbucketleakrate`	0 to 65535	Sets leak rate of the single-bit bucket in minutes (one entry removed per leak-rate).
`eccbucketsize`	0 to 255	Sets size of ECC single-bit-error bucket (logs event when full).
`eghs state`	`on|off`	Enables/disables the commissioning of otherwise incompatible global hot spare drives as Emergency Hot Spare (EHSP) drives.
`eghs smarter`	`on|off`	Enables/disables the commissioning of Emergency Hot Spare (EHSP) drives for Predictive Failure (PFA) events.
`eghs eug`	`on|off`	Enables/disables the commissioning of Unconfigured Good drives as Emergency Hot Spare (EHSP) drives.

Table 4 Properties for Show and Set Commands (Continued)

Property Name	Set Command Range	Description
backplane mode	0: Use autodetect logic of backplanes, such as SGPIO and I2C SEP using GPIO pins. 1: Disable autodetect SGPIO. 2: Disable I2C SEP autodetect. 3: Disable both the autodetects.	Configures enclosure detection on a non-SES/expander backplane.
backplane expose	on|off	Enables/disables device drivers to expose enclosure devices; for example, expanders, SEPs.
dimmerswitch|ds	See Dimmer Switch Commands.	See Dimmer Switch Commands.
foreignautoimport	on|off	Imports foreign configuration automatically, at boot.
jbod	on|off	Enables/disables JBOD mode; by default, drives become system drives.Not supported by all controllers.
loadbalancemode	on|off	Enables/disables automatic load balancing between SAS phys or ports in a wide port configuration.
maintainpdfailhistory	on|off	Maintains the physical drive fail history.
migraterate	0 to 100	Sets data migration rate in percentage.
patrolread|pr	See Patrol Read.	See Patrol Read.
perfmode	0: Tuned to provide best IOPS, currently applicable to non-FastPath 1: Tuned to provide least latency, currently applicable to non-FastPath	Performance tuning setting for the controller.
pi	on|off	Enables/disables data protection on the controller.
pi import	on|off	Enables/disables import data protection drives on the controller.
prcorrectunconfiguredareas	on|off	Corrects media errors during PR by writing 0s to unconfigured areas of the disk.
prrate	0 to 100	Sets patrol read rate of the virtual drives in percentage.
rebuildrate	0 to 100	Sets rebuild rate of the drive in percentage.
reconrate	0 to 100	Sets reconstruction rate for a drive in percentage.
restorehotspare	on|off	Becomes a hot spare on insertion of a failed drive.
smartpollinterval	0 to 65535	Set time for polling of SMART errors in seconds.
spinupdrivecount	0 to 255	Sets number of drives that are spun up at a time.
spinupdelay	0 to 255	Sets spin-up delay between a group of drives or a set of drives, in seconds.
stoponerror	on|off	Stops the ServeRAID BIOS during POST, if any errors are encountered.
time	Valid time in *yymmdd hh:mm:ss* format or systemtime	Sets the controller time to your input value or the system time (local time in 24-hour format).
usefdeonlyencrypt	on|off	Enables/disables FDE drive-based encryption.
DPM	on|off	Enables/disables drive performance monitoring
supportssdpatrolread	on|off	Enables/disables patrol read for SSD drives.

Table 4 Properties for Show and Set Commands (Continued)

Property Name	Set Command Range	Description
SGPIOforce	on|off	Forces SGPIO status per port only for four drives; affects HPC controllers.
immediateio	on|off	Enables or disables Immediate IO.
driveactivityled	on|off	Activates or deactivates the LED that indicates drive activity.
sesmonitoring	on|off	Enables or disables SES monitoring.

4.4.2.2 Controller Show Commands

The Storage Command Line Tool supports the following show commands:

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

The detailed description for each command follows.

storcli /*cx* show

This command shows the summary of the controller information. The summary includes basic controller information, foreign configurations, drive groups, virtual drives, physical drives, enclosures, and BBU information.

Input example:

```
storcli /c1 show
```

storcli /*cx* show all

This command shows all controller information, which includes basic controller information, bus information, controller status, advanced software options, controller policies, controller defaults, controller capabilities, scheduled tasks, miscellaneous properties, foreign configurations, drive groups, virtual drives, physical drives, enclosures, and BBU information.

Input example:

```
storcli /c0 show all
```

NOTE The PCI information displayed as a part of `storcli /cx show` and `storcli /cx show all` commands is not applicable for the FreeBSD operating system. Hence, the PCI information fields are displayed as N/A.

storcli /*cx* show freespace

This command shows the usable free space in the controller.

Input example:

```
storcli /c0 show freespace
```

4.4.2.3 Controller Background Tasks Operation Commands

4.4.2.3.1 Rebuild Rate

```
storcli /cx set rebuildrate=<value>
storcli /cx show rebuildrate
```

The detailed description for each command follows.

storcli /c*x* set rebuildrate=<*value*>

This command sets the rebuild task rate of the specified controller. The input value is in percentage.

Input example:

```
storcli /c0 set rebuildrate=30
```

NOTE A high rebuild rate slows down I/O processing.

storcli /c*x* show rebuildrate

This command shows the current rebuild task rate of the specified controller in percentage.

Input example:

```
storcli /c0 show rebuildrate
```

4.4.2.3.2 Patrol Read

The Storage Command Line Tool supports the following patrol read commands:

```
storcli /cx resume patrolread
storcli /cx set patrolread ={{on mode=<auto|manual>}|{off}}
storcli /cx set patrolread [starttime=<yyyy/mm/dd hh>] [maxconcurrentpd=<value>]
[includessds=<on|off>] [uncfgareas=<on|off>]
storcli /cx set patrolread  delay=<value>
storcli /cx show patrolread
storcli /cx start patrolread
storcli /cx stop patrolread
storcli /cx pause patrolread
```

NOTE A patrol read operation is scheduled for all the physical drives of the controller.

The detailed description for each command follows.

storcli /c*x* resume patrolread

This command resumes a suspended patrol read operation.

Input example:

```
storcli /c0 resume patrolread
```

storcli /c*x* set patrolread {=on mode=<auto|manual>}|{off}

This command turns the patrol read scheduling on and sets the mode of the patrol read to automatic or manual.

Input example:

```
storcli /c0 set patrolread=on mode=manual
```

storcli /c*x* set patrolread [starttime=<*yyyy/mm/dd hh*>] [maxconcurrentpd=<*value*>] [includessds=<on|off>] [uncfgareas=on|off]

This command schedules a patrol read operation. You can use the following options for patrol read command

Table 5 Set Patrolread Input Options

Option	Value Range	Description
`starttime`	A valid date and hour in 24 hours format	Sets the start time in *yyyy/mm/dd hh* format.
`maxconcurrentpd`	Valid number of physical drives present	Sets the number of physical drives that can be patrol read at a single time.
`includessds`	—	Include SSDs in the patrol read.
`uncfgareas`	—	Include the areas not configured in the patrol read.

NOTE Controller time is taken as a reference for scheduling a patrol read operation.

Input example:

```
storcli /c0 set patrolread=on starttime=2012/02/21 00
```

storcli /c*x* set patrolread [delay=<*value*>]

This command delays the scheduled patrol read in hours.

Input example:

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

storcli /c*x* show patrolRead

This command shows the current state of the patrol read along with other details such as the **PR Mode**, **PR Execution Delay**, **PR iterations completed**, and **PR on SSD**. This command also shows the start time and the date when the patrol read started.

The values shown for the current state of the patrol read are **Ready**, **Active**, **Paused**, **Aborted**, **Stopped**, or **Unknown**.

If the state of the patrol read is active, a numeric value is shown along with the state which depicts the number of physical drives that have completed the patrol read operation. As an example, `Active 1` means that the one physical drive has completed the patrol read operation.

Input example:

```
storcli /c0 show patrolread
```

storcli /c*x* start patrolread

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

Input example:

```
storcli /c0 start patrolread
```

storcli /c*x* stop patrolread

This command stops a running patrol read operation.

Input example:

```
storcli /c0 stop patrolread
```

NOTE You cannot resume a stopped patrol read.

storcli /c*x* pause patrolread

This command pauses a running patrol read operation.

Input example:

```
storcli /c0 pause patrolread
```

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

4.4.2.3.3 Consistency Check

The Storage Command Line Tool supports the following commands to schedule, perform, and view the status of a consistency check (CC) operation:

```
storcli /cx set consistencycheck|cc=[off|seq|conc] [delay=value]
starttime=yyyy/mm/dd hh [excludevd=x-y,z]
storcli /cx show cc
storcli /cx show ccrate
```

The detailed description for each command follows.

storcli /*cx* set consistencycheck|cc=[off|seq|conc][delay=*value*] starttime=*yyyy/mm/dd hh* [excludevd=*x-y,z*]

This command schedules a consistency check (CC) operation. You can use the following options with the consistency check command.

Table 6 Set CC Input Options

Option	Value Range	Description
`cc`	`seq`: Sequential mode. `conc`: Concurrent mode. `off`: Turns off the consistency check.	Sets CC to either sequential mode, or concurrent mode, or turns off the CC. **NOTE** The concurrent mode slows I/O processing.
`delay`	`-1` and any integer value.	Delay a scheduled consistency check. The value is in hours. A value of `0` makes the CC runs continuously with no delay (in a loop). **NOTE** Only scheduled consistency checks can be delayed.
`starttime`	A valid date and hour in 24-hours format.	Start time of a consistency check is *yyyy/mm/dd hh* format.
`excludevd`	The range should be less than the number of virtual drives.	Excludes virtual drives from the consistency checks. To exclude particular virtual drives, you can provide list of virtual drive names (V*x*,V*y* … format) or the range of virtual drives that you want to exclude from a consistency check (V*x*-V*y* format). If this option is not specified in the command, no virtual drives are excluded.

Input example:

```
storcli /c0 set CC=on starttime=2012/02/21 00 excludevd v0-v3
```

storcli /*cx* show cc

This command shows the consistency check schedule properties for a controller.

Input example:

```
storcli /c0 show cc
```

storcli /*cx* show ccrate

This command checks the status of a consistency check operation. The CC rate appears in percentage.

Input example:

```
storcli /c0 show ccrate
```

NOTE A high CC rate slows I/O processing.

4.4.2.4 Premium Feature Key Commands

The Storage Command Line Tool supports the following commands for premium feature keys:

```
storcli /cx set advancedsoftwareoptions(aso) key=<value> [preview]
storcli /cx aso [transfertovault][rehostcomplete][deactivatetrialkey]
storcli /cx show safeid
```

The detailed description for the command follows.

storcli /*cx* set advancedsoftwareoptions(aso) key=<*value*> [preview]

This command activates advanced software options (ASO) for a controller. You can use the following options with the advanced software options command.

Table 7 Set Advanced Software Options Input Options

Option	Value Range	Description
key	40 alpha numeric characters.	Key to activate ASO on the controller. **NOTE** After they are activated, ASOs cannot be removed from the controller.
deactivatetri-alkey	—	Deactivates the trial key applied on the specified controller.
rehostcomplete	—	Enables rehosting on the specified controller.
transfertovault	—	Transfers the ASO key to the vault and disables the ASO.

Input example:

```
storcli /c0 set Aso key=LSI0000
```

storcli /*cx* show safeid

This command shows the Safe ID of the specified controller.

Input example:

```
storcli /c0 show safeid
```

4.4.2.5 Controller Security Commands

The Storage Command Line Tool supports the following controller security commands:

```
storcli /cx compare securitykey=ssssss
storcli /cx delete securitykey
storcli /cx set securitykey keyid=kkkk
storcli /cx set securitykey=sssss [passphrase=sssss][keyid=sssss]
storcli /cx set securitykey=sssss  oldsecuritykey=ssss [passphrase=sssss]
[keyid=sssss]
```

The detailed description for each command follows.

storcli /*cx* show securitykey keyid

This command shows the security key on the controller.

Input example:

```
storcli /c0 show securityKey keyid
```

storcli /*cx* compare securitykey=*ssssss*

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

storcli /*cx* delete securitykey

This command deletes the security key of the controller.

Input example:

```
storcli /c0 delete securitykey
```

storcli /*cx* set securitykey keyId=*kkkk*

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

storcli /*cx* set securitykey=sssss [passphrase=sssss][keyid=*sssss*]

This command sets the security key for the controller. You can use the following options with the set security key command.

Table 8 Set Security Key Input Options

Option	Value Range	Description
passphrase	Should have a combination of numbers, upper case letters, lower case letters and special characters. Minimum of 8 characters and maximum of 32 characters.	String that is linked to the controller and is used in the next bootup to encrypt the lock key. If the passphrase is not set, the controller generates it by default.
keyid	—	Unique ID set for different controllers to help you specify a pass-phrase to a specific controller.

Input example:

```
storcli /c0 set securitykey=Lsi@12345 passphrase=Lsi@123456 keyid=1
```

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

This command changes the security key for the controller.

Input example:

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

4.4.2.6 Controller Cache Command

The following command flushes the controller cache.

storcli /*cx* flush|flushcache

This command flushes the controller cache.

Input example:

```
storcli /c0 flushcache
```

4.4.3 Drive Commands

This section describes the drive commands, which provide information and perform actions related to physical drives. The following table describes frequently used virtual drive commands.

Table 9 Physical Drives Commands Quick Reference Table

Commands	Value Range	Description
`set`	`missing`: Sets the drive status as missing. `good`: Sets the drive status to unconfigured good. `offline`: Sets the drive status to offline. `online`: Sets the drive status to online.	Sets physical drive properties.
`show`	`all`: shows all properties of the physical drive. See Drive Show Commands.	Shows virtual drive information.

4.4.3.1 Drive Show Commands

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

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

NOTE If enclosures are used to connect physical drives to the controller, specify the enclosure ID in the command. If no enclosures are used, you must specify the controller ID and slot ID.

The detailed description for each command follows.

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

This command shows the summary of the physical drive for a specified slot in the controller.

Input example:

```
storcli /c0/e0/s4,5 show
```

storcli /*cx*[/eall]/sall show

This command shows the summary information for all the enclosures and physical drives connected to the controller.

Input example:

```
storcli /c0/eall/sall show
```

storcli /*cx*[/*ex*]/*sx*|sall show all

This command shows all information of a physical drive for the specified slot in the controller. If you use the `all` option, the command shows information for all slots on the controller. *x* stands for a number, a list of numbers, a range of numbers, or all numbers.

This command also shows the NCQ (Native Command Queuing) status (**Enabled**, **Disabled**, or **N/A**) which is applicable only to SATA drives. If the controller to which the SATA drive is connected supports NCQ and NCQ is enabled on the SATA drive, the status is shown as **Enabled**; else it is shown as **Disabled**. If NCQ is not a supported drive operation on the controller, the status is shown as **N/A**.

Input examples:

```
storcli /c0/e3/s0-3 show all
```

```
storcli /c0/e35/sall show all
```

NOTE The `storcli /cx/sx show all` command shows tape drives information.

storcli /*cx*[/*ex*]/*sx* show smart

This command displays the SMART information of a SATA drive.

Input example:

```
storcli /c0/e5/s1 show smart
```

4.4.3.2 Missing Drives Commands

The Storage Command Line Tool supports the following commands to mark and replace missing physical drives:

```
storcli /cx[/ex]/sx insert dg=A array=B row=C
storcli /cx[/ex]/sx set missing
storcli /cx[/ex]/sx set offline
storcli /cx/dall
```

The detailed description for each command follows.

storcli /*cx*[/*ex*]/*sx* insert dg=A array=*B* row=*C*

This command replaces the configured drive that is identified as missing, and then starts an automatic rebuild.

Input example:

```
storcli /c0/e25/s3 insert dg=0 array=2 row=1
```

storcli /*cx*[/*ex*]/*sx* set missing

This command marks a drive as missing.

Input example:

```
storcli /c0/s4 set missing
```

storcli /cx/dall

This command is used to find the missing drives.

storcli /*cx*[/*ex*]/*sx* set offline

This command marks the drive in an array as offline.

NOTE To set a drive that is part of an array as *missing*, first set it as *offline*. After the drive is set to *offline*, you can then set the drive to *missing*.

4.4.3.3 Set Drive State Commands

The Storage Command Line Tool supports the following commands to set the status of physical drives:

```
storcli /cx[/ex]/sx set jbod
storcli /cx[/ex]/sx set good [force]
storcli /cx[/ex]/sx set offline
storcli /cx[/ex]/sx set online
storcli /cx[/ex]/sx set missing
storcli /cx[/ex]/sx set bootdrive=<on|off>
```

The detailed description for each command follows.

storcli /*cx*[/*ex*]/*sx* set jbod

This command sets the drive state to JBOD.

Input example:

```
storcli /c1/e56/s3 set jbod
```

storcli /*cx*[/*ex*]/*sx* set good [force]

This drive changes the drive state to unconfigured good. If the drive has the operating system in it, use the `force` option.

Input example:

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

storcli /*cx*[/*ex*]/*sx* set offline

This command changes the drive state to offline.

Input example:

```
storcli /c1/e56/s3 set offline
```

storcli /*cx*[/*ex*]/*sx* set online

This command changes the drive state to online.

Input example:

```
storcli /c1/e56/s3 set online
```

storcli /*cx*[/*ex*]/*sx* set missing

This command marks a drive as missing.

Input example:

```
storcli /c1/e56/s3 set missing
```

storcli /*cx*[/*ex*]/*sx* set bootmode=<on|off>

This command sets or unsets a physical drive as a boot drive.

Input example:

```
storcli /c1/e56/s3 set bootmode=on
```

4.4.3.4 Drive Initialization Commands

When you initialize drives, all the data from the drives is cleared. The Storage Command Line Tool supports the following commands to initialize drives:

```
storcli /cx[/ex]/sx show initialization
storcli /cx[/ex]/sx start initialization
storcli /cx[/ex]/sx stop initialization
```

The detailed description for each command follows.

storcli /*cx*[/*ex*]/*sx* show initialization

This command shows the current progress of the initialization progress in percentage.

The estimated time (in minutes) left to complete the operation is also shown.

Input example:

```
storcli /c0/e31/s4 show initialization
```

storcli /*cx*[/*ex*]/*sx* start initialization

This command starts the initialization process on a drive.

Input example:

```
storcli /c0/e31/s4 start initialization
```

storcli /*cx*[/*ex*]/*sx* stop initialization

This command stops an initialization process running on the specified drive. A stopped initialization process cannot be resumed.

Input example:

```
storcli /c0/e56/s1 stop initialization
```

4.4.3.5 Drive Firmware Download Commands

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

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

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

Input example:

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

4.4.3.6 Locate Drives Commands

The Storage Command Line Tool supports the following commands to locate a drive and activate the physical disk activity LED:

```
storcli /cx[/ex]/sx start locate
storcli /cx[/ex]/sx stop locate
```

The detailed description for each command follows.

storcli /*cx*[/*ex*]/*sx* start locate

This command locates a drive and activates the drive's LED.

Input example:

```
storcli /c0/e56/s1 start locate
```

storcli /*cx*[/*ex*]/*sx* stop locate

This command stops a locate operation and deactivates the drive's LED.

Input example:

```
storcli /c0/e56/s1 stop locate
```

4.4.3.7 Prepare to Remove Drives Commands

The Storage CLI supports the following commands to prepare the physical drive for removal:

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

The detailed description for each command follows.

storcli /*cx*[/*ex*]/*sx* spindown

This command spins down an unconfigured drive and prepares it for removal. The drive state is unaffiliated and it is marked offline.

Input example:

```
storcli /cx/e34/s4 spindown
```

storcli /*cx*[/*ex*]/*sx* spinup

This command spins up a spun-down drive and the drive state is unconfigured good.

Input example:

```
storcli /cx/e34/s4 spinup
```

4.4.3.8 Drive Security Commands

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

```
storcli /cx[/ex]/sx show securitykey keyid
```

storcli /*cx*[/*ex*]/*sx* show securitykey keyid

This command shows the security key for secured physical drives.

Input example:

```
storcli /c0/e252/s1 show SecurityKey keyid
```

storcli /*cx*[/*ex*]/*sx* set security = on

This command enables security key on a JBOD.

Input example:

```
storcli /c0/e252/s1 set security = on
```

4.4.3.9 Drive Secure Erase Commands

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

```
storcli /cx[/ex]/sx secureerase [force]
storcli /cx[/ex]/sx show erase
storcli /cx[/ex]/sx start erase [simple|normal|thorough] [patternA=<value1>]
[patternB=<value2>]
storcli /cx[/ex]/sx stop erase
```

The detailed description for each command follows.

storcli /*cx*[/*ex*]/*sx* secureerase [force]

This command erases the drive's security configuration and securely erases data on a drive. You can use the `force` option as a confirmation to erase the data on the drive and the security information.

Input example:

```
storcli /c0/e25/s1 secureerase
```

NOTE This command deletes data on the drive and the security configuration and this data is no longer accessible. This command is used for SED drives only.

storcli /*cx*[/*ex*]/*sx* show erase

This command provides the status of erase operation on non-SEDs.

Input example:

```
storcli /c0/e25/s1 show erase
```

storcli /*cx*[/*ex*]/*sx* start erase [simple|normal|thorough|standard] [patternA=<*val1*>] [patternB=<*val2*>]

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

Table 10 Drive Erase Command Options

Options	Value Range	Description
`erase`	`simple`: Single pass, single pattern write `normal`: Three pass, three pattern write `thorough`: Nine pass, repeats the normal write 3 times	Secure erase type.
`patternA`	8-bit value	Erase pattern A to overwrite the data.
`patternB`	8-bit value	Erase pattern B to overwrite the data.

Input example:

```
storcli /c0/e25/s1 start erase thorough patternA=10010011 patternB=11110000
```

4.4.3.10 Rebuild Drives Commands

The following commands rebuild drives in the Storage Command Line Tool:

```
storcli /cx[/ex]/sx pause rebuild
storcli /cx[/ex]/sx resume rebuild
storcli /cx[/ex]/sx show rebuild
storcli /cx[/ex]/sx start rebuild
storcli /cx[/ex]/sx stop rebuild
```

NOTE If enclosures are used to connect physical drives to the controller, specify the enclosure ID in the command.

The detailed description for each command follows.

storcli /*cx*[/*ex*]/*sx* pause rebuild

This command pauses an ongoing rebuild process. You can run this command only for a drive that is currently rebuilt.

Input example:

```
storcli /c0/s4 pause rebuild
```

storcli /*cx*[/*ex*]/*sx* resume rebuild

This command resumes a paused rebuild process. You can run this command only when a paused rebuild process for the drive exists.

Input example:

```
storcli /c0/s4 resume rebuild
```

storcli /*cx*[/*ex*]/*sx* show rebuild

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

The estimated time (in minutes) left to complete the operation is also shown.

Input example:

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

storcli /*cx*[/*ex*]/*sx* start rebuild

This command starts a rebuild operation for a drive.

Input example:

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

storcli /*cx*[/*ex*]/*sx* stop rebuild

This command stops a rebuild operation. You can run this command only for a drive that is currently rebuilt.

Input example:

```
storcli /c0/s4 stop rebuild
```

4.4.3.11 Drive Copyback Commands

The Storage Command Line Tool supports the following commands for drive copyback:

```
storcli /cx[/ex]/sx pause copyback
storcli /cx[/ex]/sx resume copyback
storcli /cx[/ex]/sx show copyback
storcli /cx[/ex]/sx start copyback target=eid:sid
storcli /cx[/ex]/sx stop copyback
```

The detailed description for each command follows.

NOTE In the copyback commands, `cx[/ex]/sx` indicates the source drive and `eid:sid` indicates the target drive.

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

storcli /*cx*[/*ex*]/*sx* pause copyback

This command pauses a copyback operation. You can run this command only when there is a copyback operation running.

Input example:

```
storcli /c0/e25/s4 pause copyback
```

storcli /*cx*[/*ex*]/*sx* resume copyback

This command resumes a paused copyback operation. You can run this command only when there is a paused copyback process for the drive.

Input example:

```
storcli /c0/e25/s4 resume copyback
```

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

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

The estimated time (in minutes) left to complete the operation is also shown.

Input example:

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

storcli /*cx*[/*ex*]/*sx* start copyback target=eid:sid

This command starts a copyback operation for a drive.

Input example:

```
storcli /c0/e25/s4 start copyback target=25:8?
```

storcli /*cx*[/*ex*]/*sx* stop copyback

This command stops a copyback operation. You can run this command only on drives that have the copyback operation running.

Input example:

```
storcli /c0/e25/s4 stop copyback
```

NOTE A stopped rebuild process cannot be resumed.

4.4.3.12 Hot Spare Drive Commands

The following commands create and delete hot spare drives:

```
storcli /cx[/ex]/sx add hotsparedrive
{dgs=<n|0,1,2...>}[enclaffinity][nonrevertible]
storcli /cx/[ex]/sx delete hotsparedrive
```

NOTE If enclosures are used to connect the physical drives to the controller, specify the enclosure ID in the command.

The detailed description for each command follows.

storcli /*cx*[/*ex*]/*sx* add hotsparedrive [{dgs=<*n*|0,1,2...>}] [enclaffinity][nonrevertible]

This command creates a hot spare drive. You can use the following options to create a hot spare drive:

Table 11 Add Hotsparedrive Input Options

Option	Value Range	Description
dgs	Valid drive group number	Specifies the drive group to which the hot spare drive is dedicated.
enclaffinity	Valid enclosure number	Specifies the enclosure with which the hot spare is associated. If this option is specified, affinity is set; if it is not specified, there is no affinity. **NOTE** Affinity cannot be removed after it is set for a hot spare drive.
nonrevertible	—	Sets the drive as a nonrevertible hot spare.

Input example:

```
storcli /c0/e3/s4,5 add hotsparedrive
```

This command sets the drives /c0/e3/s4,5 as Global Hot spare.

Input example:

```
storcli /c0/e3/s6,8 add hotsparedrive dgs=0,1
```

This command sets /c0/e3/s6,8 as Dedicated Hot spare for disk groups 0,1.

storcli /*cx*/[*ex*]/*sx* delete hotsparedrive

This command deletes a hot spare drive.

Input example:

```
storcli /c0/e3/s4,5 delete hotsparedrive
```

4.4.4 Virtual Drive Commands

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

Table 12 Virtual Drives Commands Quick Reference Table

Commands	Value Range	Description
`add`	See the following **Add RAID Configuration Input Options** tables.	Creates virtual drives.
`delete`	`cc` or `cachecade`: Deletes CacheCade virtual drives. `force`: Deletes the virtual drive where operating system is present.	Deletes a virtual drive.
`set`	See the following **Add RAID Configuration Input Options** tables and **Change Virtual Properties Commands** section.	Sets virtual drive properties.
`show`	`all`: Shows all properties of the virtual drive. `cc`: Shows properties of CacheCade virtual drives. See the **Virtual Drive Show Command** section.	Shows virtual drive information.

4.4.4.1 Add Virtual Drives Commands

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

```
storcli /cx add vd raid[0|1|5|6|00|10|50|60][Size=<VD1_Sz>,<VD2_Sz>,..|all]
[name=<VDNAME1>,..] drives=e:s|e:s-x,y;e:s-x,y,z [PDperArray=x][SED]
[pdcache=on|off|default][pi] [DimmerSwitch(ds)=default|automatic(auto)|
none|maximum(max)|MaximumWithoutCaching(maxnocache)]
[wt|wb|awb] [nora|ra] [direct|cached][cachevd] [Strip=<8|16|32|64|128|256|1024>]
[AfterVd=X] [EmulationType=0|1|2] [Spares = [e:]s|[e:]s-x|[e:]s-x,y]
[force][ExclusiveAccess]
```

```
storcli /cx add vd each raid0 [name=<VDNAME1>,..] [drives=e:s|e:s-x|e:s-x,y]
[SED] [pdcache=on|off|default][pi] [DimmerSwitch(ds)=default|automatic(auto)|
none|maximum(max)|MaximumWithoutCaching(maxnocache)] [wt|wb|awb] [nora|ra]
[direct|cached] [EmulationType=0|1|2]
[Strip=<8|16|32|64|128|256|1024>][ExclusiveAccess]
```

```
storcli /cx add VD cachecade|cc raid[0,1,10] drives = [e:]s|[e:]s-x|[e:]s-x,y
[WT|WB|AWB] [assignvds = 0,1,2]
```

This command creates a RAID configuration. You can use the following options to create the RAID volume:

NOTE * indicates default values.

The detailed description for each command follows.

storcli /*cx* add vd raid[0|1|5|6|00|10|50|60][Size=<*VD1_Sz*>,<*VD2_Sz*>,..|*all] [name=<*VDNAME1*>,..] drives=*e:s*|*e:s-x*|*e:s-x,y;e:s-x,y,z* [PDperArray=x][SED] [pdcache=on|off|*default][pi] [DimmerSwitch(ds)=default|automatic(auto)| *none|maximum(max)|MaximumWithoutCaching(maxnocache)][cachevd][ExclusiveAccess|SharedAccess*] [wt|*wb |awb] [nora|*ra] [*direct|cached] [Strip=<8|16|32|64|128|256|1024>] [AfterVd=*X*] [Spares = [*e*:]*s*|[*e*:]*s-x*|[*e*:]*s-x,y*] [force]**

Table 13 Add RAID Configuration Input Options

Option	Value Range	Description
raid	[0|1|5|6|00|10|50|60].	Sets the RAID type of the configuration.
size	Maximum size based on the physical drives and RAID level.	Sets the size of each virtual drive. The default value is for the capacity of all referenced disks.
name	15 characters of length.	Specifies the drive name for each virtual drive.
drives	Valid enclosure number and valid slot numbers for the enclosure.	In *e*:*s*|*e*:*s-x*|*e*:*s-x*,*y*: ■ *e* specifies the enclosure ID. ■ *s* represents the slot in the enclosure. ■ *e*:*s-x* is the range convention used to represent slots *s* to *x* in the enclosure *e* (250 characters max.). **NOTE** Make sure that the same block size (in a physical drive) is used in each [*e*:*s*] pair. As an example, if you use 4096 bytes in the *e0*:*s0* pair, use 4096 bytes in the *e1*:*s1* pair too. Mixing of block sizes between the [*e*:*s*] pairs is not supported.
pdperarray	1-16.	Specifies the number of physical drives per array. The default value is automatically chosen.
sed	—	Creates security-enabled drives.
pdcache	on|off|default.	Enables or disables PD cache.
pi	—	Enables protection information.
dimmerswitch	default: Logical device uses controller default power-saving policy. automatic (auto): Logical device power savings are managed by firmware. none: No power-saving policy. maximum (max): Logical device uses maximum power savings. MaximumWithoutCaching (maxnocache): Logical device does not cache write to maximize power savings.	Specifies the power-saving policy. Sets to default automatically.
direct|cached	cached: Cached I/O. direct: Direct I/O.	Sets the logical drive cache policy. Direct I/O is the default.
wt|wb|awb	wt: Write through.wb: Write back.awb: Always Write Back.	Enables write through. Write back is the default.
nora|ra	ra: Read ahead.nora: No read ahead.	Disables read ahead. Enabled is the default.
cachevd	—	Enables SSD caching on the created virtual drive.
strip	8, 16, 32, 64, 128, 256, 512, 1024.	Sets the strip size for the RAID configuration.

Table 13 Add RAID Configuration Input Options (Continued)

Option	Value Range	Description
aftervd	Valid virtual drive number.	Creates the VD in the adjacent free slot next to the specified VD.
spares	Number of spare physical drives present.	Specifies the physical drives that are to be assigned to a disk group for spares.
force	—	Forces a security-capable physical drive to be added to a drive group without security.

Input example:

```
storcli /c0 add vd raid10 size=2gb,3gb,4gb names=tmp1,tmp2,tmp3
drives=252:2-3,5,7 pdperarray=2
```

storcli /cx add vd cc|cachecade raid[0,1,10] drives=[*e*:]*s*|[*e*:]*s-x*|[*e*:]*s-x,y* [[wt|*wb|awb]] [assignvds=0,1,2]

This command creates CacheCade virtual drives and associates existing virtual drives to CacheCade virtual drives. You can use the following options to create the CacheCade virtual drive.

Table 14 Add RAID Configuration Input Options

Option	Value Range	Description
cachecade	—	Creates a CacheCade virtual drive.
raid	0, 1, 10	Sets the RAID type of the CacheCade virtual drive.
drives	Valid enclosure number and valid slot number	See the drivesrow in the previous table for format.
wt|*wb|awb	wt: Enables write through. wb: Enables write back. awb Enables always write back.	Enables or disables write cache.
assignvds	Valid virtual drive number (0 to 63)	Specifies the list of virtual drives associated with the new CacheCade virtual drives.

Input example:

```
storcli /c0 add vd raid10 size=2gb,3gb,4gb names=tmp1,tmp2,tmp3 drives=252:2-3, 7
```

4.4.4.2 Delete Virtual Drives Commands

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

```
storcli /cx/vx|vall del
storcli /cx/vx|vall del cachecade
storcli /cx/vx|vall del force
storcli /cx/vx del [cachecade] [discardcache] [force]
```

ATTENTION If the virtual drive has user data, you must use the force option to delete the virtual drive.
A virtual drive with a valid master boot record (MBR) and a partition table is considered to contain user data.

If you delete a virtual drive with a valid MBR without erasing the data and then create a new virtual drive using the same set of physical drives and the same RAID level as the deleted virtual drive, the old unerased MBR still exists at block0 of the new virtual drive, which makes it a virtual drive with valid user data. Therefore, you must provide the force option to delete this newly created virtual drive.

The detailed description for each command follows.

storcli /*cx*/*vx*|vall del

This command deletes a particular virtual drive or, when the `vall` option is used, all the virtual drives on the controller are deleted.

Input example:

```
storcli /c0/v2 del
```

ATTENTION This command deletes virtual drives. Data located on these drives will no longer be accessible.

storcli /*cx*/*vx*|vall del cachecade

This command deletes a specific CacheCade virtual drive on a controller, or all the CacheCade configuration for a controller.

Input example:

```
storcli /c0/vall del cachecade
```

NOTE This command deletes virtual drives. Data located on these drives will no longer be accessible.

storcli /*cx*/*vx*|vall del force

This command deletes a virtual drive only after the cache flush is completed. With the `force` option, the command deletes a virtual drive without waiting for the cache flush to complete.

Input example:

```
storcli /c0/v2 del force
```

NOTE This command deletes the virtual drive where the operating system is present. Data located on these drives and the operating system of the drive will no longer be accessible.

storcli /*cx*/*vx* del [cachecade] [discardcache] [force]

This command with the `discardCache` option deletes the virtual drive without flushing the cached data.

Input example:

```
storcli /c0/v2 delete discardcache
```

4.4.4.3 Virtual Drive Show Commands

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

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

The detailed description for each command follows.

storcli /*cx*/*vx* show

This command shows the summary of the virtual drive information.

Input example:

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

storcli /*cx*/*vx* show all

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

Input example:

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

storcli /*cx*/*vx* show hoqrebuild

This command shows the current status of the head of queue rebuild for the virtual drive.

Input example:

```
storcli /c0/v0 show hoqrebuild
```

4.4.4.4 Preserved Cache Commands

If a virtual drive becomes offline or is deleted because of missing physical disks, the controller preserves the dirty cache from the virtual disk. The Storage Command Line Tool supports the following commands for preserved cache:

```
storcli /cx/vx delete preservedCache [force]
storcli /cx show preservedCache
```

The detailed description for each command follows.

storcli /*cx*/*vx* delete preservedcache

This command deletes the preserved cache for a particular virtual drive on the controller in missing state. Use the `force` option to delete the preserved cache of a virtual drive in offline state.

Input example:

```
storcli /c0/v1 delete preservedcache
```

storcli /*cx* show preservedCache

This command shows the virtual drive that has preserved cache and whether the virtual drive is offline or missing.

Input example:

```
storcli /c0 show preservedCache
```

4.4.4.5 Change Virtual Properties Commands

The Storage Command Line Tool supports the following commands to change virtual drive properties:

```
storcli /cx/vx set accesspolicy=<rw|ro|blocked|rmvblkd>
storcli /cx/vx set iopolicy=<cached|direct>
storcli /cx/vx set name=<namestring>
storcli /cx/vx set pdcache=<on|off|default>
storcli /cx/vx set rdcache=<ra|nora>
storcli /cx/vx|vall set ssdcaching=<on|off>
storcli /cx/vx|vall set HostAccess=ExclusiveAccess|SharedAccess
storcli /cx/vx set wrcache=<wt|wb|awb>
storcli /cx/vx set emulationType=0|1
storcli /cx/vx set ds=Default|Auto|None|Max|MaxNoCache
storcli /cx/vx set autobgi=On|Off
storcli /cx/vx set pi=Off
storcli /cx/vx set bootdrive=<On|Off>
storcli /cx/vx set hidden=On|Off
storcli /cx/vx set hoqrebuild=On|Off
storcli /cx/vx set cbsize=0|1|2 cbmode=0|1|2|3|4|7
```

The detailed description for each command follows.

storcli /c*x*/v*x* set accesspolicy=<rw|ro|blocked|rmvblkd>

This command sets the access policy on a virtual drive to read write, read only, or blocked or rmvblkd (remove blocked).

Input example:

```
storcli /c0/v0 set accesspolicy=rw
```

storcli /c*x*/v*x* set iopolicy=<cached|direct>

This command sets the I/O policy on a virtual drive to cached I/O or direct I/O.

Input example:

```
storcli /c0/v0 set iopolicy=cached
```

storcli /c*x*/v*x* set name=<*namestring*>

This command names a virtual drive. The name is restricted to 15 characters

Input example:

```
storcli /c1/v0 set name=testdrive123
```

storcli /c*x*/v*x* set pdcache=<on|off|default>

This command sets the current disk cache policy on a virtual drive to on, off, or default setting.

Input example:

```
storcli /c0/v0 set pdcache=on
```

storcli /c*x*/v*x* set rdcache=<ra|nora>

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

Input example:

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

storcli /c*x*/v*x*|vall set ssdcaching=<on|off>

This command assigns CacheCade virtual drives. If `ssdcaching=off`, the CacheCade virtual drive is removed.

Input example:

```
storcli /c0/v0 set ssdcaching=on
```

storcli /cx/vx|vall set HostAccess=ExclusiveAccess|SharedAccess

This command sets the host access policy for the virtual drive. when the host access policy is exclusive access, a server has exclusive access to the virtual drive. The virtual drive cannot be shared between servers. If the host policy is shared access, then the virtual drive can be shared between servers.

Input example:

```
storcli /c0/v0 set HostAccess=ExclusiveAccess
```

storcli/cx/vx set wrcache=<wt|wb|awb>

This command sets the write cache policy on a virtual drive to write back, write through, or always write back.

Input example:

```
storcli /c0/v0 set wrcache=wt
```

storcli /cx/vx set hidden=on|off

This command hides or unhides a virtual drive. If `hidden=on`, the virtual drive is hidden.

Input example:

```
storcli /c0/v0 set hidden=on
```

storcli /cx/vx set hoqrebuild=on|off

This command enables or disables the head of the queue drive rebuild on a virtual drive.

Input example:

```
storcli /c0/v0 set hoqrebuild=on
```

storcli /cx/vx set cbsize=0|1|2 cbmode=0|1|2|3|4|7

This command sets the Cache bypass size and the Cache bypass mode on a virtual drive.

The `cbsize` option follows:

- 0 – 64k Cache bypass.
- 1 – 128k Cache bypass.
- 2 – 256k Cache bypass.

The `cbmode` option follows:

- 0 – Enable the intelligent mode Cache bypass.
- 1 – Enable the standard mode Cache bypass.
- 2 – Enable the custom mode Cache bypass 1.
- 3 – Enable the custom mode Cache bypass 2.
- 4 – Enable the custom mode Cache bypass 3.
- 7 – Disable Cache bypass.

Input example:

```
storcli /c0/v0 set cbsize=1 cbmode=2
```

4.4.4.6 Virtual Drive Initialization Commands

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

```
storcli /cx/vx show init
storcli /cx/vx start init [full][Force]
storcli /cx/vx stop init
```

NOTE If the virtual drive has user data, you must use the `force` option to initialize the virtual drive.

A virtual drive with a valid MBR and partition table is considered to contain user data.

The detailed description for each command follows.

storcli /*cx*/*vx* show init

This command shows the initialization progress of a virtual drive in percentage.

The estimated time (in minutes) left to complete the operation is also shown.

Input example:

```
storcli /c0/v2 show init
```

storcli /*cx*/*vx* start init [full]

This command starts the initialization of a virtual drive. The default initialization type is fast initialization. If the `full`option is specified, full initialization of the virtual drive starts.

Input example:

```
storcli /cx/vx start init [full]
```

storcli /*cx*/*vx* stop init

This command stops the initialization of a virtual drive. A stopped initialization cannot be resumed.

Input example:

```
storcli /c0/v0 stop init
```

4.4.4.7 Virtual Drive Erase Commands

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

```
storcli /cx/vx erase
storcli /cx/vx show erase
```

The detailed description for each command follows.

storcli /*cx*/*vx* erase

This command erases the data on the virtual drive.

Input example:

```
storcli /c0/v0 erase
```

storcli /*cx*/*vx* show erase

This command shows the status of the erase operation on the virtual drive.

Input example:

```
storcli /c0/v0 show erase
```

4.4.4.8 Virtual Drive Consistency Check Commands

The Storage Command Line Tool supports the following commands for virtual drive consistency checks:

```
storcli /cx/vx pause cc
storcli /cx/vx resume cc
storcli /cx/vx show cc
storcli /cx/vx start cc [force]
storcli /cx/vx stop cc
```

NOTE If enclosures are used to connect the physical drives to the controller, specify the IDs in the command.

The detailed description for each command follows.

storcli /*cx*/*vx* pause cc

This command pauses an ongoing consistency check process. You can resume the consistency check at a later time. You can run this command only on a virtual drive that has a consistency check operation running.

Input example:

```
storcli /c0/v4 pause cc
```

storcli /*cx*/*vx* resume cc

This command resumes a suspended consistency check operation. You can run this command on a virtual drive that has a paused consistency check operation.

Input example:

```
storcli /c0/v4 resume cc
```

storcli /*cx*/*vx* show cc

This command shows the progress of the consistency check operation in percentage.

The estimated time (in minutes) left to complete the operation is also shown.

Input example:

```
storcli /c0/v5 show cc
```

storcli /*cx*/*vx* start cc force

This command starts a consistency check operation for a virtual drive. Typically, a consistency check operation is run on an initialized virtual drive. Use the `force` option to run a consistency check on an uninitialized drive.

Input example:

```
storcli /c0/v4 start cc
```

storcli /*cx*/*vx* stop cc

This command stops a consistency check operation. You can run this command only for a virtual drive that has a consistency check operation running.

Input example:

```
storcli /c0/v4 stop cc
```

NOTE You cannot resume a stopped consistency check process.

4.4.4.9 Background Initialization Commands

The Storage Command Line Tool supports the following commands for background initialization:

```
storcli /cx/vx resume bgi
storcli /cx/vx set autobgi=<on|off>
storcli /cx/vx show autobgi
storcli /cx/vx show bgi
storcli /cx/vx stop bgi
storcli /cx/vx suspend bgi
```

The detailed description for each command follows.

storcli /*cx*/*vx* resume bgi

This command resumes a suspended background initialization operation.

Input example:

```
storcli /c0/v0 resume bgi
```

storcli /*cx*/*vx* set autobgi=<on|off>

This command sets the auto background initialization setting for a virtual drive to on or off.

Input example:

```
storcli /c0/v0 set autobgi=on
```

storcli /*cx*/*vx* show autobgi

This command shows the background initialization setting for a virtual drive.

The estimated time (in minutes) left to complete the operation is also shown.

Input example:

```
storcli /c0/v0 show autobgi
```

storcli /*cx*/*vx* show bgi

This command shows the background initialization progress on the specified virtual drive in percentage.

The estimated time (in minutes) left to complete the operation is also shown.

Input example:

```
storcli /c0/v0 show bgi
```

storcli /*cx*/*vx* stop bgi

This command stops a background initialization operation. You can run this command only for a virtual drive that is currently initialized.

Input example:

```
storcli /c0/v4 stop bgi
```

storcli /*cx*/*vx* pause bgi

This command suspends a background initialization operation. You can run this command only for a virtual drive that is currently initialized.

Input example:

```
storcli /c0/v4 pause bgi
```

4.4.4.10 Virtual Drive Expansion Commands

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

```
storcli /cx/vx expand size=<value> [expandarray]
storcli /cx/vx|vall show expansion
```

The detailed description for each command follows.

storcli /*cx*/*vx* expand size=<*value*> [expandarray]

This command expands the virtual drive within the existing array or if you replace the drives with drives larger than the size of the existing array. The value of the expand size is in GB. If the `expandarray` option is specified, the existing array is expanded. If this option is not specified, the virtual drive is expanded.

storcli /*cx*/*vx* show expansion

This command shows the expansion information on the virtual drive with and without array expansion.

Input example:

```
storcli /c0/v0 show expansion
```

4.4.4.11 Display the Bad Block Table

The Storage Command Line Tool supports the following command to check for bad block entries of virtual drives on the selected controller:

```
storcli /cx/vx show bbmt
```

Input example:

```
storcli /c0/v0 show bbmt
```

4.4.4.12 Clear the LDBBM Table Entires

The Storage Command Line Tool supports the following command to clear the LDBBM table entries:

```
storcli /cx/vx delete bbmt
```

Input example:

```
storcli /c0/v0 delete bbmt
```

4.4.5 Foreign Configurations Commands

The Storage Command Line Tool supports the following commands to view, import, and delete foreign configurations:

```
storcli /cx/fall|fall del|delete [ securitykey=sssssssssss ]
storcli /cx/fall|fall import [preview][ securitykey=sssssssssss ]
storcli /cx/fall|fall show [all] [ securitykey=sssssssssss ]
```

NOTE Provide the security key when importing a locked foreign configuration created in a different machine that is encrypted with a security key.

The detailed description for each command follows.

storcli /*cx*/*fall*|fall del| delete [securitykey=*sssssssssss*]

This command deletes the foreign configuration of a controller. Input the security key if the controller is secured.

Input example:

```
storcli /c0/fall delete
```

storcli /*cx*/*fall*|fall import [preview] [securitykey=*sssssssssss*]

This command imports the foreign configurations of a controller. The `preview` option shows a summary of the foreign configuration before importing it.

Input example:

```
storcli /c0/fall import
```

storcli /*cx*/*fall*|fall show [all][securitykey=*sssssssssss*]

This command shows the summary of the entire foreign configuration for a particular controller. The `all` option shows all the information of the entire foreign configuration.

NOTE The EID:Slot column is populated for the foreign PDs that are locked.

Input example:

```
storcli /c0/fall show preview
storcli /c0/fall import preview
storcli /c0/fall show all
```

4.4.6 BIOS-Related Commands

The Storage Command Line Tool supports the following BIOS commands:

```
storcli /cx set
bios[state=<on|off>[mode=soe|be|hcoe|hsm][abs=on|off][deviceexposure=<value>]
```

The detailed description for the command follows.

storcli /*cx* set bios [state=<on|off>][mode=soe|be|hcoe|hsm][abs=on|off][deviceexposure=<value>]

This command enables or disables the ServeRAID controller's BIOS, sets the BIOS boot mode, and enables the BIOS to select the best logical drive as the boot drive.

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

Input example:

```
storcli /c0 set bios[state=on][mode=soe][abs=on][deviceexposure=20]
```

4.4.6.1 OPROM BIOS Commands

The Storage Command Line Tool supports the following OPROM BIOS commands:

```
storcli /cx/ex/sx set bootdrive=on|off
storcli /cx/vx set bootdrive=on|off
storcli /cx show bootdrive
```

The detailed description for each command follows.

storcli /*cx*/e*x*/s*x* set bootdrive=on|off

This command sets the specified physical drive as the boot drive. During the next reboot, the BIOS looks for a boot sector in the specified physical drive.

Input example:

```
storcli /c0/e32/s4 set bootdrive=on
```

storcli /*cx*/v*x* set bootdrive=on|off

This command sets the specified virtual drive as the boot drive. During the next reboot, the BIOS looks for a boot sector in the specified virtual drive.

Input example:

```
storcli /c0/v0 set bootdrive=on
```

storcli/*cx*/v*x* show bootdrive

This command shows the boot drive for the controller. The boot drive can be a physical drive or a virtual drive.

Input example:

```
storcli /c0/v0 show bootdrive
```

4.4.7 Drive Group Commands

This section describes the drive group commands.

4.4.7.1 Drive Group Show Commands

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

```
storcli /cx/dall show
storcli /cx/dall show all
storcli /cx/dall show cachecade
storcli /cx/dx show
storcli /cx/dx show all
storcli /cx/dx set security=on
storcli /cx/dx split mirror
storcli /cx/dall show mirror
storcli /cx/dall add mirror src=<val>[force]
storcli /cx/dx set hidden=<on|off>
```

storcli /*cx*/dall show

This command shows the topology information of all the drive group.

Input example:

```
storcli /c0/dall show
```

storcli /*cx*/dall show all

This command shows all available configurations in the controller which includes topology information, virtual drive information, physical drive information, free space, and free slot information.

Input example:

```
storcli /c0/dall show all
```

storcli /*cx*/dall show cachecade

This command shows all CacheCade virtual drive information.

Input example:

```
storcli /c0/dall show cachecade
```

storcli /*cx*/dx show

This command shows the topology information of the drive group.

Input example:

```
storcli /c0/dx show
```

storcli /*cx*/dx show all

This command shows the physical drive and the virtual drive information for the drive group.

Input example:

```
storcli /c0/dx show all
```

storcli /*cx*/dx set security=on

This command enables security on the specified drive group.

Input example:

```
storcli /c0/dx set security=on all
```

storcli /*cx*/dx split mirror

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

Input example:

```
storcli /c0/dx split mirror
```

storcli /cx/dall show mirror

This command shows information about the mirror associated with the drive group.

Input example:

```
storcli /c0/dall show mirror
```

storcli /cx/dall add mirror src=<val>[force]

This command joins the virtual drive with its mirror. The possible values to be used are 0, 1, or 2.

Input example:

```
storcli /c0/dall add mirror src=<1>[force]
```

storcli /cx/dx set hidden=<on|off>

This command hides or unhides a drive group.

Input example:

```
storcli /c0/d0 set hidden=on
```

4.4.8 Dimmer Switch Commands

4.4.8.1 Change Virtual Drive Power Settings Commands

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

storcli /*cx*/v*x* set ds=<default | auto | none | max | maxnocache>

This command changes the power-saving properties on an unconfigured drive and a hot spare drive. See `dimmerswitch` in the following table for values.

Input example:

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

You can use the following combinations for the Dimmer Switch commands:

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

The following table describes the power-saving options.

Table 15 Dimmer Switch Input Options

Option	Value Range	Description
`dimmerswitch` or `ds`	`on|off`	Turns the Dimmer Switch option on.
`type`	`1`: Unconfigured `2`: Hot spare `4`: All of the drives (unconfigured drives and hot spare drives).	Specifies the type of drives that the Dimmer Switch feature is applicable. By default, it is activated for unconfigured drives and hot spare drives.
`defaultldtype`	`auto`: Logical device power savings are managed by the firmware. `none`: No power saving policy. `max`: Logical device uses maximum power savings. `maxnocache`: Logical device does not cache write to maximise power savings.	Specifies the default logical drive type that is created by the Dimmer Switch option; set to none automatically.
`properties`	`disableldps`: Interval in hours or time in *hh:mm* format `spinupdrivecount`: Valid enclosure number (0 to 255) `SpinUpEncDelay`: Valid time in seconds	Sets the interval or time in which the power-saving policy for the logical drive is turned off. Specifies the number of drives in the enclosure that are spun up. Specifies the delay of spin-up groups within an enclosure in seconds.

storcli/*cx* show DimmerSwitch(ds)

This command shows the current Dimmer Switch setting for the controller.

Input example:

```
storcli/c0 show ds
```

4.4.9 BBU Commands

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

NOTE To increase the life of a battery, the battery is not fully charged. Band Gap charging keeps the maximum battery charge within a band comfortably above the data retention time requirement instead of keeping the battery charged to the maximum level. However, when a learn cycle is required, the battery is fully charged because a learn cycle starts only once the battery is fully charged.

```
storcli /cx/bbu show
storcli /cx/bbu show all
storcli /cx/bbu set autolearnmode=<value>
storcli /cx/bbu set bbuMode=<value>
storcli /cx/bbu set learndelayinterval=<value>
storcli /cx/bbu set powermode=sleep
storcli /cx/bbu set writeaceess=sealed
storcli /cx/bbu set learnStartTime=[DDD HH|off]
storcli /cx/bbu show modes
storcli /cx/bbu show properties
storcli /cx/bbu show status
storcli /cx/bbu start learn
```

```
storcli /cx/bbu start retentiontest
```

The detailed description for each command follows.

storcli /*cx*/bbu show

This command shows the summary information for the BBU of a controller.

Input example:

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

storcli /*cx*/bbu show all

This command shows all the information of the BBU.

Input example:

```
storcli /c0/bbu show all
```

storcli /*cx*/bbu set autolearnmode=<value>

This command starts the automatic learn cycle on the battery. The possible values are **0** - Enabled, **1**- Disabled, and **2** - WarnViaEvent.

Input example:

```
storcli /c0/bbu set autolearnmode=0
```

storcli /*cx*/bbu set bbuMode=<*value*>

This command sets the BBU mode for the BBU. The following table shows the various BBU modes:

Table 16 BBU Mode

Mode	Description
0	48 hours of retention[a] at 60 °C, 1-year Service Life.
1	12 hours of retention at 45 °C, 5-year Service Life, transparent learn.[b]
2	12 hours of retention at 55 °C, 3-year Service Life, transparent learn.
3	24 hours of retention at 45 °C, 3-year Service Life, transparent learn.
4	48 hours of retention at 45 °C, 3-year Service Life.
5	48 hours of retention at 55 °C, 1-year Service Life.
6	Same as the description for BBU mode 5. The BBU mode 6 enables you to receive events when the battery capacity reaches suboptimal and critical thresholds.

a. Indicates how long the battery can hold data in the controller's memory in case of accidental system shutdown.

b. The controller's performance is not affected during the battery's learn cycle.

Input example:

```
storcli /c0/bbu set bbuMode=2
```

NOTE BBU modes are supported on any iBBU08/09 BBU/controller combo and later-generation controllers.

storcli /*cx*/bbu set learndelayinterval=<*value*>

This command sets the learn delay interval for the BBU in hours. The value must be between 0 to 168 hours (7 days).

Input example:

```
storcli /c0/bbu set learnDelayInterval=30
```

storcli /*cx*/bbu set powermode=sleep

This command places the battery in low-power storage mode. The battery automatically exits this state after 5 seconds.

Input example:

```
storcli /c0/bbu set powermode=sleep
```

storcli /*cx*/bbu set writeaccess=sealed

This command seals the gas gauge EEPROM write access.

NOTE Use the `set writeaccess=sealed` command at manufacturing time.

Input example:

```
storcli /c0/bbu set writeaccess=sealed
```

storcli /*cx*/bbu set writeaccess=sealed

This command seals the gas gauge EEPROM write access.

NOTE Use the `set writeaccess=sealed` command at manufacturing time.

Input example:

```
storcli /c0/bbu set writeaccess=sealed
```

storcli /*cx*/bbu set learnStartTime=[DDD HH| off]

This command sets the learn start time for the BBU in hours on the day specified. **DDD** refers to the day of the week (SUN,MON,....SAT), **HH** refers to the hours (0-23 hours), and **off** sets the learn start to off.

Input example:

```
storcli /c0/bbu set learnStartTime=MON 12
```

storcli /*cx*/bbu show properties

This command shows the BBU Learn properties for a controller.

Input example:

```
storcli /c0/bbu show properties
```

storcli /*cx*/bbu show status

This command shows the battery information, firmware status, and the gas gauge status.

Input example:

```
storcli /c0/bbu show status
```

storcli /*cx*/bbu start learn

This command starts the BBU learning cycle. The battery learn cycle is immediately started and no other parameters are required for this command.

Input example:

```
storcli /c0/bbu start learn
```

storcli /*cx*/bbu start retentiontest

This command starts the battery retention test. This command requires you to reboot your system.

Input example:

```
storcli /c0/bbu start retentiontest
```

4.4.10 CacheVault Commands

The Storage Command Line Tool supports the following CacheVault® command:

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

storcli /*cx*/cv show all

This command shows all the information about a CacheVault module that is connected to a controller.

NOTE This command works only when a CacheVault module is connected to the controller; otherwise, an error message appears.

Input example:

```
storcli /c0/cv show all
```

4.4.11 Enclosure Commands

The Storage Command Line Tool supports the following enclosure commands:

```
storcli /cx/ex download src=filepath[forceActivate]
storcli /cx/ex show all
storcli /cx/ex show status
```

The detailed description for each command follows.

storcli /*cx*/*ex* download src=filepath [forceactivate]

This command flashes the firmware with the file specified at the command line. The enclosure performs an error check after the operation. The following option can be used with the enclosure firmware download command.

Table 17 Enclosure Firmware Download Command Options

Option	Value Range	Description
forceactivate	—	Issues a command descriptor block (CDB) with write command with no data with command mode 0x0F (flash download already in progress). **NOTE** This option is used primarily to activate Scotch Valley Enclosures.

NOTE The firmware file that is used to flash the enclosure can be of any format. The StorCLI utility assumes that you provide a valid firmware image.

Input example:

```
storcli /c0/e0 download src=c:\file2.bin
```

storcli /*cx*/*ex* show all

This command shows all enclosure information, which includes general enclosure information, enclosure inquiry data, a count of enclosure elements, and information about the enclosure elements.

Input example:

```
storcli /c0/e0 show all
```

storcli /*cx*/*ex* show status

This command shows the enclosure status and the status of all the enclosure elements.

Input example:

```
storcli /c0/e0 show status
```

4.4.12 PHY Commands

The Storage Command Line Tool supports the following PHY commands:

```
storcli /cx/px|pall set linkspeed=0(auto)|1.5|3|6|12
storcli /cx/px|pall show
storcli /cx/px|pall show all
storcli /cx/ex show phyerrorcounters
storcli /cx[/ex]/sx show phyerrorcounters
storcli /cx[/ex]/sx reset phyerrorcounters
```

The detailed description for each command follows.

storcli /*cx*/*px*|pall set linkspeed=0(auto)|1.5|3|6|12

This command sets the PHY link speed. You can set the speed to 1.5 Gb/s, 3 Gb/s, 6 Gb/s, or 12 Gb/s. The linkspeed is set to `auto` when you specify `linkspeed = 0`.

Input example:

```
storcli /c0/p0 set linkspeed=1.5
```

storcli /*cx*/*px*|pall show

This command shows the basic PHY layer information.

Input example:

```
storcli /c1/p0 show
```

storcli /*cx*/*px*|pall show all

This command shows all the PHY layer information.

Input example:

```
storcli /c1/p0 show all
```

storcli /cx/ex show phyerrorcounters

This command shows the enclosure/expander PHY error counters.

Input example:

```
storcli /c1/e0 show phyerrorcounters
```

storcli /cx[/ex]/sx show phyerrorcounters

This command shows the drive PHY error counters.

Input example:

```
storcli /c1/e0/s0 show phyerrorcounters
```

storcli /cx[/ex]/sx reset phyerrorcounters

This command resets the drive PHY error counters.

Input example:

```
storcli /c1/e0/s0 reset phyerrorcounters
```

4.4.13 Logging Commands

The Storage Command Line Tool supports the following commands to generate and maintain log files:

```
storcli /cx clear events
storcli /cx delete termlog
storcli /cx show events file=<absolute path>
storcli /cx show eventloginfo
storcli /cx show termlog type=config|contents
storcli /cx show dequeue log file =<filepath>
```

The detailed description for each command follows.

storcli /*cx* delete events

This command deletes all records in the event log.

Input example:

```
storcli /c0 delete events
```

storcli /*cx* delete termlog

This command clears the TTY (firmware log for issue troubleshooting) logs.

Input example:

```
storcli /c0 delete termlog
```

storcli /*cx* show events file=<*absolute path*>

This command prints the system log to a text file and saves the file in the specified location.

Input example:

```
storcli /c0 show events file=C:\Users\brohan\test\eventreports
```

storcli /*cx* show eventloginfo

This command shows the history of log files generated.

Input example:

```
storcli /c0 show eventloginfo type=config
```

storcli /cx show termlog type=config|contents

This command shows the firmware logs. The `config` option shows the term log configuration (settings of TTY BBU buffering), the `contents` option shows the term log. The `contents` option is the default.

Input example:

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

storcli /cx show dequeue log =<filepath>

This command shows the debug log from the firmware.

Input example:

```
storcli /c0 show dequeue log=<c:\test\log.txt>
```

4.4.14 Automated Physical Drive Caching Commands

The Storage Command Line Tool supports the following automated physical drive caching commands:

```
storcli /cx set autopdcache=<off|r0>[immediate]
storcli /cx show autopdcache
```

The detailed description for each command follows.

storcli /cx set autopdcache=<off|r0>[immediate]

This command allows you to set the controller's automated physical drive cache policy to RAID 0. When set to RAID-0, all un-configured physical drives are configured as a single RAID 0 drive, until the maximum virtual drive limit is reached. The `immediate` option lets this command execute the conversion (to RAID0) operation only on all the existing physical drives. Any newly physical drives connected in the future do not get converted to RAID 0. If you omit the `immediate` option in this command, conversion to RAID 0 takes place on newly connected physical drives too. Automatic conversion to RAID0 can be turned off by setting the autopdcache policy to `off`.

Input example:

```
storcli /c0 set autopdcache=r0 immediate
```

storcli /cx show autopdcache

This command lets you view the automatic physical drive caching property.

Input example:

```
storcli /c0 show autopdcache
```

4.5 Frequently Used Tasks

4.5.1 Showing the Version of the Storage Command Line Tool

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

```
Storcli -v
```

4.5.2 Showing StorCLI Help

The following command shows the command line tool help:

```
Storcli -h
```

Help appears for all the StorCLI commands.

4.5.3 Showing System Summary Information

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

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

4.5.4 Showing Free Space in a Controller

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

```
Storcli /cx show freespace
```

4.5.5 Adding Virtual Drives

The following command creates a virtual drive:

```
Storcli /cx add vd type=raid[0|1|5|6|10|50|60][Size=<VD1_Sz>,<VD2_Sz>,..|*all]
[name=<VDNAME1>,..] drives=e:s|e:s-x|e:s-x,y [PDperArray=x|auto*]
[SED] [pdcache=on|off|*default][pi] [DimmerSwitch(ds)=default|automatic(auto)|
*none|maximum(max)|MaximumWithoutCaching(maxnocache)] [wt|*wb|awb] [nora|*ra]
[*direct|cached]
[strip=<8|16|32|64|128|256|512|1024] [AfterVd=x] [Spares=[e:]s|[e:]s-x|[e:]s-x,y]
[Cbsize = 0|1|2 Cbmode = 0|1|2]
[force]
```

The following inputs can be used when adding virtual drives:

- The controller in which the virtual drives are created.
- The RAID type of the virtual drives. The supported RAID types are 0, 1, 5, 6, 10, 50, 60.
- The size of each virtual drive.
- The drives that are used to create the virtual drives.

  drives = *e*:*s*|*e*:*s*-*x*|*e*:*s*-*x*,*y*

  Where:
  - *e* specifies the enclosure id.
  - *s* represents the slot in the enclosure.
  - *e*:*s*-*ex* is the range conventions used to represents slots s to x in the enclosure e.
- The physical drives per array. The physical drives per array can be set to a particular value.
- The `SED` option creates security-enabled drives.
- The `PDcache` option can be set to `on` or `off`.
- The `pi` option enables protection information.
- The Dimmer Switch is the power save policy. It can be set to default or `automatic *`,`none`, `maximum(max)`, or `MaximumWithoutCaching(maxnocache)`.
- The `wt` option disables write back.
- The `nora` option disables read ahead.

- The `cached` option enables the cached memory.
- The `strip` option sets the strip size. It can take the values `8`, `16`, `32`, `64`, `128`, `256`, `512`, `1024`.
- The `AfterVdX` option creates the virtual drives in the adjacent free slot next to the specified virtual drives.

NOTE The * indicates default values used in the creation of the virtual drives. If values are not specified, the default values are taken.

Example: `/cxadd vd type=r1 drives=0:10-15 WB Direct strip=64`

This command creates a RAID volume of RAID 1 type from drives in slots 10 to slot 15 in enclosure 0. The strip size is 64kb.

4.5.6 Setting the Cache Policy in a Virtual Drive

The following command sets the write cache policy of the virtual drive:

`Storcli /cx/v(x|all) set wrcache=wt|wb|awb`

The command sets the write cache to write back, write through, or always write back.

4.5.7 Showing Virtual Drive Information

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

`storcli /cx show [all]`

4.5.8 Deleting Virtual Drives

The following command deletes virtual drives:

`storcli /cx/v(x|all) del [cc|cachecade]`

The following inputs are required when deleting a virtual drive:

- The controller on which the virtual drive or virtual drives is present.
- The virtual drives that must be deleted; or you can delete all the virtual drives on the controller using the `vall` option.
- The `cc` or `cachecade` option to confirm that the deleted drive is a CacheCade drive.

4.5.9 Flashing Controller Firmware

The following command is used to flash the controller firmware.

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

For more information, see section 4.5.9 Flashing Controller Firmware.

Chapter 5: MegaRAID Storage Manager Overview and Installation

MegaRAID Storage Manager (MSM) is a configuration and monitoring utility used with ServeRAID C100 and ServeRAID C105. This chapter provides a brief overview of the MegaRAID Storage Manager utility and explains how to install it on the supported operating systems.

5.1 Overview

MegaRAID Storage Manager may be installed either on the local system which contains the RAID 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 ServeRAID controller(s) you are using.

MegaRAID Storage Manager enables you to configure, monitor, and maintain storage configurations created under ServeRAID C105. 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 ServeRAID controllers. Some MegaRAID Storage Manager features are not applicable for ServeRAID C105.

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 the 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, for the procedures used to create storage configurations.

5.1.2 Monitoring Storage Devices

MegaRAID Storage Manager displays the status of virtual disks, drives, 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

The hardware requirements for MegaRAID Storage Manager are as follows:

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

NOTE It is recommended that you use at least 1 Gbyte of system memory.

- Physical drive with at least 50 Mbytes available free space

The supported operating systems for the MegaRAID Storage Manager are as follows:

- Microsoft® Windows® Server versions 2008, 2008R2, 2012, and 2012R2
- Microsoft Windows Workstation version 7 & 8
- Red Hat® Enterprise Linux™ (RHEL) versions 5 & 6
- SuSE® Linux Enterprise Server (SLES) versions 10 & 11

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, you can either install the additional libraries or run a native 64-bit version of MSM.

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 the 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 1 Customer Information Screen

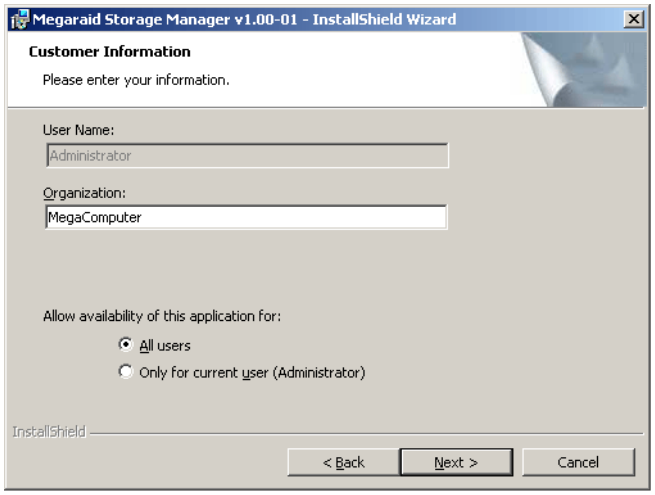

5. Enter your user name and organization name. In the bottom part of the screen, select an installation option:
   - If you select **All users**, any user with administrative privileges can use this version of MegaRAID Storage Manager to view or change storage configurations.
   - If you select **Only 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 2 Setup Type Screen

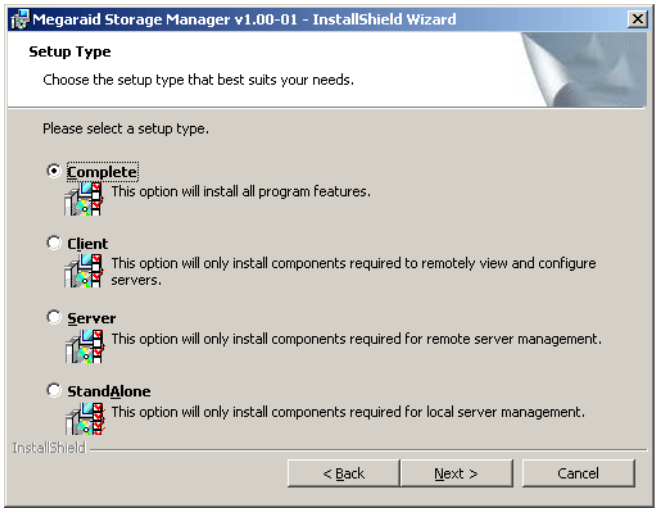

8. Select one of the Setup options. The options are fully explained in the screen text.
   - Normally, you would select **Complete** if you are installing MegaRAID Storage Manager on a server. This option installs the complete MSM.
   - Select **Client** if you are installing MegaRAID Storage Manager 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 Confgurator, 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 Confgurator, and DebugLog are not available.

— Select **StandAlone** to 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 if you need to install MegaRAID Storage Manager on a system running 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.

5.3.3 Linux Installation Messages

One or more of the following messages may 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 Screens and Menus

This chapter explains how to start MegaRAID Storage Manager and describes the MegaRAID Storage Manager screens 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 RHEL system, select **Applications > System Tools > MegaRAID Storage Manager StartupUI**.
   - To start MegaRAID Storage Manager on a SUSE system, select **Start > System > More Programs -> MegaRAID Storage Manager**.

   When the program starts, the Select Server window appears, as shown in the following figure.

Figure 1 Select Server Window

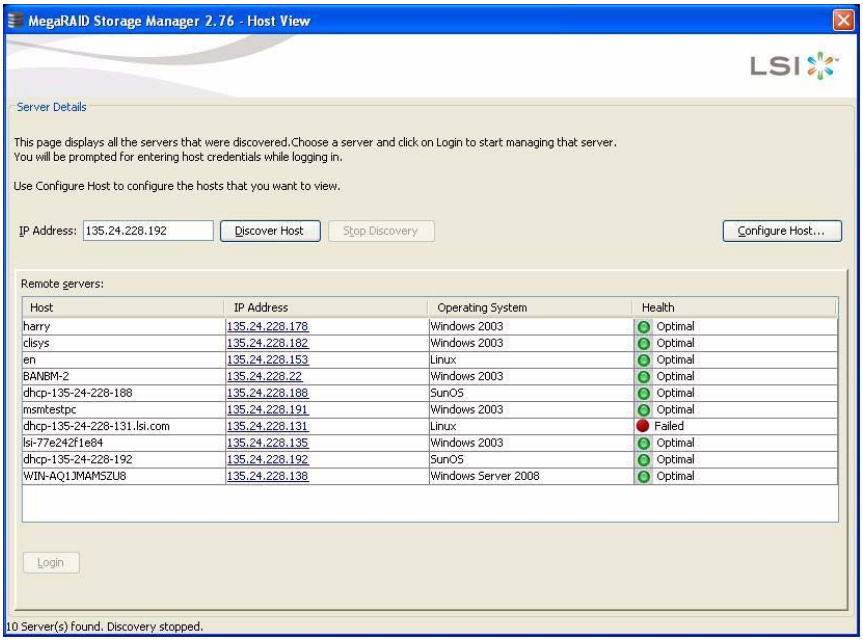

If the circle in the server icon is yellow instead of green, it means that the server is running in a degraded state—for example, because a disk drive used in a virtual disk has failed. If the circle is red, the storage configuration in the server has failed.

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 Framework** box, you can also access a standalone 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 2 Server Login Window

3. Enter your user name and your 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.
   - Select **Full Access** if you need to both view the current configuration and change the configuration.
   - Select **View Only** if you need to only view and monitor the configuration.

5. Click **Login**.

NOTE If the computer is networked, this is the login to the computer itself, not the network login.

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

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

6.2 MegaRAID Storage Manager Screen

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

Figure 3 MegaRAID Storage Manager Main Menu Screen

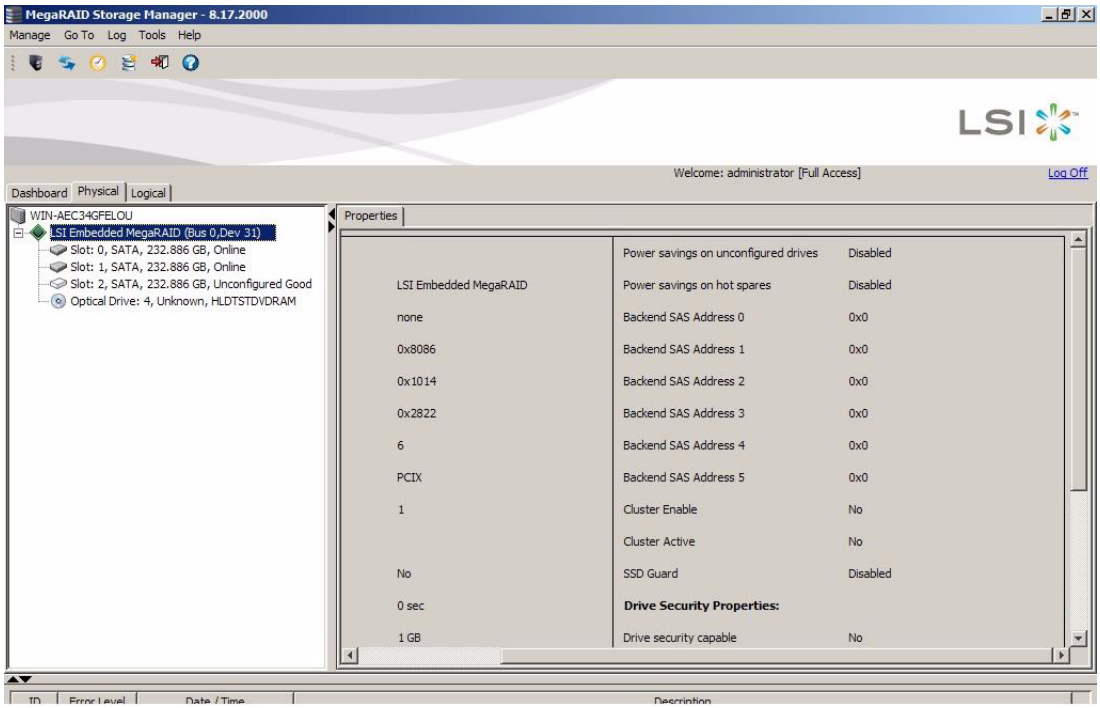

The following topics describe the panels and menu options that appear on this screen.

6.2.1 Dashboard/Physical View/Logical View

The left panel of the ServeRAID C105 main menu screen displays the *Dashboard* view, the *Physical* view, or the *Logical* view of the system and the attached 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 4 MegaRAID Storage Manager Dashboard View

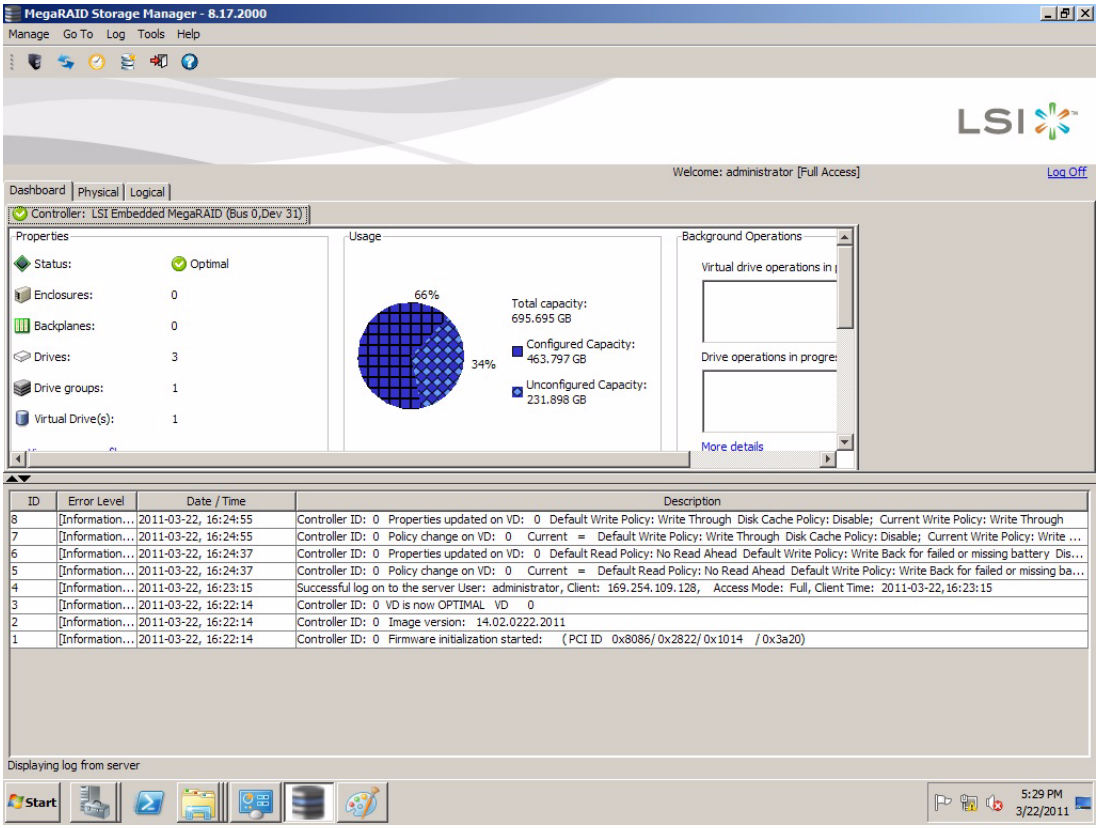

6.2.1.2 Physical view

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 item appear in the right panel of the screen under the Properties tab.

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

Figure 5 MegaRAID Storage Manager Physical View

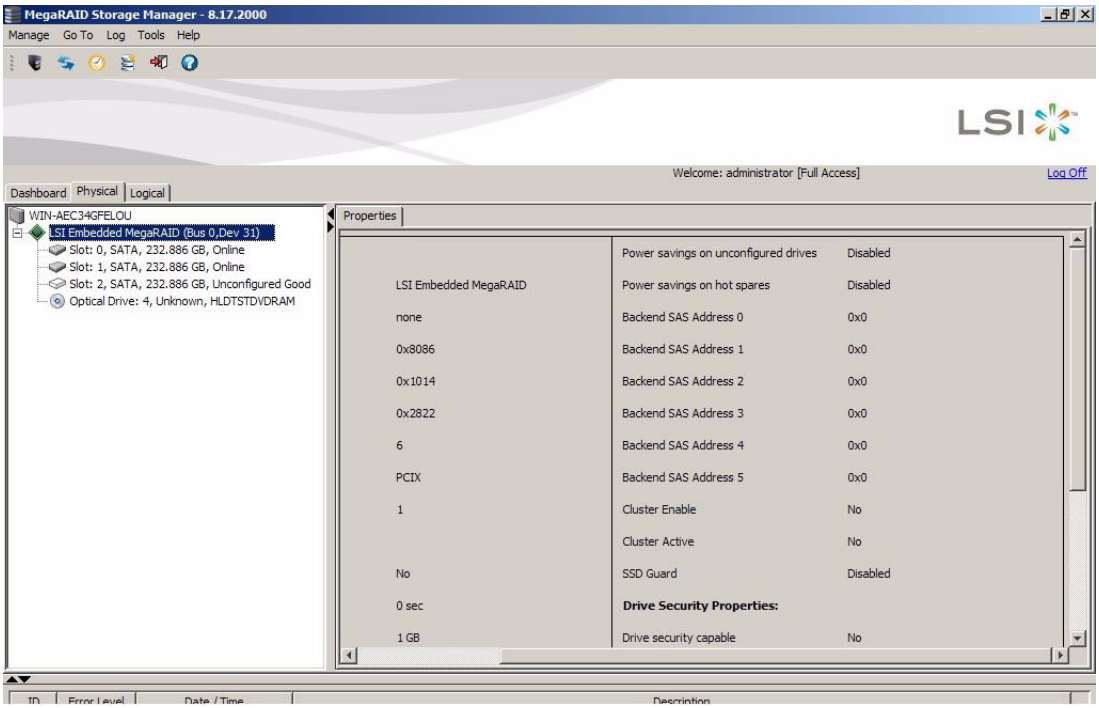

6.2.1.3 Logical View

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 and the Properties tab.

Figure 6 MegaRAID Storage Manager Logical View

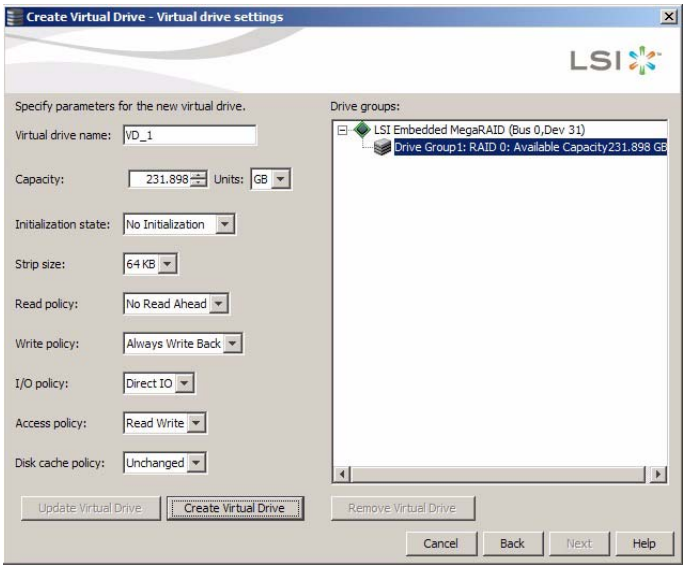

The following icons in the left panel represent the controllers, disk drives, and other devices:

-  - Status
-  - System
-  - Controller
-  - Backplane
-  - Port
-  - Drive group
-  - Virtual drive
-  - Online drive
-  - Global hotspare
-  - Tape drive
-  - CD-ROM
-  - Foreign drive
-  - Unconfigured foreign drive
-  - Unconfigured drive

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

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

A yellow circle to the right of an icon indicates that a device is running in a degraded state. For example, this icon indicates that a virtual disk is running in a degraded state because a disk 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, as shown in Figure 3. New event log entries appear during the session. Each entry has a timestamp and date, an error level indicating the severity of the event, and a brief description of the event.

For more information about the event log, see Section 8.1, Monitoring System Events. For more information about the event log entries, see Appendix A, ServeRAID C100 and ServeRAID C105 MegaRAID Storage Manager Events and Messages.

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 7, Chapter 8, and Chapter 9 of this manual.

6.2.3.1 Manage Menu

The Manage menu has a Refresh option for updating the display in the ServeRAID C105 main menu screen (refresh is seldom required; the display normally updates automatically) and an Exit option to end your session on ServeRAID C105. 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 ServeRAID C105 main menu. The options also vary depending on the current state of the selected device. For example, if you select an offline drive, the Make Drive Online option appears in the Physical Drive menu.

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

6.2.3.3 Log Menu

The Log menu includes options for saving and clearing the message log. For more information, see Appendix A, ServeRAID C100 and ServeRAID C105 MegaRAID Storage Manager Events and Messages.

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.

6.2.3.5 Help Menu

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

NOTE When you use the Online Help for MegaRAID Storage Manager, 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 ServeRAID controllers.

The controllers support RAID 0, RAID 1, and RAID 10 storage configurations. The Configuration Wizard allows you to easily create new storage configurations and modify the configurations. To learn more about RAID and RAID levels, see Section 1.5, RAID Overview.

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 Configuration Wizard to create new storage configurations. To open the MegaRAID Storage Manager Configuration Wizard, select a controller in the left panel of the MegaRAID Storage Manager main menu screen and then select **Operations > Advanced Operations > Configuration > Configuration Wizard**.

You can use the ServeRAID C105 to create new storage configurations on systems with ServeRAID controllers. You can create the following types of configurations:

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

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

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. You should change these parameters only if you have a specific reason for doing so. It is usually best to leave them at their default settings.

- **Initialization state**: Initialization prepares the storage medium for use. Specify the initialization status:
  - *No Initialization*: (the default) The new configuration is not initialized and the existing data on the drives is not overwritten.
  - *Fast Initialization*: 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.
- **Stripe size**: The setting is 64 KB only.
- **Access policy**: Select the type of data access that is allowed for this virtual drive.
  - *Read/Write*: (the default) Allow read/write access. This is the default.
  - *Read Only*: Allow read-only access.
  - *Blocked*: Do not allow access.
- **Disk cache policy**: Select a cache setting for this drive:
  - *Enabled*: Enable the disk cache.
  - *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 created spanned drives using the simple configuration procedure. To create spanned drives, use the advanced configuration procedure described in Section 7.1.3, Creating a Virtual Drive Using Advanced Configuration.

Follow these steps to create a new storage configuration in simple 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 1 Virtual Drive Creation Menu

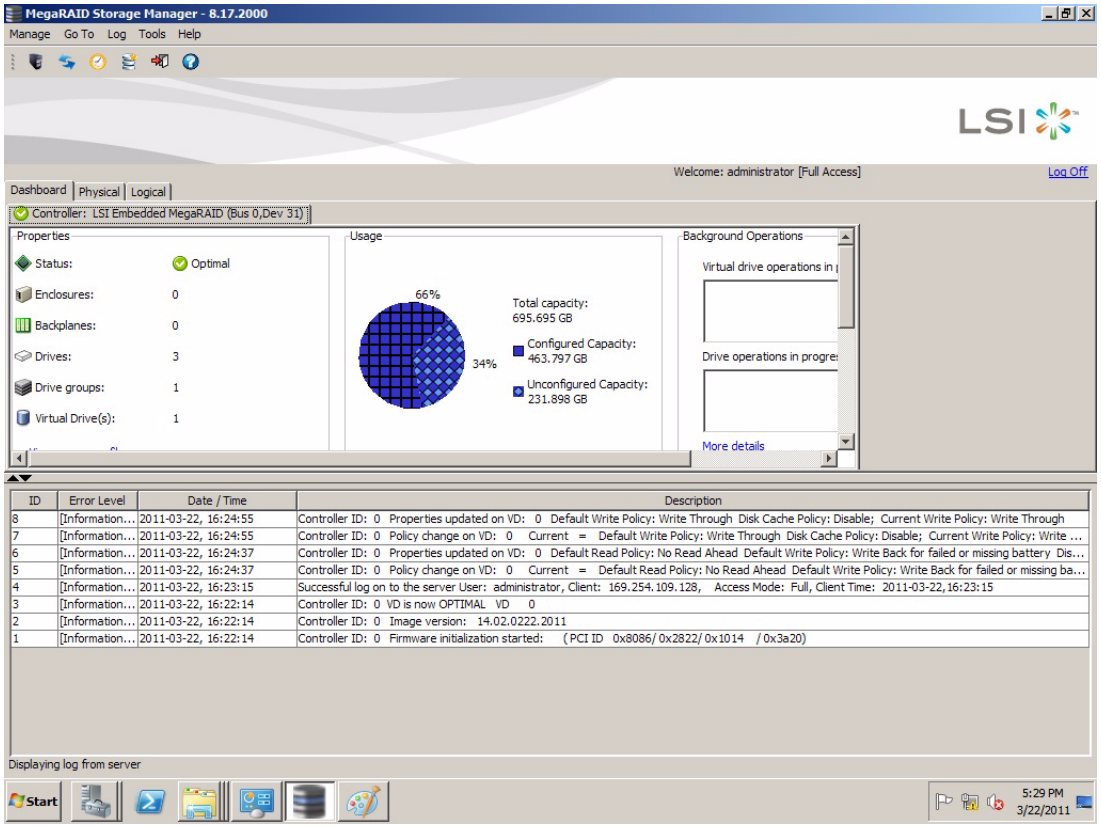

2. Click **Create Virtual Drive** in the Actions section of the screen.
   The dialog box for the configuration mode (Simple or Advanced) appears, as shown in the following figure.

Figure 2 Virtual Drive Simple Configuration Mode

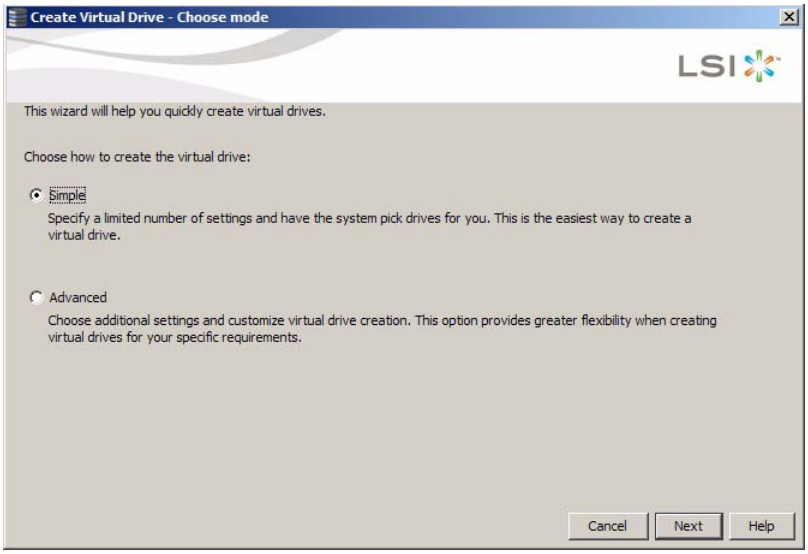

3. Click **Simple** and press **Next**.

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

Figure 3 Create Virtual Drive Screen

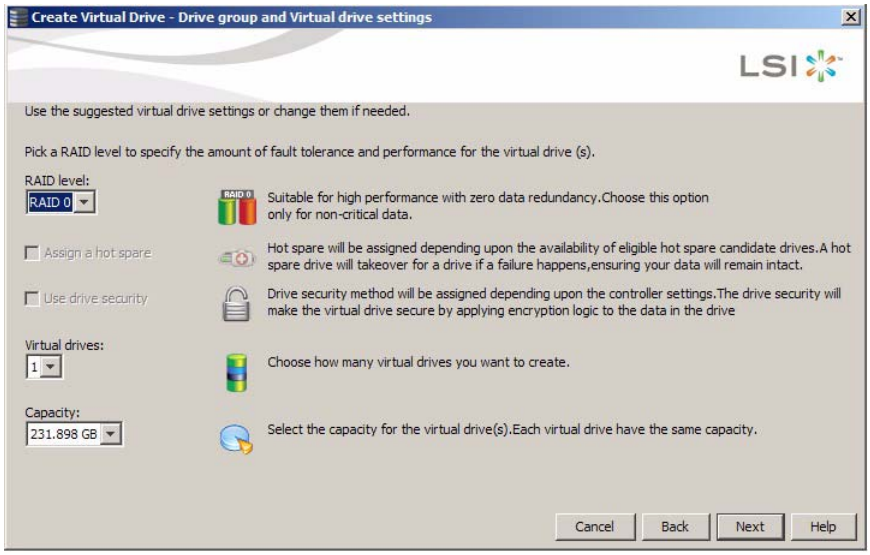

4. Select the RAID level desired for the virtual drive.

   When you use simple configuration, the RAID controller supports RAID level 1. In addition, it supports independent drives (configured as RAID 0). The screen text gives a brief description of the RAID level that you select. The RAID levels that you can choose depend on the number of drives available. To learn more about RAID levels, see Section 1.5, RAID Overview.
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 4 Create Virtual Drive - Summary Window

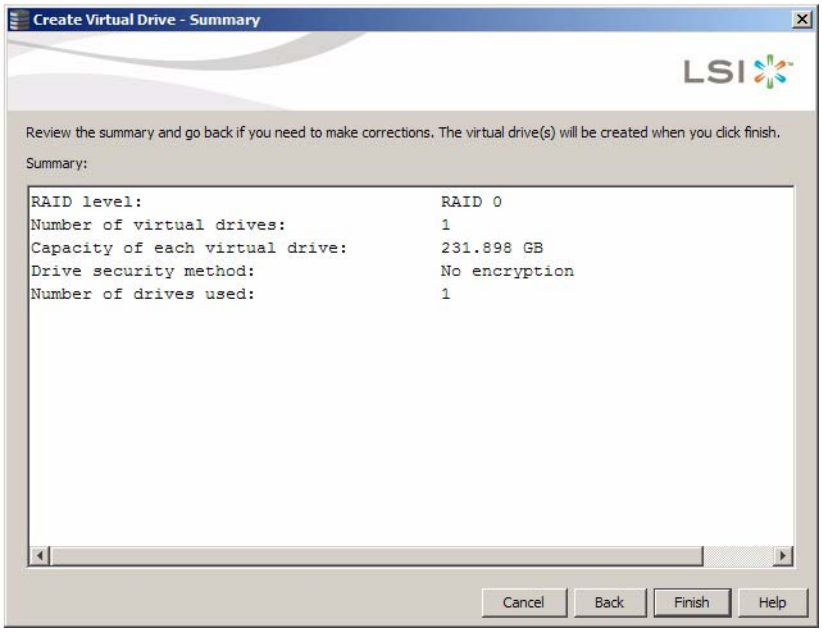

8. Click **Back** to return to the previous screen to change any selections or click **Finish** to accept and complete the configuration.

   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. This example shows the configuration of a spanned drive group.

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 5 Virtual Drive Creation Menu

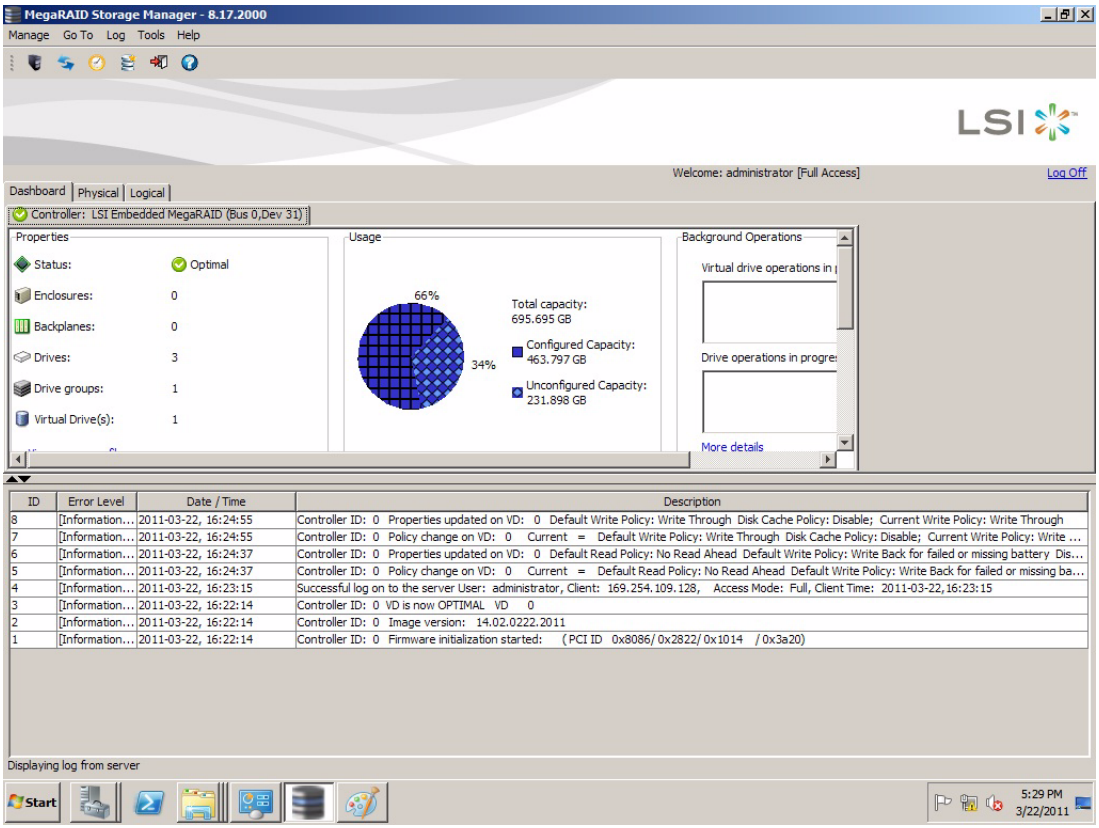

2. Click **Create Virtual Drive** in the Actions section of the screen.
   The dialog box for the configuration mode (Simple or Advanced) appears, as shown in the following figure.

Figure 6 Virtual Drive Advanced Configuration Mode

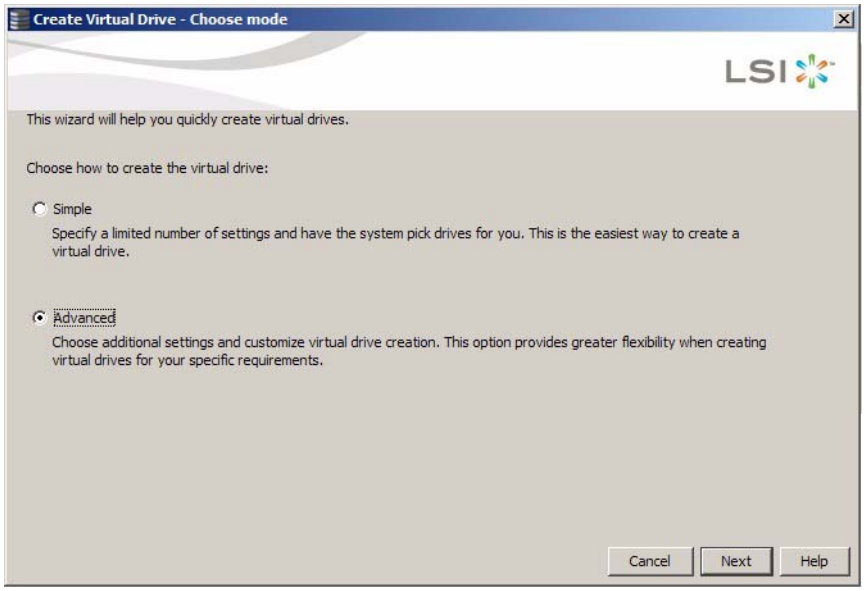

3. Click **Advanced** and then click **Next**.

   The Create Drive Group Settings dialog box appears, as shown in the following figure.

Figure 7 Create Drive Group Settings Dialog Box

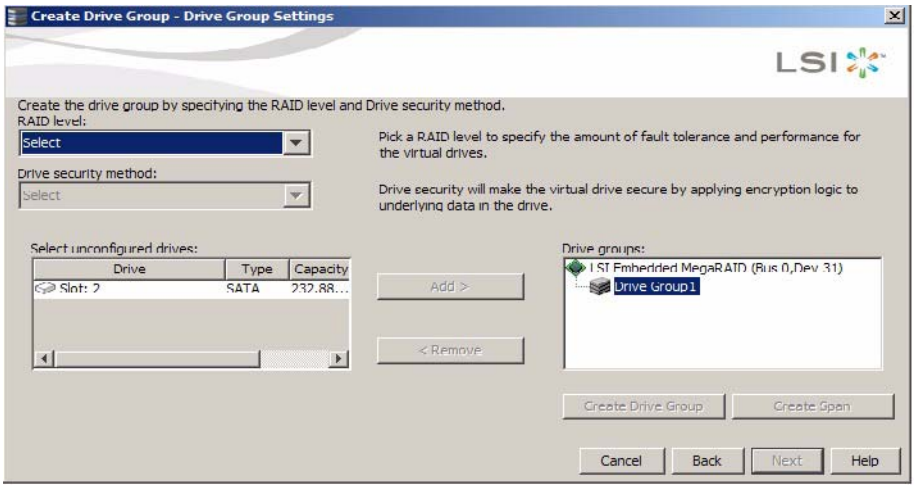

4. Select the following items on the Create Drive Group Settings dialog box:
   a. Select the RAID level desired for the drive group from the drop-down menu.

      Drive Group 1 appears in the **Drive groups** field.

      The RAID controller supports RAID levels 1 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. 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 **Span {number}** below **Drive Group {number}**.
   c. If you are creating a RAID 10 drive group, perform the following steps; if not, go to step g.

      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.

   d. Click **Create Span** to create a second span in the drive group.
   e. Select *unconfigured* drives from the list of drives and click **Add>** to add them to the second drive group.

      The selected drives appear in the second span under **Span {number}** below the second drive group, **Drive Group {number}**.
   f. Click **Create Drive Group** to make a drive group with the spans.
   g. Click **Next** to complete the steps for RAID 10 configuration.
   h. Click **Create Drive Group** to make a drive group.
   i. 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 8 Virtual Drive Settings Dialog Box

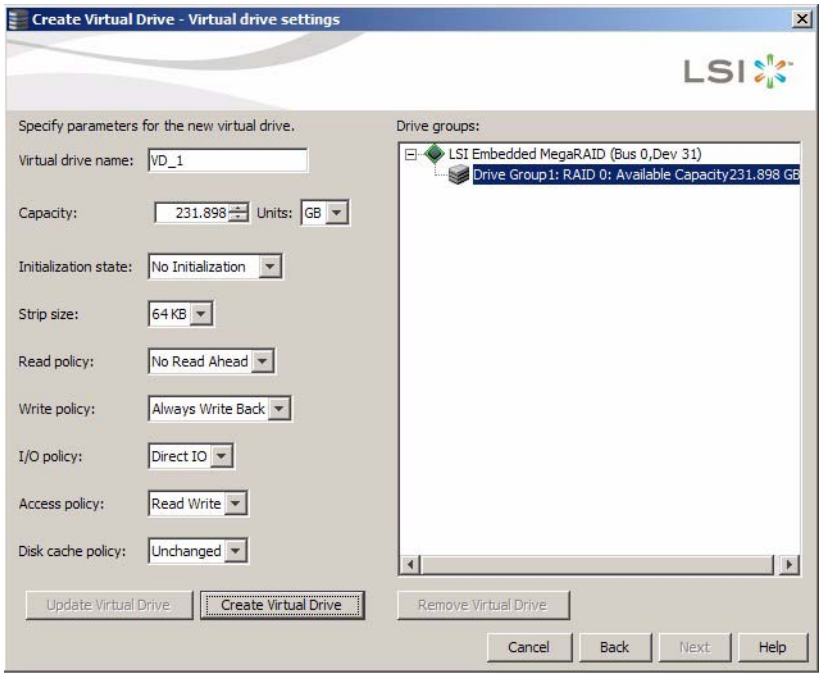

j. Change any virtual drive settings, if desired.

See Section 7.1.1, Selecting Virtual Drive Settings, for more information about the virtual drive settings.

5. Click **Create Virtual Drive**.

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

6. Click **Next**.

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

Figure 9 Virtual Drive Summary Window

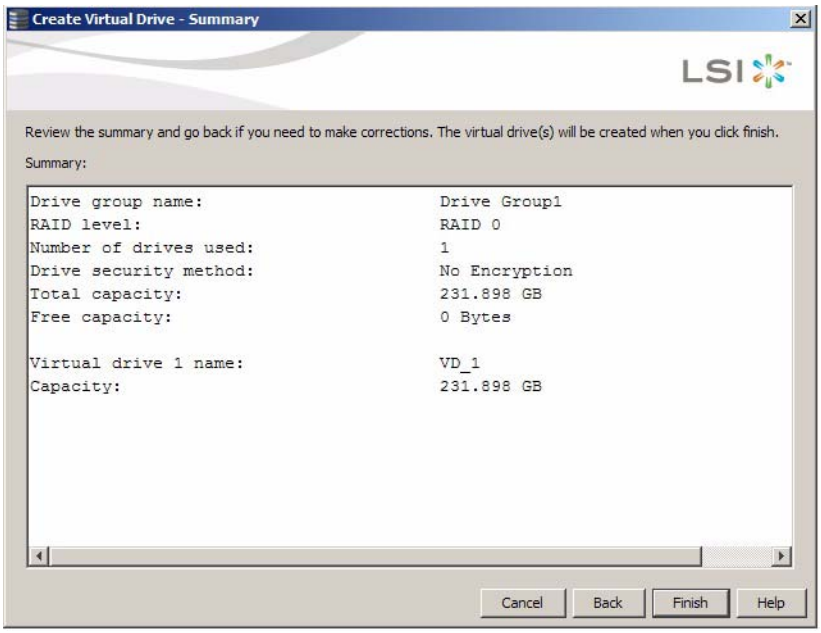

7. Click **Back** to return to the previous screen to change any selections or click **Finish** to accept and complete the configuration.

   After you click **Finish**, the new storage configuration is created and initialized.

   After the configuration is completed, a dialog box notifies you that the virtual drives were created successfully. If more drive capacity exists, the dialog box asks whether you want to create more virtual drives. If no more drive capacity exists, you are prompted to close the configuration session.
8. Select **Yes** or **No** to indicate whether you want to create additional virtual drives.

   If you select **Yes**, the system takes you to the Create Virtual Drive screen, as shown in Figure 3. If you select **No**, the utility asks whether you want to close the wizard.
9. If you selected **No** in the previous step, select **Yes** or **No** to indicate whether you want to close the wizard.

   If you select **Yes**, the configuration procedure closes. If you select **No**, the dialog box closes and you remain on the same page.

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

NOTE Select a controller icon in the **Physical** tab or the **Logical** tab in the left panel of the ServeRAID C105 main menu screen.

1. Select **Go To > Controller > Set Adjustable Task Rates** from the menu bar, as shown in the following figure.

Figure 10 Set Adjustable Task Rates Menu

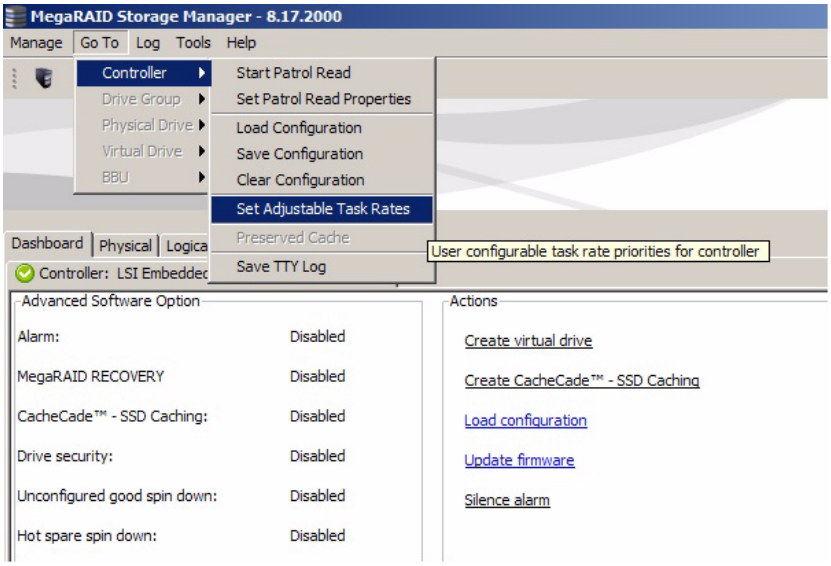

The Set Adjustable Task Rates dialog box appears, as shown in the following figure.

Figure 11 Set Adjustable Task Rates Dialog Box

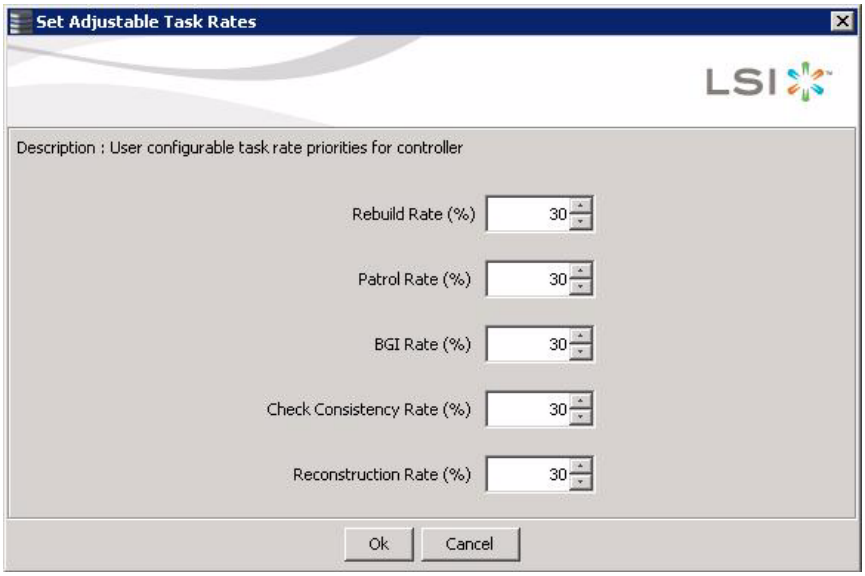

2. Enter changes, as needed, to the following task rates:
   - Rebuild Rate
   - Patrol Read
   - Background Initialization (BGI) (for fast initialization)
   - Check Consistency (for consistency checks).
   - Reconstruction

   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.
3. Click **OK** to accept the new task rates.
4. 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, Write Policy, and other virtual drive properties at any time after a virtual drive is created. Follow these steps to change the virtual drive properties.

1. Select a virtual drive icon in the **Physical** tab or the **Logical** tab in the left panel of the ServeRAID C105 main menu screen.
2. Select **Go To > Virtual Drive > Set Virtual Drive Properties** from the menu bar, as shown in the following figure.

Figure 12 Set Virtual Drive Properties Menu

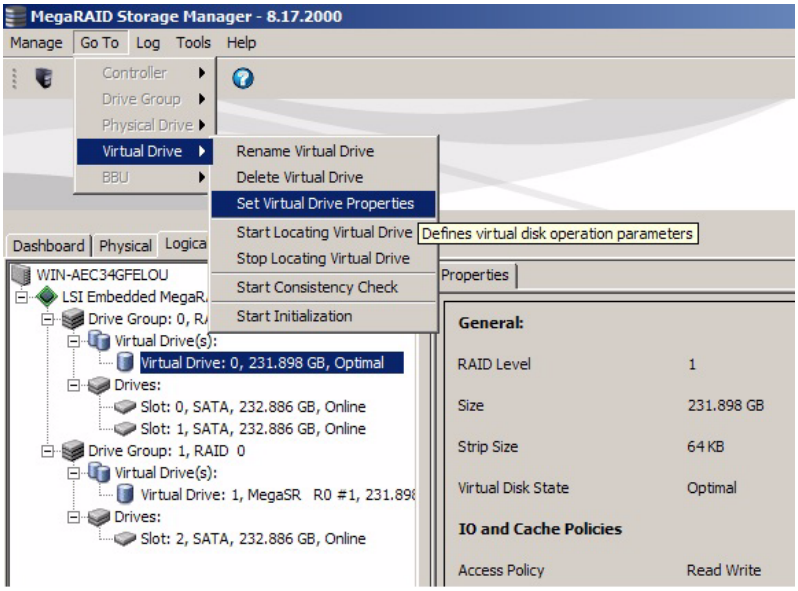

The Set Virtual Drive Properties dialog box appears, as shown in the following figure.

Figure 13 Set Virtual Drive Properties Screen

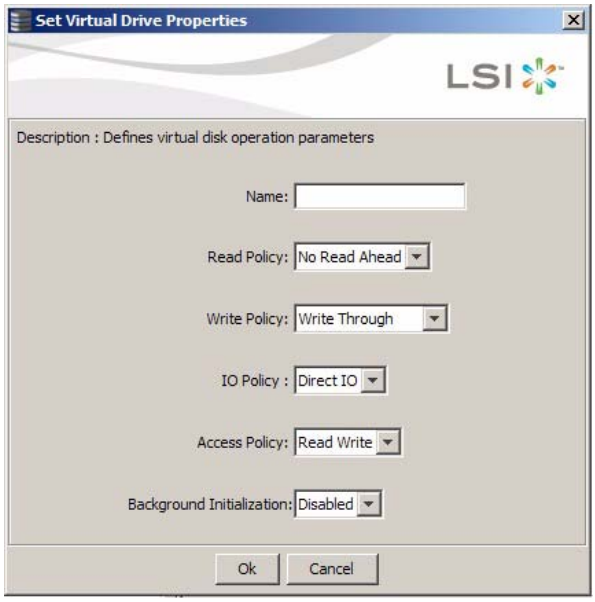

3. Change the virtual drive properties as needed.
   For information about these properties, see Section 7.1.1, Selecting Virtual Drive Settings.
4. Click **Ok** to accept the changes.

The virtual drive settings are updated.

7.4 Deleting a Virtual Drive

NOTE 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.
   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

The MegaRAID Storage Manager enables you to monitor the status of disk drives, virtual disks, and other storage devices.

8.1 Monitoring System Events

MegaRAID Storage Manager 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 disk or the removal of a physical drive) an event message appears in the log displayed 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 are for the delivery of alerts, the severity level of events, exceptions, and email settings.

Figure 1 Event Information Window

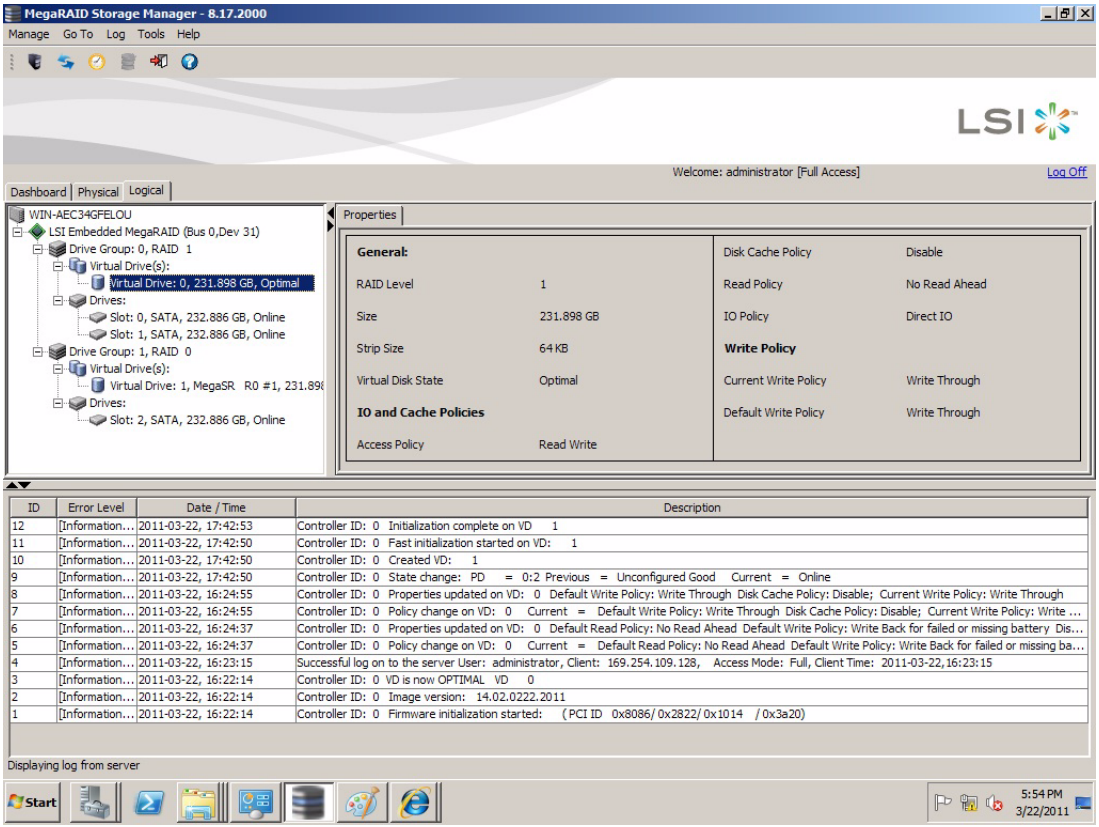

Each message that appears in the event log has a severity level that indicates the importance of the event, a date and timestamp, and a brief description. You can click an event to display the same information in a window. (For a list of all events, see Appendix A, ServeRAID C100 and ServeRAID C105 MegaRAID Storage Manager Events and Messages.)

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 2 Log Menu

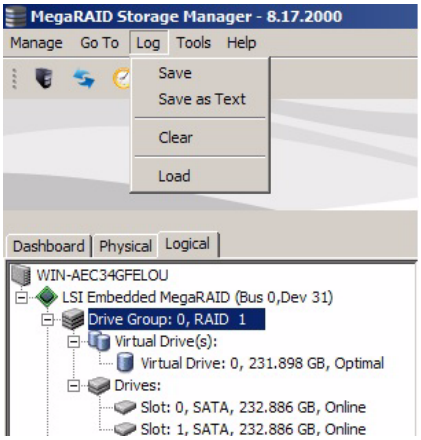

8.2 Configuring Alert Notifications

The Alert Notification Configuration feature allows you to control and configure the alerts that ServeRAID C105 sends when various system events occur.

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

Figure 3 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 4 Alerts Notification Configuration Screen

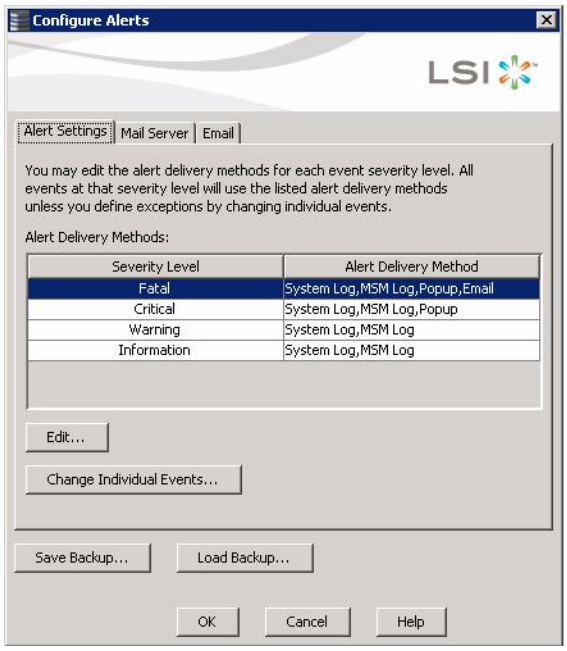

You can select the **Alert Settings** tab to perform the following actions:

- 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.

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

You can select the **Mail Server** tab to perform the following actions:

- Enter or edit the sender e-mail address.
- Enter the SMTP server.
- Require 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.

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

You can select the **Email** tab to perform the following actions:

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

CAUTION 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 **Alerts Setting** tab.
2. Under the **Alerts Delivery Methods** heading, select one of the severity levels.
3. Click **Edit**.

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

Figure 5 Alert Notification Delivery Methods Dialog Box

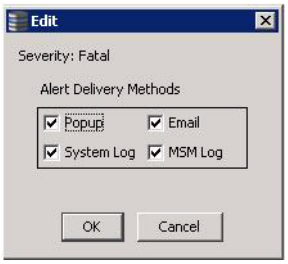

4. Select the desired alert delivery methods for alert notifications at the event severity level.
5. Click **OK** to set the delivery methods used for the severity level that you selected.

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 **Alerts Setting** tab.
   The the **Alerts Setting** portion of the screen appears, as shown in the following figure.
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 6 Change Individual Events Dialog Box

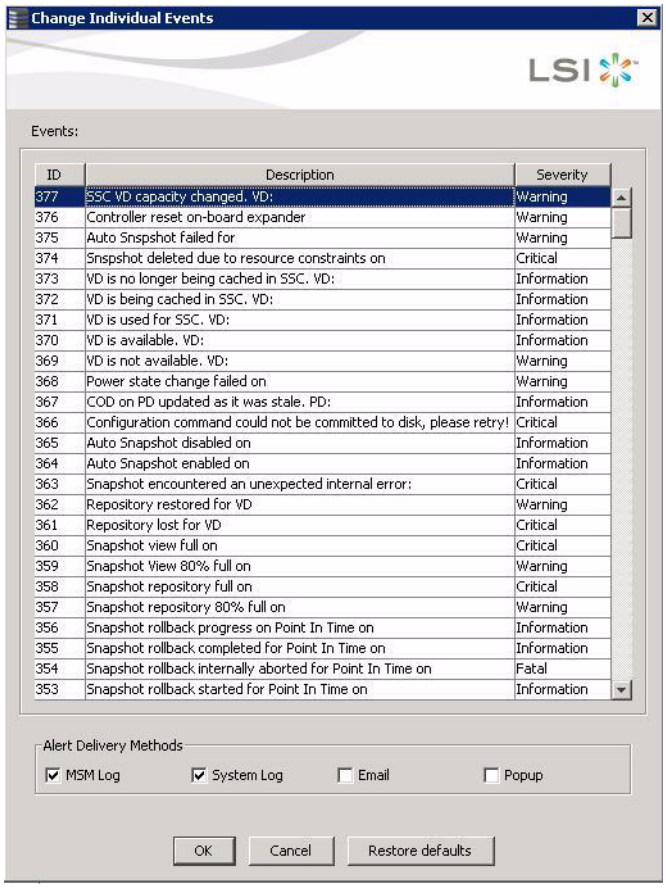

3. Click an event in the list to select it.
   The current alert delivery methods appear for the selected event under the **Alert Delivery Methods** heading.
4. Select the desired alert delivery methods for the event.
5. Press ESC to return to the **Alerts Notification Configuration** screen.
6. Click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

8.3 Changing the Severity Level for Individual Events

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

1. On the Alerts Notification Configuration screen, click the **Alerts Setting** tab.
   The **Alerts Setting** portion of the screen appears.
2. Click **Change Individual Events**.
   The **Change Individual Events** dialog box appears, as shown in the following figure. 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 7 Change Individual Events Severity Level Menu

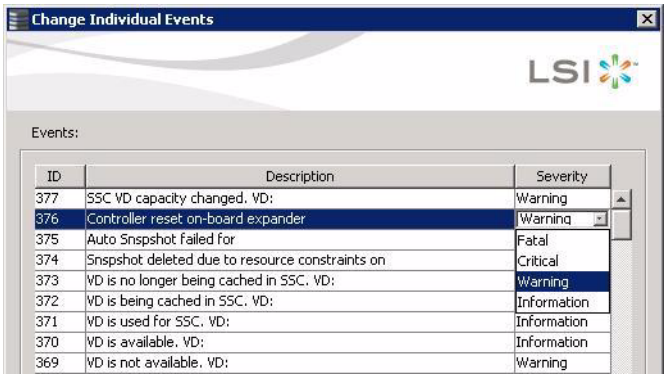

5. Select a different severity level for the event from the menu.
6. Press ESC to return to the **Alerts Notification Configuration** screen.
7. Click **OK** to save all of the changes made to the events.

8.3.1 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 8 Pop-up for Multiple Events

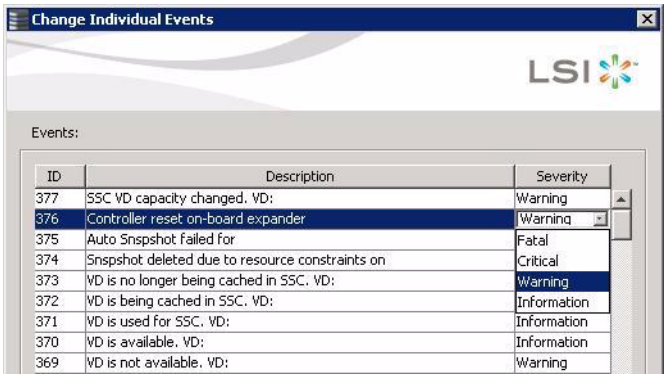

8.3.2 Entering or Editing the Sender Email Address and SMTP Server

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 9 Mail Server Options

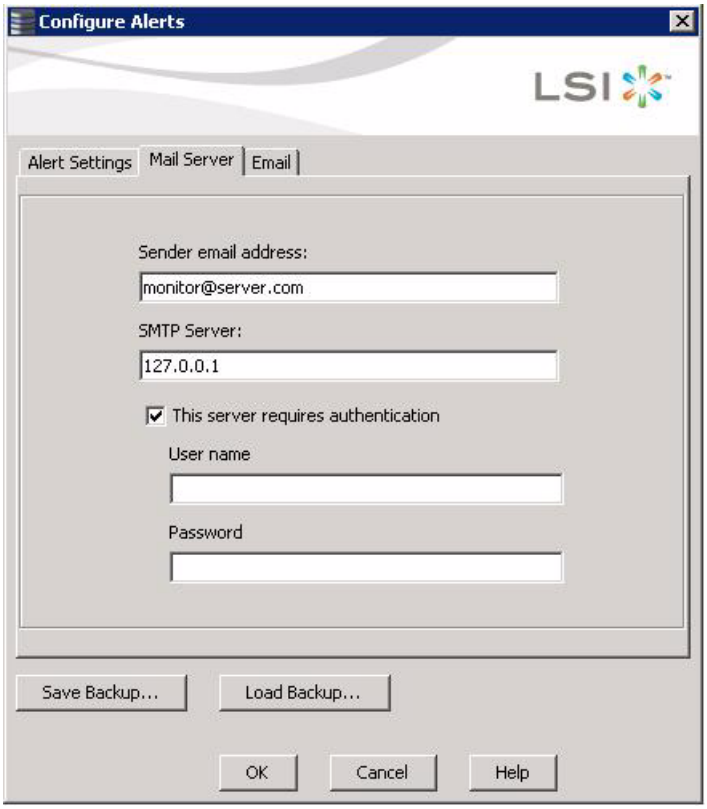

2. Enter a new sender email address in the **Sender email address** field or edit the existing sender email address.
3. Click **OK**.

8.3.3 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 the following figure. 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**.

8.3.4 Saving Backup Configurations

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

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

8.3.5 Loading Backup Configurations

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

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

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

8.3.6 Adding Email Addresses of Recipients of Alert Notifications

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

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

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

Figure 10 Email Settings

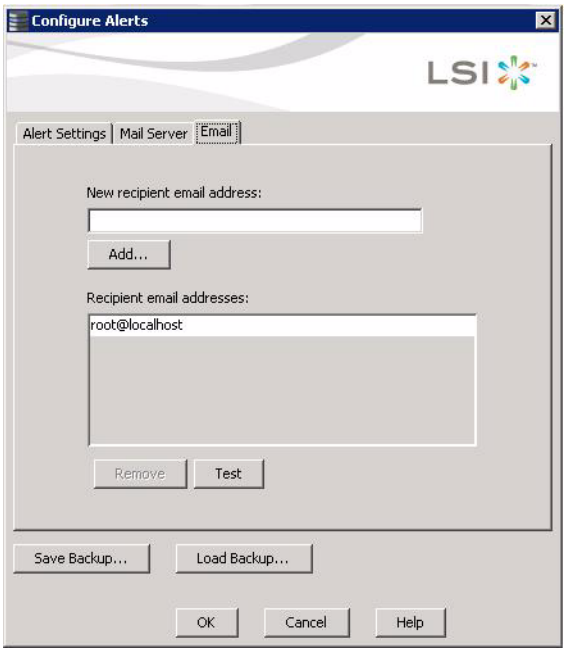

2. Enter the email address you want to add in the **New recipient email address** field.
3. Click **Add**.
   The new email address appears in the **Recipient email addresses** field.

8.3.7 Testing Email Addresses of Recipients of Alert Notifications

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

1. Click the **E-mail** tab on the Event Notification Configuration screen.
   The **E-mail** section of the screen appears, as shown in 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.

8.3.8 Removing Email Addresses of Recipients of Alert Notifications

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

1. Click the **E-mail** tab on the Event Notification Configuration screen.
   The **E-mail** section of the screen appears, as shown in the following figure.
2. Click an email address in the **Recipient email addresses** field.
   The **Remove** button, which was grayed out, is now active.
3. Click **Remove**.
   The email address is deleted from the list.

8.4 Monitoring Controllers

NOTE The ServeRAID C105 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 screen below 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 View, for a complete list of device icons.

To display complete controller information, click a controller icon in the left panel of the MegaRAID Storage Manager main menu screen, and click the **Properties** tab in the right panel. The following figure shows the Controller Information properties.

Figure 11 Controller Information Properties

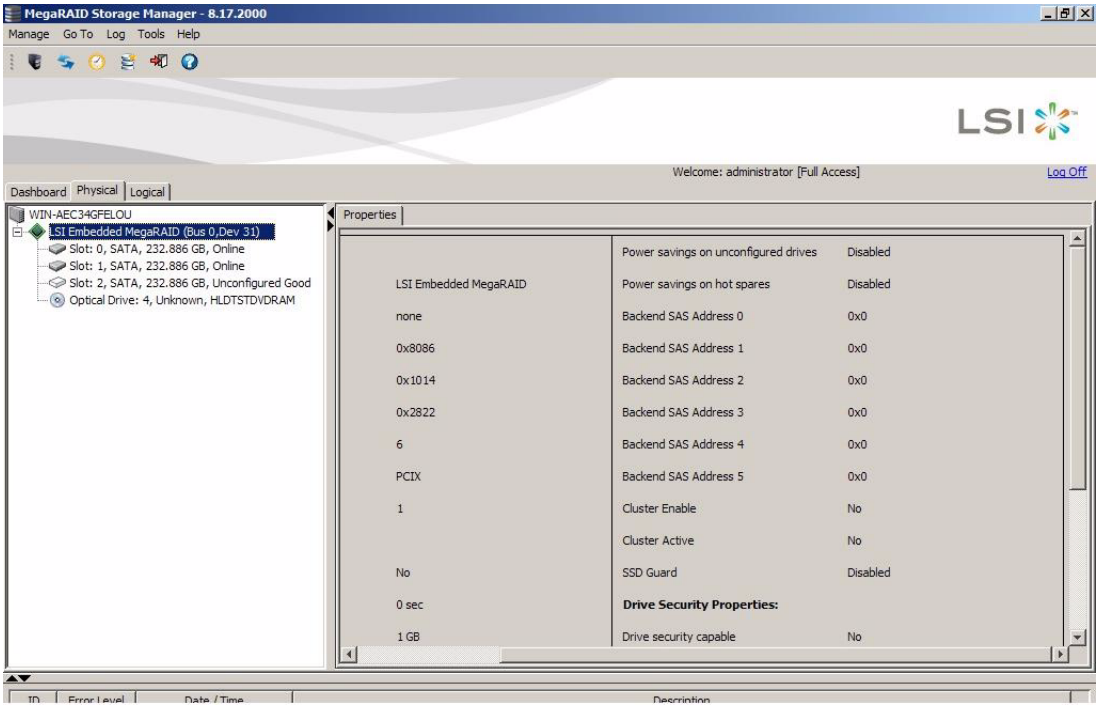

Most of the information on this screen is self-explanatory. Note that the *Rebuild Rate*, *Consistency Check Rate*, and *BGI Rate* (background initialization) options are all user selectable. For more information, see Section 7.2, Changing Adjustable Task Rates.

8.5 Monitoring Disk 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 disk drive is operating normally, its icon looks like this: . If the disk 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 View, for a complete list of device icons.

To display complete disk drive information, click a disk drive icon in the left panel of the MegaRAID Storage Manager main menu screen, and click the **Properties** tab in the right panel.

The following figure shows the Properties panel for a physical drive.

Figure 12 Physical Drive Information

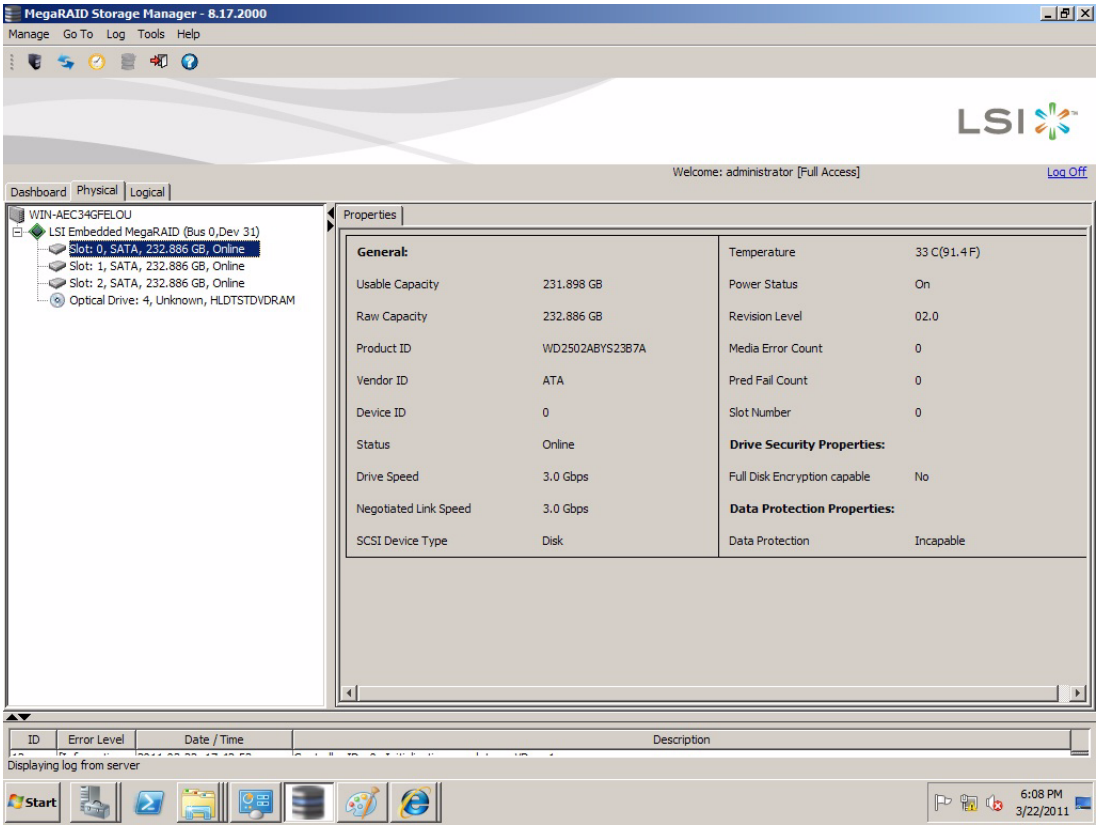

The information on this panel is self-explanatory. There are no user-selectable properties for physical devices. Icons for other storage devices such as CD-ROM drives and DAT drives may also appear in the left panel.

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 hotspares 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 ServeRAID C105 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, and reserved space is red. When you select a virtual drive from the drop-down menu, the drive space used by that virtual drive is displayed in green.

8.6 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 hotspare 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 ServeRAID C105 main menu.
2. Select **Go To > Controller > Set Patrol Read Properties**.

Figure 13 Start Patrol Read Menu

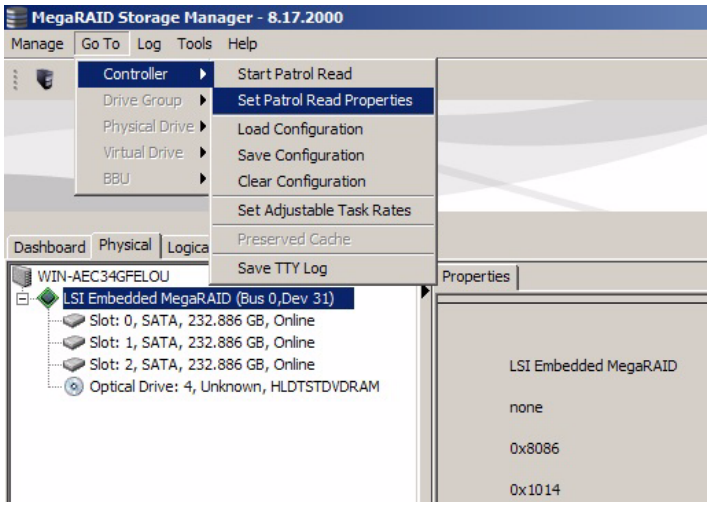

The Patrol Read - Set properties screen displays, as shown in the following figure.

Figure 14 Patrol Read Configuration

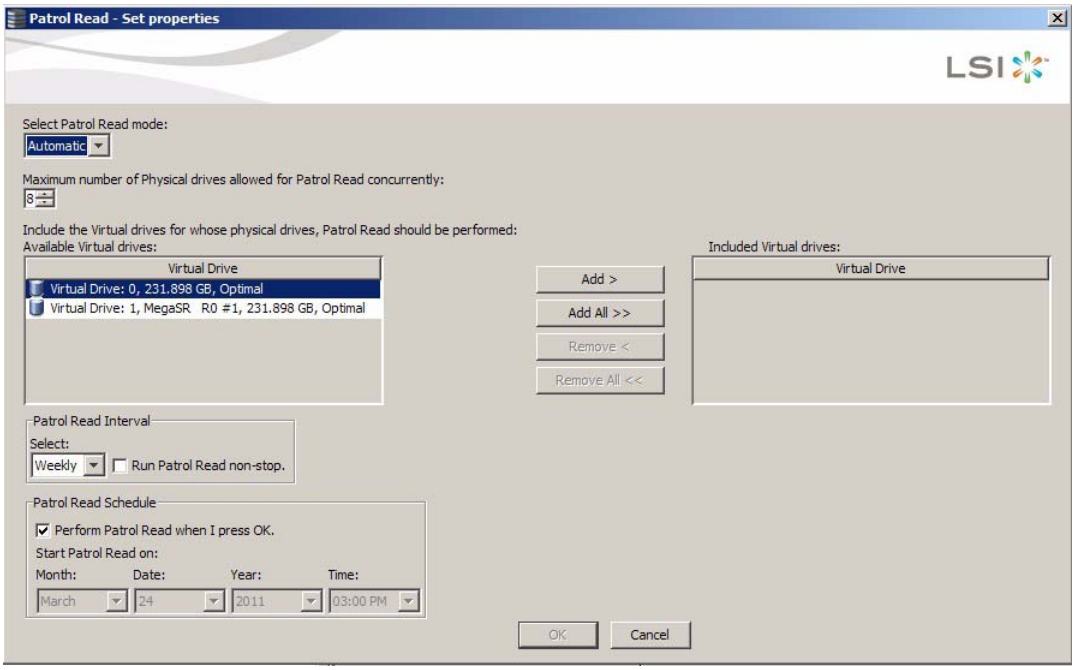

3. Select a mode for a patrol read from these options:
   - **Automatic**: Patrol read runs automatically at the time interval you specify on this screen.
   - **Manual**: Patrol read runs only when you manually start it by selecting **Start Patrol Read** from the controller Options panel.
   - **Disabled**: Patrol read does not run.
4. Specify a maximum count of drives to include in the patrol read.

   The count must be from 1 to 8 drives. Eight drives is the maximum supported by ServeRAID C105.
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 It is recommended 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 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 ServeRAID C105 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.

8.6.1 Patrol Read Task Rates

You have the option to change the patrol read *task rate*. The task rate determines the amount of system resources that are dedicated to a patrol read when it is running. It is recommended, 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.

8.7 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: .

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

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

Figure 15 Virtual Drive Properties

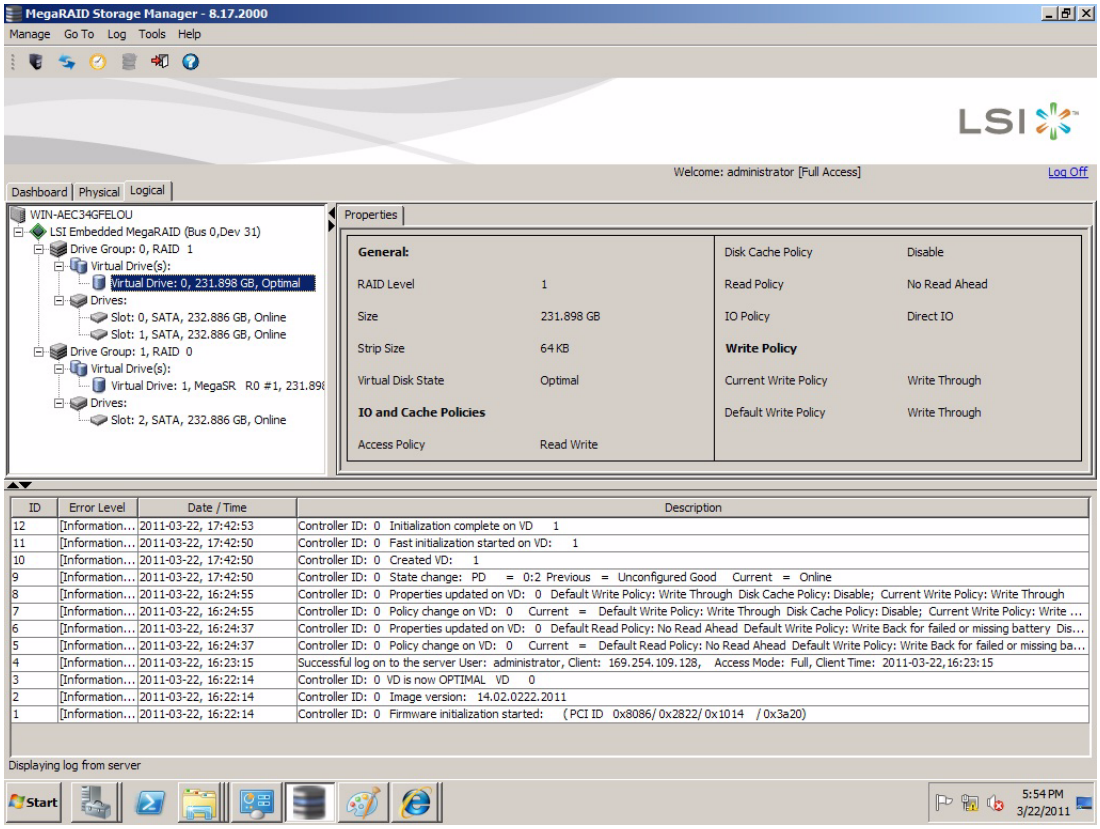

The RAID level, stripe size, and access policy of the virtual disk are set when the virtual drive is configured.

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 hotspare drives).
3. To stop the LEDs from blinking, click **Go To > Virtual Drive > Stop Locating Virtual Drive**.

8.8 Monitoring Rebuilds and Other Processes

MegaRAID Storage Manager allows you to monitor the progress of rebuilds and other lengthy processes 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 to open this window, as shown in the following figure.

Figure 16 Group Show Progress Menu

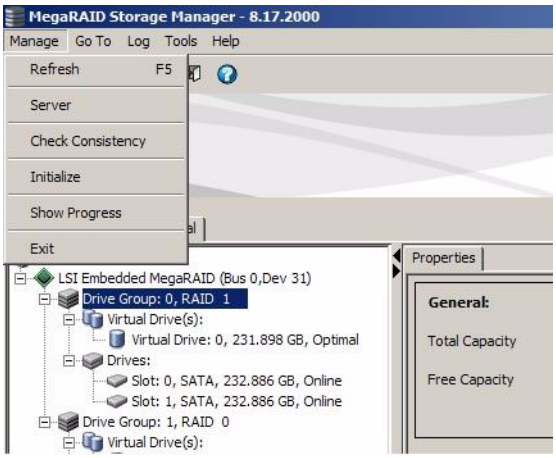

The Group Show Progress window appears, as shown in the following figure.

Figure 17 Group Show Progress Window

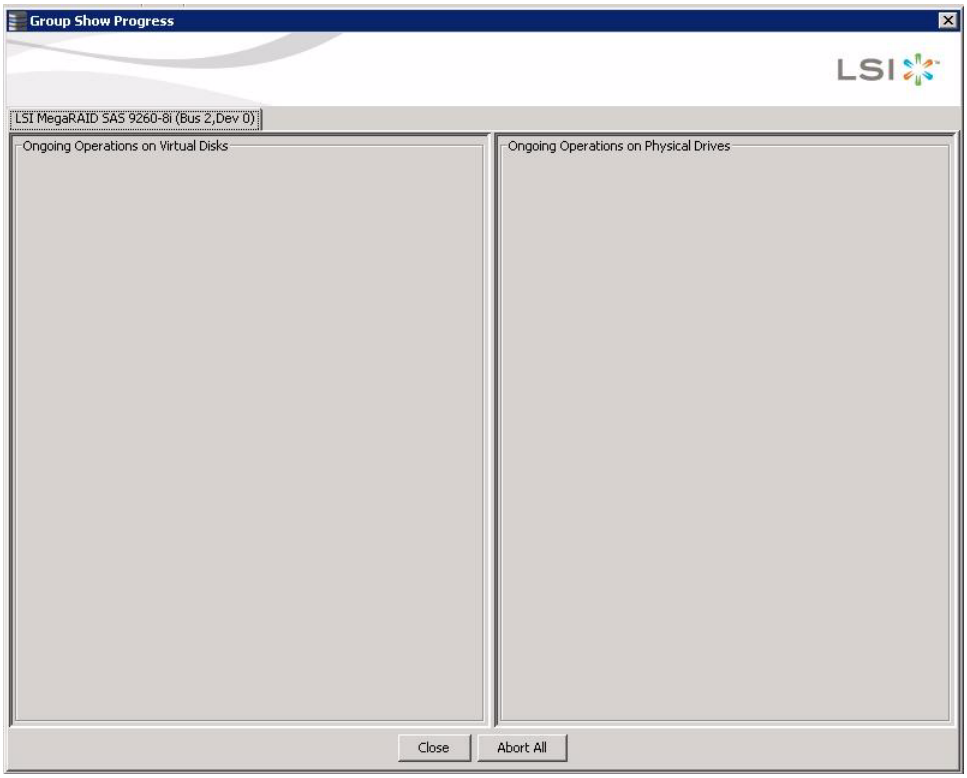

Operations on virtual disks appear in the left panel of the Group Show Progress window, and operations on physical drives appear in the right panel. The following operations appear in this window:

- Background or foreground initialization of a virtual disk
- Consistency check (see Section 9.2, Running a Consistency Check)
- Modify drive group
- Rebuild (see Section 9.3, Rebuilding a Drive)

2. (Optional) To abort an ongoing process, click the **Abort** button next to the status indicator.
3. (Optional) Click **Abort All** to abort all ongoing processes.
4. Click **Close** to close the window.

Chapter 9: Maintaining and Managing Storage Configurations

This section explains how to use MegaRAID Storage Manager 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 disk(s) to initialize.

CAUTION Initialization erases all data on the virtual drive. Make sure you 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 disk. (For more information, see Section 7.1.1, Selecting Virtual Drive Settings.)
5. Click **Start** to begin the initialization.

   You can monitor the progress of the initialization. See Section 8.8, Monitoring Rebuilds and Other Processes, for more information.

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 1 Group Initialization Dialog Box

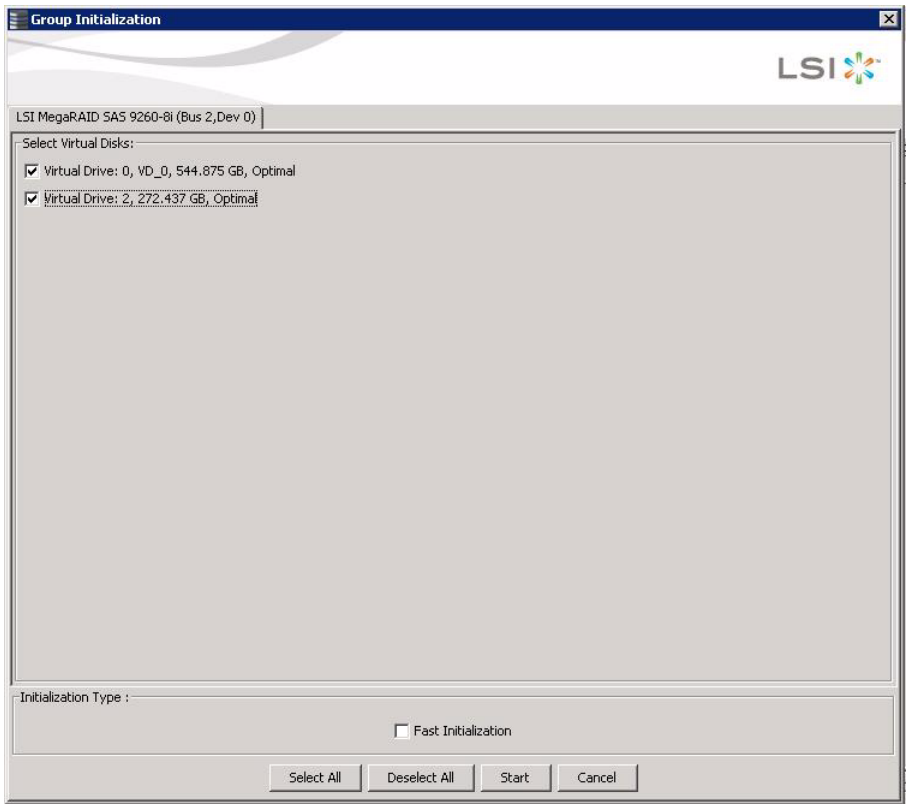

2. Check the virtual drives to run the initialization on or click **Select All** to select all of the virtual drives.
3. Click **Start**.

   You can monitor the progress of the group initialization. See Section 8.8, Monitoring Rebuilds and Other Processes, 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 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 2 Group Consistency Check Dialog Box

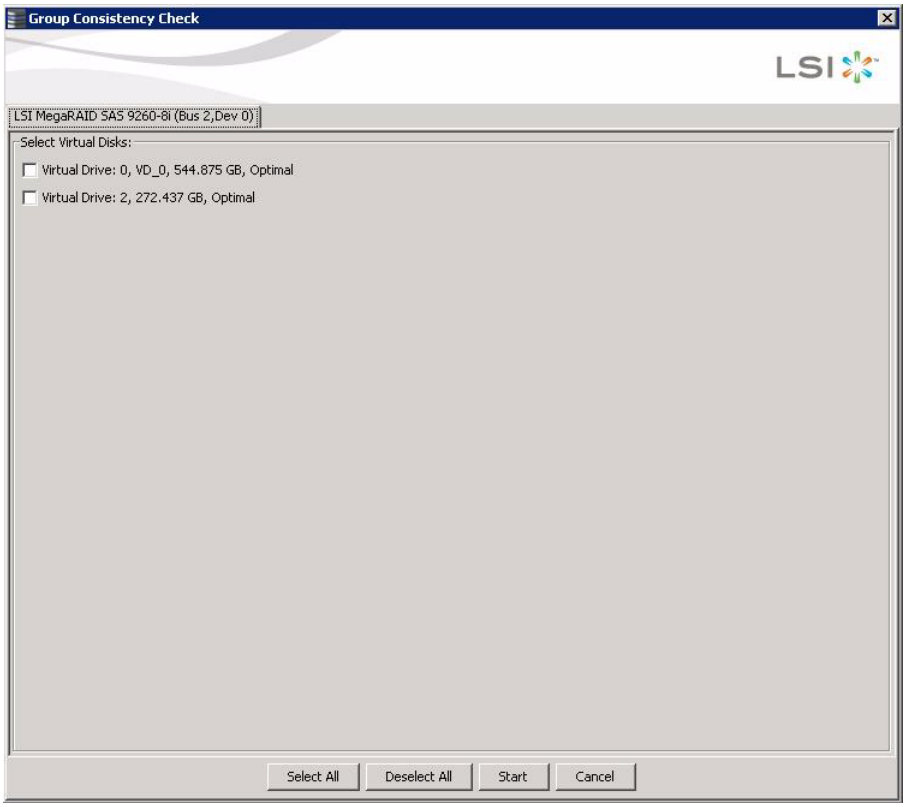

2. Check the virtual drives to run the consistency check on or click **Select All** to select all of the virtual drives.
3. Click **Start**.

You can monitor the progress of the group consistency check. See Section 8.8, Monitoring Rebuilds and Other Processes, for more information.

9.2.2 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 drives that you install in a computer system. In addition, if one or more drives are removed from a configuration, by a cable pull or drive removal, for example, the configuration on those drives is considered a foreign configuration by the RAID controller. Drives that are foreign are listed on the physical drives list with a special symbol in ServeRAID C100 and ServeRAID C105.

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.

ServeRAID C100 and ServeRAID C105 normally detect newly installed drives and display icons for them in the main menu screen. If for some reason, ServeRAID C100 and ServeRAID C105 do 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 main menu screen.
2. Select **Go To > Controller > Scan for Foreign Configuration**.

   If ServeRAID C105 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.

9.3 Rebuilding a Drive

If a single drive in a RAID 1 or RAID 10 virtual disk fails, the system is protected from data loss. The failed drive must be replaced, and the drive's data 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 hotspare disks 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 disk drive icon: . A small yellow circle appears to the right of the icon of the virtual drive that uses this drive: . This icon 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 physical 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 disk 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**.

9.4 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 ServeRAID C100 and ServeRAID C105 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. On the main menu screen, 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 hotspare drive for the virtual drive from which you have removed the drive. (See Section 7.2, Changing Adjustable Task Rates.)
   When a hotspare is available, the data on the virtual drive will be rebuilt. You can now use the removed drive for another configuration.

   CAUTION If ServeRAID C100 and ServeRAID C105 detect that a drive in a virtual drive has failed, they make the drive offline. If this happens, you must remove the drive and replace it. You cannot make the drive usable for another configuration by using the **Mark physical disk as missing** and **Rescan** commands.

Appendix A: ServeRAID C100 and ServeRAID C105 MegaRAID Storage Manager Events and Messages

This appendix lists the MegaRAID Storage Manager events and their related messages that can appear in the event log.

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.

NOTE MegaRAID Storage Manager can be used to manage a wide range of ServeRAID controllers. Some of the events and messages listed in this appendix are not applicable to ServeRAID C100 and ServeRAID C105.

Each message that appears in the event log has an error level that indicates the severity of the event, as shown in the following table .

Table 1 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.

The following table lists all of the MegaRAID Storage Manager event messages. The event message descriptions include placeholders for specific values that are determined when the event is generated. Some of the error messages are relevant only for hardware RAID.

Table 2 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

(Sheet 1 of 13)

Table 2 Event Messages

Number (Hex)	Number (Decimal)	Type	Event Text
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
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)

(Sheet 2 of 13)

Table 2 Event Messages

Number (Hex)	Number (Decimal)	Type	Event Text
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
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

(Sheet 3 of 13)

Table 2 Event Messages

Number (Hex)	Number (Decimal)	Type	Event Text
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
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

(Sheet 4 of 13)

Table 2 Event Messages

Number (Hex)	Number (Decimal)	Type	Event Text
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	Global Hot Spare created on %s (%s)
0x0088	136	Information	Global Hot Spare %s (%s) disabled
0x0089	137	Caution	Global Hot 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

(Sheet 5 of 13)

Table 2 Event Messages

Number (Hex)	Number (Decimal)	Type	Event Text
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
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

(Sheet 6 of 13)

Table 2 Event Messages

Number (Hex)	Number (Decimal)	Type	Event Text
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
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

(Sheet 7 of 13)

Table 2 Event Messages

Number (Hex)	Number (Decimal)	Type	Event Text
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
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

(Sheet 8 of 13)

Table 2 Event Messages

Number (Hex)	Number (Decimal)	Type	Event Text
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
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

(Sheet 9 of 13)

Table 2 Event Messages

Number (Hex)	Number (Decimal)	Type	Event Text
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 VDs: %s
0x0145	325	Information	Controller cache discarded by user for VDs: %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
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

(Sheet 10 of 13)

Table 2 Event Messages

Number (Hex)	Number (Decimal)	Type	Event Text
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	Snaphot 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
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

(Sheet 11 of 13)

Table 2 Event Messages

Number (Hex)	Number (Decimal)	Type	Event Text
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	397	Warning	Cache-vault microcode update required
0x018e	398	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
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

(Sheet 12 of 13)

Table 2 Event Messages

Number (Hex)	Number (Decimal)	Type	Event Text
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

(Sheet 13 of 13)

Appendix B: MDRAID Migration

This appendix describes the process to migrate from ServeRAID to MDRAID.

NOTE The following instructions are for ServeRAID C105 MDRAID migration and are not applicable for ServeRAID C100.

Perform the following steps to run MDRAID and check if the DDF created using ServeRAID is recognized on the Romley platform:

1. Set AHCI in AHCI/Compatibility/Enhanced mode (non LSI mode).
2. Set SCU in LSI's RAID mode.
3. Install the OS on AHCI controller.
4. Create RAID 0/RAID 1 volumes using pre-boot applications.
5. Boot into the OS using the `nodmraid` option, the native iSCI driver would claim the SCU controller and discover the drives.
6. Use MDADM to import the VOLUMES created using ServeRAID's pre-boot applications by performing the following steps:
   a. `fdisk -l` (show all connected disk)
   b. `mdadm --examine /dev/devicename` (examine all connected hard disks)
   c. `mdadm --examine --scan --config=mdadm.conf > /etc/mdadm.conf` (populate the mdadm.conf file)
   d. `mdadm --assemble –scan` (scan all the PD's and create the mdraid volume)
   e. `mdadm /proc/mdstat` or `mdadm --detail /dev/md126`
7. Create file system over the imported MDRAID volume and perform the I/O operations.

Figure 1 Setup After Migration

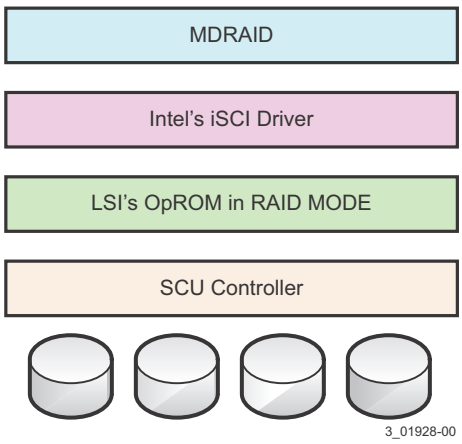

NOTE The following instructions are for ServeRAID C100 MDRAID migration and are not applicable for ServeRAID C105.

Perform the following steps to run MDRAID and check if the DDF created using ServeRAID is recognized on the Denlow Platform:

1. Install the OS on AHCI controller in AHCI mode.
2. Halt the system and remove the OS PD.
3. Set AHCI in LSI RAID mode.
4. Create RAID 0/RAID 1 volumes using pre-boot applications.
5. Reboot and set the AHCI in AHCI mode.
6. Connect the OS PD into any one of the available AHCI port.
7. Boot into the OS with `nodmraid` option, the native AHCI driver would claim the AHCI controller and discover the drives.
8. Use MDADM to import the VOLUMES created using LSI's SWR pre-boot applications. Use the following steps –
9. fdisk -l (this will shows all the connected disk)
   a. `mdadm --examine /dev/devicename` (examine all connected hard disks)
   b. `mdadm --examine --scan --config=mdadm.conf > /etc/mdadm.conf` (populate the mdadm.conf file)
   c. `mdadm --assemble –scan` (scan all the PD's and create the mdraid volume)
   d. `mdadm /proc/mdstat` or `mdadm --detail /dev/md126`
10. Create file system over the imported MDRAID volume and perform the I/O operations.

Appendix C: MegaCLI Commands to StorCLI Command Conversion

C.1 System Commands

Table 1 System Commands

Description	MegaCLI Command	StorCLI Command
Show the software version.	`MegaCLI -v`	`storcli -v`
Show help information.	`MegaCLI -help|-h|?`	`storcli -help|-h|?`
Show the number of controllers connected.	`MegaCLI -adpCount`	`storcli show ctrlcount`

C.2 Controller Commands

Table 2 Controller Commands

Description	MegaCLI Command	StorCLI Command
Show the status of properties related to the controllers.	`MegaCli -AdpGetProp <`*`PropertyName`*`>-aN|-a0,1,2|-aALL`	`/cx show <`*`propertyName`*`>`
	The following properties can be used with this command:	The following properties can be used with this command:
	`abortcconerror`	`abortcconerror`
	`alarmdsply`	`alarm`
	`autodetectbackplanedsbl`	`backplane`
	`autoenhancedimportdsply`	`foreignautoimport`
	`autosnapshotspace`	
	`batwarndsbl`	`batterywarning`
	`bgirate`	`bgirate`
	`bootwithpinnedcache`	`bootwithpinnedcache`
	`cachebypass`	`cachebypass`
	`ccrate`	`ccrate`
	`clusterenable`	
	`coercionmode`	`coercion`
	`copybackdsbl`	`copyback`
	`defaultldpspolicy`	`ds`
	`defaultsnapshotspace`	
	`defaultviewspace`	
	`disableldpsinterval`	`ds`
	`disableldpstime`	`ds`
	`disableocr`	`ocr`
	`eccbucketcount`	`eccbucketsize`
	`eccbucketleakrate`	`eccbucketleakrate`
	`enableeghsp`	`eghs`
	`enableesmarter`	`eghs`

Table 2 Controller Commands (Continued)

Description	MegaCLI Command	StorCLI Command
	enableeug	eghs
	enablejbod	Jbod
	enblspindownunconfigdrvs	ds
	loadbalancemode	loadbalancemode
	maintainpdfailhistoryenbl	maintainpdfailhistory
	ncqdsply	ncq
	patrolreadrate	prrate
	perfmode	perfmode
	predfailpollinterval	smartpollinterval
	rebuildrate	rebuildrate
	reconrate	migraterate
	rstrhotspareoninsert	restorehotspare
	smartcpybkenbl	copyback
	spindowntime	ds
	spinupencdelay	ds
	spinupdelay	spinupdelay
	spinupencdrvcnt	spinupdrivecount
	ssdsmartcpybkenbl	copyback
	usediskactivityforlocate	activityforlocate
	usefdeonlyencrypt	usefdeonlyencrypt
Set properties on the selected controllers.	Megacli -AdpSetProp <*propertyname*>-an|-a0,1,2|-aall	/*cx* set <*property1*>
	The following properties can be set using this command:	The following properties can be set using this command:
	abortcconerror	abortcconerror=<on|off>
	alarmdsply	alarm=<on|off| silence>
	autodetectbackplanedsbl	backplane=<*value*>
	autoenhancedimportdsply	foreignautoimport=<on|off>
	batwarndsbl	batterywarning=<on|off>
	bgirate	bgirate=<*value*>
	bootwithpinnedcache	bootwithpinnedcache=<on|off>
	cachebypass	cachebypass=<on|off>
	ccrate	ccrate=<*value*>
	clusterenable	
	coercionmode	coercion=<*value*>
	copybackdsbl	copyback=<on|off> type=<smartssd|smarthdd|all>
	defaultldpspolicy	ds=<*value*>
	defaultsnapshotspace	
	defaultviewspace	
	disableldpsinterval	ds=<*value*>

Table 2 Controller Commands (Continued)

Description	MegaCLI Command	StorCLI Command
	disableldpstime	ds=<*value*>
	disableocr	ocr=<*value*>
	eccbucketcount	eccbucketsize=<*value*>
	eccbucketleakrate	eccbucketleakrate=<*value*>
	enableeghsp	eghs [state=<on|off>]
	enableesmarter	eghs [smarter=<on|off>]
	enableeug	eghs [eug=<on|off>
	enablejbod	jbod=<on|off>
	enblspindownunconfigdrvs	ds=<*value*>
	loadbalancemode	loadbalancemode=<*value*>
	maintainpdfailhistoryenbl	maintainpdfailhistory=<on|off>
	ncqdsply	ncq=<on|off>
	patrolreadrate	prrate=<*value*>
	perfmode	perfmode=<*value*>
	predfailpollinterval	smartpollinterval=<*value*>
	rebuildrate	rebuildrate=<*value*>
	reconrate	migraterate=<*value*>
	rstrhotspareoninsert	restorehotspare=<on|off>
	smartcpybkenbl	copyback=<on|off> type=<smartssd|smarthdd|all>
	spindowntime	ds=<on|off>
	spinupdelay	spinupdelay=<*value*>
	spinupdrivecount	spinupdrivecount=<*value*>
	spinupencdelay	ds
	spinupencdrvcnt	ds
	sdsmartcpybkenbl	copyback=<on|off> type=<smartssd|smarthdd|all>
	usediskactivityforlocate	activityforlocate=<on|off>
	usefdeonlyencrypt	usefdeonlyencrypt=<on|off>
Show the number of controllers connected.	MegaCLI -adpCount	storcli show ctrlcount
Show all information about the adapter, such as cluster state, BIOS, alarm, firmware, version, and so on.	MegaCli -AdpAllInfo -aN|-a0,1,2|-aALL	storcli /c*x* show all
Show the freespace available in the controller.	MegaCLI -CfgFreeSpaceinfo -aN|-a0,1,2|-aALL	storcli /c*x* show freespace
Download the controller firmware.	MegaCli -AdpFwFlash -f *filename* [-NoSigChk] [-NoVerChk] [-ResetNow] -aN|-a0,1,2|-aALL	storcli /c*x* download file=<*filepath*> [fwtype=<*val*>] [nosigchk] [noverchk] [resetnow]

Table 2 Controller Commands (Continued)

Description	MegaCLI Command	StorCLI Command
Show the preserved cache status.	MegaCLI-GetPreservedCacheList -aN|-a0,1,2|-aALL	storcli /cx show preservedcache
Set the controller time	MegaCLI –AdpSetTime *yyyymmdd hh:mm:ss* -aN|-a0,1,2|-aALL	storcli /c(x|all) set time=<*yyyymmdd hh:mm:ss* | systemtime>
Show the controller time.	MegaCLI –AdpGetTime -aN	storcli /cx show time

C.3 Patrol Read Commands

Table 3 Patrol Read Commands

Description	MegaCLI Command	StorCLI Command
Show the patrol read status and patrol read parameters, if any in progress.	MegaCli -AdpPR -info -aN|-a0,1,2|-aALL	storcli/cx show patrolRead
Set the patrol read options on a single adapter, multiple adapters, or all adapters. (x = single controller).	MegaCli -AdpPR –Dsbl|EnblAuto|EnblMan|Start|Stop | Info|Suspend|Resume|Stop| SSDPatrolReadEnbl | SSDPatrolReadDsbl |{SetDelay *Val*}|{-SetStartTime *yyyymmdd hh*}|{maxConcurrentPD Val} -aN|-a0,1,2|-aALL	storcli /cx set patrolread {=on mode=<auto|manual>}|{off} storcli /cx set patrolread [starttime=<*yyyy/mm/dd hh*>] [maxconcurrentpd=<*value*>] [includessds=<on|off>] [uncfgareas=on|off] storcli /cx set patrolread delay=<*value*>
Disable patrol read.	MegaCli -AdpPR -Dsbl -aN|-a0,1,2|-aALL	storcli /cx set patrolread=off
Enable automatic patrol read.	MegaCli -AdpPR -EnblAuto -aN|-a0,1,2|-aALL	storcli /cx set patrolread=on mode=auto
Enable manual patrol read.	MegaCli -AdpPR -EnblMan -aN|-a0,1,2|-aALL	storcli /cx set patrolread=on mode=manual
Start patrol read.	MegaCli -AdpPR -Start -aN|-a0,1,2|-aALL	storcli /cx start patrolRead
Suspend a running patrol read.	MegaCli -AdpPR -Suspend -aN|-a0,1,2|-aALL	storcli /cx suspend patrolRead
Resume a suspended patrol read.	MegaCli -AdpPR -Resume -aN|-a0,1,2|-aALL	storcli /cx resume patrolRead
Stop a running patrol read.	MegaCli -AdpPR -Stop -aN|-a0,1,2|-aALL	storcli /cx stop patrolRead
Include SSD drives in patrol read.	MegaCli -AdpPR -SSDPatrolReadEnbl -aN|-a0,1,2|-aALL	storcli /cx set patrolRead includessds=on | onlymixed
Exclude SSD drives in patrol read.	MegaCli -AdpPR -SSDPatrolReadDsbl -aN|-a0,1,2|-aALL	storcli /cx set patrolRead includessds=off
Delay a patrol read,	MegaCli -AdpPR -SetDelay Val -aN|-a0,1,2|-aALL	storcli /cx set patrolread delay=<*value*>
Schedule a patrol read.	MegaCli -AdpPR -SetStartTime *yyyymmdd hh* -aN|-a0,1,2|-aALL	storcli /cx set patrolread=on starttime=*YYYY/MM/DD HH*
Set the value for maximum concurrent physical drives for the patrol read.	MegaCli -AdpPR -maxConcurrentPD Val -aN|-a0,1,2|-aALL	storcli /cx set patrolread maxconcurrentpd=*xx*

C.4 Consistency Check Commands

Table 4 Consistency Check Commands

Description	MegaCLI Command	StorCLI Command
Schedule a consistency check.	MegaCLI -AdpCcSched -Dsbl|-Info| {-ModeConc | -ModeSeq [-ExcludeLD -LN|-L0,1,2] [-SetStartTime *yyyymmdd hh*] [-SetDelay *val*] } -aN|-a0,1,2|-aALL	storcli /c*x* set consistencycheck|cc=[off|seq|conc] [delay=value] starttime=*yyyy/mm/dd hh* [excludevd=x-y,z]
Show consistency check status and consistency parameters, in progress, if any.	MegaCLI -AdpCcSched -Info	storcli /c*x* show cc/ConsistencyCheck

C.5 OPROM BIOS Commands

Table 5 OPROM BIOS Commands

Description	MegaCLI Command	StorCLI Command
Schedule a consistency check.	MegaCli -AdpBIOS -Dsply -aN|-a0,1,2|-aALL	storcli /c*x* show bios
Show consistency check status and consistency parameters, if any in progress.	MegaCli -AdpBootDrive -{-Set {-Lx | -physdrv[E0:S0]}} -aN|-a0,1,2|-aALL	storcli /c*x*/e*x*/s*x* set bootdrive=on|off storcli /c*x*/v*x* set bootdrive=on|off
Sets the BIOS properties for the controller.	MegaCli -AdpBIOS -Enbl | -Dsbl | -Dsply | SOE | BE EnblAutoSelectBootLd | DsblAutoSelectBootLd -aN|-a0,1,2|-aALL	storcli /c*x* set bios=<on|off> storcli /c*x* set stoponerror|soe=<on|off> storcli /c*x* set autobootselect(abs)=<on|off>

C.6 Battery Commands

Table 6 Battery Commands

Description	MegaCLI Command	StorCLI Command
Show battery-related information.	MegaCli -AdpBbuCmd -aN|-a0,1,2|-aALL	storcli /c*x*/bbu show storcli /c*x*/bbu show all
Show the battery learn properties.	MegaCli -AdpBbuCmd -GetBbuProperties -aN|-a0,1,2|-aALL	storcli /c*x*/bbu show properties
Show the battery information, firmware status, and the gas gauge status.	MegaCli -AdpBbuCmd -GetBbuStatus -aN|-a0,1,2|-aALL	storcli /c*x*/bbu show status
Show battery capacity information.	MegaCli -AdpBbuCmd -GetBbuCapacityInfo -aN|-a0,1,2|-aALL	storcli /c*x*/bbu show all
Show battery design information.	MegaCli -AdpBbuCmd -GetBbuDesignInfo -aN|-a0,1,2|-aALL	storcli /c*x*/bbu show all

Table 6 Battery Commands (Continued)

Description	MegaCLI Command	StorCLI Command
Set battery properties	MegaCli -AdpBbuCmd -SetBbuProperties -f <*fileName*> -aN|-a0,1,2|-aALL	storcli /c*x*/bbu set learnDelayInterval=<*value*> storcli /c*x*/bbu set bbuMode=<*value*> storcli /c*x*/bbu set autolearnmode=<*value*> where x= 0 – Enabled, 1 – Disabled, 2 – Warn though event.
Start battery learn cycle.	MegaCli -AdpBbuCmd -BbuLearn -aN|-a0,1,2|-aALL	storcli /c*x*/bbu start learn
Set the battery to low power storage mode.	MegaCli -AdpBbuCmd -BbuMfgSleep -aN|-a0,1,2|-aALL	storcli /c*x*/bbu set powermode=sleep
Seal the gas gauge EEPROM write access	MegaCli -AdpBbuCmd -BbuMfgSeal -aN|-a0,1,2|-aALL	storcli /c*x*/bbu set writeaccess=sealed

C.7 RAID Configuration Commands

Table 7 RAID Configuration Commands

Description	MegaCLI Command	StorCLI Command
Create a RAID configuration of RAID type 0, 1, 5, and 6.	MegaCli –CfgLDAdd -R0|-R1|-R5|-R6[E0:S0,E1:S1,...] [WT | WB] [NORA | RA | ADRA] [Direct | Cached] [CachedBadBBU|NoCachedBadBBU] [-sz*XXXXXXXX* [-sz*YYYYYYYY* [...]]] [-strpszM] [–Hsp[E5:S5,...]] [–afterLdX] -aN	storcli /c*x* add vd type=raid[0|1|5|6] [Size=<*VD1_Sz*>,< *VD2_Sz*>,..|*all] [name=<*VDNAME1*>,..] drives=*e*:*s*|*e*:*s*-*x*|*e*:*s*-*x*,*y*;*e*:*s*-*x*,*y*,*z* [PDperArray=*x*] [SED] [pdcache=on|off|*default][pi] [DimmerSwitch(ds)=default|automatic(auto)|*none|maximum(max) |MaximumWithoutCaching(maxnocache)] [wt|*wb|awb] [nora|*ra] [*direct|cached] [strip=<8|16|32|64|128|256|512|1024] [AfterVd=*X*] [Spares=[*e*:]*s*|[*e*:]*s*-*x*|[*e*:]*s*-*x*,*y*] [force]
Create a CacheCade virtual drive.	MegaCLI -CfgCacheCadeAdd [-rX] -Physdrv[E0:S0,...] {-Name LdNamestring} [WT|WB|ForcedWB] [-assign -LX|L0,2,5..|LALL] -aN|-a0,1,2|-Aall	storcli /c*x* add vd cachecade|cc Type=raid[0,1,10] drives=[*e*:]*s*|[*e*:]*s*-*x*|[*e*:]*s*-*x*,*y* [< WT| WB>] [assignvds=0,1,2]*e*:]
Create a RAID configuration of RAID type 10, 50, and 60.	MegaCli –CfgSpanAdd -aN|-a0,1,2|-aALL -R10|-R50|R60 –Array0[E0:S0,E1:S1,...] –Array1[E0:S0,E1:S1,...] [...] [WT | WB] [NORA | RA | ADRA] [Direct | Cached] [CachedBadBBU|NoCachedBadBBU] [-sz*XXXXXXXX*[-sz*YYYYYYYY* [...]]] [-strpszM] [–afterLdX] -aN	storcli /c*x* add vd type=raid[10|50|60] [Size=<*VD1_Sz*>,<*VD2_Sz*>,..|*all] [name=<*VDNAME1*>,..] drives=*e*:*s*|*e*:*s*-*x*|*e*:*s*-*x*,*y*;*e*:*s*-*x*,*y*,*z* [PDperArray=*x*] [SED] [pdcache=on|off|*default][pi] [DimmerSwitch(ds)=default|automatic(auto)|*none|maximum(max) |MaximumWithoutCaching(maxnocache)] [wt|*wb|awb] [nora|*ra] [*direct|cached] [strip=<8|16|32|64|128|256|512|1024] [AfterVd=*X*] [Spares=[*e*:]*s*|[*e*:]*s*-*x*|[*e*:]*s*-*x*,*y*] [force]
Clear the complete configuration.	MegaCli -CfgClr [-Force] -aN|-a0,1,2|-aALL	storcli /c*0*/vall delete [force]

Table 7 RAID Configuration Commands (Continued)

Description	MegaCLI Command	StorCLI Command
Show the topology information of the drive group.	MegaCLI -CfgDsply -aN|-a0,1,2|-Aall	storcli /cx/dall show [all]
Show information for a CacheCade virtual drive.	MegaCLI -CfgCacheCadeDsply -aN|-a0,1,2|-Aall	storcli /cx/dall show CacheCade(cc)
Delete a virtual drive hosting the operating system.	MegaCLI -CfgLdDel -LX|-L0,2,5...|-LALL [-Force] -aN|-a0,1,2|-aALL	storcli /cx/v/vx [all] delete -force
Delete a CacheCade virtual drive.	MegaCLI -CfgCacheCadeDel -LX|-L0,2,5...|-LALL -aN|-a0,1,2|-Aall	storcli /cx/vx [all] delete CacheCade(cc)
Show, delete, and import the foreign configuration commands.	MegaCli –CfgForeign –Scan | {-Preview | –Dsply| -Import | -Clear[FID]} -aN|-a0,1,2|-aALL"	storcli /cx/f(x|all) show [all] [securityKey=*xxx*] storcli /cx/f(x|all) del|delete [securityKey=*xxx*] storcli /cx/f(x|all) import [preview] [securityKey=*xxx*]"

C.8 Security Commands

Table 8 Security Commands

Description	MegaCLI Command	StorCLI Command
Set the key ID for the controller.	MegaCli -CreateSecurityKey -SecurityKey *sssssssssss* | [-Passphrase *sssssssssss*] |[-KeyID *kkkkkkkkkkk*] -aN	storcli /*cx* set SecurityKey=*XXXXXX* [passphrase=*yyyyy*] [keyId=*zzzz*]
Change the security key for the controller.	MegaCli -ChangeSecurityKey -OldSecurityKey *sssssssssss* | -SecurityKey *sssssssssss* | [-Passphrase *sssssssssss*] | [-keyID *kkkkkkkkkkk*] -aN	storcli /*cx* set SecurityKey=*XXXXXX* OldSecurityKey=*yyyyy*
Compare and verify the security key for the controller.	MegaCli -VerifySecurityKey -SecurityKey *sssssssssss* -aN	storcli /*cx* compare SecurityKey=*xxxxxx*
Delete the security key.	MegaCLI -DestroySecurityKey | [-Force] -aN	storcli /*cx* delete SecurityKey
Set the security key for the controller.	MegaCli -SetKeyID -KeyID *kkkkkkkkkkk* -aN	storcli /*cx* set SecurityKey KeyId=*xxxx*

C.9 Virtual Drive Commands

Table 9 Virtual Drive Commands

Description	MegaCLI Command	StorCLI Command
Show the virtual drive information.	MegaCli –LDInfo –Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /cx/v(x|all) show storcli /cx/v(x|all) show all
Set virtual drive properties.	MegaCli –LDSetProp WT | WB|NORA |RA | ADRA|-Cached|Direct|CachedBadBBU| NoCachedBadBBU} | -RW|RO|Blocked | {-Name nameString} |-EnDskCache|DisDskCache –Lx| -L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /cx/v(x|all) set wrcache=WT|WB|AWB storcli /cx/v(x|all) set rdcache=RA|NoRA storcli /cx/v(x|all) set iopolicy=Cached|Direct storcli /cx/v(x|all) set accesspolicy=RW|RO|Blocked|RmvBlkd storcli /cx/v(x|all) set pdcache=On|Off|Default storcli /cx/v(x|all) set name=<*NameString*>
Set power-saving (dimmer switch) properties.	MegaCli -LDSetPowerPolicy -Default| -Automatic| -None| -Maximum| -MaximumWithoutCaching -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /cx/v(x|all) set ds=Default|Auto|None|Max|MaxNoCache
Show virtual drive expansion information.	MegaCli -getLdExpansionInfo -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /cx/v(x|all) show expansion
Expand the virtual drive within the existing array; also use if you replace the drives with larger drives, beyond the size of the existing array.	MegaCli -LdExpansion -pN -dontExpandArray -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /cx/v(x|all) expand Size=<*value*> [expandarray]
Secure the virtual drive.	MegaCLI --LDMakeSecure -Lx|-L0,1,2,...|-Lall –An	storcli /cx/vx set security=on
Show specific properties of virtual drives.	MegaCli –LDGetProp -Cache | -Access | -Name | -DskCache -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /cx/vx show
Start virtual drive initialization.	MegaCli –LDInit –Start [Fast|Full] -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /cx/v(x|all) start init[Full]
Stop a running virtual drive initialization.	MegaCli –LDInit -Abort -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /cx/v(x|all) stop init
Show the initialization progress.	MegaCli –LDInit –ShowProg -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /cx/v(x|all) show init
Start a consistency check on an uninitialized virtual drive.	MegaCli –LDCC –Start –Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /cx/v(x|all) start cc[Force]

Table 9 Virtual Drive Commands (Continued)

Description	MegaCLI Command	StorCLI Command
Start, stop, suspend, resume, and show the progress of a consistency check operation.	MegaCli -LDCC -Start|-Abort| -Suspend|-Resume|-ShowProg| -ProgDsply -Lx|-L0,1,2|-LALL -aN|-a0,1,2|-aALL	storcli /cx/v(x|all) start cc storcli /cx/v(x|all) stop cc storcli /cx/v(x|all) pause cc storcli /cx/v(x|all) resume cc storcli /cx/v(x|all) show cc
Enable/disable automatic background initialization. Show, stop, pause, resume, and show the progress of the background initialization.	MegaCLI -LDBI -Enbl|-Dsbl| -getSetting|-Abort|-Suspend| -Resume|-ShowProg|-ProgDsply -Lx|-L0,1,2|-LALL -aN|-a0,1,2|-Aall	storcli /cx/v(x|all) set autobgi=On|Off storcli /cx/v(x|all) show autobgi storcli /cx/v(x|all) stop bgi storcli /cx/v(x|all) pause bgi storcli /cx/v(x|all) resume bgi storcli /cx/v(x|all) show bgi
Start and show progress for a migrate operation.	MegaCli –LDRecon {–Start –Rx [Add | Rmv PhysDrv[E0:S0,E1:S1,...]] } | –ShowProg|-ProgDsply –Lx –aN	storcli /cx/vx start migrate type=raidx [option=add|remove drives=[e:]s|[e:]s-x|[e:]s-x,y] [Force] storcli /cx/v(x|all) show migrate
Delete preserved cache.	MegaCLI -DiscardPreservedCache -Lx|-L0,1,2|-Lall -force -aN|-a0,1,2|-aALL	storcli /cx/v(x|all) delete preservedcache[force]
Assign the CacheCade virtual drive.	MegaCLI -Cachecade -assign|-remove -Lx|-L0,1,2|-LALL -aN|-a0,1,2|-aALL	storcli /cx/vx|all set ssdCaching=on|off

C.10 Physical Drive Commands

Table 10 Physical Drive Commands

Description	MegaCLI Command	StorCLI Command
Show drive information.	MegaCli -pdInfo -PhysDrv[E0:S0,E1:S1,...] -aN|-a0,1,2|-aALL	storcli /cx/ex/sx show storcli /cx/ex/sx show all
Start, stop, pause, resume, or show the progress of a rebuild operation.	MegaCLI PDRbld -Start|-Stop |-Suspend|-Resume|-ShowProg |-ProgDsply -PhysDrv [E0:S0,E1:S1,...] -aN|-a0,1,2|-aALL	storcli /cx/ex/sx start rebuild storcli /cx/ex/sx stop rebuild storcli /cx/ex/sx pause rebuild storcli /cx/ex/sx resume rebuild storcli /cx/ex/sx shnow rebuild
Start, stop, pause, resume, or show the progress of a copyback operation.	MegaCLI PDCpyBk -Start|-Stop |-Suspend|-Resume|-ShowProg |-ProgDsply -PhysDrv [E0:S0,E1:S1,...] -aN|-a0,1,2|-aALL	storcli /cx/ex/sx start copyback target = exx:sxx storcli /cx/ex/sx stop copyback storcli /cx/ex/sx pause copyback storcli /cx/ex/sx resume copyback storcli /cx/ex/sx show copyback
Mark a drive as missing.	MegaCli -PdMarkMissing -physdrv[E0:S0,E1:S1,...] -aN|-a0,1,2|-aALL	storcli /cx/ex/sx set missing

Table 10 Physical Drive Commands (Continued)

Description	MegaCLI Command	StorCLI Command
Show missing drive information.	`MegaCli -PdGetMissing -aN|-a0,1,2|-aALL`	`storcli /cx/ex/sx show all` **NOTE** This information is shown as part of the show all command.
Replace the configured drive that is identified as missing, and then start an automatic rebuild.	`MegaCli -PdReplaceMissing -physdrv[E0:S0] -arrayA, -rowB -aN`	`storcli /cx/ex/sx insert array=x row=y`
Set the drive state to online	`MegaCli –PDOnline -PhysDrv[E0:S0,E1:S1....] -aN|-a0,1,2`	`storcli /cx/ex/sx set online`
Set the drive state to offline.	`MegaCli –PDOffline -PhysDrv[E0:S0,E1:S1....] -aN|-a0,1,2|-aALL`	`storcli /cx/ex/sx set offline`
Set the drive state to JBOD	`MegaCli –PDMakeGood -PhysDrv[E0:S0,E1:S1....] -aN|-a0,1,2|-aALL`	`storcli /cx/ex/sx set good [force]`
Set the drive state to JBOD	`MegaCli -PDMakeJBOD -PhysDrv[E0:S0,E1:S1,...] -aN|-a0,1,2|-aALL`	`storcli /cx/ex/sx set jbod`
Add and delete hot spare drives.	`MegaCli –PDHSP {-Set [{-Dedicated -ArrayN |-Array0,1...}] [-EnclAffinity] [-nonRevertible] } | -Rmv -PhysDrv[E0:S0,E1:S1,...] -aN|-a0,1,2|-aALL`	`storcli /cx/ex/sx add hotsparedrive [dgs=<N|0,1,2..>] enclaffinity nonrevertible` `storcli /cx/ex/sx delete hotsparedrive`

Start, stop, pause, resume or show the progress of an initialization process.	`MegaCli –PDClear -Start |-Stop| -ShowProg |-ProgDsply -PhysDrv[E0:S0,E1:S1....] -aN|-a0,1,2|-aALL`	`storcli /cx/ex/sx start initialization`
		`storcli /cx/ex/sx stop initialization`
		`storcli /cx/ex/sx pause initialization`
		`storcli /cx/ex/sx resume initialization`
		`storcli /cx/ex/sx show initialization`

Start a drive locate and activate the drive's LED or stop a drive locate and deactivate the drive's LED.	`MegaCli –PDLocate {[-start] | -stop} -physdrv[E0:S0,E1:S1,...] -aN|-a0,1,2|-aALL`	`storcli /cx/ex/sx start locate` `storcli /cx/ex/sx stop locate`
Spin down an unconfigured drive and prepare it for removal or spin up spun-down drive and mark the drive state as unconfigured good.	`MegaCli –PDPrpRmv [-Undo] –PhysDrv[E0:S0,E1:S1....] -aN|-a0,1,2|-aALL`	`storcli /cx/ex/sx spindown` `storcli /cx/ex/sx spinup`
Show physical drive information of all connected drives.	`MegaCli –PDList -aN|-a0,1..|-aAll`	`storcli /cx/eall/sall show [all]` **NOTE** This command does not show drives whose enclosure device ID is not available.
Flash the physical drive firmware.	`MegaCLI PdFwDownload[offline] [ForceActivate] {[-SataBridge] -PhysDrv[0:1]}|{-EncdevId[devId1]} -f <filename> -aN|-a0,1,2|-Aall`	`storcli /cx[/ex]/sx download src=<filepath> [satabridge]` `storcli /cx/ex download src=<filepath> [forceActivate]`
Erase the drive's security configuration and securely erase data on a drive.	`MegaCli -PDInstantSecureErase -PhysDrv[E0:S0,E1:S1,...] | [-Force] -aN|-a0,1,2|-aALL`	`storcli /cx/ex/sx secureerase [force]`

Table 10 Physical Drive Commands (Continued)

Description	MegaCLI Command	StorCLI Command
Show the security key for secured physical drives	MegaCli -GetKeyID [-PhysDrv[E0:S0]] -aN	storcli /cx/ex/sx securitykey keyid
Start, stop, and show the progress of a secure erase operation	MegaCli -SecureErase Start[Simple| [Normal [|ErasePattern ErasePatternA|ErasePattern ErasePatternA ErasePattern ErasePatternB]]|[Thorough [|ErasePattern ErasePatternA|ErasePattern ErasePatternA ErasePattern ErasePatternB]]] | Stop| ShowProg| ProgDsply [-PhysDrv [E0:S0,E1:S1,...] | -Lx|-L0,1,2|-LALL] -aN|-a0,1,2|-aALL	storcli /cx[/ex]/sx start erase [simple | normal | thorough] [erasepatternA=<*val*>]\n[erasepatternB=<*val*>] Examples: storcli /cx/ex/sx start erase simple storcli /cx/ex/sx start erase normal erasepatterna=10101010 storcli /cx/ex/sx start erase thorough erasepatterna=10101010 erasepatternb=10101111 storcli /cx/ex/sx stop erase
Enable/disable the direct physical drive mapping mode.Show the current state of the direct physical drive mapping.	MegaCLI DirectPdMapping -Enbl|-Dsbl|-Dsply -aN|-a0,1,2|-Aall	storcli /cx set directpdmapping=<on | off> storcli /cx show directpdmapping

C11 Enclosure Commands

Table 11 Enclosure Commands

Description	MegaCLI Command	StorCLI Command
Show enclosure information.	MegaCli –EncInfo -aN|-a0,1,2|-aALL	storcli /cx/ex show storcli /cx/ex show all
Show enclosure status.	MegaCli –EncStatus -aN|-a0,1,2|-aALL	storcli /cx/ex show status

C.12 PHY Commands

Table 12 PHY Commands

Description	MegaCLI Command	StorCLI Command
Show PHY information.	MegaCli –PHYInfo -phyM -aN|-a0,1,2|-aALL	storcli /cx/px(x|all) show storcli /cx/px(x|all) show all
Set PHY link speed.	MegaCLI PhySetLinkSpeed -phyM -speed -aN|-a0,1,2|-aALL	storcli /cx/px(x|all) set linkspeed=0(auto)|1.5|3|6|12
Show the PHY error counters.	Megacli PhyErrorCounters -An	storcli /cx/px(x|all) show storcli /cx/px(x|all) show all

C.13 Alarm Commands

Table 13 Alarm Commands

Description	MegaCLI Command	StorCLI Command
Show alarm properties.	MegaCli -AdpGetProp AlarmDsply -aN|-a0,1,2|-aALL	storcli /c*x*(*x*|all) show alarm
Set alarm properties.	MegaCli -AdpSetProp AlarmEnbl | AlarmDsbl | AlarmSilence -aN|-a0,1,2|-aALL	storcli /c*x*(*x*|all) set alarm=<on|off|silence>

C.14 Event Log Properties Commands

Table 14 Event Log Properties Commands

Description	MegaCLI Command	StorCLI Command
Show event logs.	MegaCli -AdpEventLog -GetEventLogInfo -aN|-a0,1,2|-aALL	storcli /c*x* show eventloginfo
Show the specified type of event logs.	MegaCli -AdpEventLog -GetEvents {-info -warning -critical -fatal} {-f <*fileName*>} -aN|-a0,1,2|-aALL	storcli /c*x* show events [[type=<sincereboot| sinceshutdown| includedeleted|latest=*x*| ccincon vd=<0,1,...>] filter=<info| warning| critical|fatal>] file=<*filepath*>
Show the specified event logs.	MegaCli -AdpEventLog -GetSinceShutdown {-info -warning -critical -fatal} {-f <*fileName*>} -aN|-a0,1,2|-aALL	storcli /c*x* show events [type=[latest=*x*|ccincon vd=|[sincereboot|sinceshutdown|includedeleted|latest|ccincon]]] [filter=[info|warning|critical|fatal]] file=*xyz*.txt
Delete the event logs.	MegaCli -AdpEventLog -Clear -aN|-a0,1,2|-aALL	storcli /c*x* delete events

C.15 Premium Feature Key Commands

Table 15 Premium Feature Key Commands

Description	MegaCLI Command	StorCLI Command
Show the Safe ID of the controller.	MegaCli -ELF -GetSafeId -a0	storcli /c*x*(*x*|all) show safeid
Show the Advanced Software Options that are enabled on the controller, including the ones in trial mode.	MegaCli -ELF –ControllerFeatures -a0	storcli /c*x*(*x*|all) show all **NOTE** This information shows as part of the controller show all.
Apply the Activation Key in preview mode.	MegaCli -ELF -Applykey key –val -preview -a0	storcli /c*x*(*x*|all) set aso key=<*key value*> preview
Apply the Activation Key.	MegaCli -ELF -Applykey key –val -a0	storcli /c*x*(*x*|all) set aso key=<*key value*>

Table 15 Premium Feature Key Commands (Continued)

Description	MegaCLI Command	StorCLI Command
Deactivate the trial key.	`MegaCli -ELF –DeactivateTrialKey -a0`	`storcli /cx(x|all) set aso deactivatetrialkey`
Show the re-host information and, if re-hosting is necessary, show the controller and key vault serial numbers.	`MegaCli -ELF -ReHostInfo -a0`	`storcli /cx(x|all) show rehostinfo`
Indicate to the controller that the re-host is complete.	`MegaCli -ELF -ReHostComplete -a0`	`storcli /cx(x|all) set aso rehostcomplete`

Appendix D: Getting Help and Technical Assistance

If you need help, service, or technical assistance or just want more information about Lenovo products, you will find a wide variety of sources available from Lenovo to assist you.

Use this information to obtain additional information about Lenovo and Lenovo products, and determine what to do if you experience a problem with your Lenovo system or optional device.

Note: This section includes references to IBM web sites and information about obtaining service. IBM is Lenovo's preferred service provider for the BladeCenter, System x, Flex System, and NeXtScale System products.

D.1 Before you call

Before you call, make sure that you have taken these steps to try to solve the problem yourself.

If you believe that you require warranty service for your Lenovo product, the service technicians will be able to assist you more efficiently if you prepare before you call.

- Check all cables to make sure that they are connected.
- Check the power switches to make sure that the system and any optional devices are turned on.
- Check for updated software, firmware, and operating-system device drivers for your Lenovo product. The Lenovo Warranty terms and conditions state that you, the owner of the Lenovo product, are responsible for maintaining and updating all software and firmware for the product (unless it is covered by an additional maintenance contract). Your service technician will request that you upgrade your software and firmware if the problem has a documented solution within a software upgrade.
- If you have installed new hardware or software in your environment, check http://www.ibm.com/systems/info/x86servers/serverproven/compat/us to make sure that the hardware and software is supported by your product.
- Go to http://www.ibm.com/supportportal to check for information to help you solve the problem.
- Gather the following information to provide to the service technician. This data will help the service technician quickly provide a solution to your problem and ensure that you receive the level of service for which you might have contracted.
  - Hardware and Software Maintenance agreement contract numbers, if applicable
  - Machine type number (Lenovo 4-digit machine identifier)
  - Model number
  - Serial number
  - Current system UEFI and firmware levels
  - Other pertinent information such as error messages and logs
- Go to http://www.ibm.com/support/entry/portal/Open_service_request to submit an Electronic Service Request. Submitting an Electronic Service Request will start the process of determining a solution to your problem by making the pertinent information available to the service technicians. The IBM service technicians can start working on your solution as soon as you have completed and submitted an Electronic Service Request.

You can solve many problems without outside assistance by following the troubleshooting procedures that Lenovo provides in the online help or in the Lenovo product documentation. The Lenovo product documentation also describes the diagnostic tests that you can perform. The documentation for most systems, operating systems, and programs contains troubleshooting procedures and explanations of error messages and error codes. If you suspect a software problem, see the documentation for the operating system or program.

D.2 Using the documentation

Information about your Lenovo system and preinstalled software, if any, or optional device is available in the product documentation. That documentation can include printed documents, online documents, readme files, and help files.

See the troubleshooting information in your system documentation for instructions for using the diagnostic programs. The troubleshooting information or the diagnostic programs might tell you that you need additional or updated device drivers or other software. Lenovo maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. To access these pages, go to http://www.ibm.com/supportportal.

D.3 Getting help and information from the World Wide Web

On the World Wide Web, the IBM website has up-to-date information about IBM systems, optional devices, services, and support. You can find service information for IBM systems and optional devices at http://www.ibm.com/supportportal/.

D.4 Software service and support

Through IBM Support Line, you can get telephone assistance, for a fee, with usage, configuration, and software problems with your Lenovo products.

For more information about Support Line and other IBM services, see http://www.ibm.com/services or see http://www.ibm.com/planetwide for support telephone numbers. In the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

D.5 Hardware service and support

IBM is Lenovo's preferred service provider for the BladeCenter, System x, Flex System and NeXtScale System products.

You can receive hardware service through your Lenovo reseller or from IBM. To locate a reseller authorized by Lenovo to provide warranty service, go to http://www.ibm.com/partnerworld and click **Business Partner Locator**. For IBM support telephone numbers, see http://www.ibm.com/planetwide. In the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

In the U.S. and Canada, hardware service and support is available 24 hours a day, 7 days a week. In the U.K., these services are available Monday through Friday, from 9 a.m. to 6 p.m.

D.6 Taiwan product service

IBM is Lenovo's preferred service provider for the BladeCenter, System x, Flex System and NeXtScale System products. Use this information to contact IBM Taiwan product service.

台灣 IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路 7 號 3 樓
電話：0800-016-888

IBM Taiwan product service contact information:

IBM Taiwan Corporation
3F, No 7, Song Ren Rd.
Taipei, Taiwan
Telephone: 0800-016-888

Appendix E: Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area.

Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

E.1 Trademarks

Lenovo, the Lenovo logo, BladeCenter, Flex System, NeXtScale System, and System x are trademarks of Lenovo in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

E.2 Important notes

Processor speed indicates the internal clock speed of the microprocessor; other factors also affect application performance.

Processor speed indicates the internal clock speed of the microprocessor; other factors also affect application performance.

CD or DVD drive speed is the variable read rate. Actual speeds vary and are often less than the possible maximum.

When referring to processor storage, real and virtual storage, or channel volume, KB stands for 1 024 bytes, MB stands for 1 048 576 bytes, and GB stands for 1 073 741 824 bytes.

When referring to hard disk drive capacity or communications volume, MB stands for 1 000 000 bytes, and GB stands for 1 000 000 000 bytes. Total user-accessible capacity can vary depending on operating environments.

Maximum internal hard disk drive capacities assume the replacement of any standard hard disk drives and population of all hard-disk-drive bays with the largest currently supported drives that are available from Lenovo.

Maximum memory might require replacement of the standard memory with an optional memory module.

Each solid-state memory cell has an intrinsic, finite number of write cycles that the cell can incur. Therefore, a solid-state device has a maximum number of write cycles that it can be subjected to, expressed as `total bytes written` (TBW). A device that has exceeded this limit might fail to respond to system-generated commands or might be incapable of being written to. Lenovo is not responsible for replacement of a device that has exceeded its maximum guaranteed number of program/erase cycles, as documented in the Official Published Specifications for the device.

Lenovo makes no representations or warranties with respect to non-Lenovo products. Support (if any) for the non-Lenovo products is provided by the third party, not Lenovo.

Some software might differ from its retail version (if available) and might not include user manuals or all program functionality.

E.3 Recycling information

Lenovo encourages owners of information technology (IT) equipment to responsibly recycle their equipment when it is no longer needed. Lenovo offers a variety of programs and services to assist equipment owners in recycling their IT products. For information on recycling Lenovo products, go to: http://www.lenovo.com/recycling.

E.4 Telecommunications regulatory statement

This product may not be certified in your country for connection by any means whatsoever to interfaces of public telecommunications networks. Further certification may be required by law prior to making any such connection. Contact a Lenovo representative or reseller for any questions.

E.5 Electronic emission notices

When you attach a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices that are supplied with the monitor.

E.5.1 Federal Communications Commission (FCC) statement

Note: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Lenovo is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that might cause undesired operation.

E.5.2 Industry Canada Class A emission compliance statement

This Class A digital apparatus complies with Canadian ICES-003.

Avis de conformité à la réglementation d'Industrie Canada

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

E.5.3 Australia and New Zealand Class A statement

Attention: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

E.5.4 European Union EMC Directive conformance statement

This product is in conformity with the protection requirements of EU Council Directive 2004/108/EC on the approximation of the laws of the Member States relating to electromagnetic compatibility. Lenovo cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the installation of option cards from other manufacturers.

This product has been tested and found to comply with the limits for Class A Information Technology Equipment according to European Standard EN 55022. The limits for Class A equipment were derived for commercial and industrial environments to provide reasonable protection against interference with licensed communication equipment.

Lenovo, Einsteinova 21, 851 01 Bratislava, Slovakia

E.5.5 Germany Class A Statement

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2004/108/EG (früher 89/336/EWG) zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaaten und hält die Grenzwerte der EN 55022 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der Lenovo empfohlene Kabel angeschlossen werden. Lenovo übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung der Lenovo verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung der Lenovo gesteckt/eingebaut werden.

Deutschland:

Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Betriebsmitteln Dieses Produkt entspricht dem „Gesetz über die elektromagnetische Verträglichkeit von Betriebsmitteln" EMVG (früher „Gesetz über die elektromagnetische Verträglichkeit von Geräten"). Dies ist die Umsetzung der EU-Richtlinie 2004/108/EG (früher 89/336/EWG) in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Betriebsmitteln, EMVG vom 20. Juli 2007 (früher Gesetz über die elektromagnetische Verträglichkeit von Geräten), bzw. der EMV EG Richtlinie 2004/108/EC (früher 89/336/EWG), für Geräte der Klasse A.

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen. Verantwortlich für die Konformitätserklärung nach Paragraf 5 des EMVG ist die Lenovo (Deutschland) GmbH, Gropiusplatz 10, D-70563 Stuttgart.

Informationen in Hinsicht EMVG Paragraf 4 Abs. (1) 4: **Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 Klasse A.**

Nach der EN 55022: „Dies ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funkstörungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen durchzuführen und dafür aufzukommen."

Nach dem EMVG: „Geräte dürfen an Orten, für die sie nicht ausreichend entstört sind, nur mit besonderer Genehmigung des Bundesministers für Post und Telekommunikation oder des Bundesamtes für Post und Telekommunikation betrieben werden. Die Genehmigung wird erteilt, wenn keine elektromagnetischen Störungen zu erwarten sind." (Auszug aus dem EMVG, Paragraph 3, Abs. 4). Dieses Genehmigungsverfahren ist nach Paragraph 9 EMVG in Verbindung mit der entsprechenden Kostenverordnung (Amtsblatt 14/93) kostenpflichtig.

Anmerkung: Um die Einhaltung des EMVG sicherzustellen sind die Geräte, wie in den Handbüchern angegeben, zu installieren und zu betreiben.

E.5.6 Japan VCCI Class A statement

この装置は、クラス A 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。 VCCI-A

This is a Class A product based on the standard of the Voluntary Control Council for Interference (VCCI). If this equipment is used in a domestic environment, radio interference may occur, in which case the user may be required to take corrective actions.

E.5.7 Korea Communications Commission (KCC) statement

이 기기는 업무용(A급)으로 전자파적합기기로서 판매자 또는 사용자는 이 점을 주의하시기 바라며, 가정외의 지역에서 사용하는 것을 목적으로 합니다.

This is electromagnetic wave compatibility equipment for business (Type A). Sellers and users need to pay attention to it. This is for any areas other than home.

E.5.8 Russia Electromagnetic Interference (EMI) Class A statement

ВНИМАНИЕ! Настоящее изделие относится к классу А.
В жилых помещениях оно может создавать радиопомехи, для снижения которых необходимы дополнительные меры

E.5.9 People's Republic of China Class A electronic emission statement

中华人民共和国 "A类" 警告声明

声 明

此为A级产品，在生活环境中，该产品可能会造成无线电干扰。在这种情况下，可能需要用户对其干扰采取切实可行的措施。

E.5.10 Taiwan Class A compliance statement