ThinkSystem Storage Adapter Software User Guide

Note

Before using this information and the product it supports, read the general information in Appendix A “Getting help and technical assistance”, Appendix B “Notices”, the safety information, warranties, and licenses information on the Lenovo Web site at: https://support.lenovo.com/documents/LNVO-DOCS.

First Edition (August 2017)

Contents

Chapter 1: Overview

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

If you want to use a different software application to perform these procedures, refer to the LSI Storage Authority (LSA) software. The LSI Storage Authority software is a web-based application that enables you to monitor, maintain, troubleshoot, and configure the LSI MegaRAID products. The LSI Storage Authority graphical user interface (GUI) helps you to view, create, and manage storage configurations.

1.1 SAS Technology

The ThinkSystem 12Gb/s SAS RAID controllers are high-performance intelligent PCI Express-to-SAS/SATA controllers with RAID control capabilities. The ThinkSystem 12Gb/s SAS RAID controllers provide reliability, high performance, and fault-tolerant disk subsystem management. These controllers are an ideal RAID solution for the internal storage of workgroup, departmental, and enterprise systems., and they offer a cost-effective way to implement RAID in a server.

SAS technology brings a wealth of options and flexibility with the use of SAS devices, and SATA II and SATA III devices within the same storage infrastructure. These devices bring individual characteristics that make each of these more suitable choice depending on your storage needs. ThinkSystem gives you the flexibility to combine these two similar technologies on the same controller, within the same enclosure, and in the same virtual drive.

NOTE Carefully assess any decision to combine SAS drives and SATA drives within the same virtual drives. Avoid mixing drives.

The ThinkSystem 12Gb/s SAS RAID controllers are based on the Broadcom's first-to-market SAS IC technology and proven ThinkSystem technology. As third-generation PCI Express RAID controllers, the ThinkSystem SAS RAID controllers address the growing demand for increased data throughput and scalability requirements across midrange and enterprise-class server platforms. Lenovo offers a family of ThinkSystem SAS RAID controllers addressing the needs for both internal and external solutions.

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

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

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

1.2 Serial-Attached SCSI Device Interface

SAS is a serial, point-to-point, enterprise-level device interface that leverages the proven SCSI protocol set. SAS is a convergence of the advantages of SATA, SCSI, and Fibre Channel, and is the future mainstay of the enterprise and high-end workstation storage markets. SAS offers a higher bandwidth per pin than parallel SCSI, and it improves the signal and data integrity.

The SAS interface uses the proven SCSI command set to ensure reliable data transfers, while providing the connectivity and flexibility of point-to-point serial data transfers. The serial transmission of SCSI commands eliminates clock-skew challenges. The SAS interface provides improved performance, simplified cabling, smaller connectors, lower pin count, and lower power requirements when compared to parallel SCSI.

SAS controllers leverage a common electrical and physical connection interface that is compatible with Serial ATA technology. The SAS and SATA protocols use a thin, 7-wire connector instead of the 68-wire SCSI cable or 26-wire ATA cable. The SAS/SATA connector and cable are easier to manipulate, allow connections to smaller devices, and do not inhibit airflow. The point-to-point SATA architecture eliminates inherent difficulties created by the legacy ATA master-slave architecture, while maintaining compatibility with existing ATA firmware.

1.3 Serial ATA III Features

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

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

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

1.4 Solid State Drive Features

The ThinkSystem firmware supports the use of SSDs as standard drives. SSD drives are expected to behave like SATA or SAS HDDs except for the following:

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

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

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

1.4.1 SSD Guard

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

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

1.5 UEFI 2.0 Support

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

Chapter 2: Introduction to RAID

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

RAID Description

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

RAID Benefits

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

RAID Functions

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

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

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

2.1 Components and Features

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

2.1.1 Drive Group

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

2.1.2 Virtual Drive

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

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

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

2.1.3 Fault Tolerance

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

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

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

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

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

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

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

2.1.3.1 Multipathing

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

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

Multipathing provides the following features:

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

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

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

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

2.1.3.2 True Multipathing

A device, connected in multi-path, configured as JBOD, has each of the individual paths exposed directly to the host. The host handles multipathing to the device and manages them. The firmware presents the drivers with a unique target ID per device path, allowing the host to discover both paths as distinct SCSI devices. The firmware also presents the drivers with a unique device handle for each path, enabling the driver to issue fast path I/Os to either path of the device.

NOTE True multipath is not supported on SATA devices.

2.1.4 Consistency Check

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

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

2.1.5 Replace

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

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

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

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

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

Order of Precedence

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

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

2.1.6 Background Initialization

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

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

New RAID 5 virtual drives and new RAID 6 virtual drives require a minimum number of drives for the background initialization to start. If fewer drives exist, the background initialization does not start. The background initialization needs to be started manually by initiating a consistency check.

The following number of drives are required to start a background initialization:

- New RAID 5 virtual drives must have at least five drives.
- New RAID 6 virtual drives must have at least seven drives.

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

2.1.7 Patrol Read

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

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

You can use the LSI Storage Authority Software to select the patrol read options, which you can use to set automatic or manual operation, or disable patrol read. See Patrol Read.

2.1.8 Disk Striping

Disk striping lets you write data across multiple drives instead of just one drive. Disk striping involves partitioning each drive storage space into stripes that can vary in size from a minimum of 64 KB to 1 MB or be a fixed 64 KB, depending on your ThinkSystem RAID adapter. These stripes are interleaved in a repeated sequential manner. The combined storage space is composed of stripes from each drive. It is recommended that you keep stripe sizes the same across RAID drive groups.

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

Figure 1 Example of Disk Striping (RAID 0)

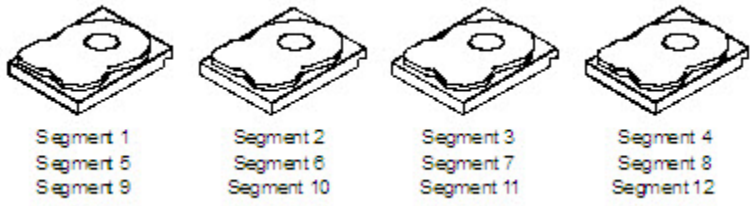

Stripe Width

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

Stripe Size

The stripe size is the length of the interleaved data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 1 MB of drive space and has 64 KB of data residing on each drive in the stripe. In this case, the stripe size is 1 MB and the strip size is 64 KB.

Strip Size

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

2.1.9 Disk Mirroring

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

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

Figure 2 Example of Disk Mirroring (RAID 1)

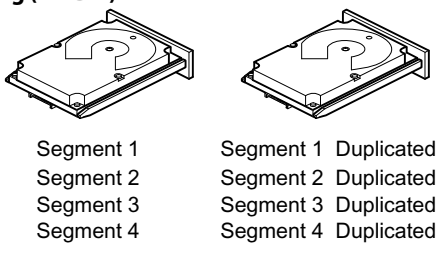

2.1.10 Parity

Parity generates a set of redundancy data from two or more parent data sets. The redundancy data can be used to reconstruct one of the parent data sets in the event of a drive failure. Parity data does not fully duplicate the parent

data sets, but parity generation can slow the write process. In a RAID drive group, this method is applied to entire drives or stripes across all of the drives in a drive group. The types of parity are described in the following table.

Table 1 Types of Parity

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

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

Figure 3 Example of Distributed Parity (RAID 5 Drive Group)

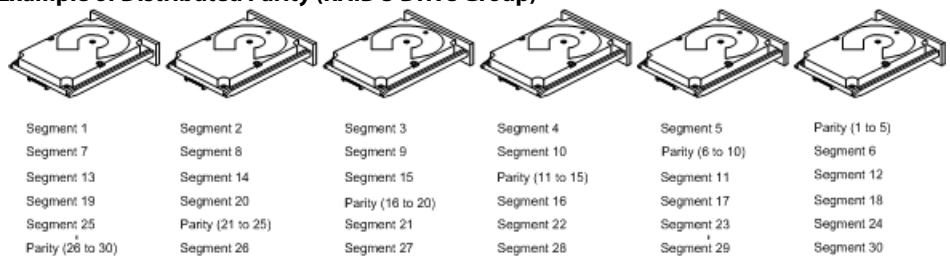

Note: Parity is distributed across all drives in the drive group.

3_01081-00

2.1.11 Disk Spanning

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

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

Figure 4 Example of Disk Spanning

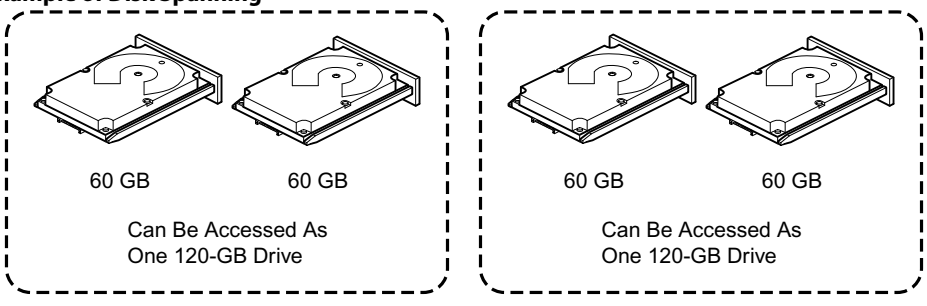

3_01082-00

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

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

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

See StorCLI for detailed procedures for configuring drive groups and virtual drives, and spanning the drives.

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

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

NOTE In a spanned virtual drive (RAID 10, RAID 50, RAID 60) the span numbering starts from Span 0, Span 1, Span 2, and so on.

2.1.12 Hot Spares

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

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

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

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

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

The hot spare can be of two types:

- Global hot spare
- Dedicated hot spare

Global Hot Spare

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

Dedicated Hot Spare

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

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

Observe the following parameters when using hot spares:

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

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

2.1.13 Disk Rebuilds

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

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

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

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

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

2.1.14 Rebuild Rate

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

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

2.1.15 Hot Swap

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

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

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

2.1.16 Drive States

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

Table 3 Drive States

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

2.1.17 Virtual Drive States

The virtual drive states are described in the following table.

Table 4 Virtual Drive States

State	Description
Optimal	The virtual drive operating condition is good. All configured drives are online.
Degraded	The virtual drive operating condition is not optimal. One of the configured drives has failed or is offline.
Partial Degraded	The operating condition in a RAID 6 virtual drive is not optimal. One of the configured drives has failed or is offline. A RAID 6 drive group can tolerate up to two drive failures.
Failed	The virtual drive has failed.
Offline	The virtual drive is not available to the RAID controller.
Foreign	A foreign configuration is a RAID configuration that already exists on a replacement set of drives that you install in a computer system. WebBIOS Configuration Utility and the LSI Storage Authority Software allows you to import the existing configuration to the RAID controller or clear the configuration so you can create a new one.

2.1.18 Enclosure Management

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

2.2 RAID Levels

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

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

2.2.1 Summary of RAID Levels

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

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

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

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

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

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

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

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

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

2.2.2 Selecting a RAID Level

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

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

2.2.3 RAID 0 Drive Groups

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

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

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

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

Table 5 RAID 0 Drive Group Overview

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

Figure 5 RAID 0 Drive Group Example with Two Drives

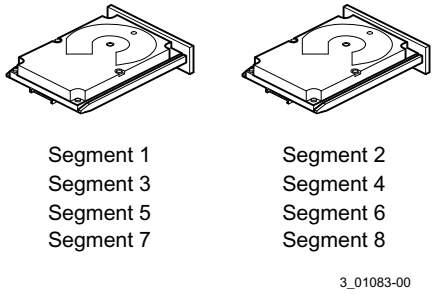

2.2.4 RAID 1 Drive Groups

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

Table 6 RAID 1 Drive Group Overview

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

Figure 6 RAID 1 Drive Group

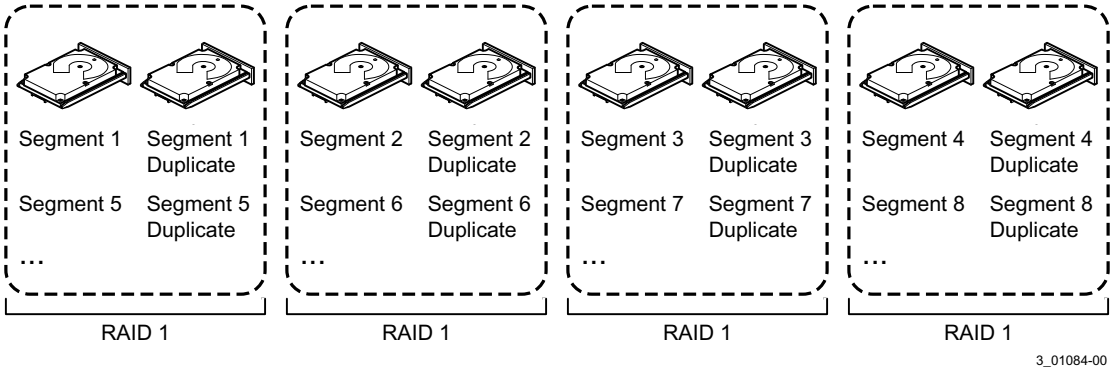

2.2.5 RAID 5 Drive Groups

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

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

Table 7 RAID 5 Drive Group Overview

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

Figure 7 RAID 5 Drive Group with Six Drives

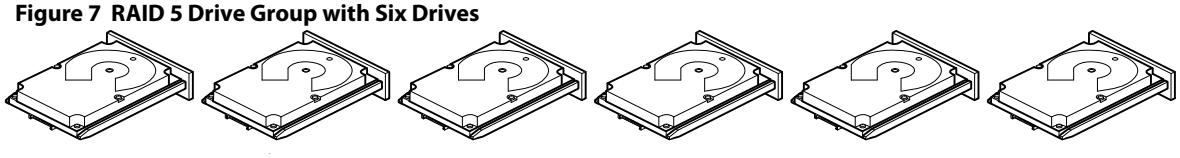

Segment 1	Segment 2	Segment 3	Segment 4	Segment 5	Parity (1–5)
Segment 7	Segment 8	Segment 9	Segment 10	Parity (6–10)	Segment 6
Segment 13	Segment 14	Segment 15	Parity (11–15)	Segment 11	Segment 12
Segment 19	Segment 20	Parity (16–20)	Segment 16	Segment 17	Segment 18
Segment 25	Parity (21–25)	Segment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

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

3_01085-00

2.2.6 RAID 6 Drive Groups

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

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

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

Table 8 RAID 6 Drive Group Overview

Uses	Use for any application that has high read request rates but low random or small block write rates.
Strong points	Provides data redundancy, high read rates, and good performance in most environments. Can survive the loss of two drives or the loss of a drive while another drive is being rebuilt. Provides the highest level of protection against drive failures of all of the RAID levels. Performance is similar to that of a RAID 5 drive group.
Weak points	Not well-suited to tasks requiring a lot of small and/or random write operations. A RAID 6 virtual drive must generate two sets of parity data for each write operation, which results in a significant decrease in performance during write operations. Drive performance is reduced during a drive Rebuild operation. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes. A RAID 6 drive group costs more because of the extra capacity required by using two parity blocks per stripe.
Drives	4 through 32.

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

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

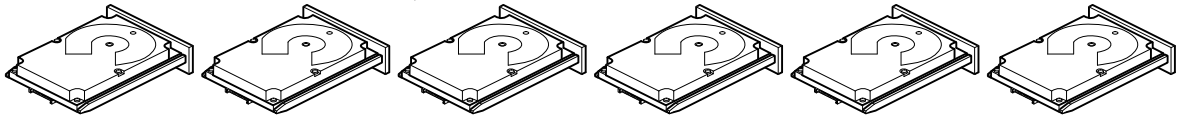

Segment 1	Segment 2	Segment 3	Segment 4	Parity (P1–P4)	Parity (Q1–Q4)
Segment 6	Segment 7	Segment 8	Parity (P5–P8)	Parity (Q5–Q8)	Segment 5
Segment 11	Segment 12	Parity (P9–P12)	Parity (Q9–Q12)	Segment 9	Segment 10
Segment 16	Parity (P13–P16)	Parity (Q13–Q16)	Segment 13	Segment 14	Segment 15
Parity (P17–P20)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

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

3_01086-00

2.2.7 RAID 00 Drive Groups

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

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

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

Table 9 RAID 00 Drive Group Overview

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

Figure 9 RAID 00 Drive Group Example with Two Drives

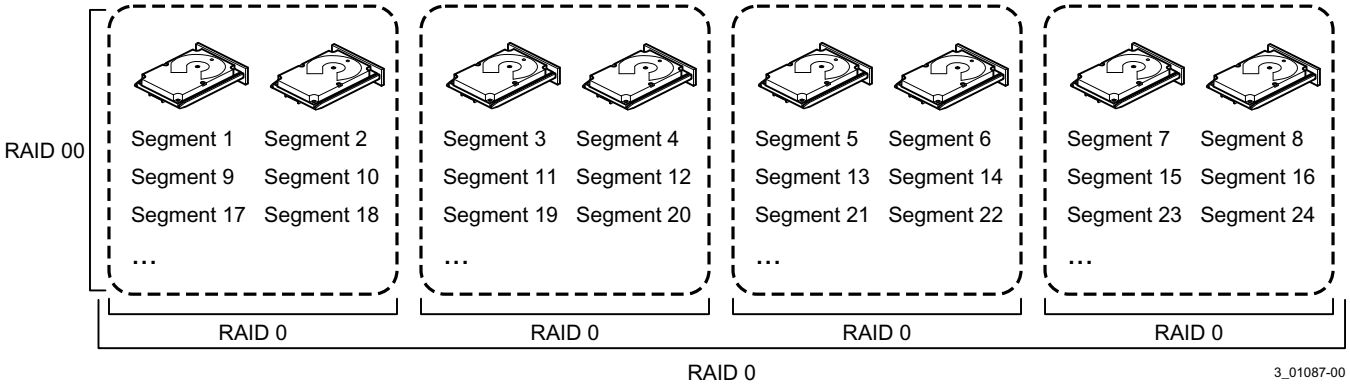

2.2.8 RAID 10 Drive Groups

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

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

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

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

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

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

Table 10 RAID 10 Drive Group Overview

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

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

Figure 10 RAID 10 Level Virtual Drive

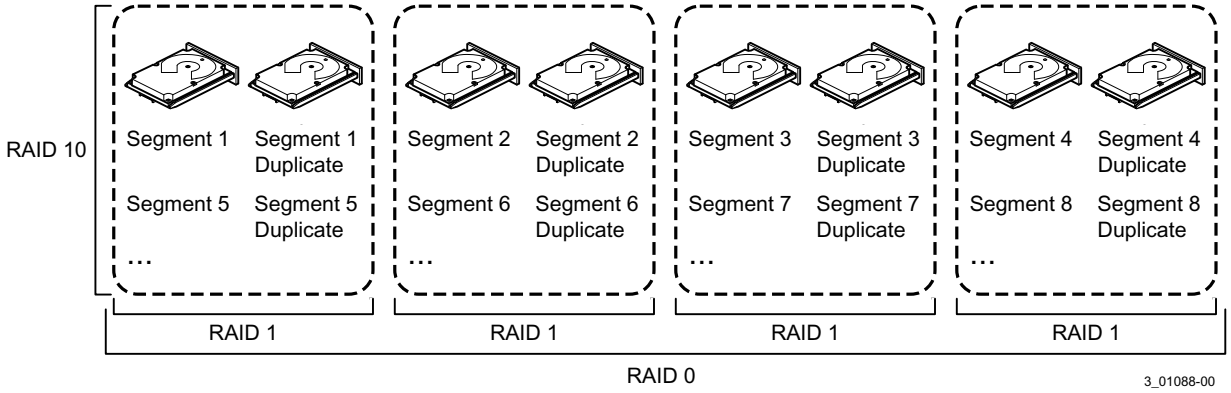

2.2.9 RAID 50 Drive Groups

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

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

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

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

Table 11 RAID 50 Drive Group Overview

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

Figure 11 RAID 50 Level Virtual Drive

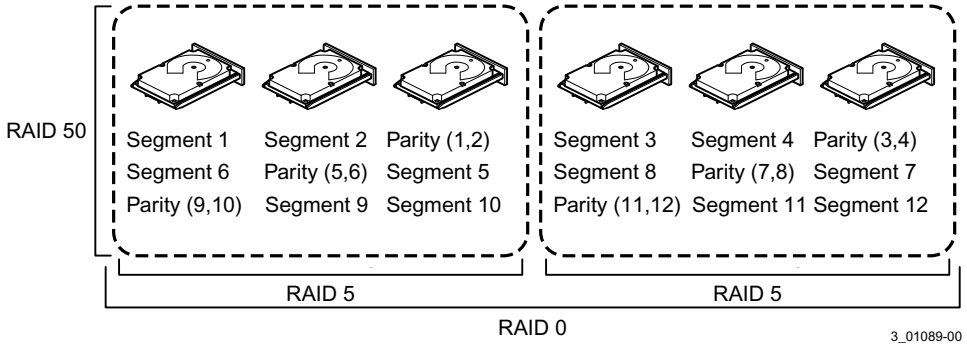

2.2.10 RAID 60 Drive Groups

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

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

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

Table 12 RAID 60 Drive Group Overview

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

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

Figure 12 RAID 60 Level Virtual Drive

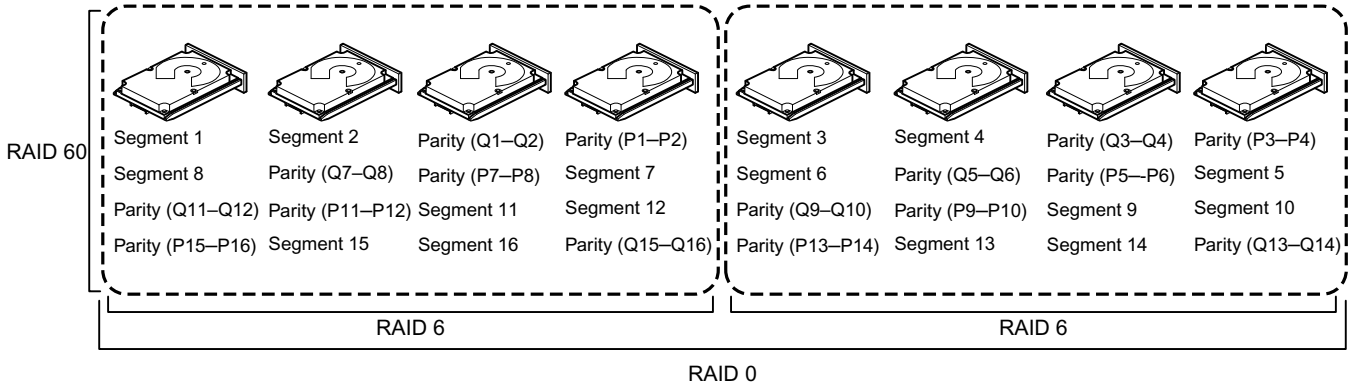

Note: Parity is distributed across all drives in the drive group.

3_01090-00

2.3 RAID Configuration Strategies

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

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

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

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

2.3.1 Maximizing Fault Tolerance

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

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

Table 13 RAID Levels and Fault Tolerance

RAID Level	Fault Tolerance
0	Does not provide fault tolerance. All data is lost if any drive fails. Disk striping writes data across multiple drives instead of just one drive. It involves partitioning each drive storage space into stripes that can vary in size. A RAID 0 drive group is ideal for applications that require high performance but do not require fault tolerance.
1	Provides complete data redundancy. If one drive fails, the contents of the other drive in the drive group can be used to run the system and reconstruct the failed drive. The primary advantage of disk mirroring is that it provides 100 percent data redundancy. Because the contents of the drive are completely written to a second drive, no data is lost if one of the drives fails. Both drives contain the same data at all times. A RAID 1 drive group is ideal for any application that requires fault tolerance and minimal capacity.
5	Combines distributed parity with disk striping. Parity provides redundancy for one drive failure without duplicating the contents of entire drives. If a drive fails, the RAID controller uses the parity data to reconstruct all missing information. In a RAID 5 drive group, this method is applied to entire drives or stripes across all drives in a drive group. Using distributed parity, a RAID 5 drive group offers fault tolerance with limited overhead.
6	Combines distributed parity with disk striping. A RAID 6 drive group can sustain two drive failures and still maintain data integrity. Parity provides redundancy for two drive failures without duplicating the contents of entire drives. If a drive fails, the RAID controller uses the parity data to reconstruct all missing information. In a RAID 6 drive group, this method is applied to entire drives or stripes across all of the drives in a drive group. Using distributed parity, a RAID 6 drive group offers fault tolerance with limited overhead.
00	Does not provide fault tolerance. All data in a virtual drive is lost if any drive in that virtual drive fails. Disk striping writes data across multiple drives instead of just one drive. It involves partitioning each drive storage space into stripes that can vary in size. A RAID 00 drive group is ideal for applications that require high bandwidth but do not require fault tolerance.
10	Provides complete data redundancy using striping across spanned RAID 1 drive groups. A RAID 10 drive group works well for any environment that requires the 100 percent redundancy offered by mirrored drive groups. A RAID 10 drive group can sustain a drive failure in each mirrored drive group and maintain data integrity.
50	Provides data redundancy using distributed parity across spanned RAID 5 drive groups. A RAID 50 drive group includes both parity and disk striping across multiple drives. If a drive fails, the RAID controller uses the parity data to re-create all missing information. A RAID 50 drive group can sustain one drive failure per RAID 5 drive group and still maintain data integrity.
60	Provides data redundancy using distributed parity across spanned RAID 6 drive groups. A RAID 60 drive group can sustain two drive failures per RAID 6 drive group and still maintain data integrity. It provides the highest level of protection against drive failures of all of the RAID levels. A RAID 60 drive group includes both parity and disk striping across multiple drives. If a drive fails, the RAID controller uses the parity data to re-create all missing information.

2.3.2 Maximizing Performance

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

Table 14 RAID Levels and Performance

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

2.3.3 Maximizing Storage Capacity

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

Table 15 RAID Levels and Capacity

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

2.4 RAID Availability

2.4.1 RAID Availability Concept

Data availability without downtime is essential for many types of data processing and storage systems. Businesses want to avoid the financial costs and customer frustration associated with failed servers. RAID helps you maintain data availability and avoid downtime for the servers that provide that data. RAID offers several features, such as spare drives

and rebuilds, that you can use to fix any drive problems, while keeping the servers running and data available. The following subsections describe these features.

Spare Drives

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

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

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

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

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

Rebuilding

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

2.5 Configuration Planning

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

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

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

2.6 Number of Drives

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

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

Drive Group Purpose

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

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

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

Table 16 Factors to Consider for Drive Group Configuration

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

Chapter 3: SafeStore Disk Encryption

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

Overview

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

With the SafeStore encryption service, data is encrypted by the drives. You can designate which data to encrypt at the individual virtual drive (VD) level.

Any encryption solution requires management of the encryption keys. The security service provides a way to manage these keys. Both the WebBIOS Configuration Utility and the LSI Storage Authority Software offer procedures that you can use to manage the security settings for the drives.

Purpose and Benefits

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

In addition, SafeStore local key management removes the administrator from most of the daily tasks of securing data, thereby reducing user error and decreasing the risk of data loss. Also, SafeStore local key management supports instant secure erase of drives that permanently removes data when repurposing or decommissioning drives. These services provide a much more secure level of data erasure than other common erasure methods, such as overwriting or degaussing.

Terminology

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

Table 17 Terms Used in the SafeStore Encryption Feature

Option	Description
Authenticated Mode	The RAID configuration is keyed to a user password. The password must be provided on system boot to authenticate the user and facilitate unlocking the configuration for user access to the encrypted data.
Key backup	You need to provide the controller with a lock key if the controller is replaced or if you choose to migrate secure virtual disks. To do this task, you must back up the security key.
Re-provisioning	Re-provisioning disables the security system of a device. For a controller, it involves destroying the security key. For SafeStore encrypted drives, when the drive lock key is deleted, the drive is unlocked and any user data on the drive is securely deleted. This situation does not apply to controller-encrypted drives, because deleting the virtual disk destroys the encryption keys and causes a secure erase. See Instant Secure Erase, for information about the instant secure erase feature.
Security Key	A key based on a user-provided string. The controller uses the security key to lock and unlock access to the secure user data. If the security key is unavailable, user data is irretrievably lost. You must take all precautions to never lose the security key.
Un-Authenticated Mode	This mode allows controller to boot and unlock access to user configuration without user intervention.

3.1 Workflow

Overview

The SafeStore workflow follows:

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

3.1.1 Enable Security

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

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

- Create the security key identifier
- Create the security key
- Create a password (optional)

Create the Security Key Identifier

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

Create the Security Key

You need to enter the security key to perform certain operations. You can choose a strong security key that the controller suggests. The security key must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (for example, < > @ +).

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

Create a Password

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

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

3.1.2 Change Security

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

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

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

Change the Security Key Identifier

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

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

Change the Security Key

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

Add or Change the Password

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

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

3.1.3 Create Secure Virtual Drives

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

Simple Configuration

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

Advanced Configuration

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

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

3.1.4 Import a Foreign Configuration

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

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

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

3.2 Instant Secure Erase

Instant Secure Erase is a feature used to erase data from encrypted drives. After the initial investment for an encrypted disk, there is no additional cost in dollars or time to erase data using the Instant Secure Erase feature.

You can change the encryption key for all ThinkSystem RAID controllers that are connected to encrypted drives. All encrypted drives, whether locked or unlocked, always have an encryption key. This key is set by the drive and is always active. When the drive is unlocked, the data to host from the drive (on read operations) and from the host to the drive cache (on write operations) is always provided. However, when resting on the drive platters, the data is always encrypted by the drive.

You might not want to lock your drives because you must manage a password if they are locked. Even if you do not lock the drives, a benefit still exists to using encrypted disks.

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

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

Consider the following reasons for using instant secure erase.

To repurpose the hard drive for a different application

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

To replace drives

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

To return a disk for warranty activity

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

Chapter 4: Ctrl-R Utility

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

4.1 Overview

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

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

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

NOTE Only 26 characters display for the name of the controller. The product name displays in the VD Mgmt page, Properties page, Foreign tab, controller section menu, and on top of the CTRL-H page.

4.2 Starting the Ctrl-R Utility

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

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

   ```
   Copyright© Broadcom
   Press <Ctrl><R> for Ctrl-R
   ```

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

4.3 Exiting the Ctrl-R Utility

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

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

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

4.4 Ctrl-R Utility Keystrokes

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

Table 18 Ctrl-R Utility Keystrokes

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

4.5 Ctrl-R Utility Menus

The Ctrl-R Utility contains the following menus:

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

4.5.1 Virtual Drive Management Menu

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

Figure 13 Virtual Drive Management Screen

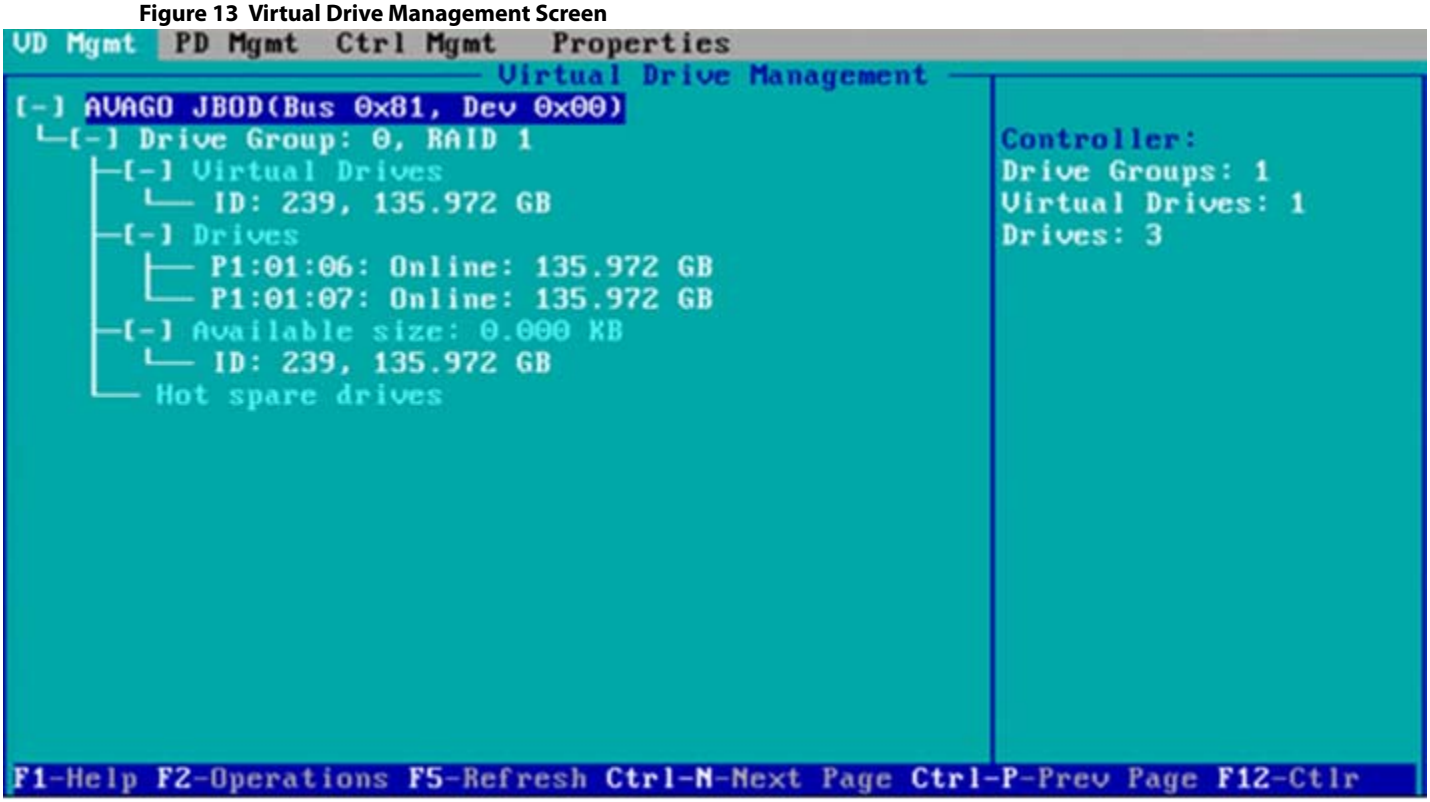

This screen shows information on the configuration of controllers, drive groups, and virtual drives.

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

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

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

4.5.2 Physical Drive Management Menu

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

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

The following figure displays the summary information.

Figure 14 Physical Drive Management Summary Screen

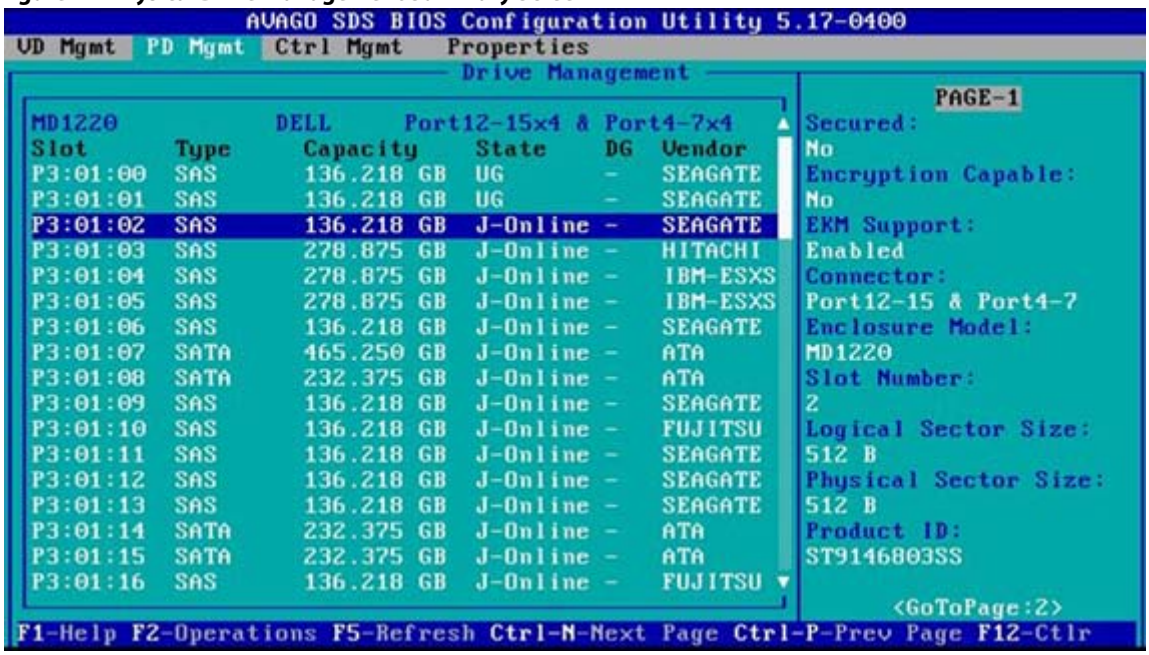

The right panel of the screen shows additional attributes of the selected device.

4.5.3 Controller Management Menu

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

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

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

Figure 15 Controller Settings – First Screen

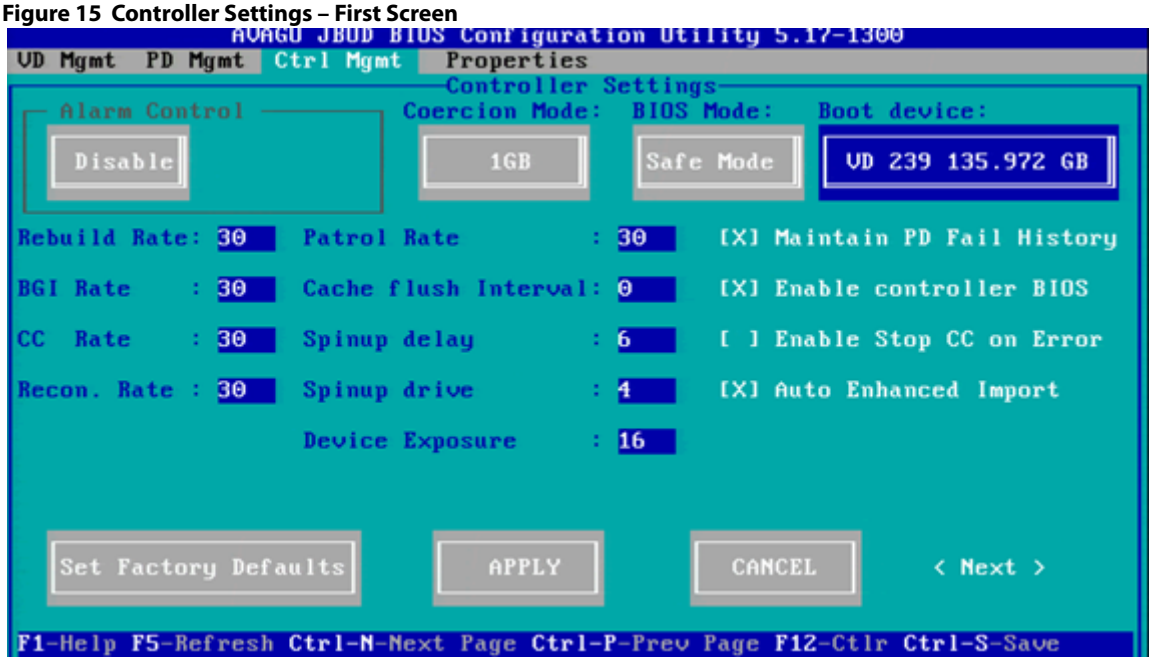

Click **Next** to open the **Controller Settings** screen (as shown in the following figure). You can manage the Link Speed, Power Save, manage battery settings, manage Mode and Parameters, begin a Start Manual Learn Cycle, enable or disable Write Verify, and enable or disable large I/O support.

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

Figure 16 Controller Settings – Second Screen

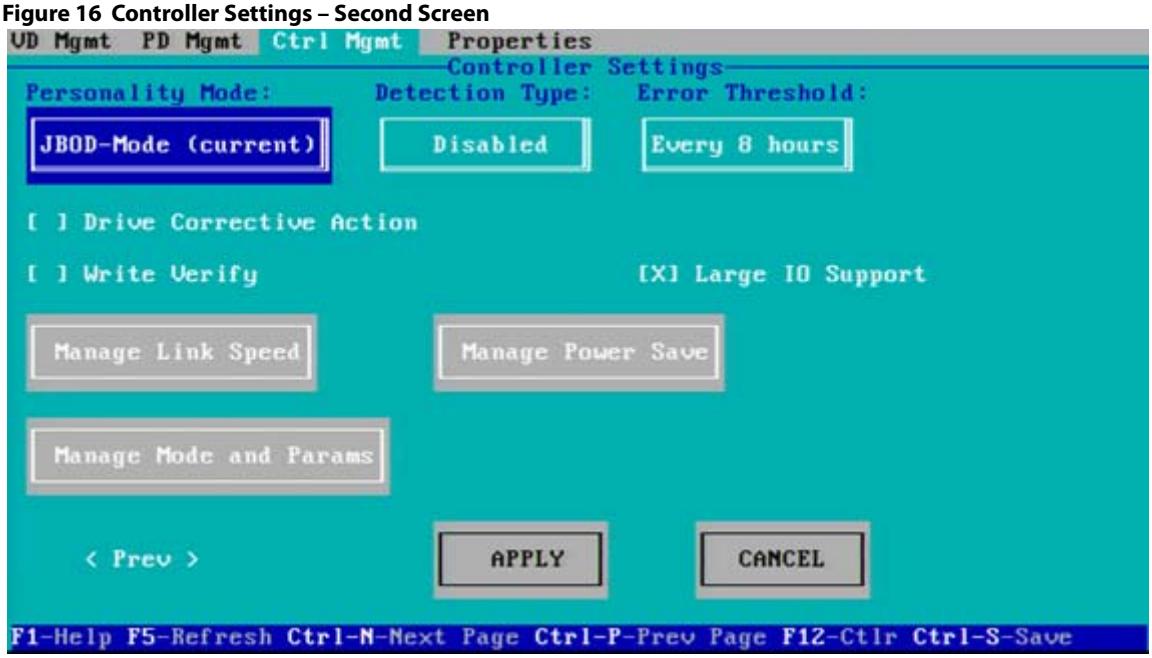

4.5.4 Properties Menu

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

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

Figure 17 Properties

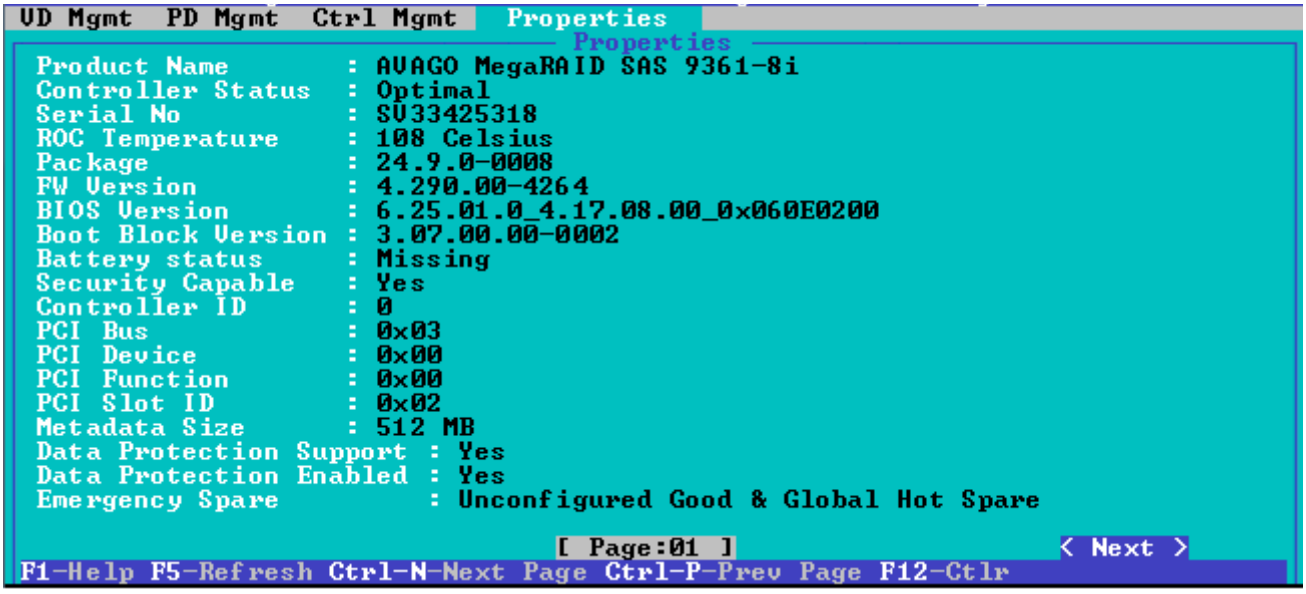

To view additional properties, you can navigate to **Next** and press Enter. The second **Properties** screen shows information, such as maximum cache size, drive standby time, battery status, and power saving properties.

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

4.5.5 Foreign View Menu

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

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

Figure 18 Foreign View Menu

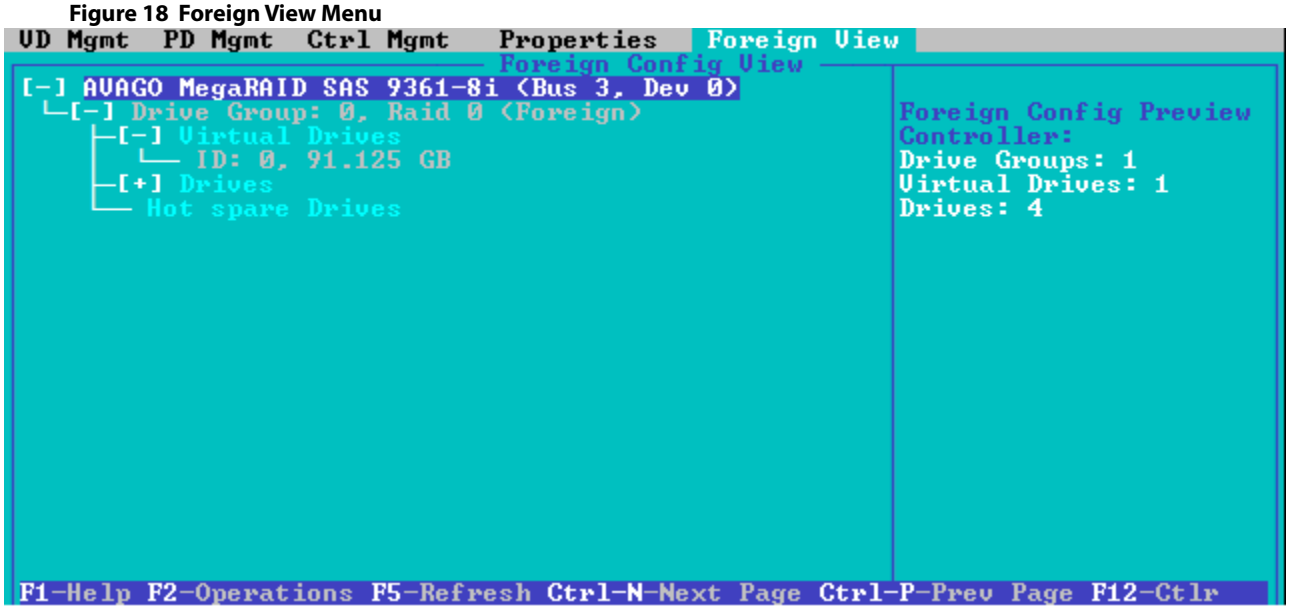

You can use the **Foreign Config View** screen to view information about the foreign configuration, such as drive groups, virtual drives, physical drives, and hot spares.

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

4.6 Managing Software Licensing

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

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

4.6.1 Managing Advanced Software Options

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

1. In the **VD Mgmt** screen, navigate to the controller and press the **F2** key.
2. Navigate to **Advanced Software Options**, and press **Enter**.

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

Figure 19 Manage MegaRAID Advanced Software Options

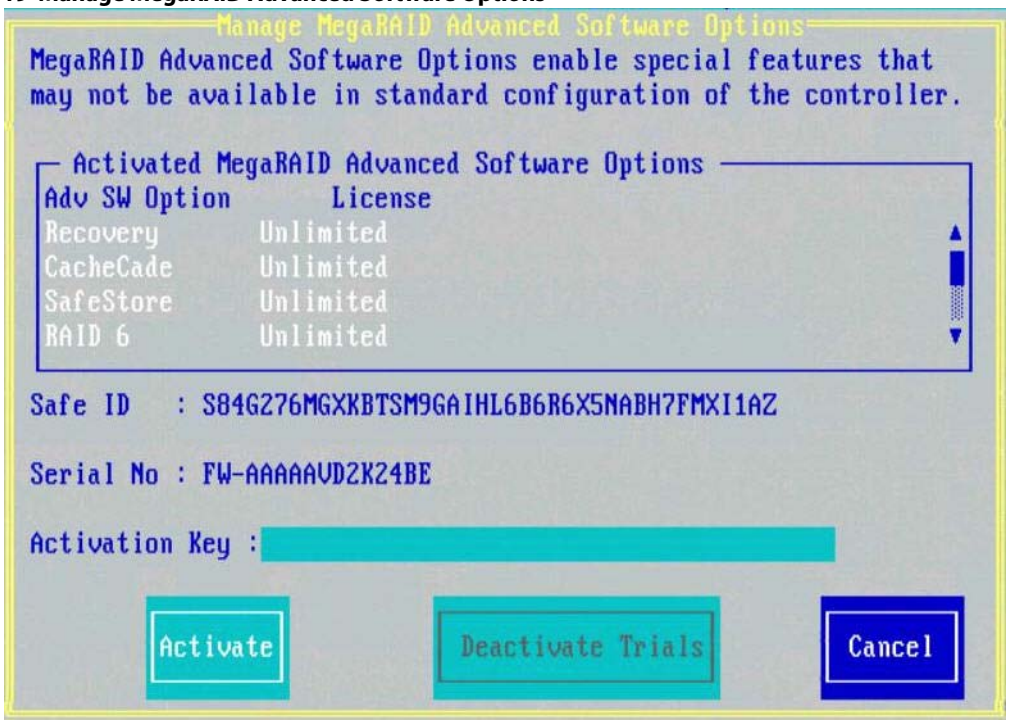

The **Activated MegaRAID Advanced Software Options** box contains the **Adv SW Option** and **License** columns.

— The **Adv SW Option** column shows the list of advanced software features available in the controller.

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

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

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

   The **Advanced Software Options Summary** dialog appears, as shown in Figure 24 on page 49.
5. Click **Deactivate Trials**.

   The **Deactivate Trial** dialog appears, as shown in the following figure.

Figure 20 Deactivate Trial Dialog

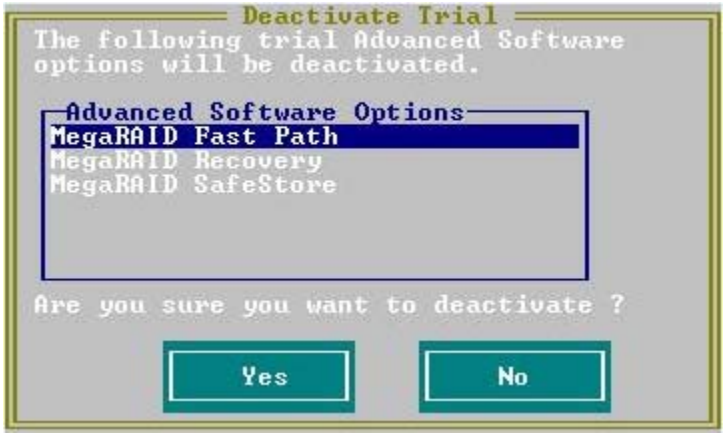

6. Perform one of these actions:
   — If you want to *deactivate* the software that is being used with a trial key, press **Yes**.
   — If you do not want to deactivate the software, press **No**.

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

- Scenario 1

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

Figure 21 Invalid Activation Key Message

- Scenario 2

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

Figure 22 Activation Key Left Blank

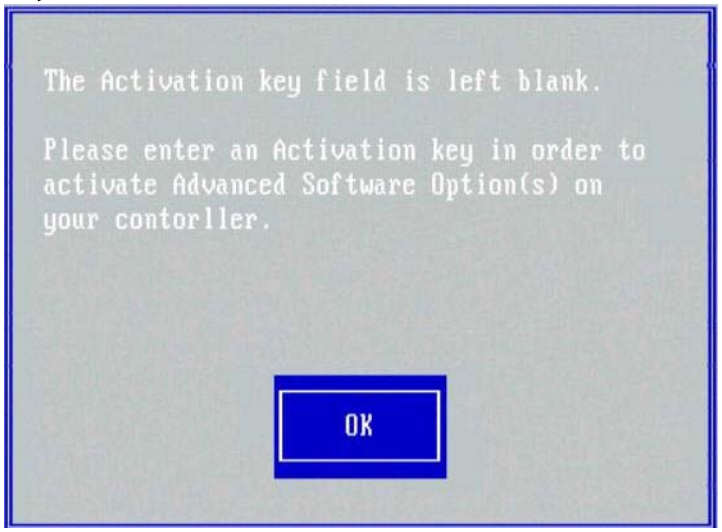

- Scenario 3

  If you enter an *incorrect* activation key, and if there is a mismatch between the activation key and the controller, the following message appears.

Figure 23 Activation Key Mismatch Message

4.6.2 Managing Advanced Software Summary

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

Figure 24 Advanced Software Options Summary

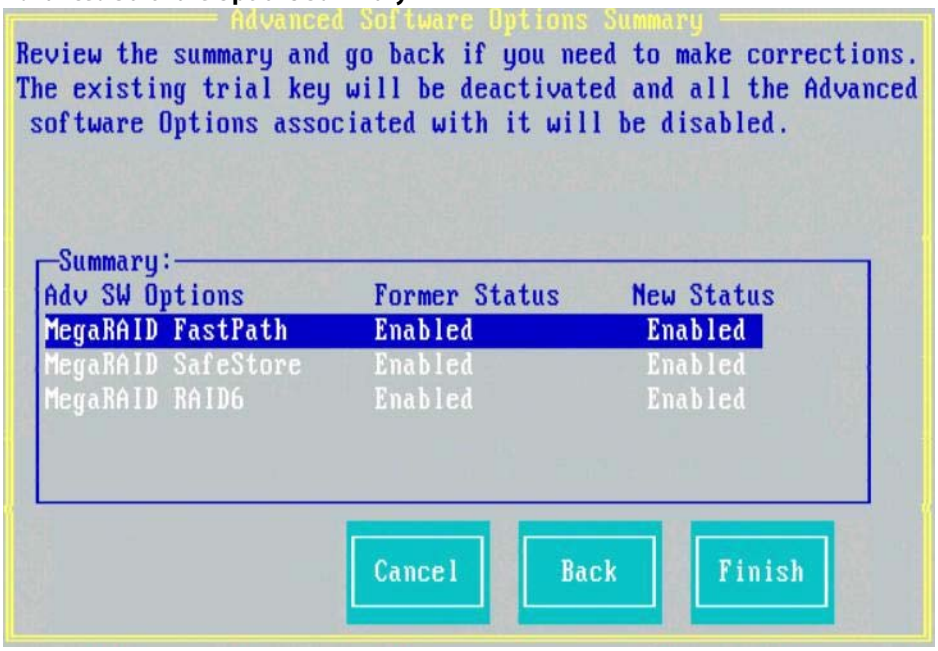

The **Summary** box shows the list of the advanced software options along with their former status and new status.

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

4.6.3 Activating an Unlimited Key over a Trial Key

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

Figure 25 Activating an Unlimited Key over a Trial Key

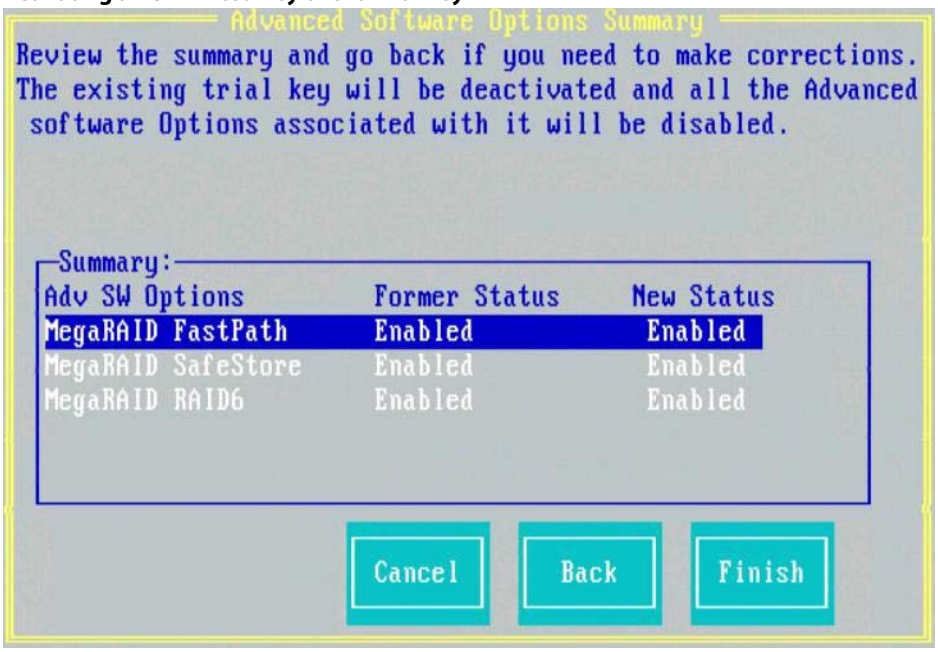

4.6.4 Activating a Trial Software

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

Figure 26 Activating a Trial Software Application

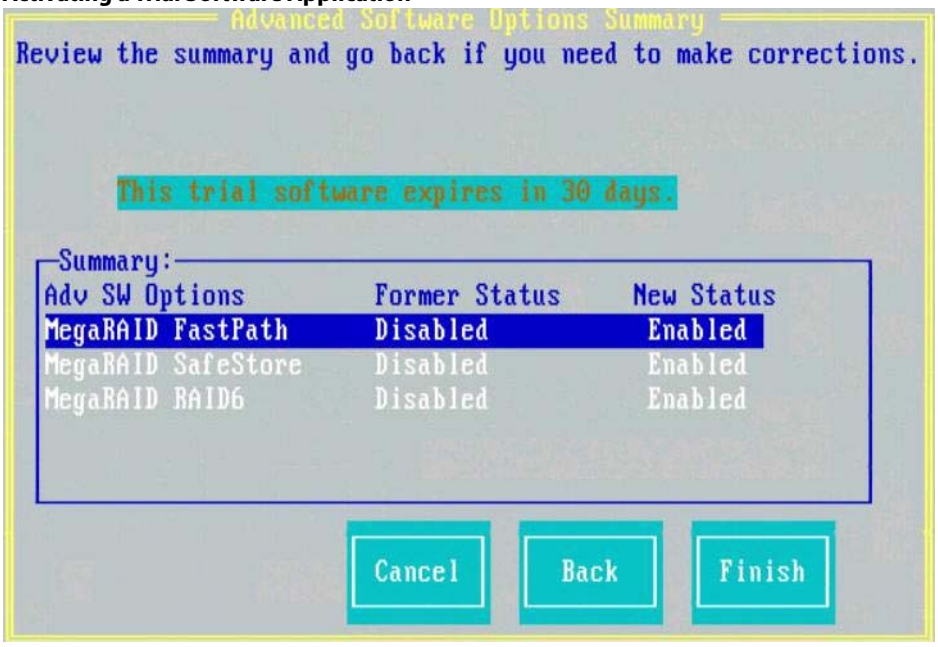

4.6.5 Activating an Unlimited Key

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

Figure 27 Activating an Unlimited Key

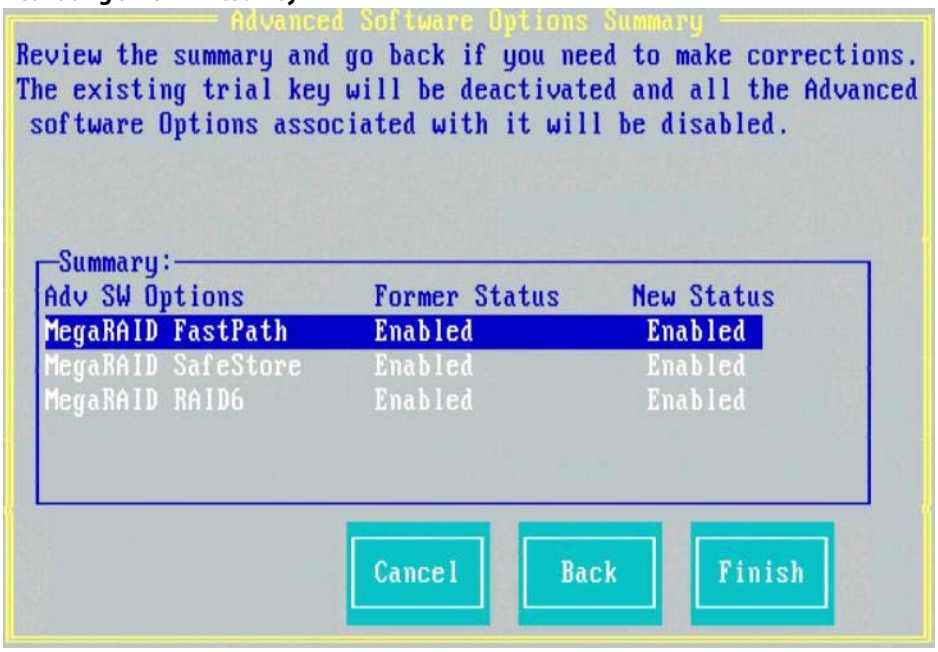

4.7 Creating a Storage Configuration

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

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

Table 19 RAID Levels

Level	Description
RAID 0	Uses data striping on two or more drives to provide high data throughput, especially for large files in an environment that requires no data redundancy.
RAID 1	Uses data mirroring on pairs of drives so that data written to one drive is simultaneously written to the other drive. RAID 1 works well for small databases or other small applications that require complete data redundancy.
RAID 5	Uses data striping and parity data across three or more drives (distributed parity) to provide high data throughput and data redundancy, especially for applications that require random access.
RAID 6	Uses data striping and parity data across three or more drives (distributed parity) to provide high data throughput and data redundancy, especially for applications that require random access. RAID 6 can survive the failure of two drives.
RAID 00	Is a spanned drive group that creates a striped set from a series of RAID 0 drive groups to provide high data throughput, especially for large files.

Table 19 RAID Levels (Continued)

Level	Description
RAID 10	A combination of RAID 0 and RAID 1 that uses data striping across two mirrored drive groups. It provides high data throughput and complete data redundancy.
RAID 50	A combination of RAID 0 and RAID 5 that uses data striping across two drive groups with parity data. It provides high data throughput and complete data redundancy.
RAID 60	A combination of RAID 0 and RAID 6 that uses data striping across two drive groups with parity data. It provides high data throughput and complete data redundancy. RAID 60 can survive the failure of two drives in each RAID set in the spanned drive group.

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

   The **Create New VD** screen appears.

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

Figure 28 Create a New Virtual Drive

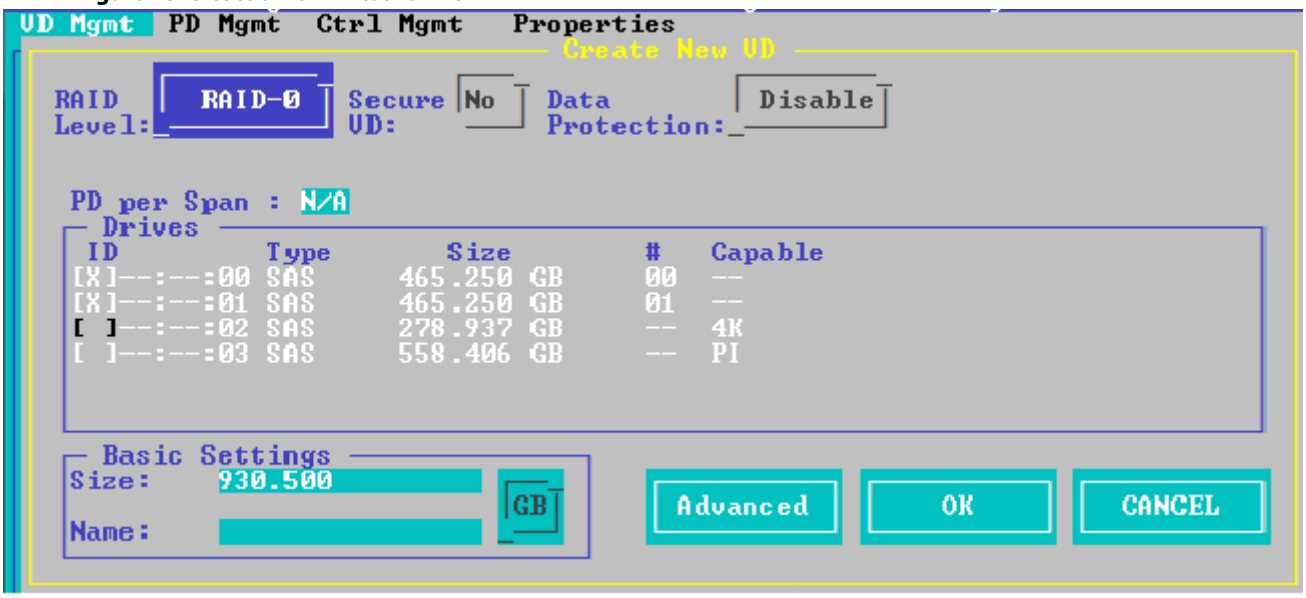

NOTE If your system detects any JBODs, the **Convert JBOD to Unconfigured Good** dialog (Figure 37 on page 63) appears before the **Create New VD** dialog. The **Convert JBOD to Unconfigured Good** dialog lets you convert the JBOD drives to Unconfigured Good, to then configure these drives as VDs.

3. Select a RAID level for the drive group from the **RAID Level** field. For more information, see Table 19, RAID Levels.
4. Select a power save mode for the drive group from the **Power save mode** field.

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

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

   Power Save (Dimmer Switch feature) is a technology that conserves energy by placing certain unused drives into a Power Save mode. In Power-Save mode, the drives use less energy. The fan and the enclosure require less energy

to cool and house the drives, respectively. Also, this technology helps avoid application time-outs caused by spin-up delays and drive wear caused by excessive spin-up/down cycles.

5. You can encrypt data and use drive-based key management for your data security solution.

   This option protects the data in the event of theft or loss of drives. Select a value from the **Secure VD** field. The options available are **Yes** and **No**.

6. You can choose whether you want to use the data protection feature on the newly created virtual drive.

   Select a value from the **Data Protection** field. The options available are **Yes** and **No**. The **Data Protection** field is enabled only if the controller has data protection physical drives connected to it.

   If you have enabled data protection while creating a virtual drive, you must select either **Full Initialization** or the **Background Initialization**, otherwise you cannot create a virtual drive.

   If there is no drive that is capable of protection, then the **Data Protection**field is grayed out. Also, if the controller does not support Virtual Drive Protection, then the **Data Protection** field will be suppressed.

7. You can change the sequence of the physical drives in the **Drives** box.

   All the available unconfigured good drives appear in the **Drives** box. Select the physical drives in the sequence that you prefer. Based on your selection, the sequence number appears in the **#** column. The **Type** column shows the drive type; for example, SAS, SATA, IDE, and so on. The **Capable** column shows the capability of the drive.

8. You can select a size lesser than the maximum size of the drive group, if you want to create other virtual drives on the same drive group.

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

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

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

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

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

10. Press **OK**.

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

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

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

12. Press **Advanced**.

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

Figure 29 Create Virtual Drive – Advanced

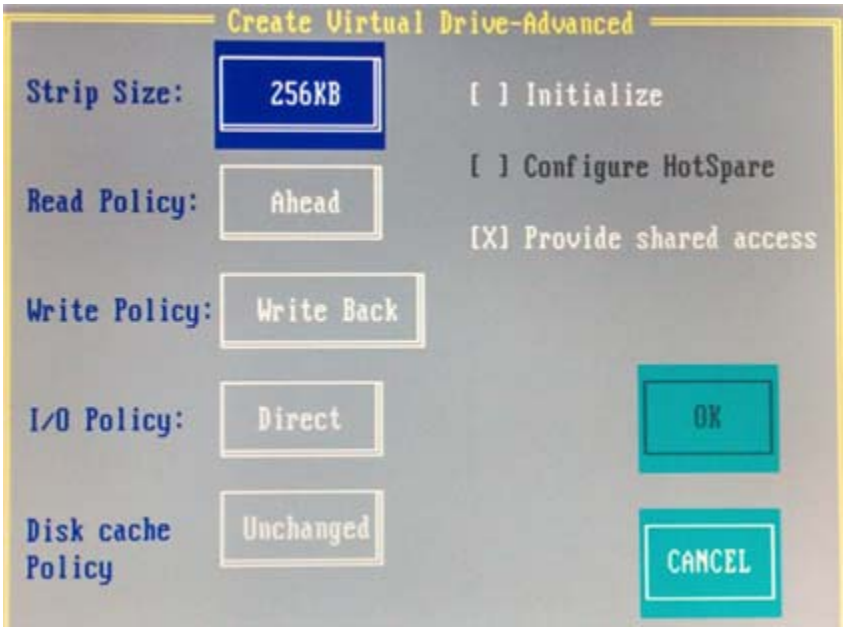

NOTE The **Provide shared access** check box appears only if the controller supports High Availability DAS.

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

    The new virtual drive is created and initialized.

4.7.1 Selecting Additional Virtual Drive Properties

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

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

  NOTE The Integrated RAID controller supports only the 64 KB stripe size.

- **Read Policy** – A virtual drive property that indicates whether the default read policy is **Always Read Ahead** or **No Read Ahead**:
  - **Always Read Ahead** – Permits the controller to read sequentially ahead of the requested data and allows the controller to store the additional data in the cache memory. Here, the controller anticipates that the data is required frequently. Even though Always Read Ahead policy speeds up the reads for sequential data, but little improvement is seen when accessing the random data.
  - **No Read Ahead** – Disables the Always Read Ahead capability of the controller.
- **Write Policy** – Select one of the following options to specify the write policy for this virtual drive:
  - **Write Back** – In this mode, the controller sends a data transfer completion signal to the host when the controller cache receives all of the data in a transaction. If you select the Write Back policy and the battery is absent, the firmware disables the Write Back policy and defaults to the Write Through policy.
  - **Write Through** – In this mode, the controller sends a data transfer completion signal to the host when the drive subsystem receives all the data in a transaction.

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

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

4.8 Clearing the Configuration

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

Perform the following steps to clear configuration:

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

   The following dialog appears.

Figure 30 Clear Configuration

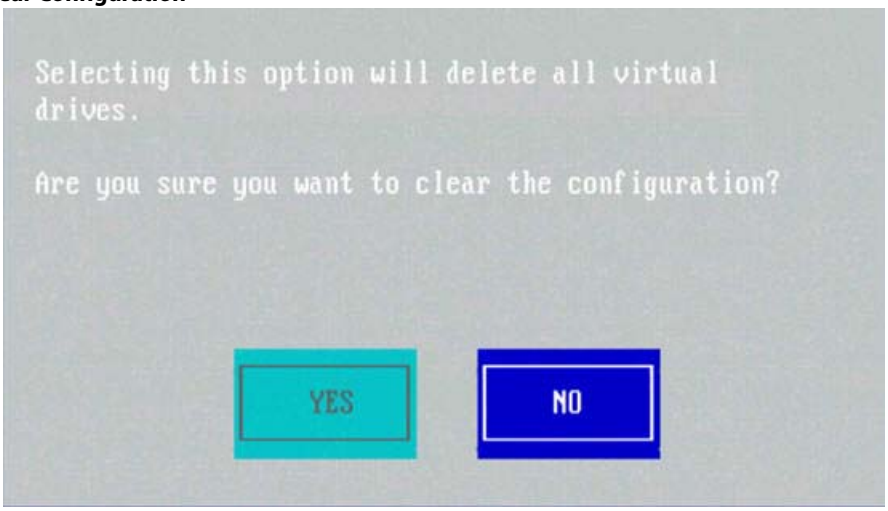

3. Press **Yes** to delete all the virtual drives.

4.9 Broadcom SafeStore Encryption Services

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

4.9.1 Enabling Drive Security

This section describes how to enable, change, and disable the drive security, and how to import a foreign configuration by using the SafeStore Encryption Services advanced software.

To enable security on the drives, you need to perform the following actions to set drive security:

- Enter a security key identifier.

  A security key identifier appears whenever you have to enter a security key.

- Enter a security key.

  After you create a security key, you can create secure virtual drives by using the key. You must use the security key to perform certain operations.

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

Perform the following steps to enable drive security.

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

   The **Create Security Key** dialog appears.

Figure 31 Create Security Key

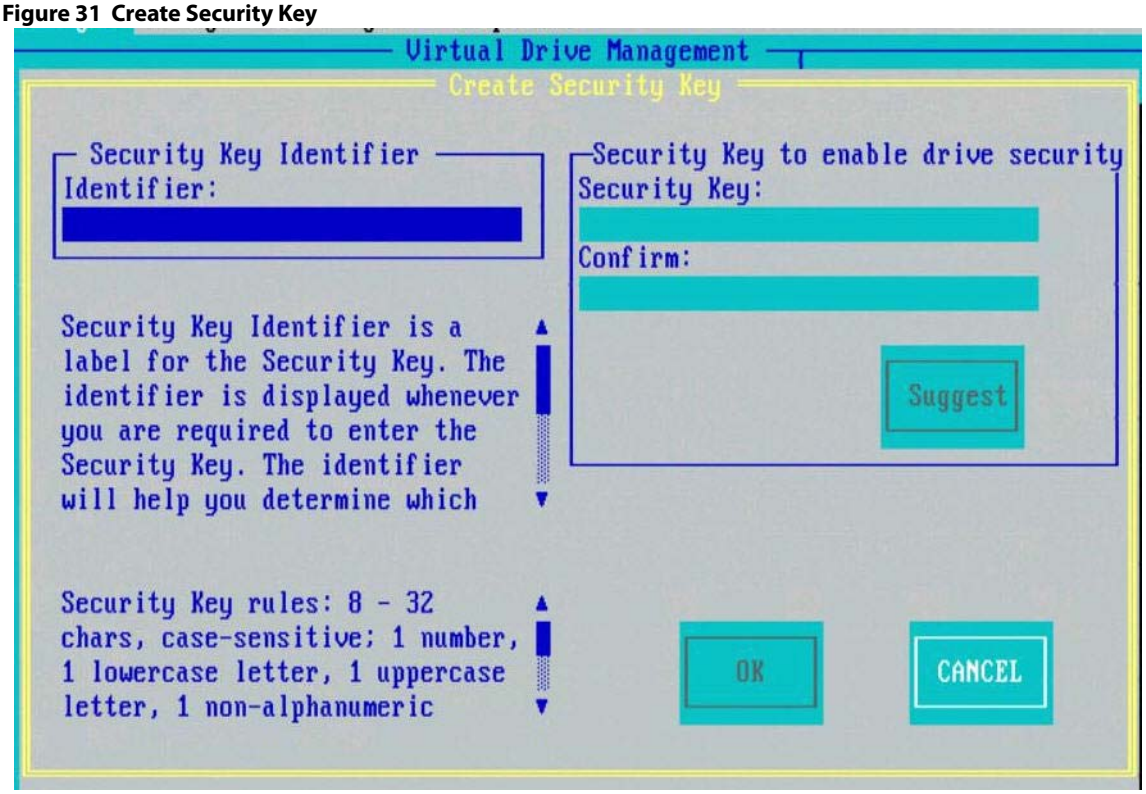

4. Either use the default security key identifier, or enter a new security key identifier.

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

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

   ATTENTION **If you forget the security key, you lose access to your data.** Be sure to record your security key information. You might need to enter the security key to perform certain operations.

   The security key is case sensitive. It must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non alphanumeric character (a symbol, for example, < > @ +). The space character is not permitted.

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

4.9.2 Changing Security Settings

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

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

   The **Change Security Key** dialog appears.

Figure 32 Change Security Key

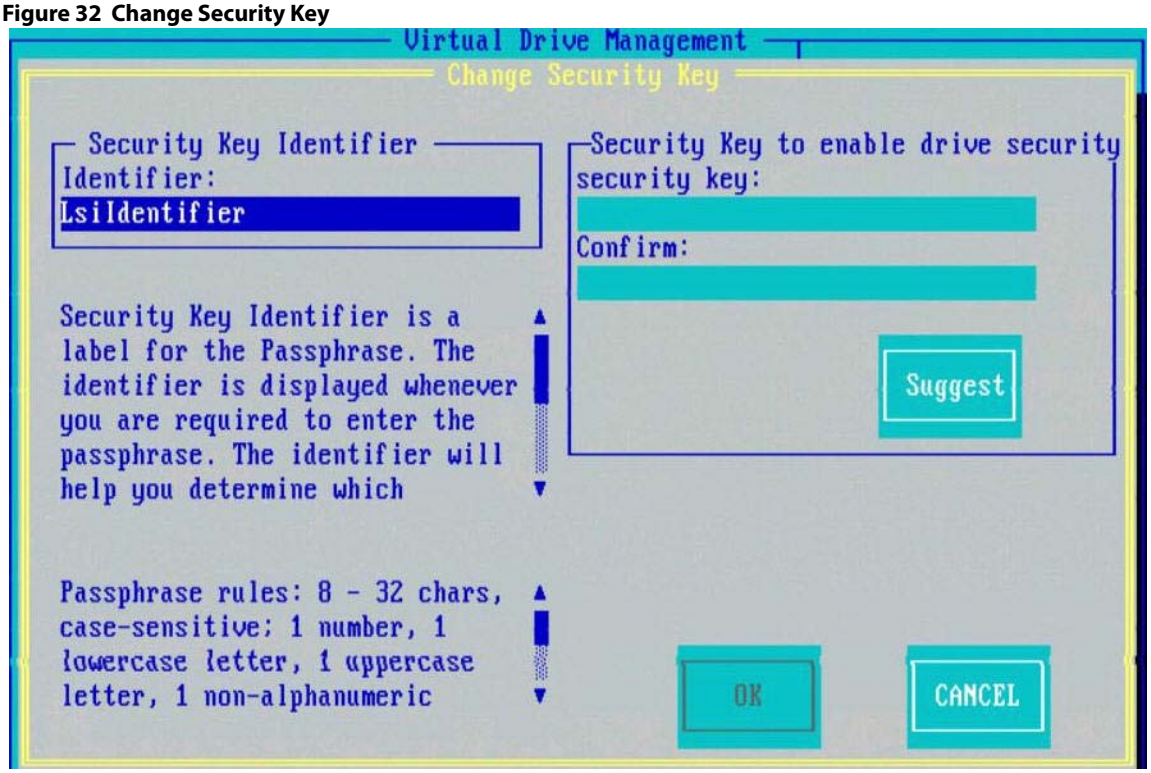

4. Either keep the existing security key identifier, or enter a new security key identifier.

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

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

   ATTENTION **If you forget the security key, you lose access to your data.** Be sure to record your security key information. You might need to enter the security key to perform certain operations.

   The security key is case sensitive. It must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non alphanumeric character (for example, < > @ +). The space character is not permitted.

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

4.9.3 Disabling Drive Security

If you disable drive security, your existing data is not secure and you cannot create any new secure virtual drives. Disabling drive security does not affect data security on foreign drives. If you removed any drives that were previously secured, you still need to enter the password when you import them. Otherwise, you cannot access the data on those drives.

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

Madurai Kamaraj University,

s on all the secure drive groups.

Perform the following steps to disable drive security:

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

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

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

4.9.4 Importing or Clearing a Foreign Configuration

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

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

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

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

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

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

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

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

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

   The foreign configuration options **Import** and **Clear** appear.

Figure 33 Foreign Configuration Options

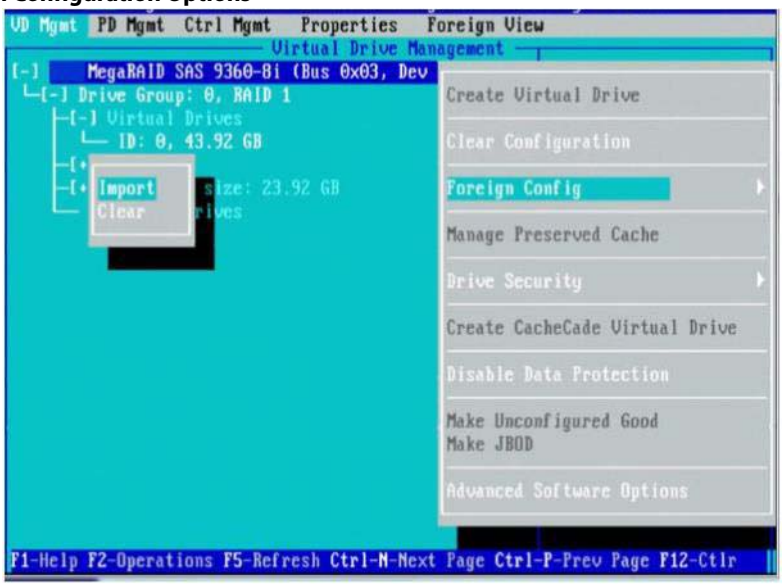

3. Navigate to the command you want to perform.
   - To import a foreign configuration, go to step 4.
   - To clear a foreign configuration, go to step 6.
4. To import a foreign configuration, select **Import**, and press **Enter**.

   The following dialog appears.

Figure 34 Foreign Configuration – Import

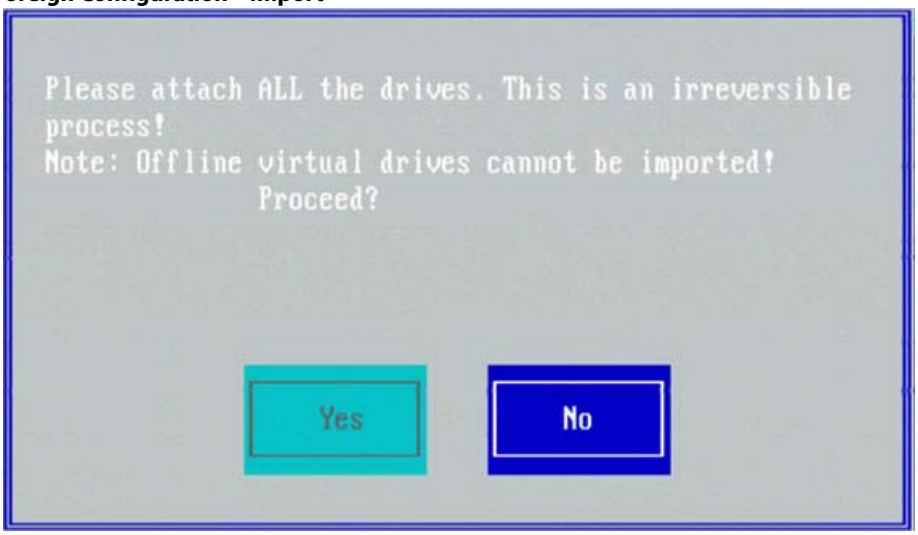

5. Press **Yes** to import the foreign configuration from all the foreign drives. Repeat the import process for any remaining drives.

   Because locked drives can use different security keys, you must verify whether there are any remaining drives to be imported.

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

existing configurations, you first must clear the configuration on those drives.

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

   The following dialog appears.

Figure 35 Foreign Configuration – Clear

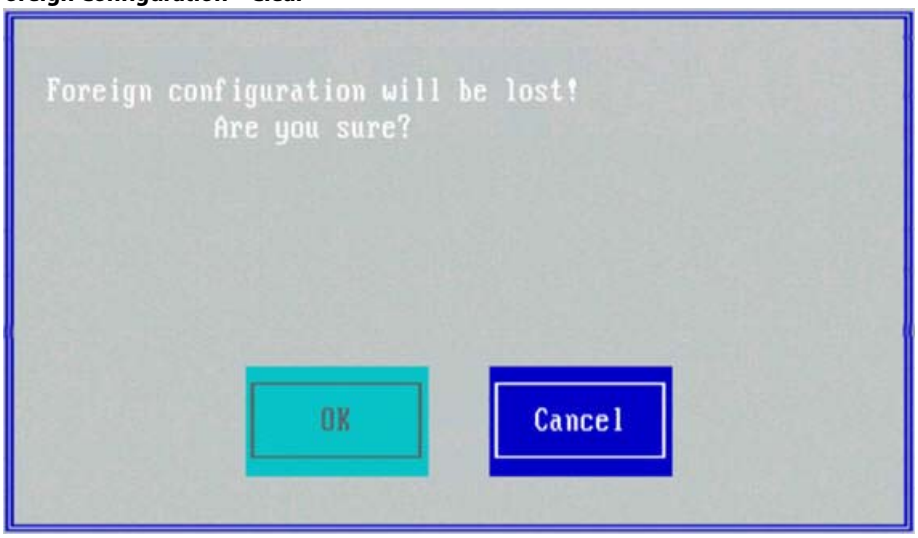

7. Press **OK** to clear a foreign configuration.

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

4.9.4.1 Foreign Configurations in Cable Pull and Drive Removal Scenarios

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

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

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

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

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

  NOTE Start a consistency check immediately after the rebuild is complete, to ensure data integrity for the virtual drives.

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

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

  NOTE Start a consistency check immediately after the rebuild is complete, to ensure data integrity for the virtual drives.

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

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

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

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

4.10 Discarding Preserved Cache

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

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

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

Perform the following steps to discard the preserved cache:

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

   The **Manage Preserved Cache** dialog appears.

Figure 36 Manage Preserved Cache

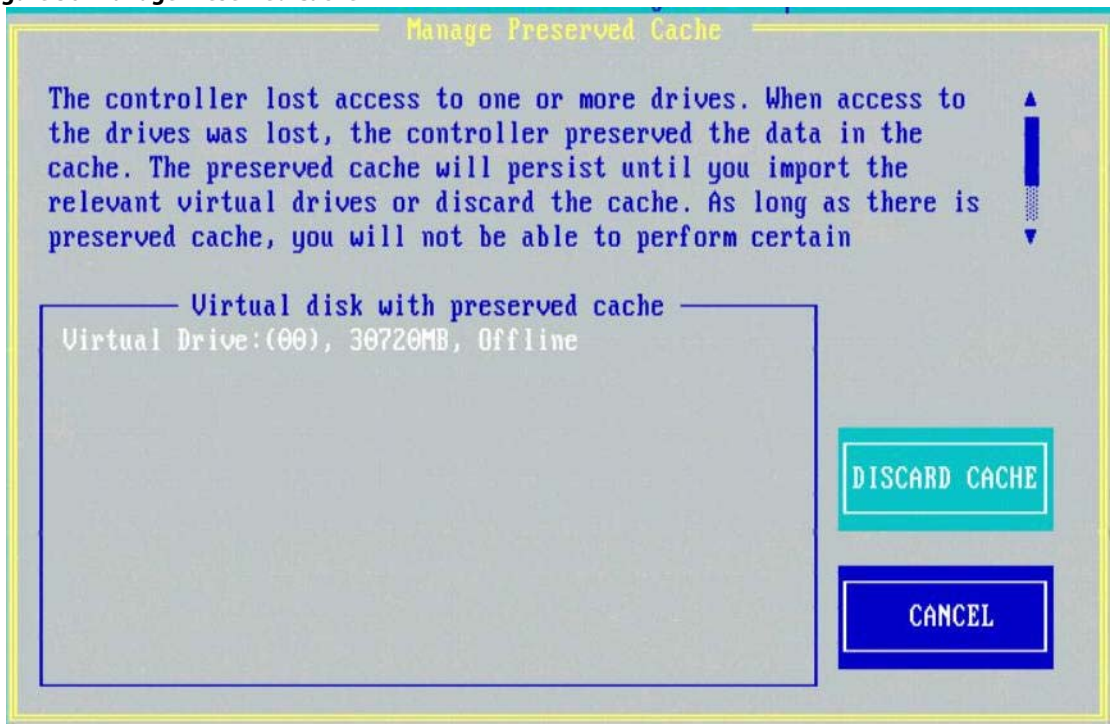

3. Click **Discard Cache** to discard the preserved cache from the virtual drive.
   A message box appears, asking you to confirm your choice.
4. Click **OK** to continue.

4.11 Converting JBOD Drives to Unconfigured Good Drives

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

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

1. In the **VD Mgmt** screen, navigate to the controller, and press the **F2** key.
2. Navigate to **Make Unconfigured Good**, and press **Enter**.
   The **Convert JBOD to Unconfigured Good** dialog appears, which shows all JBODs available in the system.

Figure 37 Convert JBOD to Unconfigured Good

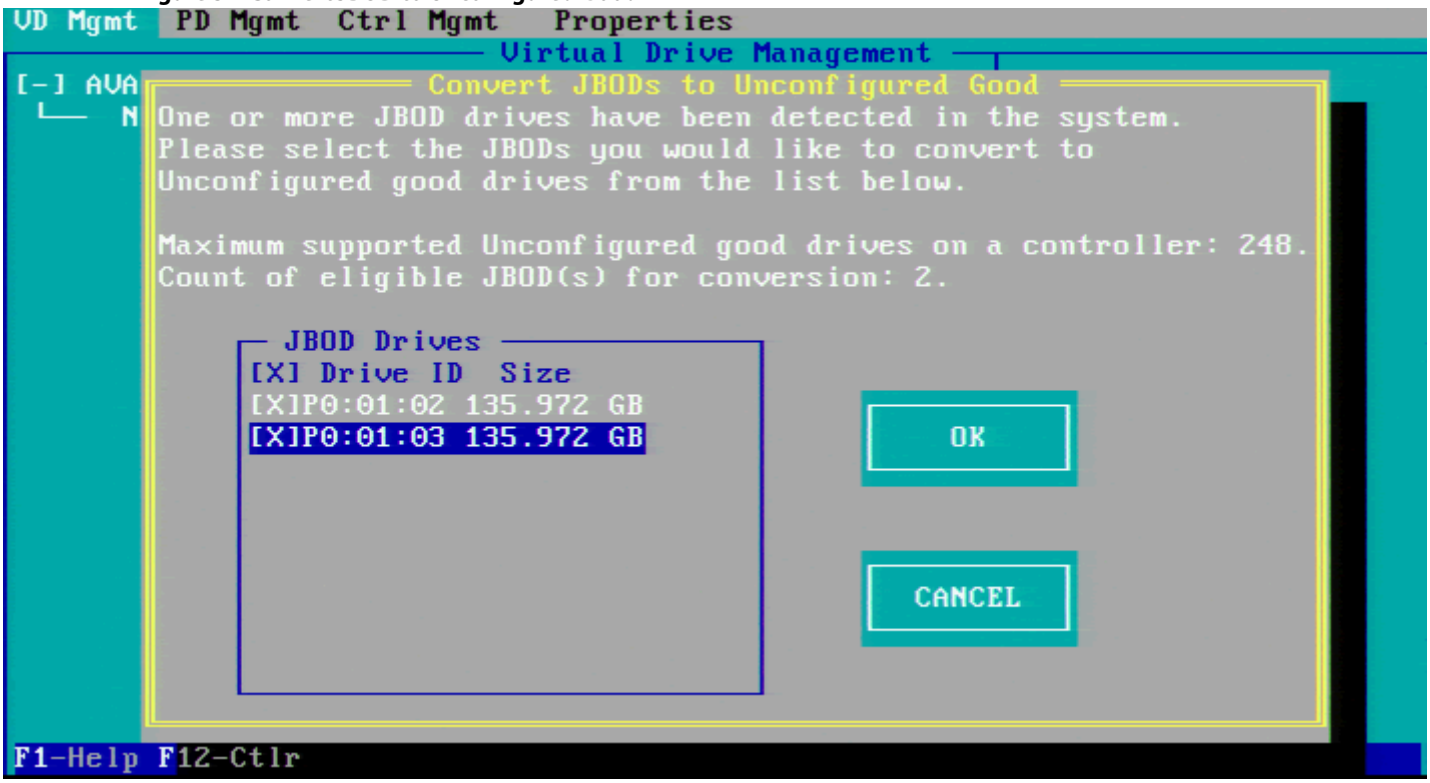

3. Select the JBODs that you want configured as Unconfigured Good drives.
   To select or deselect all the JBODs at one go, select the top most square brackets in the **JBOD Drives** box.

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

4. Click **OK**.

   The selected JBODS are converted to Unconfigured Good drives.

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

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

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

4.12 Converting Unconfigured Good Drives to JBOD Drives

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

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

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

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

Figure 38 Convert Unconfigured Good to JBOD

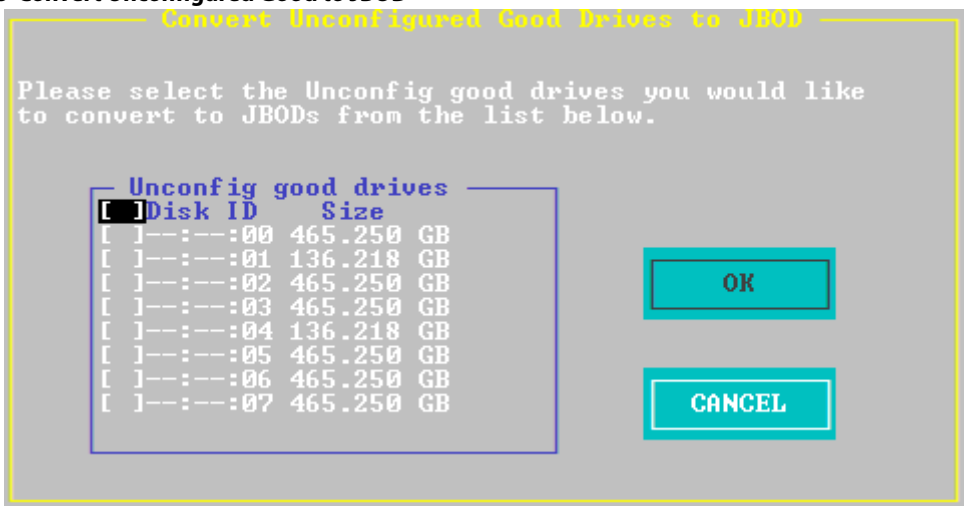

3. Select the Unconfigured Good drives that you want configured as JBODs.

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

   The selected Unconfigured Good drives are converted to JBOD drives.

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

1. In the **Drive Management** screen, navigate to a Unconfigured Good drive, and press the **F2** key.

2. Navigate to **Make JBOD**, and press **Enter**.
3. Click **OK** in the message confirmation box to continue.

4.13 Enabling Security on a JBOD

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

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

This feature is dependent on the mode that you have selected. On JBOD mode, enabling security on JBOD is not available.

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

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

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

Figure 39 Enable Security on JBOD

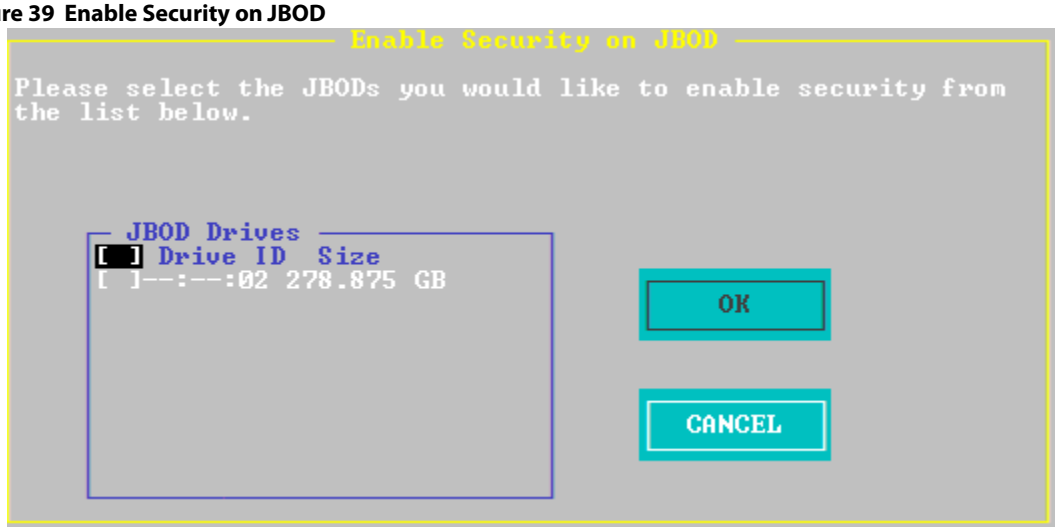

3. Select the JBOD drives for which you want to enable security.

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

   The security is enabled on all of the selected JBOD drives.

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

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

4.14 Viewing and Changing Device Properties

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

4.14.1 Viewing Controller Properties

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

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

The information in the **Properties** screen (Figure 17 on page 45) is read only. Most of this information is self-explanatory. To view additional properties, navigate to **Next**, and press **Enter**.

4.14.2 Modifying Controller Properties

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

Perform the following steps to change the controller properties:

1. Navigate to the **Ctrl Mgmt** menu to view the first **Controller Settings** screen.
2. You can change the values of the properties for the editable fields.

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

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

Table 20 Controller Settings

Options	Descriptions
Alarm Control	Select this option to enable, disable, or silence the onboard alarm tone generator on the controller.
Coercion Mode	Use this option to force drives of varying capacities to the same size so they can be used in a drive group. The coercion mode options are None, 128MB-way, and 1GB-way. The number you choose depends on how much the drives from various vendors vary in their actual size.
BIOS Mode	Specifies the following options to set the BIOS boot mode: ■ **Stop on Error**: Shows the errors encountered during boot up and waits for your input. The firmware does not proceed with the boot process until you take some action. ■ **Pause on Error**: The firmware might halt because of hardware faults. If the firmware encounters no hardware faults, the boot up continues. ■ **Ignore Error**: Ignores errors and the firmware proceeds with boot. ■ **SafeMode Error**: Boots the controller to run on safe mode.
Boot Device	Use this option to select the boot device from the list of virtual drives and JBODs. This property is applicable only for legacy BIOS.
Rebuild Rate	Use this option to select the rebuild rate for drives connected to the selected controller. The rebuild rate is the percentage of system resources dedicated to rebuilding a failed drive. The higher the number, the more system resources are devoted to a rebuild. The range of rebuild rate is between 0 and 100 percent.

Table 20 Controller Settings (Continued)

Options	Descriptions
BGI Rate	Use this option to select the amount of system resources dedicated to background initialization of virtual drives connected to the selected controller. The range of background initialization (BGI) rate is between 0 and 100 percent.
CC Rate	Use this option to select the amount of system resources dedicated to consistency checks of virtual drives connected to the selected controller. The range of Consistency Check (CC) rate is between 0 and100 percent.
Recon. Rate	Use this option to select the amount of system resources dedicated to reconstruction of drives connected to the selected controller. The range of Recon rate is between 0 and100 percent.
Patrol Rate	Use this option to select the rate for patrol reads for drives connected to the selected controller. The patrol read rate is the percentage of system resources dedicated to running a patrol read. The range of patrol read is between 0 to 100 percent.
Cache Flush Interval	Use this option to control the interval at which the contents of the onboard data cache are flushed. The range of Cache Flush Interval is between 0 to100 seconds.
Spinup Delay	Use this option to control the interval (in seconds) between the spin-up of drives connected to this controller. The delay prevents a drain on the system's power supply that would occur if all drives spun up at the same time. The range of the Spinup Delay is between 0 to 255 seconds.
Spinup Drive	Use this option to control the interval at which the contents of the onboard data cache are flushed. The range of Spinup Drive is between 0 to 255 seconds.
Maintain PD Fail History	Use this option to maintain the history of all drive failures.
Device Exposure	Displays the actual number of devices to be exposed to the host. You can assign the following values: ■ 0 and 1 = Exposes all drives to the host ■ 2 to 255 = The actual number of devices to be exposed. For example, 4 = 4 devices, 10 = 10 devices exposed, 100 = 100 devices exposed and so on.
Enable Controller BIOS	Use this option to enable or disable the BIOS for the selected controller. If the boot device is on the selected controller, the BIOS must be enabled. Otherwise, the BIOS should be disabled, or you might be unable to use a boot device elsewhere.
Enable Stop CC on Error	Use this option to stop a consistency check when the controller BIOS encounters an error.
Auto Enhanced Import	Use this option to import automatically at boot time.
Set Factory Defaults	Use this option to load the default Ctrl-R Utility settings.
Manage Link Speed	Use this option to change the link speed between the controller and the expander, or between a controller and a drive that is directly connected to the controller.
Manage Power Save	Use this option to reduce the power consumption of drives that are not in use, by spinning down the unconfigured good drives, hot spares, and configured drives.
Start Manual Learn Cycle	The manual learn cycle re-calibrates the battery integrated circuit so that the controller can determine whether the battery can maintain the controller cache for the prescribed period of time in the event of a power loss.
Manage Battery	Use this option to view information about the BBU, if the selected controller has a BBU.
Emergency Spare	Use this option to commission unconfigured good drives or global hot spares as emergency spare drives. You can select from the options **None**, **UG** (Unconfigured Good), **GHS** (Global Hot spare), or **UG and GHS** (Unconfigured Good and Global Hot spare).
Enable Emergency for SMARTer	Use this option to commission emergency hot spare drives for predictive failure analysis events.
Write Verify	Use this option to verify if the data was written correctly to the cache before flushing the controller cache.

Table 20 Controller Settings (Continued)

Options	Descriptions
Large I/O Support	Use this option to enable or disable large I/O support feature. By default, large I/O support is disabled. A reboot is required if this property is changed. When this property is changed, `The controller property change has been performed successfully. Reboot the machine for the change to take effect` message is displayed.
Detection Type	Drives tend to develop media errors over time, which can slow down performance of the drive as well as the system as a whole. The firmware attempts to detect drives that consistently perform poorly. The available options are **High Latency**, **Aggressive**, and **Default**. Depending on your requirements, use these options to set appropriate controller properties.
Drive Error Threshold	Use these options to set appropriate controller properties. The available options follow: ■ **Every 8 hours**. ■ **Every 1 hour**. ■ **Every 15 minutes**. ■ **Every 5 minutes**.
Drive Corrective Action	Drives tend to develop media errors over time, which can slow down the performance of the drive as well as the system as a whole. If a drive has certain amount of affected media leading to consistently poor I/O latency, then the firmware fails that particular drive, so that the drive rebuild/copyback process can start on that drive. The firmware also logs the appropriate events to alert the user. You can either enable or disable this option.

4.14.3 Viewing and Changing Virtual Drive Properties

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

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

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

   The **Virtual Drive Properties** dialog appears.

Figure 40 Virtual Drive Properties

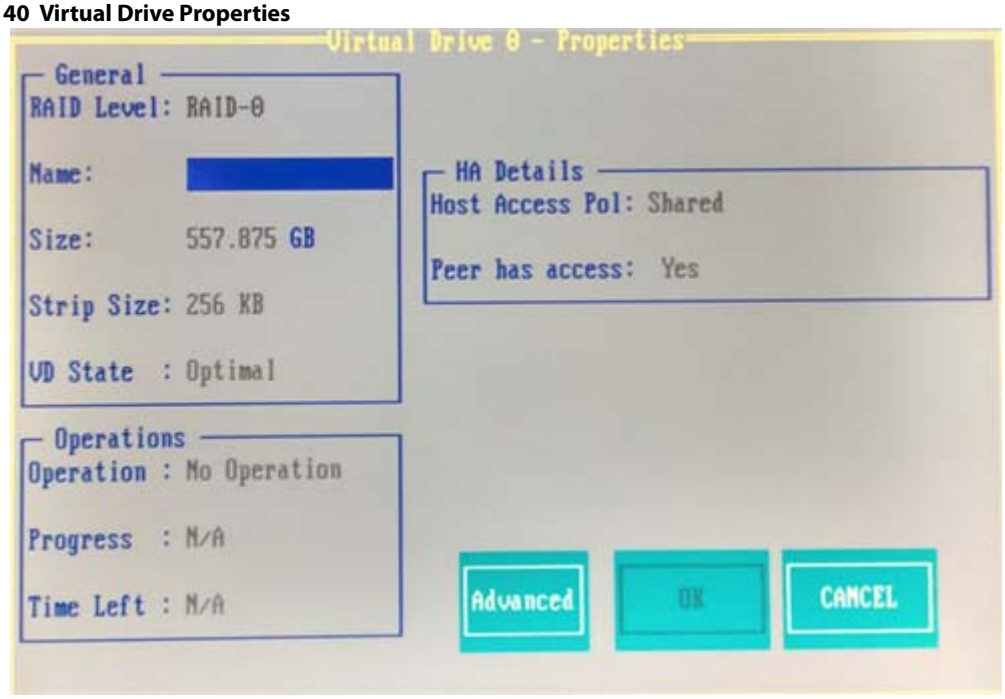

The **General** box shows the virtual drive's RAID level, name, state, size, and strip size.

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

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

- **Host access policy**

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

- **Peer has access**

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

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

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

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

   The **Advanced Properties** dialog appears.

Figure 41 Virtual Drive Management – Advanced Properties

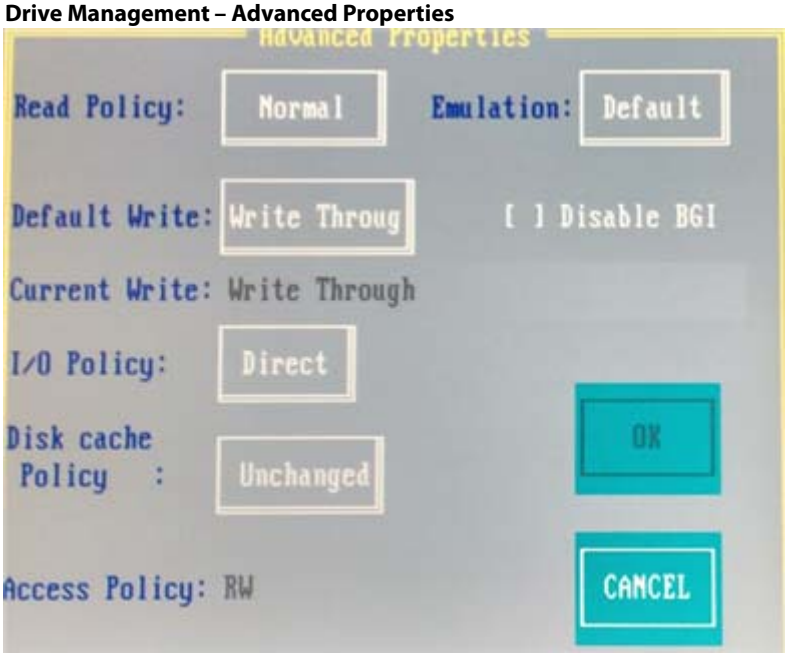

You can view the virtual drive policies that were defined when the storage configuration was created.

4.14.4 Deleting a Virtual Drive

You can delete any virtual drive on the controller if you want to reuse that space for a new virtual drive. The Ctrl-R Utility lists configurable drive groups where there is space to configure. If multiple virtual drives are defined on a single drive group, you can delete a virtual drive without deleting the entire drive group.

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

Perform the following steps to delete a virtual drive:

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

4.14.5 Deleting a Virtual Drive Group

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

Perform the following steps to delete a drive group:

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

4.14.6 Expanding a Virtual Drive

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

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

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

   The **Expand Virtual Drive** dialog appears.

Figure 42 Expand Virtual Drive

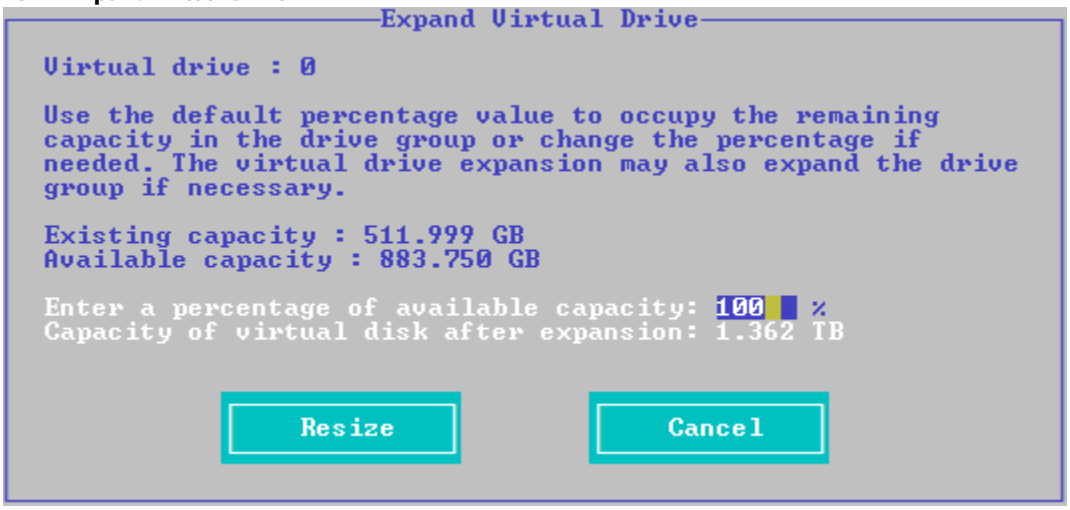

3. Enter the percentage of the available capacity that you want the virtual drive to use.

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

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

4.14.7 Erasing a Virtual Drive

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

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

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

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

   A menu appears displaying the following modes:

   — **Simple**

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

   — **Normal**

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

   — **Thorough**

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

   — **Stop Erase**

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

3. Select a mode and press **Enter**.

   A message box appears.

Figure 43 Erase Virtual Drive

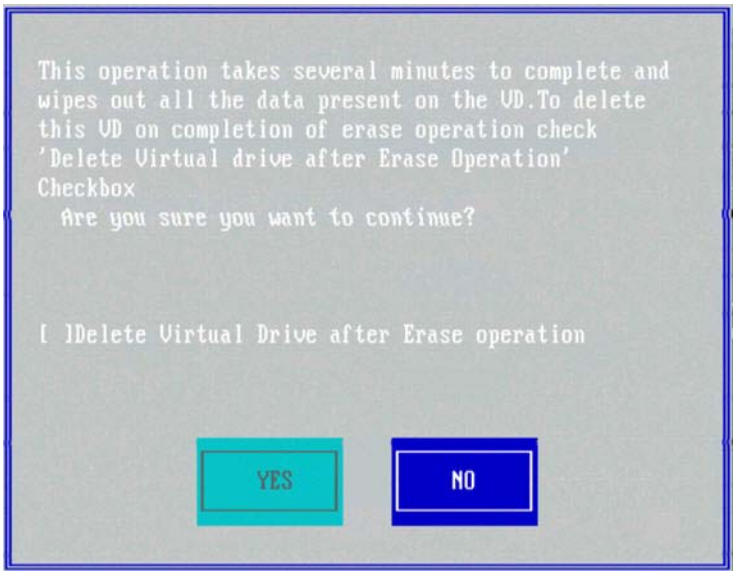

4. To delete the virtual drive after the erase operation has been completed, select the **Delete Virtual Drive after Erase operation** check box.
5. Click **Yes** for the erase operation to start.

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

4.14.8 Managing Link Speed

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

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

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

Perform the following steps to change the link speed:

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

Figure 44 Manage Link Speed

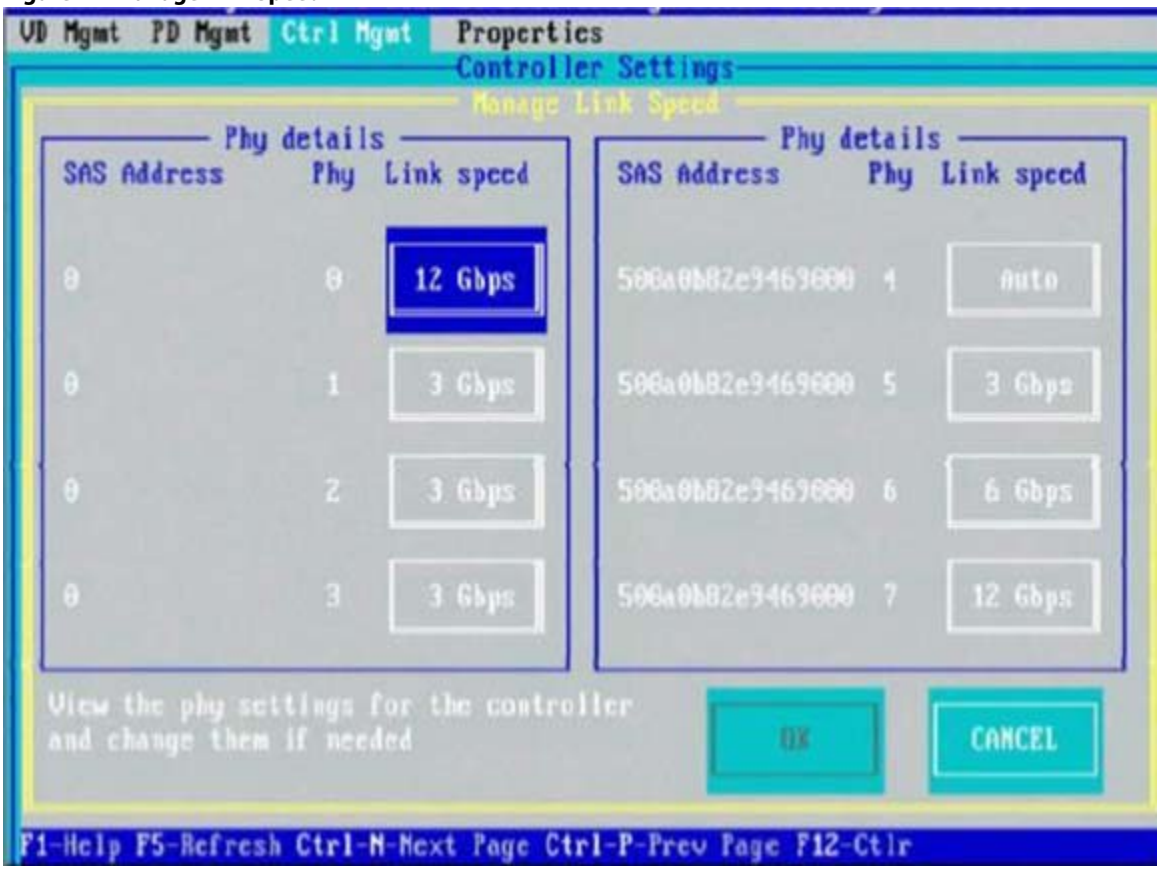

   — The **SAS Address** column shows the SAS address that uniquely identifies a device in the SAS domain.
   — The **Phy** column shows the system-supported phy link values. The phy link values are from 0 through 7.
   — The **Link Speed** column shows the phy link speeds.
3. Select the desired link speed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

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

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

4.14.9 Managing Power Save Settings for the Controller

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

Perform the following steps to change the power save settings:

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

Figure 45 Manage Power Save

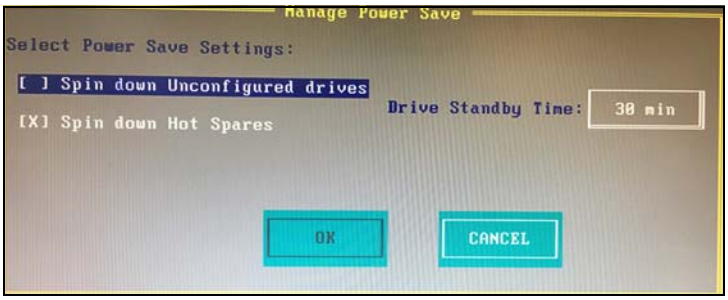

3. Select the **Spin down Unconfigured drives** check box to let the controller enable the unconfigured drives to enter the Power-Save mode.
4. Select the **Spin down Hot Spares** check box to let the controller enable the Hot spare drives to enter the Power-Save mode.
5. Select the drive standby time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

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

6. Click **OK**.
7. Click **Yes** to save the settings.

4.14.10 Start Manual Learn Cycle

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

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

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

Figure 46 Controller Settings – Second Screen

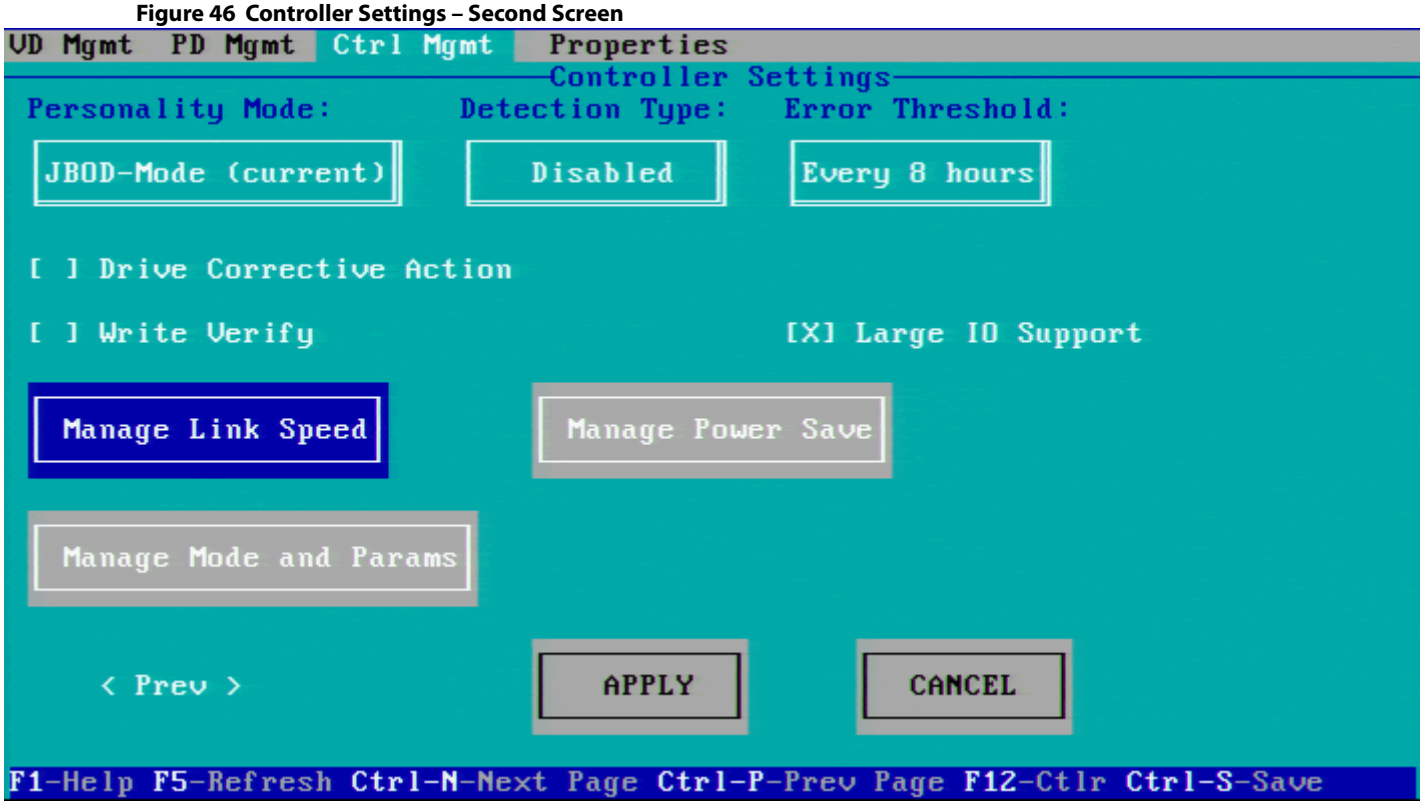

2. Click **Start Manual Learn Cycle**.
   An information box appears stating that the battery learn cycle will take a few minutes.

Figure 47 Manual Learn Cycle Warning

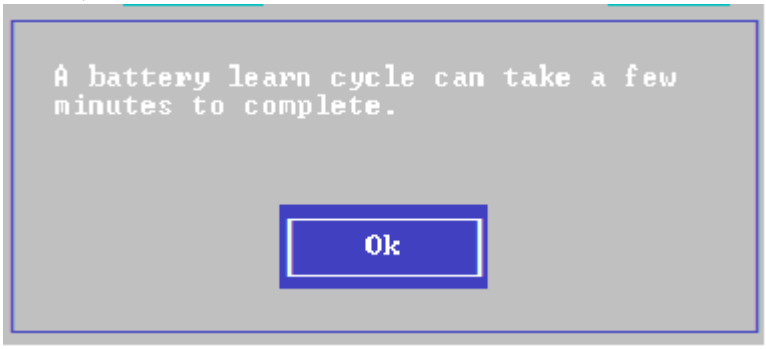

3. Click **Ok** to continue.

4.14.11 Managing Power Save Settings for the Drive Group

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

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

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

Figure 48 Manage Power Save Settings – Drive Group

3. Select a power save mode from the **Select power save mode** drop-down list.

   A description of the selected mode appears in the dialog.
4. Click **OK**.

4.14.12 Managing Dedicated Hot Spares

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

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

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

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

Figure 49 Dedicated Hotspare

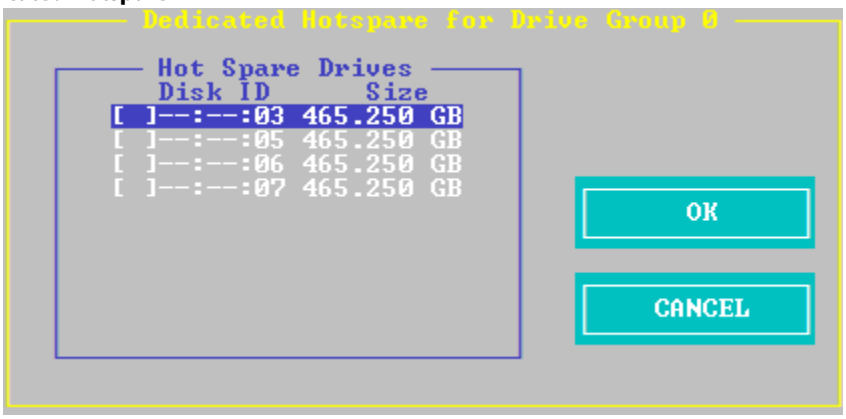

3. Perform one of these steps:
   — To create a dedicated hot spare, select a drive and click **OK**.
   — To delete a dedicated hot spare, deselect the hot spare and click **OK**.

4.14.13 Securing a Drive Group

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

Perform the following steps to secure a drive group:

1. Navigate to the **VD Mgmt** screen, navigate to the drive group that you want to secure, and press the **F2** key.
2. Navigate to **Secure Drive Group**, and press **Enter**.
   A message box appears asking for your confirmation.
3. Click **Yes** to secure the drive group.

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

4.14.14 Setting LED Blinking

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

Perform the following steps to start or stop LED blinking:

1. Navigate to the **Drive Management** screen (in the **PD Mgmt** menu).
2. Select a physical drive, and press the **F2** key.
3. Navigate to **Locate**, and press **Enter**.
   The **Start** and the **Stop** options appear.
4. Perform one of these actions:
   — Select **Start**, and press **Enter** to start LED blinking.
   — Select **Stop**, and press Enter to stop LED blinking.

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

4.14.15 Performing a Break Mirror Operation

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

Perform the following steps to perform a break mirror operation:

1. Navigate to the **VD Mgmt** screen, navigate to a drive group on which you want to perform the break mirror operation, and press the **F2** key.
2. Navigate to **Break Mirror**, and press **Enter**.
   The following message box appears, asking for your confirmation.

Figure 50 Break Mirror

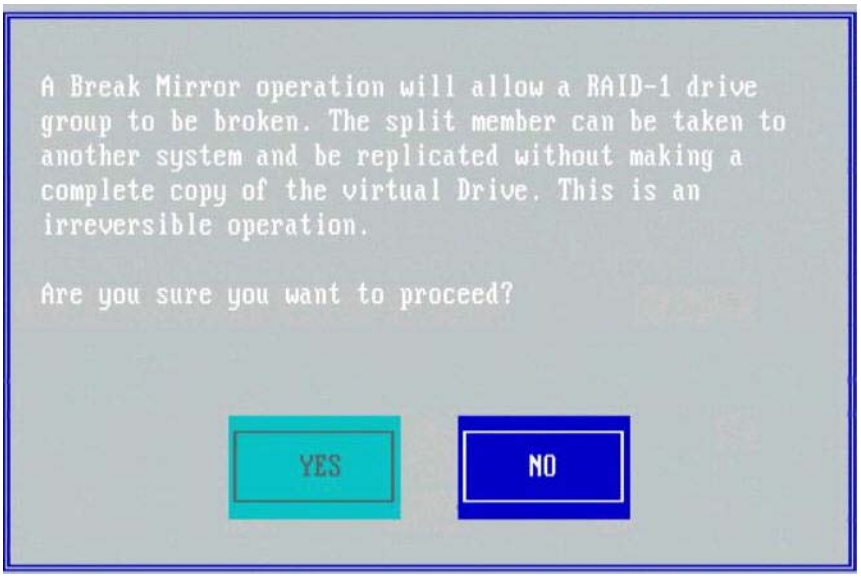

3. Click **Yes** to proceed.

4.14.16 Performing a Join Mirror Operation

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

Perform the following steps to perform a join mirror operation:

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

   The following dialog appears.

Figure 51 Join Mirror – Choose Option

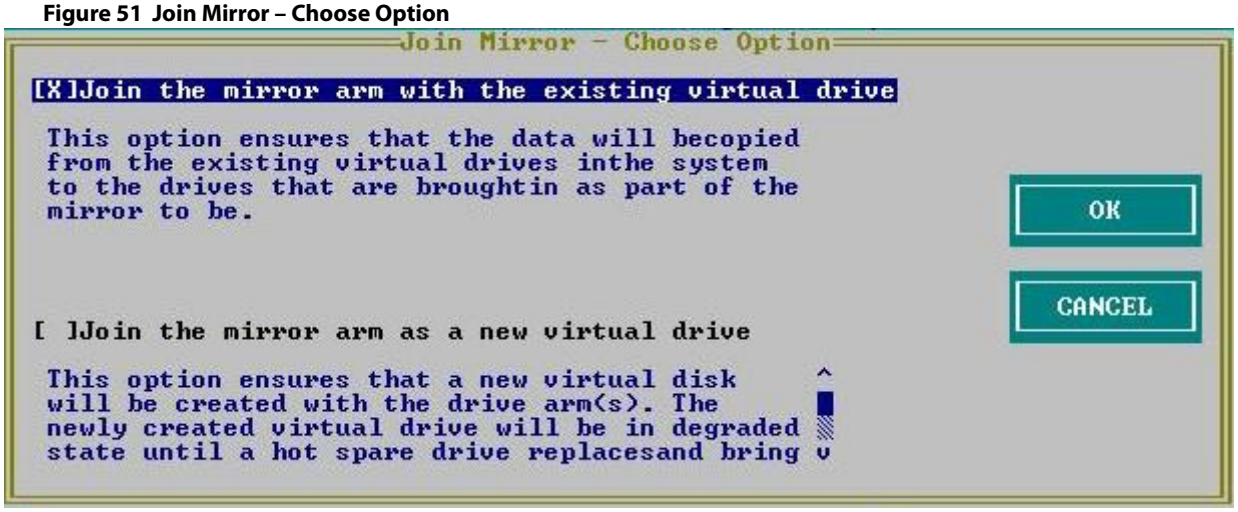

3. Select one of the options and click **OK**.
   — If you select **Join the mirror arm with the existing virtual drive**, the following confirmation dialog appears.

Figure 52 Confirmation Message

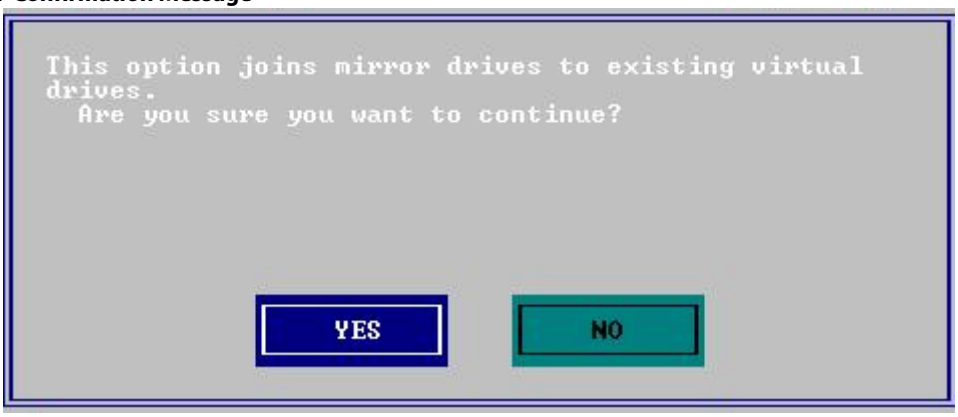

   — If you select **Join the mirror arm as a new virtual drive**, the following confirmation dialog appears.

Figure 53 Confirmation Message

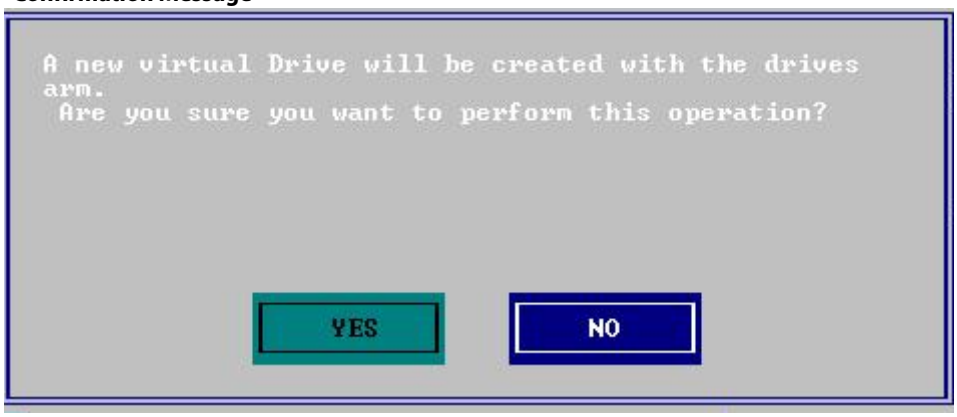

4. Click **Yes** to proceed.
   The following dialog appears.

Figure 54 Join Mirror – Choose Option

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

4.14.17 Hiding a Virtual Drive

You can hide a virtual drive on the controller.

Perform the following steps to hide a virtual drive:

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

4.14.18 Unhiding a Virtual Drive

You can unhide a virtual drive on the controller.

Perform the following steps to unhide a virtual drive:

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

4.14.19 Hiding a Drive Group

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

Perform the following steps to hide a drive group:

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

4.14.20 Unhiding a Drive Group

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

Perform the following steps to unhide a drive group:

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

4.15 Managing Storage Configurations

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

4.15.1 Initializing a Virtual Drive

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

Perform the following steps to initialize a virtual drive:

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

Figure 55 Initialize a Virtual Drive

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

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

4.15.2 Running a Consistency Check

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

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

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

Perform the following steps to run a consistency check:

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

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

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

Figure 56 Consistency Check

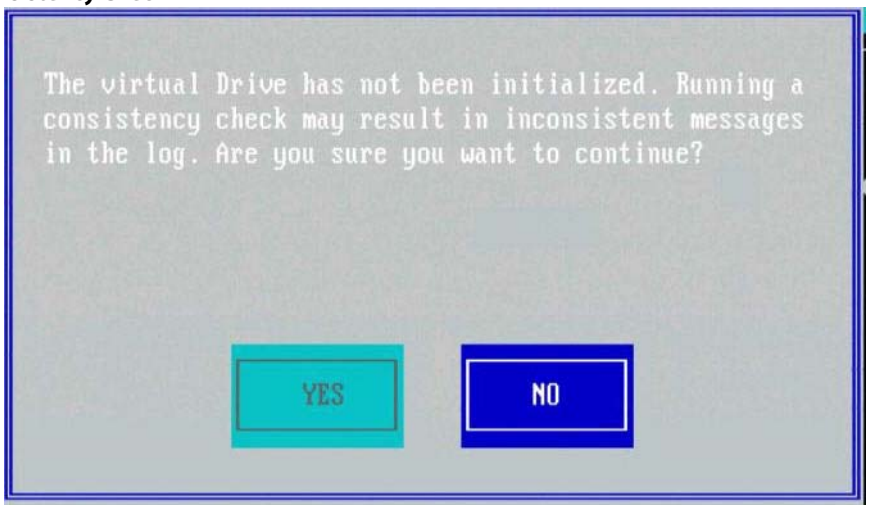

4. Click **Yes** to run the consistency check.

4.15.3 Rebuilding a Physical Drive

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

Perform the following steps to rebuild a physical drive:

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

   The rebuild operation starts.

4.15.4 Performing a Copyback Operation

You can perform a copyback operation on a selected drive.

The copyback operation copies data from a source drive of a virtual drive to a destination drive that is not a part of the virtual drive. The copyback operation often creates or restores a specific physical configuration for a drive group (for example, a specific arrangement of drive group members on the device I/O buses).

Perform the following steps to perform the copyback operation:

1. Navigate to the **Drive Management** screen, navigate to a physical drive, and press the **F2** key.

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

Figure 57 Copyback Operation

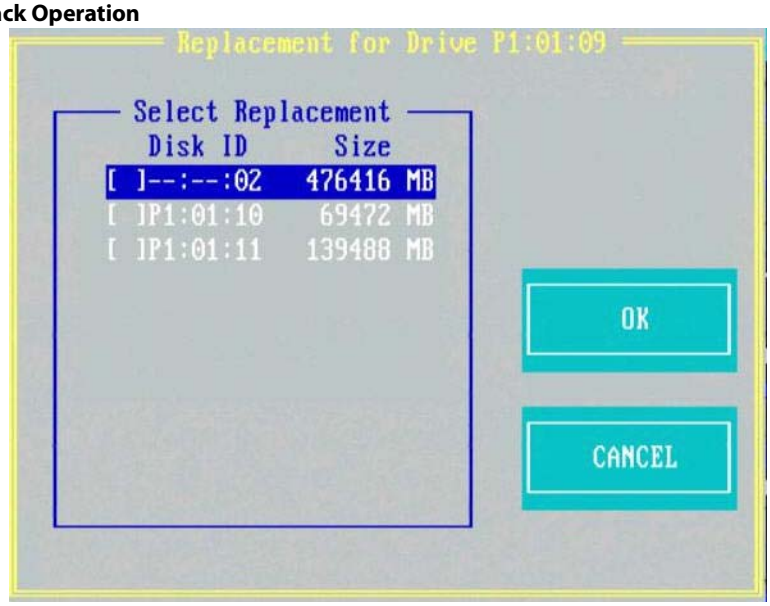

3. Select the replacement drive to which you want the data copied.
4. Click **OK**.
   The copyback operation is performed on the selected drive.

4.15.5 Removing a Physical Drive

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

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

1. Navigate to the **Drive Management** screen, and press the **F2** key.
2. Select **Prepare for Removal**, and press **Enter**.
   The physical drive is now in a power save mode.
   If you change your mind and do not want to remove the drive, navigate to **Undo Removal**, and press **Enter**.

4.15.6 Creating Global Hot Spares

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

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

Perform the following steps to create global hot spares:

1. Navigate to the **Drive Management** screen, navigate to a physical drive that you want to change to a hot spare, and press the **F2** key.
2. Select **Make Global HS**, and press **Enter**.
   The physical drive is changed to a global hot spare. The status of the physical drive as a global hot spare appears in the **Drive Management** screen.

4.15.7 Removing a Hot Spare Drive

Perform these steps to remove a hot spare drive:

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

4.15.8 Making a Drive Offline

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

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

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

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

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

4.15.9 Making a Drive Online

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

Perform the following steps to make a physical drive online:

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

4.15.10 Instant Secure Erase

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

Perform the following steps to erase data on SED drives:

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

4.15.11 Erasing a Physical Drive

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

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

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

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

   A menu appears displaying the following modes:

   — **Simple**

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

   — **Normal**

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

   — **Thorough**

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

   — **Stop Erase**

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

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

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

NOTE **Thorough** erase is not recommended for SSDs. If an erase of SSDs is required, a **Simple** erase is recommended.

Chapter 5: HII Configuration Utility

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

In addition to Intel and AMD, the controllers can also be used on the following 64-bit ARM platform with limited operating system support:

- Fedora
- Ubuntu
- CentOS

5.1 Behavior of HII

The Human Interface Infrastructure (HII) Configuration Application is used to configure controllers, physical disks, and virtual disks, and to perform other configuration tasks in a pre-boot environment.

Some of the HII Graphical User Interface keys are provided by the system BIOS. HII RAID management screens are tightly controlled by independent hardware vendors. OEMs or independent browser vendors will have no knowledge about independent hardware vendor features and their screen controls.

The following figure is an example of some of the HII GUI keys.

Figure 58 HII Keys

If the keys shown in the preceding figure do not work as expected, contact your system vendor.

For example, you may press the **F2** key and then press the <**ESC**> key to exit from the HII RAID Management screen. However, this action does not save the previous values you specified to the controller. To save the specified values, you must use the controls present in the form or screen provided by your independent hardware vendor.

Similarly, when you want to load controller defaults, you can achieve this by clicking on the **Set Factory Default** option on the first page of the **Controller Management > Advanced Controller Management > Set Factory Defaults** menu. Pressing **F3** (Optimized Defaults) will not restore the controller default settings.

5.2 Starting the HII Configuration Utility

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

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

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

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

   The **System Settings** dialog appears.

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

   The **Controller Selection** menu appears.

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

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

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

Figure 59 Dashboard View

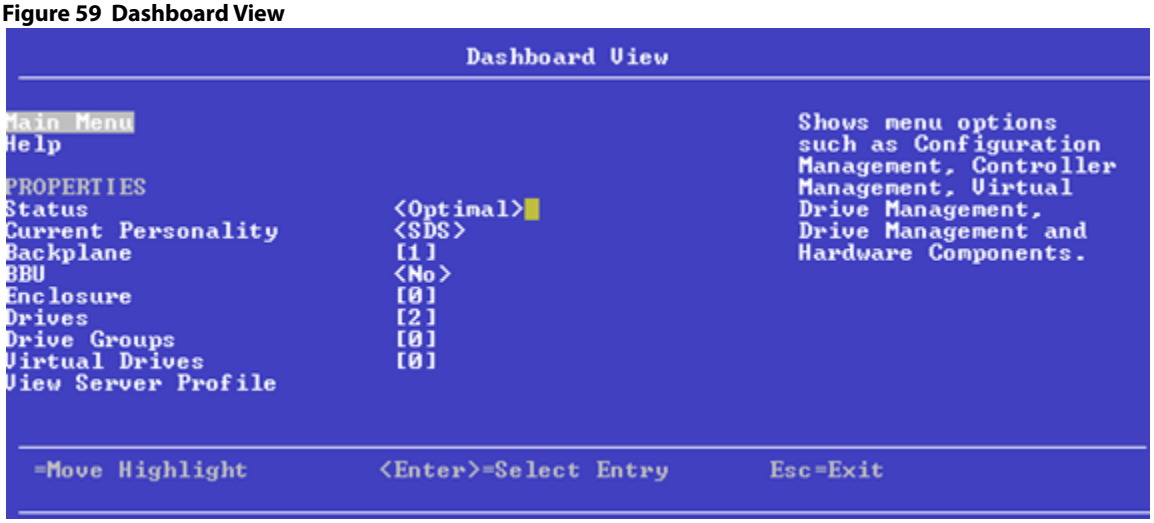

5.3 HII Dashboard View

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

5.3.1 Main Menu

When you select the **Main Menu** option in the **Dashboard View**, the **Main Menu** dialog appears. The **Main Menu** provides various menu options to configure and manage controllers, virtual drives, drive groups, and hardware components.

Figure 60 Main Menu

When the controller is running in **Safe Mode**, the **Main Menu** includes the warning message.

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

5.3.2 HELP

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

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

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

5.3.3 PROPERTIES

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

Figure 61 Dashboard View – PROPERTIES

- **Status**
  Displays the status of the controller.
- **Backplanes**
  Displays the total number of backplanes connected to the controller.
- **BBU**
  Displays whether the battery backup unit is present.
- **Enclosures**
  Displays the total number of enclosures connected to the controller.
- **Drives**
  Displays the total number of drives connected to the controller.
- **Drive Groups**
  Displays the number of drives groups.
- **Virtual Drives**
  Displays the number of virtual drives.
- **View Server Profile**
  Clicking on **View Server Profile** displays the UEFI specification version that the system supports and the following menu options, as shown in the following figure.

Figure 62 Dashboard View – PROPERTIES – Server Profile

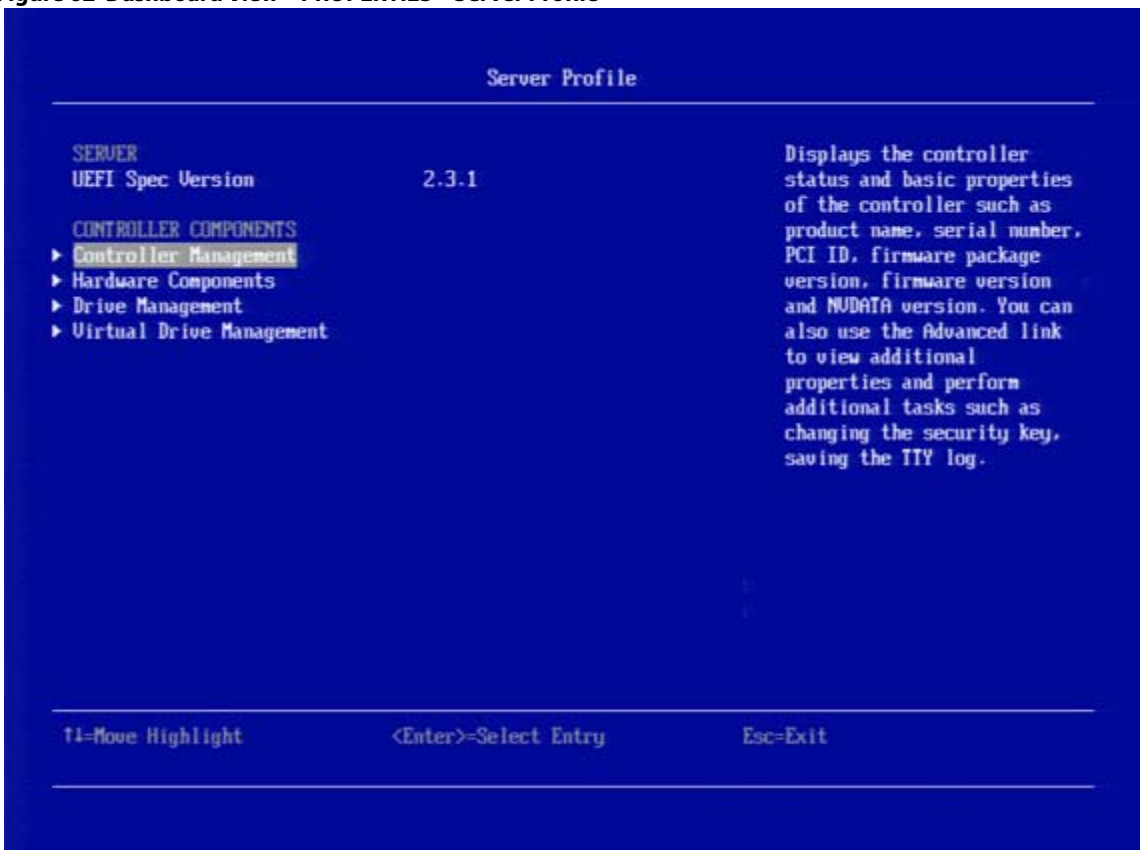

- Select **Controller Management** to view and manage controller properties and to perform tasks, such as clearing configurations, scheduling and running controller events, and running patrol reads.

  For more information, see Managing Controllers.
- **Hardware Components** to view battery properties, manage batteries, and manage enclosures.

  For more information, see Managing Hardware Components.
- **Drive Management** to view physical drive properties and to perform tasks, such as locating drives, initializing drives, and rebuilding a drive after a drive failure.

  For more information, see Managing Physical Drives.
- **Virtual Drive Management** to perform tasks, such as viewing virtual drive properties, locating virtual drives, and running a consistency check.

  For more information, see Managing Virtual Drives.

5.3.4 ACTIONS

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

Figure 63 Dashboard View – ACTIONS

- **Discard Preserved Cache**

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

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

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

- **View Foreign Configuration**

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

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

- **Configure**

  Displays configuration options. See Managing Configurations.

- **Set Factory Defaults**

  Resets the controller to its factory settings.

- **Update Firmware**

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

- **Silence Alarm**

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

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

5.3.5 BACKGROUND OPERATIONS

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

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

Figure 64 Dashboard View – BACKGROUND OPERATIONS

5.4 Critical Boot Error Message

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

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

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

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

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

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

5.5 Managing Configurations

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

Figure 65 Configuration Management

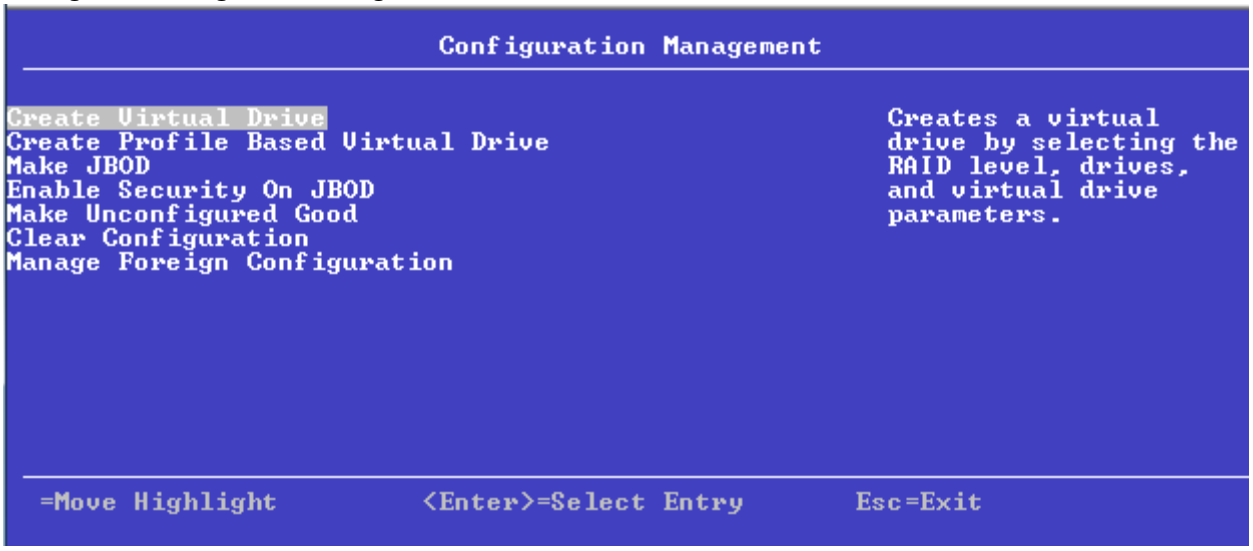

The Manage Foreign Configuration option is included for some configurations. (See Managing Foreign Configurations.)

The HII utility supports 240 VD creation. For more information, see Support Limitations.

5.5.1 Creating a Virtual Drive from a Profile

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

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

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

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

   NOTE The red arrow appears when there is too much information to display in one window. The amount of information that can be displayed in one window depends on the capabilities of the HII browser. The **Save Configuration** option is not displayed in the previous figure.
4. Choose an option from the **Drive Selection Criteria** field (if more than one option exists).
5. Select **Save Configuration** to create the chosen profile.
6. Highlight **Confirm** and press the spacebar, then highlight **Yes** and press **Enter**.

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

Table 21 Virtual Drive Creation Profile Options

Option	Description
Drive Selection Criteria	You need to select one of the various combinations of options that exist. If only one option is possible, only one option appears.
Profile Parameters:	
Virtual Drive Name	Displays the name of the virtual drive.
RAID Level	Displays the RAID level based on the profile selected. For example, if the profile selected is Generic RAID 0, **RAID 0** is displayed.
Virtual Drive Size	Displays the amount of virtual drive storage space. By default, the maximum capacity available for the virtual drive is displayed. **NOTE** Virtual drive size of floating data type up to three decimal places is supported. Some of the screens in this chapter may not reflect this feature.
Power Save Mode	Displays the selected Power Save Mode of the five available options: **None**, **Auto**, **Max**, **Max without Cache**, and **Controller Defined**.
Strip Size	Displays the strip element size for the virtual drive. Drive Stripping involves partitioning each physical drive storage space in strips of the following sizes: **64 KB**, **128 KB**, **256 KB**, **512 KB**, **1 MB**.

Table 21 Virtual Drive Creation Profile Options (Continued)

Option	Description
Read Policy	Displays the read cache policy for the virtual drive. For any profile, if the drive is an SSD drive, the **No Read Ahead** and **Always Read Ahead** options are displayed. However, No Read Ahead is the **default** read policy. The possible options follow: ■ **Default** A virtual drive property that indicates whether the default read policy is **Always Read Ahead** or **No Read Ahead**. ■ **Always Read Ahead** - Permits the controller to read sequentially ahead of the requested data and allows the controller to store the additional data in the cache memory. Here, the controller anticipates that the data is required frequently. Even though Always Read Ahead policy speeds up the reads for sequential data, but little improvement is seen when accessing the random data. ■ **No Read Ahead** - Disables the Always Read Ahead capability of the controller.
Write Policy	Displays the write cache policy for the virtual drive. For any profile, if the drive is an SSD drive, the **Write Through** option is displayed. Otherwise, the **Always Write Back** option is displayed. The possible options follow: ■ **Write Back** The controller sends a data transfer completion signal to the host when the controller cache receives all of the data in a transaction. If you select the Write Back policy and the battery is absent, the firmware disables the Write Back policy and defaults to the Write Through policy. ■ **Write Through** The controller sends a data transfer completion signal to the host when the drive subsystem receives all the data in a transaction. ■ **Always Write Back** The controller sends a data transfer completion signal to the host when the controller cache receives all the data in a transaction. If you select the Always Write Back policy and the battery is absent, the firmware is forced to use the Write Back policy.
I/O Policy	Displays the Input/Output policy for the virtual drive. For any profile, if the drive is an SSD drive, the **Direct** option is displayed. The possible options follow: ■ A virtual drive property that indicates whether the default I/O policy is Direct IO or Cached IO. ■ **Direct IO** Data read operations are not buffered in the cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from the cache memory. (The I/O policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.) ■ **Cached IO** All read operations are buffered in cache.
Access Policy	The access policy for the virtual drive. The options are **Read/Write** and **Read Only**.

Table 21 Virtual Drive Creation Profile Options (Continued)

Option	Description
Disk Cache Policy	Displays the virtual drive cache setting. The possible options are **Unchanged**, **Enable**, and **Disable**.
Default Initialization	Displays the virtual drive initialization setting. The Default Initialization displays the following options: ■ **No** Do not initialize the virtual drive. ■ **Fast** Initializes the first 100 MB on the virtual drive. ■ **Full** Initializes the entire virtual drive.
Save Configuration	Saves the configuration that the wizard created.

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

Table 22 Profile Based Virtual Drive Creation Requirements

Properties	Generic RAID0	Generic RAID1	Generic RAID5	Generic RAID6
HDD	Supported	Supported	Supported	Supported
SSD	Supported	Supported	Supported	Supported
SAS	Supported	Supported	Supported	Supported
SATA	Supported	Supported	Supported	Supported
PCIe	Supported	Supported	Supported	Not supported
SED	Supported	Supported	Supported	Supported
NonSED	Supported	Supported	Supported	Supported
Protected Information (PI)	Supported	Supported	Supported	Supported
NonProtected Information (NonPI)	Supported	Supported	Supported	Supported
Sector Size (logical block format size) – 4 KB	Supported	Supported	Supported	Supported
Sector Size (logical block format size) – 512 B	Supported	Supported	Supported	Supported
Link speed – 3Gb/s	Supported	Supported	Supported	Supported
Link speed – 6Gb/s	Supported	Supported	Supported	Supported
Link speed – 12Gb/s	Supported	Supported	Supported	Supported
Direct attached	Supported	Supported	Supported	Supported
Backplane	Supported	Supported	Supported	Supported
Enclosure	Supported	Supported	Supported	Supported
Minimum number of PDs	1	2	3	4
Maximum number of PDs	0xFF	2	0xFF	0xFF
Power-save mode	Controller-defined	Controller-defined	Controller-defined	Controller-defined

Table 22 Profile Based Virtual Drive Creation Requirements (Continued)

Properties	Generic RAID0	Generic RAID1	Generic RAID5	Generic RAID6
Strip Size	256 KB	256 KB	256 KB	256 KB
Read Policy	If the drive is an SSD drive, then the No Read Ahead option appears. Else, the Default options appears.	If the drive is an SSD drive, then the No Read Ahead option appears. Else, the Default option appears.	If the drive is an SSD drive, then the No Read Ahead option appears. Else, the Default option appears.	If the drive is an SSD drive, then the No Read Ahead option appears. Else, the Default option appears.
Write Policy	If the drive is an SSD drive, then the Write Through option appears. Else, the Write Back option appears.	If the drive is an SSD drive, then the Write Through option appears. Else, the Write Back option appears.	If the drive is an SSD drive, then the Write Through option appears. Else, the Write Back option appears.	If the drive is an SSD drive, then the Write Through option appears. Else, the Write Back option appears.
IO Policy	If the drive is an SSD drive, then the Direct IO option appears. Else, the Default options appears.	If the drive is an SSD drive, then the Direct IO option appears. Else, the Default options appears.	If the drive is an SSD drive, then the Direct IO option appears. Else, the Default options appears.	If the drive is an SSD drive, then the Direct IO option appears. Else, the Default options appears.
Access policy	Read/Write	Read/Write	Read/Write	Read/Write
Disk Cache Policy	Enable	Unchanged	Unchanged	Unchanged
Initialization	Fast	Fast	Full	Full
Dedicated Hot Spare	Not supported	Supported	Supported	Supported
Mixing of Media HDD and SSD drives	Not supported	Not supported	Not supported	Not supported
Mixing of Interface Type SAS and SATA drives	Not supported	Not supported	Not supported	Not supported
Mixing of PI and NonPI drives	Not supported	Not supported	Not supported	Not supported
Mixing SED and NonSED drives	Not supported	Not supported	Not supported	Not supported
Mixing of 1.5Gb/s, 3Gb/s, 6Gb/s, and 12Gb/s link speeds	Not supported	Not supported	Not supported	Not supported

5.5.2 Manually Creating a Virtual Drive

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

Figure 66 Create Configuration Window

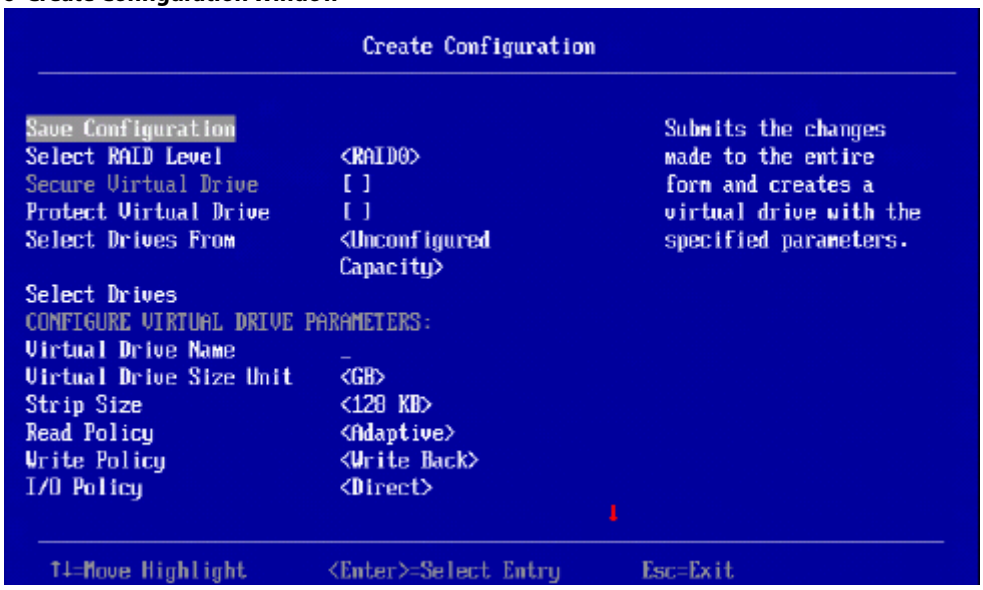

The small red arrow at the bottom of the window indicates that you can scroll down to view more information.

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

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

1. Highlight the **Select RAID Level** field and press **Enter**.
2. Select a RAID level for the virtual drive from the popup menu.

   The available RAID levels are listed in the help text of the **Create Configuration** dialog. Some system configurations do not support all these RAID levels. See Table 24 for brief descriptions of the RAID levels.
3. To view the **Secure Virtual Drive** field, enable security and attach an FDE drive. If either is missing, the field is grayed out.
4. To view the **Protect Virtual Drive** field, enable protection and attach a protected drive.

   If you have chosen all PI drives, the **Protect Virtual Drive** check box is selected only if the protection feature is supported and enabled. If you do not want this protection feature on the virtual drive, you must clear the **Protect Virtual Drive** check box.

   If you have enabled data protection while creating a virtual drive, you must select either **Full Initialization** or the **Background Initialization**, otherwise you cannot create a virtual drive.

   If there is no drive that is capable of protection, then the **Protect Virtual Drive** field is grayed out. Also, if the controller does not support Virtual Drive Protection, then the **Protect Virtual Drive** field will be suppressed.
5. If the security key is enabled, highlight the **Secure Virtual Drive** field to secure the new virtual drive.

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

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

   Free capacity means the new virtual drive is created from unused (free) drive capacity that is already part of a virtual drive. *Unconfigured capacity* means the new virtual drive is created on previously unconfigured drives.
8. Highlight the **Virtual Drive Name** field, press **Enter**, and enter a name for the new virtual drive.

9. (Optional) Change the **Virtual Drive Size Unit** value by highlighting this field, pressing **Enter**, and selecting a value from the popup menu.

   The options are MB, GB, and TB.

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

    See Table 23 for descriptions of these options.

11. Highlight **Select Drives** and press **Enter**.

    The following dialog appears.

Figure 67 Select Drives Window

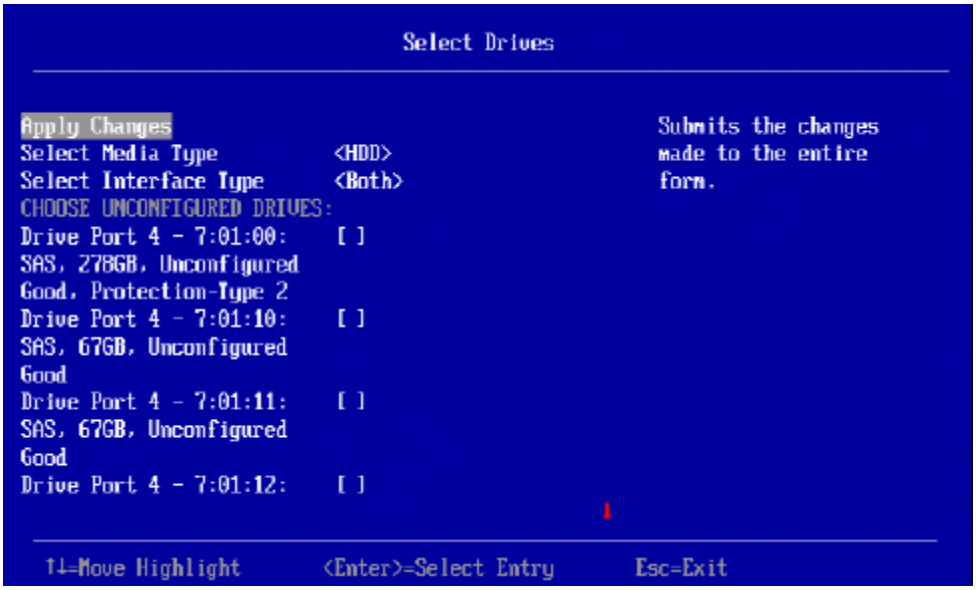

Follow these steps to select physical drives for the new virtual drive.

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

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

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

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

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

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

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

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

   NOTE Be sure to select the number of drives required by the specified RAID level, or the HII utility will return you to the root menu when you try to create the virtual drive. For example, RAID 1 virtual drives use exactly

two drives, and RAID 5 virtual drives use three or more virtual drives. See Table 24 for more information.

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

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

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

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

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

   A message appears confirming that the configuration is being created.

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

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

Table 23 Virtual Drive Policies

Property	Description
Strip Size	The virtual drive strip size per DDF. The possible values are as follows: ■ **7**: 64 KB ■ **8**: 1 MB
Read Policy	Displays the read cache policy for the virtual drive. For any profile, if the drive is an SSD drive, the **No Read Ahead**and **Always Read Ahead**options are displayed. However, No Read Ahead is the **default** read policy. The possible options follow: ■ **Default** A virtual drive property that indicates whether the default read policy is **Always Read Ahead** or **No Read Ahead**. ■ **Always Read Ahead** - Permits the controller to read sequentially ahead of the requested data and allows the controller to store the additional data in the cache memory. Here, the controller anticipates that the data is required frequently. Even though Always Read Ahead policy speeds up the reads for sequential data, but little improvement is seen when accessing the random data. ■ **No Read Ahead** - Disables the Always Read Ahead capability of the controller.
Write Policy	The write cache policy for the virtual drive. The possible values are as follows: ■ **Write Back** The controller sends a data transfer completion signal to the host when the controller cache receives all of the data in a transaction. If you select the Write Back policy and the battery is absent, the firmware disables the Write Back policy and defaults to the Write Through policy. ■ **Write Through** The controller sends a data transfer completion signal to the host when the drive subsystem receives all the data in a transaction. ■ **Always Write Back** The controller sends a data transfer completion signal to the host when the controller cache receives all the data in a transaction. If you select the Always Write Back policy and the battery is absent, the firmware is forced to use the Write Back policy.

Table 23 Virtual Drive Policies (Continued)

Property	Description
I/O Policy	The I/O policy for the virtual drive. The possible values are as follows: ■ **Direct** Data reads are not buffered in cache memory. Data is transferred to cache and the host concurrently. If the same data block is read again, it comes from cache memory. (The I/O policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.) ■ **Cached** All reads are buffered in cache.
Access Policy	The access policy for the virtual drive. The options are **Read/Write**, **Read Only**, and **Blocked**.
Drive Cache	The disk cache policy for the virtual drive. The possible values are **Unchanged**, **Enable**, and **Disable**.
Disable Background Initialization (BGI)	Specifies whether background initialization is enabled or disabled. When BGI is enabled, the firmware runs the initialization process in the background. When BGI is disabled, the initialization process does not start automatically and does not run in the background.
Default Initialization	Allows choice of virtual drive initialization option. The possible options are **No**, **Fast**, and **Slow**.

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

Table 24 RAID Levels

Level	Description
RAID 0	Uses data striping on two or more drives to provide high data throughput, especially for large files in an environment that requires no data redundancy.
RAID 1	Uses data mirroring on pairs of drives so that data written to one drive is simultaneously written to the other drive. RAID 1 works well for small databases or other small applications that require complete data redundancy.
RAID 5	Uses data striping and parity data across three or more drives (distributed parity) to provide high data throughput and data redundancy, especially for applications that require random access.
RAID 6	Uses data striping and parity data across three or more drives (distributed parity) to provide high data throughput and data redundancy, especially for applications that require random access. RAID 6 can survive the failure of two drives.
RAID 00	Is a spanned drive group that creates a striped set from a series of RAID 0 drive groups to provide high data throughput, especially for large files.
RAID 10	A combination of RAID 0 and RAID 1 that uses data striping across two mirrored drive groups. It provides high data throughput and complete data redundancy.
RAID 50	A combination of RAID 0 and RAID 5 that uses data striping across two drive groups with parity data. It provides high data throughput and complete data redundancy.
RAID 60	A combination of RAID 0 and RAID 6 that uses data striping across two drive groups with parity data. It provides high data throughput and complete data redundancy. RAID 60 can survive the failure of two drives in each RAID set in the spanned drive group.

5.5.3 Viewing Drive Group Properties

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

Figure 68 View Drive Group Properties Window

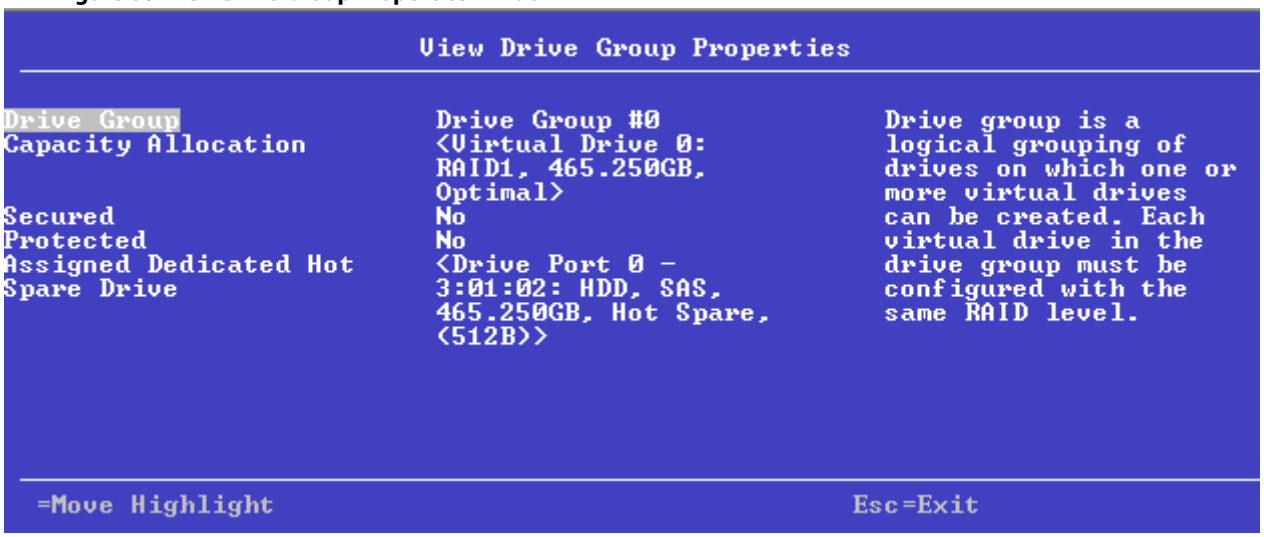

A drive group is a logical grouping of drives attached to a RAID controller on which one or more virtual drives can be created. Each virtual drive in the drive group must be configured with the same RAID level. This figure shows information for one drive group.

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

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

5.5.4 Viewing Global Hot Spare Drives

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

Figure 69 View Global Hot Spare Drives

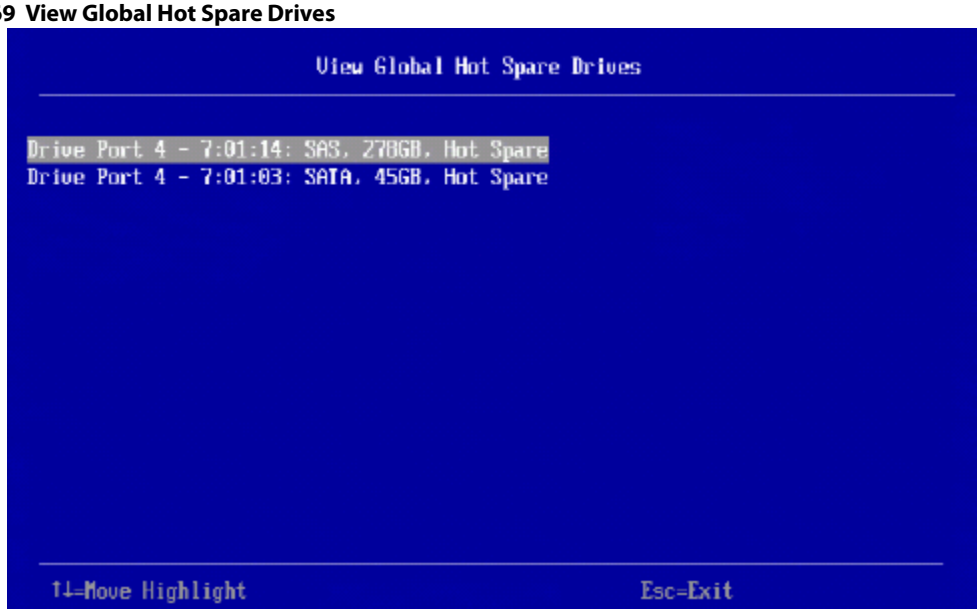

Press **Esc** to exit this window when you are finished viewing information.

5.5.5 Clearing a Configuration

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

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

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

Perform the following steps to clear configuration:

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

NOTE If your system is in JBOD personality mode, Clear Configuration clears the existing virtual drives and JBOD.

5.5.6 Managing Foreign Configurations

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

Figure 70 Manage Foreign Configuration

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

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

5.5.6.1 Previewing and Importing a Foreign Configuration

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

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

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

Figure 71 Enter Security Key for Locked Drives

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

Perform these steps to preview and import a foreign configuration.

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

Figure 72 Preview Configuration Window 1

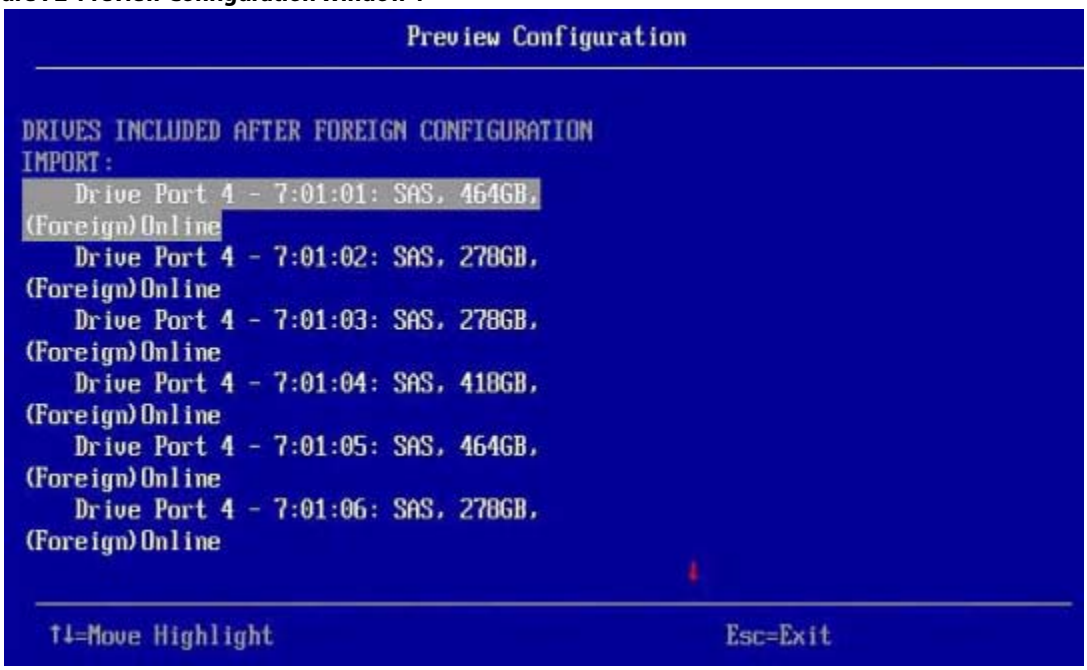

2. Scroll down, if needed, to view more information about the drives in the foreign configuration, as shown in the following figure.

Figure 73 Preview Configuration Window 2

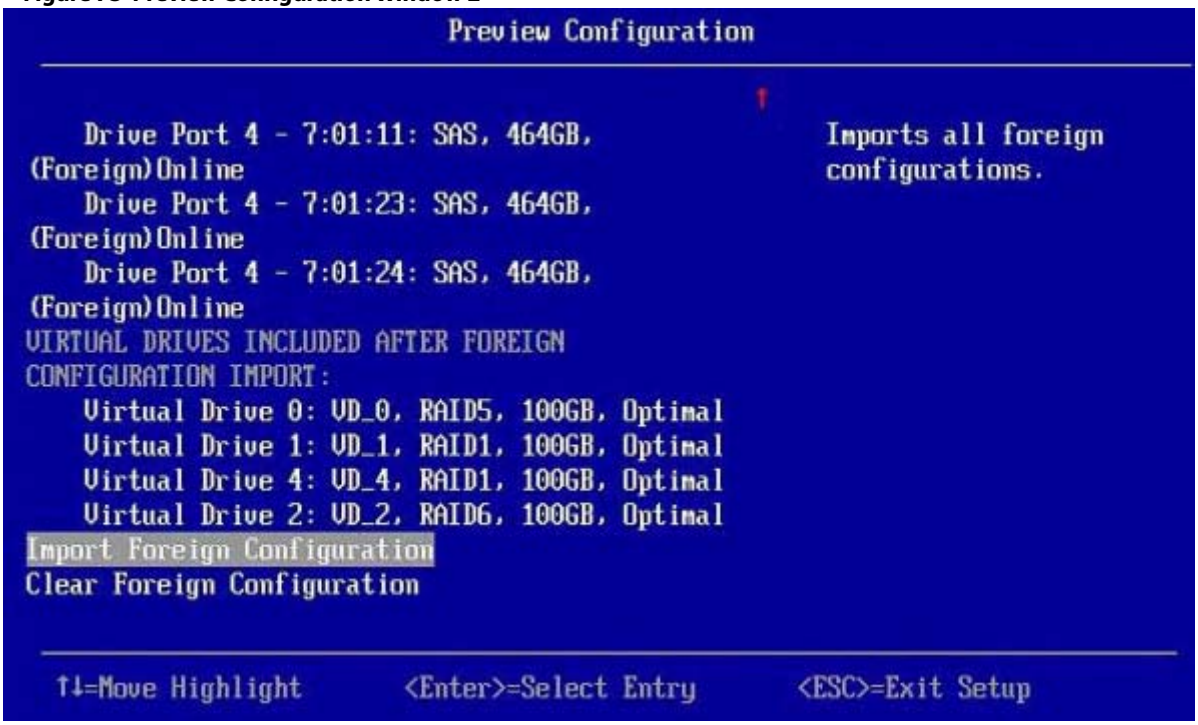

3. Review the information listed on the window.
4. Highlight **Import Foreign Configuration** and press **Enter**.

   A warning message appears that indicates the foreign configuration from the physical drives will merge with the existing configuration.
5. To confirm the import, highlight **Confirm** and press the spacebar.
6. Highlight **Yes** and press **Enter**.

   The foreign configuration is imported.

5.5.6.2 Clearing a Foreign Configuration

Perform these steps to clear a foreign configuration.

1. Highlight **Clear Foreign Configuration** on the **Manage Foreign Configuration** menu and press **Enter**.

   A warning message appears that indicates all of the foreign VDs will be deleted.
2. To confirm clearing the foreign configuration, highlight **Confirm** and press the spacebar.
3. Highlight **Yes** and press **Enter**.

   The foreign configuration is deleted.

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

5.6 Managing Controllers

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

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

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

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

  For more information, see Advanced Controller Properties.

Figure 74 Controller Management

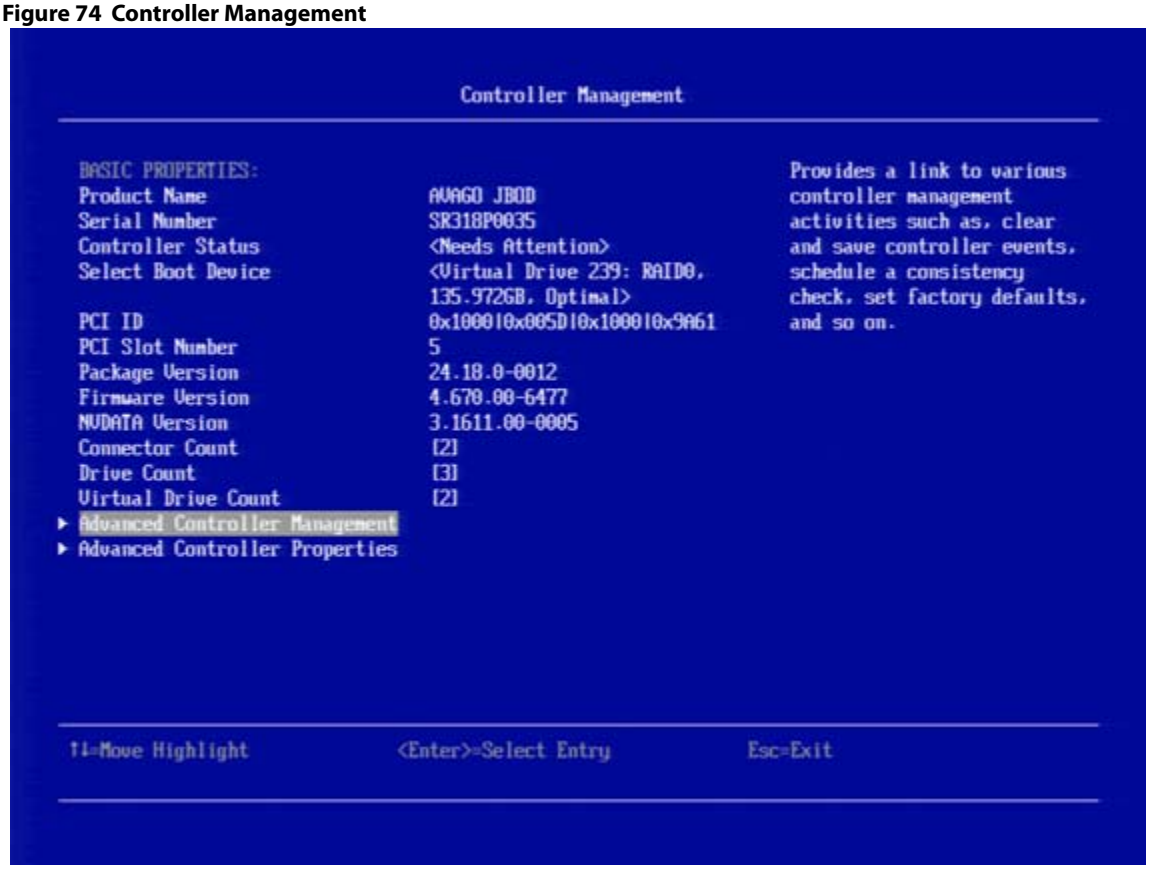

The **Controller Management** dialog lists the following basic controller properties.

Table 25 Basic Controller Properties

Property	Description
Product Name	The marketing name of the controller.
Serial Number	The serial number of the controller.
Controller Status	The cumulative status of virtual drives and physical drives connected to the controller, plus the backup battery, the enclosure and the NVDATA. The status is one of the following: ■ **Optimal**, if all components are operating normally. ■ **Needs Attention**, if any component needs attention. ■ **Safe Mode**, if the controller encountered critical errors. Most features are disabled and the controller requires user attention.
Select Boot Device	This field selects the primary boot device. The boot device may contain virtual drives and JBODs. The system here is assuming that you are selecting an offline virtual drive as a boot device because you are planning to bring this virtual drive to an optimal or a degraded state (if it is a redundant RAID volume). User's discretion is advised while assigning a boot device. **NOTE** This property is applicable for legacy BIOS. **NOTE** You will not be able to set 4K block-size devices such as JBOD or VD as boot devices; hence, preeboot utilities such as CTRL-R or HII does not allow you to choose 4K devices as boot devices. Instead, you can use the UEFI environment to boot 4K devices.
PCI ID	The PCI ID of the controller.
PCI Slot Number	The slot ID number of the PCI slot where the controller is installed.
Package Version	The version number of the package.
Expander Firmware Version	This field shows the firmware version of the expander that is connected to the controller. **NOTE** This field only appears when an expander is connected to the controller.
Firmware Version	The version number of the controller firmware.
NVDATA Version	The version number of the controller NVDATA.
Connector Count	Number of host data ports, connectors, or both currently in use on this controller.
Drive Count	Number of physical drives attached to this controller.
Virtual Drive Count	Number of virtual drives defined on this controller.

5.6.1 Advanced Controller Management Options

The **Advanced Controller Management** dialog lists all the controller management options and various actions that you can perform on the controller.

Figure 75 Advanced Controller Management

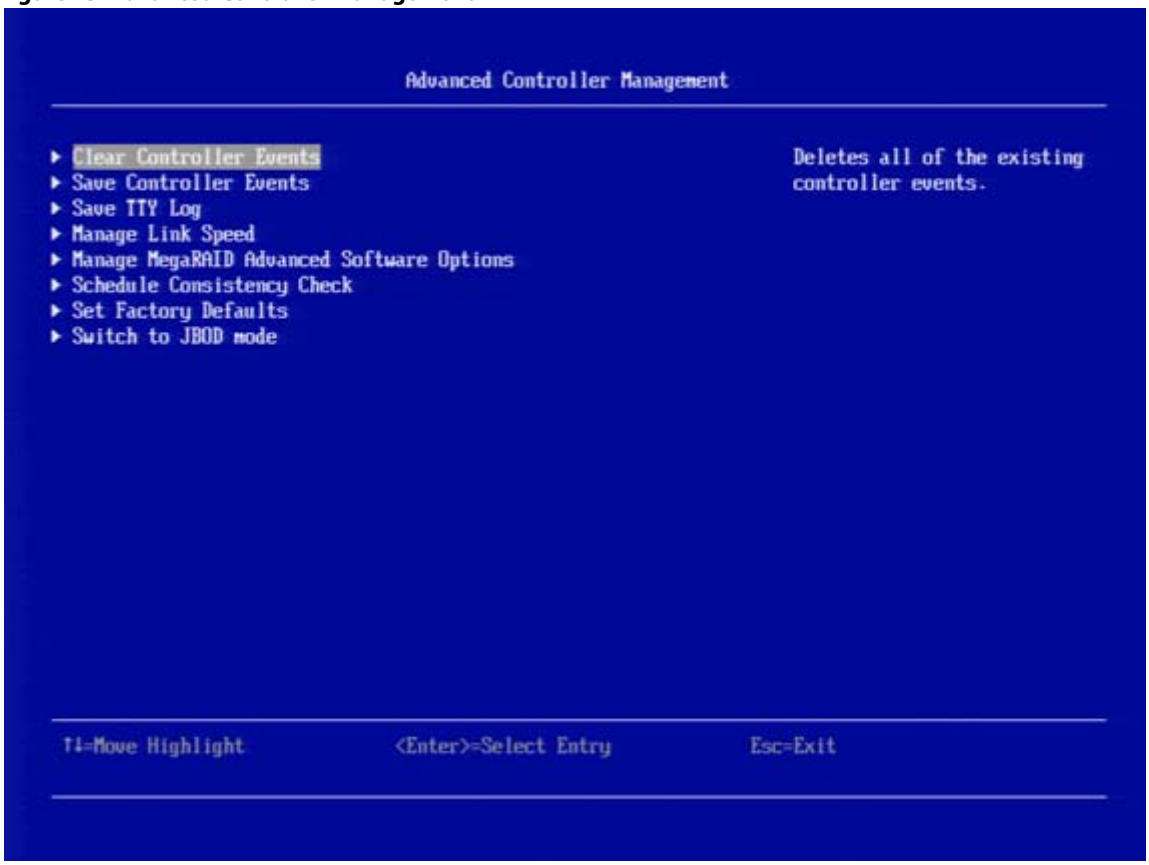

The following table describes all of the entries on the **Advanced Controller Management** dialog, including the ones that are not visible.

Table 26 Controller Management Options

Property	Description
Clear Controller Events	Clears entries from the log.
Save Controller Events	Saves the controller log entries to a file.
Save TTY Log	Saves a copy of the firmware's terminal log entries for the controller.
Enable Drive Security	Enables drive security to protect the data on your system from unauthorized access or use.
Disable Drive Security	Disables drive security.
Change Security Key	Changes the security key or switch between drive security modes on the controller.
Manage Link Speed	Enables you to change the link speed between the controller and an expander or between the controller and a drive that is directly connected to the controller. For more information, see Managing and Changing Link Speeds.

Table 26 Controller Management Options (Continued)

Property	Description
Manage ThinkSystem Advanced Software Options	DIsplays the activated ThinkSystem Advanced Software Options on the controller and lets you configure these options to use the advanced features in the controller. You need to activate the activation key to use the advanced features. **NOTE** The ThinkSystem Advanced Software Options are displayed only if the controller supports ThinkSystem software licensing.
Schedule Consistency Check	Schedules a consistency check operation to verify and correct the mirror and parity data for fault tolerant virtual drives.
Set Factory Defaults	Resets the controller to its factory settings.

5.6.1.1 Saving or Clearing Controller Events

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

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

Figure 76 Save Controller Events

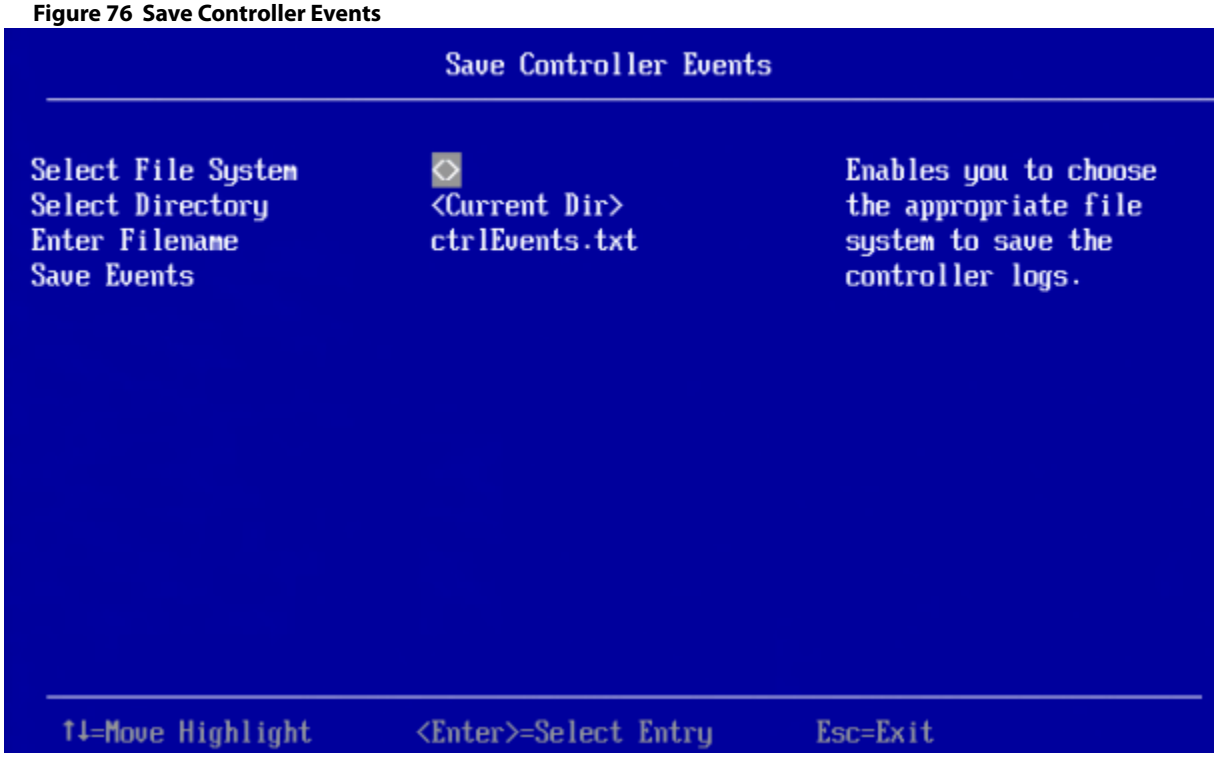

Perform these steps to save controller event log entries to a file.

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

   An error message appears if there is no file system.
2. Select a file system from the popup menu and press **Enter**.
3. To save the controller events file to a different directory from the one listed in the **Select Directory** field, highlight the current directory name and press **Enter**.
4. Select a directory name from the popup menu and press **Enter**.
5. To enter a different name for the controller event log file, highlight the current file name and press **Enter**.

6. Type the new file name in the popup dialog and press **Enter**.
7. Highlight **Save Events**, and press **Enter** to save the event log entries to the file.

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

5.6.1.2 Saving the TTY Log

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

Figure 77 Save TTY Log

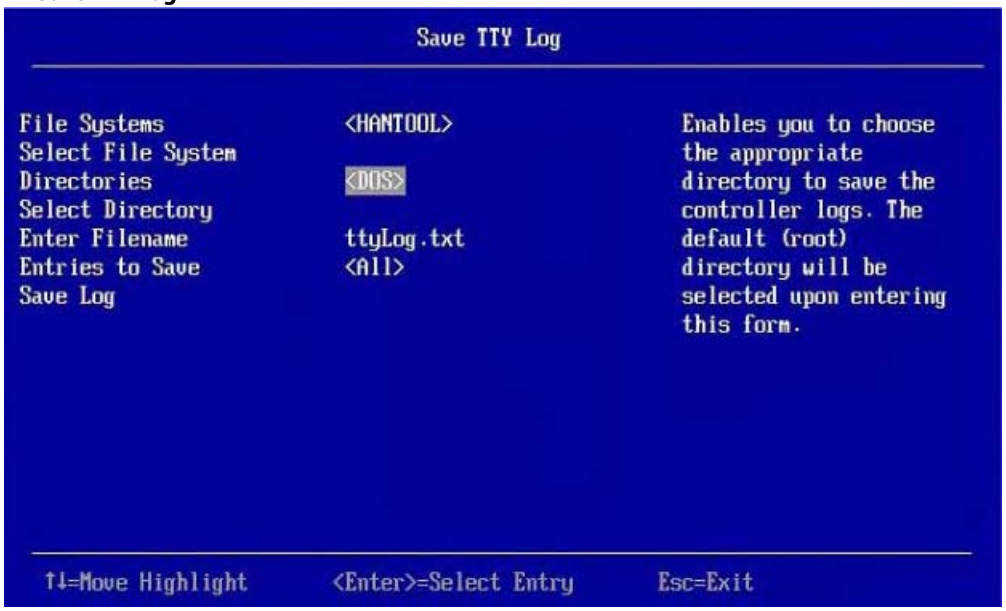

Follow these steps to save the TTY log entries to a file.

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

   An error message appears if there is no file system.
2. Select a file system from the popup menu, and press **Enter**.
3. Highlight **Select File System** and press **Enter**.
4. To save the TTY log events file to a different directory from the one listed in the **Directories** field, highlight the current directory name, and press **Enter**.
5. Select a directory name from the popup menu, and press **Enter**.
6. Highlight **Select Directory**, and press **Enter**.
7. To enter a different name for the TTY log file, highlight the current file name, and press **Enter**.
8. Type the new file name in the popup window, and press **Enter**.
9. To select how many TTY log entries to save, highlight the **Entries to Save** field, and press **Enter**.
10. Select an option from the popup menu, and press **Enter**.

    Your choices are **2 KB**, **4 KB**, **8 KB**, **16 KB**, or **All**.
11. Highlight **Save Log**, and press **Enter** to save the log entries to the file.

5.6.1.3 Enabling or Disabling Drive Security

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

Figure 78 Enable Drive Security (Choose Drive Security Mode)

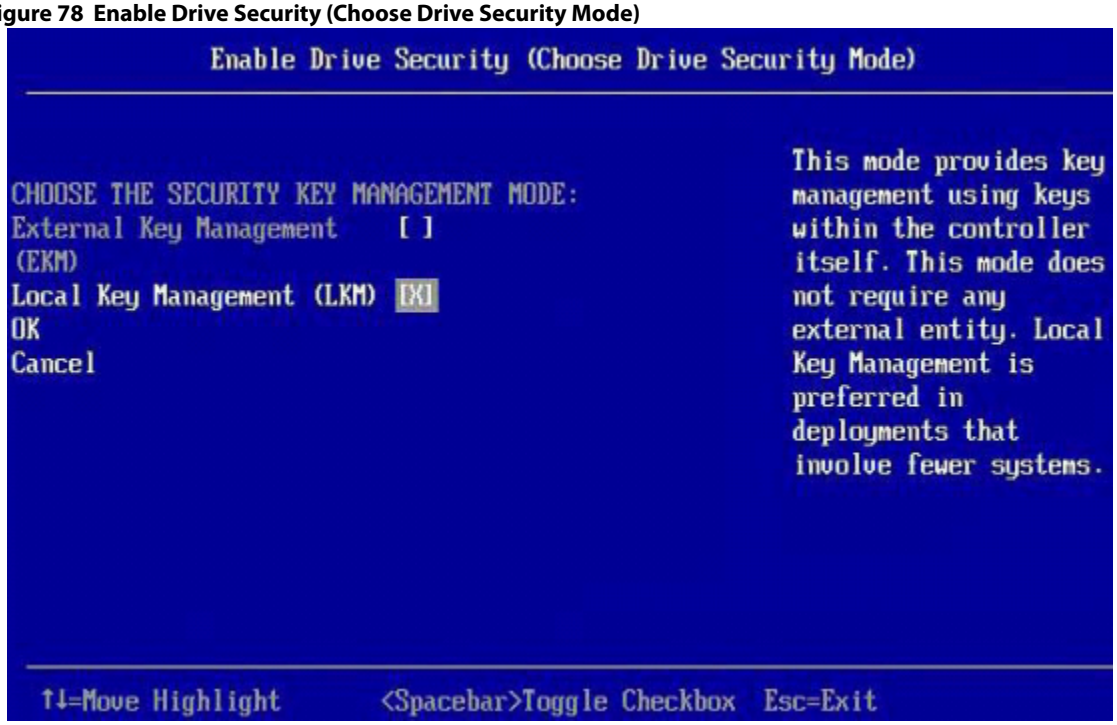

Enable drive security to protect the data on your system from unauthorized access or use. Local Key Management (LKM) is the method that the HII utility provides to manage drive security. LKM uses security keys within the controller and does not require any external entity to implement. Therefore, it is the preferred security mode for configurations that involve a smaller number of computer systems.

Follow these steps to enable LKM security on your configuration.

1. Highlight the **Local Key Management (LKM)** field and, if required, press the spacebar to enter an X in this field.
2. Highlight **OK** and press **Enter**.
   The following dialog appears.

Figure 79 Enable Drive Security

The highlighted field is the security key identifier, which appears whenever you need to enter the security key. If you have more than one security key, the identifier helps you determine which security key to enter.

3. To change the security key identifier, press **Enter** and enter the new identifier in the popup window.
4. To request the controller to suggest a drive security key, highlight **Suggest Security Key** and press **Enter**.
5. To enter your own security key, highlight the **Security Key** field, press **Enter**, and type the security key.

   The **Security Key** field is case-sensitive. The security key must be between eight and thirty-two characters and must contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (for example, > @ +).
6. After entering the security key, highlight **Confirm** and press **Enter**. Enter the security key again to confirm it.

   The security key must match exactly the characters you entered in the **Security Key** field.
7. If you do not want the controller to require a password at boot time, deselect the **Pause for Password at Boot** option by highlighting it and pressing the spacebar.

   This option is selected by default.
8. To enforce strong password restrictions, highlight **Enforce Strong Password Security** and press the spacebar.

   A strong password must be between eight and thirty-two characters and must contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (for example, > @ +).
9. Highlight the **Password** field, press **Enter**, and type the boot time password.
10. Highlight **Confirm** and re-enter the password.

    The password must match exactly the characters you entered in the **Password** field.
11. Record the drive security information and store it in a safe place.
12. Highlight the **I Recorded The Security Settings** field and press the spacebar to select it.
13. Highlight **Enable Drive Security** and press **Enter**.
14. When the popup window appears, confirm that you want to enable drive security and select **Yes**.

    Drive security is enabled for the drives connected to this controller.

Follow these steps to disable LKM drive security:

1. Select **Disable Drive Security** from the **Advanced Controller Management** menu.

   The following warning appears.

Figure 80 Disable Drive Security Warning

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

5.6.1.4 Changing a Security Key

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

Perform these steps to change the security key settings.

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

Figure 81 Change Security Key

By default, the same security key identifier is retained.

2. To change the security key identifier, press the spacebar to deselect **Use the Existing Security Key Identifier**.
3. Highlight the **Enter a New Security Key Identifier** field, press **Enter**, and enter the new security key identifier in the popup window.
4. Highlight the **Enter Existing Security Key** field and press **Enter**.

   You are required to enter the security key to prevent unauthorized changes to the security settings.
5. Type the current security key in the popup window and press **Enter**.
6. Highlight **Suggest Security Key** and press **Enter** to have the system create a new security key.
7. To enter your own new security key, highlight the **Security Key** field, press **Enter**, and type the new security key.

   The **Security Key** field is case-sensitive. The security key must be between eight and thirty-two characters and must contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (for example, > @ +).
8. After entering the new security key, highlight **Confirm** and press **Enter**. Enter the security key again to confirm it.

   The security key must match exactly the characters you entered in the **Security Key** field.
9. If you do not want the controller to require a password at boot time, deselect the **Pause for Password at Boot** option by highlighting it and pressing the spacebar.

   This option is selected by default.
10. To enforce strong password restrictions, highlight **Enforce Strong Password Security** and press the spacebar.

    A strong password must be between eight and thirty-two characters and must contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (for example, > @ +).
11. Highlight the **Password** field, press **Enter**, and type the new boot time password.
12. Highlight **Confirm** and reenter the new password.

    The password must match exactly the characters you entered in the **Password** field.
13. Record the drive security information and store it in a safe place.
14. Highlight the **I Recorded The Security Settings**field and press the spacebar to select it.
15. Highlight **Change Security Key** and press **Enter**.
16. When the popup window appears, confirm that you want to change the security settings and select **Yes**.

    The security changes are entered for the drives connected to this controller.

5.6.1.5 Managing and Changing Link Speeds

The Manage Link Speed feature lets you change the link speed between the controller and an expander or between the controller and a drive that is directly connected to the controller. The following dialog appears when you select **Manage Link Speed** on the **Advanced Controller Management** dialog. The default settings for all phys is **Auto**.

Figure 82 Manage Link Speed

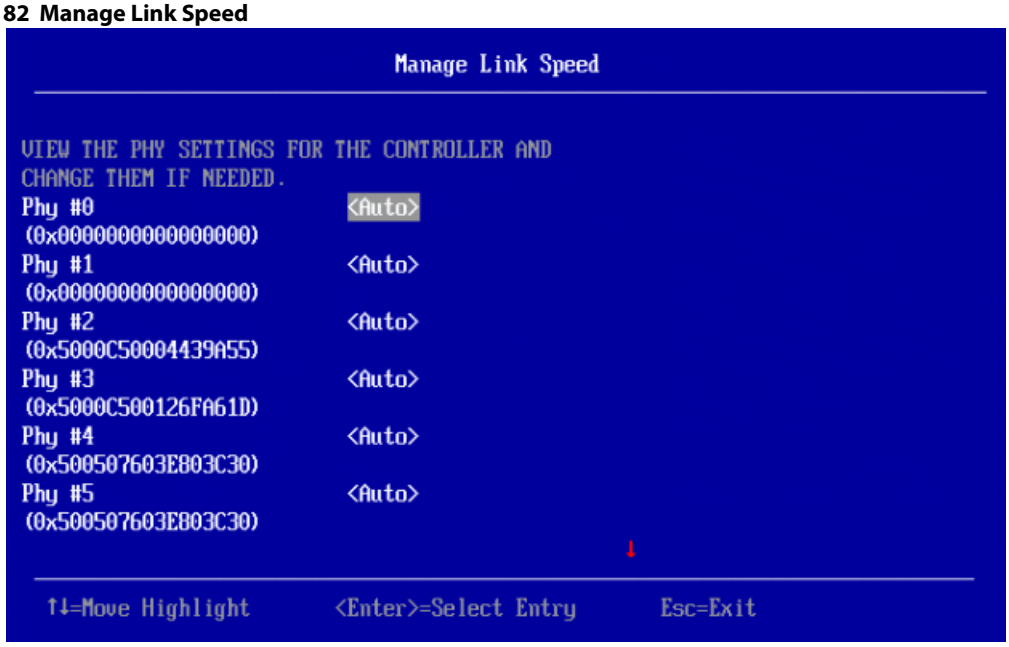

Follow these steps to change the link speed for one or more phys:

1. Highlight the field to the right of the phy number and press **Enter**.
2. Select an option from the popup menu.

   The link speed values are Auto, 1.5Gb/s, 3Gb/s, or 6Gb/s.
3. Scroll to the bottom of the phy list, highlight **OK**, and press **Enter**.

5.6.1.6 Managing Advanced Software Options

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

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

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

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

Figure 83 Manage Advanced Software Options

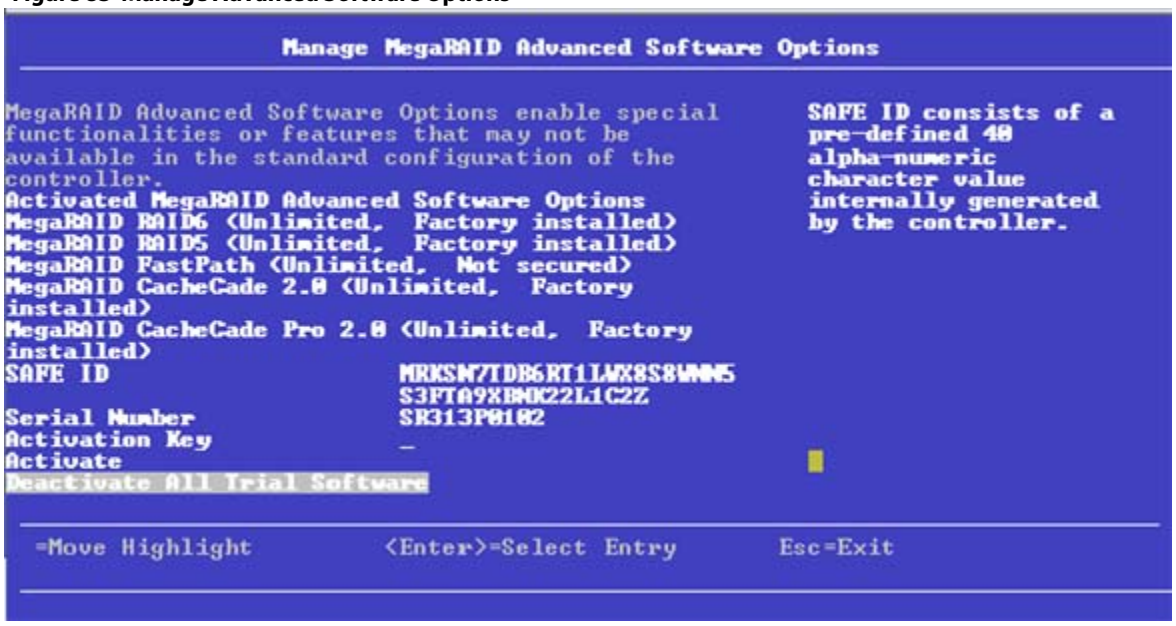

This dialog lists fields that cannot all be shown in one dialog. Scroll down to view all of the fields.

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

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

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

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

5.6.1.7 Scheduling a Consistency Check

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

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

Follow these steps to change the consistency check settings.

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

   A selectable popup menu appears.

Figure 84 Scheduling a Consistency Check

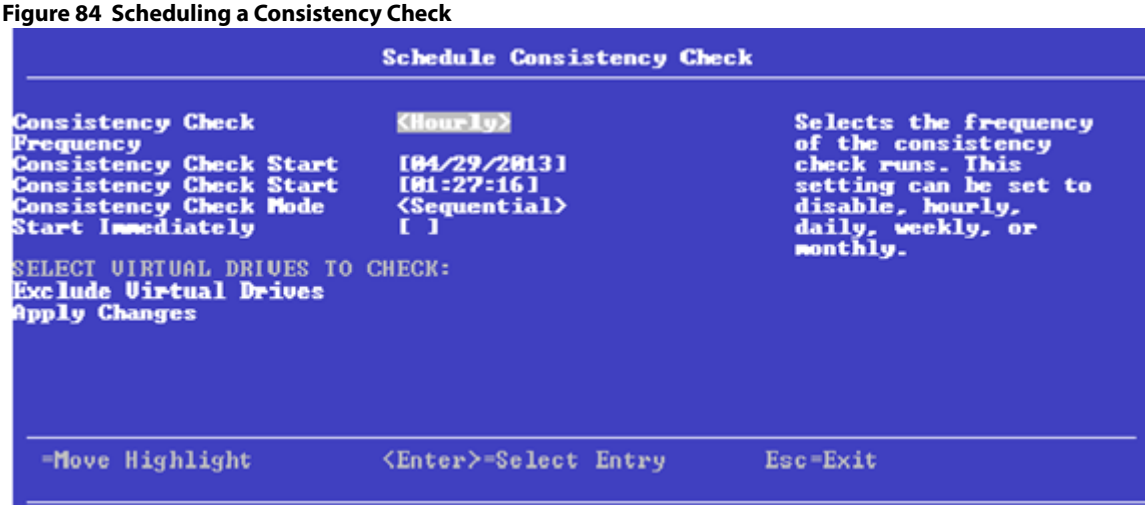

2. Select the desired interval at which to run consistency checks.

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

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

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

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

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

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

   The consistency check changes are now registered.

5.6.2 Advanced Controller Properties

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

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

- To view and modify the controller cache, highlight **Cache and Memory** and press **Enter**.

  For more information, see Setting Cache and Memory Properties.
- To view and set patrol read properties, highlight **Patrol Read**, press **Enter**.

  For more information, see Running a Patrol Read.

- To view and modify physical drive power settings, highlight **Power Settings** and press **Enter**.

  For more information, see Changing Power Save Settings.
- To view and modify properties related to replacing a drive, an emergency spare, or a hot spare, highlight **Spare** and press **Enter**.

  For more information, see Setting Emergency Spare Properties.
- To modify the rebuild rate and other task rates for a controller, highlight **Task Rates**.

  For more information, see Changing Task Rates.

Figure 85 Advanced Controller Properties

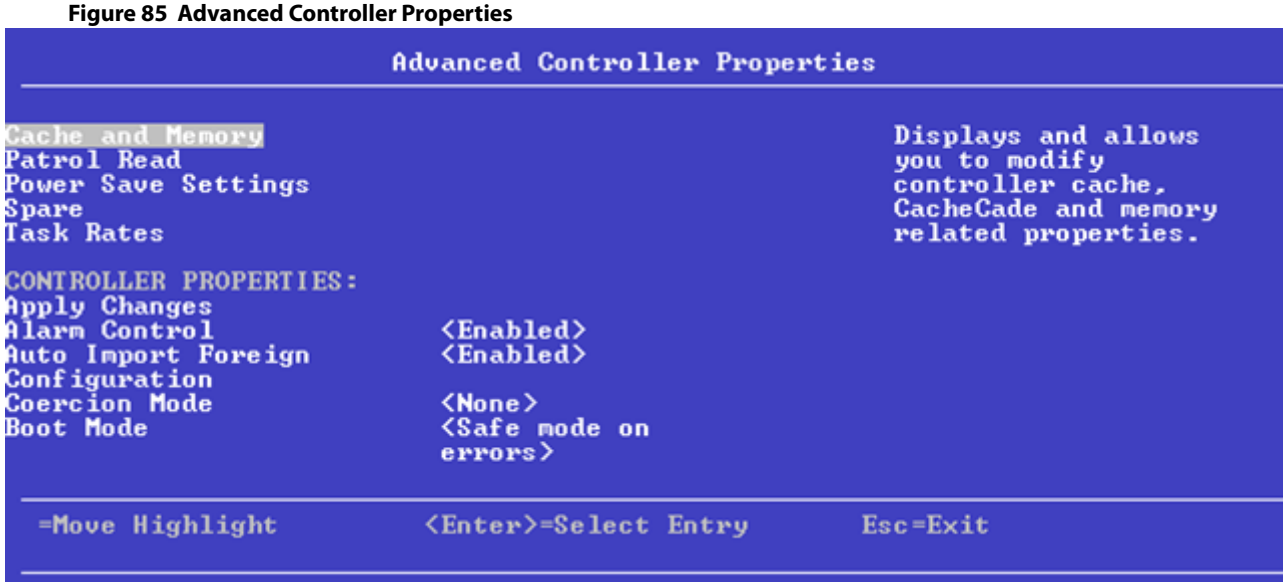

This dialog lists various properties, all of them cannot b e shown in one dialog. Scroll down to view all of the options.

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

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

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

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

   NOTE Some systems permit you to enter numeric values directly, without using the **+** and **–** keys.
3. When you finish changing the controller properties, scrolling up and down on the menu as needed, move the highlight to **Apply Changes** and press **Enter**.

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

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

Table 27 Advanced Controller Properties

Property	Description
Alarm Control	Enables or disables the controller alarm.
Auto Import Foreign Configuration	Enables or disables the automatic import of foreign configurations without any user intervention.
Boot Mode	Specifies the option to handle errors that the firmware might encounter during the boot process. The errors might require you to take action or to acknowledge the error and permit the boot process to continue. The options are *Stop on error*, *Pause on error*, *Ignore errors*, and *Safe mode*.
Controller BIOS	Enables or disables the controller BIOS. The controller BIOS should be enabled if the boot device is connected to the selected RAID controller.
Controller Temperature	Indicates the temperature of the controller.
ROC Temperature	Current temperature of the RAID-on-a-chip (ROC) on the controller, in degrees Celsius.
Shield State Supported	Indicates whether the controller supports shield state.
Drive Security	Indicates the drive security (encryption) feature status on the controller.
Extended Virtual Drive Support	Indicates whether extended virtual drive is supported.
T10-PI	Indicates the status of the data protection feature on the controller.
Maintain Drive Fail History	Enables or disables the option to track bad physical drives through a reboot.
SMART Polling	Determines the interval, in seconds, at which the controller polls for drives reporting a Predictive Drive Failure. The default is 300 seconds. To change the value, use the + and – keys on the keypad. **NOTE** Some systems let you edit the numeric value directly, without using the + and – keys.
Stop Consistency Check on Error	Enables or disables the option of stopping a consistency check operation on a redundant virtual drive if a data inconsistency is detected.
JBOD Mode	Enables or disables the JBOD mode. **NOTE** When the JBOD mode is enabled, the drive comes up as a JBOD; otherwise, it comes up as an Unconfigured Good drive. **NOTE** When the JBOD mode is disabled, if one or more selected JBODs contain an operating system or a file system, a warning message appears indicating that the listed JBOD drives have an operating system or a file system and any data on them would be lost if you proceed. If you want to disable the JBOD mode, highlight **Confirm** and press the spacebar, then highlight **Yes** and press **Enter**. Else, highlight **No**.
Write Verify	Enables or disables the write verify feature during controller cache flush. This feature verifies if the data was written correctly to the cache before flushing the cache.
Drive Detection Type	Drives tend to develop media errors over time, which can slow down performance of the drive as well as the system as a whole. The firmware attempts to detect drives that consistently perform poorly. The Drive Detection Type options available here are **High Latency**, **Aggressive**, and **Default**. Depending on your requirement, use these options to set appropriate controller properties.
Drive Corrective Action	Drives tend to develop media errors over time, which can slow down the performance of the drive as well as the system as a whole. If a drive has certain amount of affected media leading to consistently poor I/O latency, then the firmware fails that particular drive, so that the drive rebuild/copyback process can start on that drive. The firmware also logs the appropriate events to alert the user. .

Table 27 Advanced Controller Properties (Continued)

Property	Description
Drive Error Threshold	The Drive Error Threshold options available here are: ■ Every 8 hours. ■ Every 1 hour. ■ Every 15 minutes. ■ Every 5 minutes.
Large IO Support	Enables or disables the large I/O support feature. By default, large I/O support is disabled. A reboot is required if this property is changed. When this property is changed,`The controller property change has been performed successfully. Reboot the machine for the change to take effect` message is displayed.
Coercion Mode	Enables you to set the coercion mode. The available options are **None**, **128 MB**, and **1 GB**.

5.6.2.1 Setting Cache and Memory Properties

The following dialog appears when you select **Cache and Memory** from the **Advanced Controller Properties** dialog.

Figure 86 Cache and Memory

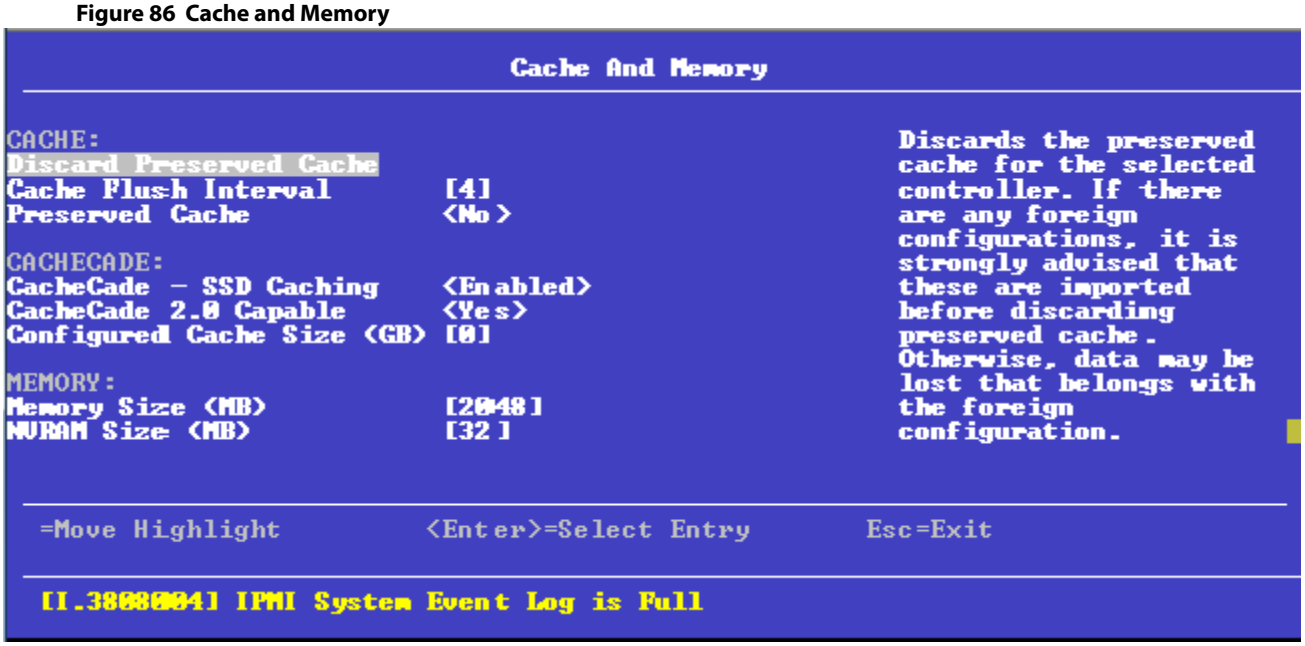

Follow these steps to set cache and memory properties:

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

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

2. To change the interval, in seconds, at which the contents of the onboard data cache are flushed, highlight **Cache Flush Interval** and press Enter. Specify a numeric value and press Enter.
3. If you want the controller to preserve cache because of missing or offline virtual drives (the cache is preserved until the virtual drive is imported or the cache is discarded), highlight **Preserved Cache**, and press Enter. Select either **Yes** or **No** and press Enter.

4. Highlight **Apply Changes** and press Enter.

   The new settings are saved in the controller properties.

5.6.2.2 Running a Patrol Read

The following dialog appears when you select **Patrol Read** from the **Advanced Controller Properties** dialog.

Figure 87 Patrol Read

A patrol read operation scans and resolves potential problems on configured physical drives.

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

Follow these steps to set patrol read properties:

NOTE You can only view the properties/options supported by your controller.

1. To select a mode for the patrol read operation, highlight **Mode** and press **Enter**. Select any of the following modes and press **Enter**.
   - **Auto**: Patrol read runs continuously on the controller based on a schedule. You do not need to start it manually.
   - **Manual**: Patrol read can be started or stopped manually.
   - **Disabled**: Patrol read does not run.
2. To specify a rate for the percentage of system resources dedicated to perform a patrol read operation on configured drives, highlight **Rate**, specify a rate as a numeric value and press **Enter**.

   100 is the maximum numeric value that you can enter as the rate.
3. To select a patrol read setting for unconfigured space, highlight **Setting for Unconfigured Space**, and press **Enter**. Select either **Enabled** or **Disabled** and press **Enter**.
4. Highlight **Apply Changes** and press **Enter**.

   The new settings are saved in the controller properties.

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

1. Highlight **Start** in the **Patrol Read** dialog and press **Enter**.
2. A message box appears stating that the operation has been successful. Click **OK** to return to the **Patrol Read** dialog.

   Suspend and **Stop** are now active.

5.6.2.3 Changing Power Save Settings

The following dialog appears when you select **Power Save Settings** from the **Advanced Controller Properties** dialog.

Figure 88 Power Save Settings

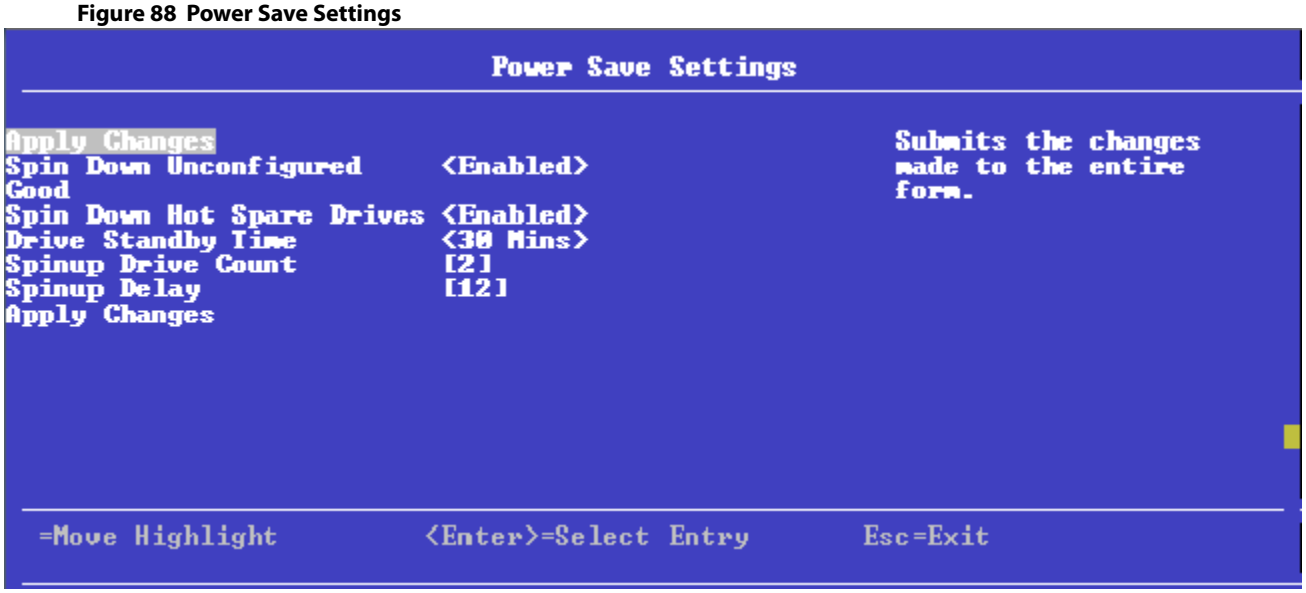

The above dialog lets you choose if you want unconfigured drives, hot spares, and configured drives to enter the power-save mode. When the unconfigured drives, hot spares, and configured drives are in power-save mode, they can be spun down.

Follow these steps to change the power-save settings:

NOTE You can only view the properties/options supported by your controller.

1. To enable or disable spinning down of unconfigured good drives, highlight **Spin Down Unconfigured Good** and press **Enter**. Select **Enable** or **Disable** and press **Enter**.
2. To enable or disable spinning down of hot spares, highlight **Spin Down Hot Spare Drives** and press Enter. Select **Enable** or **Disable** and press **Enter**.
3. To specify a drive's idle time, after which the drive goes into the power save mode, highlight **Drive Standby Time** and press **Enter**. Specify the time duration and press **Enter**.

   The drive standby time can be 30 minutes, 1 hour, 1.5 hours, or 2 hours through 24 hours.
4. To select the desired power-save mode, highlight **Power Save Mode** and press **Enter**. Select a mode (**None**, **Auto**, **Max**, and **Max without Cache**) and press **Enter**.
5. To specify the maximum number of drives that spin up simultaneously, highlight **Spinup Drive Count** and press **Enter**. Specify a numeric value and press **Enter**.
6. To control the interval (in seconds) between spin up of drives connected to the controller, highlight **Spinup Delay** and press **Enter**. Specify the time in seconds and press **Enter**.

   The delay prevents a drain on the system's power supply that would occur if all drives spun up at the same time.
7. If you do not want to schedule the drive active time, highlight **Do Not Schedule Drive Active Time** and press **Enter**.
8. To specify the Quality of Service window start time, highlight **Qos Window Start Time** and press **Enter**. Specify a start time and press **Enter**.
9. To specify the Quality of Service window end time, highlight **Qos Window End Time** and press **Enter**. Specify a end time and press **Enter**.
10. Highlight **Apply Changes** and press **Enter**.

    The new settings are saved in the controller properties.

5.6.2.4 Setting Emergency Spare Properties

The following dialog appears when you select **Spare** from the **Advanced Controller Properties** dialog.

Figure 89 Spare

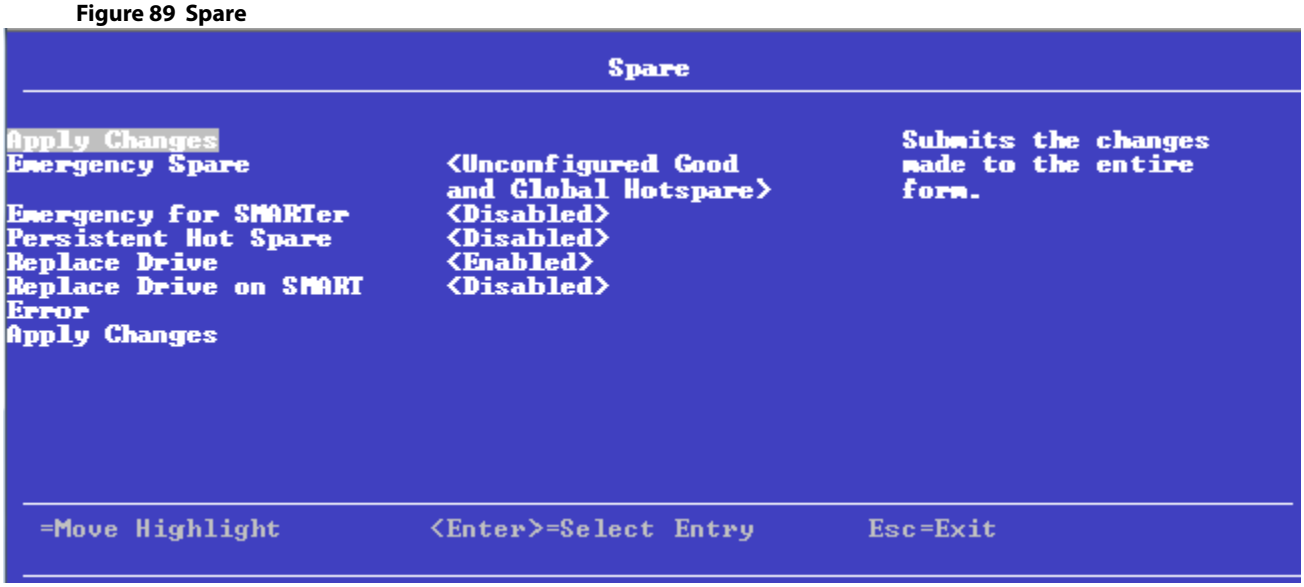

When a drive within a redundant virtual drive fails or is removed, the firmware automatically rebuilds the redundancy of the virtual drive by providing a emergency spare drive, even if no commissionable dedicated drive or global hot spare drive is present.

Follow these steps to set emergency spare properties:

1. To specify whether it is acceptable to commission otherwise incompatible global hot spare drive and/or unconfigured good drives as emergency hot spare drives, highlight **Emergency Spare** and press **Enter**. Select any of the following modes and press **Enter**.
   - **Global Hotspare**
   - **Unconfigured Good**
   - **Unconfigured Good and Global Hotspare**
   - **None**
2. To specify whether it is acceptable to commission emergency hot spare drives for PFA events, highlight **Emergency for SMARTer** and press **Enter**. Select an option (**Enabled** or **Disabled**) and press **Enter**.
3. To enable or disable the ability to have drive slots in the system backplane or in a storage enclosure dedicated as hot spare slots, highlight **Persistent Hot Spare** and press **Enter**. Select either **Enabled** or **Disabled** and press **Enter**.

   If enabled, replacement of a hot spare drive in the same slot automatically configures the drive as a hot spare.
4. To enable or disable the option to copy data back from a hot spare drive to a physical drive, highlight **Replace Drive** and press **Enter**. Select either **Enabled** or **Disabled** and press **Enter**.
5. To enable or disable the option to start a Drive Replace operation, if a Self-Monitoring Analysis and Report Technology (SMART) error is detected on a physical drive, highlight **Replace Drive on SMART Error** and press **Enter**. Select either **Enabled** or **Disabled** and press **Enter**.
6. Highlight **Apply Changes** and press **Enter**.

   The new settings are saved in the controller properties.

5.6.2.5 Changing Task Rates

The following dialog appears when you select **Task Rates** from the **Advanced Controller Properties** dialog.

Figure 90 Task Rates

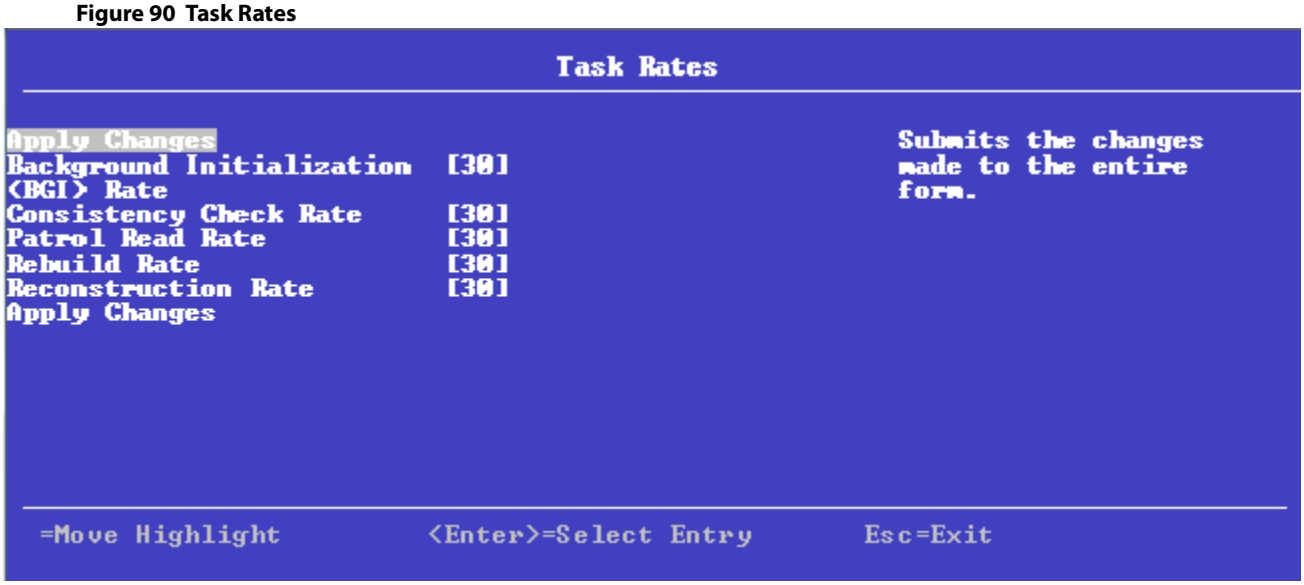

You can change the Rebuild rate and other task rates for a controller in the above dialog.

Follow these steps to change the task rates:

NOTE You can only view the properties/options supported by your controller.

1. To change the percentage of system resources dedicated to performing a BGI on a redundant virtual drive, highlight **Background Initialization <BGI> Rate** and press Enter. Specify a number from 0 to 100 and press Enter.

   The BGI rate is the percentage of the compute cycles dedicated to running a background initialization of drives on this controller. You can configure the BGI rate between 0 percent and100 percent. At 0 percent, the initialization operation runs only if the firmware is not doing anything else. At 100 percent, the initialization operation has a higher priority than I/O requests from the operating system. For best performance, use an initialization rate of approximately 30 percent.

2. To specify a rate for the percentage of system resources dedicated to performing a consistency check operation on a redundant virtual drive, highlight **Consistency Check Rate**, and press Enter. Specify a number from 0 to 100 and press Enter.

   The consistency check rate is the percentage of the compute cycles dedicated to running a consistency check on drives on this controller. You can configure the consistency check rate between 0 percent and 100 percent. At 0 percent, the consistency check operation runs only if the firmware is not doing anything else. At 100 percent, the consistency check operation has a higher priority than I/O requests from the operating system. For best performance, use a consistency check rate of approximately 30 percent.

3. To specify a rate for the percentage of system resources dedicated to performing a patrol read operation on configured physical drives, highlight **Patrol Read Rate** and press Enter. Specify a number from 0 to 100 and press Enter.

   The patrol read rate is the percentage of the compute cycles dedicated to running a patrol read on drives on this controller. You can configure the patrol read rate between 0 percent and 100 percent. At 0 percent, the patrol read runs only if the firmware is not doing anything else. At 100 percent, the patrol read has a higher priority than I/O requests from the operating system. For best performance, use a patrol read rate of approximately 30 percent.

4. To specify a rate for the percentage of system resources dedicated to rebuilding data on a new drive after a storage configuration drive has failed, highlight **Rebuild Rate** and press Enter. Specify a number from 0 to 100 and press Enter.

   The rebuild rate is the percentage of the compute cycles dedicated to rebuilding failed drives in virtual drives on this controller. You can configure the rebuild rate between 0 percent and 100 percent. At 0 percent, the rebuild runs only if the firmware is not doing anything else. At 100 percent, the rebuild operation has a higher priority than I/O requests from the operating system. For best performance, use a rebuild rate of approximately 30 percent.

5. To specify a rate for the percentage of system resources dedicated to performing a RAID Level Migration (RLM) or an Online Capacity Expansion (OCE) on a virtual drive, highlight **Reconstruction Rate** and press Enter. Specify a number from 0 to 100 and press Enter.

   The reconstruction rate is the percentage of the compute cycles dedicated to reconstructing data on drives on this controller. You can configure the reconstruction rate between 0 percent and100 percent. At 0 percent, the reconstruction operation runs only if the firmware is not doing anything else. At 100 percent, the reconstruction operation has a higher priority than I/O requests from the operating system. For best performance, use a reconstruction rate of approximately 30 percent.

6. Highlight **Apply Changes** and press Enter.

   The new settings are saved in the controller properties.

5.6.2.6 Upgrading the Firmware

The following dialog appears when you select **Update Firmware** from the **Dashboard View**. For a list of limitations, see Online Firmware Upgrade and Downgrade.

Figure 91 Controller Firmware Update

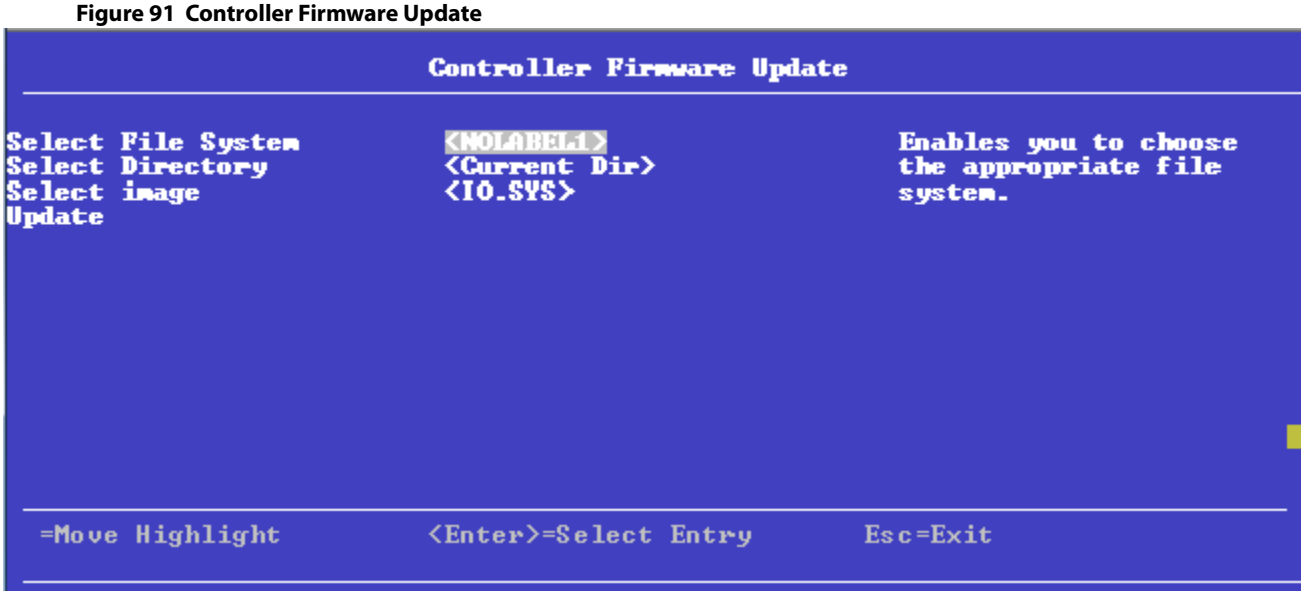

Follow these steps to upgrade the firmware:

1. To specify the file system where the `.rom` update file resides, highlight **Select File System** and press **Enter**. Select the file system and press **Enter**.
2. To specify the directory where the `.rom` file resides, highlight **Select Directory** and press **Enter**. Browse to the required the directory and press **Enter**.

   The current directory is normally highlighted. You can browse to only one level higher or one level lower.
3. To specify the `.rom` file, highlight **Select Image** and press **Enter**. Select the `.rom` file and press **Enter**.
4. Highlight **Update** and press **Enter**.

   The following Warning dialog appears.

Figure 92 Firmware Update Warning

5. Highlight the **Confirm** check box and press the spacebar to select the check box.
6. Click **Yes** to continue with the firmware update.

   After the controller is successfully updated with the new firmware code, a message box appears stating the same. Highlight **OK** and click **Enter** in the message box to return to the **Controller Management** dialog.

5.7 Managing Virtual Drives

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

Figure 93 Virtual Drive Management

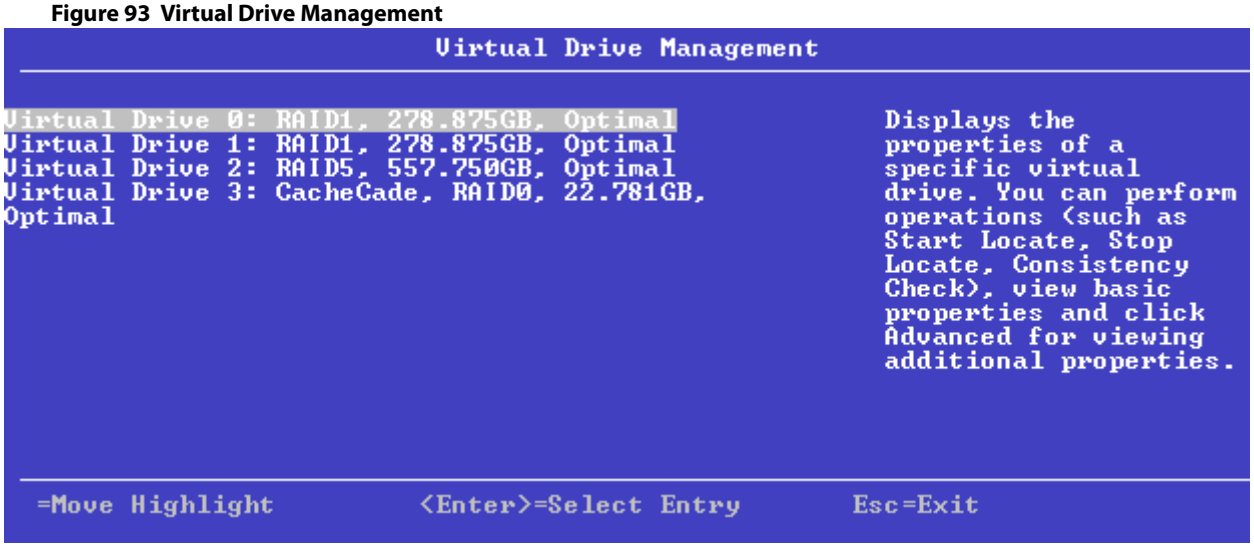

The menu lists all the virtual drives that currently exist on the controller. Highlight the virtual drive you want to manage and press Enter. The following dialog appears.

Figure 94 Virtual Drive Management

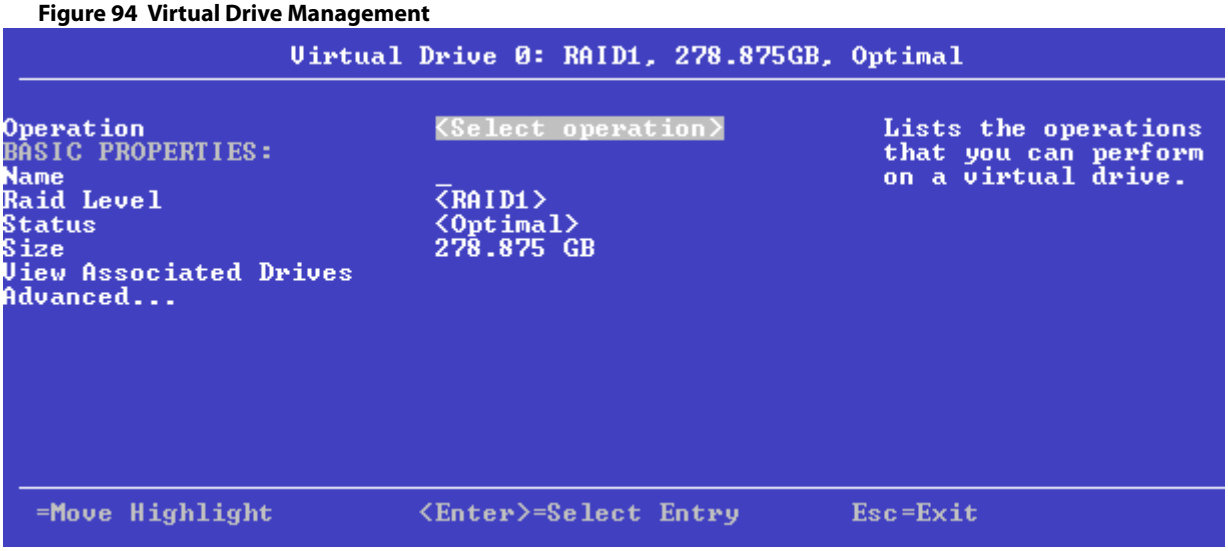

This dialog lists the following basic virtual drive properties.

Table 28 Basic Virtual Drive Properties

Property	Description
Name	The name assigned to the virtual drive. To assign a name or to change the name, highlight the field, press Enter, and type the new name in the popup window.
RAID Level	The RAID level of the virtual drive.
Status	The current status of the virtual drive.
Size	The capacity of the virtual drive, in MB or GB. **NOTE** Virtual drive size of floating data type up to three decimal places is supported. Some of the screens in this chapter may not show this feature.

For information on how to perform virtual drive operations, see Selecting Virtual Drive Operations.

For information on how to view the physical drives associated with the virtual drive, see Viewing Associated Drives.

For information on how to view and change advanced virtual drive settings, see Viewing and Managing Virtual Drive Properties and Options.

5.7.1 Selecting Virtual Drive Operations

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

Figure 95 Virtual Drive Operations Menu

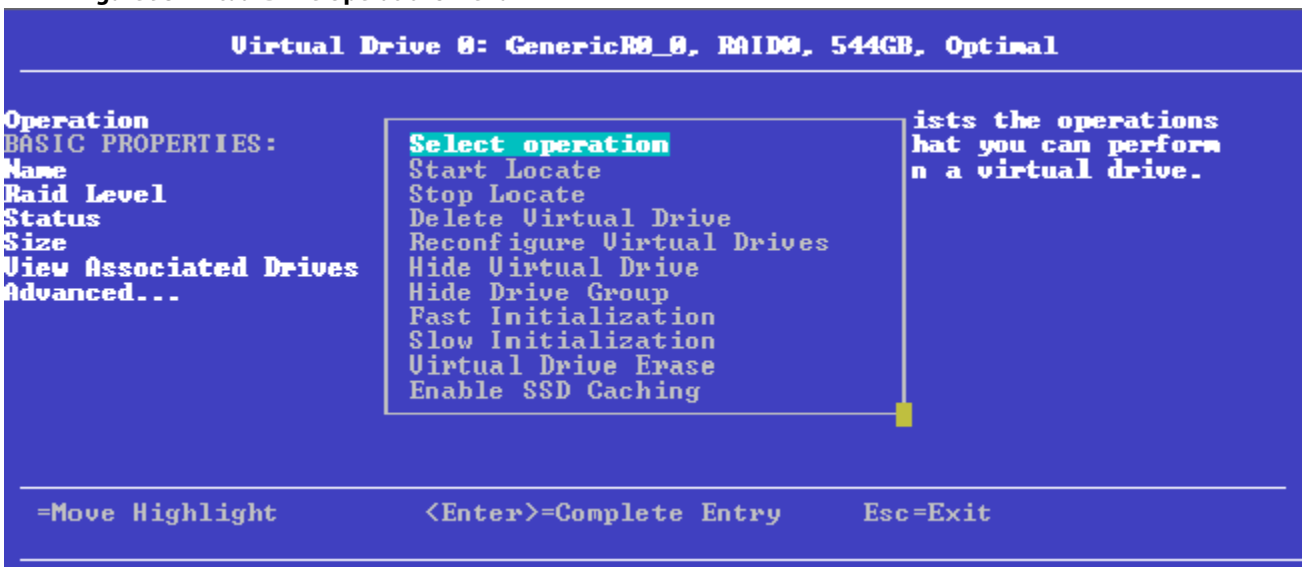

Other options, such as **Enable/Disable SSD Caching**, **Secure Virtual Drive**, **Check Consistency**, and **Expand Virtual Drive**, might also appear, depending on the current configuration of the system.

Highlight the operation you want to select and press **Enter**. Then highlight the word **Go** that appears beneath **Operation** and press **Enter** to start the operation for the currently selected virtual drive.

The following sections explain how to run the operations.

5.7.1.1 Locating Physical Drives in a Virtual Drive

To locate the physical drives in a virtual drive by flashing their LEDs, perform these steps:

1. Highlight **Start Locate** on the popup menu and press **Enter**.
2. Highlight the word **Go** that appears beneath **Operation** and press **Enter**.
   A Success message appears.
3. Highlight **OK** and press **Enter** to return to the **Virtual Drive** dialog.
   The LEDs on the physical drives start flashing, if the drive firmware supports this feature.
4. Observe the location of the drives with the flashing LEDs.
5. To stop the LEDs from flashing, access the popup menu again, highlight **Stop Locate**, and press **Enter**.
6. Highlight the word **Go** that appears beneath **Operation** and press **Enter**.
   A Success message appears.
7. Highlight **OK** and press **Enter** to return to the **Virtual Drive** dialog.
   The LEDs on the physical drives stop flashing.

5.7.1.2 Deleting a Virtual Drive

CAUTION All data on a virtual drive is lost when you delete it. Back up data you want to keep before you delete a virtual drive.

The delete virtual drive action is performed on the currently selected virtual drive. To select a different virtual drive for deletion, press Esc to return to the **Virtual Drive Selection** dialog and select the virtual drive.

To delete a virtual drive, perform these steps:

1. Highlight **Delete Virtual Drive** on the popup menu and press **Enter**.
2. Highlight the word **Go** that appears beneath **Operation** and press **Enter**.
   The **Delete Virtual Drive** warning message appears.
3. Highlight **Confirm** and press the spacebar to confirm the deletion, then highlight **Yes** and press **Enter**.
   The virtual drive is deleted.

NOTE The group initialization process is time-consuming when it is performed simultaneously on multiple drives when I/O transactions are in progress.

5.7.1.3 Hiding a Virtual Drive

To hide a virtual drive, perform these steps:

1. Highlight **Hide Virtual Drive** on the popup menu and press **Enter**.
2. Highlight the word **Go** that appears beneath **Operation** and press **Enter**.
   The **Hide Virtual Drive** warning message appears.
3. Highlight **Confirm** and press the spacebar to confirm the deletion, and then highlight **Yes** and press **Enter**.
   The virtual drive is hidden.

5.7.1.4 Unhiding a Virtual Drive

To unhide a virtual drive, perform these steps:

1. Highlight **Un-Hide Virtual Drive** on the popup menu and press **Enter**.
2. Highlight the word **Go** that appears beneath **Operation** and press **Enter**.
   The **Un-Hide Virtual Drive** warning message appears.
3. Highlight **Confirm** and press the spacebar to confirm the deletion, and then highlight **Yes** and press **Enter**.
   The virtual drive is unhidden.

5.7.1.5 Hiding a Drive Group

To hide a drive group to which the virtual drive is associated, perform these steps:

1. Highlight **Hide Drive Group** on the popup menu and press **Enter**.
2. Highlight the word **Go** that appears beneath **Operation** and press **Enter**.
   The **Hide Drive Group** warning message appears.
3. Highlight **Confirm** and press the spacebar to confirm the deletion, and then highlight **Yes** and press **Enter**.
   The drive group is hidden.

5.7.1.6 Unhiding a Drive Group

To unhide a drive group to which the virtual drive is associated, perform these steps:

1. Highlight **Un-Hide Drive Group** on the popup menu and press **Enter**.
2. Highlight the word **Go** that appears beneath **Operation** and press **Enter**.
   The **Un-Hide Drive Group** warning message appears.
3. Highlight **Confirm** and press the spacebar to confirm the deletion, and then highlight **Yes** and press **Enter**.
   The drive group is unhidden.

5.7.1.7 Reconfiguring a Virtual Drive

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

To reconfigure a virtual drive, perform these step:

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

Figure 96 Reconfigure Virtual Drives

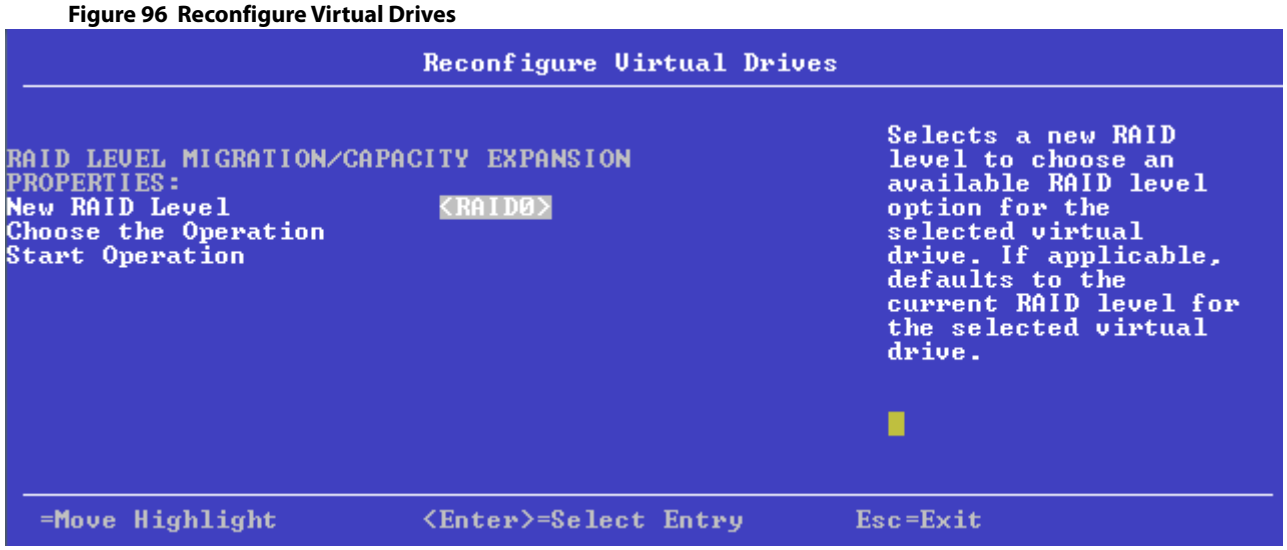

3. To change the RAID level of the selected virtual drive, highlight **New RAID Level** and press **Enter**.
4. Select a RAID level from the popup menu.
5. Depending on the source and the target RAID levels, you can either add drives or remove drives. Highlight **Choose the Operation** and press **Enter**.
6. Choose either **Add Drives** or **Remove Drives**

5.7.1.7.1 Adding Drives to a Configuration

Perform the following steps to add unconfigured drives to a configuration while reconfiguring a virtual drive.

1. If you select the **Add Drives** option and press **Enter**, the following dialog appears.

Figure 97 Select Drives – Add Drives

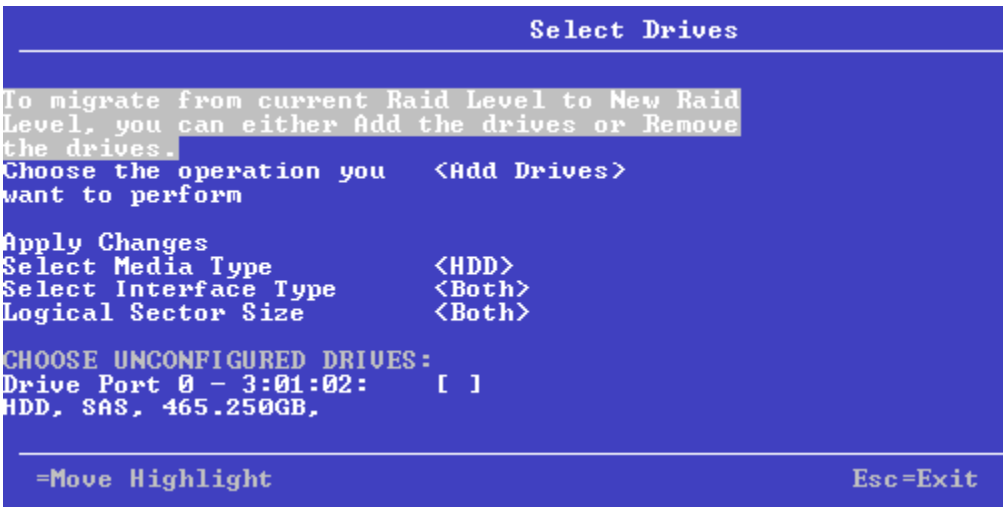

2. (Optional) To change the default **Select Media Type** value, highlight this field, press **Enter**, and select an option from the popup menu.

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

3. (Optional) To change the default **Select Interface Type** value, highlight this field, press **Enter**, and select an option from the popup menu.

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

4. To select unconfigured drives to add to the configuration, highlight the drives and press the spacebar. A small red arrow at the bottom of the dialog indicates you can scroll down to view more drives.

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

   Alternatively, use the **Check All** and **Uncheck All** options at the bottom of the list of drives to select or deselect all available drives.

   NOTE Be sure to select the number of drives required by the specified RAID level; otherwise, the HII utility displays an error message when you try to create the virtual drive. For example, RAID 1 virtual drives use exactly two drives and RAID 5 virtual drives use three or more drives. See for more information.

5. When you have selected the unconfigured drives to add, highlight **Apply Changes** and press **Enter**.

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

   The HII utility returns you to the **Reconfigure Virtual Drives** dialog.

5.7.1.7.2 Removing Drives from a Configuration

Perform the following steps to remove drives from a configuration while reconfiguring a virtual drive.

1. If you select the **Remove Drives** option and press **Enter**, the following dialog appears.

Figure 98 Select Drives – Remove Drives

2. To select the drives to remove from the configuration, highlight the drives and press the spacebar. A small red arrow at the bottom of the dialog indicates you can scroll down to view more drives.

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

Alternatively, use the **Check All** and **Uncheck All** options at the bottom of the list of drives to select or deselect all available drives.

3. When you have selected the drives to remove, highlight **Apply Changes** and press **Enter**.

The HII utility returns you to the **Reconfigure Virtual Drives** dialog.

5.7.1.8 Initializing a Virtual Drive

To initialize a virtual drive, perform these steps:

> **ATTENTION** All data on the virtual drive is lost when you initialize it. Before you start this operation, back up any data that you want to keep.

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

A fast initialization overwrites the first and last 8 MB of the virtual drive, clearing any boot records or partition information. A slow (full) initialization overwrites all blocks and destroys all data on the virtual drive.

2. Highlight the word **Go** that appears beneath **Operation** and press **Enter**.

The **Initialize Virtual Drive Warning** dialog appears.

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

A progress indicator shows the percentage completion of the initialization process. This indicator refreshes automatically.

5.7.1.9 Erasing a Virtual Drive

To erase data on a virtual drive, perform these steps:

ATTENTION All data on the virtual drive is lost when you erase it. Before you start this operation, back up any data that you want to keep.

NOTE After the data is erased, you have the option to keep the blank virtual drive, which you can use to store other data, or to delete the virtual drive completely.

1. Highlight **Virtual Drive Erase** on the popup menu and press **Enter**.
   Two additional fields appear.
2. Highlight **Erase Mode** and press **Enter**.
3. Select **Simple**, **Normal**, or **Thorough** from the popup menu.
   A Simple erase writes a pattern to the virtual drive in a single pass. The other erase modes make additional passes to erase the data more thoroughly.
4. (Optional) Highlight **Delete After Erase** and press the spacebar to select it.
5. Highlight **Go** and press **Enter**.
   The **Virtual Drive Erase** warning message appears.
6. Highlight **Confirm** and press the spacebar to confirm the operation, then highlight **Yes** and press **Enter**.
   A progress indicator shows the percentage completion of the operation. This indicator refreshes automatically. After the completion of the operation, the virtual drive is erased.

5.7.1.10 Securing a Virtual Drive

A Secure Virtual Drive operation enables security on a virtual drive. You can only disable the security by deleting the virtual drive. Perform these steps to secure a virtual drive.

1. Highlight **Secure Virtual Drive** on the popup menu and press **Enter**.
   The **Secure Virtual Drive** warning appears.
2. Highlight **Confirm** and press the spacebar to confirm the operation, then highlight **Yes** and press **Enter**.
   The virtual drive is secured.

5.7.1.11 Running a Consistency Check

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

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

   NOTE The **Check Consistency** option does not appear on the menu if the currently selected virtual drive is either RAID 0 or RAID 00 (nonredundant).

2. Highlight **Go** and press **Enter**.
   The **Consistency Check Success** dialog appears.
   As the message indicates, the consistency check is now running.
3. Highlight **OK** and press **Enter**.
   The Progress indicator in the dialog shows the percentage progress of the consistency check. To refresh the indicator, exit the dialog and re-enter it.
4. To stop or suspend the consistency check, highlight **Stop** or **Suspend** and press **Enter**.

5. To resume a suspended consistency check, highlight **Resume** and press **Enter**.

   A progress indicator shows the percentage completion of the operation. This indicator refreshes automatically.

For more information about consistency checks, see Scheduling a Consistency Check.

5.7.1.12 Expanding a Virtual Drive

Expanding a virtual drive means increasing its capacity. Existing data on the virtual drive is not impacted by the expansion. Follow these steps to expand the currently selected virtual drive.

1. Select **Expand Virtual Drive** from the popup menu.

   The **Expand Virtual Drive** dialog appears.

   The dialog shows the current capacity of the selected virtual drive, the available capacity that can be added to it, and the capacity of the expanded virtual drive, if all available capacity is added.
2. To change the amount of available capacity, highlight the **Enter a Percentage of Available Capacity** field and use the minus key (–) on the keypad to reduce percentage.

   NOTE Some systems permit you to enter numeric values directly, without using the **+** and **–** keys.
3. When you have set the capacity to the desired level, highlight **OK** and press **Enter**.

   The capacity of the virtual drive is expanded.

5.7.1.13 Disabling Protection on a Virtual Drive

To disable data protection on virtual drives, perform these steps:

1. Highlight **Disable Protection** on the popup menu and press **Enter**.
2. Highlight the word **Go** that appears beneath Operation and press **Enter**.

   Data protection is disabled on virtual drives.

5.7.2 Viewing Associated Drives

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

The dialog lists all the physical drives associated with the currently selected virtual drive. Follow these steps to view information about the associated drives.

1. To select a different virtual drive, highlight **Selected Virtual Drive**, press **Enter**, and select an entry from the popup menu.
2. Highlight one of the associated drives and press the spacebar to select it.
3. Highlight **View Drive Properties** and press **Enter**.

   The **View Drive Properties** window for the drive appears.
4. View the information on the **View Drive Properties** window.

   For more information, see Viewing Advanced Drive Properties.

5.7.3 Viewing and Managing Virtual Drive Properties and Options

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

NOTE The properties and options shown in the dialog apply to the currently selected virtual drive. To manage properties for a different virtual drive, press **Esc** until you return to the **Virtual Drive Selection** menu, select the desired virtual drive, and navigate back to this dialog.

Figure 99 Advanced Virtual Drive Properties 1

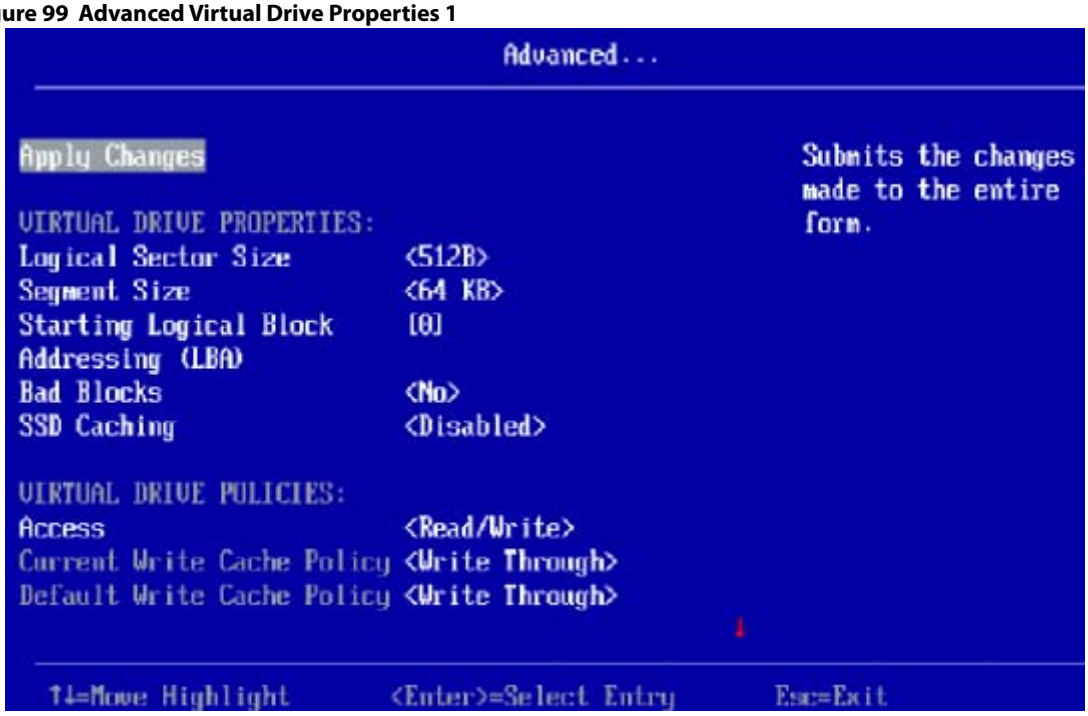

The small red arrow at the bottom of the dialog indicates that you can scroll down to view more virtual drive properties and virtual drive policies, as shown in the preceding figure.

Figure 100 Advanced Virtual Drive Properties 2

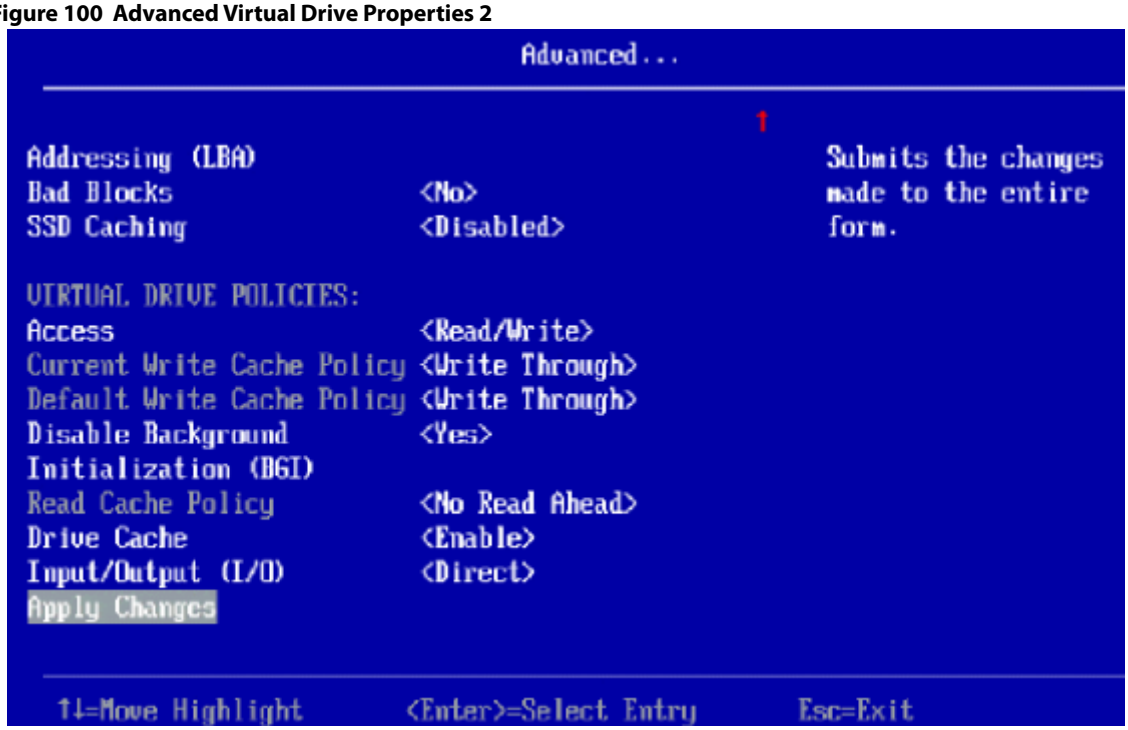

NOTE The red arrow appears when there is too much information to display in one dialog. The amount of information that can be displayed in one dialog depends on the capabilities of the HII browser.

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

Table 29 Virtual Drive Properties

Property	Description
Logical Sector Size	The logical sector size of this virtual drive. The possible options are **4 KB** and **512 B**.
Segment Size	The segment size used on this virtual drive.
Starting Logical Block	The address of the first location of a block of data stored on the virtual drive.
Addressing (LBA)	Indicates whether the virtual drive is secured.
Bad Blocks	Indicates whether the virtual drive has bad blocks.

Following the virtual drive properties listed in the dialog are virtual drive policies that you can select and change. To change any policy, highlight the field, press **Enter**, and select a value from the popup menu. When you finish changing policy settings, highlight **Apply Changes** at the top or the bottom of the selections and press **Enter**.

The following table describes the virtual drive policies.

Table 30 Virtual Drive Policies

Property	Description
Access	The access policy for the virtual drive. The options are **Read/Write**, **Read Only**, and **Blocked**.
Current Write Cache Policy	Displays the current write cache policy. The possible values are as follows: ■ **Write-Through** (**WThru**) The controller sends a data transfer completion signal to the host when the virtual drive has received all of the data and has completed the write transaction to the drive. ■ **Write-Back** (**WBack**) The controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a drive write transaction. Data is written to the virtual drive in accordance with policies set up by the controller. These policies include the amount of dirty and clean cache lines, the number of cache lines available, and the elapsed time from the last cache flush. ■ **Force Write Back**.
Default Write Cache Policy	Displays the default write cache policy of the virtual drive.
Disable Background Initialization (BGI)	Specifies whether background initialization is enabled or disabled. When BGI is enabled, the firmware runs the initialization process in the background. When BGI is disabled, the initialization process does not start automatically and does not run in the background.

Table 30 Virtual Drive Policies (Continued)

Property	Description
Read Cache Policy	Displays the read cache policy for the virtual drive. For any profile, if the drive is an SSD drive, the **No Read Ahead**and **Always Read Ahead**options are displayed. However, No Read Ahead is the **default** read policy. The possible options follow: ■ **Default** A virtual drive property that indicates whether the default read policy is **Always Read Ahead** or **No Read Ahead**. ■ **Always Read Ahead** - Permits the controller to read sequentially ahead of the requested data and allows the controller to store the additional data in the cache memory. Here, the controller anticipates that the data is required frequently. Even though Always Read Ahead policy speeds up the reads for sequential data, but little improvement is seen when accessing the random data. ■ **No Read Ahead** - Disables the Always Read Ahead capability of the controller.
Drive Cache	The disk cache policy for the virtual drive. The possible values are **Unchanged**, **Enable**, and **Disable**.
Input/Output (I/O)	The I/O policy for the virtual drive. The possible values are as follows: ■ **Direct**: Data reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. (The I/O policy applies to reads on a specific virtual drive. It does not affect the read- ahead cache.) ■ **Cached**: All reads are buffered in cache.
Emulation Type	Displays the current type of emulation policy. The available options are Default, Disable, and Force.

5.8 Managing Physical Drives

When you select **Drive Management** on the **Main Menu**, the **Drive Management Selection** dialog appears.

The menu lists all the physical drives that are connected to the controller. Highlight the drive you want to manage and press **Enter**. The following dialog appears.

Figure 101 Drive Management

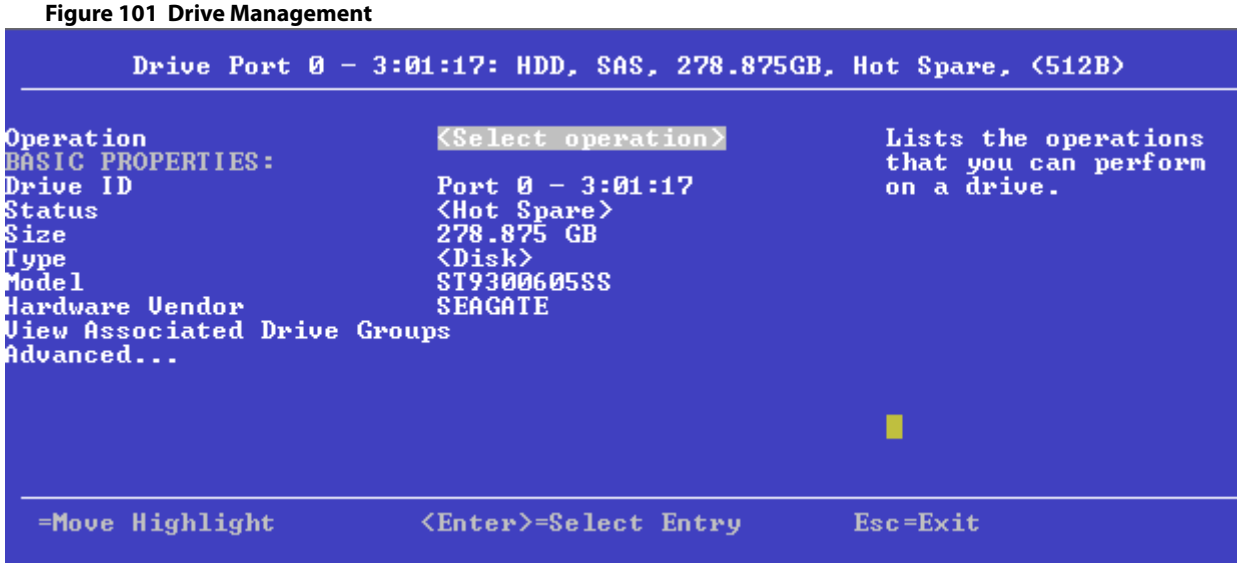

The preceding dialog lists the following basic drive properties for the selected drive:

Table 31 Basic Physical Drive Properties

Property	Description
Drive ID	The ID of the currently selected drive. The format of the Drive ID is **Connector Name + Wide Port Information: Enclosure Position: Slot Number**. If the drive is not installed in an enclosure, the format of the Drive ID is **Connector Name + Wide Port Information**.
Status	The status of the drive, such as **Online**, **Ready**, **Available**, or **Failed**.
Size	The drive capacity, in GB. A drive size of a floating data type up to three decimal places is supported. Some of the screens in this chapter may not show this feature.
Type	The device type of the drive, which is usually **Disk**.
Model	The model number of the drive.
Hardware Vendor	The hardware vendor of the drive.
Associated Virtual Drive	If this physical drive is currently used in a virtual drive, this field lists information about the virtual drive. Highlight this field and press **Enter** to view a popup window with additional information about the virtual drive.
Associated Drive Groups	If this physical drive is associated with drive groups, this field lists information about the drive groups. Highlight this field and press **Enter** to view a popup window with a list of associated drive groups. Highlight a drive from the list and press **Enter** to view additional information about the drive group, such as associated virtual drives, the capacity allocation, and the assigned dedicated hot spare drives, if any.

For information on performing drive operations, see Performing Drive Operations.

For information on viewing and changing drive settings and properties, see Viewing Advanced Drive Properties.

5.8.1 Performing Drive Operations

When you highlight the **Select operation** field, press **Enter** and a popup drive operations menu appears.

Figure 102 Select Drive Operations Menu

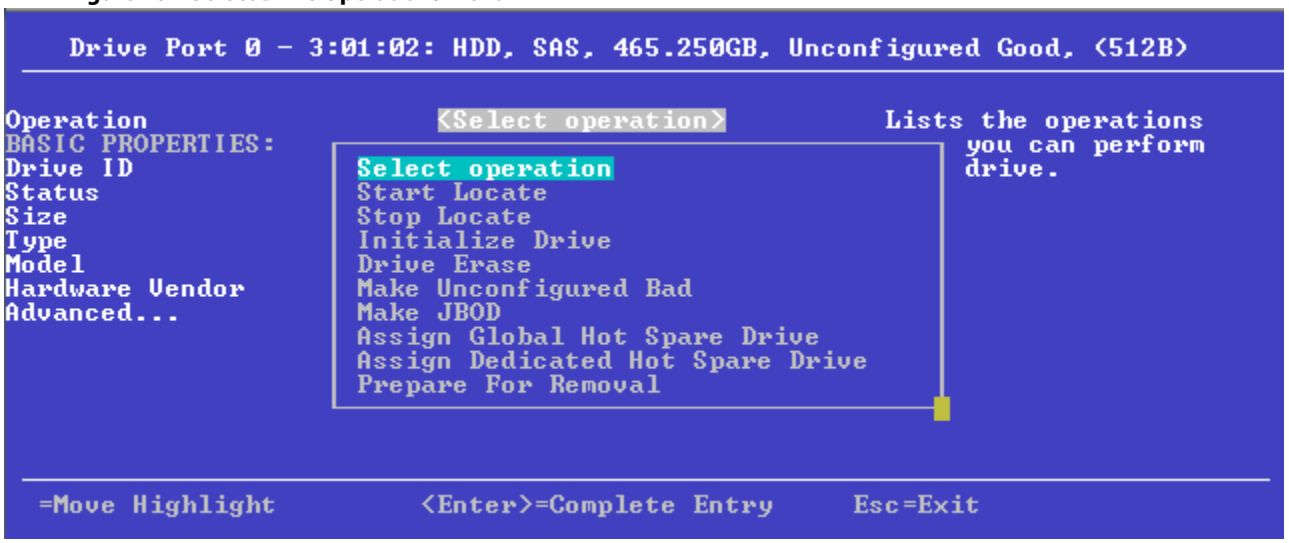

Start Locate and **Stop Locate** are the available options for any selected drive. The other menu options vary based on the status of the drive or on the selected personality mode, which can be **Online**, **Offline**, **JBOD**, **Unconfigured Good**, **Unconfigured Bad**, **Global Hot Spare**, and **Dedicated Hot Spare**. If your system is in JBOD personality mode, the **Make JBOD**, **Delete JBOD**, and other personality mode related options appear.

The drive operations run on the currently selected drive. To run an operation on a different drive, press **Esc** to return to the **Drive Selection** menu, highlight the drive you want to select, press **Enter** to select it, and return to this dialog.

5.8.1.1 Locating a Drive

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

1. Open the popup drive operations menu, highlight **Start Locate**, and press **Enter**.
2. Highlight **Go**, which appears beneath **Operation**, and press **Enter**.
   A success message appears.
3. Highlight **OK** on the success message and press **Enter**.
   The LED on the selected drive starts flashing, if the drive firmware supports this feature.
4. Observe the location of the drive with the flashing LED.
5. To stop the LED from flashing, highlight **Stop Locate** on the popup menu and press **Enter**.
6. Highlight **Go**, which appears beneath **Operation**, and press **Enter**.
   A success message appears.
7. Highlight **OK** on the success message and press **Enter**, to exit the message dialog.

5.8.1.2 Making a Drive Unconfigured Bad, Unconfigured Good, or JBOD

When you force a drive offline, it enters the *Unconfigured Bad* state.

When you power down a controller and insert a new physical drive, if the inserted drive does not contain valid DDF metadata, the drive status is Just a Bunch of Disks (*JBOD*) when you power the system again. A new drive in the JBOD drive state is exposed to the host operating system as a stand-alone drive. You cannot use the JBOD drives to create a

RAID configuration, because they do not have valid DDF records. You must first convert the drives into *Unconfigured Good*.

If a drive contains valid DDF metadata, its drive state is *Unconfigured Good*.

A drive must be in *Unconfigured Good* status before you can use it as a hot spare or use it as a member of a virtual drive. Follow these steps to change the status of an Unconfigured Bad, or Unconfigured Good, or JBOD drive.

1. Open the popup drive operations menu, either highlight the **Make Unconfigured Good**, **Make Unconfigured Bad**, depending on the personality mode that you have selected, **Make JBOD**, and press **Enter**.
2. Highlight **Go**, which appears beneath **Operation**, and press **Enter**.

   ATTENTION If you have selected the Make Unconfigured Good operation, and if the JBOD that you have selected has an operating system or a file system on it, a warning message appears indicating that the JBOD has an operating system or a file system and any data on it would be lost if you proceed with the conversion. If you want to proceed, highlight **Confirm** and press the spacebar, then highlight **Yes** and press `Enter`. Otherwise, highlight **No** and press `Enter` to return to the previous screen. To run this operation on a different drive, press Esc to return to the **Drive Selection** menu and select another drive.

   A message appears indicating that the operation was successful.
3. Highlight **OK** on the success message and press `Enter`.

   NOTE To refresh the status of the drive displayed in the dialog, exit back to the **Main Menu**, then re-enter the **Drive Management** dialog.

5.8.1.3 Enabling Security on JBOD

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

NOTE Enabling Security on JBOD can only be performed in RAID mode. If you are in JBOD mode, this option is not available.

1. Open the popup drive operations menu, highlight **Enable Security on JBOD** and press **Enter**.
2. Highlight **Go**, which appears beneath **Operation**, and press **Enter**.

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

5.8.1.4 Replacing a Drive

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

Follow these steps to replace a drive.

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

   The following dialog appears.

Figure 103 Replace Drive Window

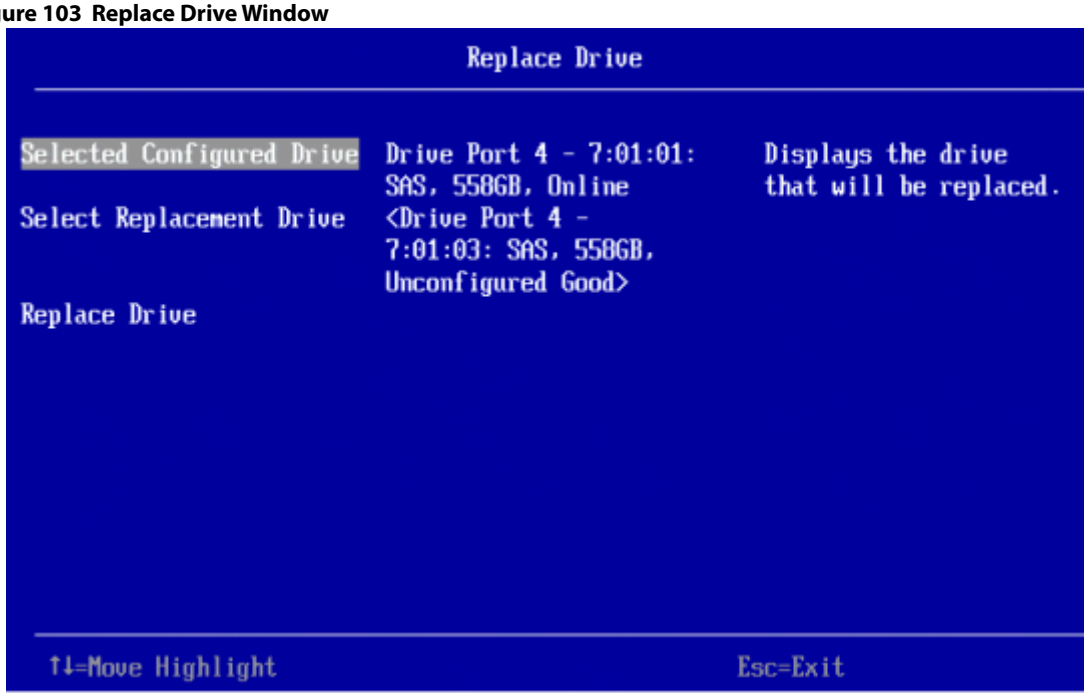

3. Highlight **Select Replacement Drive** and press **Enter**.

   A popup list of available replacement drives appears. In this example, only one replacement drive is available.
4. Select the replacement drive and press **Enter**.
5. Highlight **Replace Drive** and press **Enter**.

   A success message appears, and the replacement process begins as the data on the drive is rebuilt on the replacement drive.
6. Click **OK**.

   You are returned to the **Drive Management** menu. The status of the drive changes from **Online** to **Replacing**. You can perform other tasks in the HII utility while the replacement operation runs.

5.8.1.5 Placing a Drive Offline

Perform these steps to force a physical drive offline. If you perform this operation on a good drive that is part of a redundant virtual drive with a hot spare, the drive rebuilds to the hot spare drive. The drive you force offline goes into the Unconfigured Bad state.

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

   The Place Drive Offline message appears.
3. Highlight **Confirm**, and press the spacebar to confirm the operation.
4. Highlight **Yes**, and press **Enter**.

   The selected drive is forced offline.

5.8.1.6 Placing a Drive Online

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

1. Open the pop-up drive operations menu, highlight **Place Drive Online**, and press **Enter**.
2. Highlight **Go** and press **Enter**.

   The **Place Drive Online** warning appears.

ATTENTION Forcing a drive online that is part of a redundant array is *not* recommended.

3. Highlight **Confirm** and press the spacebar to confirm the operation.
4. Highlight **Yes** and press **Enter**.
   A message appears indicating that the action has been completed.
5. Highlight **Yes** and press **Enter** to return to the previous dialog.
   The drive is now online.

5.8.1.7 Marking a Drive Missing

Perform the following steps to mark a drive missing.

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

1. Open the popup drive operations menu, highlight **Mark Drive as Missing**, and press **Enter**.
2. Highlight **Go** and press **Enter**.
   A warning message appears.
3. Highlight **Confirm** and press the space bar to confirm the operation.
4. Highlight **Yes** and press **Enter**.
   A message appears indicating that the action has been completed.
5. Highlight **OK** and press **Enter** to return to the previous dialog.
   The drive is marked as missing.

5.8.1.8 Replacing a Missing Drive

Perform the following steps to replace the drive that is marked as missing.

1. Open the popup drive operations menu, highlight **Replace Missing Drive**, and press **Enter**.
2. Highlight **Go** and press **Enter**.
   A warning message appears.
3. Highlight **Confirm** and press the space bar to confirm the operation.
4. Highlight **Yes** and press **Enter**.
   A message appears indicating that the action has been completed.
5. Highlight **OK** and press **Enter** to return to the previous dialog.
   The drive that was marked as missing is replaced.

5.8.1.9 Assigning a Global Hot Spare Drive

Global hot spare drives provide protection to redundant virtual drives on the controller. If you select an Unconfigured Good drive, you have the option to assign it as a global hot spare drive. Perform these steps to assign a global hot spare.

1. Open the popup drive operations menu, highlight **Assign Hot Spare Drive**, and press **Enter**.
2. Highlight **Go**, which appears beneath **Operation**, and press **Enter**.
   The hot spare selection dialog appears.

3. Highlight **Assign Global Hot Spare Drive** and press **Enter**.

   The status of the selected drive changes to hot spare.

   NOTE To refresh the status of the drive displayed in the dialog, exit back to the **Main Menu**, then re-enter the **Drive Management** dialog.

5.8.1.10 Assigning a Dedicated Hot Spare Drive

Dedicated hot spare drives provide protection to one or more specified redundant virtual drives on the controller. If you select an Unconfigured Good drive, you have the option to assign it as a dedicated spare drive. Perform these steps to assign a dedicated hot spare.

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

   The following dialog appears.

Figure 104 Associate Virtual Drives to the Dedicated Hot Spare Drive

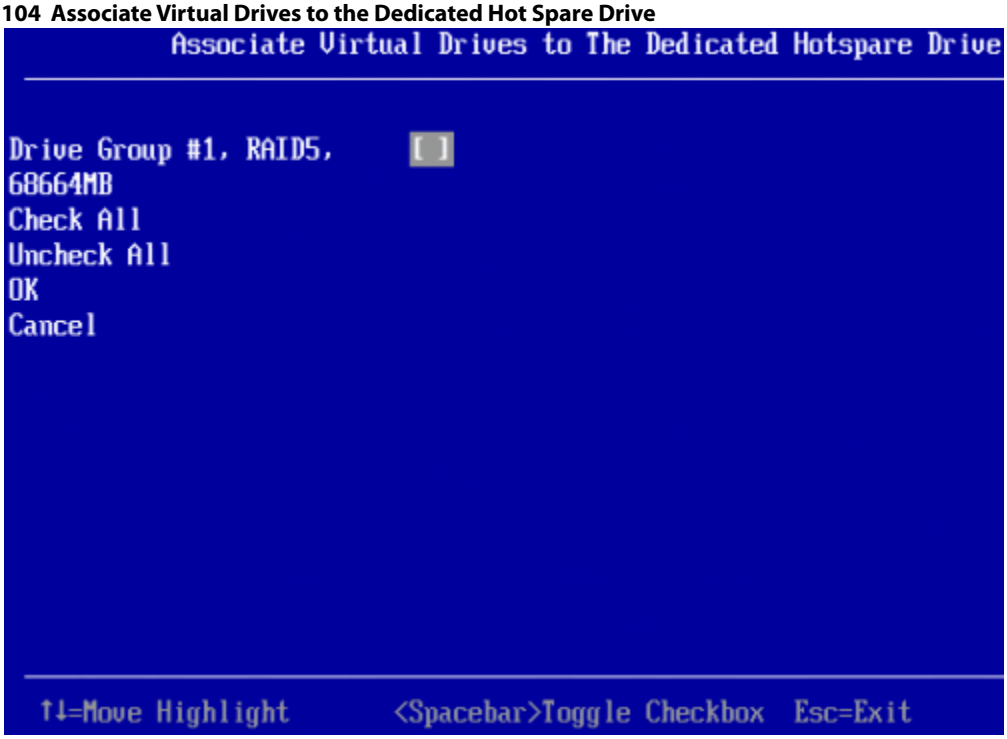

The preceding figure lists a single entry for each existing drive group. If you create a partial virtual drive on the same drive group, you can view a single entry with the cumulative size.

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

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

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

   When you return to the previous dialog, the status of the selected drive changes to hot spare.

   NOTE To refresh the status of the drive displayed in the dialog, exit back to the **Main Menu** and then re-enter the **Drive Management** dialog.

5.8.1.11 Unassigning a Hot Spare Drive

If the currently selected drive is a hot spare drive, you can unassign it and return it to Unconfigured Good status.

Perform these steps to unassign a hot spare drive.

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

1. Open the popup drive operations menu, highlight **Unassign Hot Spare Drive**, and press Enter.
2. Highlight **Go**, which appears beneath **Operation**, and press Enter.
   The **Unassign Hotspare Drive** warning message appears.
3. Highlight **Confirm** and press the spacebar to confirm the operation.
4. Highlight **Yes** and press Enter.
   A confirmation message appears.
5. Click **OK** to return to the **Drive Management** menu.
   The drive that was formerly a hot spare now appears as Unconfigured Good.

NOTE To refresh the status of the drive displayed in the dialog, exit back to the **Main Menu** and then re-enter the **Drive Management** dialog.

5.8.1.12 Initializing or Erasing a Drive

Follow these steps to initialize or erase the currently selected drive. An initialize operation fills the drive with zeroes. An erase operation initializes the drive with a pattern of zeros and ones.

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

1. Open the popup drive operations menu, highlight **Initialize Drive** or **Erase Drive**, and press Enter.
2. If you select **Drive Erase**, highlight the **Erase Mode** field and press Enter.
3. Select **Simple**, **Normal**, or **Thorough** from the popup menu and press Enter.
4. Highlight **Go** and press Enter.
   The **Initialize Virtual Drive** message appears. (The message is similar to that of erasing a drive.)
5. Highlight **Confirm** and press the spacebar to confirm the operation.
6. Highlight **Yes** and press Enter.
   A message appears indicating that the initialization or erase operation has started.
7. Highlight **Yes** and press Enter to return to the previous window.
   This dialog displays a progress indicator that shows the percentage completion of the operation. It also displays a `Stop` command, as shown in the following figure.

Figure 105 Initialize Progress Indicator

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

NOTE The progress indicator refreshes automatically.

5.8.1.13 Rebuilding a Drive

The manual Rebuild option is available only under certain conditions, as described here. If a hot spare drive is available, a rebuild starts automatically if a physical drive in a redundant array fails or is forced offline. If the Emergency Spare controller property is set to **Unconfigured Good** or **Unconfigured Good** and **Global Hotspare**, HII firmware automatically uses an Unconfigured Good drive to rebuild a failed or offline drive if no hot spares are available.

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

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

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

   A progress indicator shows the percentage completion of the rebuild operation. This indicator refreshes automatically, and the **Rebuild Drive Success** message appears.

5.8.1.14 Securely Erasing a Drive

Perform these steps to securely erase the currently selected FDE-capable drive. This option is available only if the controller supports security and if security is configured.

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

Perform these steps to securely erase an FDE-capable drive:

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

   The **Secure Erase** warning message appears.
3. Highlight **Confirm** and press the spacebar to confirm the operation.

4. Highlight **Yes** and press **Enter**.
   A message appears indicating that the secure erase operation has started.
5. Highlight **Yes** and press **Enter** to return to the previous dialog.
   This dialog now displays a progress bar and a `Stop` command.
6. To stop the secure erase process, highlight **Stop**, and press **Enter**.

NOTE A progress indicator shows the percentage completion of the operation. This indicator refreshes automatically.

5.8.1.15 Removing a Physical Drive

Perform these steps to remove a physical drive:

1. Open the popup drive operations menu, highlight **Prepare for Removal**, and press **Enter**.
2. Highlight **Go** and press **Enter**.
   A warning message appears.
3. Highlight **Confirm** and press the spacebar to confirm the operation.
4. Highlight **Yes** and press **Enter**.
   A message appears indicating that the action has been completed.
5. Highlight **Yes** and press **Enter** to return to the previous dialog.
   The drive is removed.

5.8.2 Viewing Advanced Drive Properties

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

Figure 106 Advanced Drive Properties

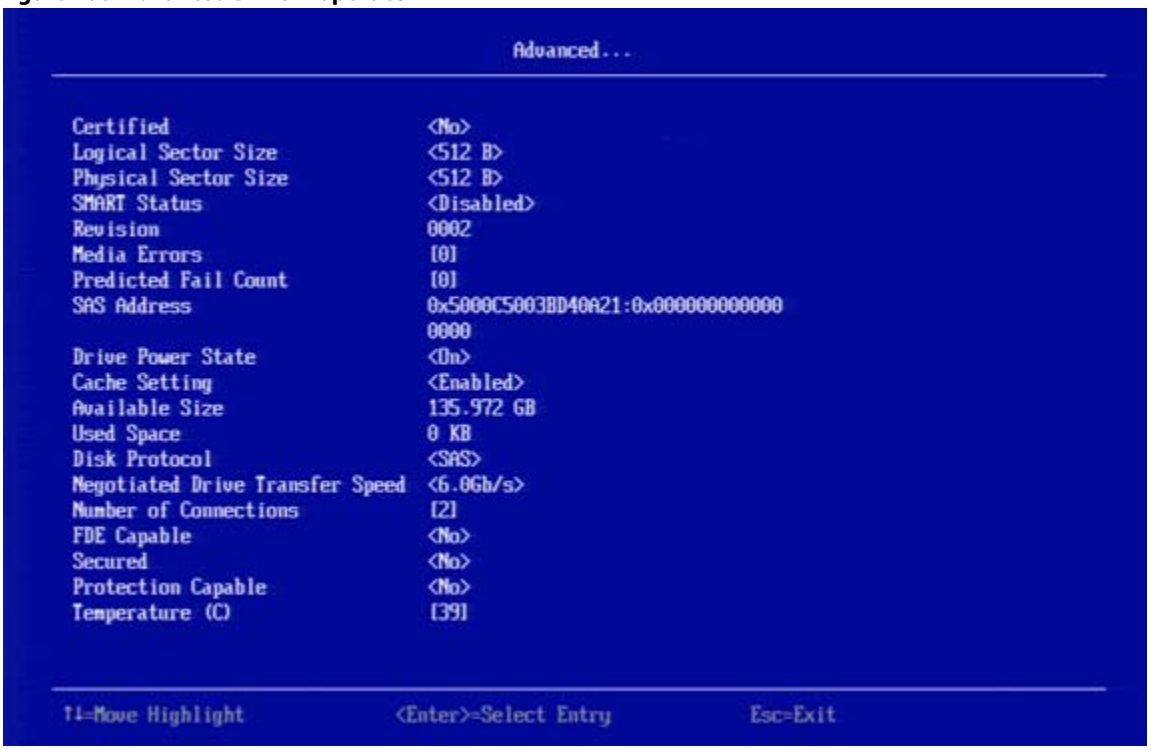

The following table describes all the entries listed on the **Advanced Drive Properties** dialog.

Table 32 Advanced Drive Properties

Property	Description
Certified	Indicates whether the selected drive is vendor-certified. In some configurations you can only use certified drives to create configurations.
Logical Sector Size	The logical sector size of this drive. The possible options are **4 KB** or **512 B**.
Physical Sector Size	The physical sector size of this drive. The possible options are **4 KB** or **512 B**.
SMART Status	Indicates whether the Self-Monitoring Analysis and Reporting Technology (SMART) feature is enabled or disabled on the drive. The SMART feature monitors the internal performance of all motors, heads, and drive electronics to detect predictable drive failures.
Revision	The firmware revision level of the drive.
Connected Port	The port on which the drive is connected.
Media Errors	The number of physical errors detected on the disk media.
Predicted Fail Count	A property indicating the number of errors that have been detected on the disk media.
SAS Address	The World Wide Name (WWN) for the drive.
Emergency Spare	Indicates whether the drive is commissioned as an emergency spare.
Driver Power State	Indicates whether the selected drive is powered on.
Commissioned Hot Spare	Indicates if any hot spare drive (dedicated, global, or emergency) has actually been commissioned.
Cache Setting	Indicates if the drive cache is enabled or disabled.
Available Size (GB)	The available size of the drive, in GB.
Used Space	The configured space of the drive, in GB.
Disk Protocol	Indicates whether the drive uses SAS or SATA protocol.
Negotiated Drive Transfer Speed	The negotiated link speed for data transfer to and from the drive.
Number of Connections	The number of connection on the drive. SAS drives have two ports.
FDE Capable	Indicates whether the drive is capable of encryption.
Secured	Indicates whether the drive is secured.
Protection Capable	Indicates whether the drive can be protected.
Temperature	Indicates the temperature for the physical drive in Celsius.

5.9 Managing Hardware Components

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

Figure 107 Hardware Components Menu

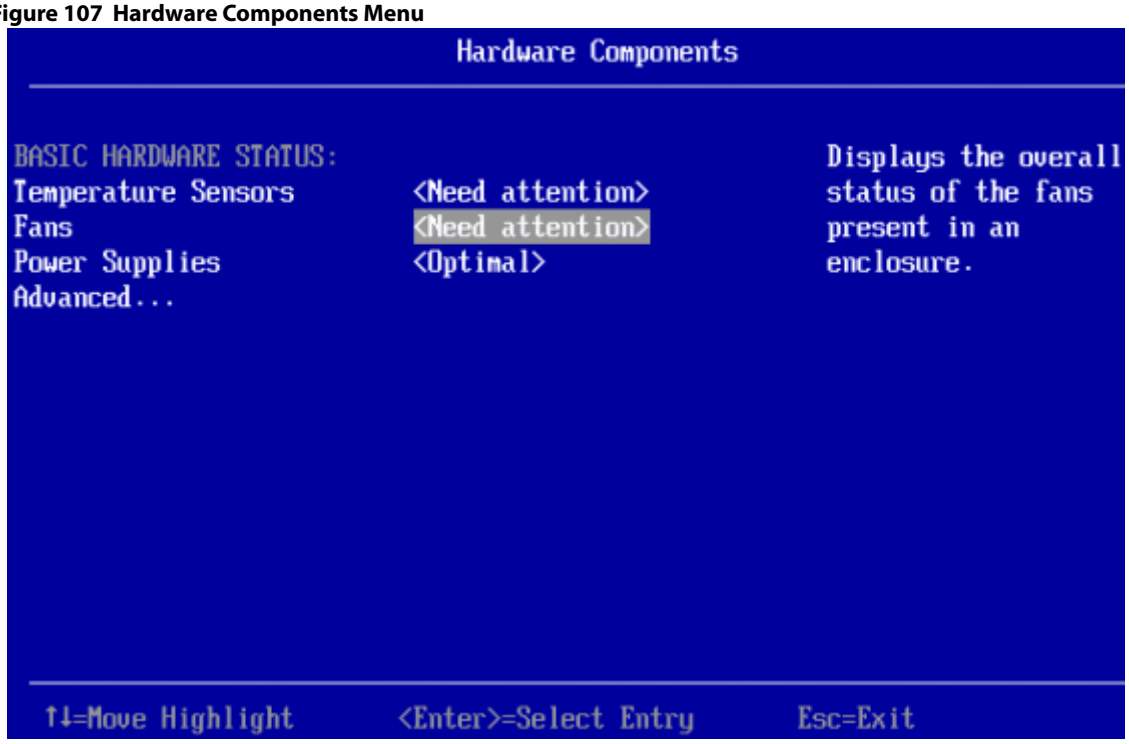

The preceding figure lists the status of the temperature sensors, fans, power supplies, and other hardware components (such as batteries) installed in the system.

Select **Advanced** and press Enter to view more detailed information about the installed hardware components. The following dialog appears.

Figure 108 Advanced Hardware Components Menu

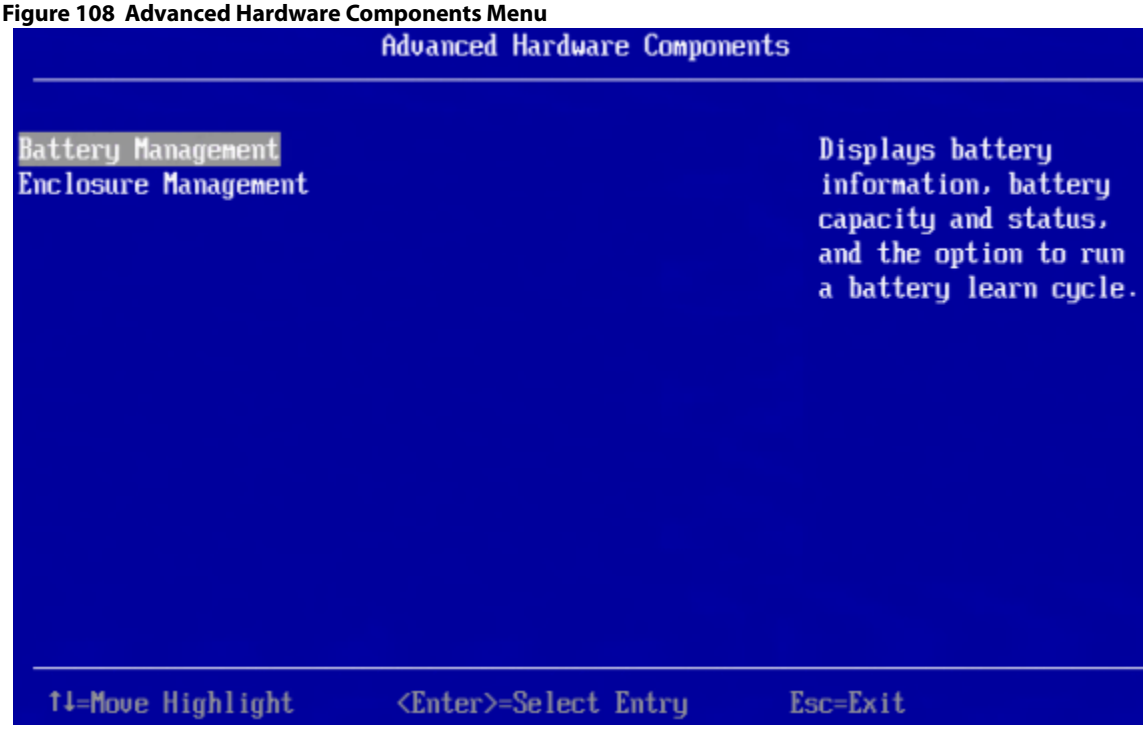

Select **Battery Management** or **Enclosure Management** to view more detailed information.

5.9.1 Managing Batteries

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

Figure 109 Battery Management

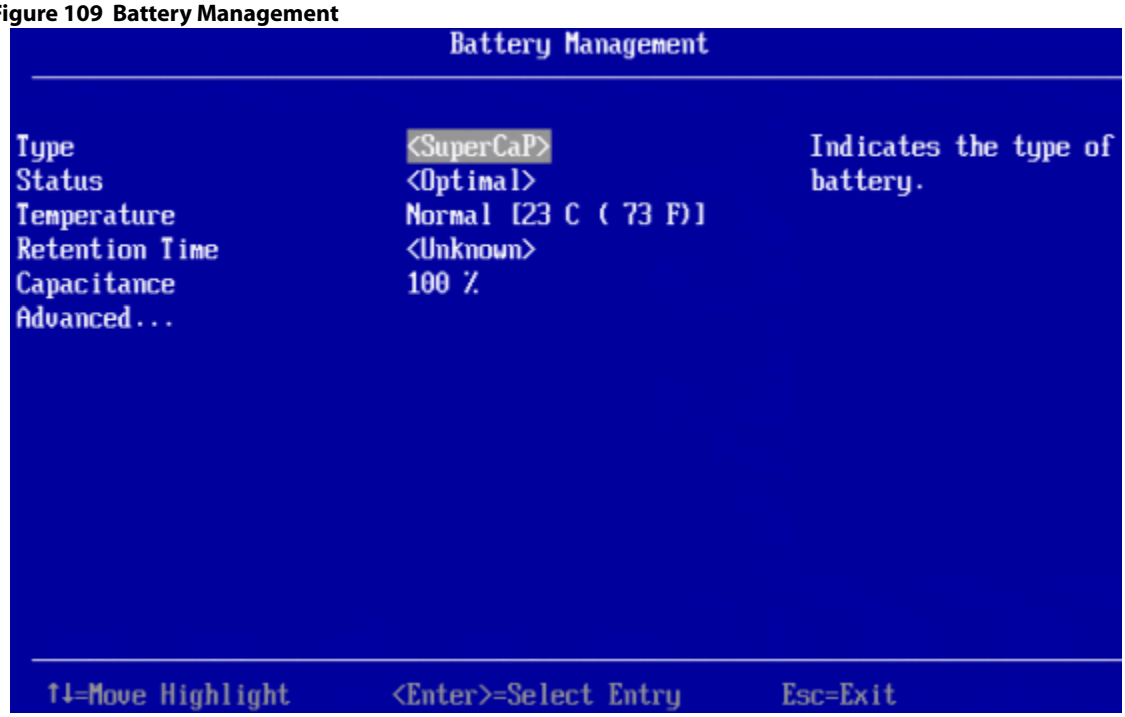

The following table describes the basic battery properties.

Table 33 Basic Battery Management Properties

Property	Description
Type	Type of the battery, such as Super Cap.
Status	Current status of the battery, such as **Optimal**. The battery status field has six states. If the battery operation is normal, the state is Optimal. ■ Optimal ■ Missing ■ Failed ■ Degraded ■ Degraded [Needs Attention] ■ Unknown
Temperature	Indicates the current temperature of the battery. Also indicates whether the current temperature of the battery is normal or high.
Retention Time	The number of hours the battery can support with the capacity it now has. The possible values are **48+ hours**, **Unknown**, or an exact number of hours between 1 and 48.
Capacitance	Available capacitance of the battery, stated as a percentage.

To view advanced battery properties, highlight **Advanced** and press Enter. The following dialog appears.

Figure 110 Advanced Battery Management

The small red arrow at the bottom of the dialog indicates that you can scroll down to view more Advanced Battery Management properties.

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

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

Table 34 Advanced Battery Management Properties

Property	Description
Start Manual Learn Cycle*	Highlight this field and press Enter to start a manual battery learn cycle.
Set Automatic Learn Cycle Properties*	Highlight this field and press Enter to set the properties for an automatic battery learn cycle.
Manufacturer	Manufacturer of the battery.
Serial Number	Serial number of the battery.
Date of Manufacture	Manufacturing date of the battery.
Firmware Version	Firmware version of the battery.
Status	Status of the battery. If the status is Learning, Degraded, or Failed, a reason is listed for the status.
Voltage	Voltage level of the battery, in mV. Also indicates if the current battery voltage is normal or low.
Current	Current of the battery, in mA.
Design Capacity	Theoretical capacity of the battery.
Full Capacity	Full charge capacity of the battery.
Remaining Capacity	Remaining capacity of the battery.
Auto-learn Mode	Indicates whether auto-learn mode is enabled or disabled. A learn cycle is a battery calibration operation that the controller performed periodically to determine the battery condition. This operation cannot be disabled.
Next Learn Cycle Time	Date and hour of the next scheduled learn cycle.

5.9.1.1 Setting Automatic Learn Cycle Properties

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

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

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

To generate an event as a reminder to start a learn cycle manually, highlight the field next to **Generate an event...**, and press the spacebar.

To enable or disable automatic learn cycle mode, highlight the field next to **Learn Cycle**, press Enter, and make a selection from the popup menu.

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

Use the **Apply**, **OK**, and **Cancel** fields at the bottom of the selections (not visible in this figure) to apply, confirm or cancel any changes to the learn cycle options.

5.9.2 Managing Enclosures

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

The **Enclosure Management** dialog shows the Vendor ID, Enclosure ID, Enclosure Model, Enclosure Location, Product Revision Level, Number of slots for the selected enclosure.

Figure 111 Enclosure Management

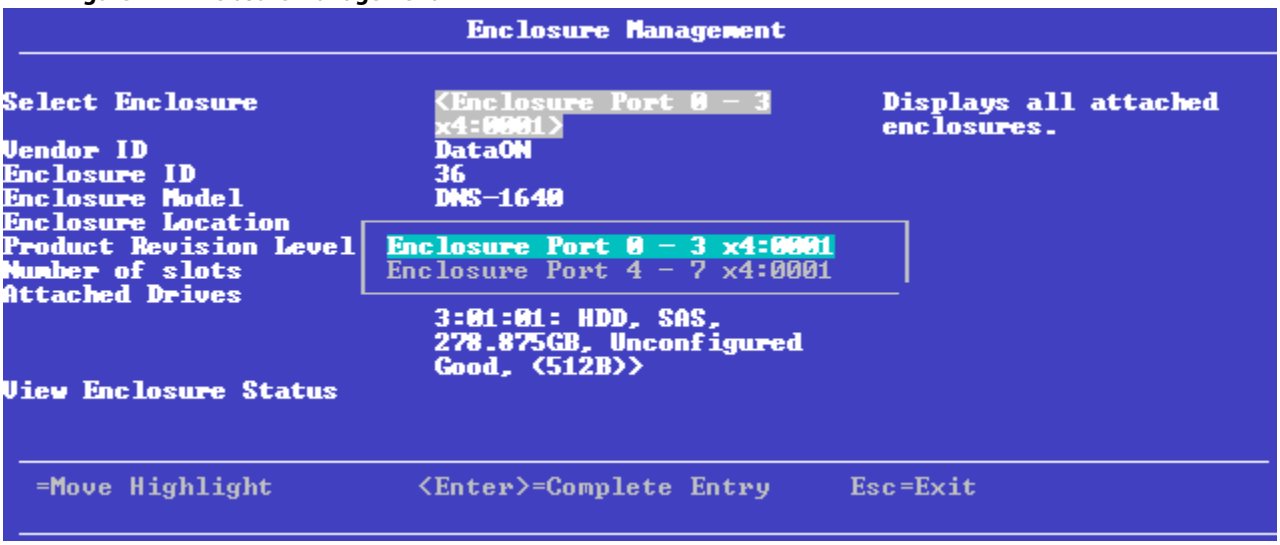

To select a different enclosure, highlight the **Select Enclosure** field, press Enter, and select the enclosure from the popup menu.

To view a popup menu of drives connected to the enclosure, highlight the **Attached Drives** field and press Enter.

To view more information about the enclosure status, highlight **View Enclosure Status** and press Enter. The following dialog appears.

Figure 112 View Enclosure Status

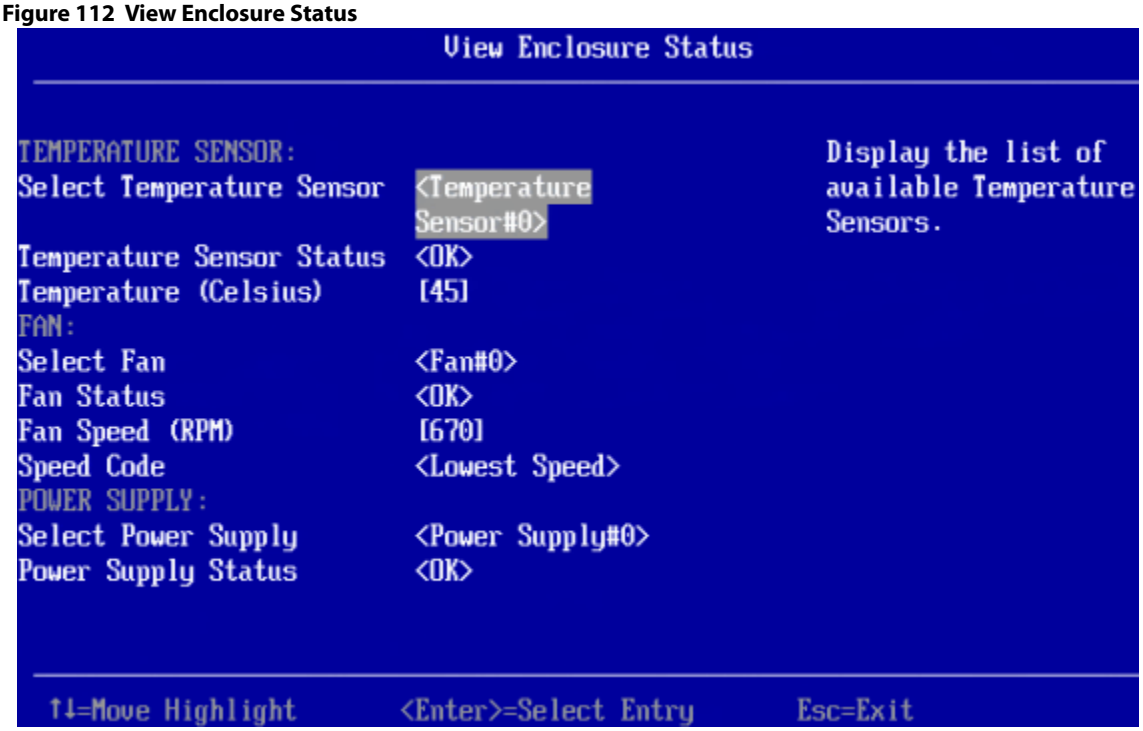

The **View Enclosure Status** dialog shows information about the temperature sensors, fans, and power supplies installed in the selected enclosure. To view a selectable popup menu of all of the installed sensors, fans, or power supplies, highlight the appropriate **Select** field, and press Enter.

Chapter 6: StorCLI

6.1 Overview

The Storage Command Line Interface (StorCLI) tool is the command line management software designed for the ThinkSystem product line. The StorCLI tool is a command line interface that is designed to be easy to use, consistent, and easy to script. This chapter provides information on how to install and use the StorCLI tool and explains the various features of the StorCLI tool.

NOTE The legacy commands are deprecated from this guide.

6.2 Support for MegaCLI Commands

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

6.3 Controllers Supported by the StorCLI Tool

The StorCLI tool is designed to work with the ThinkSystem product line. The StorCLI tool supports the all ThinkSystem 12Gb/s SAS RAID Controllers.

6.4 Operating System Installation

To check for the latest list of supported operating systems and to download the device drivers for those operating systems, go to http://support.lenovo.com.

6.5 StorCLI Tool Command Syntax

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

NOTE In large configurations, running two instances of the StorCLI tool in parallel (at the same time) is not recommended.

NOTE To get the output in JSON format, add J at the end of the command syntax. For example:

```
storcli /cx show <property1>|<property2> J
```

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

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

The StorCLI tool syntax uses the following general format:

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

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

Table 35 Object Identifiers in the StorCli Command Syntax

Object Identifier	Description
No object identifier specified	If no object identifier exists, the command is a system command.
/c*x*	This object identifier is for controller *x*.
/c*x*/v*x*	This object identifier is for a virtual drive *x* on controller *x*.
/c*x*/vall	This object identifier is for all virtual drives on controller *x*.
/c*x*/e*x*	This object identifier is for an enclosure x on controller *x*.
/c*x*/eall	This object identifier is for all enclosures on controller *x*.
/c*x*/f*x*	This object identifier is for a foreign configuration x on controller *x*.
/c*x*/fall	This object identifier is for all foreign configurations on controller *x*.
/c*x*/e*x*/s*x*	This object identifier is for the drive is slot *x* on enclosure *x* on controller *x*.
/c*x*/s*x*	This object identifier represents the drives that are directly attached to controller *x*.
/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 36 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.

Table 36 Verbs in the StorCli Command Syntax (Continued)

Verbs	Description
`stop`	This verb stops an operation that is in progress. A stopped process cannot be resumed.
`import`	This verb imports the foreign configuration into the drive.
`expand`	This verb expands the size of the virtual drive.
`insert`	This verb replaces the configured drive that is identified as missing, and starts an automatic rebuild.
`flasherase`	This verb erases the flash memory on the controller.
`transform`	This verb downgrades the firmware memory on the controller.
`restart`	This verb restarts the controller without a system reboot.
`apply`	This verb applies the activation Key to a WarpDrive card.

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

6.6 StorCLI (Storage Command Line Interface) Commands

StorCLI is a command line utility tool. StorCLI is not case sensitive. The order in which you specify the command options should be the same as in this document; otherwise, the commands may fail.

NOTE StorCLI does not support the Snapshot feature.

This section describes the commands supported by StorCLI.

6.6.1 System Commands

6.6.1.1 System Show Commands

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

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

The detailed description for each command follows.

storcli show

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

storcli show all

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

storcli show ctrlcount

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

storcli show help

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

storcli -v

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

6.6.2 Controller Commands

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

6.6.2.1 Show and Set Controller Properties Commands

Table 37 Controller Commands Quick Reference Table

Commands	Value Range	Description
`show <properties>`	See Table 38	Shows specific controller properties.
`set <properties>`	See Table 38	Sets controller properties.
`show`	`all`: Shows all properties of the virtual drive. `freespace`: Shows the free space in the controller. See Controller Show Commands.	Shows physical drive information.

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

NOTE You cannot set multiple properties with a single command.

storcli /*cx* show <*property*>

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

General example output:

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

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

```
storcli /cx show abortcconerror
storcli /cx show activityforlocate
storcli /cx show alarm
storcli /cx show backplane
storcli /cx show badblocks
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 DPM
storcli /cx show eccbucketleakrate
storcli /cx show eccbucketsize
storcli /cx show eghs
storcli /cx show failpdonsmarterror
storcli /cx show flushwriteverify
storcli /cx show jbod
storcli /cx show loadbalancemode
storcli /c0 show largeiosupport
storcli /cx show maintainpdfailhistory
storcli /cx show migraterate
storcli /cx show ncq
storcli /cx show patrolread|pr
storcli /cx show perfmode
storcli /cx show pi
storcli /cx show prcorrectunconfiguredareas
storcli /cx show prrate
storcli /cx show personality
storcli /cx show rebuildrate
storcli /cx show rehostinfo
storcli /cx show restorehotspare
storcli /cx show safeid
storcli /cx show sesmultipathcfg
storcli /cx show smartpollinterval
storcli /cx show spinupdelay
storcli /cx show spinupdrivecount
storcli /cx show SGPIOforce
```

```
storcli /cx show time
storcli /cx show usefdeonlyencrypt
storcli /cx show wbsupport
```

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  backplane [mode=<0-3>][expose=<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|None ]
storcli /cx set  copyback=<on|off> type=<smartssd|smarthdd|all>
storcli /cx set  directpdmapping=<on|off>
storcli /cx set DPM=<on|off>
storcli /cx set driveactivityled=<on|off>
storcli /cx set  dimmerswitch|ds=<on|off type=1|2|4>
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  foreignautoimport=<on|off>
storcli /cx set failpdonsmarterror=<on|off>
storcli /cx set flushwriteverify=<on|off>
```

```
storcli /cx set immediateio=<on|off>
storcli /cx set  jbod=<on|off>
storcli /cx set  loadbalancemode=<value>
storcli /cx set largeiosupport=on|off
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 /cx set  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 personality=RAID|JBOD
storcli /cx set personality behavior=JBOD/None
storcli /cx set personality behavior [sesmgmt=on/off] [securesed=on/off]
[multipath=on/off] [multiinit=on/off]
storcli /cx set  rebuildrate=<value>
storcli /cx set  restorehotspare=<on|off>
storcli /cx set sesmultipathcfg=<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  supportssdpatrolread=<on|off>
storcli /cx set  SGPIOforce=<on|off>
storcli /cx set sesmonitoring=[on|off]
storcli /cx set time=yyyymmdd hh:mm:ss|systemtime
storcli /cx set  termlog[=on|off|offthisboot]
storcli /cx set  usefdeonlyencrypt=<on|off>
storcli /cx set  usefdeonlyencrypt=<on|off>
```

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

Table 38 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 or disables drive activity, drive activity locates function for systems without SGPIO/SES capabilities.
alarm	on|off|silence silence: Silences the alarm.	Enables or disables alarm on critical errors.
batterywarning	on|off	Enables or disables battery warnings.
bgirate	0 to 100	Sets background initialization rate in percentage.
backplane mode	0: Use autodetect logic of backplanes, such as SGPIO and I^2C SEP using GPIO pins. 1: Disable autodetect SGPIO. 2: Disable I^2C SEP autodetect. 3: Disable both the autodetects.	Configures enclosure detection on a non-SES/expander backplane.
backplane expose	on|off	Enables or disables device drivers to expose enclosure devices; for example, expanders, SEPs.
cachebypass	on|off	Enables or disables the cache bypass performance improvement feature.
cacheflushint	0 to 255, default value 4	Sets cache flush interval in seconds.
ccrate	0 to 100	Sets consistency check rate in percentage.
coercion	0: No coercion 1: 128 MB 2: 1 GB	Sets drive capacity in coercion mode.
consistencycheck	See Consistency Check.	See Consistency Check.
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 or disables copy back for drive types.
dimmerswitch|ds	See Dimmer Switch Commands.	See Dimmer Switch Commands.
directpdmapping	on|off	Enables or disables direct physical drive mapping. When enclosures are used, this feature is disabled; otherwise it should be enabled.
DPM	on|off	Enables or disables drive performance monitoring
driveactivityled	on|off	Activate or deactivate the Drive Activity LED.
eccbucketleakrate	0 to 65535	Sets the leak rate of the single-bit bucket in minutes (one entry removed per leak-rate).
eccbucketsize	0 to 255	Sets the size of ECC single-bit-error bucket (logs event when full).

Table 38 Properties for Show and Set Commands (Continued)

Property Name	Set Command Range	Description
eghs state	on|off	Enables or disables the commissioning of otherwise incompatible global hot spare drives as Emergency Hot Spare (EHSP) drives.
eghs smarter	on|off	Enables or disables the commissioning of Emergency Hot Spare (EHSP) drives for Predictive Failure (PFA) events.
eghs eug	on|off	Enables or disables the commissioning of Unconfigured Good drives as Emergency Hot Spare (EHSP) drives.
foreignautoimport	on|off	Imports a foreign configuration automatically, at boot.
failpdonsmarterror	on|off	Enables or disables the *Fail PD on SMARTer* property.
flushwriteverify	on|off	Enables or disables the Write Verify feature. This feature verifies if the data was written correctly to the cache before flushing the controller cache.
immediateio	on|off	Enables or disables Immediate I/O transactions.
jbod	on|off	Enables or disables JBOD mode; by default, drives become system drives. Not supported by all controllers. **NOTE** If you try to disable the JBOD mode, and if any of the JBOD has an operating system/file system, then the StorCLI tool displays a warning message indicating that the JBOD has an operating system or a file system on it and prompts you to use the force option to proceed with the disable operation.
loadbalancemode	on|off	Enables or disables automatic load balancing between SAS phys or ports in a wide port configuration.
largeiosupport	on|off	Sets the current settings on the controller for large I/O support.
maintainpdfailhistory	on|off	Maintains the physical drive fail history.
migraterate	0 to 100	Sets data migration rate in percentage.
patrolread|pr	See Patrol Read.	See Patrol Read.
perfmode	0: Tuned to provide best IOPS, currently applicable to non-FastPath 1: Tuned to provide least latency, currently applicable to non-FastPath	Performance tuning setting for the controller.
pi	on|off	Enables or disables data protection on the controller.
pi import	on|off	Enables or disables import data protection drives on the controller.
prcorrectunconfiguredareas	on|off	Correct media errors during PR by writing 0s to unconfigured areas of the disk.
prrate	0 to 100	Sets the patrol read rate of the virtual drives in percentage.
rebuildrate	0 to 100	Sets the rebuild rate of the drive in percentage.
reconrate	0 to 100	Sets the reconstruction rate for a drive, as a percentage.

Table 38 Properties for Show and Set Commands (Continued)

Property Name	Set Command Range	Description
restorehotspare	on|off	Becomes a hot spare on insertion of a failed drive.
sesmonitoring	on|off	Enables or disables SES monitoring.
sesmultipathcfg	on|off ■ 0: Type of association based on LUN. ■ 1: Type of association based on the target port.	Sets the type of association for SES in a multipath configuration.
smartpollinterval	0 to 65535	Set the time for polling of SMART errors, in seconds.
spinupdrivecount	0 to 255	Sets the number of drives that are spun up at a time.
spinupdelay	0 to 255	Sets the spin-up delay between a group of drives or a set of drives, in seconds.
stoponerror	on|off	Stops the ThinkSystem BIOS during POST, if any errors are encountered.
termlog	on | off | offthisboot offthisboot: Disables the termlog flush to ONFI only for this boot. In the next boot, the termlog will be enabled.	Enables or disables the termlog to be flushed from DDR to ONFI. offthisboot - disables the termlog flush to ONFI only for this boot. In the next boot, the termlog is enabled.
supportssdpatrolread	on|off	Enables or disables patrol read for SSD drives.
SGPIOforce	on|off	Forces the SGPIO status per port only for four drives; affects HPC controllers.
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 or disables FDE drive-based encryption.

6.6.2.2 Controller Show Commands

The StorCLI supports the following show commands:

```
storcli /cx show
storcli /cx show all [logfile[=filename]]
storcli /cx show freespace
```

The detailed description for each command follows.

storcli /*cx* show

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

Input example:

```
storcli /c1 show
```

storcli /*cx* show all [logfile[=*filename*]]

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

If you use the `logfile` option in the command syntax, the logs are written to the specified file. If you do not specify a file name, then the logs are written to the `storsas.log` file. If you do not use the `logfile` option in the command syntax, the entire log output is printed to the console.

Do not use spaces in between file names.

Input examples:

```
storcli /c0 show all [logfile[=log.txt]]
storcli /c0 show all logfile = abc.txt
```

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

storcli /c*x* show freespace

This command shows the usable free space in the controller.

Input example:

```
storcli /c0 show freespace
```

6.6.2.3 Controller Debug Commands

The StorCLI tool supports the following debug commands:

Syntax

```
storcli /c  x set debug type = <value> option = <value> level = [<value in hex>]
```

This command enables the firmware debug variables.

Where:

- `/cx` – specifies the controller where *x* is the index of the controller.
- `type` – takes the value from 0 – 128, mapping each number to a particular debug variable in the firmware.
- `option` – takes the value from 0 – 4, where;
  - 0 – NA
  - 1– SET
  - 2 – CLEAR
  - 3 – CLEAR ALL
  - 4 – DEBUG DUMP
- `level` – supports multiple levels of debugging in the firmware.

Syntax

```
storcli /c  x set debug reset all
```

This command enables the firmware debug logs from the application

Where:

`/cx` - specifies the controller where *x* is the index of the controller.

NOTE The **debug type**, the **debug value**, and the **debug level** for the following debug commands are exclusively used by the Broadcom Technical Support team to provide technical support. For assistance with these debug commands, contact an Broadcom Technical Support representative.

6.6.2.4 Controller Background Tasks Operation Commands

6.6.2.4.1 Rebuild Rate

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

The detailed description for each command follows.

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

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

Input example:

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

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

storcli /*cx* show rebuildrate

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

Input example:

```
storcli /c0 show rebuildrate
```

6.6.2.4.2 Patrol Read

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

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

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

The detailed description for each command follows.

storcli /*cx* resume patrolread

This command resumes a suspended patrol read operation.

Input example:

```
storcli /c0 resume patrolread
```

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

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

Input example:

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

storcli /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 operations.

Table 39 Set Patrol Read Input Options

Option	Value Range	Description
`starttime`	A valid date and hour in 24 – hour format	Sets the start time in *yyyy/mm/dd hh* format.
`maxconcurrentpd`	Valid number of physical drives present	Sets the number of physical drives that can be patrol read at a single time.
`includessds`	—	Include SSDs in the patrol read operation.
`uncfgareas`	—	Irrespective of the status of `uncfgareas` (on|off), patrol read always scans the entire physical drive. However, if `uncfgareas=on`, patrol read corrects the media error in the unconfigured area of the physical drive. If `uncfgareas=off`, patrol read does not correct the media error in the unconfigured area of the physical drive.

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

Input example:

```
storcli /c0 set patrolread=on starttime=2012/02/2100
```

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

This command delays the scheduled patrol read in hours.

Input example:

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

storcli /c*x* show patrolRead

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

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

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

Input example:

```
storcli /c0 show patrolread
```

storcli /c*x* start patrolread

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

Input example:

```
storcli /c0 start patrolread
```

storcli /c*x* stop patrolread

This command stops a running patrol read operation.

Input example:

```
storcli /c0 stop patrolread
```

NOTE You cannot resume a stopped patrol read.

storcli /c*x* pause patrolread

This command pauses a running patrol read operation.

Input example:

```
storcli /c0 pause patrolread
```

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

6.6.2.4.3 Consistency Check

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

```
storcli /cx set consistencycheck|cc=[off|seq|conc][delay=value]
starttime=yyyy/mm/dd hh [excludevd=x-y,z| None]
```

```
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,none]

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

Table 40 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=CONC starttime=2012/02/21 00 excludevd v0-v3
```

storcli /*cx* show cc

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

Input example:

```
storcli /c0 show cc
```

storcli /*cx* show ccrate

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

Input example:

```
storcli /c0 show ccrate
```

NOTE A high CC rate slows I/O processing.

6.6.2.5 Premium Feature Key Commands

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

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

The detailed description for the command follows.

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

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

Table 41 Set Advanced Software Options Input Options

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

Input example:

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

storcli /*cx* show safeid

This command shows the Safe ID of the specified controller.

Input example:

```
storcli /c0 show safeid
```

6.6.2.6 Controller Security Commands

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

`storcli /cx compare securitykey=ssssss`

`storcli /cx delete securitykey`

`storcli /cx set securitykey keyid=kkkk`

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

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

`storcli /c x[/ex]/s xset security=on`

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 42 Set Security Key Input Options

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

Input example:

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

storcli /*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
```

storcli /c *x*/ex/*sx* set security=on

This command sets the security on the FDE-capable JBOD drive.

Input example

```
storcli /c0/e0/s0/set security=on
```

6.6.2.7 Flashing Controller Firmware Command

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

The following command is used to flash the controller firmware.

storcli /c*x* download file=*filepath* [fwtype=<*value*>] [nosigchk] [noverchk] [resetnow]

This command flashes the firmware with the ROM file to the specified adapter from the given file location (*filepath* is the absolute file path). See Online Firmware Upgrade and Downgrade for limitations.

You can use the following options in the table to flash the firmware:

Table 43 Flashing Controller Firmware Input Options

Option	Value Range	Description
nosigchk	—	The application flashes the firmware even if the check word on the file does not match the required check word for the controller. **NOTE** You can damage the controller if a corrupted image is flashed using this option.
noverchk	—	The application flashes the controller firmware without checking the version of the firmware image.
fwtype	0: Application 1: TMMC 2: GC-Enhanced	The firmware type to be downloaded. The application downloads the firmware for the controller. The TMMC downloads the firmware for the TMMC battery only. Default is 0 (application).
resetnow	—	Invokes online firmware update on the controller; you do not need to reboot the controller to make the update effective. **NOTE** The resetnow option is not supported in the UEFI mode.

6.6.2.8 Controller Cache Command

The following command flushes the controller cache.

storcli /c*x* flush|flushcache

This command flushes the controller cache.

Input example:

```
storcli /c0 flushcache
```

6.6.2.9 Controller Configuration Commands

The following command works with the controller configuration.

storcli /cx set config file=log.txt

This command restores the controller configuration and its properties from a specified file.

NOTE You cannot load a saved configuration file over an existing configuration file when there are already existing virtual drives. You must first clear the configuration file on the target controller.

Input example:

```
storcli /c0 set config file=log.txt
```

storcli /cx get config file=file_name

This command saves the controller configuration and its properties to a specified file.

Input example:

```
storcli /c0 get config file=filename
```

6.6.3 Diagnostic Commands

The Storage Command Line Interface Tool supports the following diagnostic commands:

```
storcli /cx start diag duration
```

The detailed description for each command follows.

storcli /cx start Diag Duration=<Val>

This commands runs the diagnostic self-check on the controller for the specified time period in seconds.

Input example:

```
storcli /c0 start diag duration=5
```

6.6.4 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 44 Physical Drives Commands Quick Reference Table

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

6.6.4.1 Drive Show Commands

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

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

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

The detailed description for each command follows.

storcli /c*x*[/e*x*]/s*x* show

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

Input example:

```
storcli /c0/e0/s4 show
```

storcli /c*x*[/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 /c*x*[/e*x*]/s*x*|sall show all

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

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

Input examples:

```
storcli /c0/e3/s0-3 show all
storcli /c0/e35/sall show all
```

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

storcli /c*x*/[e*x*]/s*x* show smart

This command displays the SMART information of a SATA drive.

Input example:

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

storcli /cx/ex/sx show errorcounters

If a faulty cable or a bad drive is found, this command displays the error counters for that specific faulty cable or a bad drive. If no drive is present, this command only displays error counters for a faulty cable.

Input example:

```
storcli /c0/e5/s1 show errorcounters
```

NOTE Note that specifying <ex> or the enclosure index is optional.

storcli /cx/ex/sx reset errorcounters type=<1>|<2>

This command resets the drive/slot error counters.

- If you input the error counter type as 1, this command resets the drive error counters.
- If you input the error counter type as 2, this command resets the slot error counters.

- If no is drive present, this argument takes 2 as an input and resets only the slot error counters.

Input example:

```
storcli /c0/e5/s1 reset errorcounters type=1
```

NOTE Note that specifying <ex> or the enclosure index is optional.

6.6.4.2 Missing Drives Commands

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

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

The detailed description for each command follows.

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.

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

This command marks a drive as missing.

Input example:

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

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/dall

This command is used to find the missing drives.

6.6.4.3 Set Drive State Commands

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

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

The detailed description for each command follows.

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

This command sets the drive state to JBOD.

Input example:

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

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

This command changes the drive state to unconfigured good.

Input example:

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

NOTE If the drive has an operating system or a file system on it, the StorCLI tool displays an error message and fails the conversion. If you want to proceed with the conversion, use the `force` option as shown in the following command.

Input example:

```
storcli /c1/e56/s3 set good [force]
```

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

This command changes the drive state to offline.

Input example:

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

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

This command changes the drive state to online.

Input example:

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

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

This command marks a drive as missing.

Input example:

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

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

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

Input example:

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

6.6.4.4 Drive Initialization Commands

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

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

The detailed description for each command follows.

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

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

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

Input example:

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

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

This command starts the initialization process on a drive.

Input example:

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

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

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

Input example:

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

6.6.4.5 Drive Firmware Download Commands

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

storcli /*cx*[/*ex*]/*sx* download src=filepath [satabridge] [mode= 5|7]

This command flashes the drive firmware with the specified file.

The `satabridge` option lets you download the SATA bridge firmware in online mode.

The `mode` options specify the SCSI write buffer mode. The description follows:

- 5 – The entire drive firmware file is downloaded at once.
- 7 – The drive firmware file is downloaded in chunks of 32KB.

NOTE The default mode is 7.

Input example:

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

Input example:

```
storcli /c0/e56/s1 download src=c:\file1.bin mode=5
```

storcli /*cx*[/*ex*]/*sx* download src= <filepath>[mode= E|F]*offline*[activatenow] *[delay*=<*value*>*]*

storcli /*cx*[/*ex*]/*sx* download mode=F*offline [delay*=<*value*>*]*

These commands support the drive firmware download using `Mode E` and `Mode F`. The mode options specify the `SCSI WRITE BUFFER` mode.

The description follows:

- **Mode E** – Downloads the microcode and allows you to issue this command for multiple devices. You can only use this in an offline mode.
- **Mode F** – Activates the deferred microcode and allows you to issue this command to all devices in a safe manner. You can only use this in an offline mode. You cannot issue this command before issuing the `Mode E` command. The default delay time is `15` seconds. You can specify any delay time between `1` to `300` seconds.

NOTE You can download as well as activate the drive firmware by executing the `activatenow` command in the same command line. You can

also specify the delay time, but the delay time specified by you is applicable only for activation and not for downloading the drive firmware.

Input examples for Mode E:

```
storcli /c0/e0/s0download src=file.rom mode=E offline
```

Download successful.

```
storcli /c0/e0/sall download src=file.rom mode=E offline
```

Downloaded sequentially on the drives.

Input Examples for Mode F:

```
storcli /c0/e0/sall download mode=F offline
```

Activation of the microcode successful

```
Storcli /c0/e0/sall download mode=F offline delay=15
```

Activation completed with a 15-second delay.

6.6.4.6 Drive Firmware Update Through Parallel HDD Microcode

ThinkSystem provides an interface to update the drive firmware in both online and offline modes through host applications, such as, the StorCLI tool. Using the parallel HDD microcode update feature, firmware updates can be done simultaneously on multiple HDDs of the same family in an online mode. Also, the parallel HDD microcode update overcomes the VD tolerance level. You can use the parallel HDD microcode update feature to update up to eight devices at the same time. It is recommended to perform the parallel HDD microcode update in system maintenance mode.

The parallel HDD microcode update is not supported in the following scenarios:

- If physical drive firmware download is already in progress on any physical drive.
- If Pinned Cache is present on the controller.
- Online firmware upgrade is not supported if FEATURE SET value is enabled for DEFAULT and disabled for LOW COST.

Command Usage Examples

```
storcli /c0/ex/sall download src=drv_fw.lod [mode=5/7] [parallel] [force]
```

```
storcli /c1/e1/sall download src=drivefirmware.lod mode=5 parallel
```

Where:

- **c** – Controller number
- **x** – The index of either the controller or the enclosure
- **e** – Enclosure number
- **s** – Slot number
- **sall** – All drives
- **parallel** – Indicates firmware update is done in parallel mode
- **force** – Indicates whether you want to force this operation

storcli /c0/e1/sall download status

This command provides the current firmware download status on the specified drive list.

6.6.4.7 Locate Drives Commands

The StorCLI tool supports the following commands to locate a drive and activate the physical disk activity LED:

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

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

The detailed description for each command follows.

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

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

Input example:

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

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

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

Input example:

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

6.6.4.8 Prepare to Remove Drives Commands

The StorCLI tool supports the following commands to prepare the physical drive for removal:

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

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

The detailed description for each command follows.

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

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

Input example:

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

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

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

Input example:

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

NOTE The `spinup` command works on a physical drive only if the user had previously issued a `spindown` command on the same physical drive.

6.6.4.9 Drive Security Command

The StorCLI tool supports the following drive security commands:

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

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

This command shows the security key for secured physical drives.

Input example:

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

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

This command enables security on a JBOD.

Input example:

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

6.6.4.10 Drive Secure Erase Commands

The StorCLI 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-SED drives.

Input example:

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

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

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

Table 45 Drive Erase Command Options

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

Input example:

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

6.6.4.11 Rebuild Drives Commands

The following commands rebuild drives in the StorCLI 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
```

6.6.4.12 Drive Copyback Commands

The StorCLI 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.

6.6.4.13 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 46 Add Hot Spare Drive Input Options

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

Input example:

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

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

Input example:

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

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

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

This command deletes a hot spare drive.

Input example:

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

6.6.4.14 Drive Predictive Failure Monitoring Commands

The StorCLI tool supports the following commands for drive predictive failure monitoring:

```
Storcli /cx show pdfailevents [lastoneday] [fromSeqNum=xx] [file=filename]
Storcli /cx set pdfaileventoptions detectiontype=val correctiveaction=val
errorrthreshold=val
```

The detailed description for each command follows.

Storcli / *cx* show pdfailevents *[lastoneday] [fromSeqNum=xx][file=filename]*

This command shows all of the drive predictive failure events.

Input example 1:

```
storcli /c0 show pdfailevents
```

This command shows all of the drive predictive failure events from the oldest sequence number.

Input example 2:

```
storcli /c0 show pdfailevents lastoneday
```

This command shows all of the drive predictive failure events that occurred in the last 24 hours.

Input example 3:

```
storcli /c0 show pdfailevents fromSeqNum
```

This command shows all of the drive predictive failure events generated from the specified sequence number.

NOTE While running these commands, if you provide a file name, the events are written to the specified file as values separated by commas.

Storcli / *cx* set pdfaileventoptions detectiontype=*val* correctiveaction=*val* errorrthreshold=*val*

This command provides the current settings of the `pdfaileventoptions` set on the controller and the various options to change these settings.

Input example 1:

```
storcli /c0 set pdfaileventoptions detectiontype=x
```

Where:

- `00b` = Detection disabled
- `01b` = Detection enabled, high latency for reads is OK.
- `10b` = Detection enabled, aggressive (high latency for reads is not OK).
- `11b` = Detection enabled, use NVDATA specified value, see `recoveryTimeLimit` and `writeRetryCount`.

This command sets the detection type for the drive. The valid range is 0 to 3.

NOTE For the changes to take effect, a reboot is required.

Input example 2:

```
storcli /c0 set pdfaileventoptions correctiveaction=x
```

Where:

- `0` = Only log events
- `1` = Log events, take corrective action based on SMARTer.

This command sets the corrective actions to be taken when the media error is detected. The valid value is 0 or 1.

Input example 3:

```
storcli /c0 set pdfaileventoptions errorrthreshold=x
```

Where:

- `00b` – `1` = One error every 8 hours (least tolerant)
- `01b` – `8` = One error every 1 hour.
- `10b` – `32` = One error every 15 minutes.
- `11b` – `90` = One error every 5 minutes (most tolerant of drive with degraded media).

This command sets the error threshold for the controller. The valid range is 0 to 3.

6.6.5 Virtual Drive Commands

The StorCLI tool supports the following virtual drive commands. The following table describes frequently used virtual drive commands.

Table 47 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 virtual drives. See the **Virtual Drive Show Command** section.	Shows virtual drive information.

6.6.5.1 Add Virtual Drives Commands

The StorCLI tool supports the following commands to add virtual drives:

```
storcli /cx add vd raid[0|1|5|6|00|10|50|60][Size=<VD1_Sz>,<VD2_Sz>,..|all]
[name=<VDNAME1>,..] drives=e:s|e:s-x,y;e:s-x,y,z [PDperArray=x][SED]
[pdcache=on|off|default][pi] [DimmerSwitch(ds)=default|automatic(auto)|
none|maximum(max)|MaximumWithoutCaching(maxnocache)]

[wt|wb|awb] [nora|ra] [direct|cached][cachevd] [Strip=<8|16|32|64|128|256|1024>]
[AfterVd=X][EmulationType=0|1|2] [Spares = [e:]s|[e:]s-x|[e:]s-x,y]
[force][ExclusiveAccess]
```

NOTE The supported strip size can vary from a minimum of 64 KB to 1 MB for ThinkSystem controllers and only 64 KB for Integrated RAID controllers.

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

NOTE The supported strip size can vary from a minimum of 64 KB to 1 MB for ThinkSystem controllers and only 64 KB for Integrated RAID controllers.

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

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

NOTE * indicates default values.

The detailed description for each command follows.

storcli /cx add vd raid[0|1|5|6|00|10|50|60][Size=<*VD1_Sz*>,<*VD2_Sz*>,..|*all] [name=<*VDNAME1*>,..] drives=*e:s*|*e:s-x*|*e:s-x,y*;*e:s-x,y,z* [PDperArray=x][SED] [pdcache=on|off|*default][pi] [DimmerSwitch(ds)=default|automatic(auto)|

***none|maximum(max)|MaximumWithoutCaching(maxnocache)][cachevd][ExclusiveAccess|SharedAccess*]** [wt|*wb |awb] [nora|*ra] [*direct|cached] [EmulationType=0][Strip=<8|16|32|64|128|256|1024>] [AfterVd=*X*] [Spares = [*e*:]*s*|[*e*:]*s-x*|[*e*:]*s-x,y*] [force]**

NOTE The supported strip size can vary from a minimum of 64 KB to 1 MB for ThinkSystem controllers and only 64 KB for Integrated RAID controllers.

Table 48 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	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. **NOTE** Power savings for logical devices are not supported.
direct|cached	cached: Cached I/O. direct: Direct I/O.	Sets the logical drive cache policy. Direct I/O is the default.
EmulationType	0: Default emulation, which means if there are any 512e drives in the configured ID, then the physical bytes per sector is shown as 512e(4k). If there are no 512e drives then the physical bytes per sector will be 512n. 1: Disable, which means even though there are no 512e drives in the configured ID, the physical bytes per sector will be shown 512n. 2=Force, which means even though there are no 512e drives in the configured ID, the physical bytes per sector will be shown as 512e (4k).	
wt|wb|awb	wt: Write through.wb: Write back.awb: Always Write Back.	Enables write through. Write back is the default.

Table 48 Add RAID Configuration Input Options (Continued)

Option	Value Range	Description
nora|ra	ra: Read ahead.nora: No read ahead.	Disables read ahead. Enabled is the default.
cachevd	—	Enables SSD caching on the created virtual drive.
strip	8, 16, 32, 64, 128, 256, 512, 1024. **NOTE** The supported strip size can vary from a minimum of 64 KB to 1 MB for the ThinkSystem controllers and only 64 KB for Integrated ThinkSystem controllers.	Sets the strip size for the RAID configuration.
aftervd	Valid virtual drive number.	Creates the VD in the adjacent free slot next to the specified VD.
spares	Number of spare physical drives present.	Specifies the physical drives that are to be assigned to a disk group for spares.
force	—	Forces a security-capable physical drive to be added to a drive group without security.

Input example:

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

6.6.5.2 Delete Virtual Drives Commands

The StorCLI tool supports the following virtual drive delete commands:

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

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

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

The detailed description for each command follows.

storcli /cx/vx|vall del

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

Input example:

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

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

storcli /cx/vx|vall del force

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

Input example:

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

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

storcli /*cx*/v*x* del [cachecade] [discardcache] [force]

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

Input example:

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

6.6.5.3 Virtual Drive Show Commands

The StorCLI tool supports the following virtual drive show commands:

```
storcli /cx/vx show
storcli /cx/vx show all [logfile[=filename]]
```

The detailed description for each command follows.

storcli /*cx*/v*x* show

This command shows the summary of the virtual drive information.

Input example:

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

storcli /*cx*/v*x* show all [logfile[=*filename*]]

The `show all` command shows all of the virtual drive information, which includes the virtual drive information, physical drives used for the virtual drives, and virtual drive properties.

If you use the `logfile` option in the command syntax, the logs are written to the specified file. If you do not specify a file name, then the logs are written to the `storsas.log` file. If you do not use the `logfile` option in the command syntax, the entire log output is printed to the console.

Input example:

```
storcli /c0/v0 show all [logfile[=log.txt]]
```

6.6.5.4 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 StorCLI 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*/v*x* delete preservedcache

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

Input example:

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

storcli /*cx* show preservedCache

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

Input example:

```
storcli /c0 show preservedCache
```

6.6.5.5 Change Virtual Drive Properties Commands

The StorCLI tool supports the following commands to change virtual drive properties:

```
storcli /cx/vx set accesspolicy=<rw|ro|blocked|rmvblkd>
storcli /cx/vx set iopolicy=<cached|direct>
storcli /cx/vx set name=<namestring>
storcli /cx/vx set pdcache=<on|off|default>
storcli /cx/vx set rdcache=<ra|nora>
storcli /cx/vx|vall set ssdcaching=<on|off>
storcli /cx/vx|vall set HostAccess=ExclusiveAccess|SharedAccess
storcli /cx/vx set wrcache=<wt|wb|awb>
storcli /cx/vx set emulationType=0|1|2
storcli /cx/vx set ds=Default|Auto|None|Max|MaxNoCache
storcli /cx/vx set autobgi=On|Off
storcli /cx/vx set pi=Off
storcli /cx/vx set bootdrive=<On|Off>
storcli /cx/vx set hidden=On|Off
storcli /cx/vx set cbsize=0|1|2 cbmode=0|1|2|3|4|7
```

The detailed description for each command follows.

storcli /*cx*/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 /*cx*/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 /*cx*/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 /*cx*/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 /cx/vx 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 /cx/vx|vall set HostAccess=ExclusiveAccess|SharedAccess

This command sets the host access policy for the virtual drive. when the host access policy is exclusive access, a server has exclusive access to the virtual drive. The virtual drive cannot be shared between servers. If the host policy is shared access, the virtual drive can be shared between servers.

Input example:

```
storcli /c0/v0 set HostAccess=ExclusiveAccess
```

storcli/cx/vx set wrcache=<wt|wb|awb>

This command sets the write cache policy on a virtual drive to write back, write through, or always write back.

Input example:

```
storcli /c0/v0 set wrcache=wt
```

storcli /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
```

NOTE If the virtual drive is set to be "boot device", then it cannot be hidden

storcli /cx/vx set cbsize=0|1|2 cbmode=0|1|2|3|4|7

This command sets the Cache bypass size and the Cache bypass mode on a virtual drive.

The `cbsize` option follows:

- `0` – 64k Cache bypass.
- `1` – 128k Cache bypass.
- `2` – 256k Cache bypass.

The `cbmode` option follows:

- `0` – Enable the intelligent mode Cache bypass.
- `1` – Enable the standard mode Cache bypass.
- `2` – Enable the custom mode Cache bypass 1.
- `3` – Enable the custom mode Cache bypass 2.
- `4` – Enable the custom mode Cache bypass 3.
- `7` – Disable Cache bypass.

NOTE When `cbmode` is set to 7, the user given `cbsize` value is ignored

Input example:

```
storcli /c0/v0 set cbsize=1 cbmode=2
```

6.6.5.6 Virtual Drive Initialization Commands

The Storage Command Line Interface Tool supports the following commands to initialize virtual drives:

```
storcli /cx/vx show init
storcli /cx/vx start init [full][Force]
storcli /cx/vx stop init
```

NOTE If the virtual drive has user data, you must use the `force` option to initialize the virtual drive.

A virtual drive with a valid MBR and partition table is considered to contain user data.

The detailed description for each command follows.

storcli /*cx*/*vx* show init

This command shows the initialization progress of a virtual drive in percentage.

The estimated time (in minutes) left to complete the operation is also shown.

Input example:

```
storcli /c0/v2 show init
```

storcli /*cx*/*vx* start init [full]

This command starts the initialization of a virtual drive. The default initialization type is fast initialization. If the `full` option is specified, full initialization of the virtual drive starts.

Input example:

```
storcli /cx/vx start init [full]
```

storcli /*cx*/*vx* stop init

This command stops the initialization of a virtual drive. A stopped initialization cannot be resumed.

Input example:

```
storcli /c0/v0 stop init
```

6.6.5.7 Virtual Drive Erase Commands

The Storage Command Line Interface Tool supports the following commands to erase virtual drives:

```
storcli /cx/vx erase
storcli /cx/vx show erase
```

The detailed description for each command follows.

storcli /*cx*/vx erase

This command erases the data on the virtual drive.

Input example:

```
storcli /c0/v0 erase
```

storcli /*cx*/vx show erase

This command shows the status of the erase operation on the virtual drive.

Input example:

```
storcli /c0/v0 show erase
```

6.6.5.8 Virtual Drive Migration Commands

NOTE The virtual drive migration commands are not supported in Embedded RAID.

The Storage Command Line Interface Tool supports the following commands for virtual drive migration (reconstruction):

```
storcli /cx/vx show migrate
```

```
storcli /cx/vx start migrate <type=raidx> [option=<add|remove>
drives=[e:]s|[e:]s-x|[e:]s-x,y] [Force]
```

The detailed description for each command follows.

storcli /*cx*/*vx* show migrate

This command shows the progress of the virtual drive migrate operation in percentage.

The estimated time (in minutes) left to complete the operation is also shown.

Input example:

```
storcli /c0/v0 show migrate
```

storcli /*cx*/*vx* start migrate <type=*raidlevel*> [option=<add | remove> drives=<e1:s1,e2:s2 ...>]

This command starts the reconstruction on a virtual drive to the specified RAID level by adding or removing drives from the existing virtual drive. You can use the following options with the start migrate command.

Table 49 Virtual Drive Migration Command Options

Options	Value Range	Description
`type =RAID level`	`RAID [0|1|5|6]`	The RAID level to which the virtual drive must be migrated.
`[option=<add | remove> drives=<e1:s1,e2:s2, …>]`	`add`: Adds drives to the virtual drive and starts reconstruction. `remove`: Removes drives from the virtual drive and starts reconstruction. `drives`: The enclosure number and the slot number of the drives to be added to the virtual drive. **NOTE** Make sure that the same block size (in a physical drive) is used in each [*e*:*s*] pair. As an example, if you use 4096 bytes in the *e0*:*s0* pair, use 4096 bytes in the *e1*:*s1* pair too. Mixing of block sizes between the [*e*:*s*] pairs is not supported.	Adds or removes drives from the virtual drive.

Virtual drive migration can be done between the following RAID levels.

Table 50 Virtual Drive Migration Table

Initial RAID level	Migrated RAID level
RAID 0	RAID 1
RAID 0	RAID 5
RAID 0	RAID 6
RAID 1	RAID 0
RAID 1	RAID 5
RAID 1	RAID 6
RAID 5	RAID 0
RAID 5	RAID 6
RAID 6	RAID 0
RAID 6	RAID 5

Input example: In the following example, 252 is the enclosure number and 0, 1, and 2 are the slot numbers.

```
storcli/c0/v0 start migrate type=raid0 option=add drives=252:0,252:1,252:2
```

6.6.5.9 Virtual Drive Consistency Check Commands

The Storage Command Line Interface Tool supports the following commands for virtual drive consistency checks:

```
storcli /cx/vx pause cc
storcli /cx/vx resume cc
storcli /cx/vx show cc
storcli /cx/vx start cc [force]
storcli /cx/vx stop cc
```

NOTE If enclosures are used to connect the physical drives to the controller, specify the IDs in the command.

The detailed description for each command follows.

storcli /*cx*/*vx* pause cc

This command pauses an ongoing consistency check process. You can resume the consistency check at a later time. You can run this command only on a virtual drive that has a consistency check operation running.

Input example:

```
storcli /c0/v4 pause cc
```

storcli /*cx*/*vx* resume cc

This command resumes a suspended consistency check operation. You can run this command on a virtual drive that has a paused consistency check operation.

Input example:

```
storcli /c0/v4 resume cc
```

storcli /*cx*/*vx* show cc

This command shows the progress of the consistency check operation in percentage.

The estimated time (in minutes) left to complete the operation is also shown.

Input example:

```
storcli /c0/v5 show cc
```

storcli /*cx*/*vx* start cc force

This command starts a consistency check operation for a virtual drive. Typically, a consistency check operation is run on an initialized virtual drive. Use the `force` option to run a consistency check on an uninitialized drive.

Input example:

```
storcli /c0/v4 start cc
```

storcli /*cx*/*vx* stop cc

This command stops a consistency check operation. You can run this command only for a virtual drive that has a consistency check operation running.

Input example:

```
storcli /c0/v4 stop cc
```

NOTE You cannot resume a stopped consistency check process.

6.6.5.10 Background Initialization Commands

The Storage Command Line Interface Tool supports the following commands for background initialization:

```
storcli /cx/vx resume bgi
storcli /cx/vx set autobgi=<on|off>
storcli /cx/vx show autobgi
storcli /cx/vx show bgi
storcli /cx/vx stop bgi
storcli /cx/vx suspend bgi
```

The detailed description for each command follows.

storcli /*cx*/*vx* resume bgi

This command resumes a suspended background initialization operation.

Input example:

```
storcli /c0/v0 resume bgi
```

storcli /*cx*/*vx* set autobgi=<on|off>

This command sets the auto background initialization setting for a virtual drive to on or off.

Input example:

```
storcli /c0/v0 set autobgi=on
```

storcli /*cx*/*vx* show autobgi

This command shows the background initialization setting for a virtual drive.

The estimated time (in minutes) left to complete the operation is also shown.

Input example:

```
storcli /c0/v0 show autobgi
```

storcli /*cx*/*vx* show bgi

This command shows the background initialization progress on the specified virtual drive in percentage.

The estimated time (in minutes) left to complete the operation is also shown.

Input example:

```
storcli /c0/v0 show bgi
```

storcli /*cx*/*vx* stop bgi

This command stops a background initialization operation. You can run this command only for a virtual drive that is currently initialized.

Input example:

```
storcli /c0/v4 stop bgi
```

storcli /*cx*/*vx* pause bgi

This command suspends a background initialization operation. You can run this command only for a virtual drive that is currently initialized.

Input example:

```
storcli /c0/v4 pause bgi
```

6.6.5.11 Virtual Drive Expansion Commands

The Storage Command Line Interface Tool supports the following commands for virtual drive expansion:

```
storcli /cx/vx expand size=<value> [expandarray]
storcli /cx/vx|vall show expansion
```

The detailed description for each command follows.

storcli /*cx*/*vx* expand size=<*value*> [expandarray]

This command expands the virtual drive within the existing array or if you replace the drives with drives larger than the size of the existing array. Even though the value provided by you may be in MB, the value of the expanded size is displayed based on the nearest possible unit. Depending on the input (value) provided by you, `storcli` recognizes the size from the input provided by you and rounds up the size to the nearest percentage of free space remaining on the drive group; hence, the actual expanded size may differ from the size requested by you. If the `expandarray` option is specified, the existing array is expanded. If this option is not specified, the virtual drive is expanded.

storcli /*cx*/*vx* show expansion

This command shows the expansion information on the virtual drive with and without array expansion.

Input example:

```
storcli /c0/v0 show expansion
```

6.6.5.12 Display the Bad Block Table

The Storage Command Line Interface Tool supports the following command to check for bad block entries of virtual drives on the selected controller:

```
storcli /cx/vx show bbmt
```

Input example:

```
storcli /c0/v0 show bbmt
```

6.6.5.13 Clear the LDBBM Table Entires

The Storage Command Line Interface Tool supports the following command to clear the LDBBM table entries:

`storcli /cx/vx delete bbmt`

Input example:

```
storcli /c0/v0 delete bbmt
```

6.6.6 Clear a Configuration

To clear an existing configuration, use the `delete config` command.

storcli /cx delete config [force]

This command allows you to clear an existing configuration.

Input example

```
storcli /c0 delete config[force]
```

6.6.7 Foreign Configurations Commands

The Storage Command Line Interface Tool supports the following commands to view, import, and delete foreign configurations:

`storcli /cx/fall|fall del|delete [ securitykey=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
```

6.6.8 BIOS-Related Commands

The Storage Command Line Interface Tool supports the following BIOS commands:

```
storcli /cx set bios [state=<on|off>] [Mode=<SOE|PE|IE|SME>] [abs=<on|off>]
[DeviceExposure=<value>]
```

The detailed description for the command follows.

storcli /*cx* set bios [state=<on|off>] [Mode=<SOE|PE|IE|SME>] [abs=<on|off>] [DeviceExposure=<value>]

This command enables or disables the controller's BIOS, sets the BIOS boot mode, and enables the BIOS to select the best logical drive as the boot drive.The mode options abbreviations follow:

- SOE: Stop on Errors.
- PE: Pause on Errors.
- IE: Ignore Errors.
- SME: Safe mode on Errors.

NOTE The legacy BIOS can load a limited number of the PCI device's BIOS. Disable the BIOS to avoid issues during POST.

Input example:

```
storcli /c0 set bios[state=on][Mode=SOE][abs=on][deviceexposure=20]
```

6.6.8.1 OPROM BIOS Commands

The Storage Command Line Interface Tool supports the following OPROM BIOS commands:

```
storcli /cx/ex/sx set bootdrive=on|off
storcli /cx/vx set bootdrive=on|off
storcli /cx show bootdrive
```

The detailed description for each command follows.

storcli /*cx*/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
```

6.6.9 Drive Group Commands

This section describes the drive group commands.

6.6.9.1 Drive Group Show Commands

The Storage Command Line Interface Tool supports the following drive group commands:

```
storcli /cx/dall show
storcli /cx/dall show all
storcli /cx/dall show cachecade
storcli /cx/dx show
storcli /cx/dx show all
storcli /cx/dx set security=on
storcli /cx/dx split mirror
storcli /cx/dall show mirror
storcli /cx/dall add mirror src=<val>[force]
storcli /cx/dx set hidden=<on|off>
```

storcli /cx/dall show

This command shows the topology information of all the drive group.

Input example:

```
storcli /c0/dall show
```

storcli /cx/dall show all

This command shows all available configurations in the controller which includes topology information, virtual drive information, physical drive information, free space, and free slot information.

Input example:

```
storcli /c0/dall show all
```

storcli /cx/dx show

This command shows the topology information of the drive group.

Input example:

```
storcli /c0/dx show
```

storcli /cx/dx show all

This command shows the physical drive and the virtual drive information for the drive group.

Input example:

```
storcli /c0/dx show all
```

storcli /cx/dx set security=on

This command enables security on the specified drive group.

Input example:

```
storcli /c0/dx set security=on all
```

storcli /*cx*/dx split mirror

This command enables you to perform a break mirror operation on a drive group. The break mirror operation enables a RAID 1 configured drive group to be broken into two volumes. You can use one of the volumes in another system and replicate it without making a copy of the virtual drive.

Input example:

```
storcli /c0/dx split mirror
```

storcli /*cx*/dall show mirror

This command shows information about the mirror associated with the drive group.

Input example:

```
storcli /c0/dall show mirror
```

storcli /*cx*/dall add mirror src=<val>[force]

This command joins the virtual drive with its mirror. The possible values to be used are 0, 1, or 2.

Input example:

```
storcli /c0/dall add mirror src=<1>[force]
```

storcli /*cx*/d*x* set hidden=<on|off>

This command hides or unhides a drive group.

Input example:

```
storcli /c0/d0 set hidden=on
```

6.6.10 Dimmer Switch Commands

6.6.10.1 Change Virtual Drive Power Settings Commands

The Storage Command Line Interface Tool supports the following commands to change the Dimmer Switch settings. You can use the following combinations for the Dimmer Switch commands:

```
storcli /cx set ds=off type=1|2|4
storcli /cx set ds=on type=1|2 [properties]
storcli /cx set ds=on type=4 defaultldtype=<value> [properties]
storcli /cx set ds=on [properties]
```

The following table describes the power-saving options.

Table 51 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.
`properties`	`disableldps`: Interval in hours or time in *hh:mm* format `spinupdrivecount`: Valid enclosure number (0 to 255) `SpinUpEncDelay`: Valid time in seconds	Sets the interval or time in which the power-saving policy for the logical drive is turned off. Specifies the number of drives in the enclosure that are spun up. Specifies the delay of spin-up groups within an enclosure in seconds.

storcli/*cx* show DimmerSwitch(ds)

This command shows the current Dimmer Switch setting for the controller.

Input example:

```
storcli/c0 show ds
```

6.6.11 Enclosure Commands

The Storage Command Line Interface Tool supports the following enclosure commands:

```
storcli /cx/ex download src=filepath[forceActivate]
storcli /cx/ex show all
storcli /cx/ex show status
```

The detailed description for each command follows.

storcli /*cx*/*ex* download src=filepath [forceactivate]

This command flashes the firmware with the file specified at the command line. The enclosure performs an error check after the operation. The following option can be used with the enclosure firmware download command.

Table 52 Enclosure Firmware Download Command Options

Option	Value Range	Description
`forceactivate`	—	Issues a command descriptor block (CDB) with write command with no data with command mode 0x0F (flash download already in progress). **NOTE** This option is used primarily to activate Scotch Valley Enclosures.

NOTE The firmware file that is used to flash the enclosure can be of any format. The StorCLI utility assumes that you provide a valid firmware image.

Input example:

```
storcli /c0/e0 download src=c:\file2.bin
```

storcli /*cx*/*ex* show all

This command shows all enclosure information, which includes general enclosure information, enclosure inquiry data, a count of enclosure elements, and information about the enclosure elements.

Input example:

```
storcli /c0/e0 show all
```

storcli /*cx*/*ex* show status

This command shows the enclosure status and the status of all the enclosure elements.

Input example:

```
storcli /c0/e0 show status
```

6.6.12 PHY Commands

The Storage Command Line Interface Tool supports the following phy commands:

```
storcli /cx/px|pall set linkspeed=0(auto)|1.5|3|6|12
storcli /cx/px|pall show
storcli /cx/px|pall show all
storcli /cx/ex show phyerrorcounters
storcli /cx[/ex]/sx show phyerrorcounters
storcli /cx[/ex]/sx reset phyerrorcounters
```

The detailed description for each command follows.

storcli /*cx*/*px*|pall set linkspeed=0(auto)|1.5|3|6|12

This command sets the PHY link speed. You can set the speed to 1.5 Gb/s, 3 Gb/s, 6 Gb/s, or 12 Gb/s. The linkspeed is set to `auto` when you specify `linkspeed = 0`.

Input example:

```
storcli /c0/p0 set linkspeed=1.5
```

storcli /*cx*/*px*|pall show

This command shows the basic PHY layer information.

Input example:

```
storcli /c1/p0 show
```

storcli /*cx*/*px*|pall show all

This command shows all the PHY layer information.

Input example:

```
storcli /c1/p0 show all
```

storcli /*cx*/*ex* show phyerrorcounters

This command shows the enclosure/expander phy error counters.

Input example:

```
storcli /c1/e0 show phyerrorcounters
```

storcli /*cx*[/*ex*]/*sx* show phyerrorcounters

This command shows the drive phy error counters.

Input example:

```
storcli /c1/e0/s0 show phyerrorcounters
```

storcli /*cx*[/*ex*]/*sx* reset phyerrorcounters

This command resets the drive phy error counters.

Input example:

```
storcli /c1/e0/s0 reset phyerrorcounters
```

6.6.13 Logging Commands

The Storage Command Line Interface Tool supports the following commands to generate and maintain log files:

```
storcli /cx clear events
storcli /cx delete termlog
storcli /cx show events file=<absolute path>
storcli /cx show eventloginfo
storcli /cx show termlog type=config|contents [logfile[=filename]]
storcli /cx show dequeuelog file =<filepath>
Storcli /cx show alilog [logfile[=filename]]
```

The detailed description for each command follows.

storcli /*cx* delete events

This command deletes all records in the event log.

Input example:

```
storcli /c0 delete events
```

storcli /*cx* delete termlog

This command clears the TTY (firmware log for issue troubleshooting) logs.

Input example:

```
storcli /c0 delete termlog
```

storcli /*cx* show events file=<*absolute path*>

This command prints the system log to a text file and saves the file in the specified location.

Input example:

```
storcli /c0 show events file=C:\Users\brohan\test\eventreports
```

storcli /*cx* show eventloginfo

This command shows the history of log files generated.

Input example:

```
storcli /c0 show eventloginfo type=config
```

storcli /*cx* show termlog type=config|contents [logfile[=*filename*]]

This command shows the firmware logs. The `config` option shows the term log configuration (settings of TTY BBU buffering), the `contents` option shows the term log. The `contents` option is the default.

If you use the `logfile` option in the command syntax, the logs are written to the specified file. If you do not specify a file name, then the logs are written to the `storsas.log` file. If you do not use the `logfile` option in the command syntax, the entire log output is printed to the console.

Input example:

```
storcli /c0 show termlog=contents [logfile[=log.txt]]
```

storcli /*cx* show dequeuelog =<filepath>

This command shows the debug log from the firmware.

Input example:

```
storcli /c0 show dequeuelog=<c:\test\log.txt>
```

storcli /*cx*show alilog [logfile[=*filename*]]

This command gets the controller property, TTY logs, and events to the specified file.

Input example:

```
storcli /c0 show alilog [logfile[=log.txt]]
```

6.6.14 Automated Physical Drive Caching Commands

The Storage Command Line Interface Tool supports the following automated physical drive caching commands:

```
storcli /cx set autopdcache=<off|r0>[immediate]
storcli /cx show autopdcache
```

The detailed description for each command follows.

storcli /*cx* set autopdcache=<off|r0>[immediate]

This command lets you set the controller's automated physical drive cache policy to RAID 0. When set to RAID-0, all un-configured physical drives are configured as a single RAID 0 drive, until the maximum virtual drive limit is reached. The `immediate` option lets this command execute the conversion (to RAID 0) operation only on all the existing physical drives. Any newly physical drives connected in the future do not get converted to RAID 0. If you omit the `immediate` option in this command, conversion to RAID 0 takes place on newly connected physical drives too. Automatic conversion to RAID 0 can be turned off by setting the autopdcache policy to `off`.

Input example:

```
storcli /c0 set autopdcache=r0 immediate
```

storcli /*cx* show autopdcache

This command lets you view the automatic physical drive caching property.

Input example:

```
storcli /c0 show autopdcache
```

6.7 Frequently Used Tasks

6.7.1 Showing the Version of the Storage Command Line Interface Tool

The following command shows the version of the command line tool:

```
storcli -v
```

6.7.2 Showing the StorCLI Tool Help

The following command shows the StorCLI tool help:

```
storcli -h
```

Help appears for all the StorCLI tool commands.

6.7.3 Showing System Summary Information

The following command shows the summary of all the controller information:

```
storcli -show [all]
```

6.7.4 Showing Free Space in a Controller

The following command shows the free space available in the controller:

```
storcli /cx show freespace
```

6.7.5 Adding Virtual Drives

The following command creates a virtual drive:

```
storcli /cx add vd type=raid[0|1|5|6|10|50|60][Size=<VD1_Sz>,<VD2_Sz>,..|*all]
[name=<VDNAME1>,..] drives=e:s|e:s-x|e:s-x,y [PDperArray=x|auto*]
[SED] [pdcache=on|off|*default][pi] [DimmerSwitch(ds)=default|automatic(auto)|
*none|maximum(max)|MaximumWithoutCaching(maxnocache)] [wt|*wb|awb] [nora|*ra]
[*direct|cached]
[strip=<8|16|32|64|128|256|512|1024] [AfterVd=x] [Spares=[e:]s|[e:]s-x|[e:]s-x,y]
```

NOTE The supported strip size can vary from a minimum of 64 KB to 1 MB for controllers and only 64 KB for Integrated controllers.

```
[Cbsize = 0|1|2 Cbmode = 0|1|2]
[force]
```

The following inputs can be used when adding virtual drives:

- The controller in which the virtual drives are created.
- The RAID type of the virtual drives.
  The supported RAID types are 0, 1, 5, 6, 10, 50, 60.
- The size of each virtual drive.

- The drives that are used to create the virtual drives.

  drives = *e*:*s*|*e*:*s*-*x*|*e*:*s*-*x*,*y*

  Where:

  — *e* specifies the enclosure ID.

  — *s* represents the slot in the enclosure.

  — *e*:*s*-*ex* is the range conventions used to represents slots s to x in the enclosure e.
- The physical drives per array.

  The physical drives per array can be set to a particular value.
- The SED option creates security-enabled drives.
- The PDcache option can be set to on or off.
- The pi option enables protection information.
- The Dimmer Switch is the power save policy. It can be set to default or automatic *, none, maximum(max), or MaximumWithoutCaching(maxnocache).
- The wt option disables write back.
- The nora option disables read ahead.
- The cached option enables the cached memory.
- The strip option sets the strip size.

  It can take the values 8, 16, 32, 64, 128, 256, 512, 1024.

  NOTE The supported strip size can vary from a minimum of 64 KB to 1 MB for controllers and only 64 KB for Integrated controllers.
- The AfterVdX option creates the virtual drives in the adjacent free slot next to the specified virtual drives.

  NOTE The * indicates default values used in the creation of the virtual drives. If values are not specified, the default values are taken.

Example: /c*x*add vd type=r1 drives=0:10-15 WB Direct strip=64

This command creates a RAID volume of RAID 1 type from drives in slots 10 to slot 15 in enclosure 0. The strip size is 64kb.

6.7.6 Setting the Cache Policy in a Virtual Drive

The following command sets the write cache policy of the virtual drive:

storcli /c*x*/v(*x*|all) set wrcache=wt|wb|awb

The command sets the write cache to write back, write through, or always write back.

6.7.7 Showing Virtual Drive Information

The following command shows the virtual drive information for all the virtual drives in the controller:

storcli /c*x* show [all]

6.7.8 Deleting Virtual Drives

The following command deletes virtual drives:

storcli /c*x*/v(*x*|all) del [cc|cachecade]

The following inputs are required when deleting a virtual drive:

- The controller on which the virtual drive or virtual drives is present.
- The virtual drives that must be deleted; or you can delete all the virtual drives on the controller using the `vall` option.
- The `cc` or `cachecade` option to confirm that the deleted drive is a CacheCade drive.

6.7.9 Flashing Controller Firmware

The following command is used to flash the controller firmware.

```
storcli /cx download file=filepath [fwtype=<value>] [nosigchk]
[noverchk][resetnow]
```

For more information, see Flashing Controller Firmware Command. For limitations, see Online Firmware Upgrade and Downgrade.

Appendix A: 3ware CLI Commands to StorCLI Command Conversion

A.1 System Commands

Table 53 System Commands

Description	3ware® CLI Command	StorCLI Command
Show a general summary of all detected controllers.	`tw_cli show`	`show` `show ctrlcount`

A.2 Controller Commands

Table 54 Controller Commands

Description	3ware CLI Command	StorCLI Command
Show all information about the adapter, such as cluster state, BIOS, alarm, firmware, version, and so on.	`tw_cli /cx show all`	`/cx show all`
Download the firmware to all compatible controllers that can be flashed with the image. By default, CLI checks for signature and version.	`/cx update fw=filename_with_path [force]`	`/cx download src=filepath [nosigchk] [noverchk]`

Show the status of properties related to the controllers.	`/cx show <PropertyName>`
	
	The following properties can be used with this command:
	`a0,1,2|-aALL`
	`achip`
	`AENs [reverse]`
	`alarms [reverse]`
	`allunitstatus`
	`autocarve`
	`autorebuild`
	`bios`
		`/cx show <PropertyName>`
		
		The following properties can be used with this command:
		`abortcconerror`
		`activityforlocate`
		`alarm`
		`autorebuild`
		`backplane`
		`batterywarning`
		`bgirate`
		`bootwithpinnedcache`

Table 54 Controller Commands (Continued)

Description	3ware CLI Command	StorCLI Command
	carvesize	cachebypass
	ctlbus diag	cacheflushint
	dpmstat [type=<inst|ra|ext>	ccrate
	driver	
	drivestatus	coercion
	events [reverse]	copyback
	exportjbod firmware	directpdmapping
	memory	ds
	model	eccbucketleakrate
	monitor	eccbucketsize
	numdrives	enableeghsp
	numports	enableesmarter
	numunits	enableeug
	ondegrade	exposeencldevice
	pcb	jbod
	pchip	loadbalancemode
	phy	maintainpdfailhistory
	rebuild	migraterate
	rebuildmodel	ncq
	rebuildrate	perfmode
	selftest	pr
	serial	prcorrectunconfiguredareas
	spinup	prrate
	stagger	rebuildrate
	unitstatus	rehostinfo
	verify	restorehotspare
	verifymode	safeid
	verifyrate	smartpollinterval
		spinupdelay
		spinupdrivecount
		time
		usefdeonlyencrypt
Set properties on the selected controllers.	autocarve=<on|off>	abortcconerror=<on|off>

Table 54 Controller Commands (Continued)

Description	3ware CLI Command	StorCLI Command
	autodetect=<on|off >	activityforlocate=<on|off>
	disk=<p:-p>|all	alarm=<on|off>
	autorebuild=<on|off>	autorebuild=<on|off>
	carvesize=<1024..32768>	backplane=<*value*>
	dpmstat=<on|off>	batterywarning=<on|off>
	ondegrade=<cacheoff|follow>	bgirate=<*value*>
	rebuild=<enable|disable|><1..5>	bootwithpinnedcache=<on|off>
	rebuildmode=<adaptive|lowlatency>	cachebypass=<on|off>
	rebuildrate=<1..5>	flush|flushcache
	selftest=<enable|disable>	cacheflushinterval=<*value*>
	spinup=<*value*>	ccrate=<*value*>
	stagger=<*value*>	coercion=<*value*>
	verify=advanced|basic|<1..5>	clusterenable=<*value*>
	verify=basic [pref=*ddd*:*hh*] where *hh*=(00...23 and *ddd*={mon|tue|wed|thu|fri|sat|sun}	copyback=<on|off> type=<smartssd|smarthdd|all>
	verify=enable|disable|<1..5>	directpdmapping=<on|off>
	verifymode=<adaptive|lowlatency>	eccbucketleakrate=<*value*>
	verifyrate=<1..5>	eccbucketsize=<*value*>
		enableeghsp=<on|off>
		enableesmarter=<*value*>
		enableeug=<on|off>
		exposeencldevice=<on|off>
		foreignautoimport=<on|off>
		jbod=<on|off>
		loadbalancemode=<*value*>
		maintainpdfailhistory=<on|off>
		migraterate=<*value*>
		ncq=<on|off>
		perfmode=<*value*>
		prcorrectunconfiguredareas=<on|off>
		prrate=<*value*>
		rebuildrate=<*value*>
		restorehotspare=<on|off>
		smartpollinterval=<*value*>
		spinupdelay=<*value*>
		spinupdrivecount=<*value*>

Table 54 Controller Commands (Continued)

Description	3ware CLI Command	StorCLI Command
		stoponerror=<on|off>
		usefdeonlyencrypt=<on|off> time=yyyymmddhh:mm:ss|systemtime usefdeonlyencrypt=<on|off>

A.3 Alarm Commands

Table 55 Alarm Commands

Description	3Ware CLI Command	StorCLI Command
Set alarm properties.	/c*x*/e*x*/alm*x* set alarm=<mute|unmute|off> **NOTE** The 3ware® controllers have enclosure alarms.	/c*x* set alarm=<on|off|silence> **NOTE** The StorCLI controllers have controller alarms.
Show alarm properties.	/c*x*/e*x* show alarms **NOTE** This command applies for only 9750 and 9690SA controllers.	/c*x* show alarm

A.4 Patrol Read and Consistency Check Commands

Table 56 Patrol Read and Consistency Check Commands

Description	3ware CLI Command	StorCLI Command
Show patrol read status and patrol read parameters, if any in progress.	/c*x*/u*x* show	/c*x* show patrolRead
Set the patrol read options on a single adapter, multiple adapters, or all adapters (*x* = single controller).	/c*x*/u*x* start verify /c*x*/u*x* set autoverify=<on|off> /c*x* add verify=dddh:hh:duration	/c*x* set patrolread {=on mode=<auto|manual>}|{off} /c*x* set patrolread [starttime=<*yyyy/mm/dd hh* [maxconcurrentpd=<*value*>] [includessds=<on|off>] [uncfgareas=on|off] /c*x* set patrolread delay=<*value*>
Show consistency check status, if any in progress, and consistency check parameters.	/c*x*/u*x* show	/c*x*/v*x* show cc /c*x* show ccrate
Set consistency check options on a single adapter, multiple adapters, or all adapters (*x* = single controller).	/c*x*/u*x* start verify /c*x*/u*x* set autoverify=<on|off> /c*x* add verify=ddd:hh:duration	storcli /c*x* set consistencycheck|cc=[off|seq|conc] [delay=*value*] [starttime=*yyyy/mm/dd hh*] [excludevd=*x-y,z*|*None*]

NOTE The 3ware® CLI combines both patrol read and consistency check into a single command. The StorCLI has different commands for each.

A.5 BBU Commands

Table 57 BBU Commands

Description	3ware CLI Command	StorCLI Command
Show complete BBU information, such as status, capacity information, design information, and properties.	/cx/bbu show all	/cx/bbu show all
Show BBU summary information.	/cx/bbu show	/cx/bbu show
Show BBU properties.	/cx/bbu show batinst /cx/bbu show bootloader /cx/bbu show fw /cx/bbu show lasttest /cx/bbu show pcb /cx/bbu show serial /cx/bbu show status /cx/bbu show temp /cx/bbu show tempstat /cx/bbu show tempval /cx/bbu show volt	/cx/bbu show properties /cx/bbu show status **NOTE** Not all the properties shown in the 3ware CLI are shown in the StorCLI.
Show BBU capacity information.	/cx/bbu show cap	/cx/bbu show all
Start the learning cycle on the BBU.	/cx/bbu test [quiet]	/cx/bbu start learn

A.6 Virtual Drive Commands

Table 58 Virtual Drive Commands

Description	3Ware CLI Command	StorCLI Command
Create a RAID volume of the specified RAID type.	/c*x* add vd type=<*RaidType*> disk=<*p*:*p*|*p*-*p*|*p*:*p*-*p*>>> (where *p*=port or drive number) [strip=<*size*>] [nocache|nowrcache] [nordcache|rdcachebasic] [name=string (9000 series)] [ignoreECC] [autoverify|noautoverify] v0=*n*|vol=*a*:*b*:*c*:*d*] (*n*, *a*, *b*, *c*, *d*=size of volume in GB) [noqpolicy] [storsave=<protect|balance|perform>] [noscan] [rapidrecovery=<all|rebuild|disable>] [group=<3|4|5|6|7|8|9|10|11|12|13|14|15|16>] *RaidType*={raid0, raid1, raid5, raid10, raid50, single, spare, raid6}	/c*x* add vd type=raid[0|1|5|6|10|50|60] [[size=<*vd1_size*>,<*vd2_size*>,..|*all][name=<*vdname1*>,..] drives=*e*:*s*|*e*:*s*-*x*|*e*:*s*-*x*,*y*|*e*:*s*-*x*,*y*,*z* [pdperarray=*x*|*auto] [sed] [pdcache=on|off|*default] [pi][dimmerswitch|ds=default|automatic(auto)|*none|maximum(max) |maximumwithoutcaching(maxnocache)] [wt|*wb|awb] [nora|*ra] [*direct|cached] [strip=<8|16|32|64|128|256|512|1024] [aftervd=*x*] [spares=[*e*:]*s*|[*e*:]*s*-*x*|[*e*:]*s*-*x*,*y*|[*e*:]*s*-*x*,*y*,*z*>] [force] **NOTE** The supported strip size can vary from a minimum of 64 KB to 1 MB for ThinkSystem controllers and only 64 KB for Integrated ThinkSystem controllers.
Delete virtual drives.	/c*x*/u*x* del [quiet] **NOTE** You can delete a single unit using this command.	/c*x*/v*x* [all] delete [force] [cachecade] **NOTE** You can delete one virtual disk, multiple virtual disks, or all the selected virtual disks on selected adapters using this command.
Show drive group information.	/c*x*/u*x* show [all] **NOTE** Information of each unit is shown individually.	/c*x*/dall show [cachecade]
Scan and show available foreign configurations, provide a preview of the imported foreign configuration, show or import foreign configuration.	/c*x* rescan	c*x*/f*all* [all] show [preview] [securityKey=*sssssssssss*] c*x*/f*all* [all] import [securityKey=*sssssssssss*]
Show VD information, including name, RAID level, RAID level qualifier, size in MBs, state, strip size, number of drives, span depth, cache policy, access policy, and any ongoing activity progress, which includes initialization, background initialization, consistency check, and reconstruction.	/c*x*/u*x* show [all]	/c*x*/v*x* show all

Table 58 Virtual Drive Commands (Continued)

Description	3Ware CLI Command	StorCLI Command
Show the virtual drive properties.	/c*x*/u*x* show autoverify /c*x*/u*x* show identify /c*x*/u*x* show ignoreECC /c*x*/u*x* show initializestatus /c*x*/u*x* show name /c*x*/u*x* show parit /c*x*/u*x* show qpolicy /c*x*/u*x* show rapidrecovery /c*x*/u*x* show rdcache /c*x*/u*x* show rebuildstatus /c*x*/u*x* show serial /c*x*/u*x* show status /c*x*/u*x* show storsave /c*x*/u*x* show verifystatus /c*x*/u*x* show volumes /c*x*/u*x* show wrcache	/c*x*/v*x* show all **NOTE** The StorCLI does not have commands to show individual virtual drive properties.
Set virtual drive properties.	/c*x*/u*x* set autoverify=on|off /c*x*/u*x* set cache=on|off [quiet] /c*x*/u*x* set identify=on|off /c*x*/u*x* set ignoreECC=on|off /c*x*/u*x* set name=string /c*x*/u*x* set qpolicy=on|off /c*x*/u*x* set rapidrecovery=all|rebuild|disable /c*x*/u*x* set rdcache=basic|intelligent|off /c*x*/u*x* set storsave=protect|balance|perform [quiet] /c*x*/u*x* set wrcache=on|off [quiet]	/c*x*/v*x* set accesspolicy=<rw|ro|blocked|rmvblkd> /c*x*/v*x* set iopolicy=<cached|direct> /c*x*/v*x* set name=<*namestring*> /c*x*/v*x* set pdcache=<on|off|default> /c*x*/v*x* set rdcache=<ra|nora|adra> /c*x*/v*x* set security=<on|off> /c*x*/v*x*|vall set ssdcaching=<on|off> /c*x*/v*x* set wrcache=<wt|wb|awb>
Show cache and access policies of the virtual drive.	/c*x*/u*x* show [all] /c*x*/u*x* show autoverify /c*x*/u*x* show cache /c*x*/u*x* show identify /c*x*/u*x* show ignoreECC /c*x*/u*x* show name /c*x*/u*x* show parit /c*x*/u*x* show qpolicy /c*x*/u*x* show rapidrecovery /c*x*/u*x* show rdcache /c*x*/u*x* show rebuildstatus /c*x*/u*x* show serial /c*x*/u*x* show status intializestatus /c*x*/u*x* show storsave /c*x*/u*x* show verify status /c*x*/u*x* show volumes /c*x*/u*x* show wrcache	/c*x*/v*x* show all **NOTE** The StorCLI does not have commands to show individual virtual drive properties.

Table 58 Virtual Drive Commands (Continued)

Description	3Ware CLI Command	StorCLI Command
Start initialization (writing 0s) on the virtual drive.	/c*x*/u*x* start verify **NOTE** Only the bios can do a foreground initialization. A background initialization does otherwise. A verify starts a back ground initialization.	/c*x*/v*x* start init [Full]
Stop an ongoing initialization on the virtual drive.	/c*x*/u*x* stop verify **NOTE** Only the bios can do a foreground initialization. A background initialization does otherwise. A verify starts a back ground initialization	/c*x*/v*x* stop init
Show a snapshot of the ongoing initialization, if any.	/c*x*/u*x* show [all] **NOTE** Only the bios can do a foreground initialization. A background initialization does otherwise. A verify starts a back ground initialization.	/c*x*/v*x* show init
Start a consistency check on the virtual drive.	/c*x*/u*x* start verify	/c*x*/v*x* start cc
Stop a consistency check on the virtual drive.	/c*x*/u*x* stop verify	/c*x*/v*x* stop cc
Reconstruct the selected virtual disk to a new RAID level.	/c*x*/u*x* migrate type=<*RaidType*> [disk=<*p*:-*p*..>] [strip=<*size*>] [noscan] [nocache] [autoverify] [*group*=<3|4|5|6|7|8|9|10|11|12|13|14|15|16>] *RaidType*={ raid0, raid1, raid5, raid10, raid50, single, raid6 }	/c*x*/v*x* start migrate <type=*raidlevel*> [option=<add | remove> disk=<e1:s1,e2:s2 ..>] /c*x*/v*x* show migrate
Change the power-saving setting on the virtual drive.	/c*x*/u*x* set powersavestandbytimer=<5 to 999>	/c*x*/v*x* set ds=<default | Auto | None | Max | MaxNoCache>

A.7 Physical Drive Commands

Table 59 Physical Drive Commands

Description	3ware CLI Command	StorCLI Command
Show physical disk information.	/c*x*/p*x* show [all]	/c*x*[/e*x*]/s*x* show [all]
Start, stop, suspend, or resume an ongoing rebuild operation.	/c*x*/u*x* start rebuild disk=<p:-p...> [ignoreECC] **NOTE** Rebuilds cannot be stopped or paused.	/c*x*[/e*x*]/s*x* start rebuild /c*x*[/e*x*]/s*x* stop rebuild /c*x*[/e*x*]/s*x* pause rebuild /c*x*[/e*x*]/s*x* resume rebuild
Mark the configured physical disk drive as missing for the selected adapter.	/c*x*/p*x* remove [quiet]	/c*x*[/e*x*]/s*x* set missing
Change the physical disk drive state to offline.	/c*x*/p*x* remove [quiet]	/c*x*[/e*x*]/s*x* set offline
Add jbod.	/cx add vd type=jbod disk=<*p*> (where *p* = port or drive number)	/c*x*[/e*x*]/s*x* set jbod
Change the physical disk drive hot spare state and associate the drive to an enclosure and virtual disk.	/cx add vd type=spare disk=<*p*:*p*|*p*-*p*|*p*:*p*-*p*> (where *p* = port or drive number)	/c*x*[/e*x*]/s*x* add hotsparedrive [{dgs=<N|0,1.2...*n*,,>] [EnclAffinity][nonRevertible]

Table 59 Physical Drive Commands (Continued)

Description	3ware CLI Command	StorCLI Command
Locate the physical disk drive and activate the physical disk activity LED.	/c*x*/p*x* set identify=on|off	/c*x*[/e*x*]/s*x* start | stop locate
Prepare the unconfigured physical drive for removal.	/c*x*/p*x* remove [quiet]	/c*x*[/e*x*]/s*x* spindown
Show information about all physical disk drives and other devices connected to the selected adapters; includes drive type, size, serial number, and firmware version.	/c*x*/p*x* show [all]	/c*x*/eall/sall show [all]
Download drive or expander firmware.	/c*x*/p*x* update fw=*image.name* [force]	/c*x*[/e*x*]/s*x* download src=filepath [satabridge]

A.8 Enclosure Commands

Table 60 Enclosure Commands

Description	3ware CLI Command	StorCLI Command
Show information about the enclosure for the selected adapter.	/c*x*/e*x* show [all]	/c*x*/e*x* show [all]
Show the status of the enclosure connected to the selected adapter.	/c*x*/e*x* show [all] /c*x*/e*x* show controllers /c*x*/e*x* show slots /c*x*/e*x* show fans /c*x*/e*x* show temp /c*x*/e*x* show pwrs /c*x*/e*x* show alms	/c*x*/e*x* show status
Download enclosure firmware.	/c*x*/e*x* update fw=*image.name* [force]	/c*x*/e*x* download src=*filepath* [offline] [forceActivate]

A.9 Events and Logs

Table 61 Events and Logs

Description	3ware CLI Command	StorCLI Command
Show the total number of events, newest and oldest sequence number, shutdown sequence number, reboot sequence number, clear sequence number.	/c*x* show alarms **NOTE** This command shows AENs since last controller reset.	/c*x* show eventloginfo
Show the total event entries available at the firmware since last clear, and details of each entries of error log.	/c*x* show alarms **NOTE** This command shows AENs since last controller reset.	/c*x* show events filter=<Info | warning| critical| fatal > file=<*path of the file*>
Show the count of events starting from specified seqNum and matching category and severity	/c*x* show alarms **NOTE** This command shows AENs since last controller reset.	/c*x* show events type=<sinceShutDown | sinceReboot |ccincon vd=<0,1,2…> | includeDeleted | latest=*x* filter=<Info | warning| critical| fatal > file=<*path of the file*>
Show TTY firmware terminal log entries with details on given adapters. The information is shown as total number of entries available on the firmware side.	/c*x* show diag	/c*x* show TermLog [type=contents|Config]

A.10 Miscellaneous Commands

Table 62 Miscellaneous Commands

Description	3ware CLI Command	StorCLI Command
Show version information.	tw_cli ?	ver
Show help for all show commands at server level.	tw_cli ? tw_cli /c*x* ? tw_cli /c*x*/u*x* ? tw_cli /c*x*/p*x* ? tw_cli /c*x*/phy*x* ? tw_cli /c*x*/bbu ? tw_cli /c*x*/e*x* ? tw_cli /e*x* **NOTE** The 3ware CLI shows context-sensitive help.	show help
Show PHY connection information for physical PHY medium on the adapters.	/c*x*/phy*x* show	/c*x*/p*x* show
Set PHY link speed.	/c*x*/phy*x* set link=<0|1.5|3.0|6.0|12.0>	/cx/px set linkspeed=0(auto)|1.5|3|6|12

Appendix B: MegaCLI Commands to StorCLI Command Conversion

B.1 System Commands

Table 63 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`

B.2 Controller Commands

Table 64 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`
	`adpalilog`	alilog logfile=filename `storcli /cx show AliLog [logfile[=filename]]`
	`adpdiag`	Storcli /c0 start Diag Duration=val `storcli /cx start Diag Duration=<Val>`
	`autodetectbackplanedsbl`	`backplane`
	`autoenhancedimportdsply`	`foreignautoimport`
	`autosnapshotspace`	
	`batwarndsbl`	`batterywarning`
	`bgirate`	`bgirate`
	`bootwithpinnedcache`	`bootwithpinnedcache`
	`cachebypass`	`cachebypass`
	`ccrate`	`ccrate`
	`clusterenable`	
	`coercionmode`	`coercion`
	`copybackdsbl`	`copyback`
	`defaultldpspolicy`	`ds`
	`defaultsnapshotspace`	
	`defaultviewspace`	

Table 64 Controller Commands (Continued)

Description	MegaCLI Command	StorCLI Command
	disableldpsinterval	ds
	disableldpstime	ds
	disableocr	ocr
	eccbucketcount	eccbucketsize
	eccbucketleakrate	eccbucketleakrate
	enableeghsp	eghs
	enableesmarter	eghs
	enableeug	eghs
	enablejbod	Jbod
	enblspindownunconfigdrvs	ds
	loadbalancemode	loadbalancemode
	maintainpdfailhistoryenbl	maintainpdfailhistory
	ncqdsply	ncq
	patrolreadrate	prrate
	perfmode	perfmode
	predfailpollinterval	smartpollinterval
	rebuildrate	rebuildrate
	reconrate	migraterate
	rstrhotspareoninsert	restorehotspare
	smartcpybkenbl	copyback
	spindowntime	ds
	spinupencdelay	ds
	spinupdelay	spinupdelay
	spinupencdrvcnt	spinupdrivecount
	ssdsmartcpybkenbl	copyback
	usediskactivityforlocate	activityforlocate
	usefdeonlyencrypt	usefdeonlyencrypt
Set properties on the selected controllers.	Megacli -AdpSetProp <*propertyname*>-an|-a0,1,2|-aall	/c*x* set <*property1*>
	The following properties can be set using this command:	The following properties can be set using this command:
	abortcconerror	abortcconerror=<on|off>
	alarmdsply	alarm=<on|off| silence>
	autodetectbackplanedsbl	backplane=<*value*>
	autoenhancedimportdsply	foreignautoimport=<on|off>
	batwarndsbl	batterywarning=<on|off>
	bgirate	bgirate=<*value*>
	bootwithpinnedcache	bootwithpinnedcache=<on|off>
	cachebypass	cachebypass=<on|off>

Table 64 Controller Commands (Continued)

Description	MegaCLI Command	StorCLI Command
	ccrate	ccrate=<*value*>
	clusterenable	
	coercionmode	coercion=<*value*>
	copybackdsbl	copyback=<on|off> type=<smartssd|smarthdd|all>
	defaultldpspolicy	ds=<*value*>
	defaultsnapshotspace	
	defaultviewspace	
	disableldpsinterval	ds=<*value*>
	disableldpstime	ds=<*value*>
	disableocr	ocr=<*value*>
	eccbucketcount	eccbucketsize=<*value*>
	eccbucketleakrate	eccbucketleakrate=<*value*>
	enableeghsp	eghs [state=<on|off>]
	enableesmarter	eghs [smarter=<on|off>]
	enableeug	eghs [eug=<on|off>
	enablejbod	jbod=<on|off>
	enblspindownunconfigdrvs	ds=<*value*>
	loadbalancemode	loadbalancemode=<*value*>
	maintainpdfailhistoryenbl	maintainpdfailhistory=<on|off>
	ncqdsply	ncq=<on|off>
	patrolreadrate	prrate=<*value*>
	perfmode	perfmode=<*value*>
	predfailpollinterval	smartpollinterval=<*value*>
	rebuildrate	rebuildrate=<*value*>
	reconrate	migraterate=<*value*>
	rstrhotspareoninsert	restorehotspare=<on|off>
	smartcpybkenbl	copyback=<on|off> type=<smartssd|smarthdd|all>
	spindowntime	ds=<on|off>
	spinupdelay	spinupdelay=<*value*>
	spinupdrivecount	spinupdrivecount=<*value*>
	spinupencdelay	ds
	spinupencdrvcnt	ds
	sdsmartcpybkenbl	copyback=<on|off> type=<smartssd|smarthdd|all>
	usediskactivityforlocate	activityforlocate=<on|off>
	usefdeonlyencrypt	usefdeonlyencrypt=<on|off>
Show the number of controllers connected.	MegaCLI -adpCount	storcli show ctrlcount

Table 64 Controller Commands (Continued)

Description	MegaCLI Command	StorCLI Command
Show all information about the adapter, such as cluster state, BIOS, alarm, firmware, version, and so on.	MegaCli -AdpAllInfo –aN|–a0,1,2|-aALL	storcli /c*x* show all
Show the freespace available in the controller.	MegaCLI -CfgFreeSpaceinfo –aN|–a0,1,2|-aALL	storcli /c*x* show freespace
Download the controller firmware.	MegaCli -AdpFwFlash -f *filename* [-NoSigChk] [-NoVerChk] [-ResetNow] –aN|–a0,1,2|-aALL	storcli /c*x* download file=<*filepath*> [fwtype=<*val*>] [nosigchk] [noverchk][resetnow]
Show the preserved cache status.	MegaCLI-GetPreservedCacheList –aN|-a0,1,2|-aALL	storcli /c*x* show preservedcache
Set the controller time	MegaCLI –AdpSetTime *yyyymmdd hh*:*mm*:*ss* -aN|-a0,1,2|-aALL	storcli /c(x|all) set time=<*yyyymmdd hh*:*mm*:*ss* | systemtime>
Show the controller time.	MegaCLI –AdpGetTime -aN	storcli /c*x* show time

B.3 Patrol Read Commands

Table 65 Patrol Read Commands

Description	MegaCLI Command	StorCLI Command
Show the patrol read status and patrol read parameters, if any in progress.	MegaCli -AdpPR -info -aN|-a0,1,2|-aALL	storcli/c*x* show patrolRead
Set the patrol read options on a single adapter, multiple adapters, or all adapters. (x = single controller).	MegaCli -AdpPR –Dsbl|EnblAuto|EnblMan|Start|Stop|Info|Suspend|Resume|Stop|SSDPatrolReadEnbl | SSDPatrolReadDsbl |{SetDelay *Val*}|{-SetStartTime *yyyymmdd hh*}|{maxConcurrentPD Val} -aN|-a0,1,2|-aALL	storcli /c*x* set patrolread {=on mode=<auto|manual>}|{off} storcli /c*x* set patrolread [starttime=<*yyyy/mm/dd hh*>] [maxconcurrentpd=<*value*>] [includessds=<on|off>] [uncfgareas=on|off] storcli /c*x* set patrolread delay=<*value*>
Disable patrol read.	MegaCli -AdpPR -Dsbl -aN|-a0,1,2|-aALL	storcli /c*x* set patrolread=off
Enable automatic patrol read.	MegaCli -AdpPR -EnblAuto -aN|-a0,1,2|-aALL	storcli /c*x* set patrolread=on mode=auto
Enable manual patrol read.	MegaCli -AdpPR -EnblMan -aN|-a0,1,2|-aALL	storcli /c*x* set patrolread=on mode=manual
Start patrol read.	MegaCli -AdpPR -Start -aN|-a0,1,2|-aALL	storcli /c*x* start patrolRead
Suspend a running patrol read.	MegaCli -AdpPR -Suspend -aN|-a0,1,2|-aALL	storcli /c*x* suspend patrolRead
Resume a suspended patrol read.	MegaCli -AdpPR -Resume -aN|-a0,1,2|-aALL	storcli /c*x* resume patrolRead
Stop a running patrol read.	MegaCli -AdpPR -Stop -aN|-a0,1,2|-aALL	storcli /c*x* stop patrolRead
Include SSD drives in patrol read.	MegaCli -AdpPR -SSDPatrolReadEnbl -aN|-a0,1,2|-aALL	storcli /c*x* set patrolRead includessds=on | onlymixed

Table 65 Patrol Read Commands (Continued)

Description	MegaCLI Command	StorCLI Command
Exclude SSD drives in patrol read.	`MegaCli -AdpPR -SSDPatrolReadDsbl -aN|-a0,1,2|-aALL`	`storcli /cx set patrolRead includessds=off`
Delay a patrol read,	`MegaCli -AdpPR -SetDelay Val -aN|-a0,1,2|-aALL`	`storcli /cx set patrolread delay=<value>`
Schedule a patrol read.	`MegaCli -AdpPR -SetStartTime yyyymmdd hh -aN|-a0,1,2|-aALL`	`storcli /cx set patrolread=on starttime=YYYY/MM/DD HH`
Set the value for maximum concurrent physical drives for the patrol read.	`MegaCli -AdpPR -maxConcurrentPD Val -aN|-a0,1,2|-aALL`	`storcli /cx set patrolread maxconcurrentpd=xx`

B.4 Consistency Check Commands

Table 66 Consistency Check Commands

Description	MegaCLI Command	StorCLI Command
Schedule a consistency check.	`MegaCLI -AdpCcSched -Dsbl|-Info| {-ModeConc | -ModeSeq [-ExcludeLD -LN|-L0,1,2] [-SetStartTime yyyymmdd hh ] [-SetDelay val ] } -aN|-a0,1,2|-aALL`	`storcli /cx set consistencycheck|cc=[off|seq|conc] [delay=value] starttime=yyyy/mm/dd hh [excludevd=x-y,z]`
Show consistency check status and consistency parameters, in progress, if any.	`MegaCLI -AdpCcSched -Info`	`storcli /cx show cc/ConsistencyCheck`

B.5 OPROM BIOS Commands

Table 67 OPROM BIOS Commands

Description	MegaCLI Command	StorCLI Command
Schedule a consistency check.	`MegaCli -AdpBIOS -Dsply -aN|-a0,1,2|-aALL`	`storcli /cx show bios`
Show consistency check status and consistency parameters, if any in progress.	`MegaCli -AdpBootDrive -{-Set {-Lx | -physdrv[E0:S0]}} -aN|-a0,1,2|-aALL`	`storcli /cx/ex/sx set bootdrive=on|off` `storcli /cx/vx set bootdrive=on|off`

Sets the BIOS properties for the controller.	`MegaCli -AdpBIOS -Enbl | -Dsbl | -Dsply | SOE | BE EnblAutoSelectBootLd | DsblAutoSelectBootLd -aN|-a0,1,2|-aALL`	`storcli /cx set bios=<on|off>`
		`storcli /cx set stoponerror|SOE=<on|off>`
		`storcli /cx set autobootselect(abs)=<on|off>`

B.6 Battery Commands

Table 68 Battery Commands

Description	MegaCLI Command	StorCLI Command
Show battery-related information.	`MegaCli -AdpBbuCmd -aN|-a0,1,2|-aALL`	`storcli /cx/bbu show` `storcli /cx/bbu show all`
Show the battery learn properties.	`MegaCli -AdpBbuCmd -GetBbuProperties -aN|-a0,1,2|-aALL`	`storcli /cx/bbu show properties`
Show the battery information, firmware status, and the gas gauge status.	`MegaCli -AdpBbuCmd -GetBbuStatus -aN|-a0,1,2|-aALL`	`storcli /cx/bbu show status`
Show battery capacity information.	`MegaCli -AdpBbuCmd -GetBbuCapacityInfo -aN|-a0,1,2|-aALL`	`storcli /cx/bbu show all`
Show battery design information.	`MegaCli -AdpBbuCmd -GetBbuDesignInfo -aN|-a0,1,2|-aALL`	`storcli /cx/bbu show all`

Set battery properties	`MegaCli -AdpBbuCmd -SetBbuProperties -f <fileName> -aN|-a0,1,2|-aALL`	`storcli /cx/bbu set learnDelayInterval=<value>`
		`storcli /cx/bbu set bbuMode=<value>`
		`storcli /cx/bbu set autolearnmode=<value>` where x= 0 – Enabled, 1 – Disabled, 2 – Warn though event.

Start battery learn cycle.	`MegaCli -AdpBbuCmd -BbuLearn -aN|-a0,1,2|-aALL`	`storcli /cx/bbu start learn`
Set the battery to low power storage mode.	`MegaCli -AdpBbuCmd -BbuMfgSleep -aN|-a0,1,2|-aALL`	`storcli /cx/bbu set powermode=sleep`
Seal the gas gauge EEPROM write access	`MegaCli -AdpBbuCmd -BbuMfgSeal -aN|-a0,1,2|-aALL`	`storcli /cx/bbu set writeaccess=sealed`

B.7 RAID Configuration Commands

Table 69 RAID Configuration Commands

Description	MegaCLI Command	StorCLI Command
Create a RAID configuration of RAID type 0, 1, 5, and 6.	MegaCli –CfgLDAdd -R0|-R1|-R5|-R6[E0:S0,E1:S1,...] [WT | WB] [NORA | RA | ADRA] [Direct | Cached] [CachedBadBBU|NoCachedBadBBU] [-sz*XXXXXXXX* [-sz*YYYYYYYY* [...]]] [-strpszM] [–Hsp[E5:S5,...]] [–afterLdX] -aN	storcli /c*x* add vd type=raid[0|1|5|6] [Size=<*VD1_Sz*>,< *VD2_Sz*>,..|*all] [name=<*VDNAME1*>,..] drives=*e*:*s*|*e*:*s-x*|*e*:*s-x*,*y*;*e*:*s-x*,*y*,*z* [PDperArray=*x*] [SED] [pdcache=on|off|*default][pi] [DimmerSwitch(ds)=default|automatic(auto)|*none|maximum(max) |MaximumWithoutCaching(maxnocache)] [wt|*wb|awb] [nora|*ra] [*direct|cached] [strip=<8|16|32|64|128|256|512|1024] [AfterVd=*X*] [Spares=[*e*:]*s*|[*e*:]*s-x*|[*e*:]*s-x*,*y*] [force] **NOTE** The supported strip size can vary from a minimum of 64 KB to 1 MB for ThinkSystem controllers and only 64 KB for Integrated ThinkSystem controllers.
Create a RAID configuration of RAID type 10, 50, and 60.	MegaCli –CfgSpanAdd -aN|-a0,1,2|-aALL -R10|-R50|R60 –Array0[E0:S0,E1:S1,...] –Array1[E0:S0,E1:S1,...] [...] [WT | WB] [NORA | RA | ADRA] [Direct | Cached] [CachedBadBBU|NoCachedBadBBU] [-sz*XXXXXXXX*[-sz*YYYYYYYY* [...]]] [-strpszM] [–afterLdX] -aN	storcli /c*x* add vd type=raid[10|50|60] [Size=<*VD1_Sz*>,<*VD2_Sz*>,..|*all] [name=<*VDNAME1*>,..] drives=*e*:*s*|*e*:*s-x*|*e*:*s-x*,*y*;*e*:*s-x*,*y*,*z* [PDperArray=*x*] [SED] [pdcache=on|off|*default][pi] [DimmerSwitch(ds)=default|automatic(auto)|*none|maximum(max) |MaximumWithoutCaching(maxnocache)] [wt|*wb|awb] [nora|*ra] [*direct|cached] [strip=<8|16|32|64|128|256|512|1024] [AfterVd=*X*] [Spares=[*e*:]*s*|[*e*:]*s-x*|[*e*:]*s-x*,*y*] [force] **NOTE** The supported strip size can vary from a minimum of 64 KB to 1 MB for ThinkSystem controllers and only 64 KB for Integrated RAID controllers.
Clear the complete configuration.	MegaCli -CfgClr [-Force] -aN|-a0,1,2|-aALL	storcli /c*0*/delete config [force]
Show the topology information of the drive group.	MegaCLI -CfgDsply -aN|-a0,1,2|-Aall	storcli /c*x*/dall show [all]
Delete a virtual drive hosting the operating system.	MegaCLI -CfgLdDel -LX|-L0,2,5...|-LALL [-Force] -aN|-a0,1,2|-aALL	storcli /c*x*/v/v*x* [all] delete -force
Show, delete, and import the foreign configuration commands.	MegaCli –CfgForeign –Scan | {-Preview | –Dsply| -Import | -Clear[FID]} -aN|-a0,1,2|-aALL"	storcli /c*x*/f(*x*|all) show [all] [securityKey=*xxx*] storcli /c*x*/f(*x*|all) del|delete [securityKey=*xxx*] storcli /c*x*/f(*x*|all) import [preview] [securityKey=*xxx*]"

B.8 Security Commands

Table 70 Security Commands

Description	MegaCLI Command	StorCLI Command
Set the key ID for the controller.	MegaCli -CreateSecurityKey -SecurityKey *sssssssssss* | [-Passphrase *sssssssssss*] |[-KeyID *kkkkkkkkkk*] -aN	storcli /c*x* set SecurityKey=*XXXXXX* [passphrase=*yyyyy*] [keyId=*zzzz*]
Change the security key for the controller.	MegaCli -ChangeSecurityKey -OldSecurityKey *sssssssssss* | -SecurityKey *sssssssssss* | [-Passphrase *sssssssssss*] | [-keyID *kkkkkkkkkk*] -aN	storcli /c*x* set SecurityKey=*XXXXXX* OldSecurityKey=*yyyyy*
Compare and verify the security key for the controller.	MegaCli -VerifySecurityKey -SecurityKey *sssssssssss* -aN	storcli /c*x* compare SecurityKey=*xxxxxx*
Delete the security key.	MegaCLI -DestroySecurityKey | [-Force] -aN	storcli /c*x* delete SecurityKey
Set the security key for the controller.	MegaCli -SetKeyID -KeyID *kkkkkkkkkk* -aN	storcli /c*x* set SecurityKey KeyId=*xxxx*

B.9 Virtual Drive Commands

Table 71 Virtual Drive Commands

Description	MegaCLI Command	StorCLI Command
Show the virtual drive information.	MegaCli –LDInfo –Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) show storcli /c*x*/v(*x*|all) show all
Set virtual drive properties.	MegaCli –LDSetProp WT | WB|NORA |RA | ADRA|-Cached|Direct|CachedBadBBU| NoCachedBadBBU} | -RW|RO|Blocked | {-Name nameString} |-EnDskCache|DisDskCache –Lx| -L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) set wrcache=WT|WB|AWB storcli /c*x*/v(*x*|all) set rdcache=RA|NoRA storcli /c*x*/v(*x*|all) set iopolicy=Cached|Direct storcli /c*x*/v(*x*|all) set accesspolicy=RW|RO|Blocked|RmvBlkd storcli /c*x*/v(*x*|all) set pdcache=On|Off|Default storcli /c*x*/v(*x*|all) set name=<*NameString*>
Set power-saving (dimmer switch) properties.	MegaCli -LDSetPowerPolicy -Default| -Automatic| -None| -Maximum| -MaximumWithoutCaching -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) set ds=Default|Auto|None|Max|MaxNoCache
Show virtual drive expansion information.	MegaCli -getLdExpansionInfo -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) show expansion

Table 71 Virtual Drive Commands (Continued)

Description	MegaCLI Command	StorCLI Command
Expand the virtual drive within the existing array; also use if you replace the drives with larger drives, beyond the size of the existing array.	MegaCli -LdExpansion -pN -dontExpandArray -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) expand Size=<*value*> [expandarray]
Secure the virtual drive.	MegaCLI --LDMakeSecure -Lx|-L0,1,2,...|-Lall –An	storcli /c*x*/v*x* set security=on
Show specific properties of virtual drives.	MegaCli –LDGetProp -Cache | -Access | -Name | -DskCache -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /c*x*/v*x* show
Start virtual drive initialization.	MegaCli –LDInit –Start [Fast|Full] -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) start init[Full]
Stop a running virtual drive initialization.	MegaCli –LDInit -Abort -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) stop init
Show the initialization progress.	MegaCli –LDInit –ShowProg -Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) show init
Start a consistency check on an uninitialized virtual drive.	MegaCli –LDCC –Start –Lx|-L0,1,2|-Lall -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) start cc[Force]
Start, stop, suspend, resume, and show the progress of a consistency check operation.	MegaCli -LDCC -Start|-Abort| -Suspend|-Resume|-ShowProg| -ProgDsply -Lx|-L0,1,2|-LALL -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) start cc storcli /c*x*/v(*x*|all) stop cc storcli /c*x*/v(*x*|all) pause cc storcli /c*x*/v(*x*|all) resume cc storcli /c*x*/v(*x*|all) show cc
Enable/disable automatic background initialization. Show, stop, pause, resume, and show the progress of the background initialization.	MegaCLI -LDBI -Enbl|-Dsbl| -getSetting|-Abort|-Suspend| -Resume|-ShowProg|-ProgDsply -Lx|-L0,1,2|-LALL -aN|-a0,1,2|-Aall	storcli /c*x*/v(*x*|all) set autobgi=On|Off storcli /c*x*/v(*x*|all) show autobgi storcli /c*x*/v(*x*|all) stop bgi storcli /c*x*/v(*x*|all) pause bgi storcli /c*x*/v(*x*|all) resume bgi storcli /c*x*/v(*x*|all) show bgi
Start and show progress for a migrate operation.	MegaCli –LDRecon {–Start –Rx [Add | Rmv PhysDrv[E0:S0,E1:S1,...]] } | –ShowProg|-ProgDsply –Lx –aN	storcli /c*x*/v*x* start migrate type=raidx [option=add|remove drives=[*e*:]*s*|[*e*:]*s*-*x*|[*e*:]*s*-*x*,*y*] [Force] storcli /c*x*/v(*x*|all) show migrate
Delete preserved cache.	MegaCLI -DiscardPreservedCache -Lx|-L0,1,2|-Lall -force -aN|-a0,1,2|-aALL	storcli /c*x*/v(*x*|all) delete preservedcache[force]

B.10 Physical Drive Commands

Table 72 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`
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`

Table 72 Physical Drive Commands (Continued)

Description	MegaCLI Command	StorCLI Command
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 /c*x*/e*x*/s*x* start locate storcli /c*x*/e*x*/s*x* 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 /c*x*/e*x*/s*x* spindown storcli /c*x*/e*x*/s*x* spinup
Show physical drive information of all connected drives.	MegaCli –PDList –aN|–a0,1..|–aAll	storcli /c*x*/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 /c*x*[/e*x*]/s*x* download src=<*filepath*> [satabridge] [mode= 5|7] storcli /c*x*/e*x* 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 /c*x*/e*x*/s*x* secureerase [force]
Show the security key for secured physical drives	MegaCli -GetKeyID [–PhysDrv[E0:S0]] -aN	storcli /c*x*/e*x*/s*x* 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 /c*x*[/e*x*]/s*x* start erase [simple | normal | thorough] [erasepatternA=<*val*>]\n[erasepatternB=<*val*>] Examples: storcli /c*x*/e*x*/s*x* start erase simple storcli /c*x*/e*x*/s*x* start erase normal erasepatterna=10101010 storcli /c*x*/e*x*/s*x* start erase thorough erasepatterna=10101010 erasepatternb=10101111 storcli /c*x*/e*x*/s*x* 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 /c*x* set directpdmapping=<on | off> storcli /c*x* show directpdmapping

B.11 Enclosure Commands

Table 73 Enclosure Commands

Description	MegaCLI Command	StorCLI Command
Show enclosure information.	MegaCli –EncInfo –aN|–a0,1,2|–aALL	storcli /c*x*/e*x* show storcli /c*x*/e*x* show all
Show enclosure status.	MegaCli –EncStatus –aN|–a0,1,2|-aALL	storcli /c*x*/e*x* show status

B.12 PHY Commands

Table 74 PHY Commands

Description	MegaCLI Command	StorCLI Command
Show PHY information.	MegaCli –PHYInfo -phyM –aN|–a0,1,2|-aALL	storcli /c*x*/p*x*(*x*|all) show storcli /c*x*/p*x*(*x*|all) show all
Set PHY link speed.	MegaCLI PhySetLinkSpeed -phyM –speed -aN|-a0,1,2|-aALL	storcli /c*x*/p*x*(*x*|all) set linkspeed=0(auto)|1.5|3|6|12
Show the PHY error counters.	Megacli PhyErrorCounters -An	storcli /c*x*/p*x*(*x*|all) show storcli /c*x*/p*x*(*x*|all) show all

B.13 Alarm Commands

Table 75 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>

B.14 Event Log Properties Commands

Table 76 Event Log Properties Commands

Description	MegaCLI Command	StorCLI Command
Show event logs.	`MegaCli -AdpEventLog -GetEventLogInfo -aN|-a0,1,2|-aALL`	`storcli /cx show eventloginfo`
Show the specified type of event logs.	`MegaCli -AdpEventLog -GetEvents {-info -warning -critical –fatal} {-f <fileName>} –aN|–a0,1,2|-aALL`	`storcli /cx show events [[type=<sincereboot| sinceshutdown| includedeleted|latest=x| ccincon vd=<0,1,...>] filter=<info| warning| critical|fatal>] file=<filepath>`
Show the specified event logs.	`MegaCli -AdpEventLog –GetSinceShutdown {-info –warning -critical -fatal} {-f <fileName>} -aN|-a0,1,2|-aALL`	`storcli /cx show events [type=[latest=x|ccincon vd=|[sincereboot|sinceshutdown|includedeleted|latest|ccincon]]] [filter=[info|warning|critical|fatal]] file=xyz.txt`
Delete the event logs.	`MegaCli -AdpEventLog -Clear –aN|-a0,1,2|-aALL`	`storcli /cx delete events`

B.15 Premium Feature Key Commands

Table 77 Premium Feature Key Commands

Description	MegaCLI Command	StorCLI Command
Show the Safe ID of the controller.	`MegaCli -ELF -GetSafeId -a0`	`storcli /cx(x|all) show safeid`
Show the Advanced Software Options that are enabled on the controller, including the ones in trial mode.	`MegaCli -ELF –ControllerFeatures -a0`	`storcli /cx(x|all) show all` **NOTE** This information shows as part of the controller show all.
Apply the Activation Key in preview mode.	`MegaCli -ELF -Applykey key –val -preview -a0`	`storcli /cx(x|all) set aso key=<key value> preview`
Apply the Activation Key.	`MegaCli -ELF -Applykey key –val -a0`	`storcli /cx(x|all) set aso key=<key value>`
Deactivate the trial key.	`MegaCli -ELF –DeactivateTrialKey -a0`	`storcli /cx(x|all) set aso deactivatetrialkey`
Show the re-host information and, if re-hosting is necessary, show the controller and key vault serial numbers.	`MegaCli -ELF -ReHostInfo -a0`	`storcli /cx(x|all) show rehostinfo`
Indicate to the controller that the re-host is complete.	`MegaCli -ELF -ReHostComplete –a0`	`storcli /cx(x|all) set aso rehostcomplete`

Appendix C: Unsupported Commands in Embedded MegaRAID

The commands in the following table are not supported in Embedded MegaRAID.

Table 78 Unsupported Commands in Embedded MegaRAID

Command Group	Command
Jbod	storcli /c0 set jbod=<on|off>
	storcli /c0/s2 set jbod
	storcli /c0/s2 set bootdrive=<on|off>
DS	storcli /c*x*(*x*|all) set ds=OFF type=1|2|3|4
	storcli /c*x*(*x*|all) set ds=ON type=1|2 [properties]
	storcli /c*x*(*x*|all) set ds=ON type=3|4 DefaultLdType=<*val*> [properties]
	storcli /c*x*(*x*|all) set ds [properties]
	storcli /c*x*/v(*x*|all) set ds=Default|Auto|None|Max|MaxNoCache
Security	storcli /c*x* delete security key
	storcli /c*x* set securitykey=*xxxxxxxx* {passphrase=*xxxx*} {keyid=*xxx*}
	storcli /c*x* set securitykey keyid=*xxx*
	storcli /c*x* compare securitykey=*xxxxxxxxxx*
	storcli /c*x* set securitykey=*xxxxxxxx* oldsecuritykey=*xxxxxxxx*
ASO	storcli /c*x*(*x*|all) set aso key=<*keyvalue*> preview
	storcli /c*x*(*x*|all) set aso key=<*key value*>
	storcli /c*x*(*x*|all) set aso transfertovault
	storcli /c*x*(*x*|all) set aso rehostcomplete
	storcli /c*x*(*x*|all) set aso deactivatetrialkey
	storcli /c*x*(*x*|all) show safeid
	storcli /c*x*(*x*|all) show rehostinfo
	storcli /c0 set time =<*yyyymmdd hh*:*mm*:*ss* | system>
	storcli /c0 show cc|consistencycheck
	storcli /c0/vall show expansion
	storcli /c0 set jbod
	storcli /c*x* download src=<*filepath*> [forceActivate]
Copy back	storcli /c*x*[/e*x*]/s*x* show copyback
	storcli /c*x*[/e*x*]/s*x* start copyback target=eID:sID
	storcli /c*x*[/e*x*]/s*x* stop copyback
	storcli /c*x*[/e*x*]/s*x* pause copyback
	storcli /c*x*[/e*x*]/s*x* resume copyback
Migrate	storcli /c*x*/v(*x*|all) show migrate
	storcli /c*x*/v*x* start migrate type=*raidx* [option=add|remove drives=[*e*:]*s*|[*e*:]*s-x*|[*e*:]*s-x,y*] [Force]
Cache	storcli /c*x*/v(*x*|all) set ssdcaching=on|off
	storcli /c*x*(*x*|all) show preservedcache
	storcli /c*x*/v(*x*|all) delete preservedcache[force]

Table 78 Unsupported Commands in Embedded MegaRAID (Continued)

Command Group	Command
BBU	storcli /c*x*/bbu show
	storcli /c*x*/bbu show all
	storcli /c*x*/bbu set [learnDelayInterval=<*val*> | bbuMode=<*val*>
	storcli /c*x*/bbu start learn
Secure ease	storcli /c*x*/s*x* secureerase [force]
	storcli /c*x*/s*x* start erase [simple| normal| thorough][erasepatternA=<*val*>]
	storcli /c*x*/s*x* stop erase
	storcli /c*x*/s*x* show erase
Consistency check	storcli /c*x* show cc/ConsistencyCheck
Controller	storcli /c*x* show cc

Appendix D: CLI Error Messages

This appendix lists the software error messages for the Storage Command Line Tool (StorCLI) and the MegaCLI Configuration Utility.

The Storage Command Line Tool (StorCLI) and the MegaCLI Configuration Utility are command line interface applications you can use to manage ThinkSystem SAS RAID controllers.

D.1 Error Messages and Descriptions

Each message that appears in the event log has an error level that indicates the severity of the event, as shown in the following table.

Table 79 Error Messages and Descriptions

Decimal Number	Hex Number	Event Text
0	0x00	Command completed successfully
1	0x01	Invalid command
2	0x02	DCMD opcode is invalid
3	0x03	Input parameters are invalid
4	0x04	Invalid sequence number
5	0x05	Abort isn't possible for the requested command
6	0x06	Application 'host' code not found
7	0x07	Application already in use - try later
8	0x08	Application not initialized
9	0x09	Given array index is invalid
10	0x0a	Unable to add missing drive to array, as row has no empty slots
11	0x0b	Some of the CFG resources conflict with each other or the current config
12	0x0c	Invalid device ID / select-timeout
13	0x0d	Drive is too small for requested operation
14	0x0e	Flash memory allocation failed
15	0x0f	Flash download already in progress
16	0x10	Flash operation failed
17	0x11	Flash image was bad
18	0x12	Downloaded flash image is incomplete
19	0x13	Flash OPEN was not done
20	0x14	Flash sequence is not active
21	0x15	Flush command failed
22	0x16	Specified application doesn't have host-resident code
23	0x17	LD operation not possible - CC is in progress
24	0x18	LD initialization in progress
25	0x19	LBA is out of range
26	0x1a	Maximum LDs are already configured

Table 79 Error Messages and Descriptions (Continued)

Decimal Number	Hex Number	Event Text
27	0x1b	LD is not OPTIMAL
28	0x1c	LD Rebuild is in progress
29	0x1d	LD is undergoing reconstruction
30	0x1e	LD RAID level is wrong for requested operation
31	0x1f	Too many spares assigned
32	0x20	Scratch memory not available - try command again later
33	0x21	Error writing MFC data to SEEPROM
34	0x22	Required HW is missing (i.e. Alarm or BBU)
35	0x23	Item not found
36	0x24	LD drives are not within an enclosure
37	0x25	PD CLEAR operation is in progress
38	0x26	Unable to use SATA(SAS) drive to replace SAS(SATA)
39	0x27	Patrol Read is disabled
40	0x28	Given row index is invalid
45	0x2d	SCSI command done, but non-GOOD status was received-see mf.hdr.extStatus for SCSI_STATUS
46	0x2e	IO request for MFI_CMD_OP_PD_SCSI failed - see extStatus for DM error
47	0x2f	Matches SCSI RESERVATION_CONFLICT
48	0x30	One or more of the flush operations failed
49	0x31	Firmware real-time currently not set
50	0x32	Command issues while firmware in wrong state (i.e., GET RECON when op not active)
51	0x33	LD is not OFFLINE - IO not possible
52	0x34	Peer controller rejected request (possibly due to resource conflict)
53	0x35	Unable to inform peer of communication changes (retry might be appropriate)
54	0x36	LD reservation already in progress
55	0x37	I2C errors were detected
56	0x38	PCI errors occurred during XOR/DMA operation
57	0x39	Diagnostics failed - see event log for details
58	0x3a	Unable to process command as boot messages are pending
59	0x3b	Returned in case if foreign configurations are incomplete
61	0x3d	Returned in case if a command is tried on unsupported hardware
62	0x3e	CC scheduling is disabled
63	0x3f	PD CopyBack operation is in progress
64	0x40	Selected more than one PD per array
65	0x41	Microcode update operation failed
66	0x42	Unable to process command as drive security feature is not enabled
67	0x43	Controller already has a lock key
68	0x44	Lock key cannot be backed-up
69	0x45	Lock key backup cannot be verified
70	0x46	Lock key from backup failed verification
71	0x47	Rekey operation not allowed, unless controller already has a lock key

Table 79 Error Messages and Descriptions (Continued)

Decimal Number	Hex Number	Event Text
72	0x48	Lock key is not valid, cannot authenticate
73	0x49	Lock key from escrow cannot be used
74	0x4a	Lock key backup (pass-phrase) is required
75	0x4b	Secure LD exist
76	0x4c	LD secure operation is not allowed
77	0x4d	Reprovisioning is not allowed
78	0x4e	Drive security type (FDE or non-FDE) is not appropriate for requested operation
79	0x4f	LD encryption type is not supported
80	0x50	Cannot mix FDE and non-FDE drives in same array
81	0x51	Cannot mix secure and unsecured LD in same array
82	0x52	Secret key not allowed
83	0x53	Physical device errors were detected
84	0x54	Controller has LD cache pinned
85	0x55	Requested operation is already in progress
86	0x56	Another power state set operation is in progress
87	0x57	Power state of device is not correct
88	0x58	No PD is available for patrol read
89	0x59	Controller reset is required
90	0x5a	No EKM boot agent detected
91	0x5b	No space on the snapshot repository VD
92	0x5c	For consistency SET PiTs, some PiT creations might fail and some succeed
255	0xFF	Invalid status - used for polling command completion
93	0x5d	Secondary iButton cannot be used and is incompatible with controller
94	0x5e	PFK doesn't match or cannot be applied to the controller
95	0x5f	Maximum allowed unconfigured (configurable) PDs exist
96	0x60	IO metrics are not being collected
97	0x61	AEC capture needs to be stopped before proceeding
98	0x62	Unsupported level of protection information
99	0x63	PDs in LD have incompatible EEDP types
100	0x64	Request cannot be completed because protection information is not enabled
101	0x65	PDs in LD have different block sizes
102	0x66	LD Cached data is present on a (this) SSCD
103	0x67	Config sequence number mismatch
104	0x68	Flash image is not supported
105	0x69	Controller cannot be online-reset
106	0x6a	Controller booted to safe mode, command is not supported in this mode
107	0x6b	SSC memory is unavailable to complete the operation
108	0x6c	Peer node is incompatible
109	0x6d	Dedicated hot spare assignment is limited to array(s) with same LDs.
110	0x6e	Signed component is not part of the image

Table 79 Error Messages and Descriptions (Continued)

Decimal Number	Hex Number	Event Text
111	0x6f	Authentication failure of the signed firmware image
112	0x70	Flashing was ok but FW restart is not required, ex: No change in FW from current
113	0x71	Firmware is in some form of restricted mode, example: passive in A/P HA mode
114	0x72	The maximum number of entries are exceed.
115	0x73	Cannot start the subsequent flush because the previous flush is still active.
116	0x74	Status is ok but a reboot is need for the change to take effect.
117	0x75	Cannot perform the operation because the background operation is still in progress.
118	0x76	Operation is not possible.
119	0x77	Firmware update on the peer node is in progress.
120	0x78	Hidden policy is not set for all of the virtual drives in the drive group that contains this virtual drive.
121	0x79	Indicates that there are one or more secure system drives in the system.

Appendix E: Support Limitations

This appendix provides information about some known limitations in the MegaRAID 12Gb/s SAS RAID controller:

- Known limitations on 240 VD (240 virtual drives).
- Known limitations on BIOS.
- Known limitations on online firmware upgrade and downgrade.
- Known limitations on enclosure firmware update.

E.1 Host Software Utility

The following host software utilities support matrix provides the support information on the target IDs that are supported.

Table 80 Host Software Utilities Support Matrix

MegaRAID SAS RAID Utilities	0–63 VD Target ID's Support	240 VD Target ID's Support
StorCLI	Yes	Yes
	Yes	No
SNMP	Yes	No
Providers	Yes	No
Human Interface Infrastructure (HII)	Yes	Yes
StoreLib/StoreLib Test	Yes	Yes
StoreLib/StoreLib Test (OOB)	Yes	Yes
Legacy BIOS	Yes	Yes **NOTE** The Option ROM builds INT 13H for the boot VD, which is followed by INT 13H for the first 63 VDs reported in the VD list.

E.2 BIOS Known Limitations

The Legacy Option ROM displays only the first 64 VDs during the power-on self-test (POST). The following example describes the POST behavior when there are 90 VDs in the configuration.

Example:

- The Option ROM displays the first 64 VDs in the POST.
- 90 VDs are found on the host adapter.
- 64 VDs are handled by the BIOS.

E.3 Online Firmware Upgrade and Downgrade

The following sections and table describe some of the known limitations when using the Online Firmware Upgrade feature.

Known Limitations With Online Firmware Upgrade

- For MegaRAID 6.7 Firmware GCA and later, any attempt to directly update the firmware to an older version using the online firmware update (OFU) process is not possible. The user must reboot the server for the older version to take effect. This is because of the product name rebranding effort that has resulted in changing the current VPD data to *Broadcom*, unlike the VPD data in the older firmware version (MegaRAID 6.6 Firmware GCA, and earlier), which is *LSI*. It is important that VPD data is presented the same to the operating system. Discrepancies in the VPD data results in an operating system crash since the operating system considers this critical data. Therefore, if any attempt to directly update the firmware to an older version using the online firmware update (OFU) process results in a change in VPD data (from *Broadcom* to *LSI*) and leads to an OS crash.
- MegaRAID 6.9 Firmware GCA supports 1 MB I/Os. The operating system driver presents this capability to the operating system during the initialization of the driver. However, the operating system driver cannot reinitialize the operating system with new values if there is an online firmware update (OFU) that does not support 1 MB I/Os. For example, OFU is not supported when you downgrade the firmware from MegaRAID 6.9 Firmware GCA to MegaRAID 6.8 Firmware GCA. Due to this operating system driver limitation, downgrading the firmware to an older version (for example, MegaRAID 6.8 Firmware GCA) using the OFU process is not possible when both the firmware and the driver have established 1 MB I/O support. However, firmware flash is allowed.
- If you are doing an online firmware update from a previous version to MegaRAID 6.9 Firmware GCA with large I/O support enabled, you need to reboot the system to enable large I/O support. Until you reboot the system, your operating system will be running with only those features that were available to it when it was initially booted.

Known Limitations With Reconstruction Operation

- From MegaRAID 6.6 Firmware GCA and later, you must back up the logical drive before initiating a reconstruction operation on the logical drive.
- You must not perform any firmware upgrade or downgrade when the reconstruction operation is in progress.
- When you flash a new firmware, you should not start a reconstruction operation until the system reboots or an Online Controller Reset (OCR) is performed.

  NOTE The user must reboot the system for the flashed firmware to take effect.
- When a reconstruction operation is in progress, all virtual drives on the controller, not just the virtual drive on which reconstruction operation is in progress, will go to the Write Through mode. Irrespective of whether there is an optimal CacheVault or BBU, the Write Cache settings for all VDs on the controller will go the Write Through mode. You will not be able to enable the Write Back mode on any VDs on the controller until the reconstruction operation is complete. Depending on the capacity of the drives, the number of drives, and other factors, the time to complete a reconstruction operation can take from hours to days or possibly weeks. A reconstruction operation cannot be aborted once it is started. Running the VD with Write Through cache will greatly decrease the Write performance for all VDs attached to the controller until the reconstruction operation is complete.

Consistency Check, Background Initialization, and Secure Erase Limitation

When you downgrade from a 240-virtual drive supported firmware (MegaRAID 6.6 and later) to a non-240-virtual drive supported firmware (MegaRAID 6.5 and earlier), **Consistency Check**, **Background Initialization**, and **Secure Erase** operations are not resumed.

Downgrading the Driver from 240-VD Support to 64 VD Support (Limitation)

You will be able to create more than 64 VDs even though non-240-VD driver and the new 240-VD firmware are installed on the same system. When more than 64 virtual drives are configured, downgrading the driver to an older version (for example, from MegaRAID 6.6 to MegaRAID 6.5) can cause the virtual drives with target IDs greater than 64 virtual drives to be masked to the host.

Auto-Rebuild Operation Limitation

When you upgrade from a non-240-virtual drive supported firmware (MegaRAID 6.5 and earlier) to a 240-virtual drive supported firmware (MegaRAID 6.6 and later), the auto-rebuild operation may not occur.

Table 81 Online Firmware Upgrade and Downgrade Support Matrix

Release	OFU Downgrade Support	OFU Upgrade Support
MegaRAID 6.6 Firmware GCA and earlier	Yes (MegaRAID 6.6 and earlier)	Yes (MegaRAID 6.6 and later)
MegaRAID 6.7 Firmware GCA	No (MegaRAID 6.6 and earlier)	Yes (MegaRAID 6.7 and later)
MegaRAID 6.8 Firmware GCA	No (MegaRAID 6.7 and earlier)	Yes (MegaRAID 6.8 and later)
MegaRAID 6.9 Firmware GCA	No (MegaRAID 6.7 and earlier)	Yes (MegaRAID 6.8 and later)

E.4 Enclosure Firmware Update

If multiple enclosures are connected in a daisy chain mode, and the enclosure firmware is being flashed on the first enclosure while I/Os are running on the physical/virtual drives on the other daisy-chained enclosure, the firmware may encounter the following issues:

- The controller firmware might encounter Montask if the Write Back volumes exist on the enclosure.
- All the enclosures might get dropped and re-discovered when the first ESM (Enclosure Services Management) firmware update completes.
- Physical drives on the daisy-chained enclosures can go into a shield state.

  To avoid these issues, it is recommended to:

  — Stop the I/Os running on the daisy-chained enclosures before you update the enclosure firmware.
  — Execute the Enclosure Firmware Update in maintenance mode.
  — Import the drives once again.

Appendix F: Boot Messages and BIOS Error Messages

This appendix provides the boot messages and BIOS error messages present in the ThinkSystem firmware.

F.1 Displaying Boot Messages

In platforms that load the UEFI driver first, the noncritical boot messages are discarded. To display a critical boot message, the platform should support driver health, and it should load the driver health formset when the Broadcom UEFI driver returns health status as `configuration required`.

In some systems, the platform supports the driver health protocol and calls the `GetHealthStatus` function automatically during boot time. In such platforms, if a critical boot problem exists, the platform shows a critical message dialog.

In some systems, you have to turn on the option in the system BIOS setup to enable the platform to call the `GetHealthStatus` function during boot time to check the health of the controller. To ensure that the platform supports driver health protocol and checks health during boot time, perform the following steps:

1. Set the controller's boot mode to SOE using CLI or RAID management/configuration application.
2. Connect one drive to the controller.
3. Create a RAID 0 volume.
4. Shut down the system, and remove the drive.
5. Boot the system.

   The following dialog should appear.

Figure 113 Driver Health Protocol Dialog

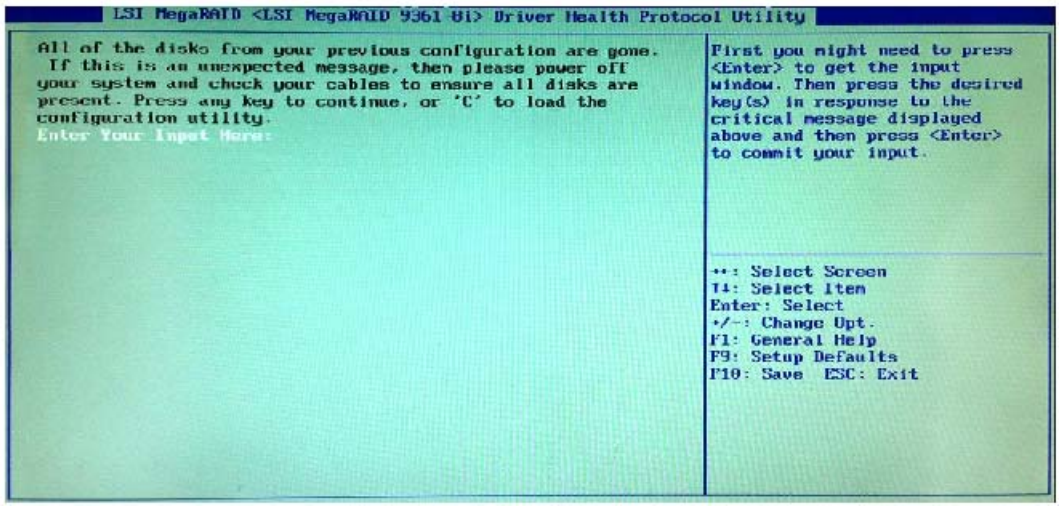

6. Press C.

   The following dialog appears.

Figure 114 Critical Message Completion Dialog

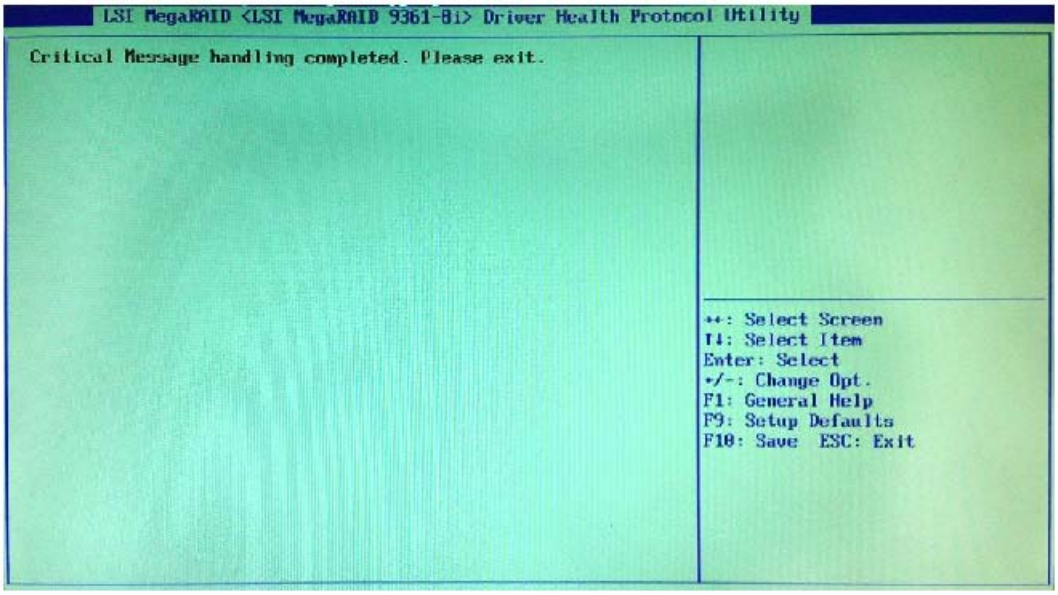

7. Press the Esc key to exit the browser.

The critical message handling completion, the security password, and the confirmation message displayed on the screen are all part of boot messages handled by the controller firmware. The password validation is also done by the controller firmware. The maximum attempt to enter the password is also handled by the firmware.

F.2 Differences in the System Boot Mode

There is a behavioral differences in the controller boot mode (SOE, COE, HCOE, and HSM) and system boot mode (legacy or UEFI). Critical boot messages are reported through events for HSM. Both critical messages and warnings are reported in HCOE mode. The behavioral differences of system boot mode is because of the following:

- Some platforms might load both OpROMs (UEFI and legacy)
- Some platforms might load legacy first, and then the UEFI driver, or vice versa
- Some platforms might load only one OpROM depending upon the system boot mode (legacy versus UEFI)

On a hybrid system that loads the UEFI driver first, the noncritical boot messages are discarded and cannot be read if controller boot mode is set to SOE or COE. If the boot mode is set to HCOE or HSM, you can see the messages in the event log.

The following table describes the boot error messages present in the ThinkSystem firmware.

- **Boot Message Type**: Name or type of the boot message on the firmware.
- **Wait Time**: A time value in seconds where the system waits for the user's input. If the wait time is elapsed, BIOS continues with default options.
  - For example, `BOOT_WAIT_TIME`, where the BIOS waits for the user's input for a default period of time (in seconds) and then continues with the default option if no user input is received.
  - For example, `BOOT_TIME_CRITICAL`, where the BIOS waits for the user's input until an input from the user is received.
- **Event Log**: When any event occurs, the firmware logs that particular event in its database.
- **Boot Message Description**: Boot message displayed on the console.

- **Comments**: Whether the message is associated with any specific controller settings or configuration settings related to the firmware.
- **Troubleshooting Actions**: If applicable, the user can take action to identify, diagnose, and resolve problems associated with the firmware. This can also be best practices, recommendations, and so on.

Table 82 Boot Messages

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
1	BOOT_MSG_CACHE_DISCARD	BOOT_TIME_WAIT	MR_EVT_CTRL_CACHE_DISCARDED	Memory or battery problems were detected. The adapter has recovered, but cached data was lost. Press any key to continue, or press C to load the configuration utility.	—	**Cause**: The cached data is lost and cannot be retrieved. **Action**: Perform memory and battery test. If needed, replace the memory card or the battery.
2	BOOT_MSG_TEST	5	Test boot message	This is a test message. You can press a key to ignore it, or you can wait five seconds. No further action is required. Press any key to continue, or press C to load the configuration utility.	—	N/A
3	BOOT_MSG_CACHE_VERSION	BOOT_TIME_WAIT	MR_EVT_CTRL_CACHE_VERSION_MISMATCH	Firmware version inconsistency was detected. The adapter has recovered, but cached data was lost. Press any key to continue, or press C to load the configuration utility.	—	**Causes**: The cached data is lost and cannot be retrieved. This boot message is displayed when dirty data needs to be flushed during boot. The version of the cache header with which dirty data was generated is different from the current version of the cache header. The version of the cache header is incremented when the cache layout is changed. On a single controller, during firmware upgrade, firmware ensures that there is no dirty data.

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
						This message occurs only when dirty cache or pinned cache is migrated and is stored by ONFI from one controller to another controller where firmware versions on the both the controllers are different. **Action**: Ensure that the other controller also has the same firmware version.
4	BOOT_MSG_DDF_FOREIGN_ FOUND	10	MR_EVT_FOREIGN_CFG_ IMPORTED	Foreign configuration(s) found on adapter. Press any key to continue or press C to load the configuration utility or press F to import foreign configuration(s) and continue.	Use property autoEnhancedImport.	**Cause**: A storage device was inserted with the metadata that does not belong to any RAID volumes recognized by the controller. **Cause**: Either import the configuration settings of the inserted storage device or delete the RAID volume.
5	BOOT_MSG_DDF_IMPORT	10	NULL	Previous configuration cleared or missing. Importing configuration created on %02d/%02d %2d:%02d. Press any key to continue, or press C to load the configuration utility.	Not supported.	**Cause**: The controller is not able to recognize the current RAID volume configuration. **Action**: Either import the configuration settings or delete the foreign configuration found on storage device.
6	BOOT_MSG_PACKAGE_VERSION	0	MR_EVT_PACKAGE_ VERSION	Firmware package: %s	—	N/A
7	BOOT_MSG_FIRMWARE_ VERSION	0	NULL	Firmware version: %s	—	N/A
8	BOOT_MSG_FIRMWARE_TEST	1	NULL	This firmware is a TEST version. It has not completed any validation.	—	**Cause**: The controller is not able to recognize the current RAID volume configuration. **Action**: Update the firmware to the correct version.

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
9	BOOT_MSG_FIRMWARE_ALPHA	1	NULL	This firmware is an ALPHA version – It has not completed all validation. The validation stamp is: %s"""	—	**Cause**: The controller is not able to recognize the current RAID volume configuration. **Action**: Update the firmware to the correct version.
10	BOOT_MSG_FIRMWARE_BETA	1	NULL	This firmware is BETA version – It has not completed all validation. The validation stamp is: %s"""	—	**Cause**: The controller is not able to recognize the current RAID volume configuration. **Action**: Update the firmware to the correct version.
11	BOOT_MSG_SAS_SATA_MIXING_VIOLATION	BOOT_TIME_WAIT	MR_EVT_ENCL_SAS_SATA_MIXING_DETECTED	An enclosure was found that contains both SAS and SATA drives, but this controller does not allow mixed drive types in a single enclosure. Correct the problem then restart your system. Press any key to continue, or press C to load the configuration utility.	—	**Cause**: A single enclosure that has both SAS and SATA drives cannot be used as the controller does not support mixed drive types in a single enclosure. **Actions**: Use only one type of drive, either SAS or SATA drive. Replace the controller with a controller that supports mixed drive types in a single enclosure. Contact Technical Support to enable this feature.
12	BOOT_MSG_SAS_NOT_SUPPORTED	BOOT_TIME_WAIT	SAS drives are not supported.	SAS drives were detected, but this controller does not support SAS drives. Remove the SAS drives then restart your system. Press any key to continue, or press C to load the configuration utility.	—	**Cause**: This controller does not support SAS drives. **Action**: Replace the SAS drives with SATA drives and restart the system.

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
13	BOOT_MSG_SATA_NOT_SUPPORTED	BOOT_TIME_WAIT	SATA drives are not supported.	SATA drives were detected, but this controller does not support SATA drives. Remove the SATA drives then restart your system. Press any key to continue, or press C to load the configuration utility.	—	**Cause**: This controller does not support SATA drives. **Action**: Replace the SATA drives with SAS drives and restart the system.
14	BOOT_MSG_ENCL_COUNT_PER_PORT_EXCEEDED	BOOT_TIME_WAIT	MR_EVT_ENCL_MAX_PER_PORT_EXCEEDED	There are %d enclosures connected to connector %s, but only maximum of %d enclosures can be connected to a single SAS connector. Remove the extra enclosures then restart your system.	—	**Cause**: This controller supports only a particular number of enclosures. **Action**: Remove extra enclosures or insert a controller that supports your enclosure requirements.
15	BOOT_MSG_SAS_TOPOLOGY_ ERROR	BOOT_TIME_WAIT	SAS discovery error	Invalid SAS topology detected. Check your cable configurations, repair the problem, and restart your system.	—	**Cause**: The controller has detected an invalid SAS topology. **Action**: Check the cables or reconfigure the attached devices to create a valid SAS topology.
16	BOOT_MSG_BBU_BAD	10	NULL	The battery is currently discharged or disconnected. Verify the connection and allow 30 minutes for charging. If the battery is properly connected and it has not returned to operational state after 30 minutes of charging then contact technical support for additional assistance.	Not supported.	**Actions**: Check the battery cable to ensure that it is connected properly. Ensure that the battery is charging properly. Contact Technical Support to replace the battery if the battery is draining out.

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
17	BOOT_MSG_BBU_MSG_DISABLE	10	MR_EVT_BBU_NOT_PRESENT	The battery hardware is missing or malfunctioning, or the battery is unconnected, or the battery could be fully discharged. If you continue to boot the system, the battery-backed cache will not function. If battery is connected and has been allowed to charge for 30 minutes and this message continues to appear, contact technical support for assistance. Press D to disable this warning (if your controller does not have a battery)	Use property disableBatteryWarning	**Action**: Check the battery cable to ensure that it is connected properly. Ensure that the battery is charging properly. Contact Technical Support to replace the battery if the battery is draining out.
18	BOOT_MSG_BAD_MFC_SASADDRESS	10	MFC data error! Invalid SAS address	Invalid SAS Address present in MFC data. Program a valid SAS Address and restart your system.	—	**Cause**: Invalid SAS address may be present. **Actions**: 1. Power off the system and remove the controller. 2. Find the SAS address label and re-program the SAS address. Contact Technical Support if you are unable to re-program the SAS address. OEMs can access the StorCLI and re-program the SAS address.

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
19	BOOT_MSG_PDS_MISSING	BOOT_TIME_WAIT	MR_EVT_CTRL_BOOT_MISSING_PDS	Some configured disks have been removed from your system, or are no longer accessible. Check your cables and also make sure all disks are present. Press any key to continue, or press C to load the configuration utility.	—	**Cause**: The controller is unable to find the configured drives. **Actions**: Check if the configured drives are present and they are properly connected. Go to BIOS and check if the devices are displayed. Ensure that the drives are spun-up and have power supplied to them. If there is a backplane, check the connector to ensure that power is being supplied to the drive.
20	BOOT_MSG_LDS_OFFLINE	BOOT_TIME_WAIT	MR_EVT_CTRL_BOOT_LDS_WILL_GO_OFFLINE	The following VDs have missing disks: %s. If you proceed (or load the configuration utility), these VDs will be marked OFFLINE and will be inaccessible. Check your cables and make sure all disks are present. Press any key to continue, or press C to load the configuration utility.	—	**Cause**: The controller is unable to find the configured drives. **Actions**: Check if the configured drives are present and they are properly connected. Go to BIOS and check if the devices are displayed. Ensure that the drives are spun-up and have power supplied to them. If there is a backplane, check the connector to ensure that power is being supplied to the drive.

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
21	BOOT_MSG_LDS_MISSING	BOOT_TIME_WAIT	MR_EVT_CTRL_BOOT_LDS_MISSING	The following VDs are missing: %s. If you proceed (or load the configuration utility), these VDs will be removed from your configuration. If you wish to use them at a later time, they will have to be imported. If you believe these VDs should be present, power off your system and check your cables to make sure all disks are present. Press any key to continue, or press C to load the configuration utility.	—	**Cause**: The controller is unable to find the configured drives. **Actions**: Check if the configured drives are present and they are properly connected. Go to BIOS and check if the devices are displayed. Ensure that the drives are spun-up and have power supplied to them. If there is a backplane, check the connector to ensure that power is being supplied to the drive.
22	BOOT_MSG_LDS_MISSING_SPANS	BOOT_TIME_WAIT	MR_EVT_CTRL_BOOT_LDS_MISSING	The following VDs are missing complete spans: %s. If you proceed (or load the configuration utility), these VDs will be removed from your configuration and the remaining drives marked as foreign. If you wish to use them at a later time, restore the missing span(s) and use a foreign import to recover the VDs. If you believe these VDs should be present, please power off your system and check your cables to make sure all disks are present. Press any key to continue, or press C to load the configuration utility.	—	**Cause**: The controller is unable to find the configured drives. **Actions**: Check if the configured drives are present and they are properly connected. Go to BIOS and check if the devices are displayed. Ensure that the drives are spun-up and have power supplied to them. If there is a backplane, check the connector to ensure that power is being supplied to the drive.

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
23	BOOT_MSG_CONFIG_MISSING	BOOT_TIME_WAIT	MR_EVT_CTRL_BOOT_CONFIG_MISSING	All of the disks from your previous configuration are gone. If this is an unexpected message, power off your system and check your cables to make sure all disks are present. Press any key to continue, or press C to load the configuration utility.	Headless mode – should not appear, if autoEnhancedImport is set.	**Cause**: The controller is unable to find the configured drives. **Actions**: Check if the configured drives are present and they are properly connected. Go to BIOS and check if the devices are displayed. Ensure that the drives are spun-up and have power supplied to them. If there is a backplane, check the connector to ensure that power is being supplied to the drive.
24	BOOT_MSG_CACHE_FLUSH_ NOT_POSSIBLE	BOOT_TIME_CRITICAL	NULL	The cache contains dirty data, but some VDs are missing or will go offline, so the cached data can not be written to disk. If this is an unexpected error, power off your system and check your cables to make sure all disks are present. If you continue, the data in cache will be permanently discarded. Press X to acknowledge and permanently destroy the cached data.	Not supported	**Cause**: The controller is unable to find the configured drives. **Actions**: Check if the configured drives are present and they are properly connected. Go to BIOS and check if the devices are displayed. Ensure that the drives are spun-up and have power supplied to them. If there is a backplane, check the connector to ensure that power is being supplied to the drive.

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
25	BOOT_MSG_LDS_WILL_RUN_ WRITE_THRU	5	NULL	Your VDs that are configured for Write-Back are temporarily running in Write-Through mode. This is caused by the battery being charged, missing, or bad. Allow the battery to charge for 24 hours before evaluating the battery for replacement. The following VDs are affected: %s Press any key to continue.	No event is logged, information for the user	**Actions**: Check the battery cable to ensure that it is connected properly. Ensure that the battery is charging properly. Contact Technical Support to replace the battery if the current supplied by the battery is draining out.
26	BOOT_MSG_MEMORY_INVALID	BOOT_TIME_CRITICAL	NULL	Invalid memory configuration detected. Contact your system support. System has halted.	Not supported	**Action**: Reseat or replace the DIMM.
27	BOOT_MSG_CACHE_DISCARD_WARNING	BOOT_TIME_WAIT	MR_EVT_CTRL_CACHE_ DISCARDED	Cache data was lost due to an unexpected power-off or reboot during a write operation, but the adapter has recovered. This could be because of memory problems, bad battery, or you might not have a battery installed. Press any key to continue or C to load the configuration utility.	Posted only when disableBatteryWarning is set, same as BOOT_MSG_CACHE_DISCARD	**Actions**: Check the battery cable to ensure that it is connected properly. Ensure that the battery is charging properly. Contact Technical Support to replace the battery if power supplied by the battery is draining out.

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
28	BOOT_MSG_CONFIG_CHANGE_WARNING	BOOT_TIME_CRITICAL	NULL	Entering the configuration utility in this state will result in drive configuration changes. Press Y to continue loading the configuration utility or power off your system and check your cables to make sure all disks are present and reboot the system.	Posted from other messages like BOOT_MSG_LDS_MISSING, when the user clicks C.	**Cause**: The controller is unable to find the configured drives. **Actions**: Check if the configured drives are present and they are properly connected. Go to BIOS and check if the devices are displayed. Ensure that the drives are spun-up and have power supplied to them. If there is a backplane, check the connector to ensure that power is being supplied to the drive. If the controller is being used to create a new configuration by reusing the drives, purge the existing data and then continue.
29	BOOT_MSG_EMBEDDED_MULTIBIT_ECC_ERROR	BOOT_TIME_CRITICAL	Multibit ECC error - memory or controller needs replacement.	Multibit ECC errors were detected on the RAID controller. If you continue, data corruption can occur. Contact technical support to resolve this issue. Press X to continue, otherwise power off the system, replace the controller, and reboot.	OEM Specific, see BOOT_MSG_HBA_MULTIBIT_ECC_ERROR for Broadcom Generic message	**Action**: 1. Reseat or replace the DIMM. 2. Restart system. If the problem persists, contact Technical Support.
30	BOOT_MSG_EMBEDDED_SINGLE_BIT_ECC_ERROR	BOOT_TIME_CRITICAL	MR_EVT_CTRL_MEM_ECC_SINGLE_BIT_CRITICAL or WARNING	Single-bit ECC errors were detected on the RAID controller. Contact technical support to resolve this issue. Press X to continue or else power off the system, replace the controller, and reboot.	OEM Specific, see BOOT_MSG_HBA_SINGLE_BIT_ECC_ERROR for Broadcom Generic message	**Action**: 1. Reseat or replace the DIMM. 2. Restart system. If the problem persists, contact Technical Support.

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
31	BOOT_MSG_EMBEDDED_ SINGLE_BIT_OVERFLOW_EC C_ ERROR	BOOT_TIME_CRITI CAL	NULL	Single-bit overflow ECC errors were detected on the RAID controller. If you continue, data corruption can occur. Contact technical support to resolve this issue. Press X to continue or else power off the system, replace the controller, and reboot.	Not supported	**Action**: 1. Reseat or replace the DIMM. 2. Restart system. If the problem persists, contact Technical Support.
32	BOOT_MSG_HBA_MULTIBIT _ ECC_ERROR	BOOT_TIME_CRITI CAL	Multibit ECC error – memory or controller needs replacement.	Multibit ECC errors were detected on the RAID controller. The DIMM on the controller needs replacement. Contact technical support to resolve this issue. If you continue, data corruption can occur. Press X to continue, otherwise power off the system and replace the DIMM module and reboot. If you have replaced the DIMM press X to continue.	—	**Action**: 1. Reseat or replace the DIMM. 2. Restart system. If the problem persists, contact Technical Support.
33	BOOT_MSG_HBA_SINGLE_B IT_ ECC_ERROR	BOOT_TIME_CRITI CAL	MR_EVT_CTRL_MEM_ ECC_SINGLE_BIT_CRITI CAL or WARNING	Single-bit ECC errors were detected during the previous boot of the RAID controller. The DIMM on the controller needs replacement. Contact technical support to resolve this issue. Press X to continue, otherwise power off the system and replace the DIMM module and reboot. If you have replaced the DIMM press X to continue.	—	**Action**: 1. Reseat or replace the DIMM. 2. Restart system. If the problem persists, contact Technical Support.

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
34	BOOT_MSG_HBA_SINGLE_BIT_OVERFLOW_ECC_ERROR	BOOT_TIME_CRITICAL	NULL	Single-bit overflow ECC errors were detected during the previous boot of the RAID controller. The DIMM on the controller needs replacement. Contact technical support to resolve this issue. If you continue, data corruption can occur. Press X to continue, otherwise power off the system and replace the DIMM module and reboot. If you have replaced the DIMM press X to continue.	Not supported	**Action**: 1. Reseat or replace the DIMM. 2. Restart system. If the problem persists, contact Technical Support.
35	BOOT_MSG_ENCL_VIOLATION_MODE	BOOT_TIME_CRITICAL	MR_EVT_CTRL_CRASH	The attached enclosure does not support in controller's Direct mapping mode. Contact your system support. The system has halted because of an unsupported configuration.	Should be able to enter HSM	**Causes**: Too many chained enclosures may be present. May also be related to a security feature in the drive. **Actions**: Remove the drives that are not supported. Reduce the number of drives. Replace the enclosure with an other one. Ensure that the firmware version is updated. Contact Technical Support if the problem persists.
36	BOOT_MSG_EXP_VIOLATION_ FORCE_REBOOT	10	MR_EVT_CTRL_CRASH	Expander detected in controller with direct mapping mode. Reconfiguring automatically to persistent mapping mode. Automatic reboot would happen in 10 seconds.	OEM Specific action, see BOOT_MSG_ENCL_VIOLATION_MODE for LSI generic	**Action**: No action required. The controller will configure itself to a persistent mapping mode and then reboot. Contact Technical Support if problem persists.

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
37	BOOT_MSG_8033X_ATU_ISSUE	BOOT_TIME_CRITICAL	NULL	Your controller's I/O processor has a fault that can potentially cause data corruption. Your controller needs replacement. Contact your system support. To continue, press Y to acknowledge.	DEPRECATED	**Action**: Contact Technical Support for replacement of the controller.
38	BOOT_MSG_MAX_DISKS_EXCEEDED	BOOT_TIME_CRITICAL	MR_EVT_PD_NOT_SUPPORTED	The number of disks exceeded the maximum supported count of %d disks. Remove the extra drives and reboot system to avoid losing data. Press Y to continue with extra drives.	—	**Actions**: Power off the system and remove the controller. Remove the extra drives to reduce the size of the topology. Replace the controller with a controller that supports a larger topology.
39	BOOT_MSG_MAX_DISKS_EXCEEDED_PER_QUAD	BOOT_TIME_CRITICAL	NULL	The number of devices exceeded the maximum limit of devices per quad. Remove the extra drives and reboot the system to avoid losing data System has halted due to unsupported configuration.	Not supported	**Actions**: Power off the system and remove the controller. Remove the extra drives to reduce the size of the topology. Replace the controller with a controller that supports a larger topology.
40	BOOT_MSG_DISCOVERY_ERROR	BOOT_TIME_CRITICAL	Discovery errors – power cycle system and drives, and try again.	A discovery error has occurred, power cycle the system and all the enclosures attached to this system.	—	**Actions**: Shutdown and restart the system as well as all the enclosures attached to the system. Ensure that all the cables are connected and connected properly. Reduce the topology in case of a bad drive. If the problem persists, collect the logs of the system, driver, and firmware and contact Technical Support.

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
41	BOOT_MSG_CTRL_SECRET_KEY_FIRST	BOOT_TIME_WAIT	NULL	Drive security is enabled on this controller and a pass phrase is required. Enter the pass phrase.	Requires user input, if undesired, change Security binding	**Action**: Enter the pass phrase.
42	BOOT_MSG_CTRL_SECRET_KEY_RETRY	BOOT_TIME_WAIT	NULL	Invalid pass phrase. Enter the pass phrase.	opRom must be enabled for user input, if undesired, change Security binding	**Action**: Enter the pass phrase.
43	BOOT_MSG_CTRL_LOCK_KEY_ INVALID	BOOT_TIME_WAIT	MR_EVT_CTRL_LOCK_KEY_FAILED	There was a drive security key error. All secure drives will be marked as foreign. Press any key to continue, or C to load the configuration utility.	—	**Action**: Check if the controller supports self-encrypting drives.
44	BOOT_MSG_KEY_MISSING_REBOOT_OR_CONTINUE	BOOT_TIME_WAIT	MR_EVT_CTRL_LOCK_KEY_FAILED	Invalid pass phrase. If you continue, a drive security key error will occur and all secure configurations will be marked as foreign. Reboot the machine to retry the pass phrase or press any key to continue.	—	**Action**: Restart the system to retry the pass phrase or press any key to continue.
45	BOOT_MSG_KEY_EKMS_FAILURE	BOOT_TIME_WAIT	MR_EVT_CTRL_LOCK_KEY_EKM_FAILURE	Unable to communicate to EKMS. If you continue, there will be a drive security key error and all secure configurations will be marked as foreign. Check the connection with the EKMS, reboot the machine to retry the EKMS or press any key to continue.	—	**Action**: Check the connection of EKMS, restart the system to re-establish the connection to EKMS.
46	BOOT_MSG_REKEY_TO_EKMS_ FAILURE	BOOT_TIME_WAIT	MR_EVT_CTRL_LOCK_KEY_REKEY_FAILED	Unable to change security to EKMS as not able to communicate to EKMS. If you continue, the drive security will remain to existing security mode. Check the connection with the EKMS, reboot the machine to retry the EKMS or press any key to continue.	—	**Action**: Check the connection of EKMS, restart the system to re-establish the connection to EKMS.

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
47	BOOT_MSG_KEY_EKMS_ FAILURE_MERCURY	20	MR_EVT_CTRL_LOCK_ KEY_EKM_FAILURE	DKM existing key request failed; existing secure configurations will be labeled foreign and will not be accessible. Reboot the server to retry.	OEM Specific, see BOOT_MSG _KEY_EKMS _FAILURE for Broadcom generic	**Action**: Check the connection of EKMS, restart the system to re-establish the connection to EKMS.
48	BOOT_MSG_REKEY_TO_EK MS_ FAILURE_MERCURY	BOOT_TIME_CRITI CAL	MR_EVT_CTRL_LOCK_ KEY_REKEY_FAILED	DKM new key request failed; controller security mode transition was not successful. Reboot the server to retry request, or press any key to continue.	OEM Specific, see BOOT_MSG _REKEY_TO _EKMS_FAIL URE for Broadcom generic	**Action**: Check the connection of EKMS, restart the system to re-establish the connection to EKMS.
49	BOOT_MSG_NVDATA_IMAG E_ MISSING	BOOT_TIME_WAIT	NVDATA image is invalid – reflash NVDATA image	Firmware did not find valid NVDATA image. Program a valid NVDATA image and restart your system. Press any key to continue.	—	**Actions**: Flash the correct firmware package that has proper NV Data image. Check the current firmware version, and if needed, updated to the latest firmware version. Updating to the latest firmware version may require importing foreign volumes.
50	BOOT_MSG_IR_MR_MIGRA TION_FAILED	BOOT_TIME_WAIT	IR to MR migration failed.	IR to MR Migration failed. Press any key to continue with MR defined NVDATA values	—	N/A
51	BOOT_MSG_DUAL_BAT_PR SNT	10	NULL	Two BBUs are connected to the adapter. This is not a supported configuration. Battery and caching operations are disabled. Remove one BBU and reboot to restore battery and caching operations. If dirty cache is lost in this boot, that could have been because of dual battery presence.	Not supported	**Actions**: Remove one BBU and restart the system to restore battery and caching operations. Due to the presence of a dual battery, you may lose the data in dirty cache while restarting the system.

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
52	BOOT_MSG_LDS_CACHE_PINNED	10	MR_EVT_CTRL_BOOT_LDS_CACHE_PINNED	Offline or missing virtual drives with preserved cache exist. Check the cables and make sure that all drives are present. Press any key to continue, or C to load the configuration utility.	Use property allowBootWithPinnedCache	**Cause**: The controller is unable to find the configured drives. **Actions**: Check if the configured drives are present and they are properly connected. Go to BIOS and check if the devices are displayed. Ensure that the drives are spun-up and have power supplied to them. If there is a backplane, check the connector to ensure that power is being supplied to the drive. Cache offload occurs if the missing drive is restored.
53	BOOT_MSG_LDS_CACHE_PINNED_HALT	BOOT_TIME_CRITICAL	MR_EVT_CTRL_BOOT_LDS_CACHE_PINNED	Offline or missing virtual drives with preserved cache exist. Check the cables and make sure that all drives are present. Press any key to enter the configuration utility.	If property allowBootWithPinnedCache is disabled	**Cause**: The controller is unable to find the configured drives. **Actions**: Check if the configured drives are present and they are properly connected. Go to BIOS and check if the devices are displayed.Ensure that the drives are spun-up and have power supplied to them. If there is a backplane, check the connector to ensure that power is being supplied to the drive. Cache offload occurs if the missing drive is restored.

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
54	BOOT_MSG_BAD_SBR_ SASADDRESS	BOOT_TIME_CRITI CAL	NULL	Invalid SAS Address present in SBR. Contact your system support. Press any key to continue with Default SAS Address.	Not supported	**Cause**: Invalid SAS address present in the SBR. **Action**: Contact Technical Support to restore to the factory default values.
55	BOOT_MSG_INCOMPATIBLE _ SECONDARY_IBUTTON	BOOT_TIME_CRITI CAL	Incompatible secondary iButton detected	Incompatible secondary iButton present! Insert the correct iButton and restart the system. Press any key to continue but OEM specific features will not be upgraded!	—	**Actions**: Insert the correct iButton or key-vault and restart the system. If problem persists, contact Technical Support for replacement of the iButton or key-vault.
56	BOOT_MSG_CTRL_ DOWNGRADE_DETECTED	BOOT_TIME_CRITI CAL	NULL	Upgrade Key Missing! An upgrade key was present on a previous power cycle, but it is not connected. This can result in inaccessible data unless it is addressed. Re-attach the upgrade key and reboot.	Not supported	**Cause**: An upgrade key that was present on a previous power cycle may not be connected. **Actions**: Reattach the upgrade key and restart the system. If the problem persists, contact Technical Support for replacement of the upgrade key.
57	BOOT_MSG_DDF_MFC_ INCOMPATIBLE	BOOT_TIME_WAIT	Native configuration is not supported, check MFC.	The native configuration is not supported by the controller. Check the controller, iButton or key-vault. If you continue the configuration will be marked foreign. Press any key to continue.	—	**Actions**: Insert the correct iButton or key-vault and restart the system. If problem persists, contact Technical Support for replacement of the iButton or key-vault.

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
58	BOOT_MSG_BBU_MSG_DISABLE_PERC	10	MR_EVT_BBU_NOT_PRESENT or REMOVED	The battery is currently discharged or disconnected. Verify the connection and allow 30 minutes for charging. If the battery is properly connected and it has not returned to operational state after 30 minutes of charging, contact technical support for additional assistance. Press D to disable this warning (if your controller does not have a battery).	Use property disableBatteryWarning, OEM Specific, also see BOOT_MSG_BBU_MSG_DISABLE	**Actions**: Check the battery cable to ensure that it is connected properly. Ensure that the battery is charging properly. Contact Technical Support to replace the battery if power supplied by the battery is draining out.
59	BOOT_MSG_LDS_WILL_RUN_ WRITE_THRU_PERC	5	NULL	The battery is currently discharged or disconnected. VDs configured in Write-Back mode will run in Write-Through mode to protect your data and will return to the Write-Back policy when the battery is operational. If VDs have not returned to Write-Back mode after 30 minutes of charging then contact technical support for additional assistance. The following VDs are affected: %s. Press any key to continue.	No event is logged, information for the user	**Actions**: Check the battery cable to ensure that it is connected properly. Ensure that the battery is charging properly. Contact Technical Support to replace the battery if the battery is draining out.
60	BOOT_MSG_CACHE_DISCARD_WARNING_PERC	BOOT_TIME_WAIT	MR_EVT_CTRL_CACHE_ DISCARDED	Cache data was lost, but the controller has recovered. This could be because your controller had protected cache after an unexpected power loss and your system was without power longer than the battery backup time. Press any key to continue or C to load the configuration utility.	Property disableBatteryWarning is set	**Actions**: Check the memory and the battery. Check the voltage levels and cache offload timing in case of power loss. If necessary, replace the memory or battery.

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
61	BOOT_MSG_ROLLBACK_ACTIVE	BOOT_TIME_CRITICAL	NULL	A snapshot rollback is in progress on VDs %s, the controller cannot boot until the rollback operation completes. Press any key to enter the configuration utility.	opRom must be enabled, if undesired, do not request rollback. Not supported in ThinkSystem 12Gb/s SAS RAID controllers	**Actions**: Wait for some time until the rollback is complete.
62	BOOT_MSG_ROLLBACK_ACTIVE_REPOSITORY_MISSING	BOOT_TIME_CRITICAL	Rollback requested, but repository is missing	The following VDs: %s have Rollback active and the corresponding Repository is missing. If you continue to boot the system or enter the configuration utility, these VDs will become unusable. Press any key to Continue.	Not supported in ThinkSystem 12Gb/s SAS RAID controllers	**Cause**: This may be related to the snapshot feature, which is not supported on ThinkSystem 12Gb/s SAS RAID controllers. **Action**: Wait for some time until the rollback is complete.
63	BOOT_MSG_REPOSITORY_MISSING	BOOT_TIME_WAIT	Snapshot repository is missing, snapshot disabled	Snapshot Repository VDs %s have been removed from your system, or are no longer accessible. Check the cables and make sure all disks are present. If you continue to boot the system, the snapshot related data will be lost. Press any key to continue, or C to load the configuration utility.	Not supported in ThinkSystem 12Gb/s SAS RAID controllers	**Cause**: The controller is unable to find the configured drives. **Actions**: Check if the configured drives are present and they are properly connected. Go to BIOS and check if the devices are displayed. Ensure that the drives are spun-up and have power supplied to them. If there is a backplane, check the connector to ensure that power is being supplied to the drive.

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
64	BOOT_MSG_CFG_CMD_LOST	BOOT_TIME_WAIT	MR_EVT_CFG_CMD_LOST	The most recent configuration command could not be committed and must be retried. Press any key to continue, or C to load the configuration utility.	—	N/A
65	BOOT_MSG_CFG_CHANGES_ LOST	10	Configuration command was not committed, please retry	Firmware could not synchronize the configuration or property changes for some of the VD's/PD's. Press any key to continue, or C to load the configuration utility.	—	**Actions**: If the same problem persists, contact Technical Support.
66	BOOT_MSG_CFG_ONBOARD_ EXP_NOT_DETECTED	BOOT_TIME_CRITICAL	On-board expander FW or mfg image is corrupted – reflash image	On-board expander firmware or manufacturing image is corrupted. The flash expander firmware and manufacturing image use the recovery tools.	—	**Actions**: Contact Technical Support for factory-only tools to assist in recovery of the expander.
67	BOOT_MSG_PFK_INCOMPATIBLE	BOOT_TIME_WAIT	MFC record not found, ensure you have the correct FW version	The native configuration is not supported by the current firmware. Make sure that the correct controller firmware is being used. If you continue, the configuration will be marked as foreign. Press any key to continue.	—	**Actions**: Collect the logs of the system, driver, and firmware. Ensure that the firmware version corrected and is updated to the latest version. Contact Technical Support if the problem persists.
68	BOOT_MSG_INVALID_FOREIGN_CFG_IMPORT	5	MR_EVT_FOREIGN_CFG_ AUTO_IMPORT_NONE	Foreign configuration import did not import any drives. Press any key to continue.	—	**Actions**: Check the firmware version of the controller. Replace the controller and try again. If the problem persists, contact Technical Support.
69	BOOT_MSG_UPGRADED_IMR_ TO_MR	2	Reboot required to complete the iMR to MR upgrade	Valid memory detected. Firmware is upgraded from iMR to MR. Reboot the system for the MR firmware to run.	—	N/A
70	BOOT_MSG_PFK_ENABLED_AT_BOOT_TIME	BOOT_TIME_WAIT	BOOT_MSG_EVENT_USE_ BOOT_MSG	Advanced software options keys were detected, features activated – %s.	—	N/A

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
71	BOOT_MSG_PFK_DISABLED_ AT_BOOT_TIME	BOOT_TIME_WAIT	BOOT_MSG_EVENT_USE_ BOOT_MSG	Advanced software options keys were missing, features deactivated – %s.	—	**Actions**: Check the cable connection. Check for the Advanced Software Options key. If the problem persists, contact Technical Support.
72	BOOT_MSG_EEPROM_ERROR_ FEATURES_DISABLED	BOOT_TIME_CRITICAL	Cannot communicate with iButton, possible extreme temps.	Cannot communicate with iButton to retrieve premium features. This is probably because of extreme temperatures. The system has halted!	—	**Actions**: Check the cable connection. Ensure that iButton is present. Check the ambient temperature near the iButton. If the problem persists, contact Technical Support.
73	BOOT_MSG_DC_ON_ DEGRADED_LD	BOOT_TIME_CRITICAL	Multiple power loss detected with I/O transactions to non optimal VDs.	Consecutive power loss detected during I/O transactions on non-optimal write-back volumes. This might have resulted in data integrity issues. Press 'X' to proceed.	—	**Actions**: Check if the controller is securely locked in the PCI slot. Check the power supply, battery, and Supercap. If you find any hardware defect, contact Technical Support.

Table 82 Boot Messages (Continued)

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
74	BOOT_MSG_USB_DEVICE_ERROR	BOOT_TIME_CRITICAL	USB cache device is not responding.	USB cache device is not responding. Power down the system for 2 minutes to attempt recovery and avoid cache data loss, and then power-on.	Not supported in ThinkSystem 12Gb/s SAS RAID controllers.	This message is not applicable to ThinkSystem 12Gb/s SAS RAID controllers because the 3108 controller supports ONFI-based cache offload. **Actions**: The 2208 controller supports USB cache offload. Ensure that USB cache is present and secure. Reseat and replace the USB cache. Power off the system for 2 minutes to attempt recovery and avoid cache data loss, then power on the system.
75	BOOT_MSG_DOWNGRADE_MR_TO_IMR	BOOT_TIME_CRITICAL	Bad or missing RAID controller memory module detected.	Bad or missing RAID controller memory module detected. Press D to downgrade the RAID controller to iMR mode. **Warning!** Downgrading to iMR mode, might result in incompatible Logical drives. Press any other key to continue, controller shall boot to safe mode.	—	**Actions**: 1. Reseat or replace the DIMM. 2. Restart system. If the problem persists, contact Technical Support for repair or replacement.
76	BOOT_MSG_HEADLESS_DUMMY	0	NULL	—	—	N/A
77	BOOT_MSG_LIST_TERMINATOR	0	NULL	—	—	N/A

Appendix G: Glossary

This glossary defines the terms used in this document.

A

Absolute state of charge	Predicted remaining battery capacity expressed as a percentage of Design Capacity. Note that the Absolute State of Charge operation can return values greater than 100 percent.
Access policy	A virtual drive property indicating what kind of access is allowed for a particular virtual drive. The possible values are *Read/Write*, *Read Only*, or *Blocked*.
Alarm enabled	A controller property that indicates whether the controller's onboard alarm is enabled.
Alarm present	A controller property that indicates whether the controller has an onboard alarm. If present and enabled, the alarm is sounded for certain error conditions.
Array	See *drive group*.
Auto learn mode	The controller performs the learn cycle automatically in this mode. This mode offers the following options: ■ BBU Auto Learn: Firmware tracks the time since the last learn cycle and performs a learn cycle when due. ■ BBU Auto Learn Disabled: Firmware does not monitor or initiate a learn cycle. You can schedule learn cycles manually. ■ BBU Auto Learn Warn: Firmware warns about a pending learn cycle. You can initiate a learn cycle manually. After the learn cycle is complete, the firmware resets the counter and warns you when the next learn cycle time is reached.
Auto learn period	Time between learn cycles. A learn cycle is a battery calibration operation performed periodically by the controller to determine the condition of the battery.
Average time to empty	One-minute rolling average of the predicted remaining battery life.
Average time to full	Predicted time to charge the battery to a fully charged state based on the one minute rolling average of the charge current.

B

Battery module version	Current revision of the battery pack module.
Battery replacement	Warning issued by firmware that the battery can no longer support the required data retention time.
Battery retention time	Time, in hours, that the battery can maintain the contents of the cache memory.
Battery status	Operating status of the battery. Possible values are Missing, Optimal, Failed, Degraded (need attention), and Unknown.
Battery type	Possible values are intelligent Battery Backup Unit (BBU), intelligent Battery Backup Unit (iBBU), intelligent Transportable Battery Backup Unit (iTBBU), and ZCR Legacy.
BBU present	A controller property that indicates whether the controller has an onboard battery backup unit to provide power in case of a power failure.
BGI rate	A controller property indicating the rate at which the background initialization of virtual drives will be carried out.
BIOS	Basic Input/Output System. The computer BIOS is stored on a flash memory chip. The BIOS controls communications between the microprocessor and peripheral devices, such as the keyboard and the video controller, and miscellaneous functions, such as system messages.

C

Cache	Fast memory that holds recently accessed data. Use of cache memory speeds subsequent access to the same data. When data is read from or written to main memory, a copy is also saved in cache memory with the associated main memory address. The cache memory software monitors the addresses of subsequent reads to see if the required data is already stored in cache memory. If it is already in cache memory (a cache hit), it is read from cache memory immediately and the main memory read is aborted (or not started). If the data is not cached (a cache miss), it is fetched from main memory and saved in cache memory.
Cache flush interval	A controller property that indicates how often the data cache is flushed.
Caching	The process of using a high speed memory buffer to speed up a computer system's overall read/write performance. The cache can be accessed at a higher speed than a drive subsystem. To improve read performance, the cache usually contains the most recently accessed data, as well as data from adjacent drive sectors. To improve write performance, the cache can temporarily store data in accordance with its write back policies.
Capacity	A property that indicates the amount of storage space on a drive or virtual drive.
Coerced capacity	A drive property indicating the capacity to which a drive has been coerced (forced) to make it compatible with other drives that are nominally the same capacity. For example, a 4-GB drive from one manufacturer might be 4,196 MB, and a 4-GB from another manufacturer might be 4,128 MB. These drives could be coerced to a usable capacity of 4,088 MB each for use in a drive group in a storage configuration.
Coercion mode	A controller property indicating the capacity to which drives of nominally identical capacity are coerced (forced) to make them usable in a storage configuration.
Consistency check	An operation that verifies that all stripes in a virtual drive with a redundant RAID level are consistent and that automatically fixes any errors. For RAID 1 drive groups, this operation verifies correct mirrored data for each stripe.
Consistency check rate	The rate at which consistency check operations are run on a computer system.
Controller	A chip that controls the transfer of data between the microprocessor and memory or between the microprocessor and a peripheral device such as a drive. RAID controllers perform RAID functions such as striping and mirroring to provide data protection.
Copyback	The procedure used to copy data from a source drive of a virtual drive to a destination drive that is not a part of the virtual drive. The copyback operation is often used to create or restore a specific physical configuration for a drive group (for example, a specific arrangement of drive group members on the device I/O buses). The copyback operation can be run automatically or manually.Typically, a drive fails or is expected to fail, and the data is rebuilt on a hot spare. The failed drive is replaced with a new drive. Then the data is copied from the hot spare to the new drive, and the hot spare reverts from a rebuild drive to its original hot spare status. The copyback operation runs as a background activity, and the virtual drive is still available online to the host.
Current	Measure of the current flowing to (+) or from (-) the battery, reported in milliamperes.
Current write policy	A virtual drive property that indicates whether the virtual drive currently supports Write Back mode or Write Through mode. ■ In Write Back mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. ■ In Write Through mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction.
Cycle count	The count is based on the number of times the near fully charged battery has been discharged to a level below the cycle count threshold.

D

Default write policy	A virtual drive property indicating whether the default write policy is Write Through or Write Back. In Write Back mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. In Write Through mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction.
Design capacity	Designed charge capacity of the battery, measured in milliampere-hour units (mAh).
Design charge capacity remaining	Amount of the charge capacity remaining, relative to the battery pack design capacity.
Design voltage	Designed voltage capacity of the battery, measured in millivolts (mV).
Device chemistry	Possible values are NiMH (nickel metal hydride) and LiON (lithium ion).
Device ID	A controller or drive property indicating the manufacturer-assigned device ID.
Device port count	A controller property indicating the number of ports on the controller.
Drive cache policy	A virtual drive property indicating whether the virtual drive cache is enabled, disabled, or unchanged from its previous setting.
Drive group	A group of drives attached to a RAID controller on which one or more virtual drives can be created. All virtual drives in the drive group use all of the drives in the drive group.
Drive state	A physical drive or a virtual drive property indicating the status of the appropriate drive. **Physical Drive State** A physical drive can be in any one of the following states: ■ **Unconfigured Good** – A drive accessible to the RAID controller but not configured as a part of a virtual drive or as a hot spare. In the output of the StorCLI commands, **Unconfigured Good** is displayed as **UGood**. ■ **Hot Spare** – A drive that is configured as a hot spare. ■ **Online** – A drive that can be accessed by the RAID controller and will be part of the virtual drive. In the output of the StorCLI commands, **Online** is displayed as **onln**. ■ **Rebuild** – A drive to which data is being written to restore full redundancy for a virtual drive. ■ **Failed** – A drive that was originally configured as Online or Hot Spare, but on which the firmware detects an unrecoverable error. ■ **Unconfigured Bad** – A drive on which the firmware detects an unrecoverable error; the drive was Unconfigured Good or the drive could not be initialized. In the output of the StorCLI commands, **Unconfigured Bad** is displayed as **UBad**. ■ **Missing** – A drive that was Online, but which has been removed from its location. ■ **Offline** – A drive that is part of a virtual drive but which has invalid data as far as the RAID configuration is concerned. In the output of the StorCLI commands, **Offline** is displayed as **offln**. ■ **None** – A drive with an unsupported flag set. An Unconfigured Good or Offline drive that has completed the prepare for removal operation.

Virtual Drive State

A virtual drive can be in any one of the following states:

- **Optimal** – A virtual drive whose members are all online.

  In the output of the StorCLI commands, **Optimal** is displayed as **optl**.
- **Partially Degraded** – A virtual drive with a redundant RAID level that is capable of sustaining more than one member drive failure. This state also applies to the virtual drive's member drives. Currently, a RAID 6 or RAID 60 virtual drive is the only virtual drive that can be partially degraded.

  In the output of the StorCLI commands, **Partially Degraded** is displayed as **Pdgd**.
- **Degraded** – A virtual drive with a redundant RAID level with one or more member failures and can no longer sustain a subsequent drive failure.

  In the output of the StorCLI commands, **Degraded** is displayed as **dgrd**.
- **Offline** - A virtual drive with on e or more member failures that make the data inaccessible.

  In the output of the StorCLI commands, **Offline** is displayed as **OfLn**.

Drive state drive subsystem — A collection of drives and the hardware that controls them and connects them to one or more controllers. The hardware can include an intelligent controller, or the drives can attach directly to a system I/O bus controller.

Drive type — A drive property indicating the characteristics of the drive.

E

EKM — External Key Management

Estimated time to recharge — Estimated time necessary to complete recharge of the battery at the current charge rate.

Expected margin of error — Indicates how accurate the reported battery capacity is in terms of percentage.

F

Fast initialization — A mode of initialization that quickly writes zeroes to the first and last sectors of the virtual drive. This allows you to immediately start writing data to the virtual drive while the initialization is running in the background.

Fault tolerance — The capability of the drive subsystem to undergo a single drive failure per drive group without compromising data integrity and processing capability. ThinkSystem SAS RAID controllers provides fault tolerance through redundant drive groups in RAID levels 1, 5, 6, 10, 50, and 60. They also support hot spare drives and the auto-rebuild feature.

Firmware — Software stored in read-only memory (ROM) or programmable ROM (PROM). Firmware is often responsible for the behavior of a system when it is first turned on. A typical example would be a monitor program in a system that loads the full operating system from drive or from a network and then passes control to the operating system.

Foreign configuration — A RAID configuration that already exists on a replacement set of drives that you install in a computer system. LSI Storage Authority Software allows you to import the existing configuration to the RAID controller, or you can clear the configuration so you can create a new one.

Formatting — The process of writing a specific value to all data fields on a drive, to map out unreadable or bad sectors. Because most drives are formatted when manufactured, formatting is usually done only if a drive generates many media errors.

Full charge capacity — Amount of charge that can be placed in the battery. This value represents the last measured full discharge of the battery. This value is updated on each learn cycle when the battery undergoes a qualified discharge from nearly full to a low battery level.

G

Gas gauge status — Hexadecimal value that represents the status flag bits in the gas gauge status register.

H

Hole	In the LSI Storage Authority Software, a *hole* is a block of empty space in a drive group that can be used to define a virtual drive.
Host interface	A controller property indicating the type of interface used by the computer host system: for example, *PCIX*.
Host port count	A controller property indicating the number of host data ports currently in use.
Host system	Any computer system on which the controller is installed. Mainframes, workstations, and standalone desktop systems can all be considered host systems.
Hot spare	A standby drive that can automatically replace a failed drive in a virtual drive and prevent data from being lost. A hot spare can be dedicated to a single redundant drive group or it can be part of the global hot spare pool for all drive groups controlled by the controller.When a drive fails, LSI Storage Authority Software automatically uses a hot spare to replace it and then rebuilds the data from the failed drive to the hot spare. Hot spares can be used in RAID 1, 5, 6, 10, 50, and 60 storage configurations.
I	
Initialization	The process of writing zeros to the data fields of a virtual drive and, in fault-tolerant RAID levels, generating the corresponding parity to put the virtual drive in a Ready state. Initialization erases all previous data on the drives. Drive groups will work without initializing, but they can fail a consistency check because the parity fields have not been generated.
IO policy	A virtual drive property indicating whether Cached I/O or Direct I/O is being used. In Cached I/O mode, all reads are buffered in cache memory. In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to cache and the host concurrently. If the same data block is read again, it comes from cache memory. (The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.)
L	
LDBBM	Logical drive bad block management
Learn delay interval	Length of time between automatic learn cycles. You can delay the start of the learn cycles for up to 168 hours (seven days).
Learning cycle	A battery calibration operation performed by a RAID controller periodically to determine the condition of the battery. You can start battery learn cycles manually or automatically
Learn mode	Mode for the battery auto learn cycle. Possible values are Auto, Disabled, and Warning.
Learn state	Indicates that a learn cycle is in progress.
LKM	Local Key Management
Load-balancing	A method of spreading work between two or more computers, network links, CPUs, drives, or other resources. Load balancing is used to maximize resource use, throughput, or response time.
Low-power storage mode	Storage mode that causes the battery pack to use less power, which save battery power consumption.
M	
Manufacturing date	Date on which the battery pack assembly was manufactured.
Manufacturing name	Device code that indicates the manufacturer of the components used to make the battery assembly.

Max error	Expected margin of error (percentage) in the state of charge calculation.For example, when Max Error returns 10 percent and Relative State of Charge returns 50 percent, the Relative State of charge is more likely between 50 percent and 60 percent. The gas gauge sets Max Error to 100 percent on a full reset. The gas gauge sets Max Error to 2 percent on completion of a learn cycle, unless the gas gauge limits the learn cycle to the +512/–256-mAh maximum adjustment values. If the learn cycle is limited, the gas gauge sets Max Error to 8 percent unless Max Error was already below 8 percent. In this case Max Error does not change. The gas gauge increments Max Error by 1 percent after four increments of Cycle Count without a learn cycle.
Maximum learn delay from current start time	Maximum length of time between automatic learn cycles. You can delay the start of a learn cycle for a maximum of 168 hours (7 days).
Media error count	A drive property indicating the number of errors that have been detected on the drive media.
Migration	The process of moving virtual drives and hot spare drives from one controller to another by disconnecting the drives from one controller and attaching them to another one. The firmware on the new controller will detect and retain the virtual drive information on the drives.
Mirroring	The process of providing complete data redundancy with two drives by maintaining an exact copy of one drive's data on the second drive. If one drive fails, the contents of the other drive can be used to maintain the integrity of the system and to rebuild the failed drive.
Multipathing	The firmware provides support for detecting and using multiple paths from the RAID controllers to the SAS devices that are in enclosures. Devices connected to enclosures have multiple paths to them. With redundant paths to the same port of a device, if one path fails, another path can be used to communicate between the controller and the device. Using multiple paths with load balancing, instead of a single path, can increase reliability through redundancy.
N	
Name	A virtual drive property indicating the user-assigned name of the virtual drive.
Next learn time	Time at which the next learn cycle starts.
Non-redundant configuration	A RAID 0 virtual drive with data striped across two or more drives but without drive mirroring or parity. This provides for high data throughput but offers no protection in case of a drive failure.
NVRAM	Acronym for nonvolatile random access memory. A storage system that does not lose the data stored on it when power is removed. NVRAM is used to store firmware and configuration data on the RAID controller.
NVRAM present	A controller property indicating whether an NVRAM is present on the controller.
NVRAM size	A controller property indicating the capacity of the controller's NVRAM.
O	
Offline	A drive is offline when it is part of a virtual drive but its data is not accessible to the virtual drive.
P	
Patrol read	A process that checks the drives in a storage configuration for drive errors that could lead to drive failure and lost data. The patrol read operation can find and sometimes fix any potential problem with drives before host access. This enhances overall system performance because error recovery during a normal I/O operation might not be necessary.
Patrol read rate	The user-defined rate at which patrol read operations are run on a computer system.

Predicted battery capacity status (hold 24hr charge)	Indicates whether the battery capacity supports a 24-hour data retention time.
Product info	A drive property indicating the vendor-assigned model number of the drive.
Product name	A controller property indicating the manufacturing name of the controller.

R

RAID	A group of multiple, independent drives that provide high performance by increasing the number of drives used for saving and accessing data.A RAID drive group improves input/output (I/O) performance and data availability. The group of drives appears to the host system as a single storage unit or as multiple virtual drives. Data throughput improves because several drives can be accessed simultaneously. RAID configurations also improve data storage availability and fault tolerance. Redundant RAID levels (RAID levels 1, 5, 6, 10, 50, and 60) provide data protection.
RAID 0	Uses data striping on two or more drives to provide high data throughput, especially for large files in an environment that requires no data redundancy.
RAID 00	Uses data striping on two or more drives in a spanned drive group to provide high data throughput, especially for large files in an environment that requires no data redundancy.
RAID 1	Uses data mirroring on pairs of drives so that data written to one drive is simultaneously written to the other drive. RAID 1 works well for small databases or other small applications that require complete data redundancy.
RAID 1E	Uses two-way mirroring on two or more drives. RAID 1E provides better performance than a traditional RAID 1 array.
RAID 5	Uses data striping and parity data across three or more drives (distributed parity) to provide high data throughput and data redundancy, especially for applications that require random access.
RAID 6	Uses data striping and parity data across three or more drives (distributed parity) to provide high data throughput and data redundancy, especially for applications that require random access. RAID 6 can survive the failure of two drives.
RAID 10	A combination of RAID 0 and RAID 1 that uses data striping across two mirrored drive groups. It provides high data throughput and complete data redundancy.
RAID 50	A combination of RAID 0 and RAID 5 that uses data striping across two drive groups with parity data. It provides high data throughput and complete data redundancy.
RAID 60	A combination of RAID 0 and RAID 6 that uses data striping across two drive groups with parity data. It provides high data throughput and complete data redundancy. RAID 60 can survive the failure of two drives in each RAID set in the spanned drive group.
RAID level	A virtual drive property indicating the RAID level of the virtual drive. ThinkSystem SAS RAID controllers support RAID levels 0, 1, 5, 6, 10, 50, and 60.
RAID Migration	A feature in RAID subsystems that allows changing a RAID level to another level without powering down the system.
Raw capacity	A drive property indicating the actual full capacity of the drive before any coercion mode is applied to reduce the capacity.
Read policy	A controller attribute indicating the current Read Policy mode. Always Read Ahead Permits the controller to read sequentially ahead of the requested data and allows the controller to store the additional data in the cache memory. Here, the controller anticipates that the data is required frequently. Even though Always Read Ahead policy speeds up the reads for sequential data, but little improvement is seen when accessing the random data. No Read Ahead (also known as Normal mode in WebBIOS), the Always Read Ahead capability of the controller is disabled.

Rebuild	The regeneration of all data to a replacement drive in a redundant virtual drive after a drive failure. A drive rebuild normally occurs without interrupting normal operations on the affected virtual drive, though some degradation of performance of the drive subsystem can occur.
Rebuild rate	The percentage of central processing unit (CPU) resources devoted to rebuilding data onto a new drive after a drive in a storage configuration has failed.
Reclaim virtual drive	A method of undoing the configuration of a new virtual drive. If you highlight the virtual drive in the Configuration Wizard and click Reclaim, the individual drives are removed from the virtual drive configuration.
Reconstruction rate	The user-defined rate at which a drive group modification operation is carried out.
Redundancy	A property of a storage configuration that prevents data from being lost when one drive fails in the configuration.
Redundant configuration	A virtual drive that has redundant data on drives in the drive group that can be used to rebuild a failed drive. The redundant data can be parity data striped across multiple drives in a drive group, or it can be a complete mirrored copy of the data stored on a second drive. A redundant configuration protects the data in case a drive fails in the configuration.
Relative state of charge	Predicted remaining battery capacity expressed as a percentage of Full Charge Capacity.
Remaining capacity	Amount of remaining charge capacity of the battery as stated in milliamp hours. This value represents the available capacity or energy in the battery at any given time. The gas gauge adjusts this value for charge, self-discharge, and leakage compensation factors.
Revertible hot spare	When you use the Replace Member procedure, after data is copied from a hot spare to a new drive, the hot spare reverts from a rebuild drive to its original hot spare status.
Revision level	A drive property that indicates the revision level of the drive's firmware.
Run time to empty	Predicted remaining battery life at the present rate of discharge in minutes.

S

SAS	Acronym for Serial-Attached SCSI. SAS is a serial, point-to-point, enterprise-level device interface that leverages the Small Computer System Interface (SCSI) protocol set. The SAS interface provides improved performance, simplified cabling, smaller connectors, lower pin count, and lower power requirements when compared to parallel SCSI.
SATA	Acronym for Serial Advanced Technology Attachment. A physical storage interface standard. SATA is a serial link that provides point-to-point connections between devices. The thinner serial cables allow for better airflow within the system and permit smaller chassis designs.
SCSI device type	A drive property indicating the type of the device, such as drive.
Serial no.	A controller property indicating the manufacturer-assigned serial number.
Stripe size	A virtual drive property indicating the length of the interleaved data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 1 MB of drive space and has 64 KB of data residing on each drive in the stripe. In this case, the stripe size is 1 MB and the strip size is 64 KB. The user can select the stripe size.
Striping	A technique used to write data across all drives in a virtual drive.Each stripe consists of consecutive virtual drive data addresses that are mapped in fixed-size units to each drive in the virtual drive using a sequential pattern. For example, if the virtual drive includes five drives, the stripe writes data to drives one through five without repeating any of the drives. The amount of space consumed by a stripe is the same on each drive. Striping by itself does not provide data redundancy. Striping in combination with parity does provide data redundancy.
Strip size	The portion of a stripe that resides on a single drive in the drive group.

Subvendor ID	A controller property that lists additional vendor ID information about the controller.
T	
Temperature	Temperature of the battery pack, measured in Celsius.
U	
Uncorrectable error count	A controller property that lists the number of uncorrectable errors detected on drives connected to the controller. If the error count reaches a certain level, a drive will be marked as failed.
V	
Vendor ID	A controller property indicating the vendor-assigned ID number of the controller.
Vendor info	A drive property listing the name of the vendor of the drive.
Virtual drive	A storage unit created by a RAID controller from one or more drives. Although a virtual drive can be created from several drives, it is seen by the operating system as a single drive. Depending on the RAID level used, the virtual drive can retain redundant data in case of a drive failure.
Virtual drive state	A virtual drive property indicating the condition of the virtual drive. Examples include Optimal and Degraded.
W	
Write-back	In Write-Back Caching mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a drive write transaction. Data is written to the drive subsystem in accordance with policies set up by the controller.These policies include the amount of dirty/clean cache lines, the number of cache lines available, and elapsed time from the last cache flush.
Write policy	See *Default Write Policy*.
Write-through	In Write-Through Caching mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data and has completed the write transaction to the drive.