



12Gb/s MegaRAID[®] Tri-Mode Software

**User Guide
Version 1.8**

Broadcom, the pulse logo, Connecting everything, Avago Technologies, Avago, the A logo, MegaRAID, LSI, CacheVault, SSD Guard, SafeStore, and CacheCade are among the trademarks of Broadcom and/or its affiliates in the United States, certain other countries and/or the EU.

Data subject to change. Copyright ©2016–2020 Broadcom. All Rights Reserved.

The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries. For more information, please visit www.broadcom.com

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Table of Contents

Overview.....	10
Tri-Mode Technology.....	10
Serial-Attached SCSI Device Interface.....	10
Serial ATA III Features.....	11
Nonvolatile Memory Express Technology.....	11
Solid State Drive Features.....	11
SSD Guard.....	12
Dimmer Switch Features.....	12
UEFI 2.0 Support.....	12
Configuration Scenarios.....	12
Technical Support.....	14
SnapDump Feature.....	14
Introduction to RAID.....	15
Components and Features.....	15
Drive Group.....	15
Virtual Drive.....	15
Fault Tolerance.....	16
Multipathing.....	16
Consistency Check.....	17
Replace.....	17
Background Initialization.....	18
Profile Management.....	18
Patrol Read.....	18
Disk Striping.....	19
Disk Mirroring.....	19
Parity.....	20
Disk Spanning.....	21
Hot Spares.....	22
Disk Rebuilds.....	23
Rebuild Rate.....	23
Hot Swap.....	23
Drive States.....	23
Virtual Drive States.....	24
Beep Codes.....	24
Enclosure Management.....	25
RAID Levels.....	25

Summary of RAID Levels.....	25
Selecting a RAID Level.....	26
RAID 0 Drive Groups.....	26
RAID 1 Drive Groups.....	27
RAID 5 Drive Groups.....	28
RAID 6 Drive Groups.....	28
RAID 00 Drive Groups.....	29
RAID 10 Drive Groups.....	30
RAID 50 Drive Groups.....	31
RAID 60 Drive Groups.....	32
RAID Configuration Strategies.....	33
Maximizing Fault Tolerance.....	34
Maximizing Performance.....	35
Maximizing Storage Capacity.....	36
Configuration Planning.....	37
Number of Drives.....	37
RAID Availability.....	38
RAID Availability Concept.....	38
SafeStore Disk Encryption.....	39
Workflow.....	40
Enable Security.....	40
Change Security.....	41
Create Secure Virtual Drives.....	41
Import a Foreign Configuration.....	42
Instant Secure Erase.....	42
HII Configuration Utility.....	44
Behavior of HII.....	44
Starting the HII Configuration Utility.....	45
HII Dashboard View.....	46
Main Menu.....	46
HELP.....	47
PROPERTIES.....	48
ACTIONS.....	49
BACKGROUND OPERATIONS.....	50
MegaRAID ADVANCED SOFTWARE OPTIONS.....	50
Critical Boot Messages.....	51
Managing Configurations.....	51
Creating a Virtual Drive from a Profile.....	52
Creating a RAID 10 Volume from the Database.....	56

Manually Creating a Virtual Drive.....	58
Viewing Drive Group Properties.....	63
Viewing Global Hot Spare Drives.....	64
Clearing a Configuration.....	64
Make Unconfigured Good, Make JBOD, and Enable Security on JBOD.....	64
Make Unconfigured Good.....	65
Make JBOD.....	66
Enabling Security on JBOD.....	66
Managing Foreign Configurations.....	67
Previewing and Importing a Foreign Configuration.....	68
Clearing a Foreign Configuration.....	70
Managing Controllers.....	71
Viewing Advanced Controller Management Options.....	72
Viewing Advanced Controller Properties.....	74
Managing MegaRAID Advanced Software Options.....	77
Scheduling a Consistency Check.....	77
Managing Personality Mode.....	78
Saving or Clearing Controller Events.....	79
Enabling or Disabling Drive Security.....	80
Changing Security Settings.....	83
Saving the TTY Log.....	84
Managing SAS Storage Link Speeds.....	85
Managing PCIe Storage Interface.....	86
Managing Profiles.....	87
Downgrading the Firmware When Profiles Are Selected.....	89
Setting Cache and Memory Properties.....	90
Running a Patrol Read.....	91
Changing Power Save Settings.....	93
Setting Emergency Spare Properties.....	94
Changing Task Rates.....	95
Upgrading the Firmware.....	96
Managing Virtual Drives.....	98
Selecting Virtual Drive Operations.....	99
Locating Physical Drives in a Virtual Drive.....	100
Deleting a Virtual Drive.....	100
Hiding a Virtual Drive.....	101
Unhiding a Virtual Drive.....	101
Hiding a Drive Group.....	101
Unhiding a Drive Group.....	102
Reconfiguring a Virtual Drive.....	102

Initializing a Virtual Drive.....	104
Erasing a Virtual Drive.....	104
Securing a Virtual Drive.....	105
Running a Consistency Check.....	105
Expanding a Virtual Drive.....	106
Viewing Associated Drives.....	106
Viewing and Managing Virtual Drive Properties and Options.....	106
Managing Physical Drives.....	108
Performing Drive Operations.....	110
Locating a Drive.....	110
Making a Drive Unconfigured Bad, Unconfigured Good, or JBOD.....	111
Enabling Security on JBOD.....	111
Replacing a Drive.....	112
Placing a Drive Offline.....	113
Placing a Drive Online.....	113
Marking a Drive Missing.....	113
Replacing a Missing Drive.....	114
Assigning a Global Hot Spare Drive.....	114
Assigning a Dedicated Hot Spare Drive.....	114
Unassigning a Hot Spare Drive.....	115
Initializing or Erasing a Drive.....	115
Rebuilding a Drive.....	116
Securely Erasing a Drive.....	116
Removing a Physical Drive.....	117
Making a JBOD.....	117
Erasing a JBOD.....	117
Viewing Advanced Drive Properties.....	118
Managing Hardware Components.....	119
Managing Batteries.....	119
Setting Automatic Learn Cycle Properties.....	122
Managing Enclosures.....	122
StorCLI.....	124
Overview.....	124
Controllers Supported by the StorCLI Tool.....	124
Supported Operating Systems.....	124
Installing StorCLI on MegaRAID Controllers.....	126
Installing the StorCLI Tool on Microsoft Windows Operating Systems.....	126
Installing the StorCLI Tool on the UEFI Environment.....	126
Installing the StorCLI Tool on Linux Operating Systems.....	126
Installing the StorCLI Tool on VMware Operating Systems.....	127

StorCLI Tool Command Syntax.....	127
StorCLI Default Logging.....	129
StorCLI Controller Commands.....	130
System Commands.....	130
System Show Commands.....	130
Controller Commands.....	130
Show and Set Controller Properties Commands.....	131
Controller Show Commands.....	140
Controller Debug Commands.....	141
Controller Background Task Operation Commands.....	142
Profile Management.....	145
Premium Feature Key Commands.....	146
Controller Security Commands.....	147
Flashing Controller Firmware Command while the Firmware Is Operational.....	148
Controller Cache Command.....	149
Controller Configuration Commands.....	149
SnapDump Commands.....	150
SPDM Commands.....	152
Drive Commands.....	153
Drive Show Commands.....	153
Missing Drives Commands.....	154
Set Drive State Commands.....	155
Drive Initialization Commands.....	156
NVMe Drive Commands.....	157
Configuration Scenarios.....	158
Drive Firmware Download Commands.....	159
Drive Firmware Update through Parallel HDD Microcode.....	160
Locate Drives Commands.....	161
Prepare to Remove Drives Commands.....	161
Drive Security Command.....	162
Drive Secure EraseCryptographic Erase Commands.....	162
Drive Sanitize Command.....	163
Rebuild Drives Commands.....	164
Drive Copyback Commands.....	165
Hot Spare Drive Commands.....	166
Drive Performance Monitoring Commands.....	167
Virtual Drive Commands.....	169
Add Virtual Drives Commands.....	170
Delete Virtual Drives Commands.....	173
Virtual Drive Show Commands.....	174

Preserved Cache Commands.....	174
Change Virtual Properties Commands.....	175
Virtual Drive Initialization Commands.....	177
Virtual Drive Erase Commands.....	178
Virtual Drive Migration Commands.....	178
Virtual Drive Consistency Check Commands.....	180
Background Initialization Commands.....	181
Virtual Drive Expansion Commands.....	182
Display the Bad Block Table.....	182
Clear the LDBBM Table Entries.....	182
JBOD Commands.....	183
Create JBOD Manually.....	183
JBOD Properties.....	183
JBOD Operations.....	184
Delete JBODs or Volumes.....	184
Foreign Configuration Commands.....	185
BIOS-Related Commands.....	186
OPROM BIOS Commands.....	186
Drive Group Commands.....	187
Drive Group Show Commands.....	187
Virtual Drive Power Settings Commands.....	189
Change Virtual Drive Power Settings Commands.....	189
CacheVault Commands.....	189
Enclosure Commands.....	190
PHY Commands.....	191
PCIe Storage Interface Commands.....	192
Lane Speed Commands.....	192
Link Configuration Commands.....	193
Logging Commands.....	195
Automated Physical Drive Configurations.....	197
Frequently Used Tasks.....	198
Displaying the Version of the StorCLI Utility.....	198
Displaying the StorCLI Tool Help.....	199
Displaying System Summary Information.....	199
Displaying Free Space in a Controller.....	199
Adding Virtual Drives.....	199
Setting the Cache Policy in a Virtual Drive.....	200
Displaying Virtual Drive Information.....	200
Deleting Virtual Drives.....	200
Flashing Controller Firmware.....	200

Appendix A: Events, Messages, and Behaviors.....	202
Error Levels.....	202
Event Messages.....	202
Appendix B: Unsupported Commands in Embedded MegaRAID.....	229
Appendix C: CLI Error Messages.....	231
Error Messages and Descriptions.....	231
Appendix D: 240 Virtual Drive Feature Limitations.....	235
Host Software Utility.....	235
BIOS Known Limitations.....	235
Appendix E: Online Firmware Upgrade and Downgrade.....	236
Online Firmware Upgrade and Downgrade.....	236
Appendix F: Boot Messages and BIOS Error Messages.....	238
Displaying Boot Messages.....	238
Differences in the System Boot Mode.....	239
Appendix G: Glossary.....	265
Appendix H: Revision History.....	273

Overview

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

This guide describes how to use the Storage Command Line Interface (StorCLI) tool software and the MegaRAID Human Interface Infrastructure (HII) configuration utility.

This chapter documents the SAS technology, Serial ATA (SATA) technology, SSD Guard™, Dimmer Switch, UEFI 2.0, configuration scenarios, and drive types. Other features such as Fast Path and SafeStore™ are described in other chapters of this guide.

Tri-Mode Technology

The MegaRAID 12Gb/s Tri-Mode RAID controllers are high-performance intelligent SAS/SATA/PCIe (NVMe) devices with RAID control capabilities. The MegaRAID 12Gb/s Tri-Mode RAID controllers provide reliability, high performance, and fault-tolerant disk subsystem management. They are an ideal RAID solution for the internal storage of workgroup, departmental, and enterprise systems. The MegaRAID 12Gb/s Tri-Mode RAID controllers offer a cost-effective way to implement RAID in a server.

Tri-Mode technology brings a wealth of options and flexibility with the use of SAS devices, Serial ATA (SATA) III, SATA III devices, and PCIe (NVMe) within the same storage infrastructure. These devices bring individual characteristics that make each of these more suitable choices depending on your storage needs. MegaRAID gives you the flexibility to combine these three 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, SATA drives, and PCIe (NVMe) within the same virtual drives. Avoid mixing HDD drive types.

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

The Tri-Mode 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* and *PCIe Gen 4.0 specification*. Supporting the SAS/SATA/PCIe (NVMe), the tri-mode controller is a versatile controller that provides the backbone of both server environments and high-end workstation environments.

Each port on the Tri-Mode RAID controller supports SAS/SATA/PCIe (NVMe) 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
- NVMe, which accesses storage media attached by means of a PCIe bus

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.

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 failover configuration) through an expander

Nonvolatile Memory Express Technology

Nonvolatile memory express (NVMe) is a logical device interface specification for accessing NVMe storage media attached by means of a PCI Express (PCIe) bus, which removes SCSI from the I/O stack. By its design, NVMe allows host hardware and software to utilize the parallelism found in SSDs. As a result, NVMe reduces I/O overhead and brings performance improvements to the logical device interfaces. These improvements include multiple command queues and reduced latency.

The NVMe interface is designed with following key attributes:

- Support for up to 64K I/O queues with minimal command overhead
- Each I/O queue supports 64K I/O operations
- Each I/O queue is designed for simultaneous multi-threaded processing
- NVMe protocol enables hardware automated queues
- NVMe commands and structures are transferred end-to-end
- The NVMe protocol can be transported across multiple network fabric types

Solid State Drive Features

The MegaRAID 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.

SSD Guard

SSD Guard, a feature that is unique to MegaRAID, 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), nonredundant 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 an SSD failure is imminent, the MegaRAID software starts a rebuild to preserve the data on the SSD and sends appropriate warning event notifications.

Dimmer Switch Features

Powering drives and cooling drives represent a major cost for data centers. The MegaRAID Dimmer Switch feature set reduces the power consumption of the devices connected to a MegaRAID controller, which helps to share resources more efficiently and lowers the cost.

- Dimmer Switch 1 – Spin down unconfigured disks. This feature is configurable and can be disabled.
- Dimmer Switch 2 – Spin down Hot Spares. This feature is configurable and can be disabled.

UEFI 2.0 Support

UEFI 2.0 provides MegaRAID customers with expanded platform support. The MegaRAID 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 MegaRAID management applications through a driver configuration protocol (DCP). The UEFI driver also supports driver diagnostic protocol, which allows administrators to access pre-boot diagnostics.

Configuration Scenarios

You can use the SAS RAID controllers in three scenarios:

- **Low-end, Internal SATA Configurations**

In these configurations, use the RAID controller as a high-end SATA II-compatible controller that connects up to eight disks. These configurations are mostly for low-end or entry servers. Enclosure management is provided through out-of-band Inter-IC (I²C) bus. Side bands of both types of internal SAS connectors support the SFF-8485 (SGPIO) interface.

- **Midrange Internal SAS Configurations**

These configurations are like the internal SATA configurations but with high-end disks. These configurations are more suitable for low-range to midrange servers.

- **High-end External SAS/SATA Configurations**

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

- **NVMe Configurations**

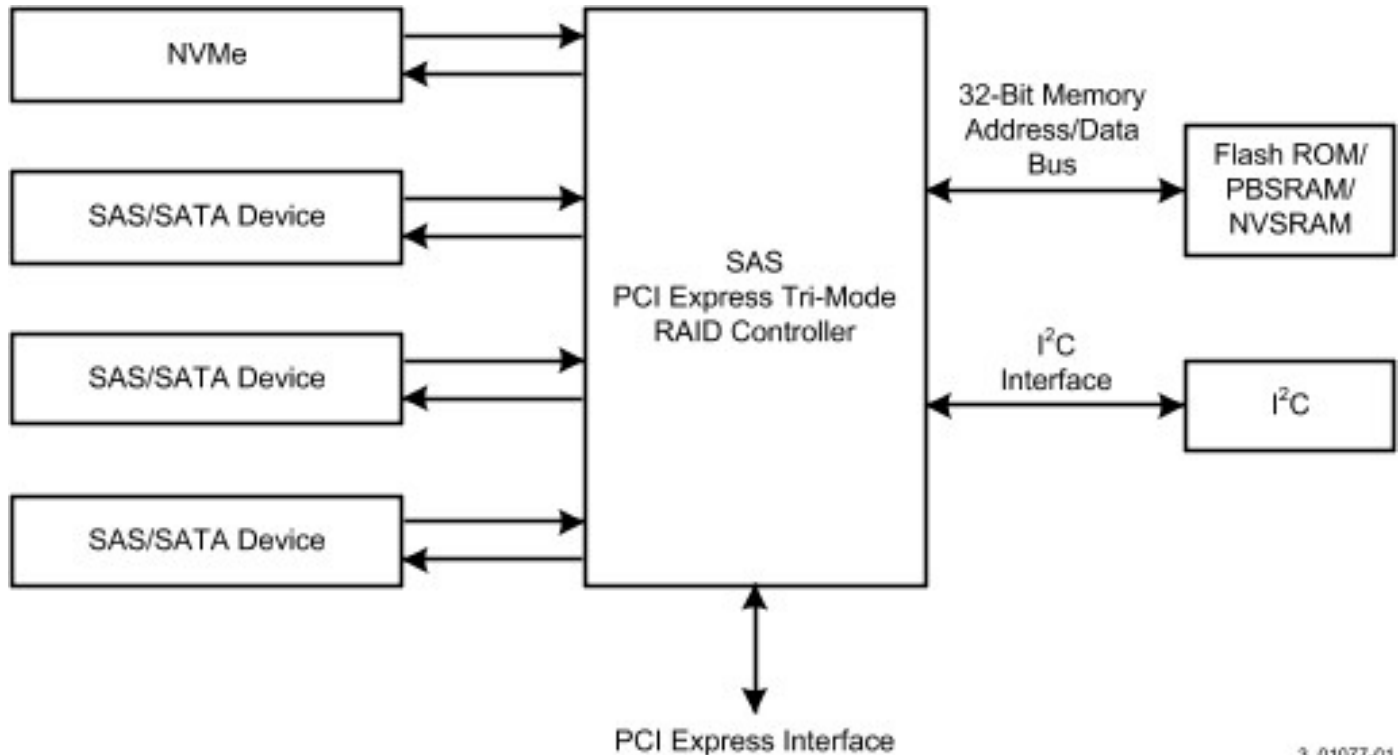
These configurations are for internal connectivity only, using NVMe, either direct connect or switch attached. NVMe configurations are suitable for low latency and high performance environments.

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

NOTE

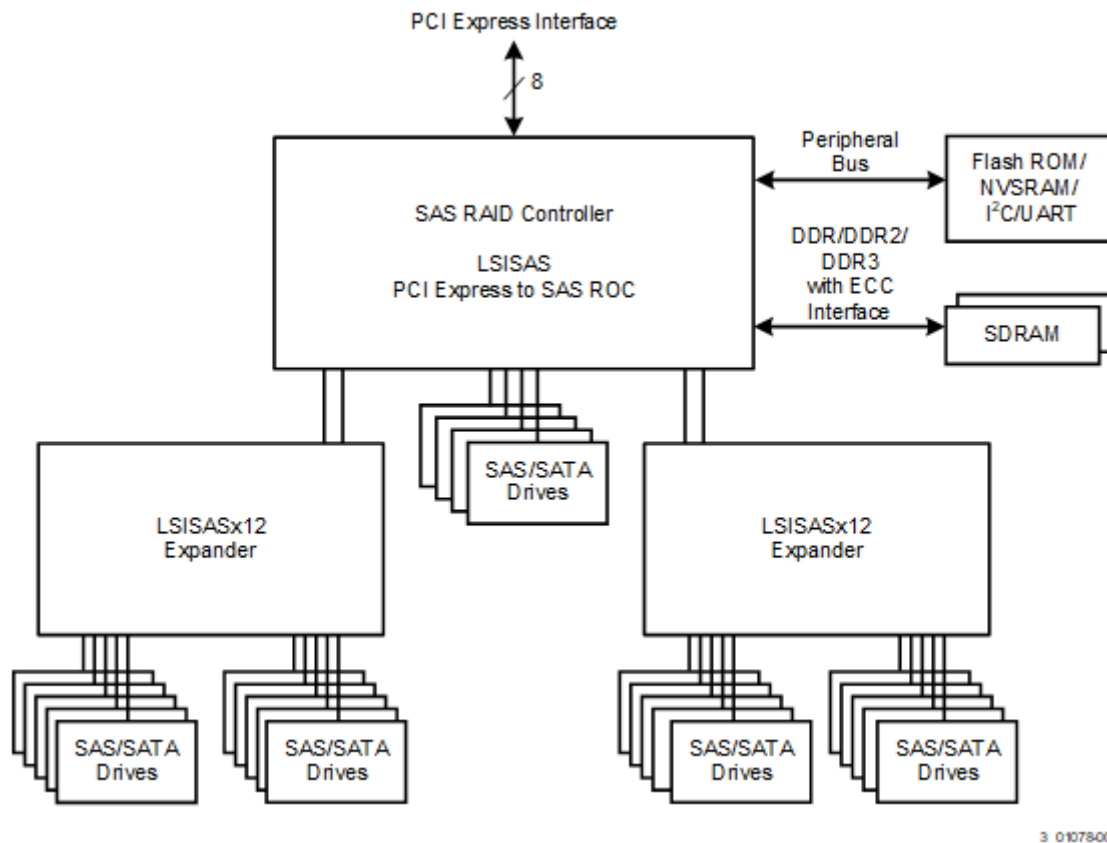
The external memory bus is 32-bit for the SAS 8704ELP and the SAS 8708ELP, and 64-bit for the SAS 8708EM2, the SAS 8880EM2, and the SAS 8888ELP.

Figure 1: Example of SAS Direct-Connect Application



3_01077-01

The following figure shows an example of a SAS RAID controller configured with an LSISASx12 expander that is connected to SAS disks, SATA II disks, or both.

Figure 2: Example of SAS RAID Controller Configured with an LSI SASx12 Expander

Technical Support

For assistance with installing, configuring, or running your Tri-Mode MegaRAID SAS RAID controllers, contact a Broadcom Technical Support representative. Click the following link to access the Broadcom Technical Support page for storage and board support:

[REQUEST TECHNICAL SUPPORT](#)

From this page, you can call a Technical Support representative, or submit a new service request and view its status.

Phone Support:

[Call Us For Storage Support](#)

1-800-633-4545 (North America)

00-800-5745-6442 (International)

+ 49 (0) 8941 352 0123 (Germany)

SnapDump Feature

SnapDump collects critical debug data such as firmware logs, events, and hardware register dumps during an initial unexpected failure. SnapDump data can be saved on the host using the Broadcom APIs, avoiding the need for an external USB-UART Dongle at customer environments.

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

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.

Drive Group

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

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

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.

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 (LSA) 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.

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.

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.

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.

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

- RAID1 virtual drives must have at least two drives to start initialization.
- RAID 5 virtual drives must have at least five drives for background initialization to start.
- RAID 6 virtual drives must have at least seven drives for background initialization to start.

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

Profile Management

Profile Management allows you to have multiple configurations supported under each personality mode. Profiles are used to customize the controller to deliver the best performance for that configuration. For example, a profile with no PCI device support can support a higher Queue Depth than a profile that supports 32 PCI devices.

When you choose profile management either through HII, StorCLI, or LSA, the firmware provides a list of profiles that you can select for the current personality.

Compatibility Check

Applications may sometime fail the profile change request for the following reasons:

- When the devices in the storage environment exceeds the maximum profile count.
- When a particular profile does not support the type of device that is discovered. For example, the storage topology has SAS/SATA devices, but the user requests an NVMe-only profile.
- When pinned cache is present.
- When critical background operations such as Rebuild, Copyback, Reconstruction, Initialization of Logical Drives, Consistency Check, Patrol Read are running.

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.

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 for MegaRAID controllers and 64 KB for Integrated MegaRAID controllers. 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 3: 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.

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 4: Example of Disk Mirroring (RAID 1)



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 5: Example of Distributed Parity (RAID 5 Drive Group)



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


3_01081-00

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 6: Example of Disk Spanning



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

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

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.

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 MegaRAID 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.

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.)

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 30 percent.

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.

Drive States

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

Table 3: Drive States

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

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.
Offline	The virtual drive is not available to the RAID controller.

Beep Codes

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

Table 5: Beep Codes, Events, and Virtual Drive States

Event	Virtual Drive State	Beep Code
RAID 0 virtual drive loses a virtual drive	Offline	3 seconds on and 1 second off
RAID 1 virtual drive loses a mirror drive	Degraded	1 second on and 1 second off
RAID 1 virtual drive loses both drives	Offline	3 seconds on and 1 second off
RAID 5 virtual drive loses one drive	Degraded	1 second on and 1 second off
RAID 5 virtual drive loses two or more drives	Offline	3 seconds on and 1 second off
RAID 6 virtual drive loses one drive	Partially Degraded	1 second on and 1 second off
RAID 6 virtual drive loses two drives	Degraded	1 second on and 1 second off

Event	Virtual Drive State	Beep Code
RAID 6 virtual drive loses more than two drives	Offline	3 seconds on and 1 second off
A hot spare completes the Rebuild process and is brought into a drive group	N/A	1 second on and 3 seconds off
A copy back occurs after a Rebuild operation completes	Optimal	1 second on and 3 seconds off

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.

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.

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.

NOTE

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

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

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 6: 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 7: RAID 0 Drive Group Example with Two Drives




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 7: 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 8: RAID 1 Drive Group



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 8: 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 9: RAID 5 Drive Group with Six Drives



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

3_01085-00

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 9: 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	3 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 10: Example of Distributed Parity across Two Blocks in a Stripe (RAID 6 Drive Group)



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

3_01086-00

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 10: 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 11: RAID 00 Drive Group Example with Two Drives



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 eight spans, with a maximum of 32 drives per span. You must use an even number of drives in each RAID 10 virtual drive in the span.

NOTE

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


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

Table 11: 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 12: RAID 10 Level Virtual Drive



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 12: 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 13: RAID 50 Level Virtual Drive

3_01089-00

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 13: RAID 60 Drive Group Overview

Uses	<p>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.</p> <p>In the case of a failure of one drive or two drives in a RAID set in a virtual drive, the RAID controller uses the parity blocks to 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.</p> <p>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.</p>
Strong points	<p>Provides data redundancy, high read rates, and good performance in most environments. Each RAID 6 set can survive the loss of two drives or the loss of a drive while another drive is being rebuilt.</p> <p>Provides the highest level of protection against drive failures of all of the RAID levels.</p>
Weak points	<p>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.</p> <p>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.</p> <p>A RAID 6 drive group costs more because of the extra capacity required by using two parity blocks per stripe.</p>
Drives	A minimum of 8.

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

Figure 14: RAID 60 Level Virtual Drive

3_01090-00

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.

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. 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 14: RAID Levels and Fault Tolerance

RAID Level	Fault Tolerance
0	Does not provide fault tolerance. All data is lost if any drive fails. Disk striping writes data across multiple drives instead of just one drive. It involves partitioning each drive storage space into stripes that can vary in size. 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.

RAID Level	Fault Tolerance
60	<p>Provides data redundancy using distributed parity across spanned RAID 6 drive groups.</p> <p>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.</p> <p>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.</p>

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 15: RAID Levels and Performance

RAID Level	Performance
0	<p>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 MegaRAID controllers and only 64 KB for Integrated MegaRAID controllers. The LSISAS210 eight controller allows strip size from 8 KB to 1 MB.</p> <p>These stripes are interleaved in a repeated sequential manner. Disk striping enhances performance because multiple drives are accessed simultaneously.</p>
1	<p>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.</p>
5	<p>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.</p> <p>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.</p>
6	<p>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.</p> <p>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.</p>
00	<p>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.</p> <p>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 MegaRAID controllers and only 64 KB for Integrated MegaRAID 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.</p>

RAID Level	Performance
10	<p>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.</p> <p>The system performance improves as the number of spans increases. (The maximum number of spans is eight.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans, and RAID performance degrades to that of a RAID 1 or RAID 5 drive group.</p>
50	<p>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.</p> <p>The system performance improves as the number of spans increases. (The maximum number of spans is eight.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans and RAID drive group performance degrades to that of a RAID 1 or RAID 5 drive group.</p>
60	<p>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 eight.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans, and RAID performance degrades to that of a RAID 1 or RAID 6 drive group.</p> <p>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.</p>

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 16: RAID Levels and Capacity

RAID Level	Capacity
0	<p>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.</p> <p>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.</p>
1	<p>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.</p>
5	<p>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.</p> <p>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.</p>

RAID Level	Capacity
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.

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.

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 17: Factors to Consider for Drive Group Configuration

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

RAID Availability

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.

SafeStore Disk Encryption

This chapter describes the SafeStore Disk Encryption service. The SafeStore Disk Encryption service is a collection of features within the Broadcom 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. The LSA software offers a procedure that you can use to manage the security settings for the drives.

Purpose and Benefits

Security is a growing market concern and requirement. MegaRAID 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.

Terminologies

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

Table 18: Terminologies 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.

Option	Description
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.

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. Configure the system with the desired password.

After the system is booted, you need not enter the password again to access the virtual drives.

If the virtual drive is moved to a different controller, the controller to which the virtual drive is moved, so that access to the data must have the following features:

- SafeStore enabled.
- Encryption enabled.
- The security key must be entered.

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 must 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.

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.

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.

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 LSA, StorCLI, or HII 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.

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 MegaRAID 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.

HII Configuration Utility

The MegaRAID Human Interface Infrastructure (HII) configuration utility configures controllers, physical disks, virtual disks, and performs other configuration tasks in a pre-boot, Unified Extensible Firmware Interface (UEFI) environment.

System BIOS should support Broadcom's private interface and add host memory address into the DMAR (DMA remapping)/RMRR (Reserved memory Region Reporting Structure) table for iMegaRAID to work seamlessly in VT-d/IOMMU (Intel Virtualization Technology for Directed I/O/input-output memory management unit) enabled system/operating system.

If you are using the iMegaRAID controller in an environment that has Windows 10 or Windows Server 2016, and if Hypervisor Code Integrity (HVCI) feature is enabled, you will see some boot issues or issues using the iMR controllers. To avoid this, the system BIOS should categorize the host memory that the UEFI driver has allocated as Read/Write.

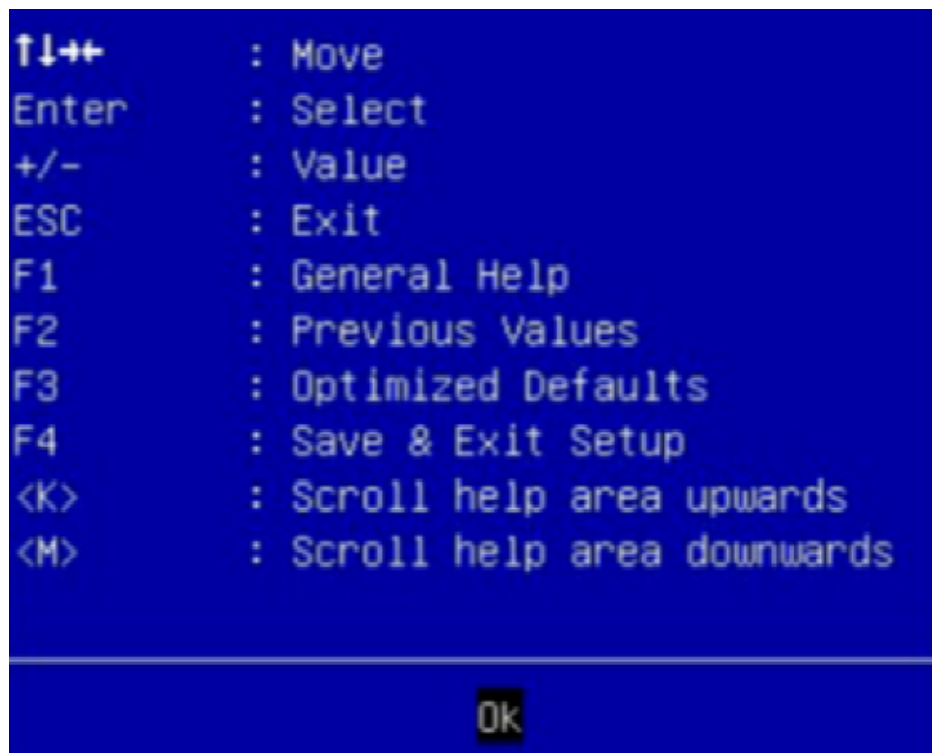
Behavior of HII

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

Some of the HII graphical user interface (GUI) 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 15: 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 the **Set Factory Default** option present on the **Dashboard View** menu or by clicking the **Controller Management > Advanced Controller Management > Set Factory Defaults** menu. Pressing **F3** (Optimized Defaults) will not restore the controller defaults.

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 bootstrap.

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. Also, the following workflow may not be the same for all OEM systems.

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 MegaRAID 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 16: Dashboard View



HII Dashboard View

While you are in the **Dashboard View**, and if HII detects any new events, HII issues various DCMDs to update the data for multiple fields present in the Dashboard; it checks and updates the controller status, updates the backplane information, updates expander/enclosure counts, updates drive group counts, updates virtual drive counts, and so on.

While you are in the **Dashboard View**, you can hot plug or unplug enclosures, and monitor those counts; you can hot plug or unplug physical drives and monitor those counts as well. You can view and preview a foreign configuration, and import and clear a foreign configuration; HII Dashboard also indicates the number of virtual drives and physical drives that are in progress.

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

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

Figure 17: Main Menu

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 [Managing Physical Drives](#).
- Select **Hardware Components** to view battery properties, manage batteries, and manage enclosures. For more information, see [Managing Hardware Components](#).

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 (if supported)

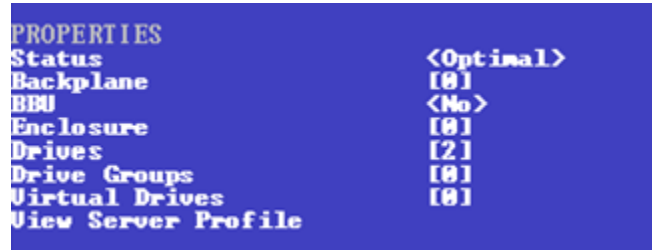
NOTE

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. Help strings for Silence Alarm is displayed only if your controller supports alarm.

PROPERTIES

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

Figure 18: Dashboard View – PROPERTIES



PROPERTIES	
Status	<Optimal>
Backplane	[0]
BBU	<No>
Enclosure	[0]
Drives	[2]
Drive Groups	[0]
Virtual Drives	[0]
View Server Profile	

- **Status**
Displays the overall status of the controller.
- **Backplane**
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**
Displays the UEFI specification version that the system supports and the following menu options, as shown in the following figure.

Figure 19: Dashboard View – 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 super capacitor properties, manage super capacitors, 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](#).

ACTIONS

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

Figure 20: 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.

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 21: Dashboard View – BACKGROUND OPERATIONS



MegaRAID ADVANCED SOFTWARE OPTIONS

This section displays the enabled advanced software options, such as the RAID levels, MegaRAID SafeStore, and MegaRAID FastPath. This section also allows you to configure and use the advanced features. See [Managing MegaRAID Advanced Software Options](#).

Figure 22: Dashboard View – MegaRAID ADVANCED SOFTWARE OPTIONS

MegaRAID ADVANCED SOFTWARE OPTIONS		
MegaRAID	RAID6	[Enabled]
MegaRAID	RAID5	[Enabled]
MegaRAID	SafeStore	[Enabled]
MegaRAID	FastPath	[Enabled]

Critical Boot Messages

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 appear 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 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 team for further assistance.

Managing Configurations

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

Figure 23: Configuration Management Screen

The **Make JBOD**, **Enable Security on JBOD**, and **Make Unconfigured Good** options are included for some controllers. (See [Make Unconfigured Good](#), [Make JBOD](#), and [Enable Security on JBOD](#)). You can enable security on the JBOD drives either from the **Configuration Management** screen or the **Drive Management** screen. The following are the prerequisites for enabling security on JBOD drives:

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

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

The HII Configuration Utility supports 240 VD creation. For more information, see [240 Virtual Drive Feature Limitations](#).

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** screen appears if you select the **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 19: Virtual Drive Creation Profile Options

Option	Description
Drive Selection Criteria	You must 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. 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 striping involves partitioning each physical drive storage space in strips of the following sizes: 64 KB , 128 KB , 256 KB , 512 KB , 1 MB .
Read Policy	<p>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 <i>default</i> read policy. The possible options follow:</p> <ul style="list-style-type: none"> • 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. Although the 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.

Option	Description
Write Policy	<p>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:</p> <ul style="list-style-type: none"> • 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	<p>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:</p> <ul style="list-style-type: none"> • 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 .
Disk Cache Policy	Displays the virtual drive cache setting. The possible options are Unchanged , Enable , and Disable .
Default Initialization	<p>Displays the virtual drive initialization setting. Default Initialization displays the following options:</p> <ul style="list-style-type: none"> • 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 20: 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

Properties	Generic RAID0	Generic RAID1	Generic RAID5	Generic RAID6
SATA	Supported	Supported	Supported	Supported
PCIe	Supported	Supported	Supported	Not supported
SED	Supported	Supported	Supported	Supported
NonSED	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	255	2	255	255
Power-save mode	Controller-defined	Controller-defined	Controller-defined	Controller-defined
Strip Size	256 KB	256 KB	256 KB	256 KB
Read Policy	If the drive is an SSD drive, the No Read Ahead option appears. Else, the Default option appears.	If the drive is an SSD drive, the No Read Ahead option appears. Else, the Default option appears.	If the drive is an SSD drive, the No Read Ahead option appears. Else, the Default option appears.	If the drive is an SSD drive, the No Read Ahead option appears. Else, the Default option appears.
Write Policy	If the drive is an SSD drive, the Write Through option appears. Else, the Write Back option appears.	If the drive is an SSD drive, the Write Through option appears. Else, the Write Back option appears.	If the drive is an SSD drive, the Write Through option appears. Else, the Write Back option appears.	If the drive is an SSD drive, the Write Through option appears. Else, the Write Back option appears.
IO Policy	If the drive is an SSD drive, the Direct IO option appears. Else, the Default option appears.	If the drive is an SSD drive, the Direct IO option appears. Else, the Default option appears.	If the drive is an SSD drive, the Direct IO option appears. Else, the Default option appears.	If the drive is an SSD drive, the Direct IO option appears. Else, the Default option 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

Properties	Generic RAID0	Generic RAID1	Generic RAID5	Generic RAID6
Mixing of Interface Type SAS, SATA, and NVMe 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

Creating a RAID 10 Volume from the Database

You can create RAID 10 volume from the Database feature. Creating RAID 10 from the Database uses drive mirroring so that data written to one drive is simultaneously written to another drive. Creating a RAID 10 volume from the Database provides you with fault tolerance and low latency for the use of the database.

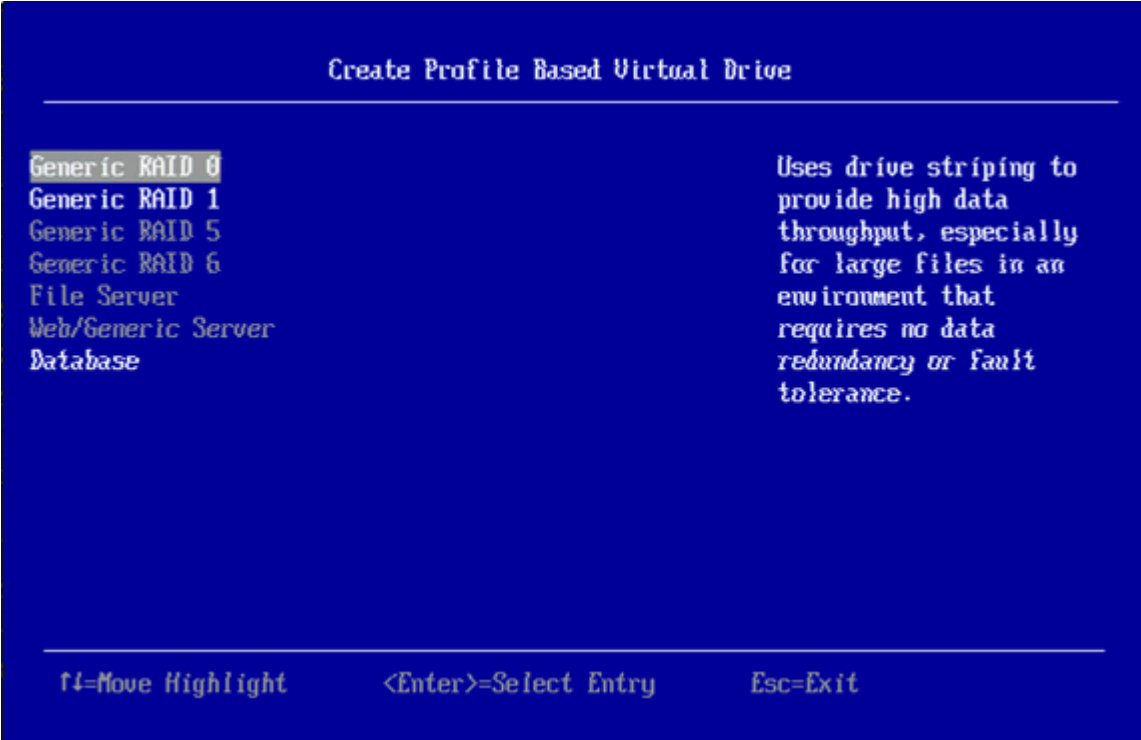
You need a minimum of four drives to create a RAID 10 volume. The profile-based virtual drive creation option allows you to create a RAID 10 volume. If you use this option, you do not choose any drives; the system automatically chooses the drives and creates a RAID 10 volume.

To create a RAID 10 volume using the profile-based virtual drive creation option, perform the following steps:

1. Select **Configuration Management** from the **Main Menu**
2. Select **Create Profile Based Virtual Drive** from the **Configuration Management** menu.

The following dialog appears.

Figure 24: Profile Based Virtual Drive Dialog



3. Highlight the **Database** option and press **Enter**.
The **Database** dialog appears.

Figure 25: Database Dialog



If a small red arrow at the bottom of the window appears, it indicates that you can scroll down to view more information. This red arrow appears when there is too much information to display in one screen. The amount of information that can be displayed in one screen depends on the capabilities of the HII browser.

4. Highlight **Save Configuration** and press **Enter**.

A message appears confirming that the configuration is being created.

5. Highlight **Confirm** and press the spacebar, then highlight **Yes** and press **Enter**.

A success message appears.

6. Highlight **OK** and press **Enter**.

The HII utility creates a RAID 10 volume and returns you to the **Configuration Management** menu.

Manually Creating a Virtual Drive

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

Figure 26: Create Virtual Drive Dialog

Create Virtual Drive		
▶ Save Configuration		
Select RAID Level	<RAID0>	Selects the desired RAID level. The RAID levels that can be configured are 0, 1, 5, 6 (if supported), 00, 10, 50, and 60 (if supported).
Protect Virtual Drive	[]	
Select Drives From	<Unconfigured Capacity>	
▶ Select Drives		
CONFIGURE VIRTUAL DRIVE PARAMETERS:		
Virtual Drive Name	-	
Virtual Drive Size	-	
Virtual Drive Size Unit	<GB>	
Strip Size	<64 KB>	
Read Policy	<No Read Ahead>	
Write Policy	<Write Through>	
I/O Policy	<Direct>	
Access Policy	<Read/Write>	
Drive Cache	<Unchanged>	
Disable Background	<No>	
Initialization		
Default Initialization	<No>	
Emulation Type	<Default>	
▶ Save Configuration		More (D/d)
↑↓=Move Highlight <Enter>=Select Entry Esc=Exit		

If a small red arrow at the bottom of the window appears, it indicates that you can scroll down to view more information. This red arrow appears when there is too much information to display in one screen. The amount of information that can be displayed in one screen depends on the capabilities of the HII browser.

NOTE

If your system detects any JBODs, the **Make Unconfigured Good** dialog appears before the **Create Configuration** window. The **Make Unconfigured Good** dialog lets you convert the JBOD drives to Unconfigured Good. See [Make Unconfigured Good](#).

NOTE

If you create a virtual drive, for example RAID 1, with different drive sizes, such as 1 TB and 2 TB, and after you have created the VD and you want to replace a small drive with a larger drive (replace 1-TB drive with a 2-TB drive), you cannot create another RAID 1 using the additional 1 TB.

NOTE

HII does not apply the mixing rule across the span when you create spanned RAID levels.

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**.

NOTE

Mixing RAID levels (R10, R50, R80) across a system is not supported.

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 of the RAID levels. See [Table 23](#) 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.

- a) If the security key is enabled, check the **Secure Virtual Drive** box to secure the new virtual drive.

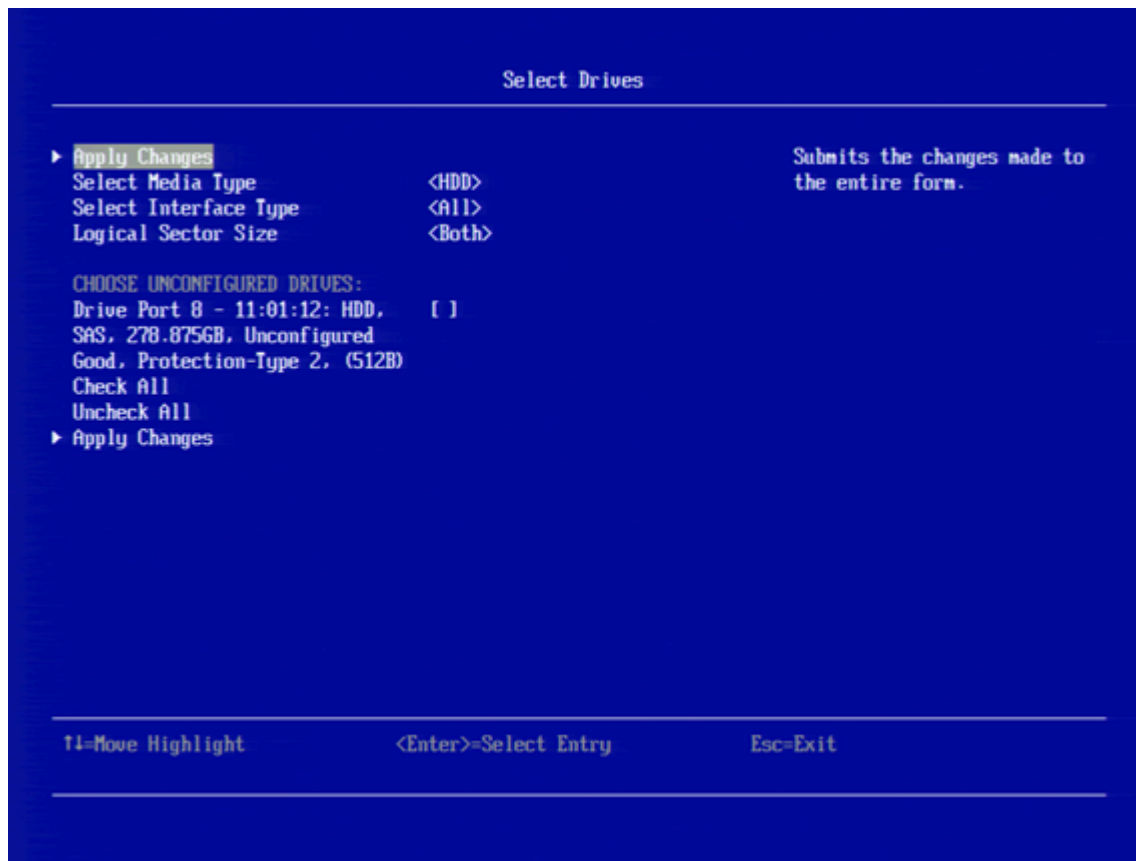
This field is not available unless the security feature is already enabled.

4. Highlight the **Select Drives From** field, press **Enter**, and select either **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 using unconfigured drives.

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

The **Select Drives** dialog appears.

Figure 27: Select Drives Dialog

6. From the **Select Drives** dialog, you can select the following options as required:

- a) (Optional) Change the default media type by highlighting the **Select Media Type** field and pressing **Enter** and then selecting an option from the popup menu.

The choices are **HDD**, **SSD**, or **Both**. However, **Both** is the default choice.

- b) (Optional) Change the default interface type by highlighting the **Select Interface Type** and pressing **Enter**, and then selecting an option from the popup menu.

The choices are **SAS**, **SATA**, **PCIe**, and **All**. Depending on the configuration of your system, combining SAS and SATA drives or drive group mixing might not be supported.

If you choose HDD for the media type, the possible options are **SAS**, **SATA**, and **Both**. PCIe is not a valid choice for HDD.

NOTE

PCIe does not appear as a valid choice if the controller does not support PCIe.

- c) (Optional) Change the default size of the logical sector by highlighting the **Logical Sector Size** and pressing **Enter**, and then selecting an option from the popup menu.

The choices are **512 B**, **4 KB**, and **Both**.

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

Alternatively, you can use the **Check All** and **Uncheck All** options at the bottom of the list of drives to either select all available drives or clear all of the selected 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

Make sure you 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 23](#) for more information.

- e) When you have selected the required drives for the new virtual drive, highlight **Apply Changes** and press **Enter** to create the virtual drive.

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

- f) 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.

- g) Highlight **Save Configuration** and press **Enter** to create the virtual drive.

A message appears confirming that the configuration is being created.

- h) Highlight **OK** and press **Enter** to acknowledge the confirmation message.

7. Highlight the **Virtual Drive Name** field, press **Enter**, and specify a name for the new virtual drive.

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

The options are **MB**, **GB**, and **TB**.

9. (Optional) Change the default values for **Strip Size**, **Read Policy**, **Write Policy**, **I/O Policy**, **Access Policy**, **Drive Cache**, **Disable Background Initialization**, **Default Initialization**, and **Emulation Type** (note that the **Emulation Type** field is suppressed for 4K virtual drives).

The following table describes the policies and their possible values or descriptions.

Table 21: Virtual Drive Policies

Property	Description
Strip Size	The virtual drive strip size per DDF. The possible values are as follows: <ul style="list-style-type: none"> • 64 KB • 128 KB • 256 KB • 512 KB • 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: <ul style="list-style-type: none"> • 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. Although the 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.

Property	Description
Write Policy	<p>The write cache policy for the virtual drive. The possible values are as follows:</p> <ul style="list-style-type: none"> • 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	<p>The I/O policy for the virtual drive. The possible values are as follows:</p> <ul style="list-style-type: none"> • 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 Full .
Emulation Type	<p>Allows you to set the emulation type on a virtual drive. The possible options are Default, Disable, or Force.</p> <p>The Force option forces the emulation to be set on a controller even when MFC settings do not support it. For more information, see Table 22, Emulation Settings.</p> <p>This field is suppressed for 4K virtual drives.</p>

The following table details the emulation settings and how the operating system reads these settings.

Table 22: Emulation Settings

Emulation Setting	Logical Sector Size in Operating System	Physical Sector Size in Operating System
Default	512 Byte	512 Byte
Disable	512 Byte	512 Byte
Force	512 Byte	4096 Byte

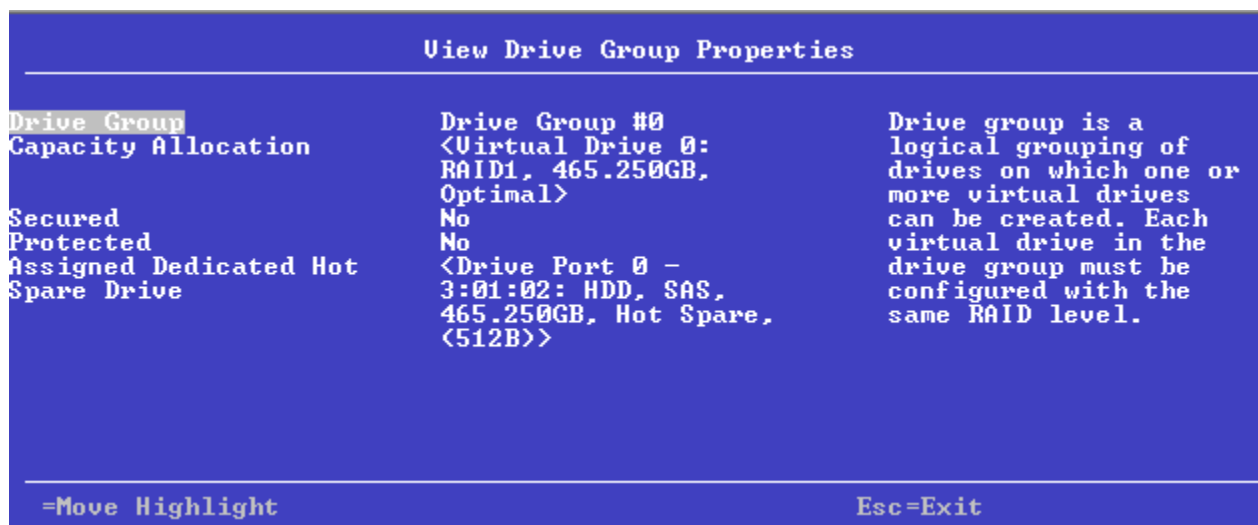
The following table describes the RAID levels that you can select when creating a new virtual drive. Some system configurations do not support RAID 6 and RAID 60.

Table 23: 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.

Viewing Drive Group Properties

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

Figure 28: 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 **Capacity Allocation** field and press **Enter**. The information appears in a popup 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.

Viewing Global Hot Spare Drives

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

Figure 29: View Global Hot Spare Drives Dialog



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

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.

To complete the command, follow these steps:

1. Highlight the brackets next to **Confirm** and press the spacebar.

An X appears in the brackets.

2. Highlight **Yes** and press **Enter**.

A success message appears.

3. Highlight **OK** and press **Enter**.

The HII Configuration Utility clears the configuration and returns you to the **Configuration Management** menu.

Make Unconfigured Good, Make JBOD, and Enable Security on JBOD

When you power off a controller and insert a new physical drive, if the inserted drive does not contain valid DDF metadata, the drive status is listed as either JBOD (Just a Bunch of Disks) or Unconfigured Good when you power on the system again.

NOTE

When the JBOD mode is enabled, the drive comes up as a JBOD drive; otherwise, it comes up as an Unconfigured Good drive.

When you power off a controller and insert a new physical drive, if the drive contains valid DDF metadata, its drive state is will be either **Undonefigured Bad** or **Foreign**. A new drive in the JBOD drive state is exposed to the host operating system as a stand-alone drive. You cannot use JBOD drives to create a RAID configuration because they do not have valid DDF records. First, the drives must be converted to the Unconfigured Good state.

If the controller supports JBOD drives, the **Configuration Management** menu of the HII Configuration Utility includes options for converting JBOD drives to Unconfigured Good, or vice versa. You can also enable security on the JBOD drives.

NOTE

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


Make Unconfigured Good

Perform these steps to change the JBOD drives to Unconfigured Good drives.

1. Highlight **Make Unconfigured Good** on the **Configuration Management** menu and press **Enter**.

The **Make Unconfigured Good** dialog appears, listing all the JBOD drives currently connected to the controller.

Figure 30: Make Unconfigured Good Dialog



Scroll down, if necessary, to view other drives that are listed.

- a) To select a specific JBOD drive and convert it to Unconfigured Good, highlight the drive and press the spacebar to select it.
- b) To select all the JBOD drives and convert them to Unconfigured Good drives, highlight **Check All** and press **Enter**.
- c) (Optional) To unselect all the drives that you have selected, highlight **Uncheck All** and press **Enter**.

ATTENTION

If one or more JBOD drives that you have selected have an operating system (OS) or a file system on them, 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 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 and unselect those JBOD drives that have an OS or a file system installed on them.

- Highlight **OK** (at the bottom of the JBOD drive list) and press **Enter** to convert the JBOD drives to Unconfigured Good drives.

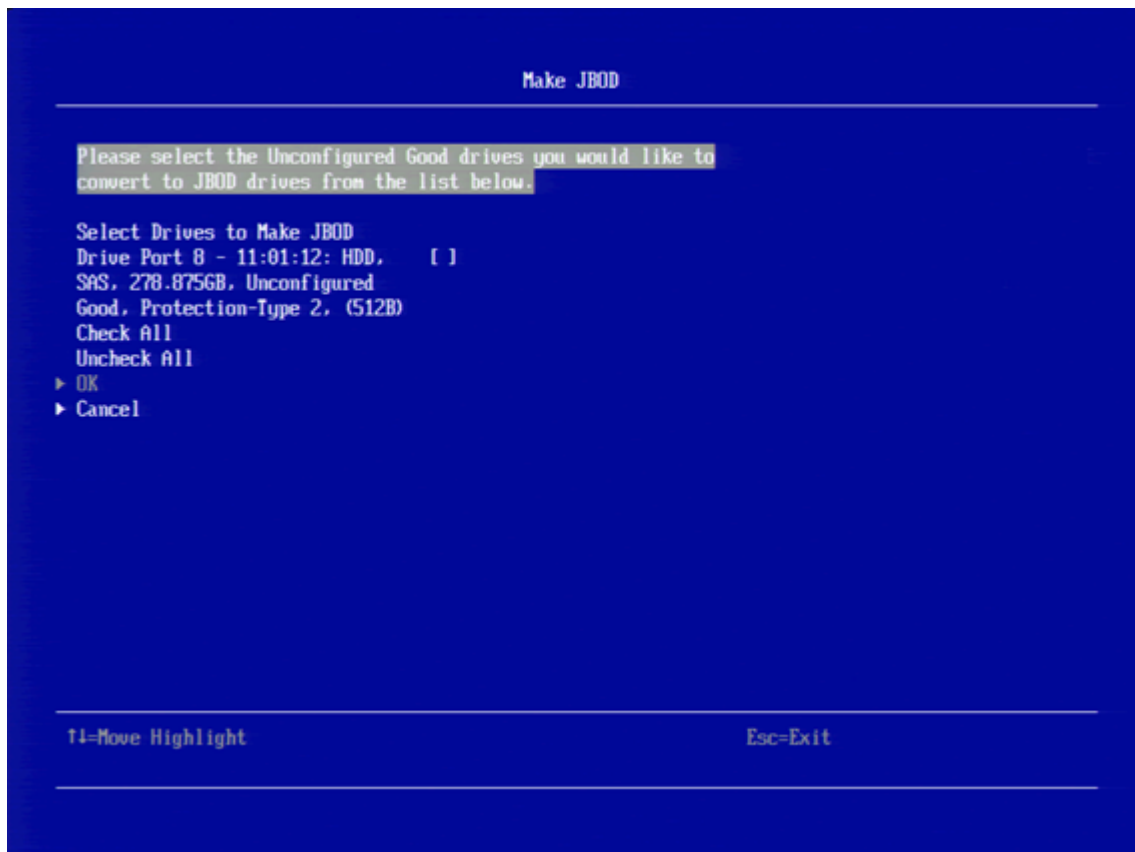
Make JBOD

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

- Highlight **Make JBOD** on the **Configuration Management** menu and press **Enter**.

The **Make JBOD** dialog appears listing all the Unconfigured Good drives currently connected to the controller.

Figure 31: Make JBOD Dialog



- To select a specific Unconfigured Good drive and convert it to JBOD, highlight the drive and press the spacebar to select it.
 - To select all the Unconfigured Good drives and convert them to JBOD drives, highlight **Check All** and press **Enter**.
 - (Optional) To clear all the drives that you have selected, highlight **Uncheck All** and press **Enter**.
- Highlight **OK** and press **Enter** to convert the Unconfigured Good drives to JBOD drives.

Enabling Security on JBOD

If you have SED-enabled JBOD drive that meets the prerequisites mentioned in [Managing Configurations](#), you can enable security on it. Follow these steps to enable the security on a JBOD drives.

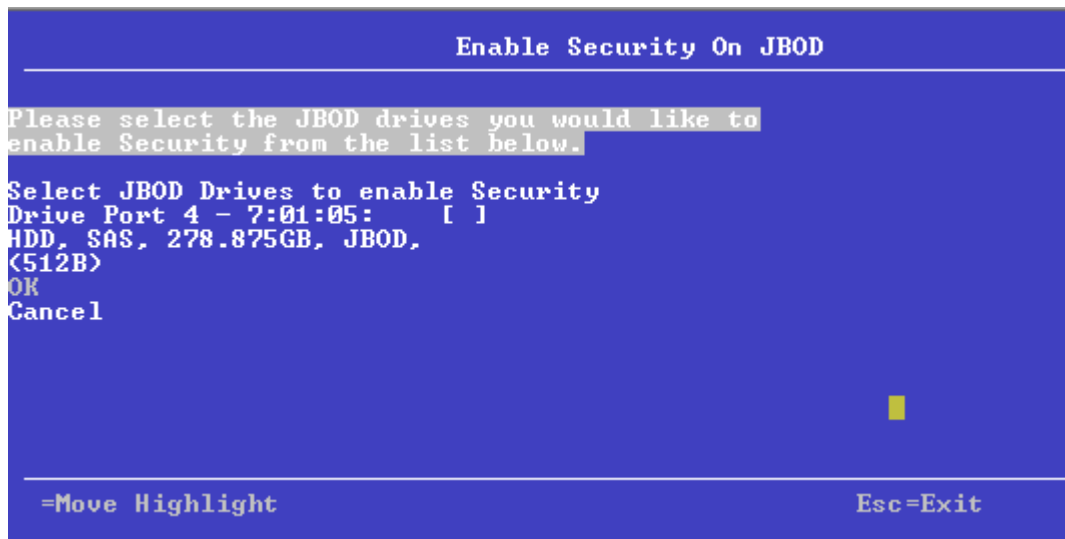
ATTENTION

All of the data on the drive is lost when you enable security on it. Therefore, back up any data that you want to keep.

1. Highlight **Enable Security on JBOD** on the **Configuration Management** menu and press **Enter**.

The **Enable Security on JBOD** dialog appears listing the SED-enabled JBOD drives currently connected to the controller.

Figure 32: Enable Security on JBOD Dialog



2. Highlight each JBOD drive to enable security on it and press the spacebar to select it.
3. Highlight **OK** and press **Enter** to enable security on the JBOD drive.
A message appears stating that the existing data in the drive would be lost if you proceed and prompting for your confirmation.
4. Highlight **Confirm** and press the spacebar, then highlight **Yes** and press **Enter**.
A success message appears.
5. Highlight **OK** and press **Enter**.
The HII Configuration Utility enables security on the JBOD drive and returns you to the **Configuration Management** menu.

Managing Foreign Configurations

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

Figure 33: Manage Foreign Configuration Dialog

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.

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 ACTIVE.
- 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.

When importing a foreign configuration, if the imported foreign virtual drive (VD) is marked as consistent in the source controller, the VD is marked as not consistent upon successful import if the target controller already has a VD configured and online. The target controller firmware has no way to determine if the import VD is consistent as other firmware operations, which can leave the drives inconsistent and shutdown might not allow the firmware to update DDF structures to reflect the state.

If a locked foreign drive is detected, then the HII Configuration Utility will display the menu option **Enter Security Key for Locked Drives** under the **Manage Foreign Configuration** menu.

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

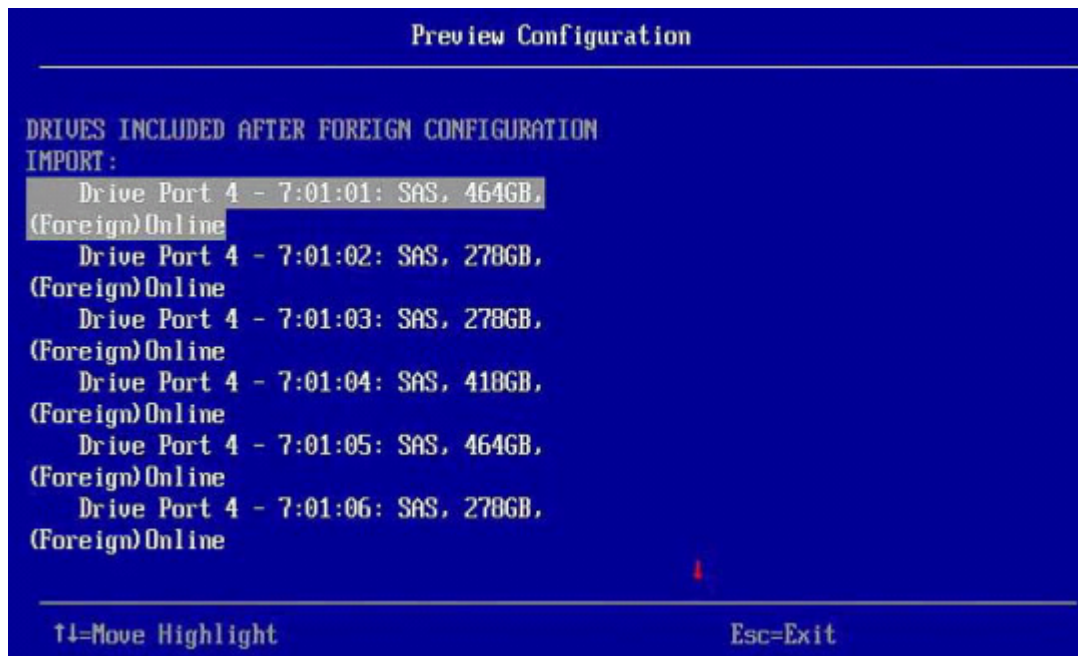
Figure 34: Enter Security Key for Locked Drives

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

Perform these steps to preview and import a foreign configuration.

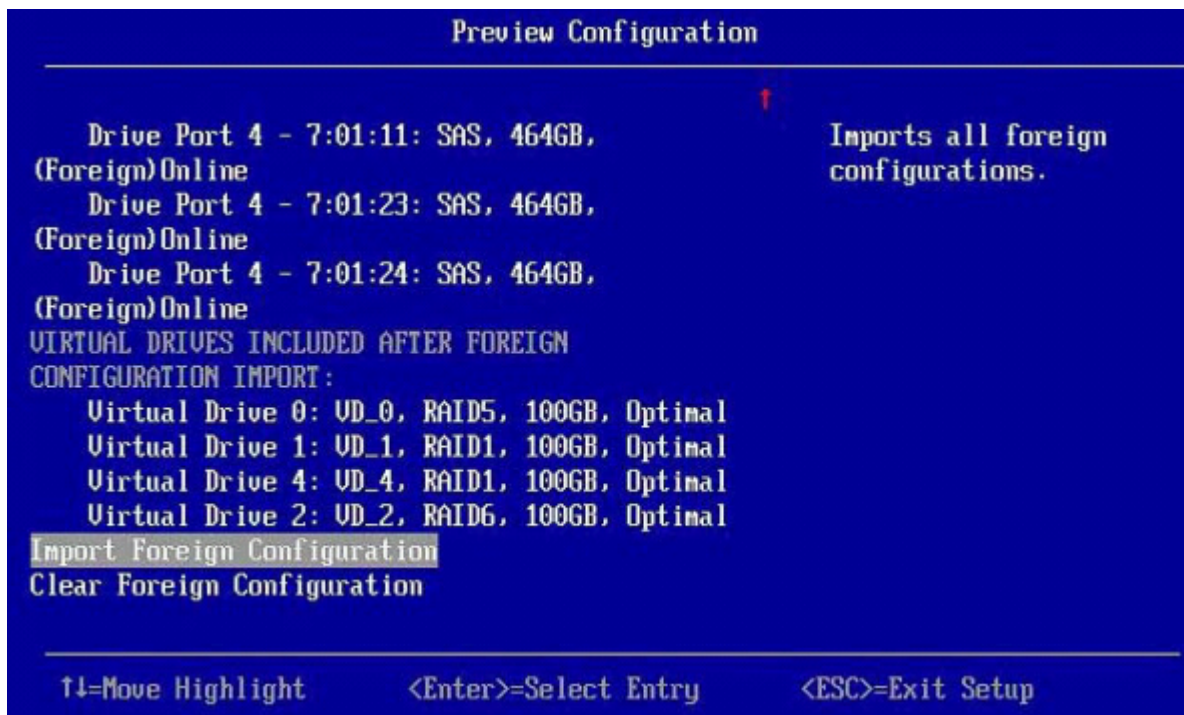
1. Highlight **Preview Foreign Configuration** on the **Manage Foreign Configuration** menu and press **Enter**.

The **Preview Configuration** dialog appears, listing information about the physical drives in the foreign configuration.

Figure 35: 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 36: 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.

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.


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 [Viewing Advanced Controller Management Options](#).
- To view additional controller properties, in the **Basic Properties** section, highlight **Advanced Controller Properties**.
For more information, see [Viewing Advanced Controller Properties](#).

Figure 37: Controller Management Dialog



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

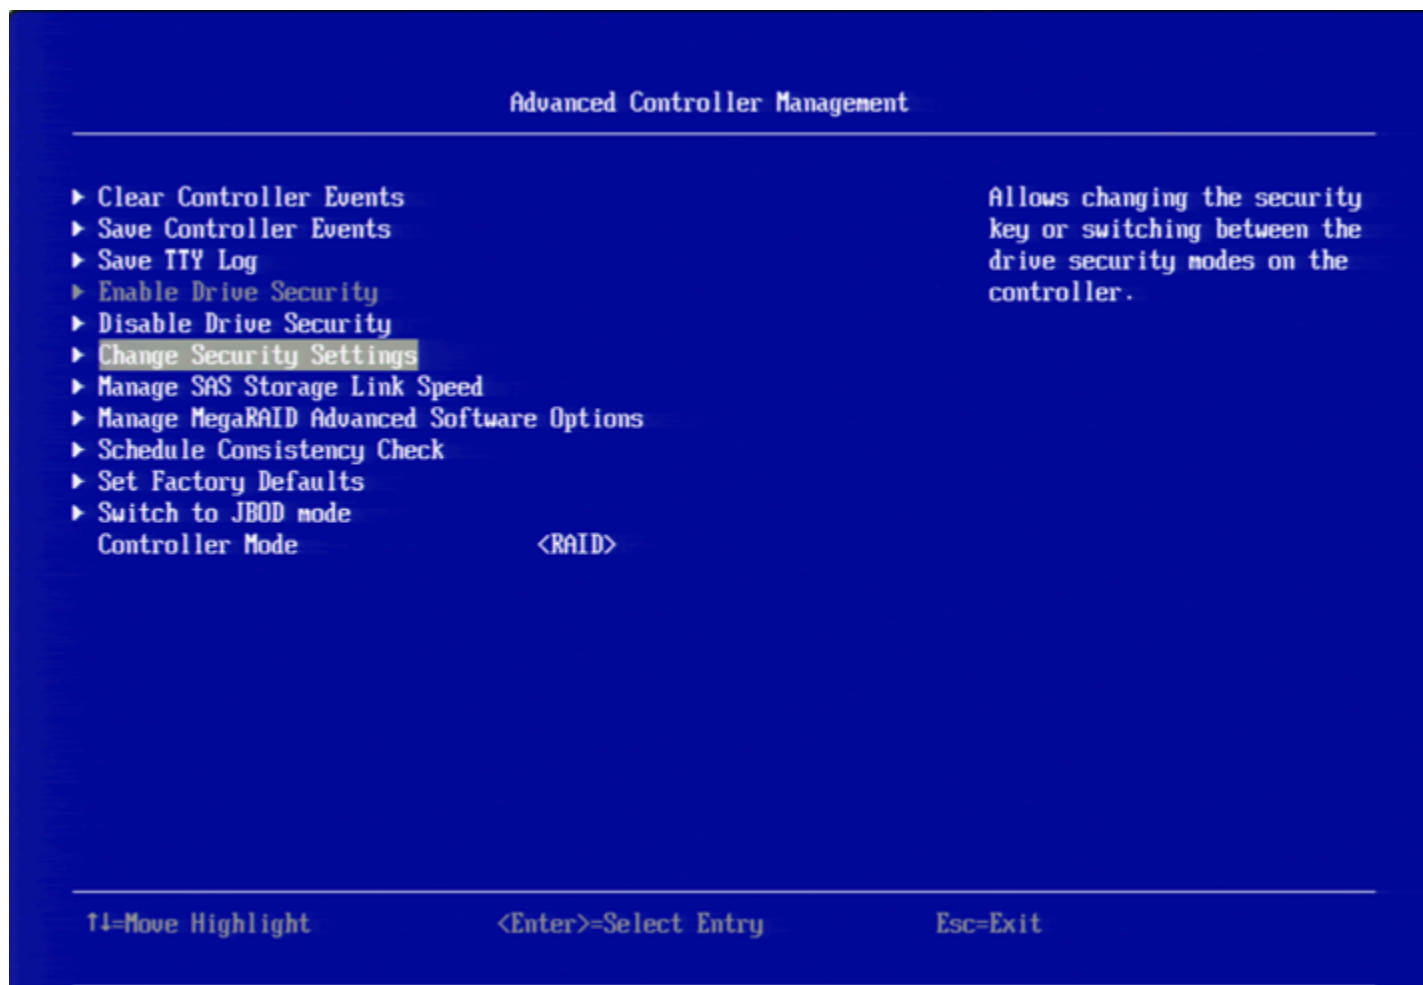
Table 24: Basic Controller Properties

Property	Description
Product Name	The marketing name of the controller.
Serial Number	The serial number of the controller.

Property	Description
Controller Status	The cumulative status of virtual drives and physical drives connected to the controller, plus the backup battery, the enclosure, and the NVDATA. The status is one of the following: <ul style="list-style-type: none"> • 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. Note: This property is applicable for legacy BIOS as legacy BIOS does not support booting from 4K block-size devices. Hence, you will not be able to see the 4K block-size devices such as JBODs or VDs that gets listed here for selection. However, you can use the UEFI environment for booting 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. 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

Viewing Advanced Controller Management Options

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

Figure 38: Advanced Controller Management Dialog

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

Table 25: 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 Settings	Changes the security settings or switches between drive security modes on the controller.

Property	Description
Manage SAS Storage 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 SAS Storage Link Speeds .
Manage PCIe Storage Interface	A lane represents a set of differential signal pairs, one pair for transmission and one pair for reception, similar to SAS phys. The Manage PCIe Storage Interface feature allows you to change the lane speed between a controller and expander or between the controller and a drive that is directly connected to the controller. MegaRAID 7.1 and later versions support both SAS/SATA topologies as well as PCIe topologies using the same device phys to manage the lane speed. For more information, see Managing PCIe Storage Interface .
Manage MegaRAID Advanced Software Options	Displays the activated MegaRAID 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. The MegaRAID Advanced Software Options are displayed only if the controller supports MegaRAID 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.
Switch to <RAID/JBOD> Mode	Used to switch between personality modes. The available personality modes are RAID and JBOD. If you switch between personality modes, for example, from RAID mode to JBOD mode, a reboot is required.
Manage Mode and Params	If your system is in a personality mode, for example, RAID mode, you can use this option to change the personality mode and its parameters. The available personality modes are RAID mode and JBOD mode.
Controller Mode	Displays the current personality of the controller.

Viewing 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 Properties** dialog lists some actions that you can perform on the controller.

- To view the details of the current profile and to select your desired profile, highlight **Profile Management** and press **Enter**. For more information, see [Managing Profiles](#).
- 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**, and 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** and press **Enter**.
For more information, see [Changing Task Rates](#).

Figure 39: Advanced Controller Properties Dialog



This dialog lists various properties, all of them cannot be 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** section, including the ones that are not visible.

Table 26: 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 on errors .
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. Applicable only for BIOS boot (see Table 24, Basic Controller Properties).
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.
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. 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. When the JBOD mode is enabled, the drive comes up as a JBOD drive; 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.
Large IO Support	By default, the large I/O support is enabled.
Enterprise Key Management	When selected, the following READ only properties are shown. <ul style="list-style-type: none"> • Capability: Indicates whether the controller firmware supports Enterprise Key Management mode. • Boot Agent: Indicates the status of the boot agent for Enterprise Key Management. • Configured: Indicates whether the Enterprise Key Management mode is configured on the controller.

Managing MegaRAID Advanced Software Options

The **Manage MegaRAID Advanced Software Options** dialog lists all the activated advance software options on the controller. You can configure the MegaRAID 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 MegaRAID Advanced Software Options** and press **Enter**.

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

Figure 40: Manage MegaRAID Advanced Software Options Dialog



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 predefined 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.

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 41: Schedule Consistency Check Dialog

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. Either check boxselect **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.

Managing Personality Mode

If your system is in personality mode (RAID or JBOD Mode), the firmware supports auto-configure options to allow the controller to function as appropriate for the user environment.


Personality mode can be configured to present a different controller name. The firmware switches the PNPID of the controller and reconfigures the controller features and usage models.

You can use the **Manage Personality Mode** setting to change the personality mode and its parameters.

1. In the **Advanced Controller Management** dialog, highlight **Manage Personality Mode** option and press **Enter**.

The **Manage Personality Mode** dialog appears.

Figure 42: Manage Personality Mode Dialog



2. Select the Personality Mode. The available options are **JBOD** and **None**.

If you choose the option **None**, you cannot enable or disable the **Auto Secure SED** option.

3. Highlight **Apply Changes** and press **Enter**.
4. Highlight **Confirm** and press the **spacebar**, then highlight **Yes** and press **Enter**.

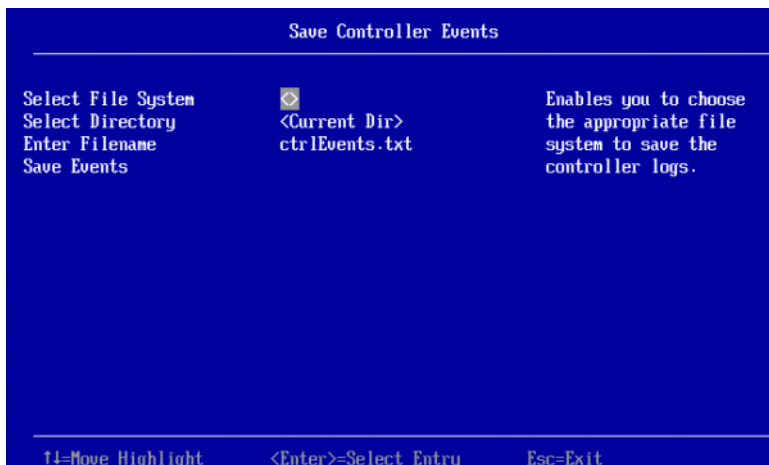
Saving or Clearing Controller Events

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 43: Save Controller Events Dialog



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 **popup**.

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**.

Enabling or Disabling Drive Security

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

Figure 44: Enable Drive Security (Choose Drive Security Mode) Dialog



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 Configuration 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.

Broadcom UEFI/HII drivers support interactive password primitive. If the OEM wants to use “Pause for password at boot” feature, which is part of the security feature, the system BIOS must support ECR 1085 and 1174; otherwise you will not be able to use this feature.

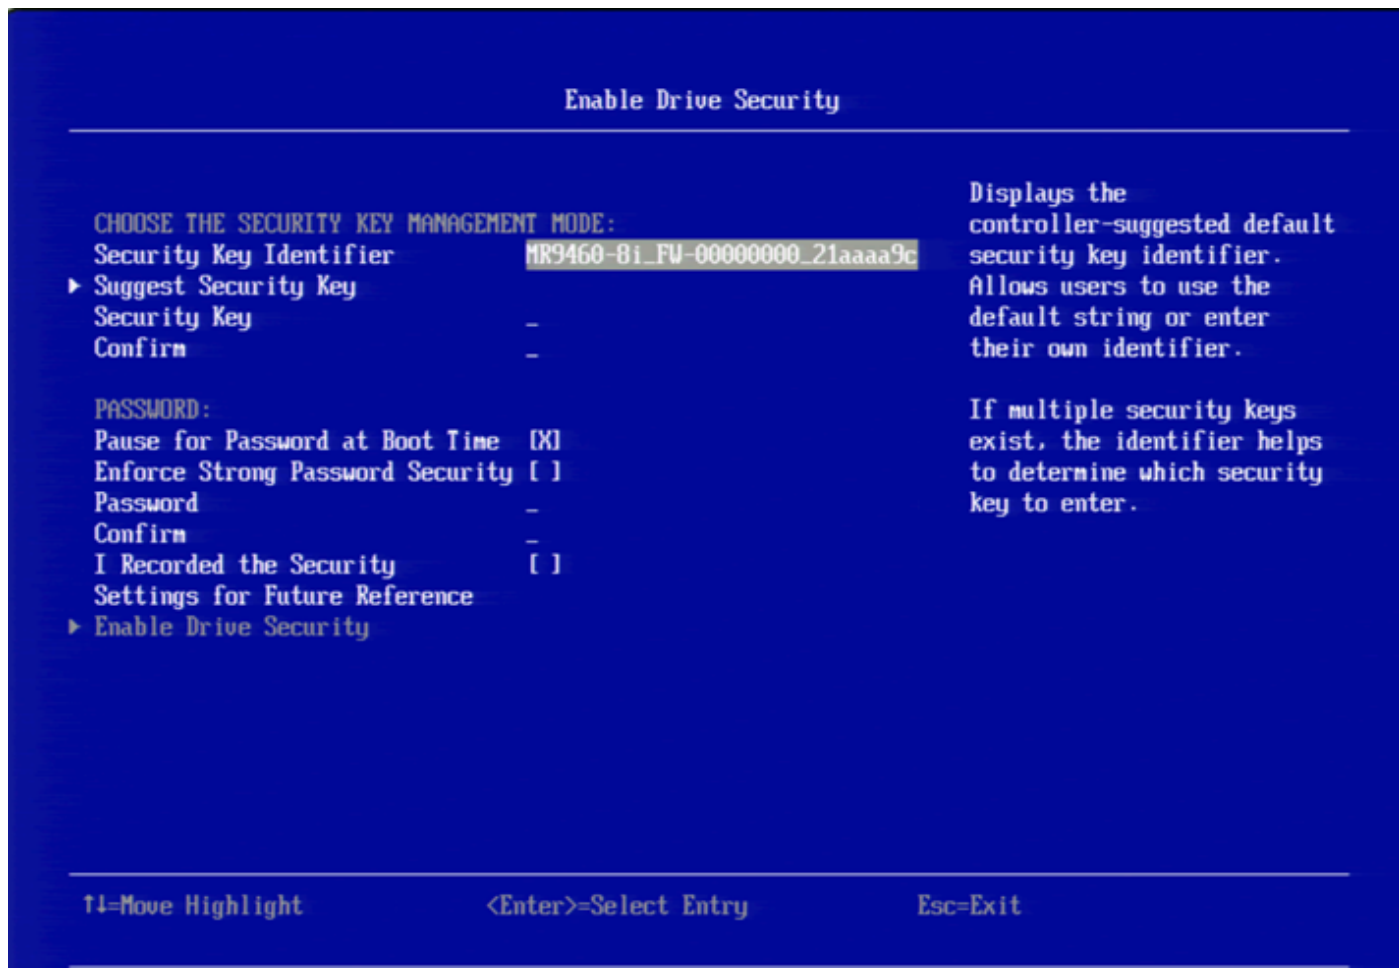
The system BIOS should use password primitive's prompt as a dialog title because this is an interactive password, and it is controlled by IHV. For example, if the password primitive's prompt is `Enter Your Input Here`, the dialog title should use the same name.

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 45: Enable Drive Security Dialog



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 8 and 32 characters and must contain at least one number, one lowercase letter, one uppercase letter, and one nonalphanumeric 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 8 and 32 characters and must contain at least one number, one lowercase letter, one uppercase letter, and one nonalphanumeric 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 46: 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.

Changing Security Settings

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 settings.

1. Highlight **OK** and press **Enter**.

The following dialog appears.

Figure 47: Change Security Settings Dialog

Change Security Settings

Enter a New Security Key Identifier	<input type="text"/>	A field that is case-sensitive; must be between 8 and 32 characters; and contains at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g., >@+).
Use the Existing Security Key Identifier	<input type="checkbox"/>	
Enter Existing Security Key	-	
Use the Existing Security Key	<input type="checkbox"/>	
▶ Suggest Security Key	-	
Enter a New Security Key	-	
Confirm	-	

PASSWORD:

Pause for Password at Boot Time ☒

Enforce Strong Password Security ☐

Password -

Confirm -

I Recorded the Security Settings for Future Reference ☐

▶ Save Security Settings

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

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 **Enter A New Security Key** field, press **Enter**, and type the new security key.

This 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 **Enter a New 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 Time** option by highlighting it and pressing the spacebar.

The contents of this field will be empty when you select this check box.

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 **Save Security Settings** 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.

Saving the TTY Log

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

Figure 48: Save TTY Log Dialog

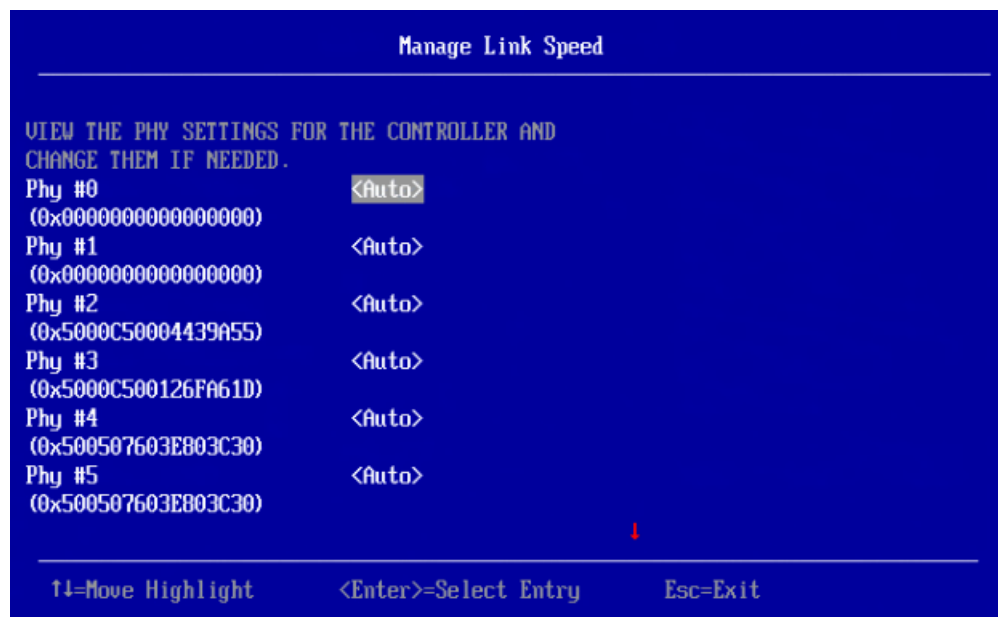
Save TTY Log		
File Systems	<HANTOOL>	Enables you to choose the appropriate directory to save the controller logs. The default (root) directory will be selected upon entering this form.
Select File System		
Directories	<DOS>	
Select Directory		
Enter Filename	ttyLog.txt	
Entries to Save	<All>	
Save Log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

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

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

Managing SAS Storage Link Speeds

The Manage SAS Storage 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 SAS Storage Link Speed** on the **Advanced Controller Management** dialog. The default settings for all phys is **Auto**.

Figure 49: Manage Link Speed Dialog

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 **popup**.
2. Select an option from the popup menu.

The link speed values are **Auto**, **3 GB/s**, **6 GB/s**, and **12 GB/s**.

3. Scroll to the bottom of the phy list, highlight **OK**, and press **Enter**.

Managing PCIe Storage Interface

The manage PCIe storage interface allows you to manage and change the lane speed and link width between a controller and an expander or between the controller and a drive that is directly connected to the controller. For managing the PCIe storage interface, navigate to **Manage PCIe Storage Interface** on the **Advanced Controller Management** dialog. By default, the lane speed in the controller is **8 GT/s** or the value last saved by you.

Figure 50: Manage PCIe Storage Interface Dialog

Follow these steps to change the lane 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 **Unknown**, **2.5 GT/s**, **5 GT/s**, and **8 GT/s**.

3. Scroll to the bottom of the phy list, highlight **Apply Changes**, and confirm by pressing the spacebar, then highlight **Yes** and press **Enter**.

Managing Profiles

Profile management allows you to have multiple configurations supported under each personality mode. Profiles customize the controller to deliver the best performance for that configuration. For example, a profile with no PCIe device support can support a higher Queue Depth than a profile that supports 32 PCIe devices. When you choose profile management through HII, the firmware provides a list of profiles that you can select for the current personality.

When you select **Profile Management** from **Advanced Controller Properties**, the following dialog is displayed.

Figure 51: Profile Management Dialog

Profile Management

Personality Mode	<JBOD>	Displays the current personality of this controller.
Number of Profiles	[7]	
Choose Profile	<20_PD240 (Default)>	

DETAILS OF THE PROFILE:

Profile ID	[20]
Requested Profile ID	[10]
Compatible	<Yes>
Max PD Count	[240]
Max VD Count	[240]
Write Back Supported	<No>

Set Profile

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

The **Profile Management** dialog lists the following details.

Table 27: Profile Management Dialog Details

Property	Description
Personality Mode	Indicates the current personality of the controller.
Number of Profiles	Indicates the number of profiles supported for the current personality of the controller.
Choose Profile	<p>Indicates the name of the profile for the current personality of the controller. It also allows you to choose an appropriate profile from the available list.</p> <p>Note: When you have selected a particular controller profile, for example PCIe4 as your profile, using the Set Factory Defaults option which is available under the Actions menu does not change your selected profile. For example, if PD 64 is the default controller profile and if you have selected PCIe4 as your controller profile, the selected profile (PCIe4) is retained by the system even if you use Set Factory Defaults option to restore factory settings on your controller.</p>

Property	Description
Profile ID	Indicates the unique identity of the selected profile. Note: The maximum number of PDs can fall in the range of 8 - 64. Check with the OEM to determine your maximum.
Requested Profile ID	Indicates the requested profile. This is displayed only after you have selected the required profile from the available list of profiles.
Compatible	Indicates whether the chosen profile is compatible with your current drive group topology.
Max PD Count	Indicates the maximum number of physical drives supported by the controller for the selected profile.
Max VD Count	Indicates the maximum number of virtual drives supported by the controller for the selected profile.
Max PCIe Count	Indicates the maximum PCIe drives supported by the controller for the selected profile.
Write Back Supported	Indicates if the Write Back functionality is supported.

Follow these steps to set or change the profile:

1. On the **Profile Management** dialog, highlight **Choose Profile** and press **Enter**.
2. From the drop-down list, highlight the required profile and press **Enter** to select it.

Table 28: Available Profiles

Personality	Profile ID	Profile Name	Description
RAID	10	PD 64	This profile supports a maximum of 64 physical and virtual drives. This profile does not support any NVMe drives.
RAID	11	PCIe4	This profile supports a maximum of 64 virtual drives and a maximum of 4 NVMe drives.
RAID	12	PD240	This profile supports a maximum of 240 physical and virtual drives. This profile does not support NVMe drives.
RAID	13	PD64-PCIe4	This profile supports a maximum of 64 physical and virtual drives and supports a maximum of 4 NVMe drives.
JBOD	20	PD 240	This profile supports a maximum of 240 SAS/SATA drives. This profile does not support NVMe drives. (RAID 0, 1, 10 only).
JBOD	21	PD 64-PCIe8	This profile supports a total of 64 SAS/SATA/NVMe drives. This profile supports a maximum of 8 NVMe drives.
JBOD	23	PD 64	This profile supports up to 240 SAS/SATA drives. This profile does not support NVMe drives.

3. Highlight **Set Profile** and press **Enter**.
4. Highlight **OK** and press **Enter** to switch to your selected profile.
5. Reboot the system for the changes to take effect.

Downgrading the Firmware When Profiles Are Selected

If you have selected a particular profile from the **Profile Management** dialog and are trying to downgrade to a previous firmware version, you may not be able to downgrade because the profile selected by you may not be available in the

previous firmware version. In these cases, you must change the profile in your current version of the firmware and then downgrade to a previous version.

For example, if you are using MegaRAID version 7.3 and you have selected **PD240** (Profile ID 12) as the profile (**PD240** supports a maximum of 240 physical and virtual drives; it does not support NVMe drives) and try to downgrade to MegaRAID 7.2 or MegaRAID 7.1, the downgrade will fail because the MegaRAID version 7.2 does not support the **PD240** profile. In this case, you must change the profile ID to 10 (**PD 64**) and then downgrade.

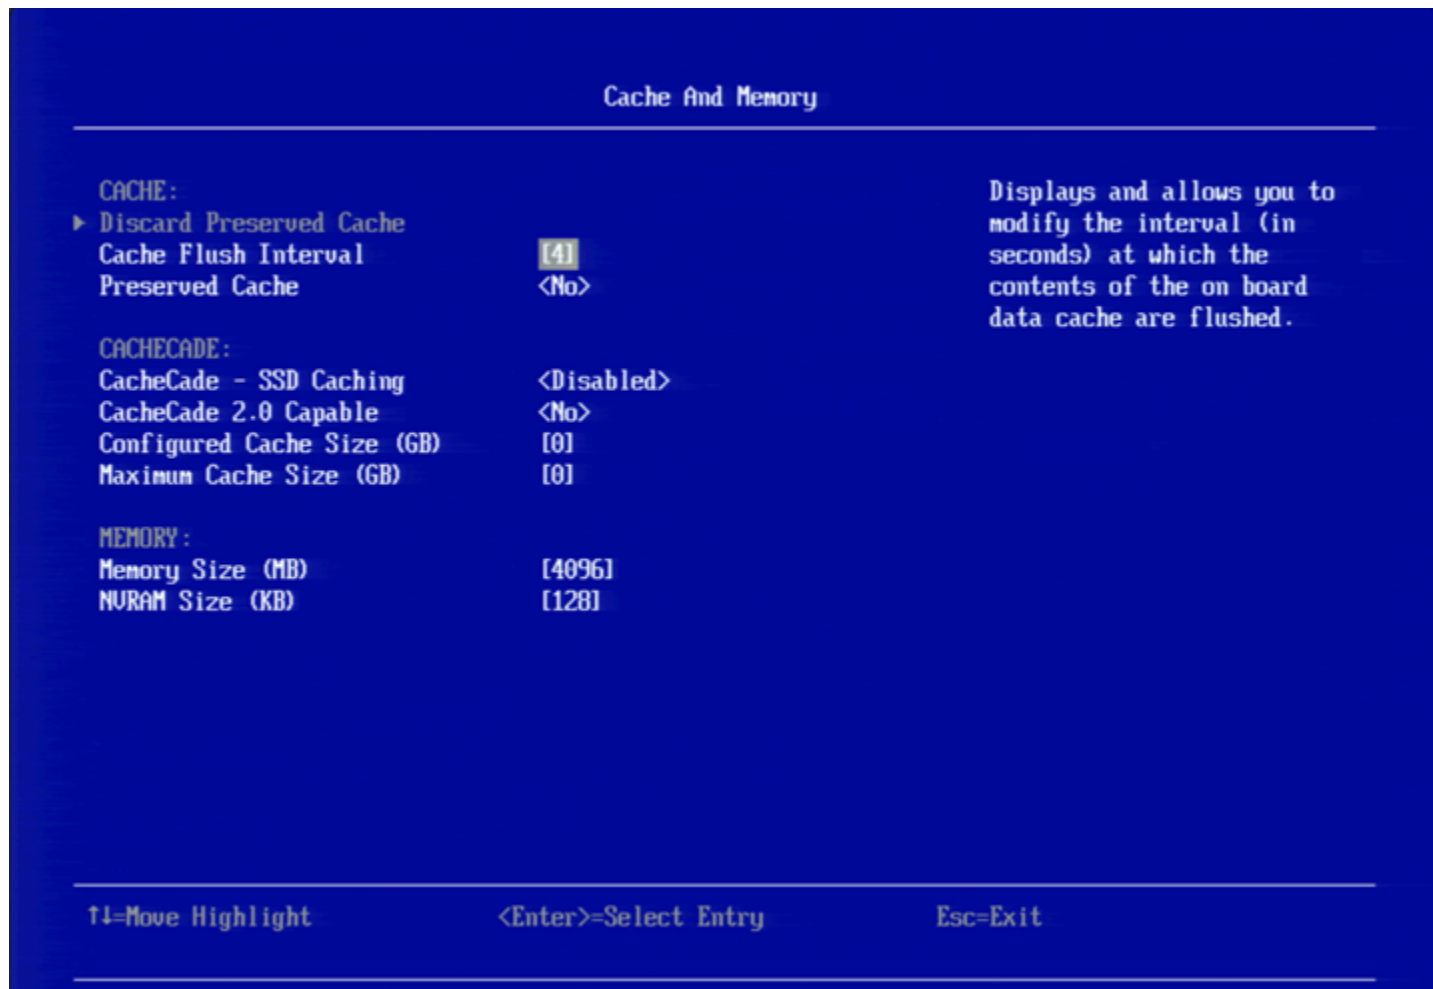
Table 29: Profiles and Drive Support

Personality	Profile ID	Profile Name	
RAID	10	PD 64	Optimized for 64 or fewer SAS/SATA drives. No NVMe drives supported. Maximum of 64 virtual drives.
RAID	11	PCIe4	Optimized for 4 or fewer NVMe drives. No SAS/SATA drives supported. Maximum of 64 virtual drives.
RAID	12	PD240	Optimized for 240 or fewer SAS/SATA drives. No NVMe drives supported. Maximum 240 Virtual Drives.
RAID	13	PD64-PCIe4	Flexibility for 64 total SAS/SATA/NVME drives. Maximum of 4 NVMe drives. Maximum 64 Virtual Drives.
JBOD	20	PD 240	Optimized for 240 or fewer SAS/SATA drives. No NVMe drives supported. Maximum 240 WT Virtual Drives. RAID (0, 1, 10 only).
JBOD	21	PD 64-PCIe8	Flexibility for 64 total SAS/SATA/NVME drives. Maximum of 8 NVMe drives. Maximum 64 WT Virtual Drives. RAID (0,1, 10 only).
JBOD	23	PD 64	Optimized for 240 or fewer SAS/SATA drives. No NVMe drives supported. Maximum of 240 WB Virtual Drives.

Setting Cache and Memory Properties

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

Figure 52: Cache and Memory Dialog



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.

Running a Patrol Read

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

Figure 53: Patrol Read Dialog

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 and options supported by your controller.

- 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.
- 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**.
The maximum numeric value that you can enter as the rate is 100.
- 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**.
- Highlight **Apply Changes** and press **Enter**.
The new settings are saved in the controller properties.

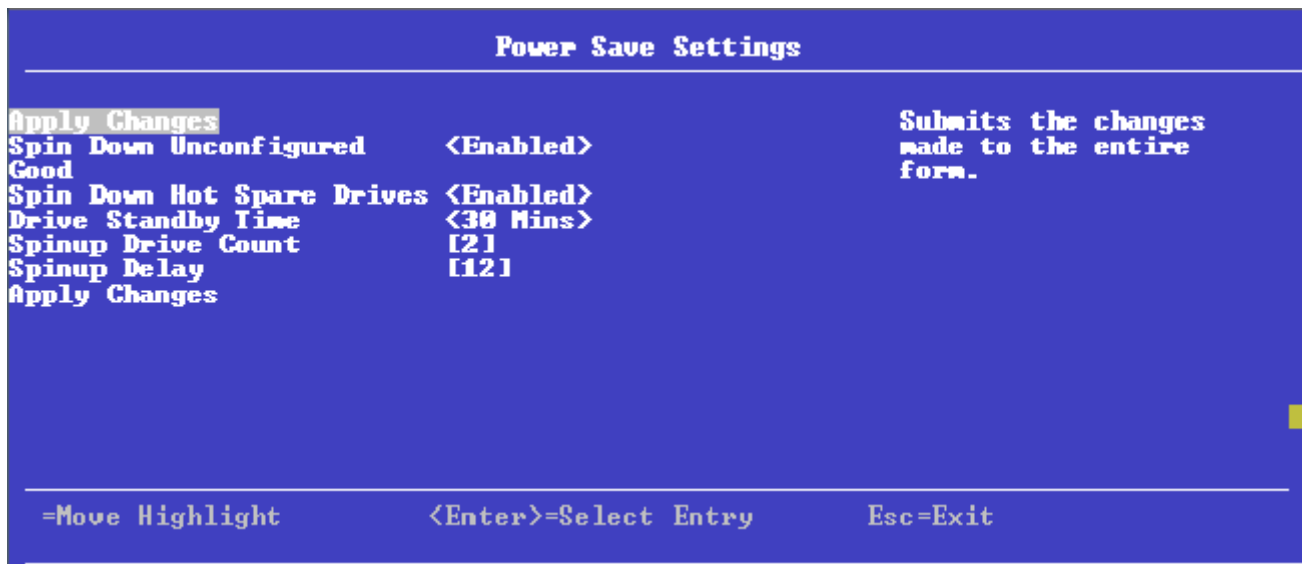
To start a patrol read without changing the patrol read properties, follow these steps:

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.

Changing Power Save Settings

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

Figure 54: Power Save Settings Dialog



The preceding 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 and 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 an end time and press **Enter**.
10. Highlight **Apply Changes** and press **Enter**.

The new settings are saved in the controller properties.

Setting Emergency Spare Properties

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

Figure 55: Spare Dialog

Spare		
Apply Changes		Submits the changes made to the entire form.
Emergency Spare	<Unconfigured Good and Global Hotspare>	
Emergency for SMARTer	<Disabled>	
Persistent Hot Spare	<Disabled>	
Replace Drive	<Enabled>	
Replace Drive on SMART Error	<Disabled>	
Apply Changes		
=Move Highlight <Enter>=Select Entry Esc=Exit		

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

Changing Task Rates

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

Figure 56: Task Rates Dialog

Task Rates	
Apply Changes	Submits the changes made to the entire form.
Background Initialization <BGI> Rate	[30]
Consistency Check Rate	[30]
Patrol Read Rate	[30]
Rebuild Rate	[30]
Reconstruction Rate	[30]
Apply Changes	

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

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

Follow these steps to change the task rates.

NOTE

You can only view the properties and 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 and 100 percent. At 0 percent, the initialization operation runs only if the firmware is not doing anything else. At 100 percent, the initialization operation has a higher priority than I/O requests from the operating system. For best performance, use an initialization rate of approximately 30 percent.
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 operation 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**.

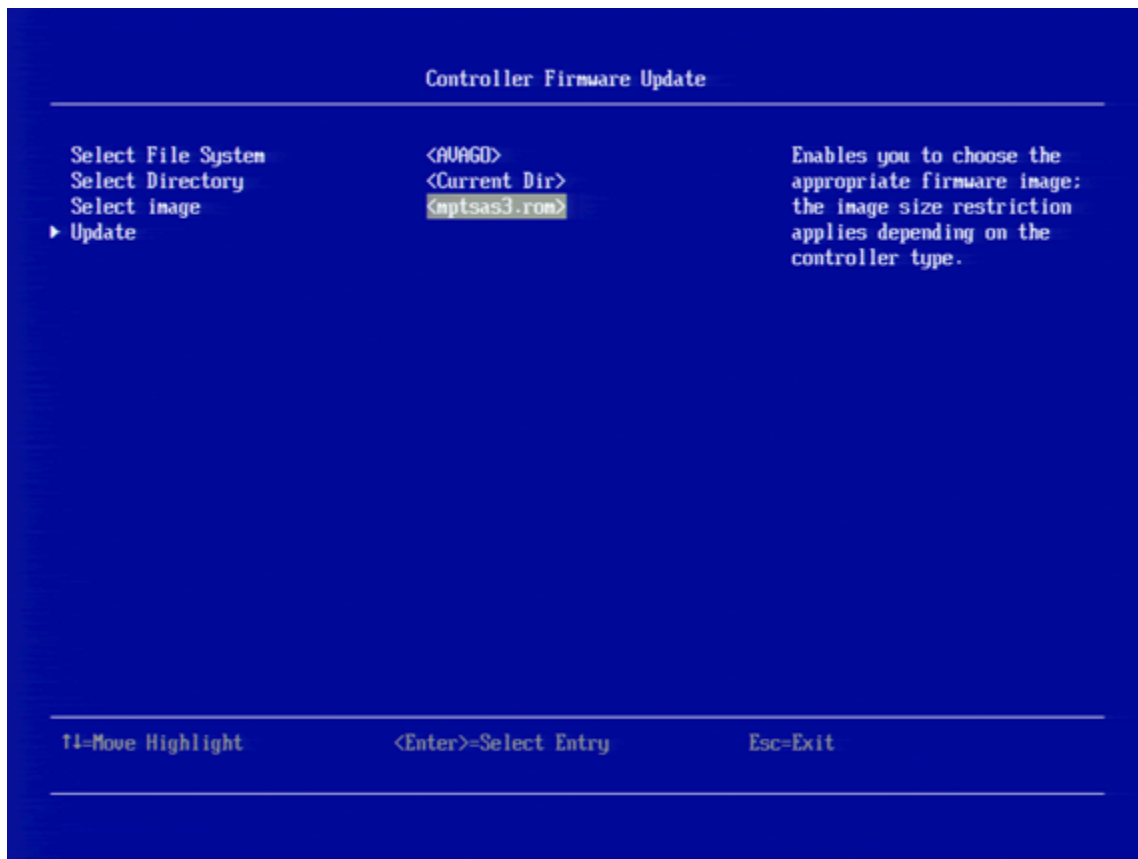
The reconstruction rate is the percentage of the compute cycles dedicated to reconstructing data on drives on this controller. You can configure the reconstruction rate between 0 percent and 100 percent. At 0 percent, the reconstruction operation runs only if the firmware is not doing anything else.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

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

The new settings are saved in the controller properties.

Upgrading the Firmware

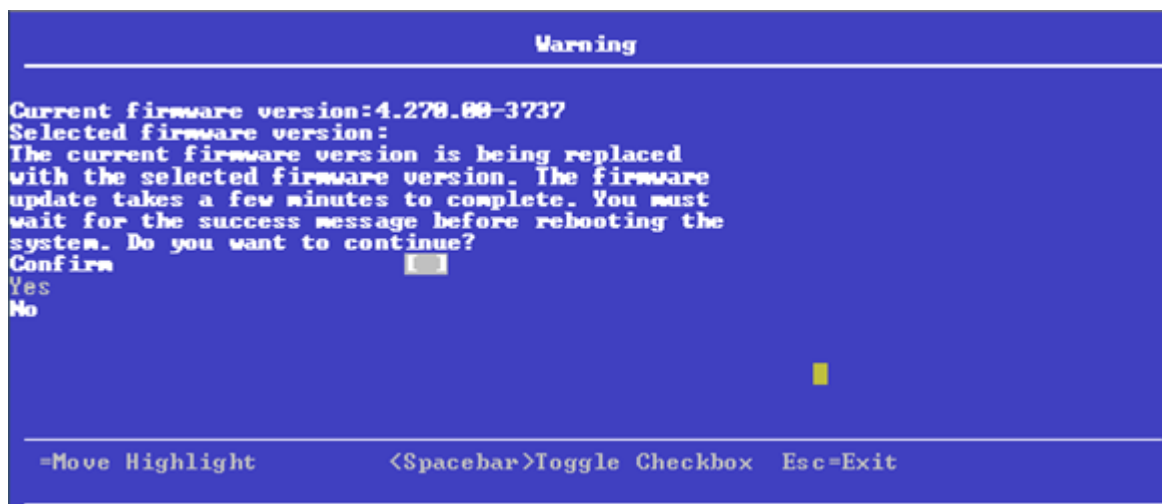
The following dialog appears when you select **Update Firmware** from the **Dashboard View**. For a list of limitations, see [Online Firmware Upgrade and Downgrade](#).

Figure 57: Controller Firmware Update Dialog

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 58: Warning Dialog



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.

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 59: Virtual Drive Management Dialog



The menu lists all of 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 60: Virtual Drive Management Dialog



This dialog lists the following basic virtual drive properties.

Table 30: 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 .Virtual drive size of floating data types up to three decimal places is supported. Some of the screens in this chapter may not show this feature.

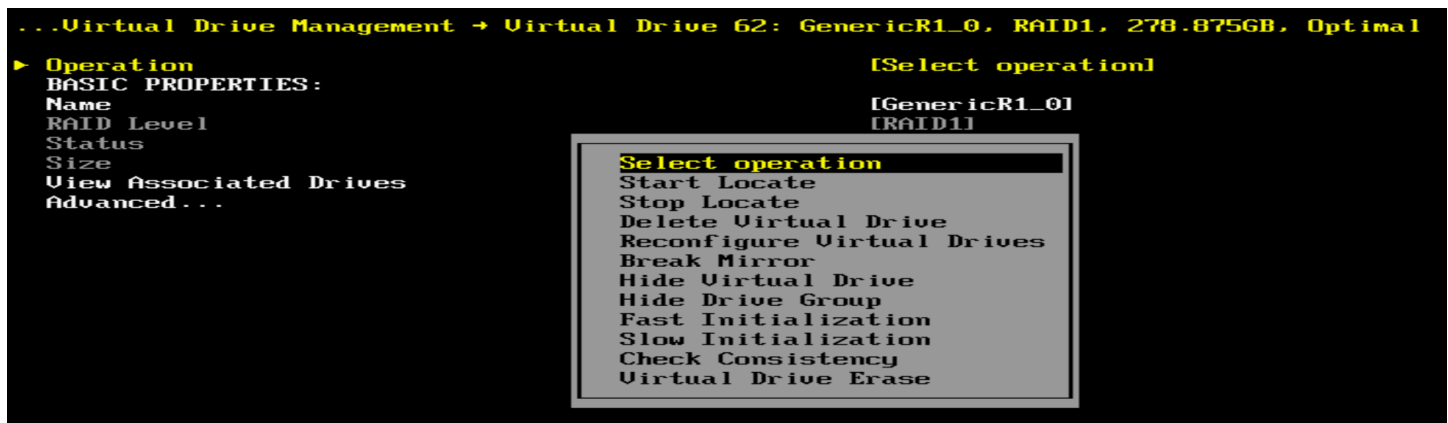
For information on how to perform virtual drive operations, see [Selecting Virtual Drive Operations](#).

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](#).

Selecting Virtual Drive Operations

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

Figure 61: Virtual Drive Operations Popup Menu

Other options, such as **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.

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.

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. You cannot close the **Group Initialization** dialog and perform any other operation on the LSA application until this process completes.

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.

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.

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.

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.

Reconfiguring a Virtual Drive

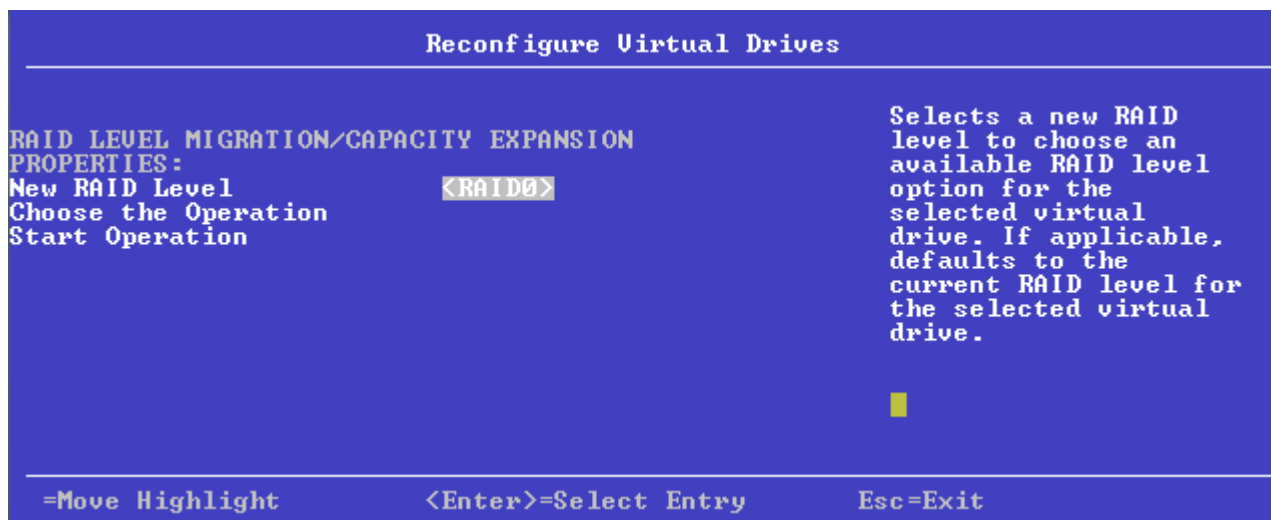
You can reconfigure a virtual drive by changing its RAID level, by adding physical drives to it or by doing both of these actions. When performing these changes, however, you must observe the maximum drive and minimum drive restrictions for the various RAID levels. See [Table 23, RAID Levels](#), for more information.

To reconfigure a virtual drive, perform these steps:

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 62: Reconfigure Virtual Drives Dialog



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**.

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**.
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**, **SSD**, or **Both**. However, **Both** is defaulted as a choice. 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**, **PCIe**, and **All**. Depending on the configuration of your system, combining SAS and SATA drives in a virtual drive might not be supported.

NOTE

PCIe does not appear as a valid choice if the controller does not support PCIe.

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 Configuration Utility displays an error message when you try to create the virtual drive. For example, RAID 1 virtual drives use exactly two drives and RAID 5 virtual drives use three or more drives. See [Table 23, RAID Levels](#) 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 Configuration Utility displays a message warning you that the remaining free capacity on the larger drives will be unusable.

The HII Configuration Utility returns you to the **Reconfigure Virtual Drives** dialog.

Removing Drives from a Configuration

Perform the following steps to remove drives from a configuration while reconfiguring a virtual drive.

NOTE

Usable Capacity Reduction is not possible without a RAID level migration.

1. If you select the **Remove Drives** option and press **Enter**, the following dialog appears.

Figure 63: Select Drives – Remove Drives Dialog

```
... Drive 62: GenericR1_0, RAID1, 278.875GB, Optimal → Reconfigure Virtual Drives → Select Drives
To migrate from current RAID Level to New RAID Level, you can either Add the drives or Remove the
drives.
▶ Choose the operation you want to perform [Remove Drives]

SELECT THE DRIVE TO REMOVE:
Drive C1 :01:14: HDD, SAS, 465.250GB, Online, [ ]
Span#0, (512B)
Drive C1 :01:11: HDD, SAS, 465.250GB, Online, [ ]
Span#0, (512B)
Check All
Uncheck All
Apply Changes
```

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 Configuration Utility returns you to the **Reconfigure Virtual Drives** dialog.

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.

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.

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.

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](#).

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 the 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.

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](#).


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 64: Advanced Virtual Drive Properties 1 Dialog



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

Table 31: 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 32: 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: <ul style="list-style-type: none"> Write-Through (WThru) The controller sends a data transfer completion signal to the host when the virtual drive has received all of the data and has completed the write transaction to the drive. Write-Back (WBack) The controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a drive write transaction. Data is written to the virtual drive in accordance with policies set up by the controller. These policies include the amount of dirty and clean cache lines, the number of cache lines available, and the elapsed time from the last cache flush. Force Write Back.
Default Write Cache Policy	Displays the default write cache policy of the virtual drive.

Property	Description
Disable Background Initialization (BGI)	Specifies whether background initialization is enabled or disabled. When BGI is enabled, the firmware runs the initialization process in the background. When BGI is disabled, the initialization process does not start automatically and does not run in the background.
Read Cache Policy	<p>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 <i>default</i> read policy. The possible options follow:</p> <ul style="list-style-type: none"> • 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)	<p>The I/O policy for the virtual drive. The possible values are as follows:</p> <ul style="list-style-type: none"> • 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.

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 65: Drive Management Dialog



The preceding dialog lists the following basic drive properties for the selected drive.

Table 33: Basic Physical Drive Properties

Property	Description
Device ID	The device ID of the currently selected drive.
Connection	The connection of the drive.
Enclosure Position/Backplane ID	The position of the enclosure or the backplane.
Slot Number	The slot number of the drive.
Status	The status of the drive, such as Online , Ready , Available , or Failed .
Size	The drive capacity, in GB. 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.
Type	The device type of the drive, which is normally Disk .
Model	The model number of the drive.

Property	Description
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.
View 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](#).

Performing Drive Operations

When you highlight the **Select operation** field and press **Enter**, a popup drive operations menu appears.

Figure 66: 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, which can be **Online**, **Offline**, **JBOD**, **Unconfigured Good**, **Unconfigured Bad**, **Global Hot Spare**, and **Dedicated Hot Spare**.

The following sections describe the available drive operations.

NOTE

The drive operations run on the currently selected drive. To run an operation on a different drive, press **Esc** to return to the **Drive Selection** menu, highlight the drive you want to select, press **Enter** to select it, and return to this dialog.

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.

Making a Drive Unconfigured Bad, Unconfigured Good, or JBOD

When you force a drive offline, it enters the *Unconfigured Bad* state.

When you power off a controller and insert a new physical drive, if the inserted drive does not contain valid DDF metadata, the drive status is listed as either JBOD (Just a Bunch of Disks) or Unconfigured Good when you power on the system again. When the JBOD mode is enabled, the drive comes up as a JBOD drive; otherwise, it comes up as an Unconfigured Good drive.

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 Bad** or **Foreign**.

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 an Unconfigured Good, or a JBOD drive.

1. Open the popup drive operations menu, highlight **Make Unconfigured Good**, **Make Unconfigured Bad**, or **Make JBOD**, and press **Enter**.
2. Highlight **Go**, which appears beneath **Operation**, and press **Enter**.

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.

Enabling Security on JBOD

If you have SED-enable JBOD that meets the prerequisites mentioned in [Managing Configurations](#), you can enable security on it. Follow these steps:

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**.

Replacing a Drive

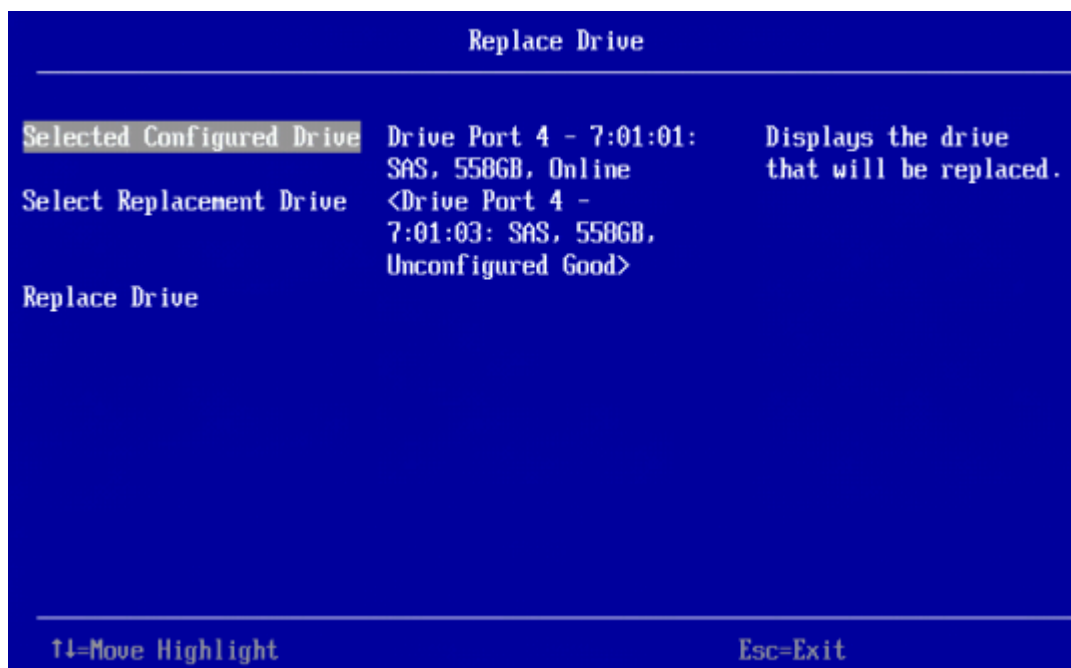
You might want to replace a drive that is a member of a redundant virtual drive connected to the controller if the drive shows signs of failing. Before you start this operation, be sure that an available Unconfigured Good replacement drive is available. The replacement drive must have at least as much 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 67: Replace Drive Dialog



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 Configuration Utility while the replacement operation runs.

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** warning appears.
3. Highlight **Confirm**, and press the spacebar to confirm the operation.
 4. Highlight **Yes** and press **Enter**.

The selected drive is forced offline.

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 popup 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.

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.

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.

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.

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 68: Associate Virtual Drives to the Dedicated HotSpare Drive Dialog

```
...465.250GB, Unconfigured Good, (512B) → Associate Drive Groups to The Dedicated Hot Spare Drive
▶ Drive Group #1, RAID1, 278.875GB [ ]
  Check All
  Uncheck All
  OK
  Cancel
```

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.

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 VD's.

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

The **Unassign Hotspare Drive** warning 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.

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.)


- Highlight **Confirm** and press the spacebar to confirm the operation.
- Highlight **Yes** and press **Enter**.

A message appears indicating that the initialization or erase operation has started.

- 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 69: Initialize Progress Indicator



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

NOTE

The progress indicator refreshes automatically.

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

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

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

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

A progress indicator shows the percentage completion of the rebuild operation. This indicator refreshes automatically, and the `Rebuild Drive Success` message appears.

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 **Cryptographic Erase** and press **Enter**.
2. Highlight **Go** and press **Enter**.

The **Cryptographic Erase** warning appears.

3. Highlight **Confirm** and press the spacebar to confirm the operation.
4. Highlight **Yes** and press **Enter**.

A message appears indicating that the cryptographic 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 cryptographic erase process, highlight **Stop**, and press **Enter**.

NOTE

A progress indicator shows the percentage completion of the operation. This indicator refreshes automatically.

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.

Making a JBOD

If your controller is in JBOD behavior mode and if you have not created any JBODs so far, the Make JBOD appears as an option when you navigate to the **<Select operation>** under **Drive Operations** dialog.

NOTE

The Make JBOD option only appears for Unconfigured Good drives.

Perform the following steps to Make a JBOD:

1. Open the popup drive operations menu, highlight **Make JBOD** and press **Enter**.
2. Highlight **Go** and press **Enter** to make an unconfigured good drive as a JBOD drive.

Erasing a JBOD

The **JBOD Erase** appears as an option when you navigate to the **<Select operation>** under **Drive Operations** dialog when the following conditions have been met:

- The controller is in JBOD behavior mode,
- you have created JBODs, and
- the selected drive is JBOD.

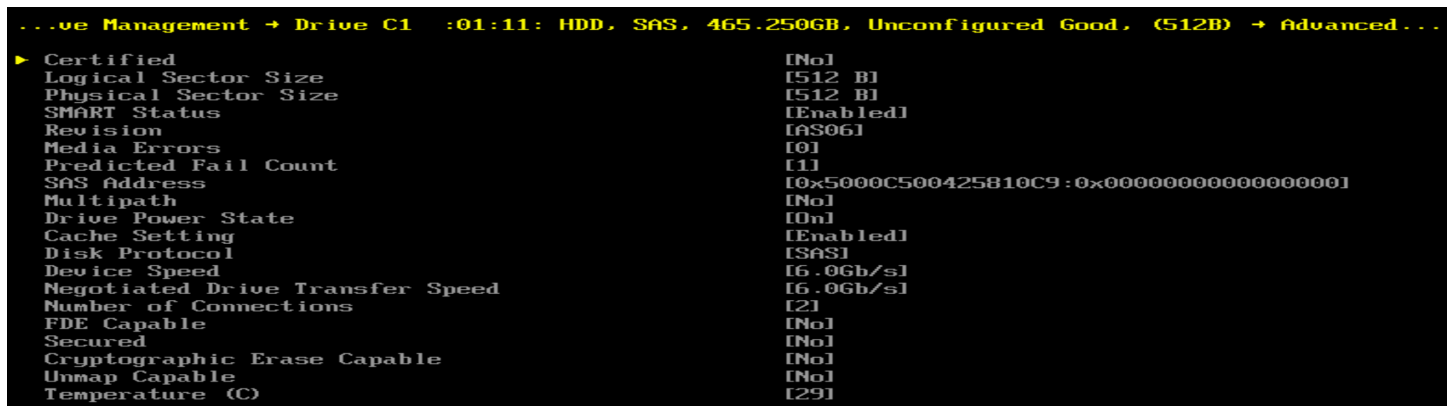
Perform the following steps to erase a JBOD:

1. Open the popup drive operations menu, highlight **JBOD Erase** and press **Enter**.
2. Highlight **Go** and press **Enter** to delete a JBOD.

Viewing Advanced Drive Properties

The following dialog appears when you select **Advanced** on the **Drive Management** menu. The property information in this dialog cannot be modified.

Figure 70: Advanced Drive Properties Dialog



NOTE

A 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 entries listed on the **Advanced Drive Properties** dialog.

Table 34: Advanced Drive Properties


Property	Description
Vendor 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	Indicates the SAS address of the connected drive. If you have connected NVMe drives, this field indicates the World Wide Identifier (WWID) of the connected NVMe drive.

Property	Description
WWID	World Wide Identifier of the connected drive in hexadecimal format.
Emergency Spare	Indicates whether the drive is commissioned as an emergency spare.
Commissioned Hot Spare	Indicates if any hot spare drive (dedicated, global, or emergency) has actually been commissioned.
Cache Setting	Indicates if the drive cache is enabled or disabled.
Available Size (GB)	The available size of the drive, in GB.
Used Space	The configured space of the drive, in GB.
Disk Protocol	Indicates whether the drive uses SAS, SATA, or PCIe 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.
Cryptographic Erase Capable	Indicates if a secured drive is erasable.
Temperature	Indicates whether the drive temperature.
Unmap Capable	Indicates whether the unmap capability is enabled or disabled for the drive.
Unmap Capability for Configuration	Indicates if the unmap capability is enabled or disabled for the current drive configuration.

Managing Hardware Components

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

Figure 71: 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 **Battery Management** or **Enclosure Management** to view more detailed information.

Managing Batteries

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

Figure 72: Battery Management Dialog

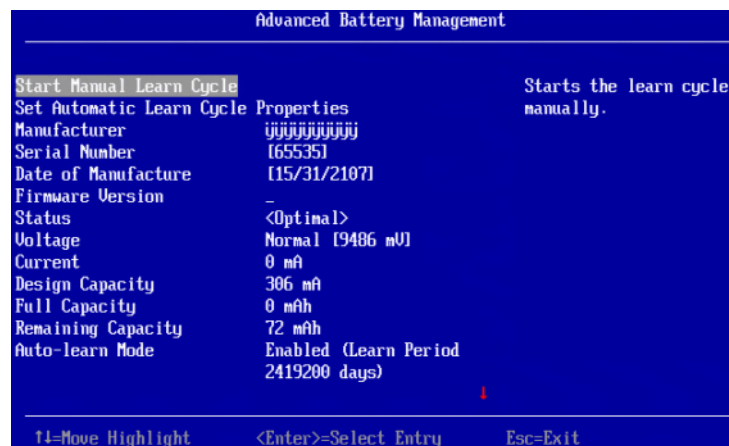


The following table describes the basic battery properties.

Table 35: 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. <ul style="list-style-type: none">• 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 73: Advanced Battery Management Dialog

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 36: 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.

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.

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 74: Enclosure Management Dialog




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 75: View Enclosure Status Dialog



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**.

StorCLI

Overview

The StorCLI tool is a command line interface that is designed to be easy to use, consistent, and easy to script. StorCLI is a unified Storage Command Line Interface, which can be used on **Software RAID (SWR)**, **MegaRAID**, and **Initiator-Target (IT)** Controllers from Gen3 onwards.

NOTE

The legacy commands are deprecated from this guide.

Controllers Supported by the StorCLI Tool

The StorCLI tool works with the MegaRAID and Software RAID (SWR) product lines. The StorCLI tool supports the following controllers:

- MegaRAID 12Gb/s SAS RAID controllers
- Software RAID (SWR) controllers
- Initiator-Target (IT) controllers

Supported Operating Systems

The following table lists the supported operating systems.

Table 37: Supported Operating Systems

Supported Operating Systems	Version/Flavors
Microsoft	Microsoft Windows Client versions <ul style="list-style-type: none"> Windows Client 10 Windows 8.1 Windows 8 Microsoft Windows Server versions <ul style="list-style-type: none"> Windows Server 2019 Windows Server 2016 Windows Server 2012 R2 Windows Server 2012
Linux	Red Hat <ul style="list-style-type: none"> Red Hat Enterprise Linux 7.4 Red Hat Enterprise Linux 7.3 Red Hat Enterprise Linux 7.2 Red Hat Enterprise Linux 7.1 Red Hat Enterprise Linux 7.0 Red Hat Enterprise Linux 6.8 Red Hat Enterprise Linux 6.7 Red Hat Enterprise Linux 6.6 SUSE <ul style="list-style-type: none"> SUSE Linux Enterprise Server 12 SP3 SUSE Linux Enterprise Server 12 SP2 SUSE Linux Enterprise Server 12 SP1 SUSE Linux Enterprise Server 12 SUSE Linux Enterprise Server 11 SP4 SUSE Linux Enterprise Server 11 SP3 SUSE Linux Enterprise Server 11 SP2 PowerPC <ul style="list-style-type: none"> Linux PowerPC for little-endian and big-endian (32 bit and 64 bit)
VMware	<ul style="list-style-type: none"> VMware ESXi 6.5 Update 1 VMware ESXi 6.5 VMware ESXi 6.0 Update 2 VMware ESXi 6.0 Update 1 VMware ESXi 6.0 VMware ESXi 5.5 Update 3 VMware ESXi 5.5 Update 2 VMware ESXi 5.5 Update 1
Citrix XenServer	<ul style="list-style-type: none"> Citrix XenServer 7.2 Citrix XenServer 7 Citrix XenServer 6.5 Citrix XenServer 6.5 SP1
CentOS	<ul style="list-style-type: none"> CentOS-7.4 CentOS-7.2 CentOS-7.1 CentOS-6.8 CentOS-6.7
Fedora	Fedora 24 Workstation

Supported Operating Systems	Version/Flavors
FreeBSD	<ul style="list-style-type: none"> FreeBSD 11 FreeBSD 10.3 FreeBSD 10.2
Ubuntu	<ul style="list-style-type: none"> Ubuntu 14.04 Ubuntu 16.04.1 LTS Ubuntu 16.10
Unified Extensible Firmware Interface	UEFI environment
ARM	Linux, Windows, and UEFI

Installing StorCLI on MegaRAID Controllers

The following topics detail the steps that are required to install the StorCLI tool for MR controllers on various operating systems.

Installing the StorCLI Tool on Microsoft Windows Operating Systems

The Windows StorCLI binary is provided in a binary format, and no separate installation is required.

1. Copy the binary file from the Broadcom website.
2. Place the binary file in the directory from which you want to run StorCLI, and run the tool.

Because Windows PowerShell is not fully supported by the StorCLI tool, use either one of the following techniques to run commands in the StorCLI tool in Windows PowerShell:

- Enclose commands in double quotation marks; for example,

```
storcli "/cx show"
```
- Launch the command prompt from within Windows PowerShell to run the StorCLI commands.

NOTE

The StorCLI tools must be run with the administrator privileges.

Installing the StorCLI Tool on the UEFI Environment

The UEFI StorCLI binary is provided in a binary format, and no separate installation is required.

NOTE

StorCLI default logging requires a minimum of 20 MB of free space.

1. Copy the binary file from the Broadcom website or from the CD provided to you on to a USB drive.
2. Using the USB drive, place the binary file in the directory from which you want to run the Storage Command Line Interface, and run the tool.

After the binaries are copied, you can start executing the StorCLI commands.

Installing the StorCLI Tool on Linux Operating Systems

To install the StorCLI tool on Linux operating systems, perform the following steps:

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

By default, the StorCLI tool will be installed in the `/opt/MegaRAID/storcli` location.

3. To upgrade the StorCLI RPM feature, run the `rpm -Uvh <StorCLI- x.xx-x.noarch.rpm>` command.

Installing the StorCLI Tool on VMware Operating Systems

To install the StorCLI tool on VMware operating systems, run the following from the command line:

```
esxcli software vib install -v=<path-to-vib-package> --no-sig-check
```

NOTE

StorCLI default logging requires a minimum of 20 MB of free space.

Example:

```
esxcli software vib install -v=/vmfs/volumes/datastore1/StorCliMN/vmware-esx-StorCli-1.01.04.vib
```

NOTE

Broadcom provides three variants of StorCLI tool for VMware to be compatible with ESXi versions and MegaRAID (MR) drivers:

VMware-NDS – This package must be used with MegaRAID/IT drivers, `lsi_mr3`, which is a native driver.

The **VMware-NDS** package can be executed with both native and legacy drivers.

StorCLI Tool Command Syntax

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

- In large configurations, running two instances of the StorCLI tool in parallel (at the same time) is not recommended.
- To get the output in JSON format, add `J` at the end of the command syntax. For example:
`storcli /cx show <property1>|<property2> J`
- Background operations are blocked in the UEFI environment, and these operations are resumed in the operating system environment.
- StorCLI discovers MegaRAID controllers in a driver-less UEFI environment only if the firmware is in operational or ready state.
- On ARM-UEFI platforms, StorCLI does not detect the controller if there is no ARM-UEFI driver present.

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 38: 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.
<code>/cx</code>	This object identifier is for controller <code>x</code> .
<code>/call</code>	This object identifier is for sending the command to all controllers.
<code>/c x/v x</code>	This object identifier is for a virtual drive <code>x</code> on controller <code>x</code> .
<code>/c x/vall</code>	This object identifier is for all virtual drives on controller <code>x</code> .
<code>/cx/ex</code>	This object identifier is for an enclosure <code>x</code> on controller <code>x</code> .
<code>/cx/eall</code>	This object identifier is for all enclosures on controller <code>x</code> .

Object Identifier	Description
/cx/fall	This object identifier is for all foreign configurations on controller x.
/cx/lnx	This is the object identifier for the lane speed x on controller x
/cx/ex/sx	This object identifier for the drive is slot x on enclosure x on controller x.
/cx/sx	This object identifier represents the drives that are directly attached to controller x.
/call/sall	This object identifier represents all the drives that are directly attached to all controllers.
/cx/ex/sall	This object identifier is for all the drives on enclosure x on controller x.
/cx/dx	This object identifier is for the drive group x on enclosure x on controller x.
/cx/dall	This object identifier is for the all drive groups on enclosure x on controller x.
/cx/px	This object identifier is for a phy operation x on controller x.
/cx/pall	This object identifier is for all phy operations on controller x.
/cx/bbu	This object identifier is for a BBU on controller x.
/cx/cv	This object identifier is for a cache vault x on controller x.
/cx/mx	This object identifier is a MUX identifier, 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 39: Verbs in the StorCLI Command Syntax

Verb	Description
add	This verb adds virtual drives, JBODs, and so on to the object identifier.
compare	This verb compares an input value with a system value.
del	This verb deletes a drive, value, or property of the object identifier.
download	This verb downloads and flashes a file to the target.
expand	This verb expands the size of the virtual drive.
erase	This verb erases a particular region on the controller, depending on the argument specified.
flush	This verb flushes a controller cache or a drive cache.
flasherase	This verb erases the flash memory on the controller.
get	This verb obtains the data from the controller.
import	This verb imports the foreign configuration into the drive.
insert	This verb replaces the configured drive that is identified as missing, and starts an automatic rebuild.
pause	This verb pauses an ongoing operation.
resume	This verb resumes paused operation.
restart	This verb restarts the controller without a system reboot.
set	This verb sets a value of the object identifier.
show	This verb shows the value and properties of the object identifier.
split	This verb enables you to perform a break mirror operation on a drive group.

Verb	Description
suspend	This verb suspends a particular operation that is being performed.
start	This verb starts an operation.
stop	This verb stops an operation that is in progress. A stopped process cannot be resumed.
spinup	This verb spins up the drives connected to the controller.
spindown	This verb spins down an unconfigured drive and prepares it for removal.
shutdown	This verb shuts down the controller.
secure erase Cryptographic Erase	This verb erases the lock key of a secure drive.
transform	This verb downgrades the firmware memory on the controller.

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

StorCLI Default Logging

Default logging functionality has been enabled in StorCLI. When a default log file is created, the file is saved as `storcli.log`. Each time default logging occurs, the information is added to the `storcli.log`. Once the log file reaches a maximum size of 3 MB, a new log file is created. There can be up to four log files at any given time. For example:

- `storcli.log`
- `storcli.log.1`
- `storcli.log.2`
- `storcli.log.3`

Due to default logging, there is a space limitation in light operating systems such as Vmware or UEFI.

NOTE

StorCLI default logging requires a minimum of 20 MB of free space.

There are two conditions under which StorCLI logging occurs.

- When the `storcliconf.ini` file is present in the same directory as the StorCLI binary.
Logging happens to the file name specified in the `ini` file. This is useful in situations where default logging will not work.
For example, a segmentation fault occurs or a crash happens in StorCLI binary. In these situations, collect a StorCLI log file by placing the `storcliconf.ini` file in the same working directory as StorCLI.
- When the `storcliconf.ini` file is not present in the same directory as the StorCLI binary.
Default logging occurs automatically.

Use the `nolog` option to disable logging for any command.

For example, include the `nolog` option in the `storcli /cx show nolog` command to prevent default logging.

StorCLI Controller 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 in order to ensure proper command execution. Incorrect or duplicate values for variables could result in the last variable being executed or in a command failure.

The version of the StorCLI and the operating system on which StorCLI is being executed are also displayed at the beginning of StorCLI output.

This section describes the commands supported by StorCLI.

System Commands

System Show Commands

StorCLI 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 connected to the system.

storcli show help

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

storcli v

This command shows the version of the StorCLI. The version of the StorCLI and the operating system on which StorCLI is being executed are also displayed at the beginning of StorCLI output.

Controller Commands

Controller commands provide information and perform actions related to a specified controller. The StorCLI Utility supports the controller commands described in this section.

Show and Set Controller Properties Commands

Table 40: Controller Commands Quick Reference Table

Commands	Value Range	Description
<code>show <properties></code>	See Table 41 .	Displays specific controller properties.
<code>set <properties></code>	See Table 41 .	Sets controller properties.
<code>show</code>	<code>all</code> : Shows all properties of the virtual drive. <code>freespace</code> : Shows the free space available in the controller. See Controller Show Commands .	Displays physical drive information.

This section provides command information for the `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:

```
storcli /c0 show bgirate
Controller = 0
Status = Success
Description = None
Controller Properties:
=====
-----
Ctrl_Prop Value
-----
BGI Rate 45%
-----
```

You can show the following properties using the `storcli /cx show <property>` command.

```
storcli /cx show abortconerror
storcli /cx show activityforlocate
storcli /cx show alarm
storcli /cx show aso
storcli /cx show autobgi
storcli /cx show autoconfig
storcli /cx show autorebuild
storcli /cx show badblocks
storcli /cx show backplane
storcli /cx show batterywarning
storcli /cx show bgi
storcli /cx show bios
storcli /cx show bootdrive
storcli /cx show bootwithpinnedcache
storcli /cx show cachebypass
storcli /cx show cacheflushint
```

```
storcli /cx show cc
storcli /cx show ccrate
storcli /cx show coercion
storcli /cx show configureautobalance
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 ds
storcli /cx show eccbucketleakrate
storcli /cx show eccbucketsize
storcli /cx show eghs
storcli /cx show erase
storcli /cx show failpdonsmarterror
storcli /cx show flush|flushcache cachecade
storcli /cx show flushwriteverify
storcli /cx show foreignautoimport
storcli /cx show hddthermalpollinterval
storcli /cx show init
storcli /cx show jbod
storcli /cx show largeiosupport
storcli /cx show loadbalancemode
storcli /cx show limitmaxratesata
storcli /cx show maintainpdfailhistory
storcli /cx show migrate
storcli /cx show migraterate
storcli /cx show ncq
storcli /cx show ocr
storcli /cx show patrolread|pr
storcli /cx show pci
storcli /cx show perfmode
storcli /cx show personality
storcli /cx show pi
storcli /cx show prcorrectunconfiguredareas
storcli /cx show preservedcache
storcli /cx show profile
storcli /cx show prrate
storcli /cx show rebuildrate
storcli /cx show rehostinfo
storcli /cx show restorehotspare
storcli /cx show safeid
storcli /cx show smartpollinterval
storcli /cx show sgpioforce
storcli /cx show spinupdelay
storcli /cx show spinupdrivecount
storcli /cx show ssdthermalpollinterval
storcli /cx show time
storcli /cx show usefdeonlyencrypt
storcli /cx show wbsupport
```

storcli /cx set <property> = <value>

General example output:

```
storcli /c0 set bgirate=40
Controller = 0
Status = Success
Description = None
Controller Properties :
=====
-----
Ctrl_Prop Value
-----
BGI Rate   40%
-----
```

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

NOTE

In the following list and table, setting a property to `on` enables that feature and setting a property to `off` disables that feature.

```
storcli /cx set abortconerror=[on|off]
storcli /cx set activityforlocate=[on|off]
storcli /cx set alarm=[on|off|silence]
storcli /cx set autoconfig=[none|R0 [immediate]|JBOD] > [usecurrent]
storcli /cx set autoconfig [sesgmt=[on|off]] [secured=[on|off]] [multipath=[on|off]]
    [multiinit=[on|off]] [discardpinnedcache=<Val>] [failPDOnReadME=[on|off]]
    [Lowlatency=low|off]]
storcli /cx set backplane [mode=[0|1|2|3]] [expose=[on|off]]
storcli /cx set batterywarning=[on|off]
storcli /cx set bgirate=<value>
storcli /cx set bootwithpinnedcache=[on|off]
storcli /cx set cachebypass=[on|off]
storcli /cx set cacheflushinterval=<value>
storcli /cx set ccraterate=<value>
storcli /cx set coercion=<value>
storcli /cx set consistencycheck|cc=[off|seq|conc] [delay=<value>]
    [starttime=<yyyy/mm/dd hh>] [excludevd=x-y,z]
storcli /cx set copyback=[on|off] type=[smartssd|smarthdd|all]
storcli /cx set dimmerswitch|ds=[on|off] [type=[1|2|4]]
storcli /cx set directpdmapping=[on|off]
storcli /cx set driveactivityled=[on|off]
storcli /cx set eccbucketleakrate=<value>
storcli /cx set eccbucketsize=<value>
storcli /cx set eghs [state=[on|off]] [smarter=[on|off]] [eug=[on|off]]
storcli /cx set failpdonsmarterror=[on|off]
storcli /cx set flushwriteverify=[on|off]
storcli /cx set foreignautoimport=[on|off]
storcli /cx set immediateio=[on|off]
storcli /cx set jbod=[on|off]
storcli /cx set loadbalancemode=<value>
storcli /cx set maintainpdfailhistory=[on|off]
storcli /cx set migraterate=<value>
```

```

storcli /cx set ncq=[on|off]
storcli /cx set patrolread|pr[ [=on mode=[auto|manual]] | [off]]
storcli /cx set perfmode=<value>
storcli /cx set personality=[RAID|HBA|JBOD]
storcli /cx set pi [state=[on|off]] [import=[on|off]]
storcli /cx set prcorrectunconfiguredareas=[on|off]
storcli /cx set profile profileid=<value>
storcli /cx set prrate=<value>
storcli /cx set rebuildrate=<value>
storcli /cx set restorehotspare=[on|off]
storcli /cx set sesmonitoring=[on|off]
storcli /cx set SGPIOforce=[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 termlog[=on|off|offthisboot]
storcli /cx set time=<yyyymmdd hh:mm:ss|systemtime>
storcli /cx set usefdeonlyencrypt=[on|off]

```

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

Table 41: Properties for Show and Set Commands

Cmd	Property Name	Set Command Range	Description
set	abortcconererror	= [on off]	Aborts consistency check when it detects an inconsistency.
set	activityforlocate	= [on off]	If set to <code>on</code> , the LEDs usually used to show drive activity are used for drive locate instead.
set	alarm	= [on off silence]	<p>Enables (<code>on</code>) or disables (<code>off</code>) alarm on critical errors. The option <code>silence</code> silences the alarm.</p> <p>This option is only supported on MegaRAID controllers that have a physical buzzer.</p> <p>If set to <code>on</code>, the alarm sounds on critical errors.</p> <p>If set to <code>off</code>, any active alarms are silenced and the alarm will not sound on any errors, including critical errors.</p> <p>If set to <code>silence</code>, any currently sounding alarm stops, but starts again after a reboot if another critical error is present and the alarm is <code>on</code>.</p>
show	aso		Displays the <code>aso</code> status.
set	aso	deactivatetrialkey	Displays the enabled Advanced Software Options
set	aso	key=<key value>	
set	aso	key=<key value> preview	
set	aso	rehostcomplete	

Cmd	Property Name	Set Command Range	Description
set	aso	transfertovault	
set	assemblynumber	=<xxxx>	
show	autobgi		Displays the autobgi status.
show	autoconfig		Displays the autoconfig status.
set	autoconfig	=[R0 JBOD NONE]	Sets the behavior to R0 , JBOD , or NONE .
show	autorebuild		Displays the autorebuild status.
set	autorebuild	=[on off]	
show	badblocks		Displays bad blocks of the CacheVault™ module.
set	backplane	expose=[on off]	Enables (on) or disables (off) device drivers to expose enclosure devices; for example, expanders, SEPs.
set	backplane	mode=0 : Use autodetect logic of backplanes, such as SGPIO and I ² C SEP using GPIO pins. mode=1 : Disable autodetect SGPIO. mode=2 : Disable I ² C SEP autodetect. mode=3 : Disable both the autodetects.	Configures enclosure detection on a non-SES/ expander backplane.
set	batterywarning	=[on off]	Enables (on) or disables (off) battery warnings.
show	bgi		Displays the bgi (background initialization) status.
set	bgirate	=0 to 100	Sets background initialization rate in percentage.
show	bios		Displays the bios status.
set	bios	[state=[on off]] [Mode=[SOE PE IE SME]] [abs=[on off]] [DeviceExposure=<value>]	
show	bootdrive		Displays the bootdrive status.
set	cachebypass	=[on off]	Enables (on) or disables (off) the cache bypass performance improvement feature.
set	cacheflushint	=0 to 255, Default value: 4	Sets the cache flush interval in seconds.
show	cc		Displays the cc (consistency check) status.
set	ccrate	=0 to 100	Sets the consistency check rate in percentage.
set	coercion	=0 : No coercion =1 : 128 MB =2 : 1 GB	Sets the drive capacity in coercion mode. The 12Gb/s MegaRAID® Tri-Mode Software User Guide (MR-TM-SW-UG107) contains more information about coercion.

Cmd	Property Name	Set Command Range	Description
set	config	file=<fileName>	Saves the controller and the controller properties to the specified file.
set	configautobalance	=[on off]	
show	configautobalance		Displays the configautobalance status.
set	consistencycheck	=[on off]	See Consistency Check .
set	copyback	=[on off] type=[smartssd smarthdd all]	Enables or disables copyback for drive types. The available choices are: smartssd – Copyback enabled for SSD drives. smarthdd – Copyback enabled for HDD drives. all – Copyback enabled for both SSD drives and HDD drives.
set	debug	type=<value> option=<value> [level=<value in hex>] reset all	
set	dimmerswitch ds	See Virtual Drive Power Settings Commands .	See Virtual Drive Power Settings Commands .
set	directpdmapping	=[on off]	When using enclosures, set the directpdmapping property to off to disable it; otherwise, set it to on to enable it.
show	dpm		Displays the dpm status.
set	driveactivityled	=[on off]	Activate or deactivate the Drive Activity LED.
show	ds	=OFF type=[1 2 3 4] =ON type=[1 2] [properties] =ON type=[3 4] DefaultLdType=<val> [properties] [properties]	Displays the Dimmer Switch status.
set	eccbucketleakrate	0 to 65535	Sets the leak rate of the single-bit bucket in minutes (one entry removed per leak-rate).
set	eccbucketsize	0 to 255	Sets the size of ECC single-bit-error bucket (logs event when full).
set	eghs	eug=[on off]	Enables (on) or disables (off) the commissioning of Unconfigured Good drives as Emergency Hot Spare (EHSP) drives.
set	eghs	smarter=[on off]	Enables (on) or disables (off) the commissioning of Emergency Hot Spare (EHSP) drives for Predictive Failure (PFA) events.
set	eghs	state=[on off]	Enables (on) or disables (off) the commissioning of otherwise incompatible global hot spare drives as Emergency Hot Spare (EHSP) drives.

Cmd	Property Name	Set Command Range	Description
show	erase		Displays the <code>erase</code> status.
set	factory defaults		
set	failpdonsmarterror	= <code>[on off]</code>	Enables (<code>on</code>) or disables (<code>off</code>) the <i>Fail PD on SMARTer</i> property.
show	flush flushcache	cachecade	Displays the CacheCade [®] flush status.
set	flushwriteverify	= <code>[on off]</code>	Enables (<code>on</code>) or disables (<code>off</code>) the Write Verify feature. This feature verifies if the data was written correctly to the cache before flushing the controller cache.
set	flushwriteverify	= <code>[on off]</code>	
show	foreignautoimport		On – Displays the <code>foreignautoimport</code> status on startup. Off – Does <i>not</i> display the <code>foreignautoimport</code> status automatically on startup.
set	foreignautoimport	= <code>[on off]</code>	Imports a foreign configuration automatically, at boot.
show	hddthermalpollinterval		Displays the <code>hddthermalpollinterval</code> status.
set	HDDThermalPollInterval	= <code><value></code>	
set	immediateio	= <code>[on off]</code>	Enables (<code>on</code>) or disables (<code>off</code>) immediate I/O transactions.
show	init		Displays the <code>init</code> status.
set	jbod	= <code>[on off]</code>	Enables (<code>on</code>) or disables (<code>off</code>) JBOD mode; by default, drives become system drives. This property is not supported by all controllers. Note: If you try to disable the JBOD mode, and if any of the JBODs have an operating system/file system, 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 <code>force</code> option to proceed with the disable operation.
set	jbodwritecache	= <code>[on off default]</code>	
show	jbodwritecache		Displays the <code>jbodwritecache</code> status.
set	largeiosupport	= <code>[on off]</code>	
show	largeiosupport		Displays the <code>largeiosupport</code> status.
set	largeQD	= <code>[on off]</code>	
show	largeQD		Displays the <code>largeQD</code> status.
set	ldlimit	= <code>[default max]</code>	
show	ldlimit		Displays the <code>ldlimit</code> status.
set	limitMaxRateSATA	= <code>[on off]</code>	
show	limitmaxratesata		Displays the <code>limitmaxratesata</code> status.

Cmd	Property Name	Set Command Range	Description
set	linkconfig	[connname=cx,cy] linkconfig=<val>	
set	loadbalancemode	= [on off]	Enables (on) or disables (off) automatic load balancing between SAS phys or ports in a wide port configuration.
show	maintainpdfailhistory		Displays the maintainpdfailhistory status.
set	maintainpdfailhistory	= [on off]	On (enabled) – The controller keeps track of drives that have previously failed on the controller. If a previously failed drive is reattached to the controller, it is displayed as a foreign configuration and will be need to be imported before it can be used. Off (disabled) – The controller does <i>not</i> track failed drives.
set	maintenance	mode=[normal nodevices]	
show	migrate		Displays the migrate status.
set	migraterate	=0 to 100	Sets the VD configuration migration rate in percentage.
show	ocr		Displays the online controller reset (ocr) status.
set	ocr	= [on off]	
set	patrolread	[starttime=<yyyy/mm/dd hh>] pr	
set	patrolread	= [[on mode=[auto manual]] [off]] pr	
set	patrolread	delay = <value>	
set	patrolread pr	See Patrol Read .	See Patrol Read .
show	pci		Displays the pci status.
set	pdfaileventoptions	[detectionType=<val>] [correctiveaction=<val>] [errorThreshold=<val>]	
set	perfmode	=<value> [maxflushlines=<value> numiostorder=<value>]	
set	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.
show	personality		Displays the current, supported, and requested personalities. It also displays the current behavior and respective behavior parameters.
set	personality	= [RAID HBA JBOD]	Sets the personality to RAID , JBOD , or HBA . If you switch personalities, you must reboot the system for the changes to take effect.

Cmd	Property Name	Set Command Range	Description
set	pi	= [on off]	Enables (on) or disables (off) Protection Information (sometimes called data protection) support on the controller.
set	pi	import= [on off]	Enables (on) or disables (off) import data protection drives on the controller.
set	prcorrectunconfiguredareas	= [on off]	Correct media errors during patrol read by writing 0s to unconfigured areas of the disk.
show	preservedcache		Displays the preservedcache status.
set	profile	profileid=<id>	Sets the profile ID. Valid profile values (set using profileid=option) differ for each MR controller model and firmware version; consult the applicable MegaRAID controller user guide and release notes document for supported profileid values.
set	prrate	=0 to 100	Sets the patrol read rate of the virtual drives in percentage.
set	rebuildrate	=0 to 100	Sets the rebuild rate of the drive in percentage.
set	reconrate	=0 to 100	Sets the reconstruction rate for a virtual drive, as a percentage.
set	restorehotspare	= [on off]	Becomes a hot spare on insertion of a failed drive.
set	sasadd	= <xxxx> [devicename] [methodport]	
set	sasaddhi	= <xxxx> [devicename] [methodport]	
set	sasaddlow	= xxxxx [devicename] [methodport]	
set	securitykey	=<xxxxxxxx> [passphrase=<xxxx>] [keyid=<xxx>]	
set	securitykey	keyid=<xxx>	
set	sesmonitoring	= [on off]	Enables (on) or disables (off) SES monitoring.
show	sgpioforce		Displays the sgpioforce status.
set	SGPIOforce	= [on off]	Forces the SGPIO status per port only for four drives; affects high performance computing (HPC) controllers.
set	smartpollinterval	=0 to 65535	Sets the time for polling of SMART errors, in seconds.

Cmd	Property Name	Set Command Range	Description
show	spinupdelay		Displays the <code>spinupdelay</code> status.
set	spinupdelay	=0 to 255	Sets the spin-up delay between a group of drives or a set of drives, in seconds.
set	spinupdrivecount	=0 to 255	Sets the number of drives that are spun up at a time.
show	ssdthermalpollinterval		Displays the <code>ssdthermalpollinterval</code> status.
set	SSDThermalPollInterval	=<value>	
set	stoponerror	=[on off]	Stops the MegaRAID BIOS during POST, if any errors are encountered.
set	supportssdpatrolread	=[on off]	Enables (on) or disables (off) patrol read for SSD drives.
set	supportssdpatrolread	=[on off] pr	
set	termlog	=[on off offthisboot]	Enables (on) or disables (off) the <code>termlog</code> to be flushed from DDR to ONFI (Open NAND Flash Interface). <code>offthisboot</code> – Disables the <code>termlog</code> flushes to ONFI only for this boot. In the next boot, the <code>termlog</code> is enabled.
set	time	Valid time in <yymmdd hh:mm:ss> format or <code>systemtime</code>	Sets the controller time to your input value or the system time (local time in 24-hour format).
set	tracernumber	= <xxxx>	
set	updatevpd	file= <filepath>	
set	usefdeonlyencrypt	=[on off]	Enables (on) or disables (off) FDE drive-based encryption.

Controller Show Commands

StorCLI supports the following `show` commands:

```
storcli /cx show
storcli /cx show all [logfile[=filename]]
```

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 `cx 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 the 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.

Ensure that the filename does not contain a blank space.

Input examples:

```
storcli /c0 show all logfile=log.txt
storcli /c0 show all logfile = abc.txt
```

NOTE

The PCI information displayed as part of the `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.

Controller Debug Commands

The Storage Command Line Tool supports the following debug commands. There should be at least 20MB of free space for StorCLI to perform debug logging.

Syntax

```
storcli /cx 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 to 128, mapping each number to a particular debug variable in the firmware.
- `option` – Takes the value from 0 to 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 /cx 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` parameters for the preceding `debug` commands are exclusively used by the Broadcom Technical Support Team to provide technical support. For assistance with these debug commands, contact Broadcom Technical Support representative.

Controller Background Task Operation Commands

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 /c1 show rebuildrate
```

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 online physical drives of the controller.

The detailed description for each command follows.

storcli /cx resume patrolread

This command resumes a suspended patrol read operation.

Input example:

```
storcli /c0 resume patrolread
```

storcli /cx set patrolread=[on [mode=[auto|manual]]][off]

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

Input example:

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

storcli /cx set patrolread [starttime=<yyyy/mm/dd hh>] [maxconcurrentpd=<value>] [includessds=[on|off]] [uncfgareas=o[n|off]]

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

Table 42: Set Patrol Read Input Options

Option	Value Range	Description
starttime	A valid date and hour in 24-hour format	Sets the start time in <i>yyyy/mm/dd hh</i> format.
maxconcurrentpd	Valid number of physical drives present	Sets the number of physical drives that can perform patrol read at a single time.
includessds	on or off	Includes SSDs in the patrol read operation.
uncfgareas	on or off	Includes the areas not configured in the patrol read process.

NOTE

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

Input example:

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

storcli /cx set patrolread [delay=<value>]

This command delays the scheduled patrol read in hours.

Input example:

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

storcli /cx show patrolRead

This command shows the 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 that 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 /cx start patrolread

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

Input example:

```
storcli /c0 start patrolread
```

storcli /cx stop patrolread

This command stops a running patrol read operation.

Input example:

```
storcli /c0 stop patrolread
```

NOTE

You cannot resume a stopped patrol read operation.

storcli /cx 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.

Consistency Check

The StorCLI Utility supports the following commands to schedule, perform, and view the status of a consistency check (CC) operation:

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

The detailed description for each command follows.

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

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

Table 43: 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. The concurrent mode slows I/O processing.
delay	-1 and any integer value.	Delays a scheduled consistency check. The value is in hours. A value of 0 makes the CC runs continuously with no delay (in a loop). Only scheduled consistency checks can be delayed.
starttime	A valid date and hour in 24-hour format.	The start time of a consistency check is yyyy/mm/dd hh format.
excludevd	The range should be less than the number of virtual drives.	Excludes virtual drives from the consistency checks. To exclude particular virtual drives, you can provide list of virtual drive names (Vx,Vy ... format) or the range of virtual drives that you want to exclude from a consistency check (Vx-Vy format). If this option is not specified in the command, no virtual drives are excluded.

Input example:

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

storcli /cx show cc

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

Input example:

```
storcli /c0 show cc
```

storcli /cx show ccrate

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

Input example:

```
storcli /c0 show ccrate
```

NOTE

A high CC rate slows I/O processing.

Profile Management

On controllers that support profile management, the StorCLI utility supports the following profile management commands:

```
storcli /cx show profile
```

```
storcli /cx set profile profileid=<value>
```

The detailed description for each command follows:

storcli /cx show profile

This command displays the profiles supported by the controller.

NOTE

Only Broadcom Tri-Mode MegaRAID controllers support profiles.

Input example:

```
storcli /c0 show profile
```

Output example:

On successful execution of the command, the output will have the following fields:

- **Mode**
The mode supported by the current controller profile (HBA, RAID, JBOD).
- **ProfileID**
Displays the current profile ID.
- **MaxPhyDrv**
Displays the maximum number of physical drives supported.
- **MaxLD**
Displays the maximum number of logical drives supported.
- **MaxNVMeDev**
Displays the maximum number of NVMe drives supported.
- **MaxAHCIDev**
Displays the maximum number of AHCI devices supported.
- **isDefault**
Displays if the displayed profile ID is the same as the default profile ID.
- **isCurrent**
Displays if the displayed profile ID is the same as the current profile ID

storcli /cx set profile profileid= <value>

This command sets the specified profile ID of the controller. You need to specify the profile ID in decimal format. For the Profile ID to change, a system reboot is required.

Input example:

```
storcli /c0 set profile profileid=11
```

NOTE

The maximum number and type of PDs supported depends on the profile ID that is selected. Check the applicable MegaRAID user guide and firmware release notes for these values.

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 show aso
storcli /cx set aso [transfertovault][rehostcomplete][deactivatetrialtkey]
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 44: Set Advanced Software Options Input Options

Option	Value Range	Description
key	40 alphanumeric characters.	The key to activate the ASO on the controller. After they are activated, ASOs cannot be removed from the controller.
deactivatetrialkey	—	Deactivates any currently active trial key.
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
```

Controller Security Commands

The StorCLI Utility supports the following controller security commands:

```
storcli /cx compare securitykey < =xxxxxxxx | file=filename >
storcli /cx delete securitykey
storcli /cx set securitykey < keyid=xxx | file=filename >
storcli /cx set securitykey < =xxxxxxxx {passphrase=xxxx} {keyid=xxx} | file=filename >
storcli /cx set securitykey < =xxxxxxxx oldsecuritykey=xxxxxxxx {passphrase=xxxx} [keyid=xxx] | file=filename>
storcli /cx [/ex]/sx set 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 < =xxxxxxxx | file=filename >

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=xxx | file=filename >

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

storcli /cx set securitykey < =xxxxxxxx {passphrase=xxxx} {keyid=xxx} | file=filename >

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

Table 45: Set Security Key Input Options

Option	Value Range	Description
passphrase	Should have a combination of numbers, uppercase letters, lowercase letters, and special characters. Minimum of 8 characters and maximum of 32 characters.	A string that is linked to the controller and is used in the next bootup to encrypt the lock key. If <code>passphrase</code> is not set, the controller generates it by default.
keyid	—	The 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 < =xxxxxxxx oldsecuritykey=xxxxxxxx {passphrase=xxxx} [keyid=xxx } | file=filename>

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

This command sets the security on the FD-capable JBOD drive.

Input example:

```
storcli /c0/e0/s0 set security=on
```

Flashing Controller Firmware Command while the Firmware Is Operational

NOTE

The Flashing Controller Firmware Command while the Firmware Is Operational is not supported in Embedded MegaRAID.

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

This command flashes the firmware with the ROM file to the specified adapter from the given file location (<filepath> is the absolute file path).

You can use the following options in the table to flash the firmware.

Table 46: Flashing Controller Firmware Input Options

Option	Value Range	Description
file	filepath	The absolute file path.
nosigchk	—	The application flashes the firmware even if the check word on the file does not match the required check word for the controller. 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. This option must be provided between phases to the downgrade firmware.
fwtype	0 : Application 1 : TMMC 2 : GG-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. The <code>resetnow</code> option is not supported in the UEFI mode.
[forceclose]	—	Used for calling the flash close before start of the FW download process.

Input example:

```
storcli /c1 download file=c:\app.rom fwtype=0
```

Controller Cache Command

The following command flushes the controller cache.

`storcli /cx flush|flushcache`

This command flushes the controller cache.

Input example:

```
storcli /c0 flushcache
```

Controller Configuration Commands

The following command works with the controller configuration.

`storcli /cx set config file=<file_name>`

This command saves the controller configuration and its properties to the specified file.

NOTE

You cannot load a saved configuration over an existing configuration when there are existing virtual drives. To load a saved configuration, you must first clear the existing configuration on the target controller.

Input example:

```
storcli /c0 set config file= log.txt
```

storcli /cx get config file=<file_name>

This command obtains the controller configuration and its properties from the specified file.

Input example:

```
storcli /c0 get config file= log.txt
```

SnapDump Commands

Snapshot dumping is a mechanism of saving a snapshot of the debug information at fault time. The intention is to collect all required information to be able to root cause the defect at the first instance of defect detection. The SnapDump command makes sure that multiple defect reproductions are not required to debug.

Windows Driver RTTrace

Options to collect the RTTrace files include:

- SnapDump zip filename – `storcli /c0 get snapdump id=<snapDumpId> file=<filename>.zip`
- To retrieve the RTT file directly from the driver – `storcli get rtt dump .`

SnapDump Support Commands

The StorCLI utility checks the SnapDump support bit and allows the user to `show` or `set` the controller on-off property for SnapDump. Using the support command, you can enable or disable the SnapDump feature on the controller using the commands that follow.

```
storcli /cx set snapdump state=on|off
```

```
storcli /cx show snapdump
```

Detailed descriptions of each command follows.

storcli /cx set snapdump state=on|off

This command enables or disables the SnapDump feature on the controller.

Input example:

```
storcli /c0 set snapdump state=on
```

storcli /cx show snapdump

The `show snapdump` command displays whether the SnapDump feature is enabled or not. If enabled, a detailed list of SnapDump properties is displayed, such as the number of dumps that firmware can save and delay for OCR. This command also displays the list of SnapDump files that the firmware currently has, the sizes of the files, and the time the SnapDump is generated, and whether it is an on-demand SnapDump or auto-generated one based on the firmware capability.

Input example:

```
storcli /c0 show snapdump
```

Modifying SnapDump Properties Command

With the SnapDump feature enabled, the user can set the various other properties of the SnapDump, such as the number of dumps that firmware can save and delay for OCR.

storcli /cx set snapdump [savecount=<value> | delayocr=<value>]

Where:

- `savecount` –
 - For the SAS3516 controller, `savecount` sets the number of times the SnapDump will persist in firmware, in case the user does not collect the data.
 - For all other controllers, `savecount` is no longer used and is displayed as N/A.
- `delayocr` – Delays the driver trigger for SnapDump before it initiates OCR; the delay is in seconds.

Input example:

```
storcli /c0 set snapdump [savecount=<value>] [delayocr=<value>]
```

Retrieving SnapDump Data Commands

The StorCLI utility supports the SnapDump commands that follow.

```
storcli /cx get snapdump ID=<val> file=<filename>
storcli /cx get snapdump ID=all
storcli /cx get snapdump
```

Detailed descriptions for each command follows.

storcli /cx get snapdump ID=<val> file=<filename>

To download a specific SnapDump ID, you must read the ID from the firmware. The StorCLI utility keeps writing the data to the file, truncating the file and adding new information.

Input example:

```
storcli /c0 get snapdump ID=<val> file=<filename>
```

Where:

- `val` – Specifies the SnapDump ID number.
- `filename` – Specifies the file, in zip format, in which to write the SnapDump data

storcli /cx get snapdump ID=all

To download all SnapDump IDs that are present on the controller, use the all option.

With this command, the filename is framed by the CLI in a specific format as shown:

```
snapdump_c#(controllerid)_id#(snapdump_id)_year_month_day_hour_min_sec.zip
```

Input example:

```
storcli /c0 get snapdump ID=all
```

storcli /cx get snapdump

To generate and download all SnapDump data when the user has not provided ID, an on-Demand request to the controller is generated and downloads all the SnapDump data present on the controller. With this command, the filename is framed by the CLI in a specific format as shown:

```
.snapdump_c#(controllerid)_id#(snapdump_id)_year_month_day_hour_min_sec.zip
```


Input example:

```
storcli /c0 get snapdump
```

Clearing SnapDump Data Commands

The StorCLI utility has an option to delete all the SnapDump data from the firmware.

NOTE

It is recommended that you save all previous snapdumps, as personality changes and flashing a new firmware package will discard all snapdumps on both DDR and flash.

```
storcli /cx delete snapdump [force]
```

A detailed description for this command follows.

storcli /cx delete snapdump [force]

To clear the SnapDump data from the firmware, use this command application to request the firmware to clear/delete the SnapDump data. If the `force` option is not specified, the StorCLI tool warns the user that this command will clear the SnapDump data and prompt the user to use the `force` option. When the `force` option is specified, the CLI requests the firmware to clear all the SnapDump data.

Input example:

```
storcli /c0 delete snapdump [force]
```

SPDM Commands

StorCLI SPDM commands display the security protocol details and allow users to configure the security protocol on a controller. The SPDM commands allow users to view the security protocol version, slot status, export and import security protocol, and invalidate a slot.

```
storcli /cx show
storcli /cx show all
storcli /cx show security spdm slotgroup=xx slot=yy
storcli /cx export security spdm slotgroup=0 slot=1 subject=subjectname file=filename
storcli /cx import security spdm slotgroup=xx slot=yy file=filename [seal]
storcli /cx set security spdm slotgroup=xx slot=yy invalidate
storcli /cx get security spdm slotgroup=xx slot=yy file=filename
```

storcli /cx show

This command displays the security protocol support and security protocol properties detailed information.

storcli /cx show all

This command displays the security protocol information.

storcli /cx show security spdm slotgroup=xx slot=yy

This command reports the status of the certificate slot chain.

storcli /cx export security spdm slotgroup=0 slot=x1 subject=subjectname file=filename

This command requests the firmware create a certificate signing request and return it. The firmware returns an error if the requested slot is already populated and sealed, if the slot group is invalid, or if the firmware cannot support the requested BaseAsymAlgo and BaseHashAlgo fields selected.

storcli /cx import security spdm slotgroup=xx slot=yy file=filename [seal]

This command supplies a certificate chain from the application to the firmware. The firmware returns an error if the requested slot and slot group fields do not match an open session.

storcli /cx set security spdm slotgroup=xx slot=yy invalidate

This command invalidates the certificate chain storage slot.

storcli /cx get security spdm slotgroup=xx slot=yy file=filename

This command reads the certificate from the chain storage slot and allows the users to validate a downloaded certificate chain.

Drive Commands

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

Table 47: 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.

Drive Show Commands

The StorCLI Utility 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 the slot ID.

The detailed description for each command follows.

storcli /cx[/ex]/sx show

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

Input example:

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

storcli /cx[/eall]/sall show

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

Input example:

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

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

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

This command also shows the NCQ (Native Command Queuing) status (**Enabled**, **Disabled**, or **N/A**), which is applicable only to SATA drives. If the controller to which the SATA drive is connected supports NCQ, and NCQ is enabled on the SATA drive, the status is shown as **Enabled**; 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/e25/s0-3 show all
storcli /c0/e25/sall show all
```

NOTE

The `storcli /cx/sx show all` command shows drive information for the drive specified.

storcli /cx[/ex]/sx show smart

This command displays the SMART information of a SATA drive.

Input example:

```
storcli /c0/e25/s4 show smart
```

Missing Drives Commands

The StorCLI Utility supports the following commands to mark and replace missing physical drives with the specified Unconfigured Good drive:

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

The detailed description for each command follows.

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[/ex]/sx set offline

This command marks the drive in an array as offline.

Input example:

```
storcli /c0/s4 set offline
```

storcli /cx/dall

This command finds the missing drives.

Set Drive State Commands

The StorCLI 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 /c0/e25/s4 set jbod
```

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

This command changes the drive state to unconfigured good.

Input example:

```
storcli /c0/e25/s4 set good
```

NOTE

If the drive has an operating system or a file system on it, the StorCLI Utility 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 /c0/e25/s4 set good [force]
```

storcli /cx[/ex]/sx set offline

This command changes the drive state to offline.

Input example:

```
storcli /c0/e25/s4 set offline
```

storcli /cx[/ex]/sx set online

This command changes the drive state to online.

Input example:

```
storcli /c0/e25/s4 set online
```

storcli /cx[/ex]/sx set missing

This command marks a drive as missing.

Input example:

```
storcli /c0/e25/s4 set missing
```

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

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

Input example:

```
storcli /c0/e25/s4 set bootdrive=on
```

Drive Initialization Commands

When you initialize drives, all the data from the drives is cleared. The StorCLI Utility 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/e25/s4 show initialization
```

storcli /cx[/ex]/sx start initialization

This command starts the initialization process on a drive.

Input example:

```
storcli /c0/e25/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/e25/s4 stop initialization
```

NVMe Drive Commands

The StorCLI Utility supports the following NVMe drive commands.

```
storcli/cx show failedNvmeDevices
storcli/cx[/ex]/sx show repair
storcli/cx[/ex]/sx start repair [force]
storcli/cx[/ex]/sx stop repair
```

Show Failed NVMe Devices

This command displays the list of initialization failed NVMe drives.

Syntax

```
storcli/cx show failedNvmeDevices
```

Show NVMe Drive Repair Status

This command displays the NVMe drive repair status.

Syntax

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

Start NVMe Drive Repair

This command starts the repair process on the requested drive.

Syntax

```
storcli/cx[/ex]/sx start repair [force]
```

Force – Deletes all data present on the drive.

Stop NVMe Drive Repair

This command stops the repair of the requested NVMe drive.

Syntax

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

Configuration Scenarios

You can use the SAS RAID controllers in three scenarios:

- **Low-end, Internal SATA Configurations**

In these configurations, use the RAID controller as a high-end SATA II-compatible controller that connects up to eight disks. These configurations are mostly for low-end or entry servers. Enclosure management is provided through out-of-band Inter-IC (I²C) bus. Side bands of both types of internal SAS connectors support the SFF-8485 (SGPIO) interface.

- **Midrange Internal SAS Configurations**

These configurations are like the internal SATA configurations but with high-end disks. These configurations are more suitable for low-range to midrange servers.

- **High-end External SAS/SATA Configurations**

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

- **NVMe Configurations**

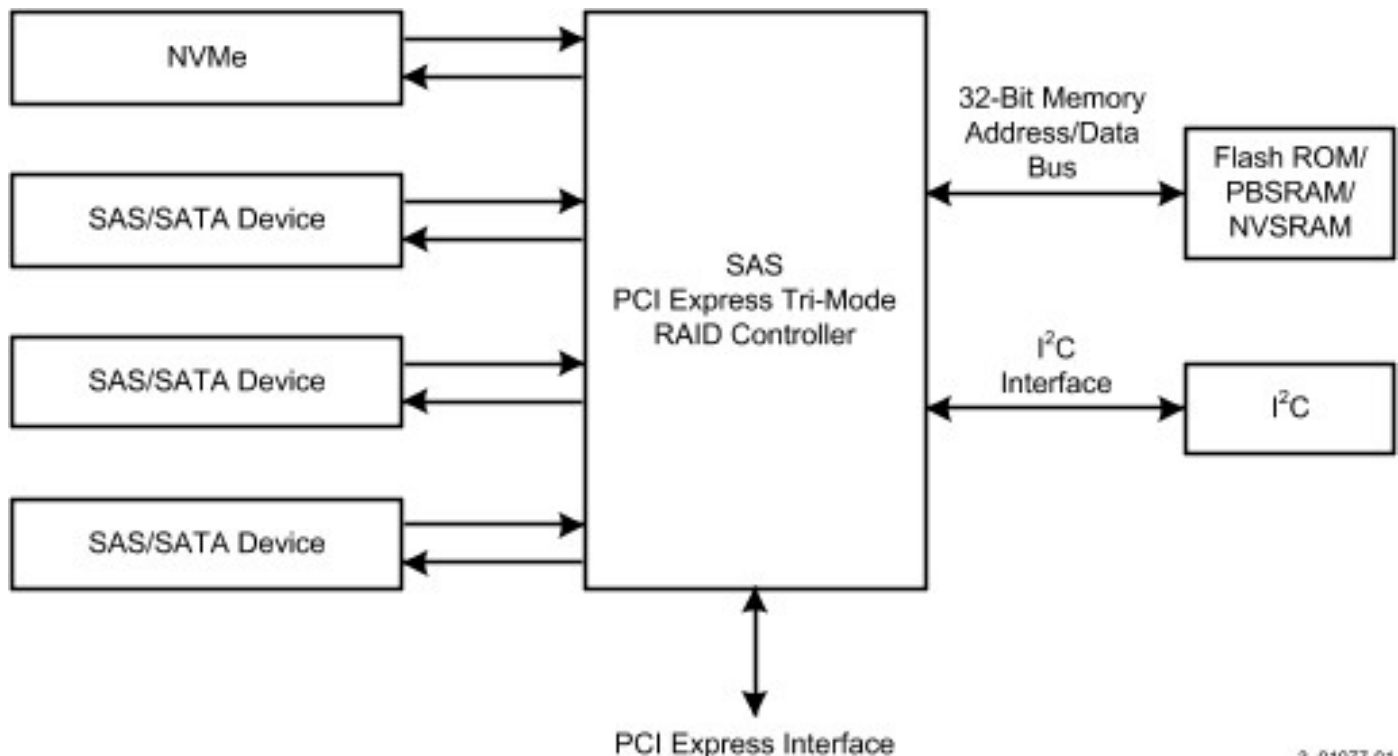
These configurations are for internal connectivity only, using NVMe, either direct connect or switch attached. NVMe configurations are suitable for low latency and high performance environments.

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

NOTE

The external memory bus is 32-bit for the SAS 8704ELP and the SAS 8708ELP, and 64-bit for the SAS 8708EM2, the SAS 8880EM2, and the SAS 8888ELP.

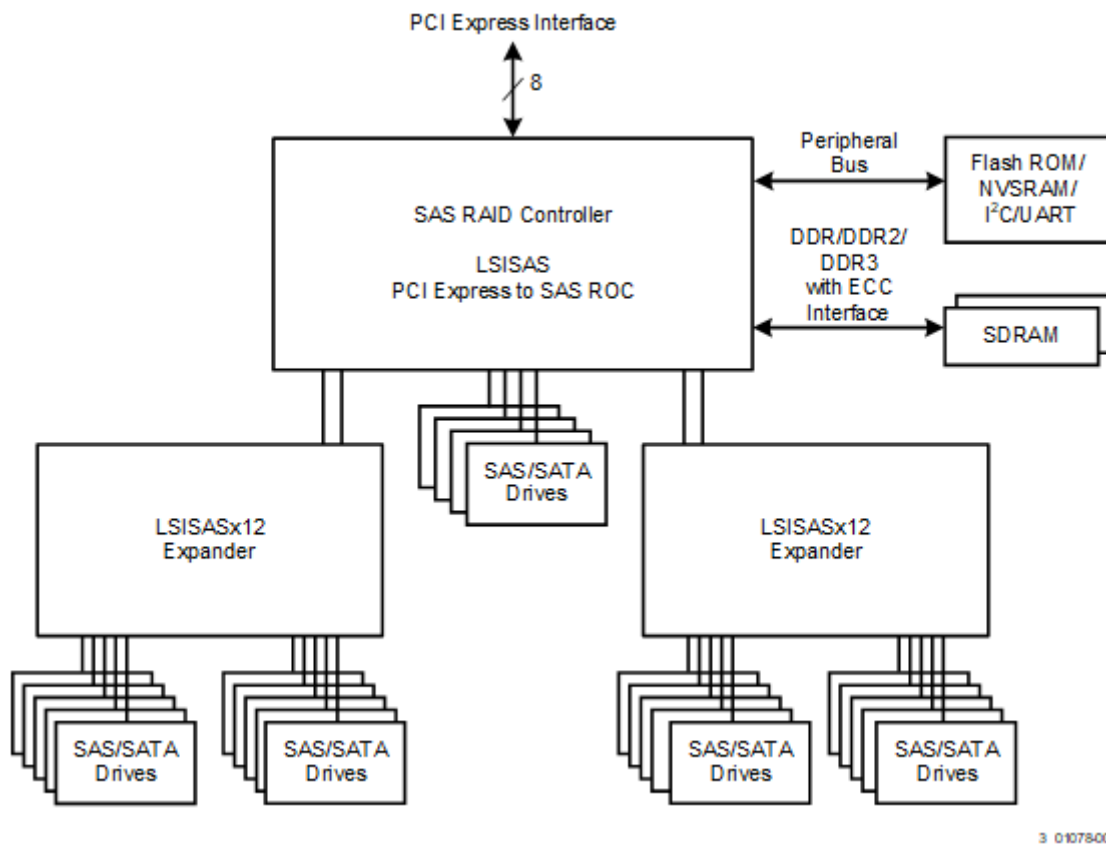
Figure 76: Example of SAS Direct-Connect Application



3_01077-01

The following figure shows an example of a SAS RAID controller configured with an LSI SASx12 expander that is connected to SAS disks, SATA II disks, or both.

Figure 77: Example of SAS RAID Controller Configured with an LSI SASx12 Expander



Drive Firmware Download Commands

The StorCLI Utility 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 32 KB.

NOTE

The default mode is 7 .

Input example:

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

Drive Firmware Update through Parallel HDD Microcode

MegaRAID provides an interface to update the drive firmware in both online and offline modes through host applications, such as StorCLI. Using the parallel HDD microcode update feature, firmware updates can be performed 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 a 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 performed 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.

Locate Drives Commands

The StorCLI Utility 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/e25/s4 start locate
```

storcli /cx[/ex]/sx stop locate

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

Input example:

```
storcli /c0/e25/s4 stop locate
```

Prepare to Remove Drives Commands

The StorCLI Utility 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/e25/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 /c0/e25/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.

Drive Security Command

The StorCLI Utility 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/[e25]/s4 show SecurityKey keyid
```

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

This command sets the security on the FDE-capable JBOD drive.

Input example:

```
storcli /c0/[e25]/s4 set security = on
```

Drive Secure Erase/Cryptographic Erase Commands

The StorCLI Utility 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|crypto|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|crypto|standard] [patternA=<va/1>] [patternB=<va/2>]**NOTE**

The erase option is supported only on UG drives and is not supported on JBOD drives.

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 48: 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 three times crypto: Performs cryptographic erase for SSD drives	Secures erase type.
patternA	8-bit value	Erases pattern A to overwrite the data.
patternB	8-bit value	Erases pattern B to overwrite the data.

Input example:

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

Drive Sanitize Command

The StorCLI Utility supports the following drive sanitize commands:

```
storcli /cx[/ex]/sx start sanitize [cryptoerase| overwrite| blockErase ] [ause]
storcli /cx[/ex]/sx show sanitize
```

storcli /cx[/ex]/sx start sanitize [cryptoerase| overwrite| blockErase] [ause]

This command lets you erase the data that resides on a physical drive. You can use the following sanitize-type arguments with the `start sanitize` command.

cryptoerase – This argument corrupts the encryption keys that may have been present on the drive.

overwrite – This argument overwrites all zeros to the data that may be present on the existing drives.

blockerase – This argument allows the drive to clear or erase the existing data drive.

ause – If, for some reason, the sanitize operation fails, the system tries to bring the drive out of the failure mode irrespective of whether the Allow Unrestricted Sanitize Exit (AUSE) argument is specified or not. However, if this argument is specified, and if the system succeeds in bringing the drive out of the failure mode, the drive is then returned as an Unconfigured Good drive. If you do not specify the `ause` argument, and if the sanitize operation fails, the system places the drive in an Unconfigured Bad state.

Example

```
storcli /c0/e0/s4 start sanitize overwrite
```

storcli /cx[/ex]/sx show sanitize

This command displays the progress of the sanitize operation in percentage.

example:

```
storcli /c0/e0/s4 show sanitize
```

Rebuild Drives Commands

The following commands rebuild drives in the StorCLI Utility:

```
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 rebuilding.

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/s4 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
```

Drive Copyback Commands

The StorCLI Utility supports the following commands for drive copyback:

```
storcli /cx[/ex]/sx pause copyback
storcli /cx[/ex]/sx resume copyback
storcli /cx show copyback
storcli /cx[/ex]/sx start copyback target=eid:sid
storcli /cx[/ex]/sx stop copyback
storcli /cx set copyback=on type=ctrl
storcli /cx set copyback=on type=smarthdd
storcli /cx set copyback=on type=ssd
storcli /cx set copyback=on type=all
```

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 a copyback operation is 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 a paused copyback process exists for the drive.

Input example:

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

storcli /cx 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.

storcli /cx set copyback=on type=ctrl

This command sets a control copyback operation.

Input example:

```
storcli /c0/e25/s4 set copyback type=ctrl
```

storcli /cx set copyback=on type=smarthdd

This command sets a smart HDD copyback operation.

Input example:

```
storcli /c0/e25/s4 show copyback type=smarthdd
```

storcli /cx set copyback=on type=smartssd

This command sets a smart SSD copyback operation.

Input example:

```
storcli /c0/e25/s4 show copyback type=smartssd
```

storcli /cx set copyback=on type=all

This command sets a copyback operation.

Input example:

```
storcli /c0/e25/s4 show copyback type=all
```

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 connected to the physical drives of 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 49: 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. 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/e25/s4,5 add hotsparedrive
```

This command sets the drives /c0/e25/s4,5 as global hot spare.

Input example:

```
storcli /c0/e25/s6,8 add hotsparedrive dgs=0,1
```

This command sets /c0/e25/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/e25/s4,5 delete hotsparedrive
```

Drive Performance Monitoring Commands

The StorCLI Utility supports the following commands for drive performance monitoring:

```
storcli /cx show pdfailevents [lastoneday] [lastseqnum=<val>] [file=filename]
storcli /cx set pdfaileventoptions detectiontype=val correctiveaction=val errorthreshold=val
```

The detailed description for each command follows.

storcli / cx show pdfailevents

This command shows all of the drive predictive failure events.

Input example:

```
storcli /c0 show pdfailevents
```


storcli / cx show pdfailevents lastoneday

This command shows all of the drive predictive failure events that occurred in the last 24 hours.

Input example:

```
storcli /c0 show pdfailevents lastoneday
```

storcli / cx show pdfailevents lastseqnum=xx]

This command shows all of the drive predictive failure events generated from the specified sequence number.

Input example:

```
storcli /c0 show pdfailevents lastseqnum=10
```

storcli / cx set pdfaileventoptions detectiontype=val correctiveaction=val errorrrthreshold=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:

- 0 – Detection disabled
- 1 – Detection enabled, high latency for reads is OK
- 2 – Detection enabled, aggressive (high latency for reads is not OK)
- 3 – 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 errorrrthreshold=x
```

Where:

- 0 = 1 – One error every 8 hours (least tolerant)
- 1 = 8 – One error every 1 hour
- 2 = 32 – One error every 15 minutes
- 3 = 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

storcli / cx show dpm

This command shows the drive performance monitoring status on the controller.

Input example:

```
storcli /c0 show dpm
```

storcli / cx start dpm

This command starts the performance monitoring on the controller.

Input example:

```
storcli /c0 start dpm
```

storcli / cx stop dpm

This command stops the performance monitoring on the controller.

Input example:

```
storcli /c0 stop dpm
```

storcli /cx delete dpmstat type = HIST | LCT | RA | EXT | All

This command deletes the drive performance monitoring statistics of the controller.

Where:

- HIST – Histogram of response time.
- LCT – Long Time commands.
- RA – Running average drive statistics.
- EXT – Extended DPM information.

storcli /cx[/ex]/sx show dpmstat type = HIST | LCT | RA | EXT [logfile[=filename]]

This command shows the drive performance monitoring statistics of the mentioned drive.

Where:

- HIST – Histogram of response time.
- LCT – Long Time commands.
- RA – Running average drive statistics.
- EXT – Extended DPM information.
- logfile – If the file name is not specified, it logs to `tostarsas.log` or to a user-specified file.

Virtual Drive Commands

The StorCLI Utility supports the following virtual drive commands. The following table describes frequently used virtual drive commands.

Table 50: Virtual Drives Commands Quick Reference Table

Commands	Value Range	Description
add	See Table 51, Add RAID Configuration Input Options and Table 52, Add RAID Configuration Input Options .	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 Table 51, Add RAID Configuration Input Options , Table 52, Add RAID Configuration Input Options , and Change Virtual Properties Commands .	Sets virtual drive properties.
show	all : Shows all properties of the virtual drive. cc : Shows properties of CacheCade virtual drives. See Virtual Drive Show Commands .	Shows virtual drive information.

Add Virtual Drives Commands

The StorCLI Utility 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>,...|remaining] [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][Cbsize=0|1|2 Cbmode=0|1|2|3|4|7]
```

NOTE

The supported strip size can vary from a minimum of 64 KB to 1 MB for MegaRAID controllers and only 64 KB for Integrated MegaRAID 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 MegaRAID controllers and only 64 KB for Integrated MegaRAID controllers.

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

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

NOTE

* indicates default values in the following commands.

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>,...|*remaining] [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 MegaRAID controllers and only 64 KB for Integrated MegaRAID controllers.

Table 51: 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 in length.	Specifies the drive name for each virtual drive.
drives	Valid enclosure number and valid slot numbers for the enclosure.	<p>In <i>e:s</i> <i>e:s-x</i> <i>e:s-x,y</i>:</p> <ul style="list-style-type: none"> <i>e</i> specifies the enclosure ID. <i>s</i> represents the slot in the enclosure. <i>e:s-x</i> is the range convention used to represent slots <i>s</i> to <i>x</i> in the enclosure <i>e</i> (250 characters maximum). <p>Make sure that the same block size (in a physical drive) is used in each <i>[e:s]</i> pair. As an example, if you use 4096 bytes in the <i>e0:s0</i> pair, use 4096 bytes in the <i>e1:s1</i> pair too. Mixing of block sizes between the <i>[e:s]</i> pairs is not supported.</p>
pdperarray	1–16.	Specifies the number of physical drives per array. The <i>default</i> value is automatically chosen.
sed	—	Creates security-enabled drives.
pdcache	on off default.	Enables or disables PD cache.
pi	—	Enables protection information.
dimmerswitch	<p><i>default</i>: Logical device uses controller default power-saving policy.</p> <p><i>automatic (auto)</i>: Logical device power savings are managed by firmware.</p> <p><i>none</i>: No power-saving policy.</p> <p><i>maximum (max)</i>: Logical device uses maximum power savings.</p> <p><i>MaximumWithoutCaching (maxnocache)</i>: Logical device does not cache write to maximize power savings.</p>	<p>Specifies the power-saving policy.</p> <p>Sets to <i>default</i> automatically.</p>
direct cached	<p><i>cached</i>: Cached I/O.</p> <p><i>direct</i>: Direct I/O.</p>	<p>Sets the logical drive cache policy.</p> <p>Direct I/O is the default.</p>

Option	Value Range	Description
EmulationType	0 : Default emulation, which means if there are any 512e drives in the configured ID, then the physical bytes per sector are shown as 512e (4k). If there are no 512e drives, the physical bytes per sector will be 512n. 1 : Disable, which means even though there are 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	<ul style="list-style-type: none"> wt : Write through. wb : Write back. awb : Always Write Back. 	Enables write through. Write back is the default.
nora ra	<ul style="list-style-type: none"> 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 MegaRAID controllers and only 64 KB for Integrated MegaRAID 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 the addition of a secured physical drive to an existing drive group without security.	Forces a security-capable physical drive to be added to a drive group without security.

Input example:

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

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

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

Table 52: Add RAID Configuration Input Options

Option	Value Range	Description
cachecade	—	Creates a CacheCade virtual drive.
raid	0, 1, 10.	Sets the RAID type of the CacheCade virtual drive.

Option	Value Range	Description
drives	Valid enclosure number and valid slot number	See the <code>drives</code> row in the previous table for format.
wt *wb awb	wt : Enables write through. wb : Enables write back. awb Enables always write back.	Enables or disables write cache.
assignvds	Valid virtual drive number (0 to 63).	Specifies the list of virtual drives associated with the new CacheCade virtual drives.

Input example:

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

Delete Virtual Drives Commands

The StorCLI Utility 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 unerasd MBR still exists at block 0 of the new virtual drive, which makes it a virtual drive with valid user data. Therefore, you must provide the `force` option to delete this newly created virtual drive.

The detailed description for each command follows.

storcli /cx/vx|vall del

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

Input example:

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

ATTENTION

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

storcli /cx/vx|vall del cachecade

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

Input example:

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

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/vx del [cachecade] [discardcache] [force]

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

Input example:

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

Virtual Drive Show Commands

The StorCLI Utility 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/vx show

This command shows the summary of the virtual drive information.

Input example:

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

storcli /cx/vx 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
```

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 Utility supports the following commands for preserved cache:

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

The detailed description for each command follows.

storcli /cx/vx delete preservedcache

This command deletes the preserved cache for a particular virtual drive on the controller in missing state. Use the `force` option to delete the preserved cache along with the virtual drive when the virtual drive is in an 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
```

Change Virtual Properties Commands

The StorCLI Utility 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 wrccache=<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/vx set accesspolicy=<rw|ro|blocked|rmvblkd>

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

Input example:

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

storcli /cx/vx set iopolicy=<cached|direct>

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

Input example:

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

storcli /cx/vx set name=<namestring>

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

Input example:

```
storcli /c0/v0 set name=testdrive123
```

storcli /cx/vx set pdcache=<on|off|default>

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

Input example:

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

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

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

Input example:

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

storcli /cx/vx|vall set ssdcaching=<on|off>

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

Input example:

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

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

This command sets the host access policy for the virtual drive. When the host access policy is exclusive access, a server has exclusive access to the virtual drive. The virtual drive cannot be shared between servers. If the host policy is shared access, 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 not exported to the OS. The OS will not be able to read or write to that virtual drive until the virtual drive is unhidden.

Input example:

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

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

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

The values for the `cbsize` options follow:

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

The values for the `cbmode` options follow:

- 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 supplied `cbsize` value is ignored.

Input example:

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

Virtual Drive Initialization Commands

The StorCLI Utility 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
```

Virtual Drive Erase Commands

The StorCLI Utility supports the following commands to erase virtual drives:

```
storcli /cx/vx start erase
```

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

The detailed description for each command follows.

storcli /cx/vx start erase

This command starts the data erase operation on the specified virtual drive.

Input example:

```
storcli /c0/v0 start 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
```

Virtual Drive Migration Commands**NOTE**

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

The StorCLI Utility 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]
```

NOTE

Usable Capacity Reduction is not possible without a RAID level migration.

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 53: Virtual Drive Migration Command Options

Options	Value Range	Description
<code>type = RAID level</code>	RAID [0 1 5 6]	The RAID level to which the virtual drive must be migrated.
<code>[option=<add remove> drives=<e1:s1,e2:s2, ...>]</code>	<p><code>add</code>: Adds drives to the virtual drive and starts reconstruction.</p> <p><code>remove</code>: Removes drives from the virtual drive and starts reconstruction.</p> <p><code>drives</code>: The enclosure number and the slot number of the drives to be added to the virtual drive. Make sure that the same block size (in a physical drive) is used in each <code>[e:s]</code> pair. As an example, if you use 4096 bytes in the <code>e0:s0</code> pair, use 4096 bytes in the <code>e1:s1</code> pair too. Mixing of block sizes between the <code>[e:s]</code> pairs is not supported.</p>	Adds or removes drives from the virtual drive.

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

Table 54: 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
```

Virtual Drive Consistency Check Commands

The StorCLI Utility 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.

Background Initialization Commands

The StorCLI Utility 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
```

Virtual Drive Expansion Commands

The StorCLI Utility 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. Although the value provided might 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 can 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
```

Display the Bad Block Table

The StorCLI Utility 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
```

Clear the LDBBM Table Entries

The StorCLI Utility supports the following command to clear the LDBBM table entries:

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

Input example:

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

JBOD Commands

The StorCLI Utility supports the switching behavior within the JBOD personality mode. StorCLI also supports configuration parameters for a personality and allows you to create and configure JBODs. You can create JBODs from all Unconfigured Good drives or specific Unconfigured Good drives. You can also delete these JBODs. You can also choose JBOD as a boot device.

The StorCLI Utility supports the following JBOD commands:

```
storcli /cx/ex/sx set jbod
storcli /cx/eall/sall show jbod [all]
storcli /cx/ex/sx set bootdrive=<on|off>
storcli /cx/ex/sx delete jbod
```

For more information, see also *set personality behavior* under [Table 1](#), .

Create JBOD Manually

The StorCLI Utility has the option to convert all specified Unconfigured Good drives as JBODs.

NOTE

The drive token is optional. If you specify the drives, the JBODs are created on those specified drives, otherwise, StorCLI creates JBODs on all available Unconfigured Good drives on the controller.

storcli /cx/ex/sx set jbod

This command allows you to add JBOD drive

Input example

```
storcli /c0/e1/s1 set jbod
```

JBOD Properties

JBOD properties are used to list all the available JBOD on the controller with their properties.

storcli /cx/eall/sall show jbod [all]

This command lists all the available JBODs on the controller with their associated properties.

Input example:

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

Table 55: Example Output of all the available JBODs on the Controller

ID	EID:SLT	DID	State	Intf	Med	Size	SeSz	Model	Vendor	Port
0	10:01	2	Online	SAS	HDD	100GB	512B	ST91000640SS	SAMSUNG	0-3
1	10:03	5	Online	SAS	HDD	123GB	4K	ST91000640SS	SAMSUNG	0-3
2	10:04	6	Online	SAS	HDD	100GB	512B	ST91000640SS	SAMSUNG	0-3

JBOD Operations

JBODs can start and stop the INIT, and also erase operations on them. JBODs can also be set as a boot volume. The commands for the respective operation to start and stop JBODs follow:

storcli /cx/jbodx show init

This command displays the initialization status

Input example

```
storcli /c0/jbod0 show init
```

storcli /cx/ex/sx start erase [simple| normal| thorough] patternA=<val>][patternB=<val>

This command allows you to securely erase non-SED drives with the specified erase patterns. 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 56: Drive Erase Command Options

Options	Value Range	Description
erase	<ul style="list-style-type: none"> 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/e10/s0 start erase through patternA=10010011 patternB=11110000
```

storcli /cx/ex/sx show erase

This command displays the erase status.

Input example

```
storcli /c0/e10/s0 show erase
```

storcli /cx/ex/sx set bootdrive=<on|off>

This command allows you to set the selected JBOD as boot volume.

Input example

```
storcli /c0/e10/s0 set bootdrive= on|off
```

Delete JBODs or Volumes

To delete JBODs, use the drive's JBOD delete command.

storcli /c x/jbodall delete

This command allows you to delete all the JBODs.

Input example

```
storcli /c0/jbodall delete
```

Delete all volumes on the controller

To delete all volumes on the controller, use the `vall delete` command.

storcli /c x/ vall delete

This command deletes all the volumes on the controller.

Input example

```
storcli /c0/vall delete
```

Foreign Configuration Commands

The StorCLI Utility supports the following commands to view, import, and delete foreign configurations:

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

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|del| delete [securitykey=ssssssssssss]

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 import [preview] [securitykey=ssssssssssss]

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 show [all][securitykey=ssssssssssss]

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
```

BIOS-Related Commands

The StorCLI Utility 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 MegaRAID 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 MegaRAID BIOS to avoid issues during POST.

Input example:

```
storcli /c0 set bios state=on Mode=SOE abs=on deviceexposure=20
```

OPROM BIOS Commands

The StorCLI Utility supports the following OPRM BIOS commands:

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

The detailed description for each command follows.

storcli /cx/ex/sx set bootdrive=on|off

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

Input example:

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

storcli /cx/vx set bootdrive=on|off

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

Input example:

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

storcli/cx/vx show bootdrive

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

Input example:

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

Drive Group Commands

This section describes the drive group commands.

Drive Group Show Commands

The StorCLI Utility supports the following drive group commands:

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

storcli /cx/dall show

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

Input example:

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

storcli /cx/dall show all

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

Input example:

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

storcli /cx/dall show cachecade

This command shows all CacheCade virtual drive information.

Input example:

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

storcli /cx/dx show

This command shows the topology information of the drive group.

Input example:

```
storcli /c0/d0 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/d0 show all
```

storcli /cx/dx set security=on[all]

This command enables security on the specified drive group.

Input example:

```
storcli /c0/d0 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.

NOTE

Break mirror is an offline operation. This command is supported only in UEFI.

Input example:

```
storcli /c0/d0 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=<va>[force]

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

Input example:

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

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

This command hides or unhides a drive group.

Input example:

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

Virtual Drive Power Settings Commands

Change Virtual Drive Power Settings Commands

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

The following table describes the power-saving options.

Table 57: Dimmer Switch Input Options

Option	Value Range	Description
dimmerswitch or ds	on off	Turns the Dimmer Switch option on.
type	1 : Unconfigured 2 : Hot spare 3 : TBD 4 : Unconfigured and hot spare drives	Specifies the type of drives for which the Dimmer Switch feature is applicable. Note: The dimmer switch is only activated for unconfigured drives and hot spare drives, but not for configured drives.
properties	disableldps : Interval in hours or time in <i>hh:mm</i> format spinupdrivecount : Valid enclosure number (0 to 255) SpinUpEncDelay : Valid time in seconds	Sets the interval or time in which the power-saving policy for the logical drive is turned off. Specifies the number of drives in the enclosure that are spun up. Specifies the delay of spin-up groups within an enclosure in seconds.

storcli /cx show DimmerSwitch (ds)

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

Input example:

```
storcli /c0 show ds
```

CacheVault Commands

The StorCLI Utility supports the following CacheVault commands:

```
storcli /cx/cv show
storcli /cx/cv show all
storcli /cx/cv show status
storcli /cx/cv start learn
```

The detailed description for each command follows.

storcli /cx/cv show

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

Input example:

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

storcli /cx/cv show all

This command shows all the information of the CacheVault.

NOTE

This command only works when a CacheVault is connected to the controller; otherwise, an error message appears. A capacitance value above 100% will be displayed as 100%.

Input example:

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

storcli /cx/cv show status

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

Input example:

```
storcli /c0/cv show status
```

storcli /cx/cv start learn

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

Input example:

```
storcli /c0/cv start learn
```

Enclosure Commands

The StorCLI Utility supports the following enclosure commands:

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

The detailed description for each command follows.

NOTE

StorCLI supports and can be used to qualify only Broadcom expanders and enclosures.

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 58: 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). This option is used primarily to activate Scotch Valley enclosures.

NOTE

The firmware file that is used to flash the enclosure can be 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/e25 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
```

PHY Commands

The StorCLI Utility 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 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 1.5Gb/s, 3Gb/s, 6Gb/s, or 12Gb/s. The link speed 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 reset phyerrorcounters

This command resets the drive phy error counters.

Input example:

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

PCIe Storage Interface Commands

The PCIe Storage Interface is the fundamental interface that connects peripheral devices to the host processor and through a memory controller to the memory architecture in the system. The PCIe interface communicates over one or more lanes that consist of one transmit and one receive serial interface for each lane.

Lane Speed Commands

The StorCLI Utility supports the following lane speed commands:

```
storcli /cx/lnx show
storcli /cx/lnall show
storcli /cx/lnx set lanespeed=0(disabled)|2.5|5|8|16
```

The detailed description for each command follows.

storcli /cx/lnx show

This command displays the lane information.

Input example:

```
storcli /c0/ln1 show
```

storcli /cx/lnall show

This command displays the summary information on all of the exiting lanes.

Input example:

```
storcli /c0/lnall show
```

storcli /cx/lnx set lanespeed=0 (disabled) | 2.5 | 5 | 8 | 16

This command sets the lane speed. You can set the speed as 0 (disabled), 2.5GT/s, 5GT/s, 8GT/s, or 16GT/s. By default, the lane speed in the controller is 8GT/s or the value last saved by you.

Input example:

```
storcli /c0/ln1 set lanespeed=2.5
```

Output example:

Figure 78: Lane Speed Output

LaneInformation :|

=====

LaneNo LaneId Enbl Conn Link CurrSpeed Wwid SupSpeed

0	65535	Yes	1	0	8GT/s	0	2.5,5,8
1	65535	Yes	0	0	8GT/s	0	2.5,5,8
2	65535	Yes	1	0	8GT/s	0	2.5,5,8
3	65535	Yes	0	0	8GT/s	0	2.5,5,8
4	65535	Yes	1	1	8GT/s	0	2.5,5,8
5	65535	Yes	0	1	8GT/s	0	2.5,5,8
6	65535	Yes	1	1	8GT/s	0	2.5,5,8
7	65535	Yes	0	1	8GT/s	0	2.5,5,8

Link Configuration Commands

The StorCLI Utility supports the following link configuration commands:

```
storcli /cx/show linkconfig
storcli /cx/set linkconfig [connname=cx,cy] linkconfig=<val>
```

The detailed description for each command follows.

storcli cx/show linkconfig

This command displays the link configuration information for the current link configuration, pending link configuration, and the available link configuration.

Input example:

```
storcli /c1/show linkconfig
```

Output example: Current Link Configuration**Figure 79: Current Link Configuration**

Current Config :

=====

```
-----
Conn ConfigID LinkConfig
-----
```

```
C1,C0      5 0-0:x1,1-1:x1,2-2:x1,3-3:x1
C3,C2      5 8-8:x1,9-9:x1,10-10:x1,11-11:x1
|-----
```

Output example – Pending Link Configuration**Figure 80: Pending Link Configuration**

Pending Config :

=====

```
-----
Conn ConfigID LinkConfig
-----
```

```
C1,C0      3 0-1:x2,2-3:x2,4-7:x4
C3,C2      3 8-9:x2,10-11:x2,12-15:x4
-----
```

Output example: Available Link Configuration

Figure 81: Available Link Configuration

Available Config :

=====

ConfigID LinkConfig

```

1 0-3:x4,4-7:x4
2 0-3:x4,4-5:x2,6-7:x2
3 0-1:x2,2-3:x2,4-7:x4
4 0-1:x2,2-3:x2,4-5:x2,6-7:x2
6 4-4:x1,5-5:x1,6-6:x1,7-7:x1
7 0-1:x2,2-2:x1,3-3:x1
8 4-5:x2,6-6:x1,7-7:x1
9 0-0:x1,1-1:x1,2-3:x2
10 4-4:x1,5-5:x1,6-7:x2
-----
```

storcli /cx set linkconfig [connname=cx,cy] linkconfig=<val>

This command helps you configure the links for different ports of a controller.

Input example:

```
storcli /c1/set linkconfig connname=c0,c1 linkconfig=x4
```

Logging Commands

The StorCLI Utility supports the following commands to generate and maintain log files:

```

storcli /cx delete 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
```

NOTE

The command output for this command cannot be JSON formatted.

storcli /cx show eventloginfo

This command shows the history of log files generated.

Input example:

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

NOTE

The command output for this command cannot be JSON formatted.

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
```

NOTE

The command output for this command cannot be JSON formatted.

storcli /cx show dequeuelog =<filepath>

This command shows the debug log from the firmware.

Input example:

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

NOTE

The command output for this command cannot be JSON formatted.

storcli /cxshow 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
```

NOTE

The command output for this command cannot be JSON formatted.

Automated Physical Drive Configurations

The StorCLI utility supports the following automated physical drive configuration commands:

```
storcli /cx set autoconfig=r0 [immediate]
storcli /cx show autoconfig
storcli /cx set autoconfig=JBOD
storcli /cx set jbod=on|off
storcli /cx set autoconfig [= <none | R0 [immediate] | JBOD> [usecurrent] ] [[sesmgmt=on|off] [secured=on|off] [multipath=on|off] [multiinit=on|off] [discardpinnedcache=Val>] [failPDOnReadME=on|off] [Lowlatency=low|off]]
```

The detailed description for each command follows.

storcli /cx set autoconfig=r0 [immediate]

This command lets you set the controller's automated configuration policy to RAID 0. When set to RAID 0, all unconfigured physical drives are configured as single RAID 0 drives, 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 Unconfigured Good drives. Any physical drives newly connected in the future do not get converted to RAID 0. If you omit the *immediate* option in this command, all attached Unconfigured Good drives will become single drive RAID 0, and the *autoconfig* setting will become R0. Conversion to RAID 0 does not take place on newly connected physical drives, and the *autoconfig* setting will become R0. There is no setting that allows newly added Unconfigured Drives to be automatically converted to RAID 0. Automatic conversion to RAID 0 can be turned off by setting the *autoconfig* policy to *none*, which causes all drives to be converted to Unconfigured Good.

Input examples:

```
storcli /c0 set autoconfig=r0 immediate [[sesmgmt=[on|off]] [secured=[on|off]] [multipath=[on|off]
[multiinit=[on|off]]]

storcli /c0 set autoconfig [= <none | R0 [immediate] | JBOD> [usecurrent] ] [[sesmgmt=on|off] [secured=on|off] [multipath=on|off] [multiinit=on|off] [discardpinnedcache=<Val>] [failPDOnReadME=on|off] [Lowlatency=low|off]]
```

storcli /cx show autoconfig

This command lets you view the automatic physical drive policy.

Input example:

```
storcli /c0 show autoconfig
```

storcli /cx set autoconfig=JBOD

This command lets you set the controller's automated physical drive policy to JBOD . When set to JBOD , all unconfigured physical drives are configured as JBODs.

NOTE

If this command fails, enable the legacy JBOD mode first and retry the command.

Input example:

```
storcli /c0 set autoconfig=JBOD
```

storcli /cx set jbod=on|off

This command lets you enable the legacy JBOD mode.

Input example:

```
storcli /c0 set jbod=on
```

storcli /cx set autoconfig [=none | R0 [immediate] | JBOD> [usecurrent]] [[sesmgmt=on|off] [secured=on|off] [multipath=on|off] [multiinit=on|off] [discardpinnedcache=Val>] [failPDOnReadME=on|off] [Lowlatency=low|off]]

This command enables or disables autoconfig on the selected adapter(s). You can use the following options with the `set autoconfig` command.

Table 59: Set Autoconfig Command Options

Option	Description
none	Disable autoconfig.
r0	Autoconfig for single PD as r0.
JBOD	Autoconfig each PD as JBOD.
immediate	Configure.
unconfigured	Pds once and will not configure any newly inserted UG drives and this option is only for R0.
usecurrent	Use current parameter values that are supported while changing the mode.
sesmgmt	Enable or disable SES management.
secured	Enable or disable Security on SED.
multiinit	Enable or disable multi init.
multipath	Enable or disable multi path.

Input example:

```
storcli /c0 set autoconfig [=R0 [immediate]> [usecurrent] ] [[sesmgmt=on] [secured=on] [multipath=on] [multiinit=on] [discardpinnedcache=Val>] [failPDOnReadME=on] [Lowlatency=low]]
```

Frequently Used Tasks

Displaying the Version of the StorCLI Utility

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

```
storcli -v
```

Displaying the StorCLI Tool Help

The following command displays the StorCLI Utility help:

```
storcli -h
```

Help appears for all the StorCLI tool commands.

Displaying System Summary Information

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

```
storcli -show [all]
```

Displaying Free Space in a Controller

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

```
storcli /cx show freespace
```

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 MegaRAID controllers and only 64 KB for Integrated MegaRAID 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, and 60.
- The size of each virtual drive.
- The drives that 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, or 1024.

NOTE

The supported strip size can vary from a minimum of 64 KB to 1 MB for MegaRAID controllers and only 64 KB for Integrated MegaRAID 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.

Input example:

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 64 KB.

Setting the Cache Policy in a Virtual Drive

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

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

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

Displaying Virtual Drive Information

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

```
storcli /cx/v(x|all) show
```

Deleting Virtual Drives

The following command deletes virtual drives:

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

The following inputs are required when deleting a virtual drive:

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

Flashing Controller Firmware

The following command is used to flash the controller firmware.

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

NOTE

The command output for this command cannot be JSON formatted.

Events, Messages, and Behaviors

This appendix lists the LSA events that can appear in the event log and event messages.

LSA software monitors the activity and performance of all controllers in the workstation and the devices attached to them. When an event occurs, such as the start of an initialization, an event message appears in the log at the bottom of the LSA main menu window. The messages are also logged in the Windows Application log (Event Viewer).

Error Levels

Each message that appears in the event log has a severity level that indicates the severity of the event, as shown in the following table.

Table 60: Event Error Levels

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

Event Messages

The following table lists all of the event messages. The event message descriptions include placeholders for specific values that are determined when the event is generated. For example, in message 0x0000 in the Event Messages table, “%s” is replaced by the firmware version, which is read from the firmware when the event is generated.

Table 61: Event Messages

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0000	Information	MegaRAID firmware initialization started (PCI ID %04x/%04x/%04x/%04x)	Logged at firmware initialization.
0x0001	Information	MegaRAID firmware version %s	Logged at firmware initialization to display firmware version.
0x0002	Fatal	Unable to recover cache data from TBBU	Currently not logged.
0x0003	Information	Cache data recovered from TBBU successfully	Currently not logged.
0x0004	Information	Configuration cleared	Logged when controller configuration is cleared.
0x0005	Warning	Cluster down; communication with peer lost	Currently not logged.
0x0006	Information	Virtual drive %s ownership changed from %02x to %02x	Currently not logged.
0x0007	Information	Alarm disabled by user	Logged when user disables alarm.
0x0008	Information	Alarm enabled by user	Logged when user enables alarm.
0x0009	Information	Background initialization rate changed to %d% %	Logged to display background initialization progress indication in percentage.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x000A	Fatal	Controller cache discarded due to memory/energy pack problems	Logged on cache discard due to hardware problems.
0x000B	Fatal	Unable to recover cache data due to configuration mismatch	Currently not logged.
0x000C	Information	Cache data recovered successfully	Logged when cache data is successfully recovered after reboot.
0x000D	Fatal	Controller cache discarded due to firmware version incompatibility	Logged when cache data discarded because of firmware version mismatch.
0x000E	Information	consistency check rate changed to %d%%	Logged to display consistency check progress indication percentage.
0x000F	Fatal	Fatal firmware error: %s	Logged in case of fatal errors and also while entering debug monitor.
0x0010	Information	Factory defaults restored	Logged while controller is reset to factory defaults.
0x0011	Information	Flash downloaded image corrupt	Logged to inform downloaded flash image is corrupt.
0x0012	Critical	Flash erase error	Logged in case of flash erase failure, generally after flash update.
0x0013	Critical	Flash timeout during erase	Logged to indicate flash erase operation timed out.
0x0014	Critical	Flash error	Generic unknown internal error during flash update flash.
0x0015	Information	Flashing image: %s	Logged to display flash image name string before getting updated to controller.
0x0016	Information	Flash of new firmware images complete	Logged to inform successful update of flash images.
0x0017	Critical	Flash programming error	Logged to notify, write failure during flash update, not being allowed usually due to internal controller settings.
0x0018	Critical	Flash timeout during programming	Logged to indicate flash write operation timed out.
0x0019	Critical	Flash chip type unknown	Logged during flash update tried with unsupported flash chip type.
0x001A	Critical	Flash command set unknown	Logged while unsupported flash command set detected, most likely because of unsupported flash chip.
0x001B	Critical	Flash verify failure	Logged when compare operation fails between written flash data and original data.
0x001C	Information	Flush rate changed to %d seconds	Logged to notify modified cache flush frequency in seconds.
0x001D	Information	Hibernate command received from host	Logged to inform about reception of hibernation command from host to controller, generally during host shutdown.
0x001E	Information	Event log cleared	Logged when controller log has been cleared.
0x001F	Information	Event log wrapped	Logged when controller log has been wrapped around, when the maximum logs are written.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0020	Fatal	Multi-bit ECC error: ECAR=%x, ELOG=%x, (%s)	Logged to notify ECC multi-bit error in memory, ELOG: ecc info (source, type, syndrome), ECAR: ecc address.
0x0021	Warning	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s)	Logged to notify ECC single-bit error in memory, ELOG: ecc info (source, type, syndrome), ECAR: ecc address.
0x0022	Fatal	Not enough controller memory	Logged to notify fatal controller condition, when you run out of memory to allocate.
0x0023	Information	Patrol Read complete	Logged when patrol read completes.
0x0024	Information	Patrol Read paused	Logged when patrol read is paused.
0x0025	Information	Patrol Read Rate changed to %d%%	Logged to indicate progress of patrol read in percentage.
0x0026	Information	Patrol Read resumed	Logged when patrol read is resumed.
0x0027	Information	Patrol Read started	Logged when patrol read is started.
0x0028	Information	Reconstruction rate changed to %d%%"	Logged to indicate progress of reconstruction in percentage.
0x0029	Information	Drive group modification rate changed to %d% %	Logged to indicate the change in drive group modification frequency.
0x002A	Information	Shutdown command received from host	Logged when shutdown command is received from host to controller.
0x002B	Information	Test event: %s	General controller event, with a generic string.
0x002C	Information	Time established as %s; (%d seconds since power on)	Logged when controller time was set from host, also displaying time since power on in seconds.
0x002D	Information	User entered firmware debugger	Logged when user enters controller debug shell.
0x002E	Warning	Background Initialization aborted on %s	Logged to inform about user aborted background initialization on displayed LD number.
0x002F	Warning	Background Initialization corrected medium error (%s at %lx	logged to inform about corrected medium error on displayed LD number, LBA number, PD number and PDLBA number in that order.
0x0030	Information	Background Initialization completed on %s	Logged to inform background initialization completion on displayed LD.
0x0031	Fatal	Background Initialization completed with uncorrectable errors on %s	Logged to inform background initialization completion with error on displayed LD.
0x0032	Fatal	Background Initialization detected uncorrectable double medium errors (%s at %lx on %s)	Logged to inform background initialization completion with double medium error on displayed PD, PDLBA, and LD in that order.
0x0033	Critical	Background Initialization failed on %s	Logged to inform background initialization failure on displayed LD.
0x0034	Progress	Background Initialization progress on %s is %s	Logged to inform background initialization progress in percentage of displayed LD.
0x0035	Information	Background Initialization started on %s	Logged to inform background initialization started for displayed LD.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0036	Information	Policy change on %s from %s to %s	Logged to inform the changed policy for displayed LD with old and new policies.
0x0038	Warning	Consistency Check aborted on %s	Logged to inform aborted consistency check for displayed LD.
0x0039	Warning	Consistency Check corrected medium error (%s at %lx	Logged when consistency check corrected medium error.
0x003A	Information	Consistency Check done on %s	Logged when consistency check has completed successfully on the LD.
0x003B	Information	Consistency Check done with corrections on %s	Logged when consistency check completed and inconsistency was found during check and was corrected.
0x003C	Fatal	Consistency Check detected uncorrectable double medium errors (%s at %lx on %s)	Logged when uncorrectable double medium error are detected while consistency check.
0x003D	Critical	Consistency Check failed on %s	Logged when consistency check failed as fatal error was found.
0x003E	Fatal	Consistency Check completed with uncorrectable data on %s	Logged when uncorrectable error occurred during consistency check.
0x003F	Warning	Consistency Check found inconsistent parity on %s at strip %lx	Logged when consistency check finds inconsistency parity on a strip.
0x0040	Warning	Consistency Check inconsistency logging disabled on %s (too many inconsistencies)	Logged when consistency check finds too many inconsistent parity (greater than 10) and the inconsistency parity logging is disabled.
0x0041	Progress	Consistency Check progress on %s is %s	Logged consistency check progress, the progress is logged only if the progress is greater than 1% at an interval of every 15 seconds.
0x0042	Information	Consistency Check started on %s	Logged when consistency check has started.
0x0043	Warning	Initialization aborted on %s	Logged when consistency check is aborted by you or for some other reason.
0x0044	Critical	Initialization failed on %s	Logged when initialization has failed.
0x0045	Progress	Initialization progress on %s is %s	Logged initialization progress, the progress is logged only if the progress is greater than 1% at an interval of every 15 seconds.
0x0046	Information	Fast initialization started on %s	Logged when quick initialization has started on an LD. The parameter to decide quick initialization or full initialization is passed by you.
0x0047	Information	Full initialization started on %s	Logged when full initialization has started.
0x0048	Information	Initialization complete on %s	Logged when initialization has completed successfully.
0x0049	Information	LD Properties updated to %s (from %s)	Logged when LD properties has been changed.
0x004A	Information	Reconstruction complete on %s	Logged when reconstruction has completed successfully.
0x004B	Fatal	Reconstruction of %s stopped due to unrecoverable errors	Logged when reconstruction has finished because of failure (unrecoverable errors).

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x004C	Fatal	Reconstruct detected uncorrectable double medium errors (%s at %lx on %s at %lx)	Logged while reconstructing if an unrecoverable double medium error is encountered.
0x004D	Progress	Reconstruction progress on %s is %s	Logs reconstruction progress, the progress is logged only if the progress is greater than 1% at an interval of every 15 seconds.
0x004E	Information	Reconstruction resumed on %s	Logged when reconstruction resumes after a power cycle.
0x004F	Fatal	Reconstruction resume of %s failed due to configuration mismatch	Logged when reconstruction resume failed due to configuration mismatch.
0x0050	Information	Reconstruction started on %s	Logged on start of reconstruction on an LD.
0x0051	Information	State change on %s from %s to %s	Logged when there is change in LD state. The event gives the new and old state. The state could be one of the following: LDS_OFFLINE, LDS_PARTIALLY_DEGRADED, LDS_DEGRADED, LDS_OPTIMAL.
0x0052	Information	Drive Clear aborted on %s	Logged when PD clear is aborted.
0x0053	Critical	Drive Clear failed on %s (Error %02x)	Logged when drive clear is failed and the even is logged along with error code.
0x0054	Progress	Drive Clear progress on %s is %s	Logs drive clear progress, the progress is logged only if the progress is greater than 1% at an interval of every 15 seconds.
0x0055	Information	Drive Clear started on %s	Logged when drive clear started on a PD.
0x0056	Information	Drive Clear completed on %s	Logged when PD clear task is completed successfully on a PD.
0x0057	Warning	Error on %s (Error %02x)	Logged if read returns with uncorrectable error or same errors on both the drives or write long returns with an error (for example, puncture operation could failed).
0x0058	Information	Format complete on %s	Logged when format has completed.
0x0059	Information	Format started on %s	Logged when format unit is started on a PD.
0x005A	Critical	Hot Spare SMART polling failed on %s (Error %02x)	Currently not logged.
0x005B	Information	Drive inserted: %s	Logged when drive is inserted and slot/enclosure fields of a PD are updated.
0x005C	Warning	Drive %s is not supported	Logged when the drive is not supported; reason could be the number of drive has exceeded the maximum number of supported drives or an unsupported drive is inserted like a SATA drive in SAS only enclosure or could be a unsupported drive type.
0x005D	Warning	Patrol Read corrected medium error on %s at %lx	Logged when patrol read has successfully completed recovery read and recovered data.
0x005E	Progress	Patrol Read progress on %s is %s	Logged only when patrol read has progress greater than 10% at an interval of every 15 seconds.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x005F	Fatal	Patrol Read found an uncorrectable medium error on %s at %lx	Logged when patrol read is unable to recover data.
0x0060	Warning	Predictive failure: CDB: %s	Logged when a failure is found during smart (predictive failure) poll.
0x0061	Fatal	Patrol Read puncturing bad block on %s at %lx	Logged when patrol read punctures a block due to unrecoverable medium error.
0x0062	Information	Rebuild aborted by user on %s	Logged when the user aborts a rebuild operation.
0x0063	Information	Rebuild complete on %s	Logged when the rebuild operation on a logical drive on a physical drive (which can have multiple LDs) is completed.
0x0064	Information	Rebuild complete on %s	Logged when rebuild operation is completed for all logical drives on a given physical drive.
0x0065	Critical	Rebuild failed on %s due to source drive error	Logged if one of the source drives for the rebuild operation fails or is removed.
0x0066	Critical	Rebuild failed on %s due to target drive error	Logged if the target rebuild drive (on which rebuild operation is going on) fails or is removed from the controller.
0x0067	Progress	Rebuild progress on %s is %s	Logged to indicate the progress (in percentage) of the rebuild operation on a given physical drive.
0x0068	Information	Rebuild resumed on %s	Logged when the rebuild operation on a physical drive resumes.
0x0069	Information	Rebuild started on %s	Logged when the rebuild operation is started on a physical drive.
0x006A	Information	Rebuild automatically started on %s	Logged when the rebuild operation kicks in on a spare.
0x006B	Critical	Rebuild stopped on %s due to loss of cluster ownership	Logged when the rebuild operation is stopped due to loss of ownership.
0x006C	Fatal	Reassign write operation failed on %s at %lx	Logged when a check condition or medium error is encountered for a reassigned write.
0x006D	Fatal	Unrecoverable medium error during rebuild on %s at %lx	Logged when the rebuild I/O encounters an unrecoverable medium error.
0x006E	Information	Corrected medium error during recovery on %s at %lx	Logged when recovery completed successfully and fixed a medium error.
0x006F	Fatal	Unrecoverable medium error during recovery on %s at %lx	Logged when the recovery for a failed I/O encounters a medium error.
0x0070	Information	Drive removed: %s	Logged when a drive is removed from the controller.
0x0071	Information	Unexpected sense: %s, CDB%s, Sense: %s	Logged when an I/O fails due to unexpected reasons and sense data needs to be logged.
0x0072	Information	State change on %s from %s to %s	Logged when the state of a drive is changed by the firmware or by you.
0x0073	Information	State change by user on %s from %s to %s	Not logged by the firmware.
0x0074	Warning	Redundant path to %s broken	Not logged by the firmware.
0x0075	Information	Redundant path to %s restored	Not logged by the firmware

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0076	Information	Dedicated Hot Spare Drive %s no longer useful due to deleted drive group	Not logged by the firmware.
0x0077	Critical	SAS topology error: Loop detected	Logged when device discovery fails for a SAS device as a loop was detected.
0x0078	Critical	SAS topology error: Unaddressable device	Logged when device discovery fails for a SAS device as an unaddressable device was found.
0x0079	Critical	SAS topology error: Multiple ports to the same SAS address	Logged when device discovery fails for a SAS device multiple ports with same SAS address were detected.
0x007A	Critical	SAS topology error: Expander error	Not logged by the firmware.
0x007B	Critical	SAS topology error: SMP timeout	Logged when device discovery fails for a SAS device due to an SMP timeout.
0x007C	Critical	SAS topology error: Out of route entries	Logged when device discovery fails for a SAS device as expander route table is out of entries.
0x007D	Critical	SAS topology error: Index not found	Logged when device discovery fails for a SAS device as expander route table out of entries.
0x007E	Critical	SAS topology error: SMP function failed	Logged when device discovery fails for a SAS device due to an SMP function failure.
0x007F	Critical	SAS topology error: SMP CRC error	Logged when device discovery fails for a SAS device due to an SMP CRC error.
0x0080	Critical	SAS topology error: Multiple subtractive	Logged when device discovery fails for a SAS device as a subtractive-to-subtractive link was detected.
0x0081	Critical	SAS topology error: Table to table	Logged when device discovery fails for a SAS device as table-to-table link was detected.
0x0082	Critical	SAS topology error: Multiple paths	Not logged by the firmware.
0x0083	Fatal	Unable to access device %s	Logged when the inserted drive is bad and unusable.
0x0084	Information	Dedicated Hot Spare created on %s (%s)	Logged when a drive is configured as a dedicated spare.
0x0085	Information	Dedicated Hot Spare %s disabled	Logged when a drive is removes as a dedicated spare.
0x0086	Critical	Dedicated Hot Spare %s no longer useful for all drive groups	Logged when an array with a dedicated spare is resized. The hot spare (dedicated to this array and possibly others) will not be applicable to other arrays.
0x0087	Information	Global Hot Spare created on %s (%s)	Logged when a drive is configured as a global hot spare.
0x0088	Information	Global Hot Spare %s disabled	Logged when a drive configured as global host spare fails or is unconfigured by you.
0x0089	Critical	Global Hot Spare does not cover all drive groups	Logged when the global hot spare is too small (or does not meet the SAS/SATA restrictions) to cover certain arrays.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x008A	Information	Created %s}	Logged as soon as the new logical drive created is added to the firmware configuration.
0x008B	Information	Deleted %s}	Logged when the firmware removes an LD from its configuration upon a user request from the applications.
0x008C	Information	Marking LD %s inconsistent due to active writes at shutdown	Logged when active writes are on one of the target disks of a RAID 5 LD at the time of shutdown.
0x008D	Information	Energy Pack Present	Logged during firmware initialization when checked if there is a energy pack present and the check turns out true. This event is also logged when a energy pack is inserted or replaced with a new one and the energy pack present check returns true.
0x008E	Warning	Energy Pack Not Present	Logged if the user has not disabled "Energy Pack Not Present" warning at the boot time or if a energy pack has been removed.
0x008F	Information	New Energy Pack Detected	Logged when there is a subsequent boot after a new energy pack has been inserted.
0x0090	Information	Energy Pack has been replaced	Logged when a new energy pack has been replaced with an old energy pack.
0x0091	Critical	Energy Pack temperature is high	Logged when detected that the energy pack temperature is high during the periodic energy pack status check.
0x0092	Warning	Energy Pack voltage low	Not logged by the firmware.
0x0093	Information	Energy Pack started charging	Logged as part of monitoring the energy pack status when the energy pack is getting charged.
0x0094	Information	Energy Pack is discharging	Logged as part of monitoring the energy pack status when the energy pack is getting discharged.
0x0095	Information	Energy Pack temperature is normal	Logged as part of monitoring the energy pack status when the temperature of the energy pack is normal.
0x0096	Fatal	Energy Pack has failed and cannot support data retention. Please replace the Energy Pack.	Logged when there is not enough capacity left in energy pack for expected data retention time. Energy Pack has to be replaced.
0x0097	Information	Energy Pack relearn started	Logged when the energy pack relearn started, initiated either by the user or automatically.
0x0098	Information	Energy Pack relearn in progress	Logged as part of monitoring the energy pack status when the energy pack relearn is in progress.
0x0099	Information	Energy Pack relearn completed	Logged as part of monitoring the energy pack status when the energy pack relearn is complete.
0x009A	Critical	Energy Pack relearn timed out	Not logged by the firmware.
0x009B	Information	Energy Pack relearn pending: Energy Pack is under charge	Logged as part of monitoring the energy pack status when the energy pack relearn is requested but yet to start.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x009C	Information	Energy Pack relearn postponed	Logged as part of monitoring the energy pack status when the energy pack relearn is requested but postponed as there is valid pinned cache present. This event can also be logged when learn delay interval has been explicitly set.
0x009D	Information	Energy Pack relearn will start in 4 days	Logged as part of providing energy pack learn cycle information when auto learn is enabled.
0x009E	Information	Energy Pack relearn will start in 2 day	Logged as part of providing energy pack learn cycle information when auto learn is enabled.
0x009F	Information	Energy Pack relearn will start in 1 day	Logged as part of providing energy pack learn cycle information when auto learn is enabled.
0x00A0	Information	Energy Pack relearn will start in 5 hours	Logged as part of providing energy pack learn cycle information when auto learn is enabled.
0x00A1	Information	Energy Pack removed	Logged as part of periodic monitoring of the energy pack status when a energy pack has been removed.
0x00A2	Information	Current capacity of the energy pack is below threshold	Logged as part of monitoring the energy pack status when the capacity of the energy pack is below threshold.
0x00A3	Information	Current capacity of the energy pack is above threshold	Logged as part of monitoring the energy pack status when the capacity of the energy pack is above threshold.
0x00A4	Information	Enclosure (SES) discovered on %s	Logged when an enclosure (SES) is discovered for the first time.
0x00A5	Information	Enclosure (SAFTE) discovered on %s	Not logged by the firmware.
0x00A6	Critical	Enclosure %s communication lost	Logged when the communication with an enclosure has been lost.
0x00A7	Information	Enclosure %s communication restored	Logged when the communication with an enclosure has been restored.
0x00A8	Critical	Enclosure %s fan %d failed	Logged when an enclosure fan has failed.
0x00A9	Information	Enclosure %s fan %d inserted	Logged when an enclosure fan has been inserted newly.
0x00AA	Critical	Enclosure %s fan %d removed	Logged when an enclosure fan has been removed.
0x00AB	Critical	Enclosure %s power supply %d failed	Not logged by the firmware.
0x00AC	Information	Enclosure %s power supply %d inserted	Logged when power supply has been inserted to an enclosure.
0x00AD	Critical	Enclosure %s power supply %d removed	Logged when power supply has been removed from an enclosure.
0x00AE	Critical	Enclosure %s SIM %d failed	Logged when the enclosure SIM has failed.
0x00AF	Information	Enclosure %s SIM %d inserted	Logged when an enclosure SIM has been inserted.
0x00B0	Critical	Enclosure %s SIM %d removed	Logged when an enclosure initialization was completed but later the SIM was removed.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x00B1	Warning	Enclosure %s temperature sensor %d below warning threshold	Logged when the enclosure services process has detected a temperature lower than a normal operating temperature or lower than the value indicated by the LOW WARNING THRESHOLD field in the Threshold In diagnostic page.
0x00B2	Critical	Enclosure %s temperature sensor %d below error threshold	Logged when the enclosure services process has detected a temperature lower than a safe operating temperature or lower than the value indicated by the LOW CRITICAL THRESHOLD field in the Threshold In diagnostic page.
0x00B3	Warning	Enclosure %s temperature sensor %d above warning threshold	Logged when the enclosure services process has detected a temperature higher than a normal operating temperature or higher than the value indicated by the HIGH WARNING THRESHOLD field in the Threshold In diagnostic page.
0x00B4	Critical	Enclosure %s temperature sensor %d above error threshold	Logged when the enclosure services process has detected a temperature higher than a safe operating temperature or higher than the value indicated by the HIGH CRITICAL THRESHOLD field in the Threshold In diagnostic page.
0x00B5	Critical	Enclosure %s shutdown	Logged when an unrecoverable condition is detected in the enclosure.
0x00B6	Warning	Enclosure %s not supported; too many enclosures connected to port	Logged when the maximum allowed enclosures per port is exceeded.
0x00B7	Critical	Enclosure %s firmware mismatch	Logged when two ESMs have different firmware versions.
0x00B8	Warning	Enclosure %s sensor %d bad	Logged when the device is present on the phy, but the status does not indicate its presence.
0x00B9	Critical	Enclosure %s phy %d bad	Logged when the status indicates a device presence, but there is no corresponding SAS address is associated with the device.
0x00BA	Critical	Enclosure %s is unstable	Logged when the enclosure services process reports the sense errors.
0x00BB	Critical	Enclosure %s hardware error	Logged when a critical or an unrecoverable enclosure failure has been detected by the enclosure services process.
0x00BC	Critical	Enclosure %s not responding	Logged when there is no response from the enclosure.
0x00BD	Information	SAS/SATA mixing not supported in enclosure; Drive %s disabled	Logged when the SAS/SATA mixing in an enclosure is being violated.
0x00BE	Information	Enclosure (SES) hotplug on %s was detected, but is not supported	Not reported to the user.
0x00BF	Information	Clustering enabled	Logged when the clustering is enabled in the controller properties.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x00C0	Information	Clustering disabled	Logged when the clustering is disabled in the controller properties.
0x00C1	Information	Drive too small to be used for auto-rebuild on %s	Logged when the size of the drive is not sufficient for auto-rebuild.
0x00C2	Information	BBU enabled; changing WT virtual drives to WB	Logged when changing WT virtual drives to WB and the BBU status is good.
0x00C3	Warning	BBU disabled; changing WB virtual drives to WT	Logged when changing WB virtual drives to WT and the BBU status is bad.
0x00C4	Warning	Bad block table on drive %s is 80% full	Logged when the bad block table on a drive is 80% full.
0x00C5	Fatal	Bad block table on drive %s is full; unable to log block %lx	Logged when the bad block table on a drive is full and not able to add the bad block in the bad block table.
0x00C6	Information	Consistency Check Aborted due to ownership loss on %s	Logged when the consistency check is aborted due to ownership is lost.
0x00C7	Information	Background Initialization (BGI) Aborted Due to Ownership Loss on %s	Logged when the background initialization (BGI) is aborted due to ownership loss.
0x00C8	Critical	Energy Pack/charger problems detected; SOH Bad	Logged when the energy pack is not presented or removed and SOH is bad.
0x00C9	Warning	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); warning threshold exceeded	Logged when the single-bit ECC errors exceeded the warning threshold.
0x00CA	Critical	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); critical threshold exceeded	Logged when the single-bit ECC errors exceeded the critical threshold.
0x00CB	Critical	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); further reporting disabled	Logged when the single-bit ECC errors exceeded all the thresholds and disable further logging.
0x00CC	Critical	Enclosure %s Power supply %d switched off	Logged when the enclosure services process has detected that the enclosure power supply is switched off and it was switched on earlier.
0x00CD	Information	Enclosure %s Power supply %d switched on	Logged when the enclosure services process has detected that the enclosure power supply is switched on and it was switched off earlier.
0x00CE	Critical	Enclosure %s Power supply %d cable removed	Logged when the enclosure services process has detected that the enclosure power supply cable is removed and it was inserted earlier.
0x00CF	Information	Enclosure %s Power supply %d cable inserted	Logged when the enclosure services process has detected that the enclosure power supply cable is inserted and it was removed earlier.
0x00D0	Information	Enclosure %s Fan %d returned to normal	Logged when the enclosure services process has detected that the current status of a fan is good and it was failed earlier.
0x00D1	Information	BBU Retention test was initiated on previous boot	Logged when the energy pack retention test was initiated on previous boot.
0x00D2	Information	BBU Retention test passed	Logged when the energy pack retention test passed successfully.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x00D3	Critical	BBU Retention test failed!	Logged when the energy pack retention test failed.
0x00D4	Information	NVRAM Retention test was initiated on previous boot	Logged when the NVRAM retention test was initiated on previous boot.
0x00D5	Information	NVRAM Retention test passed	Logged when the NVRAM retention test passed successfully.
0x00D6	Critical	NVRAM Retention test failed!	Logged when the NVRAM retention test failed.
0x00D7	Information	%s test completed %d passes successfully	Logged when the controller diagnostics test passes successfully.
0x00D8	Critical	%s test FAILED on %d pass. Fail data: errorOffset=%x goodData=%x badData=%x	Logged when the controller diagnostics test fails.
0x00D9	Information	Self check diagnostics completed	Logged when self-check diagnostics is completed.
0x00DA	Information	Foreign Configuration detected	Logged when foreign configuration is detected.
0x00DB	Information	Foreign Configuration imported	Logged when foreign configuration is imported.
0x00DC	Information	Foreign Configuration cleared	Logged when foreign configuration is cleared.
0x00DD	Warning	NVRAM is corrupt; reinitializing	Logged when NVRAM is corrupt and re-initialized.
0x00DE	Warning	NVRAM mismatch occurred	Logged when NVRAM mismatch occurs.
0x00DF	Warning	SAS wide port %d lost link on PHY %d	Logged when SAS wide port lost link on a PHY.
0x00E0	Information	SAS wide port %d restored link on PHY %d	Logged when a SAS wide port restored link on a PHY.
0x00E1	Warning	SAS port %d, PHY %d has exceeded the allowed error rate	Logged when a SAS PHY on port has exceeded the allowed error rate.
0x00E2	Warning	Bad block reassigned on %s at %lx to %lx	Logged when a bad block is reassigned on a drive from a error sector to a new sector.
0x00E3	Information	Controller Hot Plug detected	Logged when a controller hot plug is detected.
0x00E4	Warning	Enclosure %s temperature sensor %d differential detected	Logged when an enclosure temperature sensor differential is detected.
0x00E5	Information	Drive test cannot start. No qualifying drives found	Logged when disk test cannot start. No qualifying disks found.
0x00E6	Information	Time duration provided by host is not sufficient for self check	Logged when time duration provided by the host is not sufficient for self check.
0x00E7	Information	Marked Missing for %s on drive group %d row %d	Logged when a physical drive is marked missing on an array at a particular row.
0x00E8	Information	Replaced Missing as %s on drive group %d row %d	Logged when a physical drive is replaced Missing on an array at a particular row.
0x00E9	Information	Enclosure %s Temperature %d returned to normal	Logged when an enclosure temperature returns to normal.
0x00EA	Information	Enclosure %s Firmware download in progress	Logged when an enclosure firmware download is in progress.
0x00EB	Warning	Enclosure %s Firmware download failed	Logged when an enclosure firmware download failed.
0x00EC	Warning	%s is not a certified drive	Logged if the drive is not certified.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x00ED	Information	Dirty cache data discarded by user	Logged when dirty cache data is discarded by the user.
0x00EE	Information	Drives missing from configuration at boot	Logged when physical drives are missing from configuration at boot.
0x00EF	Information	Virtual drives (VDs) missing drives and will go offline at boot: %s	Logged when virtual drives are missing drives and will go offline at boot.
0x00F0	Information	VDs missing at boot: %s	Logged when virtual drives are missing at boot.
0x00F1	Information	Previous configuration completely missing at boot	Logged when the previous configuration is completely missing at boot.
0x00F2	Information	Energy Pack charge complete	Logged when an energy pack charge is completed.
0x00F3	Information	Enclosure %s fan %d speed changed	Logged when an enclosure fan speed changed.
0x00F4	Information	Dedicated spare %s imported as global due to missing arrays	Logged when a dedicated spare is imported as global due to missing arrays.
0x00F5	Information	%s rebuild not possible as SAS/SATA is not supported in an array	Logged when a rebuild is not possible because SAS/SATA is not supported in an array.
0x00F6	Information	SEP %s has been rebooted as a part of enclosure firmware download. SEP will be unavailable until this process completes.	Logged when SEP has been rebooted as part of enclosure firmware download. It will be unavailable until reboot completes.
0x00F7	Information	Inserted PD: %s Info: %s	Logged when a physical drive is inserted.
0x00F8	Information	Removed PD: %s Info: %s	Logged when a physical drive is removed.
0x00F9	Information	VD %s is now OPTIMAL	Logged when a logical drive state changes to optimal.
0x00FA	Warning	VD %s is now PARTIALLY DEGRADED	Logged when a logical drive state changes to a partially degraded state.
0x00FB	Critical	VD %s is now DEGRADED	Logged when a logical drive state changes to degraded state.
0x00FC	Fatal	VD %s is now OFFLINE	Logged when a logical drive state changes to offline state.
0x00FD	Warning	Energy Pack requires reconditioning; please initiate a LEARN cycle	Logged when a energy pack requires reconditioning; initiate a LEARN cycle.
0x00FE	Warning	VD %s disabled because RAID-5 is not supported by this RAID key	Logged when a virtual drive is disabled because RAID 5 is not supported by this RAID key.
0x00FF	Warning	VD %s disabled because RAID-6 is not supported by this controller	Logged when a virtual drive is disabled because RAID 6 is not supported by this controller.
0x0100	Warning	VD %s disabled because SAS drives are not supported by this RAID key	Logged when a virtual drive is disabled because SAS drives are not supported by this RAID key.
0x0101	Warning	PD missing: %s	Logged to provide information about the missing drive during boot.
0x0102	Warning	Puncturing of LBAs enabled	Currently not logged in the firmware.
0x0103	Warning	Puncturing of LBAs disabled	Currently not logged in the firmware.
0x0104	Critical	Enclosure %s EMM %d not installed	Logged when enclosure SIM is not installed.
0x0105	Information	Package version %s	Prints the package version number.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0106	Warning	Global affinity Hot Spare %s commissioned in a different enclosure	Logged when a hot spare that is a part of an enclosure is commissioned in a different enclosure.
0x0107	Warning	Foreign configuration table overflow	Logged when the number of GUIDs to import exceeds the total supported by the firmware.
0x0108	Warning	Partial foreign configuration imported, PDs not imported:%s	Logged when all the foreign configuration drives could not be imported.
0x0109	Information	Connector %s is active	Logged during initial boot when a SAS MUX connector is found for the controller.
0x010A	Information	Board Revision %s	Logged during boot.
0x010B	Warning	Command timeout on PD %s, CDB:%s	Logged when command to a PD times out.
0x010C	Warning	PD %s reset (Type %02x)	Logged when PD is reset.
0x010D	Warning	VD bad block table on %s is 80% full	Logged when the number of bad blocks entries is at 80% of what can be supported in the firmware.
0x010E	Fatal	VD bad block table on %s is full; unable to log block %lx (on %s at %lx)	Logged when the number of bad blocks exceed what can be supported in the firmware.
0x010F	Fatal	Uncorrectable medium error logged for %s at %lx (on %s at %lx)	Logged when an uncorrectable medium error is detected.
0x0110	Information	VD medium error corrected on %s at %lx	Logged on the corrected medium error.
0x0111	Warning	Bad block table on PD %s is 100% full	Logged when the bad block table is 100% full. Any more media errors on this physical drive will not be logged in the bad block table.
0x0112	Warning	VD bad block table on PD %s is 100% full	Logged when the bad block table is 100% full. Any more media errors on this logical drive will not be logged in the bad block table.
0x0113	Fatal	Controller needs replacement, IOP is faulty	Currently not logged in the firmware.
0x0114	Information	Replace Drive started on PD %s from PD %s	Logged when replace is started.
0x0115	Information	Replace Drive aborted on PD %s and src is PD %s	Logged when replace is aborted.
0x0116	Information	Replace Drive complete on PD %s from PD %s	Logged when replace is completed.
0x0117	Progress	Replace Drive progress on PD %s is %s	Logged to provide the progress of replace.
0x0118	Information	Replace Drive resumed on PD %s from %s	Logged when replace operation is resumed.
0x0119	Information	Replace Drive automatically started on PD %s from %s	Logged on automatic start of replace.
0x011A	Critical	Replace Drive failed on PD %s due to source %s error	Logged when the source physical drive of a replace fails. The replace stops and rebuild starts on the destination physical drive.
0x011B	Warning	Early Power off warning was unsuccessful	Currently not logged in the firmware.
0x011C	Information	BBU FRU is %s	Logged only for IBM.
0x011D	Information	%s FRU is %s	Logged if FRU data is present. Logged only for IBM.
0x011E	Information	Controller hardware revision ID %s	Currently not used in the firmware.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x011F	Warning	Foreign import shall result in a backward incompatible upgrade of configuration metadata	Currently not used in the firmware.
0x0120	Information	Redundant path restored for PD %s	Logged when new path is added for the physical drives.
0x0121	Warning	Redundant path broken for PD %s	Logged when one path is removed.
0x0122	Information	Redundant enclosure EMM %s inserted for EMM %s	Logged when an enclosure is added.
0x0123	Information	Redundant enclosure EMM %s removed for EMM %s	Logged when an enclosure is removed
0x0124	Warning	Patrol Read can't be started, as PDs are either not ONLINE, or are in a VD with an active process, or are in an excluded VD	Logged when none of the disks can start PR.
0x0125	Information	Replace Drive aborted by user on PD %s and src is PD %s	Logged when replace is aborted by the user.
0x0126	Critical	Replace Drive aborted on hot spare %s from %s, as hot spare needed for rebuild	Logged when replace is aborted on a Hotspare.
0x0127	Warning	Replace Drive aborted on PD %s from PD %s, as rebuild required in the array	Logged when replace is stopped for a higher priority rebuild operation on a drive.
0x0128	Fatal	Controller cache discarded for missing or offline VD %s When a VD with cached data goes offline or missing during runtime, the cache for the VD is discarded. Because the VD is offline, the cache cannot be saved.	Logged when pinned cache lines are discarded for an LD.
0x0129	Information	Replace Drive cannot be started as PD %s is too small for src PD %s	Logged when the destination PD is too small for replace.
0x012A	Information	Replace Drive cannot be started on PD %s from PD %s, as SAS/SATA is not supported in an array	Logged when there is a SAS/SATA mixing violation for the destination PD.
0x012B	Information	Microcode update started on PD %s	Logged when PD firmware download starts.
0x012C	Information	Microcode update completed on PD %s	Logged when PD firmware download completes.
0x012D	Warning	Microcode update timeout on PD %s	Logged when PD firmware download does not complete and times out.
0x012E	Warning	Microcode update failed on PD %s	Logged when PD firmware download fails.
0x012F	Information	Controller properties changed	Logged when any of the controller properties has changed.
0x0130	Information	Patrol Read properties changed	Currently not logged in the firmware.
0x0131	Information	CC Schedule properties changed	Logged when consistency check scheduling property has changed.
0x0132	Information	Energy Pack properties changed	Logged when any of the BBU properties have changed.
0x0133	Warning	Periodic Energy Pack Relearn is pending. Please initiate manual learn cycle as Automatic learn is not enabled	Logged when BBU periodic relearn is pending.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0134	Information	Drive security key created	Logged when controller lock key is created.
0x0135	Information	Drive security key backed up	Logged when controller lock key is backed up.
0x0136	Information	Drive security key from escrow, verified	Logged when controller lock key is verified from escrow.
0x0137	Information	Drive security key changed	Logged when controller lock key is re-keyed.
0x0138	Warning	Drive security key, re-key operation failed	Logged when controller lock re-key operation failed.
0x0139	Warning	Drive security key is invalid	Logged when the controller lock is not valid.
0x013A	Information	Drive security key destroyed	Logged when the controller lock key is destroyed.
0x013B	Warning	Drive security key from escrow is invalid	Logged when the controller escrow key is not valid. This escrow key cannot unlock any drive.
0x013C	Information	VD %s is now secured	Logged when a secure LD is created.
0x013D	Warning	VD %s is partially secured	Logged when all the drives in the array are not secure.
0x013E	Information	PD %s security activated	Logged when a PD security key is set.
0x013F	Information	PD %s security disabled	Logged when a security key is removed from an FDE drive.
0x0140	Information	PD %s is reprovisioned	Logged when a PD security is cleared.
0x0141	Information	PD %s security key changed	Logged when a PD lock key is re-keyed.
0x0142	Fatal	Security subsystem problems detected for PD %s	Logged when a PD security cannot be set.
0x0143	Fatal	Controller cache pinned for missing or offline VD %s	Logged when an LD cache is pinned.
0x0144	Fatal	Controller cache pinned for missing or offline VD: %s	Logged when a pinned cache is found during Online Controller Reset (OCR).
0x0145	Information	Controller cache discarded by user for VD: %s	Logged when an LD pinned cache is discarded by the user.
0x0146	Information	Controller cache destaged for VD %s	Logged when an LD pinned cache is recovered.
0x0147	Warning	Consistency Check started on an inconsistent VD %s	Logged when a consistency check is started on an inconsistent LD.
0x0148	Warning	Drive security key failure, cannot access secured configuration	Logged when an invalid lock key is detected.
0x0149	Warning	Drive security password from user is invalid	Not logged.
0x014A	Warning	Detected error with the remote Energy Pack connector cable	Not logged.
0x014B	Information	Power state change on PD %s from %s to %s	Logged when the PD power state (spun up, spun down, in-transition) changes.
0x014C	Information	Enclosure %s element (SES code 0x%x) status changed	Not logged.
0x014D	Information	PD %s rebuild not possible as HDD/CacheCade software mix is not supported in a drive group	Logged when mixing violation occurs due to HDD/SSD mismatch.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x014E	Information	Replace Drive cannot be started on PD %s from %s, as HDD/CacheCade software mix is not supported in a drive group	Logged when replace could not be started on a PD because HDD/CacheCade software mix was not supported in a drive group.
0x014F	Information	VD bad block table on %s is cleared	Logged when a VD bad block table was cleared.
0x0150	Caution	SAS topology error: 0x%lx	Logged when a SAS topology error occurred.
0x0151	Information	VD cluster of medium errors corrected for %s at %lx (on %s at %lx)	Logged when medium errors were corrected for a PD for an LD.
0x0152	Information	Controller requests a host bus rescan	Logged when controller requested a host bus rescan.
0x0153	Information	Controller repurposed and factory defaults restored	Logged when controller repurposed and factory defaults were restored.
0x0154	Information	Drive security key binding updated	Logged when drive security key binding was updated.
0x0159	Critical	Controller encountered a fatal error and was reset	Logged when a controller encountered a fatal error and was reset.
0x015A	Information	Snapshots enabled on %s (Repository %s)	Logged when snapshot was enabled on an LD.
0x015B	Information	Snapshots disabled on %s (Repository %s) by the user	Logged when snapshot was disabled on an LD by the user.
0x015C	Critical	Snapshots disabled on %s (Repository %s), due to a fatal error	Logged when snapshot was disabled on an LD due to a fatal error.
0x015D	Information	Snapshot created on %s at %s	Logged when snapshot was created on an LD.
0x015E	Information	Snapshot deleted on %s at %s	Logged when snapshot was deleted on an LD.
0x015F	Information	View created at %s to a snapshot at %s for %s	Logged when view was created at an LD.
0x0160	Information	View at %s is deleted, to snapshot at %s for %s	Logged when view at an LD was deleted.
0x0161	Information	Snapshot rollback started on %s from snapshot at %s	Logged when snapshot rollback was started on an LD.
0x0162	Fatal	Snapshot rollback on %s internally aborted for snapshot at %s	Logged when snapshot rollback was internally aborted.
0x0163	Information	Snapshot rollback on %s completed for snapshot at %s	Logged when snapshot rollback on an LD was completed.
0x0164	Information	Snapshot rollback progress for snapshot at %s, on %s is %s	Logged to report snapshot rollback progress on an LD.
0x0165	Warning	Snapshot space for %s in snapshot repository %s, is 80%% full	Logged when snapshot space for an LD in a snapshot repository was 80% full.
0x0166	Critical	Snapshot space for %s in snapshot repository %s, is full	Logged when snapshot space for an LD in a snapshot repository was full.
0x0167	Warning	View at %s to snapshot at %s, is 80%% full on snapshot repository %s	Logged when view at an LD to a snapshot was 80% full on a snapshot repository.
0x0168	Critical	View at %s to snapshot at %s, is full on snapshot repository %s	Logged when view at an LD to a snapshot was full on a snapshot repository.
0x0169	Critical	Snapshot repository lost for %s	Logged when snapshot repository was lost for an LD.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x016A	Warning	Snapshot repository restored for %s	Logged when snapshot repository was restored for an LD.
0x016B	Critical	Snapshot encountered an unexpected internal error: 0x%lx	Logged when a snapshot encountered an unexpected internal error.
0x016C	Information	Auto Snapshot enabled on %s (snapshot repository %s)	Logged when an auto snapshot is enabled.
0x016D	Information	Auto Snapshot disabled on %s (snapshot repository %s)	Logged when auto snapshot was disabled.
0x016E	Critical	Configuration command could not be committed to disk, please retry	Logged when a configuration command could not be committed to disk and was asked to retry.
0x016F	Information	COD on %s updated as it was stale	Logged when COD in DDF is updated due to various reasons.
0x0170	Warning	Power state change failed on %s (from %s to %s)	Logged when a power state change failed on a PD.
0x0171	Warning	%s is not available	Logged when an LD was not available.
0x0172	Information	%s is available	Logged when an LD was available.
0x0173	Information	%s is used for CacheCade with capacity 0x%lx logical blocks	Logged when an LD was used for CacheCade with the indicated capacity in logical blocks.
0x0174	Information	%s is using CacheCade %s	Logged when an LD was using CacheCade.
0x0175	Information	%s is no longer using CacheCade %s	Logged when an LD was no longer using CacheCade.
0x0176	Critical	Snapshot deleted due to resource constraints for %s in snapshot repository %s	Logged when the snapshot is deleted due to resource constraints in the snapshot repository.
0x0177	Warning	Auto Snapshot failed for %s in snapshot repository %s	Logged when auto snapshot fails for a VD in snapshot repository.
0x0178	Warning	Controller reset on-board expander	Logged when the chip reset issued to on-board expander.
0x0179	Warning	CacheCade (%s) capacity changed and is now 0x%lx logical blocks	Logged when the CacheCade capacity is changed along with the current capacity.
0x017A	Warning	Energy Pack cannot initiate transparent learn cycles	Logged when the energy pack cannot initiate transparent learn cycles.
0x017B	Information	Premium feature %s key was applied for - %s	Logged when the premium feature key was applied.
0x017C	Information	Snapshot schedule properties changed on %s	Logged when the snapshot schedule properties changed.
0x017D	Information	Snapshot scheduled action is due on %s	Logged when the snapshot scheduled action is due.
0x017E	Information	Performance Metrics: collection command 0x %lx	Logged during the performance metrics collection.
0x017F	Information	Premium feature %s key was transferred - %s	Logged when the premium feature key was transferred.
0x0180	Information	Premium feature serial number %s	Logged when displaying the premium feature serial number.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0181	Warning	Premium feature serial number mismatched. Key-vault serial num - %s	Logged when premium feature serial number mismatched.
0x0182	Warning	Energy Pack cannot support data retention for more than %d hours. Please replace the Energy Pack	Logged during the energy pack monitoring and it displays the remaining data retention time of the energy pack.
0x0183	Information	%s power policy changed to %s (from %s)	Logged when the power policy of an LD is changed.
0x0184	Warning	%s cannot transition to max power savings	Logged when LD cannot transition to maximum power savings.
0x0185	Information	Host driver is loaded and operational	This event is not reported to the user.
0x0186	Information	%s mirror broken	Logged when the mirror is broken for an LD.
0x0187	Information	%s mirror joined	Logged when joining the LD with its broken mirror.
0x0188	Warning	%s link %d failure in wide port	This event is not reported to the user.
0x0189	Information	%s link %d restored in wide port	This event is not reported to the user.
0x018A	Information	Memory module FRU is %s	This event is not reported to the user.
0x018B	Warning	Cache-vault power pack is sub-optimal. Please replace the pack	This event is not reported to the user.
0x018C	Warning	Foreign configuration auto-import did not import any drives	Logged when the foreign configuration auto-import did not import any drives.
0x018D	Warning	Cache-vault microcode update required	Logged when the BMU is not in Normal mode and CacheVault microcode update required.
0x018E	Warning	CacheCade (%s) capacity exceeds maximum allowed size, extra capacity is not used	Logged when the CacheCade capacity exceeds maximum allowed size; extra capacity is not used.
0x018F	Warning	LD (%s) protection information lost	Logged when the protection information is lost for an LD.
0x0190	Information	Diagnostics passed for %s	Logged when the SHIELD diagnostics passed for a PD.
0x0191	Critical	Diagnostics failed for %s	Logged when the SHIELD diagnostics failed for a PD.
0x0192	Information	Server Power capability Diagnostic Test Started	Logged when the server power capability diagnostic test starts.
0x0193	Information	Drive Cache settings enabled during rebuild for %s	Logged when the drive cache settings are enabled during rebuild for a PD.
0x0194	Information	Drive Cache settings restored after rebuild for %s	Logged when the drive cache settings are restored after rebuild for a PD.
0x0195	Information	Drive %s commissioned as Emergency spare	Logged when the drive is commissioned as an emergency spare.
0x0196	Warning	Reminder: Potential non-optimal configuration due to drive %s commissioned as emergency spare	Logged when the PD is being imported is an emergency spare.
0x0197	Information	Consistency Check suspended on %s	Logged when the consistency check is suspended on an LD.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0198	Information	Consistency Check resumed on %s	Logged when the consistency check is resumed on an LD.
0x0199	Information	Background Initialization suspended on %s	Logged when the background initialization is suspended on an LD.
0x019A	Information	Background Initialization resumed on %	Logged when the background initialization is resumed on an LD.
0x019B	Information	Reconstruction suspended on %s	Logged when the reconstruction is suspended on an LD.
0x019C	Information	Rebuild suspended on %	Logged when the rebuild is suspended on a PD.
0x019D	Information	Replace Drive suspended on %s	Logged when the replace is suspended on a PD.
0x019E	Information	Reminder: Consistency Check suspended on %	Logged as a reminder when the consistency check is suspended on an LD.
0x019F	Information	Reminder: Background Initialization suspended on %s	Logged as a reminder when the background initialization is suspended on an LD.
0x01A0	Information	Reminder: Reconstruction suspended on %s	Logged as a reminder when the reconstruction is suspended on an LD.
0x01A1	Information	Reminder: Rebuild suspended on %s	Logged as a reminder when the rebuild is suspended on a PD.
0x01A2	Information	Reminder: Replace Drive suspended on %s	Logged as a reminder when replace is suspended on a PD.
0x01A3	Information	Reminder: Patrol Read suspended	Logged as a reminder when the patrol read is suspended.
0x01A4	Information	Erase aborted on %s	Logged when the erase is aborted on a PD.
0x01A5	Critical	Erase failed on %s (Error %02x)	Logged when the erase is failed on a PD along with the error.
0x01A6	Progress	Erase progress on %s is %s	Logged to display the erase progress on a PD along with its current progress.
0x01A7	Information	Erase started on %s	Logged when erase is started on a PD.
0x01A8	Information	Erase completed on %s	Logged when the erase is completed on a PD.
0x01A9	Information	Erase aborted on %s	Logged when the erase is aborted on an LD.
0x01AA	Critical	Erase failed on %s	Logged when the erase is failed on an LD.
0x01AB	Progress	Erase progress on %s is %s	Logged to display the erase progress on an LD along with its current progress.
0x01AC	Information	Erase started on %s	Logged when the erase is started on an LD.
0x01AD	Information	Erase complete on %s	Logged when the erase is complete on an LD.
0x01AE	Warning	Potential leakage during erase on %s	Logged to inform the potential leakage during erase on an LD.
0x01AF	Warning	Energy Pack charging was suspended due to high Energy Pack temperature	Logged when the energy pack charging was suspended due to high energy pack temperature.
0x01B0	Information	NVCache firmware update was successful	This event is not reported to the user.
0x01B1	Warning	NVCache firmware update failed	This event is not reported to the user.

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x01B2	Fatal	%s access blocked as cached data in CacheCade is unavailable	This event is not reported to the user.
0x01B3	Information	CacheCade disassociate started on %s	This event is not reported to the user.
0x01B4	Information	CacheCade disassociate completed on %s	This event is not reported to the user.
0x01B5	Critical	CacheCade disassociate failed on %s	This event is not reported to the user.
0x01B6	Progress	CacheCade disassociate progress on %s is %s	This event is not reported to the user.
0x01B7	Information	CacheCade disassociate aborted by user on %s	This event is not reported to the user.
0x01B8	Information	Link speed changed on SAS port %d and PHY %d	Logged when the link speed changed on SAS port and PHY.
0x01B9	Warning	Advanced Software Options was deactivated for – %s	This event is not reported to the user.
0x01BA	Information	%s is now accessible	This event is not reported to the user.
0x01BB	Information	%s is using CacheCade	This event is not reported to the user.
0x01BC	Information	%s is no longer using CacheCade	This event is not reported to the user.
0x01BD	Warning	Patrol Read aborted on %s	Logged when the patrol read is aborted on a PD.
0x01C2	Information	Periodic Energy Pack Relearn was missed, and rescheduled to %s	Logged if energy pack relearn was missed at the scheduled time due to a system power off then the controller will reschedule automatically when you power on the system.
0x01C3	Information	Controller reset requested by host	Logged when the controller reset process started on the corresponding controller.
0x01C4	Information	Controller reset requested by host, completed	Logged when the controller reset process completed on the corresponding controller.
0x01C7	Warning	Controller booted in headless mode with errors	Logged when the controller is booted to safe mode due to warning errors.
0x01C8	Critical	Controller booted to safe mode due to critical errors	Logged when the controller is booted to safe mode due to critical errors.
0x01C9	Warning	Warning Error during boot – %s	Logged when a warning error occurs during booting the controller to safe mode.
0x01CA	Critical	Critical Error during boot – %s	Logged when a critical error occurs during booting the controller to safe mode.
0x01CB	Fatal	Fatal Error during boot – %s	Logged when a fatal error occurs during booting the controller to safe mode.
0x01CC	Fatal	Fatal Error during boot – %s	Logged when the peer controller has joined the HA domain
0x01CD	Information	Peer controller has left HA domain (ID: %s)" }	Logged when peer controller has left the HA domain.
0x01CE	Information	"%s is managed by peer controller" }	Logged when physical drives are being managed by the other node.
0x01CF	Information	"%s is managed by local controller" }	Logged when physical drives are being managed by the current node
0x01D0	Information	"%s is managed by peer controller" }	Logged when logical drives are not managed by the current node

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x01D1	Information	"%s is managed by local controller" }	Logged when logical drives are not being managed by the local node
0x01D2	Information	"%s has a conflict in HA domain" }	Logged when there is a mismatch of the target ID on both the nodes.
0x01D3	Information	"%s access is shared" }	Logged when access to the virtual drive is shared between both the nodes
0x01D4	Information	"%s access is exclusive" }	Logged when an exclusive access policy has been granted to the virtual drive
0x01D5	Warning	"%s is incompatible in the HA domain" }	Logged when the logical drive is not compatible
0x01D6	Warning	"Peer controller is incompatible" }	Logged when the peer controller is not compatible.
0x01D7	Warning	"Controllers in the HA domain are incompatible" }	Logged when the controller is not compatible
0x01D8	Warning	"Controller properties are incompatible between local and peer controllers" }	Logged when the controller properties are not compatible between peers and local controllers
0x01D9	Warning	"FW versions do not match in the HA domain" }	Logged when there is a mismatch between the version of the firmware on both the nodes
0x01DA	Critical	"Advanced Software Options %s do not match in the HA domain" }	Logged when the controller features are different
0x01DB	Information	"HA cache mirror is online"}	Logged when the cache mirror operation is enabled
0x01DC	Information	"HA cache mirror is offline"}	Logged when the cache mirror operation is disabled
0x01DD	Information	"%s access blocked as cached data from peer controller is unavailable" }	Logged when a peer node is not available or a particular logical drive has been blocked
0x01DE	Warning	"Cache-vault power pack is not supported. Please replace the pack" }	Logged when the CacheVault is not supported
0x01DF	Information	%s temperature (%d C) is above warning threshold }	Logged when the temperature of a physical drive is more than the normal threshold
0x01E0	Information	%s temperature (%d C) is above critical threshold }	Logged when the temperature of a physical drive is more than the critical threshold
0x01E1	Information	%s temperature (%d C) is normal }	Logged when the temperature of the physical drive is normal
0x01E2	Warning	"%s IOs are being throttled" }	Logged when the I/O of the physical drive is throttled
0x01E3	Information	"%s IOs are normal. (No throttling)" }	Logged when the I/O of the physical drive is normal
0x01E4	Information	%s has %d%% life left. Life left thresholds – warning:%d%%, critical:%d%% }	Currently not logged
0x01E5	Warning	%s life left (%d%%) is below optimal. Life left thresholds - warning:%d%%, critical:%d%% }	Currently not logged
0x01E6	Critical	%s life left (%d%%) is critical. Life left thresholds – warning:%d%%, critical:%d%% }	Currently not logged
0x01E7	Critical	%s failure, device locked-up }	Currently not logged
0x01E8	Warning	"Host driver needs to be upgraded %s" }	Logged when the host drive requires an upgrade

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x01E9	Warning	"Direct communication with peer controller(s) was not established. Please check proper cable connections" }	Logged when there is a loss in the communication between the peer and other nodes
0x01EA	Critical	"Firmware image does not contain signed component" }	Currently not logged
0x01EB	Critical	"Authentication failure of the signed firmware image" }	Logged when there is an authentication failure on the firmware image that has been signed
0x01EC	Information	"Setting %s as boot device" }	Logged when the logical drive is set as a boot device
0x01ED	Information	"Setting %s as boot device" }	Logged when the physical drive is set a boot device
0x01EE	Information	The BBU temperature is changed to %d (Celsius) }	Logged when the temperature of the energy pack is changed
0x01EF	Information	The controller temperature is changed to %d (Celsius) }	Logged when the controller temperature is changed
0x01F0	Information	"NVCache capacity is too small to support data backup. Write-back VDs will be converted to write-through" }	Logged when the ONFI status of the Bad Block exceeds the prescribed limit
0x01F1	Information	"NVCache data backup capacity has decreased. Consider replacement	Logged when the NVCache needs to be replaced
0x01F2	Warning	"NVCache device failed. cannot support data retention" }	Logged when the ONFI model is unable to sustain the offload capability of NVCache
0x01F3	Critical	"Boot Device reset, setting target ID as invalid" }	Logged when the target ID of the boot device is invalid
0x01F4	Fatal	Write back Nytro cache size mismatch between the servers. The Nytro cache size was adjusted to %ld GB}	Currently not logged
0x01F5	Information	"%s is not shared between servers but assigned for caching. Write back Nytro cache content of the VD will be mirrored"} }	Currently not logged
0x01F6	Information	Power %d watts usage base IOs throttle started }	Currently not logged
0x01F7	Information	"Power base IOs throttle stopped" }	Currently not logged
0x01F8	Information	"Controller tunable parameter(s) changed"} }	Currently not logged
0x01F9	Information	"Controller operating temperature within normal range, full operation restored"} }	Logged when the temperature of the controller is within the normal range
0x01FA	Information	"Controller temperature threshold exceeded. This may indicate inadequate system cooling. Switching to low performance mode"} }	Logged when the temperature of the controller exceeds the specified threshold
0x01FB	Information	"Controller supports HA mode, currently functioning with HA feature set"} }	Currently not logged
0x01FC	Information	"Controller supports HA mode, currently functioning with single controller feature set" }	Currently not logged

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x01FD	Warning	"Cache-vault components mismatch. Write-back VDs will be converted write-through")	Logged when there is a mismatch in the CacheVault component. In this case, Writeback virtual drives are converted to write-through virtual drives
0x01FE	Warning	"Controller has entered into maintenance mode (%d)" }	Logged when the controller goes into maintenance mode
0x01FF	Information	"Controller has returned to normal mode" }	Logged when the controller has returned to normal mode from maintenance e mode
0x0200	Information	"Topology is in (%s) mode" }	Currently not logged
0x0201	Information	"Cannot enter (%s) mode because %s VD %s would not be supported" }	Currently not logged
0x0202	Information	"Cannot enter (%s) mode because %s PD %s would not be supported" }	Currently not logged
0x0203	Information	"%s Cache Flush started" }	Logged when the controller starts the Cache Flush operation
0x0204	Information	"%s Cache Flush finished" }	Logged when controller completes the Cache Flush operation
0x0205	Information	"%s Cache Flush aborted by user" }	Logged when the user aborts the Cache Flush operation
0x0206	Information	"Controller personality changed to (0x%x) mode"}	Logged when the user changes the personality of the controller
0x0207	Information	"Configuration automatically created by %s" }	Logged when a configuration is automatically created
0x0208	Information	"Software Zone enabled"}	Currently not logged
0x0209	Information	"Software Zone disabled"}	Currently not logged
0x020A	Information	"Initialization aborted on %s due to controller reset" }	Logged when the controller reset operation aborts the initialization of the logical drive
0x020B	Warning	"Peer controller security key mismatch" }	Currently not logged
0x020C	Information	"Peer controller security key match" }	Currently not logged
0x020D	Information	"%s is now compatible in the HA domain" }	Currently not logged
0x020E	Warning	"PD %s %s delayed. Reason: %s" }	Currently not logged
0x020F	Warning	"VD %s %s delayed. Reason: %s" }	Currently not logged
0x0210	Information	"%s" }	Logged by controller to provide a generic message to the user
0x0211	Information	%s }	Logged by controller to provide a generic message to the user
0x0212	Warning	"MegaRAID Solution will shut down due to maximum temperature threshold exception. This may indicate inadequate system cooling"}	Logged when the temperature of the controller is above the normal threshold
0x0213	Warning	"Shutdown chassis command received from host" }	Currently not logged
0x0214	Warning	"Shutdown chassis command received from host" }	Currently not logged

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0215	Information	"Controller Information changed" }	Logged by controller to provide a generic message to the user
0x0216	Information	"Hidden policy not set for all VDs in the array"	Logged when a logical drive is partially hidden, where the hidden policy is not set for virtual drives in an array
0x0217	Information	"%s is not enterprise class self encrypting drive. Encryption capability of the drive will be disabled"	Currently not logged
0x0218	Information	"Controller firmware was updated with force option"	Logged when the firmware image is updated by the user using the "force" option
0x0219	Information	"%s default access policy is set to read/write mode"	Logged when the default access policy of the logical drive set to Read or Write mode
0x021A	Information	"Disabling writes to flash due to a critical error. Reboot the system to enable writes to flash again"	Logged when the flash fails
0x021B	Warning	"Disabling writes to flash as the part has gone bad"	Logged when a flash device transitions to a bad state
0x021C	Information	"Locate LED started on %s"	Logged when Locate LED operation is started on a physical drive
0x021D	Information	"Locate LED stopped on %s"	Logged when Locate LED operation is stopped on a physical drive
0x021E	Information	"Patrol read aborted on %s due to conflict with other background operations"	Logged a when patrol read operation is aborted due to a conflict with other background operations that are running
0x021F	Warning	"%s %s %d bad media events"	Logged when degraded media events occur
0x0220	Warning	"%s %s has bad perf, %s"	Logged when degraded media events occur due to poor performance
0x0221	Information	"SCAP HLTH: %d mF, %d mOhm, %d mV, %d Deg, 55C:x%x, 60C:x%x, 65C:x%x, 70C:x%x, 75C:x%x"	Logged when the Relearn operation is completed to reflect the health of the Supercap
0x0222	Information	"Controller personality will change PCI ID to %04x/%04x/%04x/%04x"	Logged when there is a change in the controller personality
0x0223	Information	"%s Inquiry info: %s" }	Logged when there is an update to the slot enclosure field on the physical drive
0x0224	Information	"%s is marked as Transport Ready" }	Logged when a logical drive is marked as Transport Ready
0x0225	Information	"%s is cleared from Transport Ready state" }	Logged when a logical drive is cleared from the Transport Ready state
0x0226	Information	"System reset required."	Logged when a user requests new personality that requires a system restart
0x0227	Warning	"Block recovery is skipped for Cache-vault. Reboot the system to recover the blocks"	Currently not logged
0x0228	Information	"Auto configuration option is set to - %s" }	Logged when the auto configuration option is set to a specific behavior mode

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0229	Information	"Auto configuration parameters changed" }	Logged when there is a change in the auto configuration parameters
0x022A	Information	"Reminder: Transport Ready Present %s" }	Logged when there is a Transport Ready flag present on the logical drive
0x022B	Information	"Profile updated from profile id %s" }	Logged when the controller profile update is successful
0x022C	Information	"Profile id %s autoselected by the firmware" }	Logged when the firmware automatically selects the controller profile
0x022D	Information	"Device Ids have been changed due to profile change. " }	Logged when there is a switch from higher number of physical drives to a lower number of physical drives
0x022E	Information	"Save controller FRU State: FRU-ID 0x%lx Size 0x%x bytes" }	Currently not logged
0x022F	Information	"Restore controller FRU State: FRU-ID 0x%lx Size 0x%x bytes" }	Currently not logged
0x0230	Information	"Delete controller FRU State: FRU-ID 0x%lx" }	Currently not logged
0x0231	Critical	"%s driveErrorCounter %d slotErrorCounter %d Pd failed due to %s issue" }	Currently not logged
0x0232	Critical	"Invalid SAS Address present in MFC data" }	Logged when the MFC data page encounters an invalid SAS address
0x0233	Warning	"SAS topology error: %s" }	Logged when an error occurs in the SAS topology
0x0234	Warning	"Invalid NVDATA" }	Logged when the NVDATA is not valid
0x0235	Information	"No configuration present on the controller" }	Logged when there is no configuration present on the controller
0x0236	Information	"Foreign configuration unsupported by current firmware version" }	Logged when the current version of the firmware does not support foreign configuration
0x0237	Critical	"Backup firmware image flash programming error" }	Logged when auto flash copy operation fails to correct the backup firmware image
0x0238	Information	"Active firmware image checksum error; backup firmware image activated" }	Logged when backup firmware is activated due to CRC error found on the active firmware image
0x0239	Critical	"%s could not be authenticated as a genuine drive" }	Logged when a drive fails to get authenticated as a genuine drive
0x023A	Critical	"One or more drives were failed or missing during boot" }	Currently not logged
0x023B	Information	"Data found in Write-Back cache during boot" }	Currently not logged
0x023C	Warning	"Incomplete writes on degraded %s due to power loss; check data integrity" }	Currently not logged
0x023D	Warning	"Cannot communicate with feature key; features disabled" }	Currently not logged
0x023E	Information	"Active firmware image flash programming successful" }	Logged when the firmware update on the active firmware region is successful
0x023F	Information	"Backup firmware image flash programming successful" }	Logged when the firmware update on the backup firmware region is successful

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0240	Warning	"Active and backup firmware image versions do not match" }	Logged when both the active and the backup firmware regions mismatch
0x0241	Information	"Backup firmware package version %s" }	Indicates the firmware package version
0x0242	Critical	"Critical error occurred while restoring the offloaded cache data" }	Logged when a critical error occurs while restoring the offloaded cache data
0x0243	Information	"Drive security key from escrow, %s is unlocked" }	Currently not logged
0x0244	Warning	%s temperature (%d C) is above fatal threshold }	Logged when the controller temperature is above the fatal threshold
0x0245	Information	"The controller was reset to recover from a memory access error" }	Logged when the controller resets the firmware to recover from a memory access error
0x0246	Warning	"Running firmware not compatible with profile id %s" }	Logged when the current firmware version is not compatible with the profile ID that is selected by the user

Unsupported Commands in Embedded MegaRAID

The commands in the following table are not supported in Embedded MegaRAID.

Table 62: Unsupported Commands in Embedded MegaRAID

Command Group	Command
ASO	storcli /cx(x all) set aso key=<keyvalue> preview
	storcli /cx(x all) set aso key=<key value>
	storcli /cx(x all) set aso rehostcomplete
	storcli /cx(x all) set aso deactivatetrialkey
	storcli /cx(x all) show safeid
	storcli /cx(x all) show rehostinfo
	storcli /c0 set time =<yyyymmdd hh:mm:ss system>
	storcli /c0 show cc consistencycheck
	storcli /c0/vall show expansion
	storcli /c0 set jbod
	storcli /cx download src=<filepath> [forceActivate]
BBU	storcli /cx/bbu show
	storcli /cx/bbu show all
	storcli /cx/bbu set [learnDelayInterval=<val> bbuMode=<val>
	storcli /cx/bbu start learn
Cache	storcli /cx/v(x all) set ssdcaching=on off
	storcli /cx(x all) show preservedcache
	storcli /cx/v(x all) delete preservedcache[force]
Controller	storcli /cx show cc
Copy back	storcli /cx[/ex]/sx show copyback
	storcli /cx[/ex]/sx start copyback target=eID:sID
	storcli /cx[/ex]/sx stop copyback
	storcli /cx[/ex]/sx pause copyback
	storcli /cx[/ex]/sx resume copyback
Consistency check	storcli /cx show cc/ConsistencyCheck
DS	storcli /cx(x all) set ds=OFF type=1 2 3 4
	storcli /cx(x all) set ds=ON type=1 2 [properties]
	storcli /cx(x all) set ds=ON type=3 4 DefaultLdType=<val> [properties]
	storcli /cx(x all) set ds [properties]
	storcli /cx/v(x all) set ds=Default Auto None Max MaxNoCache

Command Group	Command
Jbod	storcli /c0 set jbod=<on off>
	storcli /c0/s2 set jbod
	storcli /c0/s2 set bootdrive=<on off>
Migrate	storcli /cx/v(x all) show migrate
	storcli /cx/vx start migrate type=raidx [option=add remove drives=[e:]s [e:]s-x [e:]s-x,y] [Force]
Security	storcli /cx delete security key
	storcli /cx set securitykey=xxxxxxxx {passphrase=xxxx} {keyid=xxx}
	storcli /cx set securitykey keyid=xxx
	storcli /cx compare securitykey=xxxxxxxx
	storcli /cx set securitykey=xxxxxxxx oldsecuritykey=xxxxxxxx
Secure erase	storcli /cx/sx secureerase [force]
	storcli /cx/sx start erase [simple normal thorough] [erasepatternA=<val>]
	storcli /cx/sx stop erase
	storcli /cx/sx show erase

CLI Error Messages

This appendix lists the software error messages for the Storage Command Line Tool (StorCLI).

The StorCLI configuration utility is a command line interface application you can use to manage MegaRAID SAS RAID controllers.

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 63: 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

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

Decimal Number	Hex Number	Event Text
71	0x47	Rekey operation not allowed, unless controller already has a lock key
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.

Decimal Number	Hex Number	Event Text
110	0x6e	Signed component is not part of the image
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.

240 Virtual Drive Feature Limitations

This appendix provides information about limitations and known issues for the 240 virtual drives (VDs) feature in the MegaRAID 12Gb/s SAS RAID controller.

Host Software Utility

The following host software utilities support matrix provides the support information on the target IDs that are supported.

Table 64: Host Software Utilities Support Matrix

MegaRAID SAS RAID Utilities	0–63 VD Target IDs Support	240 VD Target IDs Support
StorCLI	Yes	Yes
LSA	Yes	Yes
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 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.

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.

On iMegaRAID controllers, special tasks, such as consistency check, rebuild rate, and background operations will not progress in an EFI environment. However, they still progress in pre-boot environment because you will be rebooting the system while exiting from the applications.

Online Firmware Upgrade and Downgrade

This appendix provides information about known issues when using the online firmware update feature of the MegaRAID 12Gb/s SAS RAID controller.

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:

- Once you have upgraded to any MegaRAID 7.6 firmware or later, downgrading to MegaRAID 7.5 or older version is not allowed. This is due to a minor hardware changes on the internal clock source.
- Once you have upgraded to MegaRAID 7.3 firmware or later, downgrading to MegaRAID 7.2 or MegaRAID 7.1 is not allowed.
- The UNMAP implementation was modified in MegaRAID 7.7 firmware to ensure all SSDs in the VD support UNMAP. When updating a controller that has firmware prior to MegaRAID 7.7 through an online method (only restarting the MegaRAID controller), VDs created on prior firmware where all members are not complaint will generate errors to the `syslog` until the server reboot is perform. After the reboot, the VD will not be reported as supporting UNMAP.

NOTE

UNMAP is supported on SAS SSDs only.

- From MegaRAID 7.x firmware GCA and later, it is recommended that you 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.
- To avoid data loss from the dirty cache on the controller, the utility forces the virtual disks into Write Through mode after a firmware upgrade. The virtual disks remains in this mode until the server reboots. In Write Through mode, the controller sends a data transfer completion signal to the host when the disk subsystem receives all the data in a transaction. This prevents the controller from discarding the dirty cache in a power outage.
- 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 or perform an online controller reset of the system for the flashed firmware to take effect.

Consistency Check, Background Initialization, and Secure EraseCryptographic Erase Limitation

When you downgrade from a 240-virtual drive supported firmware (MR 6.6 and later) to a non-240 virtual drive supported firmware (MR 6.5 and earlier), **Consistency Check**, **Background Initialization**, and **Secure EraseCryptographic Erase** operations are not resumed.

Table 65: Online Firmware Upgrade and Downgrade Support Matrix

Release	Downgrade Support	Upgrade Support
MegaRAID 7.0.1 Firmware	N/A	
MegaRAID 7.1 Firmware	MR 7.01 – MR 7.1	
MegaRAID 7.2 Firmware	MR 7.01 – MR 7.2	
MegaRAID 7.3 Firmware	Within MR 7.3 only	
MegaRAID 7.4 Firmware	MR 7.3 – MR 7.4	

Release	Downgrade Support	Upgrade Support
MegaRAID 7.5 Firmware	MR 7.3 – MR 7.5	
MegaRAID 7.6 Firmware	Within MR 7.6 only	
MegaRAID 7.7 Firmware	MR 7.6 – MR 7.7	
MegaRAID 7.8 Firmware	MR 7.6 – MR 7.7	

Boot Messages and BIOS Error Messages

This appendix provides the boot messages and BIOS error messages present in the MegaRAID firmware.

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 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.


If the controller's boot mode is set to Stop on Errors or Pause on Errors, and if the controller has pinned cache present, you need to correct the problem through HII before booting to the operating system. Until you resolve this problem, the UEFI driver reports the health status as "Configuration Required".

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.

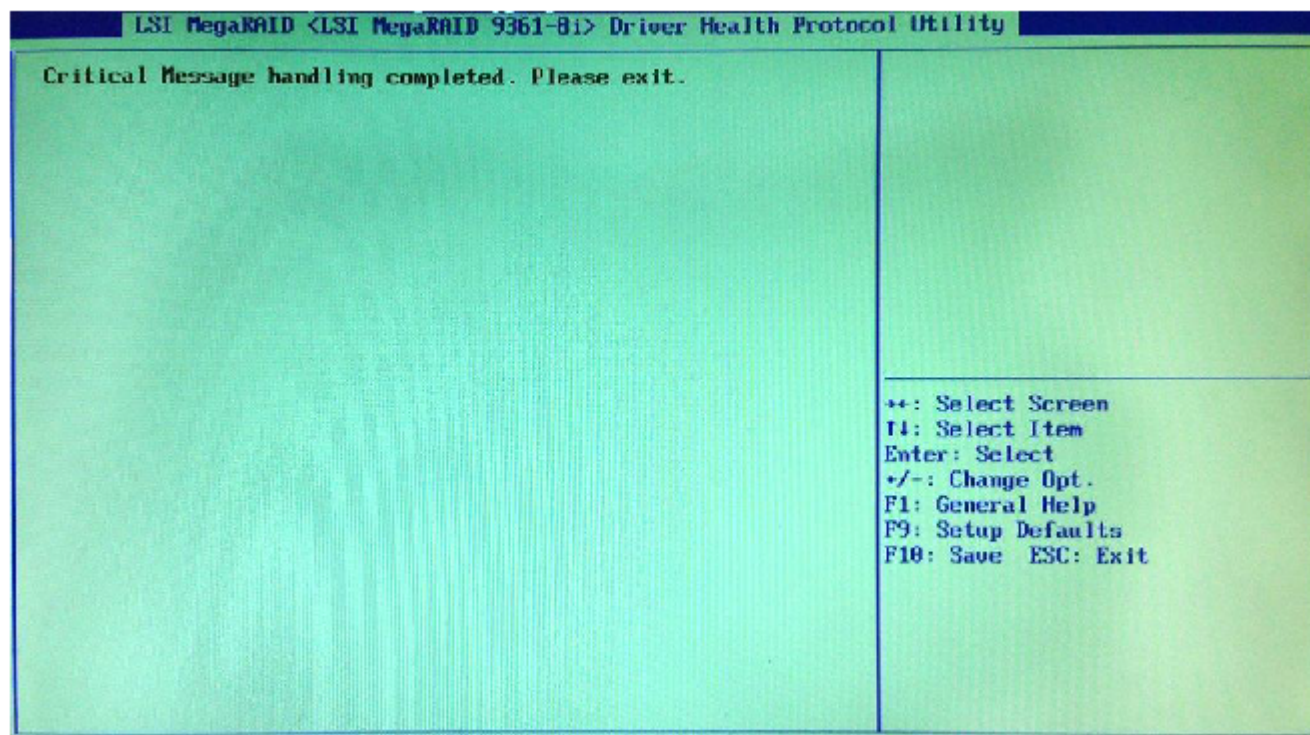
Figure 82: Driver Health Protocol Dialog



6. Press **C**.

The following dialog appears.

Figure 83: Critical Message Completion Dialog

7. Press **Esc** to exit the browser.

The critical message handling completion, the security password, and the confirm message displayed on the screen are all part of the 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.

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 MegaRAID 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.

NOTE

Starting from MegaRAID 7.0, the controller supports HII in an UEFI environment. No preboot utilities are available for legacy environment. If you are booting the system in a legacy mode, then you may see a different message as preboot utilities do not exist for legacy environment. You can resolve this by booting the system in an UEFI mode.

For example, if a preboot utility is supported by the firmware and is present in your firmware package, for boot message type `BOOT_MSG_CACHE_DISCARD`, the boot message displayed on your console may read `Memory or battery problems were detected. The adapter has recovered, but cached data was lost. Press any key to continue..`

If preboot utility is not supported by the firmware, for boot message type `BOOT_MSG_CACHE_DISCARD`, the boot message displayed on your console may read `Memory or battery problems were detected. The adapter has recovered, but cached data was lost. Press any key to continue.`

The only difference here is, with preboot utilities being present, you need to press the `C` key to continue; in the absence of preboot utilities, you need to press any key of your choice to continue.

Table 66: Boot Messages

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
0	<code>BOOT_MSG_CACHE_DISCARD</code>	<code>BOOT_TIME_WAIT</code>	<code>MR_EVT_CTRL_CACHE_DISCARDED</code>	Memory or battery problems were detected. The adapter has recovered, but cached data was lost. Press any key to continue.	—	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.
1	<code>BOOT_MSG_TEST</code>	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.	—	N/A

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
2	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.	—	<p>Causes:</p> <p>The cached data is lost and cannot be retrieved. This boot message is displayed when dirty data needs to be flushed during boot.</p> <p>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.</p> <p>On a single controller, during firmware upgrade, firmware ensures that there is no dirty data.</p>
						<p>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.</p> <p>Action: Ensure that the other controller also has the same firmware version.</p>

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
3	BOOT_MSG_DDF_FOREIGN_FOUND	10	MR_EVT_FOREIGN_CFG_IMPORTED	Foreign configuration(s) found on adapter. Press any key to continue 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.
4	BOOT_MSG_DDF_IMPORT	10	NULL	Previous configuration cleared or missing. Importing configuration created on %02d/%02d %2d: %02d. Press any key to continue.	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.
5	BOOT_MSG_PACKAGE_VERSION	0	MR_EVT_PACKAGE_VERSION	Firmware package: %s	—	N/A
6	BOOT_MSG_FIRMWARE_VERSION	0	NULL	Firmware version: %s	—	N/A
7	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.
8	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.

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
9	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.
10	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.	—	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.
11	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.	—	Cause: This controller does not support SAS drives. Action: Replace the SAS drives with SATA drives and restart the system.
12	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.	—	Cause: This controller does not support SATA drives. Action: Replace the SATA drives with SAS drives and restart the system.

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
13	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.
14	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.
15	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.

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
16	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.
17	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.

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
18	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.	—	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.
19	BOOT_MSG_LDS_OFFLINE	BOOT_TIME_WAIT	MR_EVT_CTRL_BOOT_LDS_WILL_GO_OFFLINE	The following VD's have missing disks: %s. If you proceed (or load the configuration utility), these VD's will be marked OFFLINE and will be inaccessible. Check your cables and make sure all disks are present. Press any key to continue.	—	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.

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
20	BOOT_MSG_LDS_MISS ING	BOOT_TIME_W AIT	MR_EVT_CTRL_BO OT_LDS_MISSING	The following VD's are missing: %s. If you proceed (or load the configuration utility), these VD's 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 VD's should be present, power off your system and check your cables to make sure all disks are present. Press any key to continue.	—	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.
21	BOOT_MSG_LDS_MISS ING_SPANS	BOOT_TIME_W AIT	MR_EVT_CTRL_BO OT_LDS_MISSING	The following VD's are missing complete spans: %s. If you proceed (or load the configuration utility), these VD's 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 VD's. If you believe these VD's should be present, please power off your system and check your cables to make sure all disks are present. Press any key to continue.	—	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.

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
22	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.	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.
23	BOOT_MSG_CACHE_FLUSH_NOT_POSSIBLE	BOOT_TIME_CRITICAL	NULL	The cache contains dirty data, but some VD's 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.

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
24	BOOT_MSG_LDS_WILL_RUN_WRITE_THRU	5	NULL	Your VD's 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 VD's 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.
25	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.
26	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.

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
27	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.
28	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 Avago Generic message	Action: 1. Reseat or replace the DIMM. 2. Restart system. If the problem persists, contact Technical Support.

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
29	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 Avago 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_OVERFLOW_ECC_ERROR	BOOT_TIME_CRITICAL	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.
31	BOOT_MSG_HBA_MULTIBIT_ECC_ERROR	BOOT_TIME_CRITICAL	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.

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
32	BOOT_MSG_HBA_SINGLE_BIT_ECC_ERROR	BOOT_TIME_CRITICAL	MR_EVT_CTRL_MEM_ECC_SINGLE_BIT_CRITICAL 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.
33	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.

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
34	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.
35	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.
36	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.

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
37	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.
38	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.
39	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.

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
40	BOOT_MSG_CTRL_SECURITY_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.
41	BOOT_MSG_CTRL_SECURITY_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.
42	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.
43	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.
44	BOOT_MSG_KEY_EKMS_FAILURE	BOOT_TIME_WAIT	MR_EVT_CTRL_LOCK_KEY_EKMS_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.

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
45	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.
46	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 Avago generic	Action: Check the connection of EKMS, restart the system to re-establish the connection to EKMS.
47	BOOT_MSG_REKEY_TO_EKMS_FAILURE_MERCURY	BOOT_TIME_CRITICAL	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_FAILURE for Avago generic	Action: Check the connection of EKMS, restart the system to re-establish the connection to EKMS.
48	BOOT_MSG_NVDATA_IMAGE_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.

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
49	BOOT_MSG_IR_MR_MIGRATION_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
50	BOOT_MSG_DUAL_BATTERY_PRSENT	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.
51	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.

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
52	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.
53	BOOT_MSG_BAD_SBR_SASADDRESS	BOOT_TIME_CRITICAL	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.
54	BOOT_MSG_INCOMPATIBLE_SECONDARY_IBUTTON	BOOT_TIME_CRITICAL	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.

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
55	BOOT_MSG_CTRL_DOWNGRADE_DETECTED	BOOT_TIME_CRITICAL	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.
56	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.
57	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 disable Battery Warning, 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.

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
58	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.
59	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.
60	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

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
61	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.
62	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.
63	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.
64	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.
65	BOOT_MSG_UPGRADED_I_MR_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
66	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

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
67	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.
68	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.
69	BOOT_MSG_DC_ON_DEGRADED_LD	BOOT_TIME_CRITICAL	Multiple power loss detected with I/O transactions to non optimal VD's.	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.
70	BOOT_MSG_CACHE_ERROR					
71	BOOT_MSG_SUPERCAP_CHARGING_IN_PROGRESS					
72	BOOT_MSG_SUPERCAP_CHARGING_COMPLETED					
73	BOOT_MSG_SUPERCAP_CHARGING_INCOMPLETE					

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
74	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.
75	BOOT_MSG_CACHE_OFFLOAD_DISABLE			Disable cache offload message.		
76	BOOT_MSG_NVCACHE_BACKUP_CAPACITY_OVER			The online backup capacity has exceed the limit.		
77	BOOT_MSG_NVCACHE_CONSIDER_REPLACEMENT			Consider replacing the offline NVcache.		
78	BOOT_MSG_NVCACHE_INVALID			The offline flash is no longer working properly.		
79	BOOT_MSG_SUPERCAP_LEARN_IN_PROGRESS			A supercap learn is in process.		
80	BOOT_MSG_SUPERCAP_LEARN_COMPLETED			The supercap learn is complete.		
81	BOOT_MSG_NVCACHE_COMPONENT_MISMATCH			The NVcache has an incompatible combination of DDR, NAND flash and supercap.		
82	BOOT_MSG_TEMP_THRESHOLD_ERROR			The controller maximum temperature threshold has been exceeded.		
83	BOOT_MSG_FLASH_BAD_WARNING			Warning! Flash memory is bad.		
84	BOOT_MSG_CACHE_RESTORE_ERROR			Warning! A fatal error occurred when restoring the image from flash.		

Message Number	Boot Message Type	Wait Time	Event Log	Boot Message Description	Comments	Troubleshooting Actions
85	BOOT_MSG_NVME_NOT_SUPPORTED			NVMe drivers were discovered, but are not supported.		
86	BOOT_MSG_FW_INCOMPATIBLE_WITH_PROFILE			Firmware is running that is not compatible with the current profile ID.		
87	BOOT_MSG_FIRMWARE_RELEASE_CANDIDATE			Firmware is running that is not compatible with the current profile ID.		
88	BOOT_MSG_NO_DRIVES_SUPPORTED			Drives are not supported because disable SAS is set.		
89	BOOT_MSG_HEADLESS_DUMMY	0	NULL	—	—	N/A
90	BOOT_MSG_LIST_TERMINATOR	0	NULL	—	—	N/A
91	BOOT_MSG_COUNT					

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

B

BBU present	A controller property that indicates whether the controller has an onboard supercapacitors 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 supercapacitors, reported in milliamperes.
Current write policy	<p>A virtual drive property that indicates whether the virtual drive currently supports Write Back mode (write caching enabled) or Write Through mode (write caching disabled).</p> <ul style="list-style-type: none"> • In Write Back mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. • In Write Through mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction.
Cycle count	The count is based on the number of times the near fully charged supercapacitors 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 supercapacitors, measured in milliampere-hour units (mAh).
Design charge capacity remaining	Amount of the charge capacity remaining, relative to the supercapacitors design capacity.
Design voltage	Designed voltage capacity of the supercapacitors, measured in millivolts (mV).
Device ID	A controller or drive property indicating the manufacturer-assigned device ID.
Device port count	A controller property indicating the number of ports on the controller.
Drive cache policy	A virtual drive property indicating whether the virtual drive cache is enabled, disabled, or unchanged from its previous setting.
Drive group	A group of drives attached to a RAID controller on which one or more virtual drives can be created. All virtual drives in the drive group use all of the drives in the drive group.

Drive state		<p>A physical drive or a virtual drive property indicating the status of the appropriate drive.</p> <p>Physical Drive State</p> <p>A physical drive can be in any one of the following states:</p> <ul style="list-style-type: none"> • Unconfigured Good – A drive accessible to the RAID controller but not configured as a part of a virtual drive or as a hot spare. 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. <p>Virtual Drive State</p> <p>A virtual drive can be in any one of the following states:</p> <ul style="list-style-type: none"> • 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 one 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 supercapacitors at the current charge rate.
Expected margin of error		Indicates how accurate the reported supercapacitors 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. The 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 supercapacitors. This value represents the last measured full discharge of the supercapacitors. This value is updated on each learn cycle when the supercapacitors undergo a qualified discharge from nearly full to a low level.
G	
Gas gauge status	Hexadecimal value that represents the status flag bits in the gas gauge status register.
H	
Hole	In LSI Storage Authority, a <i>hole</i> 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, <i>PCIX</i> .
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	<p>The process of making a redundant virtual drive consistent. Foreground initialization writes zeros to the data fields, erasing all existing data.</p> <p>Background Initialization (BGI) makes a virtual drive redundant by reading the other drives in the VD, calculating parity, and writing it to the drives. BGI does not erase user data. A user can use the VD while BGI is active.</p>
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.
NVMe	Acronym for nonvolatile memory express. NVMe is a logical device interface specification for accessing NVM storage media attached by means of a PCI Express (PCIe) bus, which removes SCSI from the I/O stack.
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. The 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.
Reversible 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	Degree of heat present in the supercapacitors, 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	<p>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.</p> <p>Write-back cache is used when write caching is enabled.</p>
Write policy	See <i>Default Write Policy</i> .
Write-through	<p>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.</p> <p>Write-through cache is used when write caching is disabled.</p>

Revision History

Version 1.8, September 9, 2020

The following changes were made:

- Updated [RAID 60 Drive Groups](#).
- Updated [Modifying SnapDump Properties Command](#).
- Updated [Clearing SnapDump Data Commands](#).
- Updated 0x005E, 0x0060, and 0x0071 events in [Table 61: Event Messages](#).
- Updated `storcli /cx set autoconfig` in [Automated Physical Drive Configurations](#).
- Added `forceclose` input option to [Flashing Controller Firmware Command while the Firmware Is Operational](#).
- Added new `sanitize` command in [Drive Sanitize Command](#).
- Added virtual drive information to [Previewing and Importing a Foreign Configuration](#).
- Added new `SPDM` commands to [SPDM Commands](#).
- Added new `Drive Performance Monitoring` command to [Drive Performance Monitoring Commands](#).
- Update `storcli/cx/cv show all` command in [CacheVault Commands](#).
- Update Virtual Drive Status table in [Virtual Drive States](#).

Version 1.7, March 31, 2020

The following changes were made:

- Removed JBOD object identifier from [Table 38, Object Identifiers in the StorCLI Command Syntax](#).
- Changed `personality` commands to `autoconfig` commands in [Show and Set Controller Properties Commands](#).
- Changed `personality` commands to `autoconfig` commands and added R0 option to [Table 41, Properties for Show and Set Commands](#).
- Added [SnapDump Commands](#).
- Updated `autoconfig` commands in [Automated Physical Drive Configurations](#).

Version 1.6, February 16, 2018

The following changes were made:

- Added a note regarding PDs to [Table 27, Profile Management Dialog Details](#) and [Profile Management](#).

Version 1.5, February 16, 2018

The following changes were made:

- Updated [Managing Profiles](#)
- Updated [Add Virtual Drives Commands](#).
- Updated [Virtual Drive Erase Commands](#)
- Updated [RAID Configuration Commands](#).

Version 1.4, November 30, 2017

The following changes were made:

- [Profile Management](#).
- [Manually Creating a Virtual Drive](#).
- [Viewing Advanced Controller Management Options](#).
- [Changing Security Settings](#).
- [Managing Profiles](#).
- [Adding Drives to a Configuration](#).

Version 1.3, September 11, 2017

The following changes were made:

- Added [Creating a RAID 10 Volume from the Database](#).
- Updated [Managing Profiles](#).
- Added [Downgrading the Firmware When Profiles Are Selected](#).
- Updated [Manually Creating a Virtual Drive](#) with [Table 22, Emulation Settings](#).
- Updated [Preserved Cache Commands](#).

Version 1.2, June 21, 2017

The following changes were made:

- Updated [HII Dashboard View](#).
- Updated [Displaying Boot Messages](#).

Version 1.1, March 24, 2017

The following changes were made:

- Updated [Viewing Advanced Controller Management Options](#).
- Updated [Managing SAS Storage Link Speeds](#).
- Added [Managing PCIe Storage Interface](#).
- Updated [Viewing Advanced Drive Properties](#).
- Added [PCIe Storage Interface Commands](#), and [Lane Speed Commands](#).

Preliminary, Version 1.0, October 28, 2016

Initial document release.

