



MegaRAID™ 8 Tri-Mode Software

**User Guide
Version 1.3**

Table of Contents

Overview	8
Broadcom 9600 Series Features	8
Tri-Mode Technology	8
Serial-Attached SCSI Device Interface	9
Serial ATA III Features	9
Nonvolatile Memory Express Technology	10
Configuration Scenarios	10
Technical Support	10
Snapdump Feature.....	11
Introduction to RAID	12
Components and Features	12
Drive Group.....	12
Virtual Drive.....	12
Fault Tolerance.....	13
Consistency Check.....	14
Replace.....	14
Background Initialization.....	15
Patrol Read.....	15
Disk Striping.....	15
Disk Mirroring.....	16
Parity.....	16
Disk Spanning.....	17
Hot Spares.....	18
Disk Rebuilds.....	19
Rebuild Rate.....	19
Hot Swap.....	20
Drive States.....	20
Virtual Drive States.....	20
Enclosure Management.....	21
Solid State Drive Features.....	21
SSD Guard.....	21
RAID Levels	21
Summary of RAID Levels.....	21
Selecting a RAID Level.....	22
RAID 0 Drive Groups.....	22
RAID 1 Drive Groups.....	23

RAID 5 Drive Groups.....	24
RAID 6 Drive Groups.....	25
RAID 10 Drive Groups.....	26
RAID 50 Drive Groups.....	27
RAID 60 Drive Groups.....	28
RAID Configuration Strategies.....	29
Maximizing Fault Tolerance.....	29
Maximizing Performance.....	30
Maximizing Storage Capacity.....	31
Configuration Planning.....	32
Number of Drives.....	32
RAID Availability.....	33
Drive Autoconfiguration.....	34
SafeStore Disk Encryption.....	35
Workflow.....	36
Enable Security.....	36
Change Security.....	36
Create Secure Virtual Drives.....	37
Import a Foreign Configuration.....	37
Instant Secure Erase.....	38
Managing Unconfigured Secure Drives.....	38
HII Configuration Utility.....	39
Behavior of HII.....	39
UEFI Support.....	39
Blocking Boot Events.....	39
Starting the HII Configuration Utility.....	42
HII Dashboard View.....	43
Main Menu.....	43
HELP.....	44
PROPERTIES.....	45
ACTIONS.....	46
BACKGROUND OPERATIONS.....	46
MegaRAID ADVANCED SOFTWARE OPTIONS.....	47
Managing Configurations.....	47
Creating a Virtual Drive from a Profile.....	48
Creating a RAID 10 Volume from the Database.....	52
Manually Creating a Virtual Drive.....	53
Viewing Drive Group Properties.....	57
Viewing Global Hot Spare Drives.....	58

Clearing a Configuration.....	58
Convert to Unconfigured, Convert to JBOD, and Enable Security.....	59
Convert to Unconfigured.....	59
Convert to JBOD.....	60
Enabling Security on a Controller.....	60
Managing Foreign Configurations.....	61
Viewing and Importing a Foreign Configuration.....	61
Clearing a Foreign Configuration.....	63
Managing Controllers.....	63
Viewing Advanced Controller Management Options.....	65
Viewing Advanced Controller Properties.....	66
Managing MegaRAID Advanced Software Options.....	69
Displaying the Controller Personality.....	69
Saving or Clearing Persistent Events.....	70
Enabling or Disabling Security.....	71
Changing Security Settings.....	75
Perform Cryptographic Erase on Drives.....	77
Managing Snapdump.....	78
Managing SAS Storage Link Speed.....	79
Managing PCIe Storage Interface.....	79
Setting Cache and Memory Properties.....	80
Running a Patrol Read.....	80
Setting Emergency Spare Properties.....	81
Changing Task Rates.....	83
Managing Virtual Drives.....	84
Selecting Virtual Drive Operations.....	86
Locating Physical Drives in a Virtual Drive.....	87
Deleting a Virtual Drive.....	88
Initializing a Virtual Drive.....	88
Erasing a Virtual Drive.....	88
Securing a Virtual Drive.....	89
Running a Consistency Check.....	89
Viewing Associated Drives.....	90
Viewing and Managing Virtual Drive Properties and Options.....	90
Managing Devices.....	92
Viewing Physical Drive Properties.....	93
Performing Drive Operations.....	94
Locating a Drive.....	95
Making a Drive Unconfigured Bad, Unconfigured Good, or JBOD.....	95
Enabling Security on JBOD.....	96

Replacing a Drive.....	96
Make Offline.....	97
Make Online.....	97
Mark Missing.....	98
Replacing a Missing Drive.....	98
Assigning a Global Hot Spare.....	99
Assigning a Dedicated Hot Spare.....	99
Unassigning a Hot Spare Drive.....	99
Initializing or Erasing a Drive.....	100
Rebuilding a Drive.....	101
Securely Erasing a Drive.....	101
Removing a Physical Drive.....	102
Making a JBOD.....	102
Viewing Advanced Drive Properties.....	102
Logical Unit/Namespace Information.....	103
Managing Energy Packs.....	105
HII Popup Error Protocol.....	106
StorCLI2 Utility.....	107
Supported Controllers and Operating Systems.....	107
Supported Controllers.....	107
Supported Operating Systems.....	107
StorCLI2 Default Logging.....	108
Installing StorCLI2 on MegaRAID8.....	109
Installing the StorCLI2 Tool on Microsoft Windows Operating Systems.....	109
Installing the StorCLI2 Tool on the UEFI Environment.....	110
Installing the StorCLI2 Tool on Linux Operating Systems.....	110
Uninstalling the StorCLI2 Tool on Linux Operating Systems.....	110
Installing the StorCLI2 Tool on VMware Operating Systems.....	110
Uninstalling the StorCLI2 Tool on VMware Operating Systems.....	110
StorCLI2 Commands.....	110
StorCLI2 Tool Command Syntax.....	111
System Commands.....	113
System Show Commands.....	113
Controller Help Commands.....	113
Controller Commands.....	114
Controller Show Commands.....	114
Show and Set Controller Properties Commands.....	115
Controller Background Task Operation Commands.....	122
Premium Feature Key Commands.....	127
Controller Security Commands.....	127

Flashing Controller Firmware.....	129
Snapdump Commands.....	130
Predictive Failure Monitoring Commands.....	132
Controller Replacedrive Commands.....	133
Drive Performance Monitoring Commands.....	134
SPDM Commands.....	135
Physical Drive Commands.....	136
Drive Show Commands.....	136
Missing Drives Commands.....	137
Set Drive State Commands.....	137
Drive Initialization Commands.....	139
Drive Firmware Download Commands.....	140
Locate Drives Commands.....	140
Prepare to Remove Drives Commands.....	141
Drive Security Command.....	141
Drive Secure Erase Commands.....	141
Rebuild Drives Commands.....	143
Hot Spare Drive Commands.....	144
NVMe Drive Commands.....	145
Replacedrive Commands.....	145
Spinup Drive Commands.....	146
Virtual Drive Commands.....	147
Add Virtual Drives Commands.....	147
Delete Virtual Drives Commands.....	149
Virtual Drive Show Commands.....	150
Preserved Cache Commands.....	150
Change Virtual Drive Properties Commands.....	151
Virtual Drive Initialization Commands.....	152
Virtual Drive Erase Commands.....	152
Virtual Drive Consistency Check Commands.....	153
Background Initialization Commands.....	154
Virtual Drive Expansion Commands.....	155
Foreign Configuration Commands.....	156
Drive Group Commands.....	157
Drive Group Commands.....	157
Controller Power Savings Commands.....	157
Enclosure Commands.....	158
Controller Phy Commands.....	159
Energy Pack Commands.....	160
PCIe Storage Interface Commands.....	161

Logging Commands.....	161
Automated Physical Drive Configurations.....	161
Frequently Used Tasks.....	162
Displaying the Version of the StorCLI2 Utility.....	163
Displaying the StorCLI2 Utility Help.....	163
Displaying System Summary Information.....	163
Displaying Free Space in a Controller.....	163
Adding Virtual Drives.....	163
Setting the Cache Policy in a Virtual Drive.....	164
Displaying Virtual Drive Information.....	164
Deleting Virtual Drives.....	165
Flashing Controller Firmware.....	165
SAS Address Assignment Rule.....	165
StorCLI to StorCLI2 Command Conversion.....	165
Events, Messages, and Behaviors.....	180
Error Levels.....	180
Event Messages.....	180
Glossary.....	203
Revision History.....	211
Documentation Legal Notice.....	215

Overview

This guide describes how to use the Storage Command Line Interface2 (StorCLI2) tool and the MegaRAID Human Interface Infrastructure (HII) configuration utility.

This section provides an overview of this guide, which documents the utilities that are used to configure, monitor, and maintain Tri-Mode MegaRAID™8 Serial-attached SCSI (SAS) RAID controllers with RAID control capabilities and the storage-related devices that are connected to them.

Broadcom 9600 series adapters utilize a new, unified driver that streamlines integration for users of 9600 adapters compared to prior generations that used separate drivers for Broadcom® HBAs and MegaRAID adapters.

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

Broadcom 9600 Series Features

Broadcom 9600 series adapters have a wide range of new software features compared to 9500 and 9400 Tri-Mode Storage HBAs and MegaRAID adapters. This section provides an overview on the new features, including improvements to software, drivers, and utilities.

MegaRAID has streamlined virtual drive parameters for RAID volumes. The parameters reduce stripe sizes to 64K and 256K. Reduced stripe size allows users to configure virtual drives while maintaining performance, and supporting a single stripe size and read ahead policy (no read ahead) for each virtual drive within a drive group. This maintains optimal performance for the MegaRAID solution without confusing or complicating controller management.

Previous MegaRAID products provided users manageability for latency and I/O size that required either deep knowledge of the data set, or performance testing and ongoing evaluation of the storage environment. The 9600 family of products makes these decisions for the user, reducing the need for performance tuning.

The 9600 series adapters are capable of delivering performance for all supported device types and features with no additional user input required.

The 9600 MegaRAID solutions support up to 240-single drive RAID 0 volumes per controller when the RAID volume uses the available free space on the disk. When used in a more complex configuration (for example, multiple drives or RAID 1/5/6) up to 64 virtual drives per controller are supported.

RAID volumes that were created on previous products cannot be imported to a 9600 RAID adapter. If data migration is required, the RAID volume must be recreated on the 9600 RAID adapter. Virtual drive migration from one 9600 RAID adapter to another 9600 RAID adapter automatically is supported.

Tri-Mode Technology

The MegaRAID8 Tri-Mode RAID controllers are high-performance intelligent SAS/SATA/PCIe (NVMe) devices with RAID control capabilities. The MegaRAID8 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 MegaRAID8 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 using of SAS devices, Serial ATA (SATA) II and 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. MegaRAID8 gives you the flexibility to combine these three similar technologies on the same controller, within the same enclosure, and in the same virtual drive.

The MegaRAID8 Tri-Mode RAID controllers are based on the Broadcom first-to-market SAS IC technology and proven MegaRAID8 technology. As third-generation PCI Express RAID controllers, the MegaRAID8 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 MegaRAID8 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 that is 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.

The Tri-Mode interface provides the following data transfer rates.

- SAS – 22.5Gb/s, 12Gb/s, and 6Gb/s per lane
- SATA – 6Gb/s per lane
- PCIe – up to 16.0 GT/s per lane and is backward compatible to 5.0/2.5 GT/s

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 the 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 that is attached by 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 that are created by the legacy ATA primary-secondary 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 6Gb/s
- Supports STP data transfers of 24Gb/s
- Provides a serial, point-to-point storage interface
- Simplifies cabling between devices
- Eliminates the primary-secondary construction that is 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 that are attached by a PCI Express (PCIe) bus, which removes SCSI from the I/O stack. By its design, NVMe allows the 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

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

Technical Support

For assistance with installing, configuring, or running your Tri-Mode MegaRAID8 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 section describes a Redundant Array of Independent Disks (RAID), RAID functions and benefits, RAID components, RAID levels, and configuration strategies.

In addition, this section 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.

Drives must be organized into virtual drives in a drive group. Drives must also 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
- Using Online Capacity Expansion to increase the capacity of a virtual drive

Components and Features

RAID levels describe a system for ensuring the availability and redundancy of the data that is 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 that are 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

Table 1: MegaRAID8 Array Limitations

Description	Feature eHBA	RAID eHBA	MegaRAID Adapter
Specification	9600-16i 9600-24i 9600-16e 9600W-16e	9620-16i	9660-16i 9670W-16i 9670-24i
Maximum SAS/SATA drives per controller	240	32	240
Maximum SAS/SATA enclosures per controller	20	4	20
Maximum SAS/SATA enclosures per port	10	1	10
Maximum SAS/SATA drives per enclosure	64	32	64
Maximum NVMe drives per controller	32	32	32
Maximum NVMe enclosures per controller	2	2	2
Maximum NVMe enclosures per port	1	1	1
Maximum NVMe drives per enclosure	32	32	32
Maximum drives per drive group	N/A	32	32
Maximum virtual drives per drive group	N/A	4	16
Maximum spans per virtual drive	N/A	8	8

Fault Tolerance

Fault tolerance is the capability of the subsystem to undergo a drive failure or failures without compromising the 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 a drive failure in a drive group, though performance can be degraded to some extent.

In a RAID 1 drive group, each RAID 1 group has two drives and can tolerate one drive failure. MegaRAID8 supports RAID 1 drive groups that can contain up to 32 drives, and can tolerate up to 16 drive failures (one in each pair within the 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 per group.

Each spanned RAID 10 virtual drive can tolerate multiple drive failures, if each failure is in a separate drive pair (one drive in each pair within each group). A RAID 50 virtual drive can tolerate one drive failure per drive group up to 8 failed drives, if each failure is in a separate drive group. RAID 60 drive groups can tolerate up to two drive failures in each drive group, up to 16 failed drives for an 8 span RAID 60 configuration.

NOTE

RAID level 0 is not fault tolerant. If a drive in a RAID 0 drive group fails, the entire virtual drive (all drives that are 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 if there is 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 is automatically rebuilt by hot-swapping the drive in the same drive bay. The RAID drive group continues to handle requests while the rebuild occurs.

Consistency Check

The consistency check operation verifies that the data is correct 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

You should 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, the data is rebuilt on a hot spare. Once the failed drive is replaced with a new disk, the data is copied from the hot spare to the new drive. The hot spare then reverts from a rebuild drive to its original hot spare status.

A Replace operation is also initiated when a SMART error occurs on a drive that is part of a virtual drive, as a drive with a SMART error is expected to fail. 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 a Degraded status.

The Replace operation runs as a background activity, and the virtual drive is still available online to the host.

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.

Order of Precedence

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

- If a Replace operation is already taking place on 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. Background initialization 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.

The default and recommended background initialization rate is determined by the NVDATA parameter.

Patrol Read

Patrol read involves the review of your system for possible drive errors that could lead to a 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. Patrol read 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. These stripes are interleaved in a repeated sequential manner. The combined storage space is composed of stripes from each drive. Use 64k for drive groups consisting of SSDs and 256k for drive groups consisting of HDDs.

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

Figure 1: Example of Disk Striping (RAID 0)



Stripe Width

Stripe width is the number of drives that are 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. Supported stripe sizes are 64k for drive groups consisting of SSDs. Drive groups consisting of HDDs have stripe sizes of 64k and 256k.

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 that is 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 written to a second disk, data is not lost if one disk fails. In addition, both drives always contain the same data, so either disk can act as the operational disk. If one disk fails, the contents of the other disk can run the system and can reconstruct the failed disk.

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

Figure 2: Example of Disk Mirroring (RAID 1)



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 if a drive failure occurs. 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 drives in a drive group. The types of parity are described in the following table.

Table 2: Types of Parity

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

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

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



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


3_01081-00

Disk Spanning

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

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

Figure 4: Example of Disk Spanning



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

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

The following table describes how to configure RAID 10, RAID 50, and RAID 60 drive groups by spanning. The virtual drives must have the same drive count in each drive group and the same stripe size. 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 3: Spanning for RAID Drive Groups

Level	Description
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. 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. A RAID 50 drive group supports a maximum of 8 spans.
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. A RAID 60 drive group supports a maximum of 8 spans.

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. A hot spare is usually in Standby mode, ready for service if a drive fails. Hot spares let you replace failed drives without a system shutdown or user intervention. The MegaRAID8 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 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.

Global hot spares can be created without first creating a logical drive. If all logical drives are deleted, global hot spares become unconfigured good.

NOTE

SAS/SATA drives are used for drive groups consisting of SAS/SATA devices. NVMe drives are used for drive groups consisting of NVMe devices.

Dedicated Hot Spare

Use a dedicated hot spare to replace a failed drive only in a selected drive group. Dedicated hot spares are assigned to work with one drive group or spanned 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.

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 must 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.
- A dedicated hot spare becomes a global hot spare if all the logical drives in the drive group that the hot spare is dedicated to are deleted (the drive group is deleted).

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 that is 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.

Rebuild Rate

The rebuild rate is the percentage of the compute cycles that are 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 1 percent and 100 percent. At 1 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 1 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 when one of these situations 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 that 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 4: 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.
Replacedrive	A drive that has replaced the failed drive in the RAID configuration.
Unsupported	A drive that is not supported in the RAID configuration.
Unusable	A drive that is not usable in the RAID configuration.

Virtual Drive States

The virtual drive states are described in the following table.

Table 5: 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. A RAID 6 drive group does not move to degraded until two drive failures occur.
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.

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.

Solid State Drive Features

The MegaRAID8 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 MegaRAID8, 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 MegaRAID8 software starts a rebuild to preserve the data on the SSD and sends appropriate warning event notifications.

RAID Levels

The RAID controller supports RAID levels 0, 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 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 four drives, is similar to a RAID 5 drive group. Blocks of data and parity information are written across all drives. If one or two drives fail in the drive group, the parity information is used to recover the data.

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

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 MPI3MR controller supports the standard RAID levels – RAID 0, RAID 1, RAID 5, and RAID 10. The MPI3MR 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 MPI3MR controller. One virtual drive can be created on an array (a maximum of eight if no other virtual drives are already created on the MPI3MR 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 several 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 that are 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. If any drive fails, all data is lost.
Drives	1 to 32

Figure 5: 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. 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

Figure 6: 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 numerous 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	<p>Provides high data throughput, especially for large files.</p> <p>Use RAID 5 drive groups for transaction processing applications because each drive can read and write independently.</p> <p>If a drive fails, the RAID controller uses the parity drive to re-create all missing information.</p> <p>Online customer service that requires fault tolerance.</p> <p>Any application that has high read request rates but random write request rates.</p>
Strong points	<p>Provides data redundancy, high read rates, and good performance in most environments.</p> <p>Provides redundancy with lowest loss of capacity.</p>
Weak points	<p>Not well suited to tasks requiring lots of small writes or small block write operations.</p> <p>Suffers more impact if no cache is used.</p> <p>If a drive is being rebuilt, drive performance is reduced.</p> <p>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.</p>
Drives	3 through 32

Figure 7: 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 by using a second parity block in each stripe. Use a RAID 6 drive group for data that requires a high level of protection from loss.

If there is a failure of one drive or two drives in a virtual drive, the RAID controller uses the parity blocks to re-create 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 any 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 numerous 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 that is required by using two parity blocks per stripe.
Drives	4 through 32.

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

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



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

3 0108600

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 that are 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 that are supported by RAID 10 virtual drives.

For 2-32 drives, a single span R1 is required. Multiple span RAID 10 configurations are used for 36 drives and larger configurations.


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

Table 10: RAID 10 Drive Group Overview

Uses	Appropriate when used with data storage that needs 100 percent redundancy of mirrored drive groups and that also needs the enhanced I/O performance of RAID 0 (striped drive groups.) A RAID 10 drive group works well for medium-sized databases or any environment that requires a higher degree of fault tolerance and moderate-to-medium capacity.
Strong Points	Provides both high data transfer rates and complete data redundancy.
Weak Points	Requires twice as many drives as all other RAID levels except in RAID 1 drive groups.
Drives	4 to 240 – The number of drives in each span must be an even number, and the same number of drivers in each span. 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 9: 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 can tolerate up to eight drive failures, though less than total drive capacity is available. Though multiple drive failures can be tolerated, only one drive failure can be tolerated in each RAID 5 level drive group.

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

Table 11: RAID 50 Drive Group Overview

Uses	Appropriate when used with data that requires high reliability, high request rates, high data transfer, and medium-to-large capacity. Also used when a virtual drive of greater than 32 drives is needed.
Strong points	Provides high data throughput, data redundancy, and good performance.
Weak points	Requires two times to eight times as many parity drives as a RAID 5 drive group.
Drives	A minimum of two spans of three drives per span, up to eight spans of 3 to 32 drives per span (limited by the maximum number of drives supported by the controller). Each span must contain the same number of drives. Eight spans of RAID 5 drive groups that contain 6 through 32 drives each (limited by the maximum number of devices that are supported by the controller)

Figure 10: RAID 50 Level Virtual Drive



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

Table 12: RAID 60 Drive Group Overview

Uses	<p>Provides a high level of data protection by using a second parity block in each stripe. Use a RAID 60 drive group for data that requires a high level of protection from loss.</p> <p>If there is 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 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>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 any 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 that is required by using two parity blocks per stripe.</p>
Drives	<p>Eight spans of RAID 6 drive groups that contain 8 through 32 drives each (limited by the maximum number of devices that are supported by the controller).</p>

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 11: RAID 60 Level Virtual Drive



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

3_01090-00

RAID Configuration Strategies

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

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

You cannot configure a virtual drive that optimizes all three factors. However, 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. 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 13: RAID Levels and Fault Tolerance

RAID Level	Fault Tolerance
0	Does not provide fault tolerance. If any drive fails, all data is lost. 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 written to a second drive, no data is lost if one of the drives fails. Both drives always contain the same data. 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.
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 that is offered by mirrored drive groups. A RAID 10 drive group can sustain a drive failure in each mirrored drive group and can maintain data integrity.
50	Provides data redundancy using distributed parity across spanned RAID 5 drive groups. A RAID 50 drive group includes both parity and disk striping across multiple drives. If a drive fails, the RAID controller uses the parity data to re-create all missing information. A RAID 50 drive group can sustain one drive failure per RAID 5 drive group and still maintain data integrity.
60	Provides data redundancy using distributed parity across spanned RAID 6 drive groups. A RAID 60 drive group can sustain two drive failures per RAID 6 drive group and still maintain data integrity. It provides the highest level of protection against drive failures of all of the RAID levels. A RAID 60 drive group includes both parity and disk striping across multiple drives. If a drive fails, the RAID controller uses the parity data to re-create all missing information.

Maximizing Performance

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

Table 14: RAID Levels and Performance

RAID Level	Performance
0	RAID 0 (striping) offers excellent performance. RAID 0 breaks up data into smaller blocks and then writes a block to each drive in the drive group. Disk striping writes data across multiple drives instead of just one drive. Disk striping involves partitioning each drive storage space into stripes. Virtual drives use 64 KB and 256 KB stripes. The default recommended stripe size for SSDs is 64 KB and 256 KB for HDDs based volumes. These stripes are interleaved in a repeated sequential manner. Disk striping enhances performance because multiple drives are accessed simultaneously.
1	With a RAID 1 (mirroring) drive group, each drive in the system must be duplicated, which requires more time and resources than striping. Performance is impaired during drive Rebuild operations.

RAID Level	Performance
5	<p>A RAID 5 drive group provides high data throughput, especially for large files. Use RAID 5 for any application that requires high read request rates, but low write request rates. For example, 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 good performance. However, a RAID 6 drive group is not well suited to tasks requiring numerous 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>
10	<p>A RAID 10 drive group works best for data storage that needs 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 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 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 numerous 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 15: 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 that is 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.</p> <p>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.</p> <p>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>
6	<p>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 6 drive group more expensive to implement.</p> <p>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.</p>
10	<p>A RAID 10 drive group requires twice as many drives as all other RAID levels except RAID level 1.</p> <p>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.</p> <p>Disk spanning allows multiple drives to function like one large drive. Spanning overcomes a lack of disk space and simplifies storage management by combining existing resources or adding relatively inexpensive resources.</p>
50	<p>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.</p>
60	<p>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.</p>

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 that is stored in the disk subsystem has a different frequency of read and write activity. If you know the data access requirements, you can 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 that is 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 the following questions, 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 that is stored in this drive group contain large audio or video files that must be available on demand?
Use RAID 0.
- Will this drive group contain data from an imaging system?
Use RAID 0 or RAID 10.

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

Table 16: Factors to Consider for Drive Group Configuration

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

RAID Availability

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 that is 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. The substitution can be performed while the subsystem is running (performing its normal functions). The backplane and enclosure must support hot swapping 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 is marked as *failed*. If the source drive fails, both the source drive and the hot spare drive are 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.

Drive Autoconfiguration

Autoconfiguration feature simplifies the configuration creation workflow. A newly inserted drive is automatically configured based on the auto configuration option that is selected by the user. The Auto-Configuration feature allows the following configuration option.

- **UGOOD** – The newly inserted drive is configured as an unconfigured good drive. This drive can be used as part of a configuration creation later.
- **JBOD** – The newly inserted drive is configured as JBOD.
- **SecureJBOD** – If the newly inserted drive is SED capable, it is configured as secured JBOD. If the drive is not SED capable, then it is configured as a JBOD.
- **R0** – The newly inserted drive is configured as single drive Raid 0, write through volume.
- **SecureR0** – The newly inserted drive is configured as single drive Raid 0, write through volume. If the drive is SED capable, then security/encryption is enabled on the volume.
- **R0WB** – The newly inserted drive is configured as single drive Raid 0, write back volume.
- **SecureR0WB** – The newly inserted drive is configured as single drive Raid 0, write back volume. If the drive is SED capable, then security/encryption is enabled on the volume.

NOTE

The auto configuration option is not applied to drives that are already known to the adapter.

For example, if the auto configure option is set to JBOD and an unconfigured good drive is removed and inserted back, the drive remains as unconfigured good. The drive does not change because it is already known to the adapter and the auto-configure option is not applied to the drive.

SafeStore Disk Encryption

This section 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 support local key management.

Overview

The SafeStore Disk Encryption (SED) service offers the ability to encrypt data on SED supported drives and use disk-based key management to provide data security. This solution provides data protection if there is 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 LSI Storage Authority (LSA) software offers a procedure that you can use to manage the security settings for the drives. For more information see the *LSA LSI® Storage Authority Software User Guide*.

Purpose and Benefits

Security is a growing market concern and requirement. MegaRAID8 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, 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 that are related to the SafeStore encryption feature.

Table 17: 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.
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.
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 the controller to boot and unlock access to the user configuration without user intervention.

Workflow

Overview

The SafeStore workflow follows:

- Use a compatible SED drive.
- Enable encryption when the virtual drive is created with the SED drives.
- Create a security key that conforms to the security requirements.
- Configure the system with the desired passphrase.

After the system is booted, you have the option requiring a passphrase to access the virtual drives (see [Managing Foreign Configurations](#)).

Enable Security

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

You can perform three procedures 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 you can 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 must 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 StorCLI2) 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 cannot be accessed. A total of three failed attempts are allowed on a current boot (including boot time and runtime attempts). After the maximum allowed attempts, reboot the system and try again. 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 must 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 can change the security key identifier, security key, and password. If you have previously removed any secured drives, you still must 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 can edit the security key identifier. If you plan to change the security key, change the security key identifier as well. Otherwise, you will be unable to differentiate between the security keys.

You can choose to keep the current security key identifier or enter a new key identifier. 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 enter the new security key or can accept the security key that the controller suggests.

Add or Change the Password

You can add a password or can 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 can 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 can import a locked configuration. You can import unsecured or unlocked configurations when security is disabled. A foreign configuration is a RAID configuration that exists on a replacement set of drives that you install in a computer system. The LSA, StorCLI2, 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 the 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 that is used to secure them. A total of three failed attempts are allowed on a current boot (including boot time and run time attempts). After the maximum number of allowed attempts, reboot the system and try again. Verify whether any drives

are left to import as the locked drives can use different security keys. If there are any drives that are left, repeat the import process for the remaining drives. After the drives are imported, there is no configuration to import.

Instant Secure Erase

Instant Secure Erase is a feature that is 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 MegaRAID8 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.

Managing Unconfigured Secure Drives

- **Delete security on drive** – To delete security on unconfigured (secure) drives, the drive needs to be reprovisioned.
- **Creating Non-Secure Ld on secure unconfigured drive** – If an unconfigured drive is secure, then non-secure logical drives cannot be created using the drive.
- **Delete security** – Security on all unconfigured secured drives has to be deleted before deleting controller security.

HII Configuration Utility

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

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 have no knowledge about independent hardware vendor features and their screen controls.

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 that you specified to the controller. To save the specified values, you must use the controls present in the form or screen that is provided by your independent hardware vendor.

Similarly, when you want to load controller defaults, you can achieve this by clicking the **Set Factory Defaults** option present on the **Dashboard View** menu. You can also click the **Controller Management > Advanced Controller Management > Set Factory Defaults** menu. Pressing **F3** (Optimized Defaults) will not restore the controller defaults.

UEFI Support

UEFI provides MegaRAID8 customers with expanded platform support. The MegaRAID8 UEFI driver, a boot service device driver, handles block I/O requests and SCSI pass-through (SPT) commands, and offers a pre-boot MegaRAID8 management application.

Blocking Boot Events

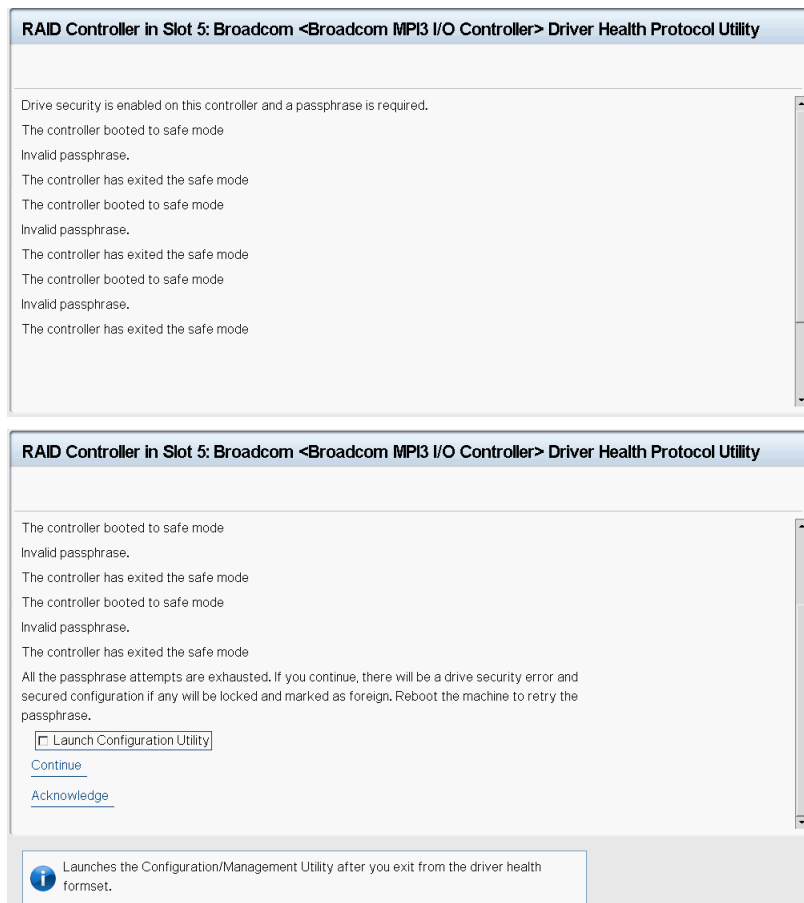
Two error modes are supported during the boot process, continue on error (COE) and safe mode on error (SME). The OEM can choose their default mode through NVDATA. If the firmware supports it, in some instances, you can override the default setting using HII or CLI.

If the controller is configured for COE, then the controller firmware attempts to boot regardless of the type of error that is encountered. In this mode, the controller firmware performs the default action for each boot message (continue, ignore, bypass) and attempts to complete the controller initialization. When a critical condition cannot be bypassed, the firmware boots in safe mode.

If the controller is configured for SME, then the controller boots in normal mode even though the configured mode is safe mode. In SME, the firmware boots without bringing the configuration online, which allows the system to boot, except when the controller is the boot device.

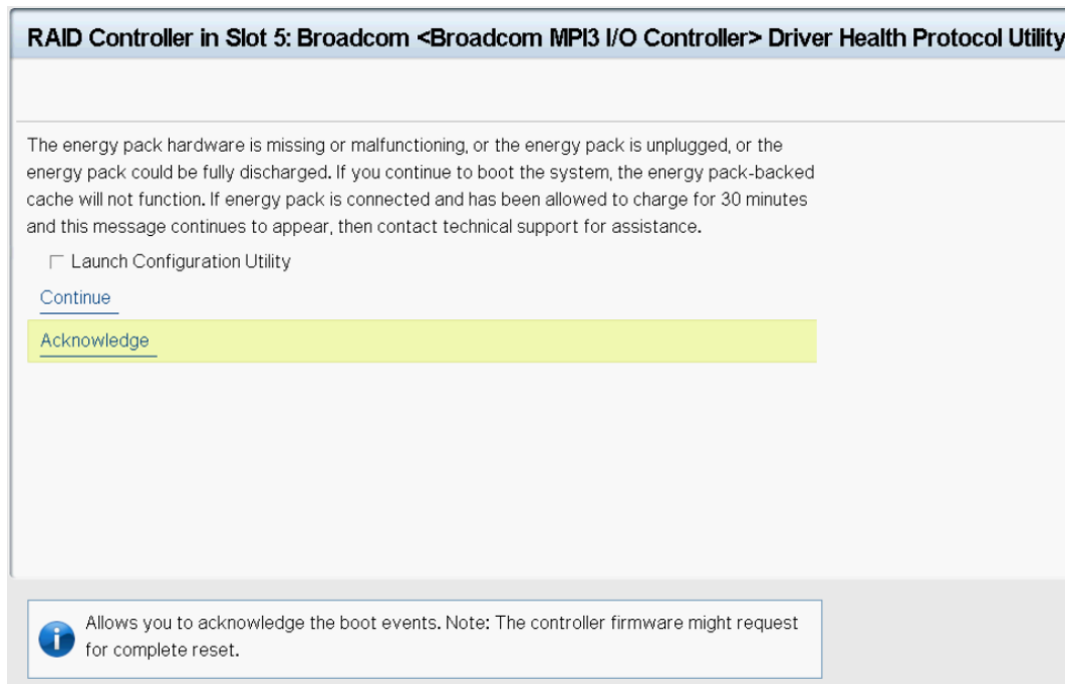
When booting in safe mode, configurations become foreign and management operations allow limited operations. Safe mode is only used for system diagnostics because a system reboot or reset may be required to exit safe mode.

The following example shows where a user configured the local key management with a boot time passphrase enabled. During the boot, the user failed to provide the correct passphrase.

Figure 12: Blocking Boot Event Example

In the following example, selecting **Acknowledge** causes the host to reboot and select **Continue** allows the controller boot sequence to continue. Certain boot events only support the **Continue** action, only host reboot action, or both. The user is presented with these options when a blocking boot event occurs.

If the user selects the **Acknowledge** option, the controller forces a host reboot. If the user selects the **Continue** option, the controller boot continues.

Figure 13: Example Blocking Boot Event Dialog**NOTE**

The following are known limitations for blocking boot events.

- Some of the messages that appear in the Blocking Boot Event screen might have spaces in them.
- If the controller firmware is waiting for the key handling (LKM passphrase, EKM controller, or drive key) after the boot event handling, then this request is honored when a reconnection occurs.
- For critical cases like a topology error, multi-bit ECC errors and so on, do not enter the configuration utility. These errors cannot be fixed through the configuration utility.

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

1. Verify the following items.
 - The cables that connect the drives to the system are well connected, and the host bus adapter (if applicable) is securely seated in its slot.
 - If your system has activity LEDs, make sure that the LEDs do not show a fault.
 - If a cabling or connection issue does not exist with the physical drives, the problem might be the driver.
2. Select **Continue** or **Acknowledge** to continue the boot process.

If you select **Acknowledge**, the firmware may ask for a system reboot.

The BSD asks the system BIOS to reboot the system by filling the health status as `rebootRequired`. The system BIOS should honor this request.

- If these steps do not fix the problem, contact the Broadcom Technical Support team for further assistance.

Starting the HII Configuration Utility

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

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

NOTE

The startup key might be **F2** or **F1** or some other key, depending on the system implementation. Refer to the on-screen text or the vendor-specific documentation for more information. Also, the following workflow may not be the same for all OEM systems.

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

The **System Settings** dialog appears.

- Highlight **Storage** and press **Enter**.

The **Controller Selection** menu appears.

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

- 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 14: Dashboard View

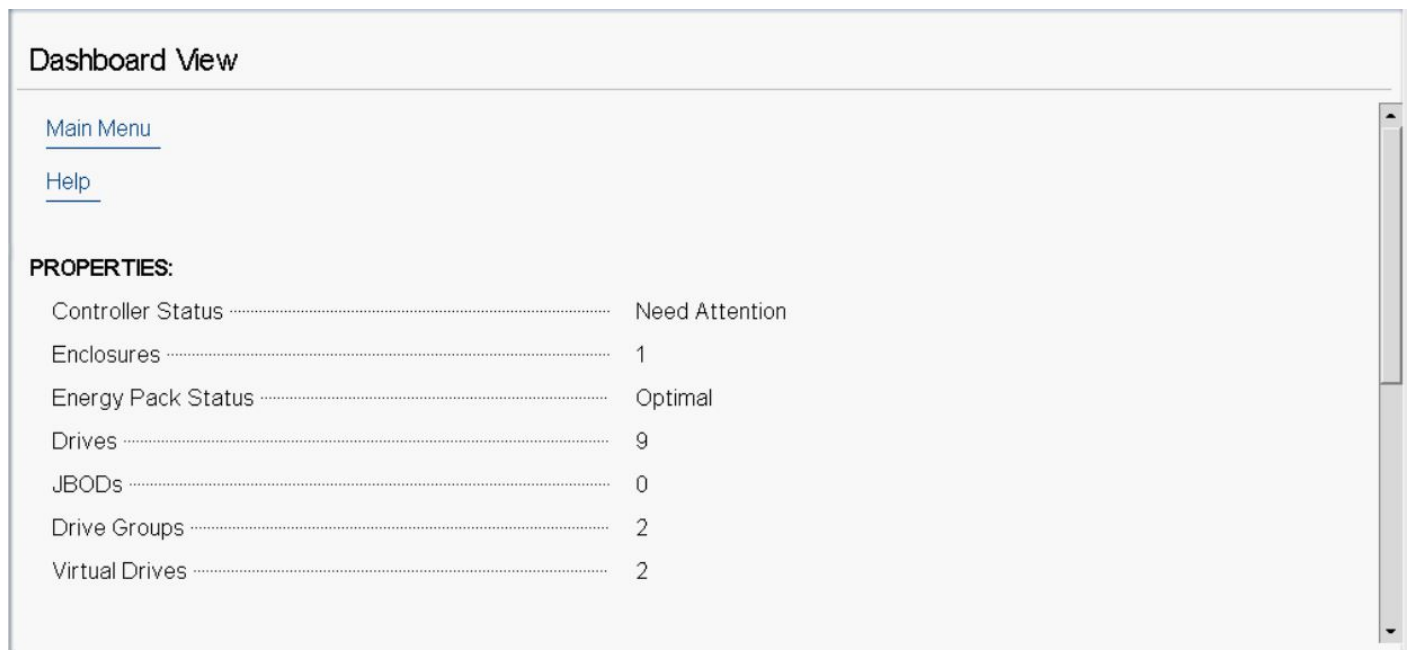


Figure 15: Dashboard View Continued**NOTE**

If you stay in this page, you may experience slow response time when a background operation is in progress. For example, a PD clear or a rebuild, a check consistency, or a full initialization.

You may experience a slow response on a configuration that has a large number of drives that are attached to the controller.

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. HII checks and updates the controller status, updates the drive counts, 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 also hot plug or unplug physical drives and monitor those counts. You can view a foreign configuration, and import and clear a foreign configuration. The HII Dashboard also indicates the number of virtual drive and physical drive operations that are in progress.

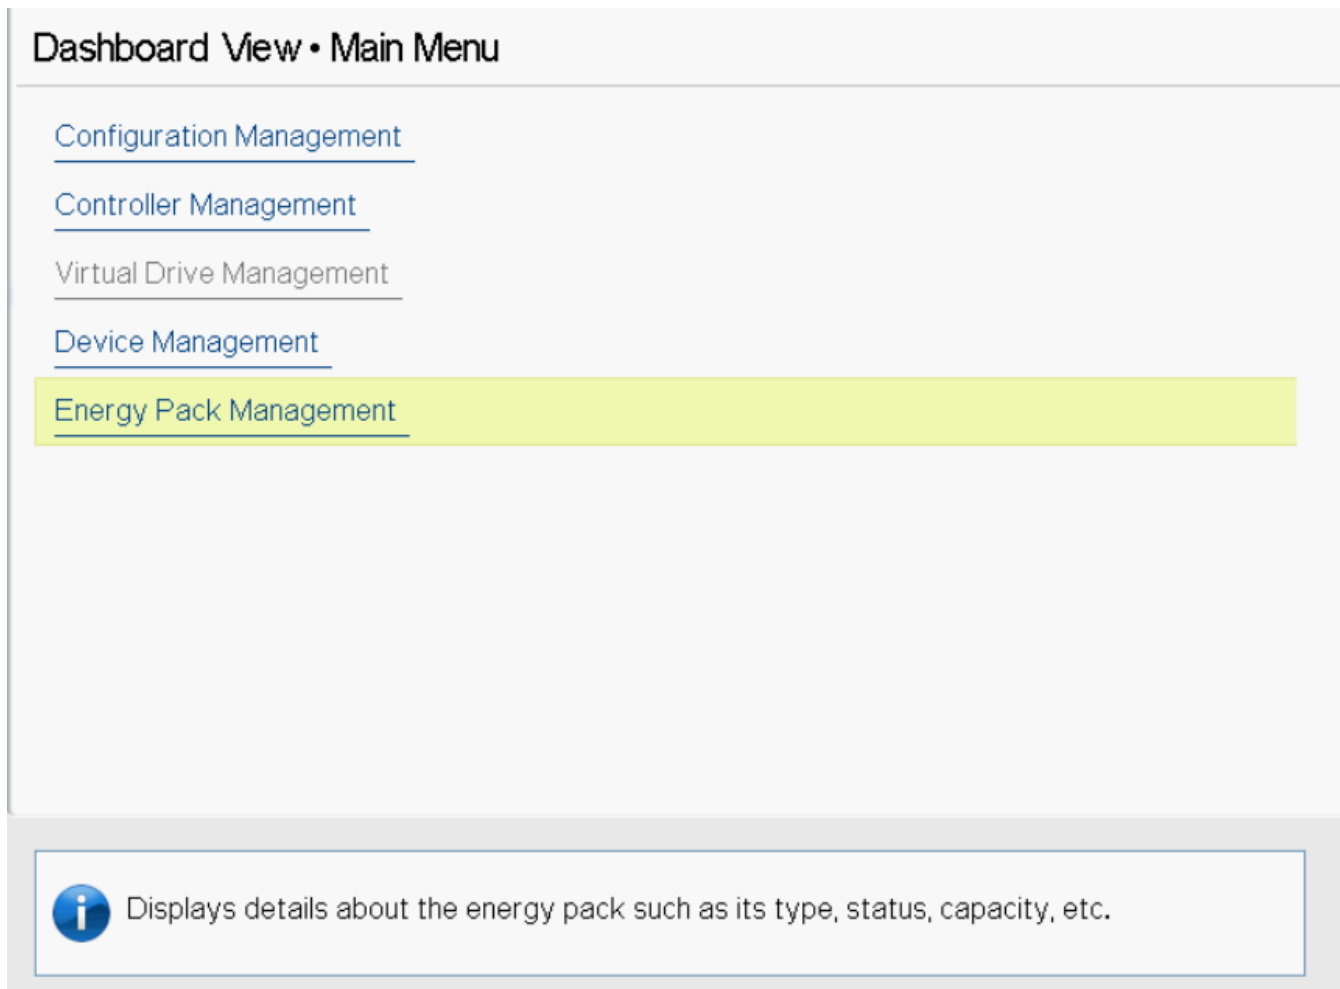
The following sections describe the **Dashboard View**.

NOTE

If you stay in this page, you may experience slow response time when an operation is in progress, for example, a PD clear or a rebuild, or a VD consistency check or a full initialization.

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 16: 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 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 **Device Management** to view enclosure details, the drives that are attached to the enclosure, the drive properties, and perform drive operations. For more information, see [Viewing Physical Drive Properties](#).
- Select **Energy Pack Management** to view energy pack properties. For more information, see [Managing Energy Packs](#).

HELP

The **HELP** section displays the HII utility context-sensitive help. The help also displays the controller status and the reason why the controller status is critical or needs attention. Help strings are displayed for the following functions:

- Discard Preserved Cache
- Foreign Configuration
- Configure

NOTE

The help strings are displayed for the Discard Preserved Cache function only if pinned cache is present.

The help strings are displayed for the Foreign Configuration function only if the foreign configuration is present.

PROPERTIES

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

Figure 17: Dashboard View – PROPERTIES

The screenshot shows the 'RAID Controller in Slot 5: Broadcom <Product Name> Configuration Utility - 08.00.04.00' interface. The 'Dashboard View' section includes a 'Main Menu' and a 'Help' link. The 'PROPERTIES:' section displays the following information:

Property	Value
Controller Status	Need Attention
Enclosures	1
Drives	8
JBODs	3
Drive Groups	2
Virtual Drives	3

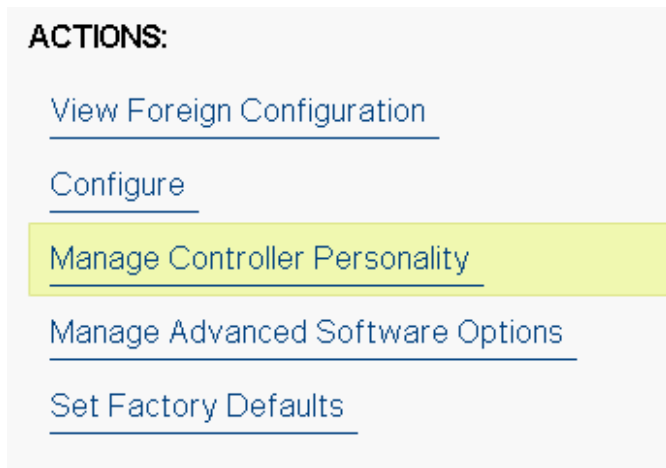
The 'ACTIONS:' section is currently empty. At the bottom, there is an information icon and the text: 'Select to manage controller configuration.'

- **Controller Status**
Displays the overall status of the controller.
- **Enclosures**
Displays the total number of logical enclosures and enclosures connected to this controller.
- **Energy Pack**
Displays the status of the energy pack.
- **Drives**
Displays the total number of drives that are connected to the controller.
- **JBODs**
Displays the number of JBODs connected to the controller.
- **Drive Groups**
Displays the number of drives groups.
- **Virtual Drives**
Displays the number of virtual drives.

ACTIONS

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

Figure 18: Dashboard View – ACTIONS



- **View Foreign Configuration**
Helps you to view and import a foreign configuration and clear a foreign configuration. See [Managing Foreign Configurations](#).
- **NOTE**
If there are secured virtual drives, make sure you enter the passphrase.
- **Configure**
Displays configuration options. See [Managing Configurations](#).
- **Manage Controller Personality**
Allows you to manage the controller personality options. See [Displaying the Controller Personality](#).
- **Manage Advanced Software Options**
Allows you to manage all the activated advance software options on the controller. See [MegaRAID ADVANCED SOFTWARE OPTIONS](#).
- **Set Factory Defaults**
Resets the controller to its factory settings.

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**, **Logical Enclosure**, or **Enclosure** window, respectively. From these windows, 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 19: Dashboard View – BACKGROUND OPERATIONS

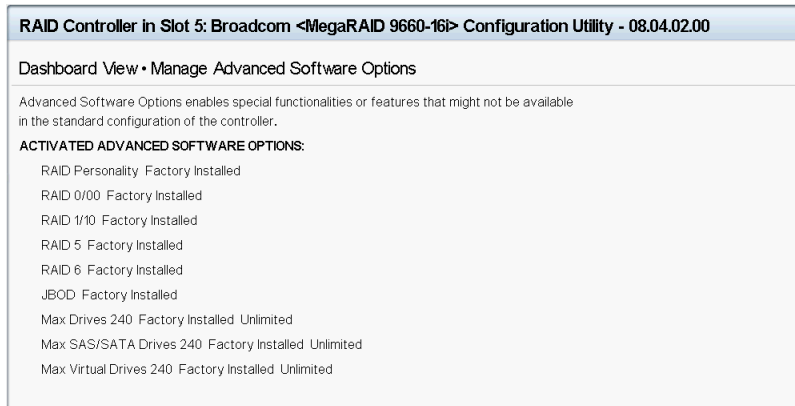
The HII **Dashboard View** has a three second interval to refresh the page. You may see a delay in the progress count on the dashboard.

To view what operations are in progress on each drive, click **Drive Operations in Progress**. The **Device Management** page opens. You can navigate to see what operation is in progress on what drive.

Progress is updated when you are in a basic property page (for example, the progress field in this page). The progress is not updated in title or in the PD or VD summary one-liner (where the VDs and PDs are listed). To see the updated progress, go back and refresh the basic property page.

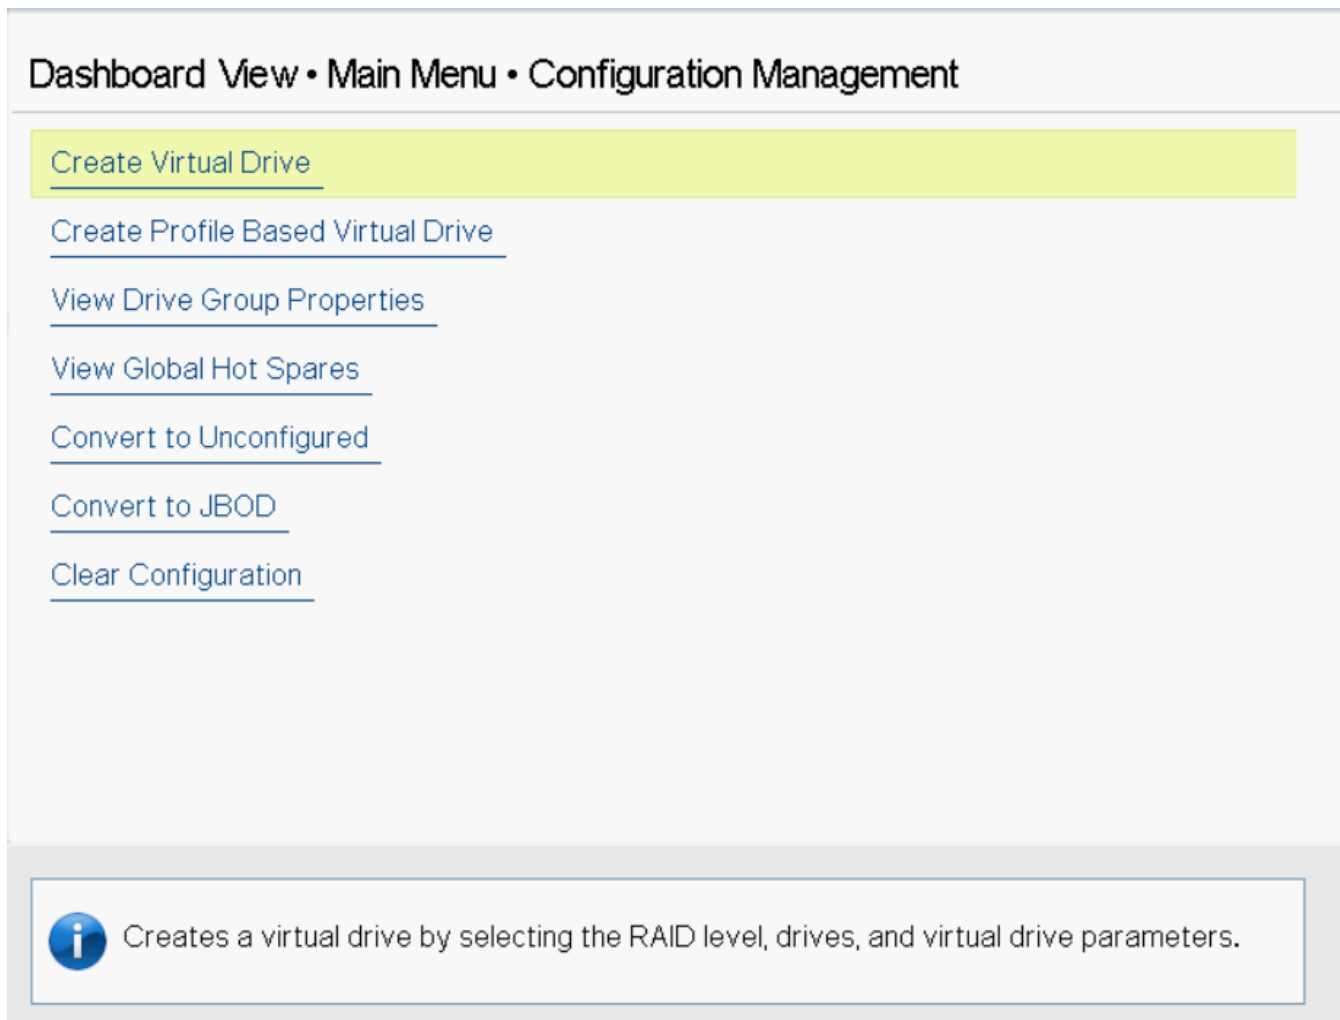
MegaRAID ADVANCED SOFTWARE OPTIONS

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

Figure 20: Dashboard View – MegaRAID ADVANCED SOFTWARE OPTIONS

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 21: Configuration Management Screen

The **Convert to JBOD**, **Enable Security on JBOD**, and **Convert to Unconfigured** options are included for some controllers. See [Convert to Unconfigured](#), [Convert to JBOD](#), and [Enable Security](#).

You can enable security on the JBOD drives either from the **Configuration Management** screen or the **Device Management** screen. The following lists 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 Foreign Configurations](#).

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 Profile Based Virtual Drive** menu. For example, select **Generic RAID 0**.

The following RAID levels are available:

- **Generic RAID 0**
- **Generic RAID 1**
- **Generic RAID 5**
- **Generic RAID 6**
- **File Server**
- **Web/Generic Server**
- **Database**

If you select the **Generic RAID 0** profile, the **Generic R0** screen appears.

NOTE

The names and options displayed may not be the same for all motherboard BIOS and OEM systems.

One or more of the above options may be grayed out if no suitable drives are found.

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 that is shown in the previous figure. The following table describes the profile options.

Table 18: 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 that is 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.
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: <ul style="list-style-type: none"> • 64 KB • 256 KB The supported strip size for SSDs is 64 KB. Both 64 KB and 256 KB (default) are supported for HDDs.
Read Cache Policy	Displays the read cache policy for the virtual drive. No Read Ahead is the default read policy.

Option	Description
Write Cache 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 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.
Drive Write Cache Policy	Displays the virtual drive write cache setting. The possible options are Default , Enable , or 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.
Create Dedicated Hot Spare	Allows you to create a dedicated hot spare.
Create Virtual Drive	Allows you to create a 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 19: Profile Based Virtual Drive Creation Requirements

Properties	Generic RAID0	Generic RAID1	Generic RAID5	Generic RAID6	File Server	Web/Generic Server	Database
HDD	Supported	Supported	Supported	Supported	Supported	Supported	Supported
SSD	Supported	Supported	Supported	Supported	Supported	Supported	Supported
SAS	Supported	Supported	Supported	Supported	Supported	Supported	Supported
SATA	Supported	Supported	Supported	Supported	Supported	Supported	Supported
PCIe	Supported	Supported	Supported	Supported	Supported	Supported	Supported
SED	Supported	Supported	Supported	Supported	Supported	Supported	Supported
NonSED	Supported	Supported	Supported	Supported	Supported	Supported	Supported
NonProtected Information (NonPI)	Supported	Supported	Supported	Supported	Supported	Supported	Supported

Properties	Generic RAID0	Generic RAID1	Generic RAID5	Generic RAID6	File Server	Web/Generic Server	Database
Sector Size (logical block format size) – 4 KB	Supported	Supported	Supported	Supported	Supported	Supported	Supported
Sector Size (logical block format size) – 512 B	Supported	Supported	Supported	Supported	Supported	Supported	Supported
Link speed – 6Gb/s	Supported	Supported	Supported	Supported	Supported	Supported	Supported
Link speed – 12Gb/s	Supported	Supported	Supported	Supported	Supported	Supported	Supported
Link speed – 24Gb/s	Supported	Supported	Supported	Supported	Supported	Supported	Supported
Direct attached	Supported	Supported	Supported	Supported	Supported	Supported	Supported
Logical Enclosure	Supported	Supported	Supported	Supported	Supported	Supported	Supported
Enclosure	Supported	Supported	Supported	Supported	Supported	Supported	Supported
Minimum number of PDs	1	2	3	4	3	3	4
Maximum number of PDs	System default	2	System default	System default	System default	System default	System default
Strip Size	System default	System default	System default	System default	System default	System default	stripSize Max
Read Policy	No Read Ahead	No Read Ahead	No Read Ahead	No Read Ahead	No Read Ahead	No Read Ahead	No Read Ahead
Write Policy	Write-Back	Write-Back	Write-Back	Write-Back	Write-Back	Write-Back	Write-Back
IO Policy	Direct I/O	Direct I/O	Direct I/O	Direct I/O	Direct I/O	Direct I/O	Direct I/O
Access policy	Read/Write	Read/Write	Read/Write	Read/Write	Read/Write	Read/Write	Read/Write
Disk Cache Policy	Enable	Unchanged	Unchanged	Unchanged	Unchanged	Unchanged	Unchanged
Initialization	Fast	Fast	Full	Full	Full	Full	Full
Dedicated Hot Spare	Not supported	Supported	Supported	Supported	Supported	Supported	Supported
Mixing of Media HDD and SSD drives	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported
Mixing of Interface Type SAS, SATA, and NVMe drives	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported

Properties	Generic RAID0	Generic RAID1	Generic RAID5	Generic RAID6	File Server	Web/Generic Server	Database
Mixing of PI and NonPI drives	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported
Mixing SED and NonSED drives	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported
Mixing of 1.5Gb/s, 3Gb/s, 6Gb/s, 12Gb/s, and 24Gb/s link speeds	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported

The following types of mixing are not supported.

- 2.5GT, 5.0GT, 8.0GT, 16.0GT, and 32.0GT
- PRP only and PRP and SGL
- 512b and 4k block size drives

Creating a RAID 10 Volume from the Database

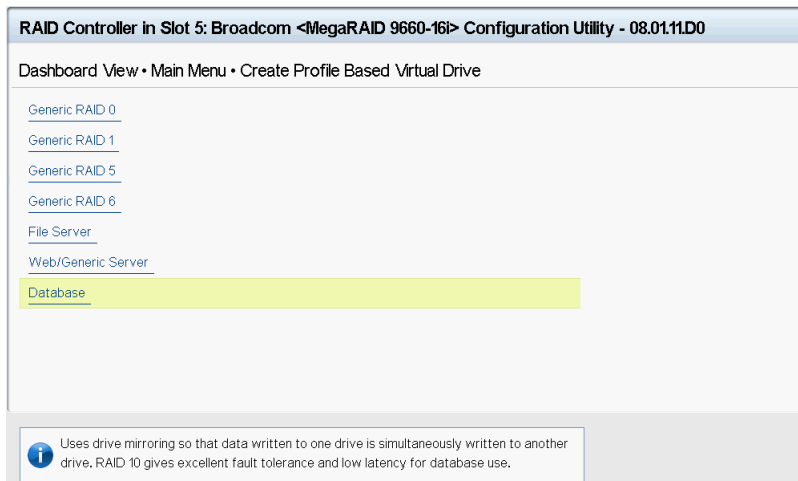
You can create RAID 10 volume from the Database feature. Creating RAID 10 from the Database uses drive mirroring so that data written to one drive is simultaneously written to another drive. Creating a RAID 10 volume from the Database provides fault tolerance and low latency for the use of the database.

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:

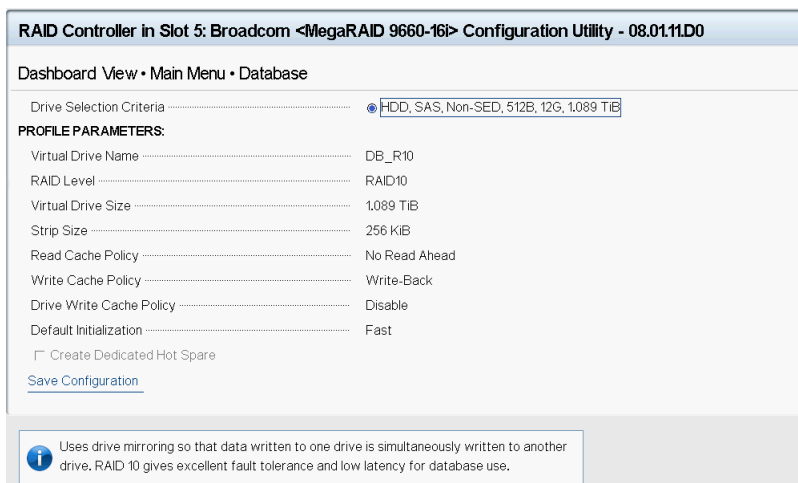
1. Select **Configuration Management** from the **Main Menu**
2. Select **Create Profile Based Virtual Drive** from the **Configuration Management** menu.

A dialog similar to the following example appears.

Figure 22: Example of a Profile Based Virtual Drive Dialog

3. Highlight the **Database** option and press **Enter**.

The **Database** dialog appears.


Figure 23: Example of a Database Dialog

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 24: Create Virtual Drive Dialog

The following limitations apply to manually creating a virtual drive.

- 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, 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.
- HII does not apply the mixing rule across the span when you create spanned RAID levels.
- HII does not maintain the drive selection order while creating VD.
- HII does not have any validation (except checking addConfig) for presenting the free capacity drive groups for creating slices. When a user attempts to create a slice on problematic free capacity (for example, a drive group having a missing drive or rebuild running or max slice already present, the drive group is offline), the firmware will fail it.
- When the maximum number of supported VDs per array is reached (even if this drive group has remaining free capacity) and there is no UG drive or any other drive group with free space, then HII will display *Virtual drive creation was successful. All of the free configurable space has been used.*
- HII uses the following format when a user wants to create VD from a drive group that has free capacity: Drive Group n, Free Space x: RAID Level, where n refers to the drive group number and x refers to the index of the free space. For example, a user may have more than one free space in a drive group or have partial VD and deleted one in the middle and there is a free space at the end.
 - A drive group with multiple free spaces.
 - Drive Group 0, Free Space 0: RAID6
 - Drive Group 0, Free Space 1: RAID6
 - A drive group with single free space.
 - Drive Group 1, Free Space 0: RAID5
 - Drive Group 2, Free Space 0: RAID1

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 RAID levels. See [Table 21](#) 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 only available when the security feature is enabled.

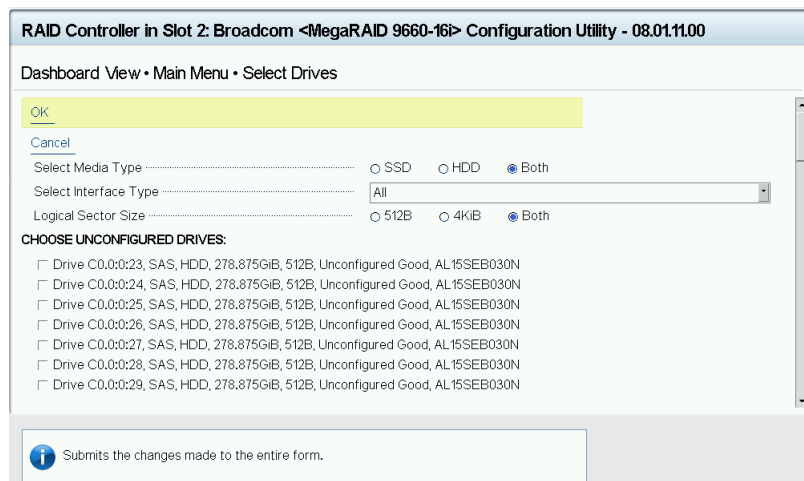
4. Highlight the **Select Drives From** field, press **Enter**, and select either **Unconfigured Capacity** or **Free Capacity**.

Free capacity means that the new virtual drive is created from unused (free) drive capacity that is already part of a virtual drive. *Unconfigured capacity* means that the new virtual drive is created using unconfigured drives.

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

The **Select Drives** dialog appears.

Figure 25: 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**, **NVMe**, 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**. NVMe is not a valid choice for HDD.

NOTE

If the controller does not support NVMe, it will not appear as a valid choice.

- 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 KiB**, 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 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

Ensure you select the number of drives that are required by the specified RAID level, or the HII utility will display an error message. Click **OK** on the error message to return to the **Select Drives** page. For example, RAID 1 virtual drives use exactly two drives, and RAID 5 virtual drives use three or more virtual drives. See [Table 21](#) for more information.

- e) When you have selected the required drives for the new virtual drive, click **OK**.
- 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 the **Virtual Drive Name** field, press **Enter**, and specify a name for the new virtual drive.
- h) (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 **GiB** and **TiB**.

- i) (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).
- j) Highlight **Save Configuration** and press **Enter** to create the virtual drive.

A message appears confirming that the configuration is being created.

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

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

Table 20: Virtual Drive Policies

Property	Description
Current Cache Status	The status of the current cache. The possible values are Write-Through , Write Back , and Always Write-Back .
Strip Size	The virtual drive strip size per DDF. The possible values are as follows: <ul style="list-style-type: none"> • 64 KB • 256 KB
Default Read Cache Policy	Displays the read cache policy for the virtual drive. <ul style="list-style-type: none"> • No Read Ahead Disables the Always Read Ahead capability of the controller.
Default Write Cache Policy	The write cache policy for the virtual drive. The possible values are as follows: <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.
Drive Write Cache Policy	The disk cache policy for the virtual drive. The possible values are Unchanged , Enable , and Disable .

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.
Initialization	Allows choice of the virtual drive initialization option. The possible options are No , Fast , and Full .

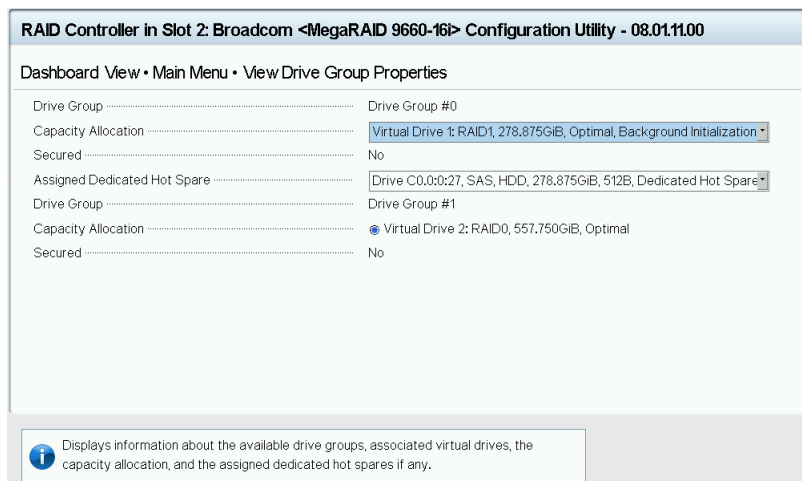
The following table describes the RAID levels that you can select when creating a new virtual drive. Some system configurations do not support RAID 6 and RAID 60.

Table 21: 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 10	A combination of RAID 0 and RAID 1 that uses data striping across two mirrored drive groups. RAID 10 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. RAID 50 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. RAID 60 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 26: View Drive Group Properties Window**NOTE**

By design, it is possible to have gaps between displayed drive group numbers (spanned drive group).

A drive group is a logical grouping of drives that are 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.

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 a single drive group.

Viewing Global Hot Spare Drives

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

Figure 27: View Global Hot Spare 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 spares that are 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.

Convert to Unconfigured, Convert to JBOD, and Enable Security

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

Convert to Unconfigured

Perform these steps to convert JBOD drives to Unconfigured drives.

1. Select **Convert to Unconfigured** on the **Configuration Management** menu.

The **Convert to Unconfigured** dialog appears, listing all the JBOD drives connected to the controller.

NOTE

JBOD drives listed as **JBOD (Bad)** are converted to Unconfigured Bad drives.

Figure 28: Convert to Unconfigured Dialog



- a) To select a specific JBOD drive and convert it to Unconfigured, select the checkbox for the drive.
- b) To select all the JBOD drives and convert them to Unconfigured drives, click **Check All**.
- c) (Optional) To unselect all the drives that you have selected, click **Uncheck All**.
- d) (Optional) To cancel, click **Cancel**.

ATTENTION

If one or more selected JBOD drives have an operating system or a file system on them, a warning message appears indicating that the listed JBOD drives have an operating system or a file system. If you proceed with the conversion, any data on these drives is lost. 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. Unselect the JBOD drives that have an OS or a file system that is installed on them.

2. Click **OK** to convert the selected JBOD drives to Unconfigured drives.

Convert to JBOD

Perform these steps to convert Unconfigured drives to JBOD drives.

1. Select **Convert to JBOD** on the **Configuration Management** menu.

The **Convert to JBOD** dialog appears, listing all the JBOD drives connected to the controller.

Figure 29: Convert to JBOD Dialog



- a) To select a specific Unconfigured drive and convert it to JBOD, select the checkbox for the drive.
 - b) (Optional) To select all the Unconfigured drives and convert them to JBOD drives, click **Select All**.
 - c) (Optional) To unselect all the drives that you have selected, click **Unselect All**.
 - d) (Optional) To cancel, click **Cancel**.
2. Click **OK** to convert the selected Unconfigured drives to JBOD drives.

Enabling Security on a Controller

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 drive.

ATTENTION

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 and lists the SED-enabled JBOD drives currently connected to the controller.

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.

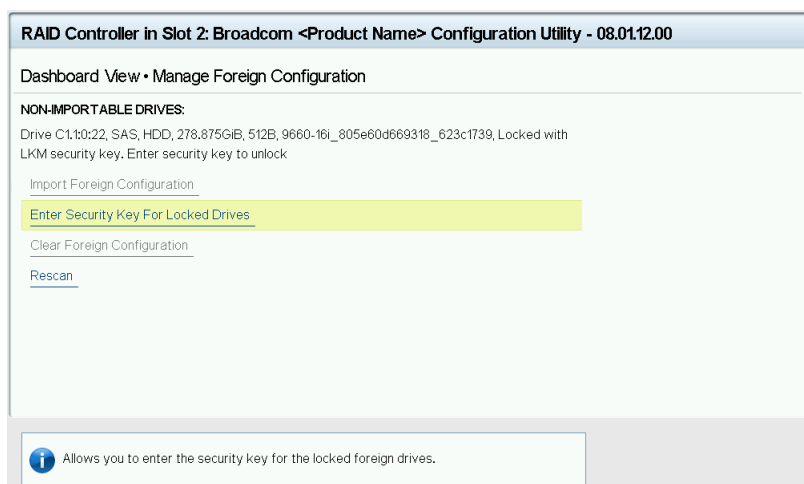
Managing Foreign Configurations

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

NOTE

A large number of foreign drives may delay loading of this page.

Figure 30: Manage Foreign Configuration Dialog



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

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

Viewing and Importing a Foreign Configuration

You can view 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. If any of the following conditions exist, you might be unable to import a foreign configuration.

- The volume state is ACTIVE.
- The volume state is either FAILED or MISSING.
- The volume uses incompatible Gen1 metadata.
- The maximum number of RAID volumes already exists on this controller.
- A virtual drive persistent ID, including pinned persistent IDs (because of pinned cache or pinned write journals), is not available. Users must wait 120 seconds before importing the foreign configuration.
- The maximum number of supported physical drives is 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.
- If a virtual drive is deleted and another one is created in a quick succession, the same persistent ID may be reused. Some operating systems may not be able to differentiate between the two IDs and may lead to inconsistent behavior. To prevent this, the persistent ID is made unavailable for a period of 120 seconds starting from the time the virtual drive

was deleted. If a configuration creation or a foreign import is attempted within the 120 second duration, the requests may fail and need to be retried after 120 seconds.

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 a 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 because 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 displays the menu option **Enter Security Key for Locked Drives** under the **Manage Foreign Configuration** menu.

Figure 31: Enter Security Key for Locked Drives

RAID Controller in Slot 2: Broadcom <Product Name> Configuration Utility - 08.01.12.00

Dashboard View • Manage Foreign Configuration • Enter Security Key for Locked Drives

LOCKED DRIVES
 Drive C1:10:22, SAS, HDD, 278.875GiB, 512B, Enterprise, Secured, Unconfigured Good, Foreign, ST300MM0026, 9660-16i_805e60d669318_623c1739

Key ID 9660-16i_805e60d669318_623c1739

Caution: The security key that you enter will not be hidden.

Security Key Lsi@1234

OK

[Cancel](#)

i Proceeds with unlocking a foreign configuration.

Perform these steps to view and import a foreign configuration.

1. Select **View Foreign Configuration** or **Manage Foreign Configuration** menu and press **Enter**.

The following dialog appears, listing information about the physical drives in the foreign configuration.

Figure 32: Manage Configuration Window

RAID Controller in Slot 2: Broadcom <MegaRAID 9660-16i> Configuration Utility - 08.01.11.00

Dashboard View • Manage Foreign Configuration

IMPORTABLE DRIVES:
 Drive Group# 0, Drive C1:10:21, SAS, HDD, 278.875GiB, 512B, Enterprise, Secured, Unconfigured Good, Foreign, ST300MM0026
 Drive Group# 0, Drive C1:10:22, SAS, HDD, 278.875GiB, 512B, Enterprise, Secured, Unconfigured Good, Foreign, ST300MM0026

VIRTUAL DRIVES:
 Drive Group# 0, Virtual Drive 1: RAID1, 278.875GiB, Optimal

Import Foreign Configuration

[Clear Foreign Configuration](#)

[Rescan](#)

i Imports all foreign configurations.

2. Review the information listed on the window.
3. 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.

4. To confirm the import, highlight **Confirm** and press the spacebar.
5. Highlight **Yes** and press **Enter**.

The foreign configuration is imported.

NOTE

After an import, the drive group numbers that are listed in **Manage Foreign Configuration** and in other pages might not be consistent. By design, it is possible to have gaps between displayed drive group numbers (spanned drive group).

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.

Managing Controllers

When you select **Controller Management** from the **Main Menu**, 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 other 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 other controller properties, in the **Basic Properties** section, highlight **Advanced Controller Properties**.
For more information, see [Viewing Advanced Controller Properties](#).

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

Table 22: Basic Controller Properties

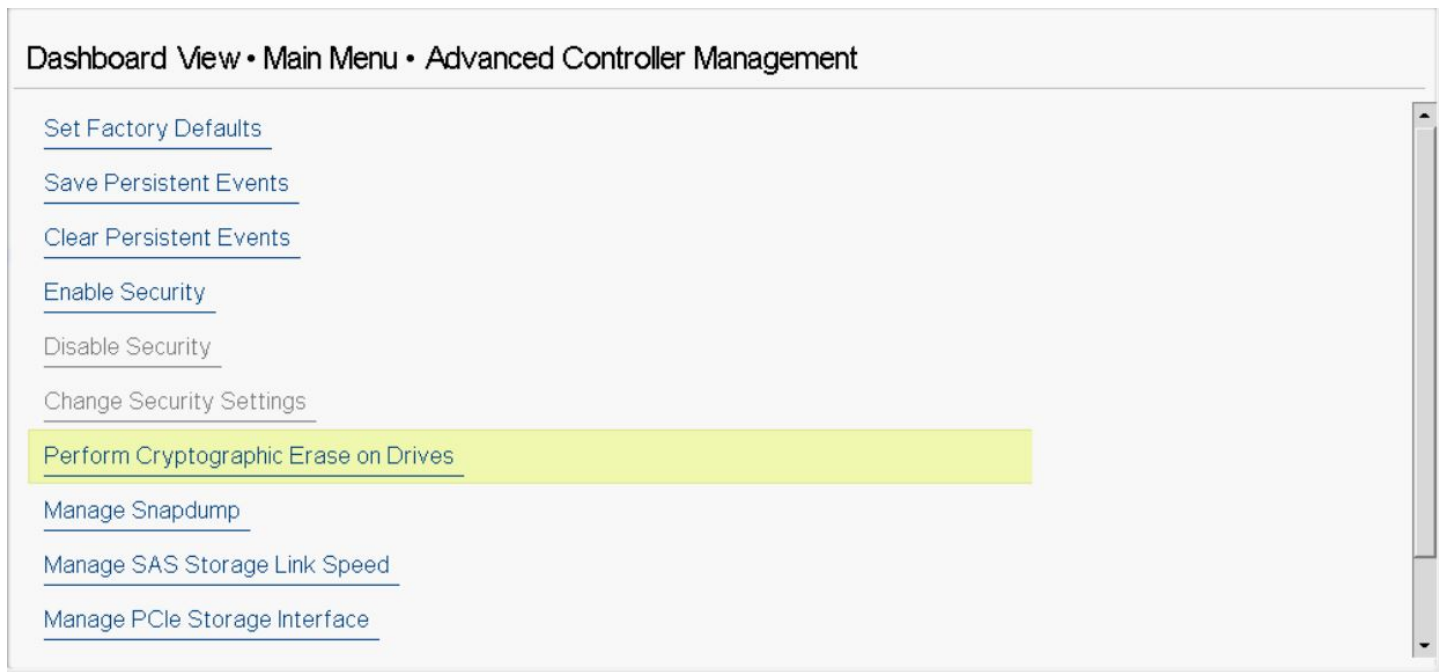
Property	Description
Product Name	The marketing name of the controller.
Serial Number	The serial number of the controller.

Property	Description
Controller Status	<p>The cumulative status of virtual drives and physical drives that are connected to the controller, plus the backup battery, the enclosure, and the NVDATA. The controller status falls into one of the following categories:</p> <ul style="list-style-type: none"> • Optimal, if all components are operating normally. • Needs Attention, if any component needs attention. • Critical, if the controller encountered critical errors. <p>Most features are disabled and the controller requires user attention.</p> <p>The controller is in a Critical state for one or more of the following reasons:</p> <ul style="list-style-type: none"> • The controller booted in safe mode. • One or more uncorrectable errors are detected in the DDR memory. • The preserved cache is present and the controller does not allow booting with the preserved cache. • The energy pack has one or more errors. <p>The controller in a Needs Attention state for one of the following reasons:</p> <ul style="list-style-type: none"> • One or more drives are in a failed state. • One or more drives are in an offline state. • One or more drives are in a bad state. • One or more unsupported drives are present. • One or more drives are in an unusable state. • One or more virtual drives are not in an optimal state. • The preserved cache is present, and the controller allows booting with the preserved cache. • One or more drives are in a foreign state. • One or more enclosures are in a fault state. • One or more NVMe drives have a negotiated port width less than the maximum port width. • The controller supports NVMe drives, but the PCIe interface is disabled. • The energy pack has one or more warnings. • The security key change (alias re-key) is pending. • A system reboot is required. • A system shutdown is required. • The controller is running a secondary firmware image. • The controller has booted in a secure debug mode. • The SAS and SATA support has been disabled because of a controller initialization error.
Controller Personality	Mode that allows the controller to function as appropriate for the user environment.
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.
Supported Device Interfaces	List of device interfaces supported.
Drive Count	Number of physical drives that are attached to this controller.
JBOD Count	Number of JBODs used by the controller.
Virtual Drive Count	Number of virtual drives defined on this controller
Chip Name	The name of the chip in the controller.
Chip Revision	The revision of the chip in the controller.
SAS Address	Unique identifier for the controller.
Advanced Controller Management	Lists all the controller management properties and options for performing various actions on the controller.
Advanced Controller Properties	Lists all the controller properties and options for performing various actions on the 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 33: Advanced Controller Management Dialog



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

Table 23: Controller Management Options

Property	Description
Set Factory Defaults	Resets the controller to its factory settings.
Save Persistent Events	Saves the log entries to a file.
Clear Persistent Events	Clears entries from the log.
Enable Security	Enables drive security to protect the data on your system from unauthorized access or use.
Disable Security	Disables drive security.
Change Security Settings	Changes the security settings or switches between drive security modes on the controller.
Perform Cryptographic Erase on Drives	Permanently erases all the data on the drive.
Manage Snapdump	Displays the Snapdump properties and operations, and allows you to manage the Snapdump features.
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 Speed .

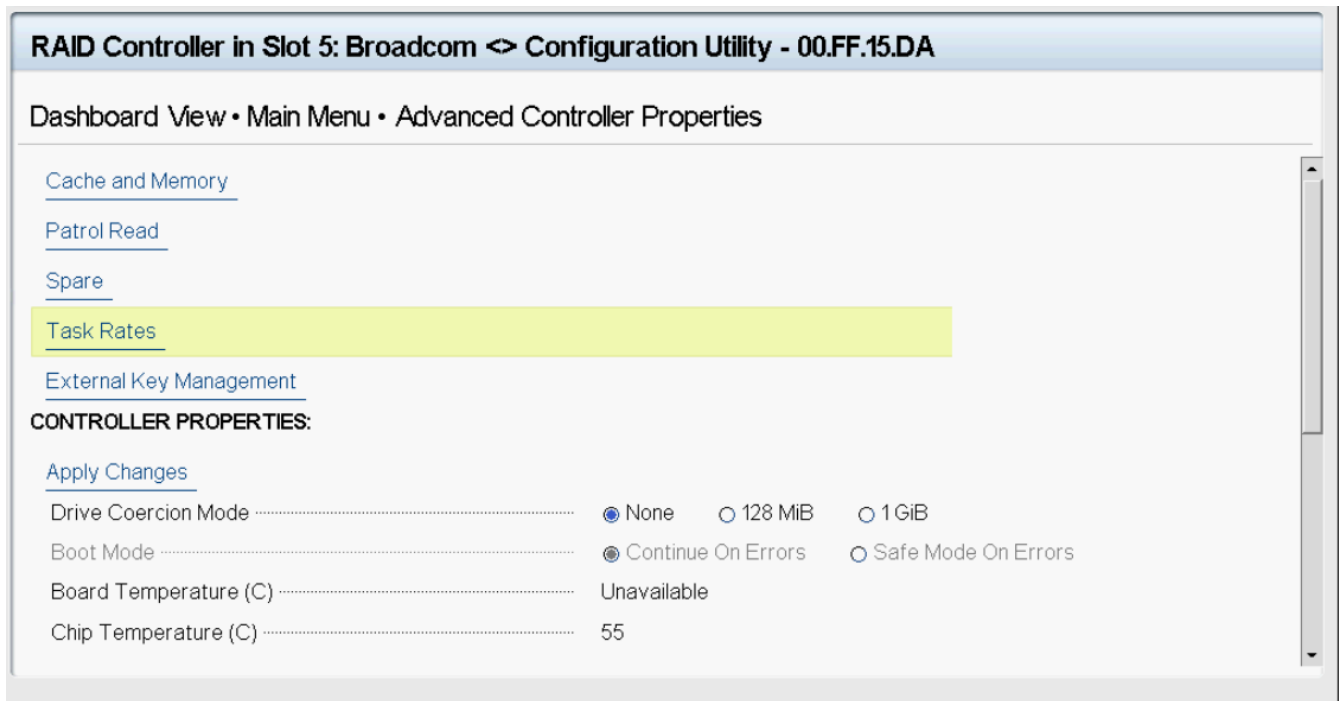
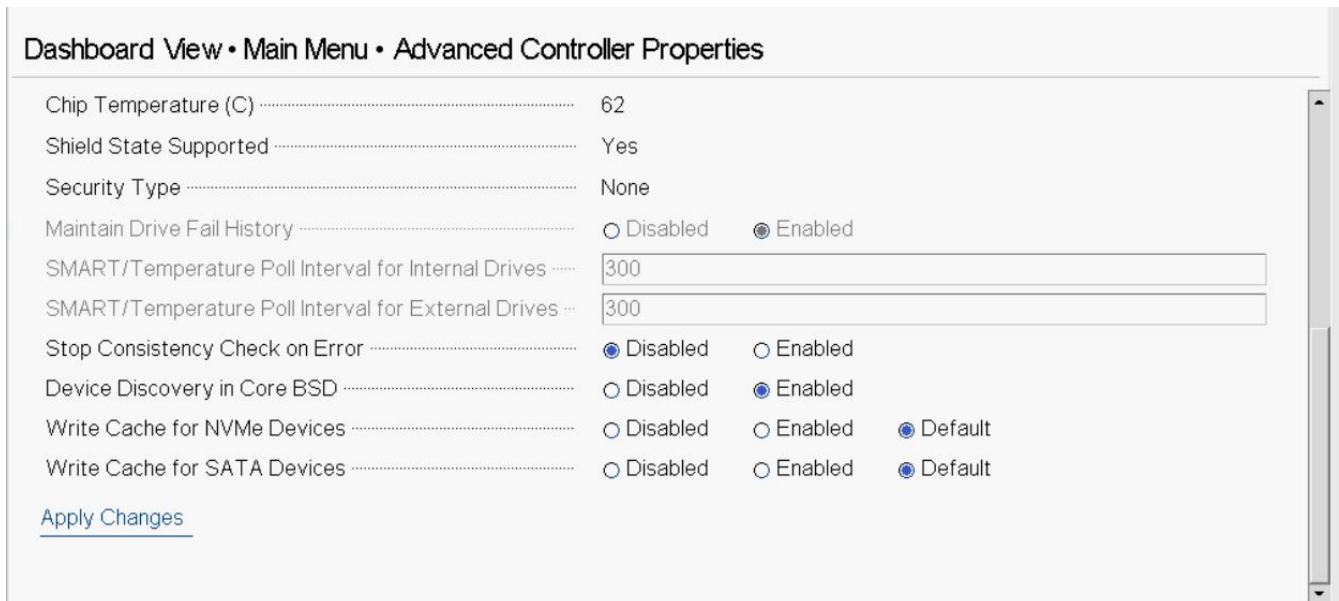
Property	Description
Manage PCIe Storage Interface	A lane represents a set of differential signal pairs, one pair for transmission and one pair for reception, similar to SAS phys. The Manage PCIe Storage Interface feature allows you to change the lane speed between a controller and expander or between 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 Advanced Software Options	Displays the activated Advanced Software Options on the controller and lets you configure these options to use the advanced features in the controller. You must activate the activation key to use the advanced features. The MegaRAID Advanced Software Options are displayed only if the controller supports MegaRAID software licensing.
Manage Controller Personality	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.

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 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 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](#).
- To view information for external keys, highlight **External Key Management** and press **Enter**.

Figure 34: Advanced Controller Properties Dialog**Figure 35: Advanced Controller Properties Dialog (continued)**

This dialog lists various properties, scroll down to view all of the options.

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.

- Highlight the value that you want and press **Enter**. For options, such as **SMART/Temperature Poll Interval for Internal Drives** 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.

- 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 that are listed in the **Advanced Controller Properties** section, including the ones that are not visible.

Table 24: Advanced Controller Properties

Property	Description
Drive Coercion Mode	Drive coercion forces the drives of varying capacities to the same size, so the drives can be used in a drive group. The coercion mode options are None , 128 MiB , and 1 GiB .
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 Continue on error and Safe mode on errors .
Board Temperature (C)	Indicates the temperature of the board.
Chip Temperature (C)	Indicates the temperature of the chip.
Shield State Supported	Indicates whether the controller supports shield state.
Security Type	Indicates the type of security on the controller.
Maintain Drive Fail History	Enables or disables the option to track bad physical drives through a reboot.
SMART Poll Interval for Internal Drives	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.
SMART Poll Interval for External Drives	Determines the interval, in seconds, at which the controller polls for JBODs 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.
Device Discovery in Core BSD	Specifies whether the UEFI BSD should register the device for block IO and EXT SCSI pass-thru access. Disabling this option results in the UEFI BSD registering this controller with the appropriate UEFI interfaces, but not registering any attached devices.
Write Cache for NVMe Devices	Enables or disables the write verify feature for NVMe during the controller cache flush. This feature verifies if the data was written correctly to the cache before flushing the cache.
Write Cache for SATA Devices	Enables or disables the write verify feature for SATA during the controller cache flush. This feature verifies if the data was written correctly to the cache before flushing the cache.

Managing MegaRAID Advanced Software Options

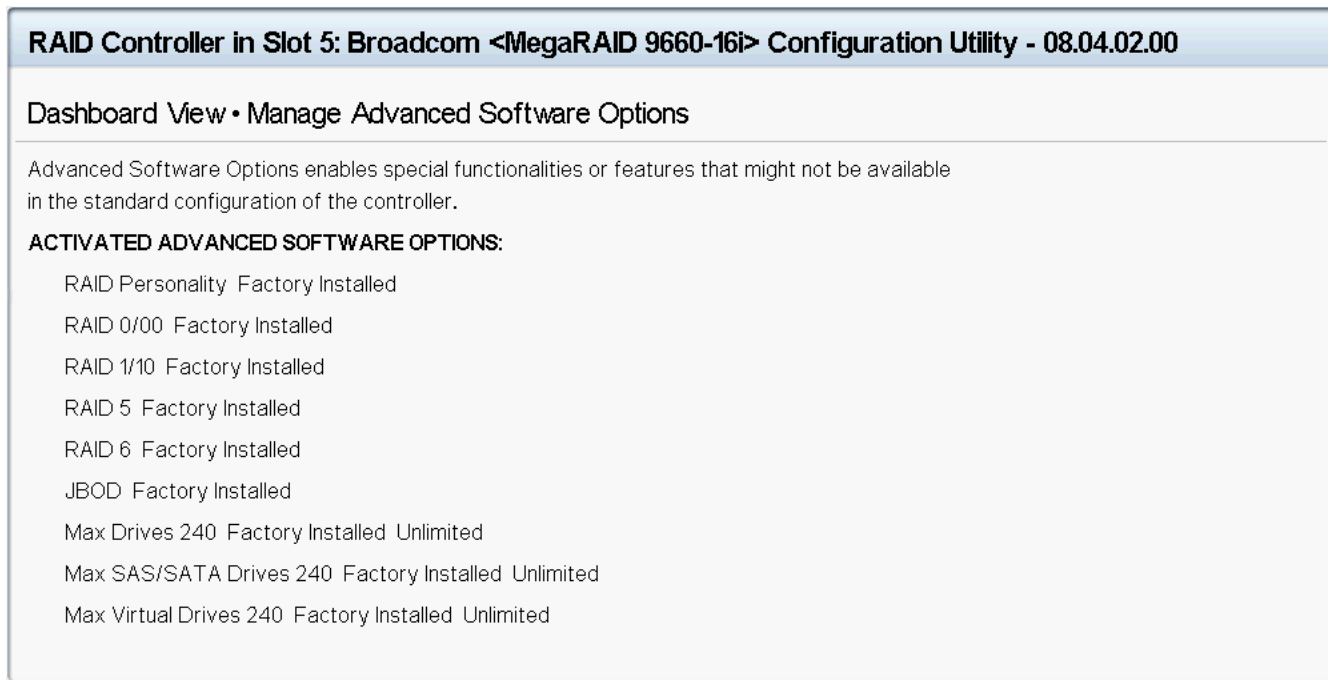
The **Manage MegaRAID Advanced Software Options** dialog lists all the activated advanced software options on the controller. You can configure the MegaRAID advanced software options to use the advanced software features.

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 36: 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 that are 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.

Displaying the Controller Personality

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 you power off a controller and insert a new physical drive, if the drive contains valid DDF metadata, its drive state is **Foreign**. A new drive in the JBOD drive state is exposed to the host operating system as a stand-alone drive if the

drive contains valid DFF metadata from another controller. 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.

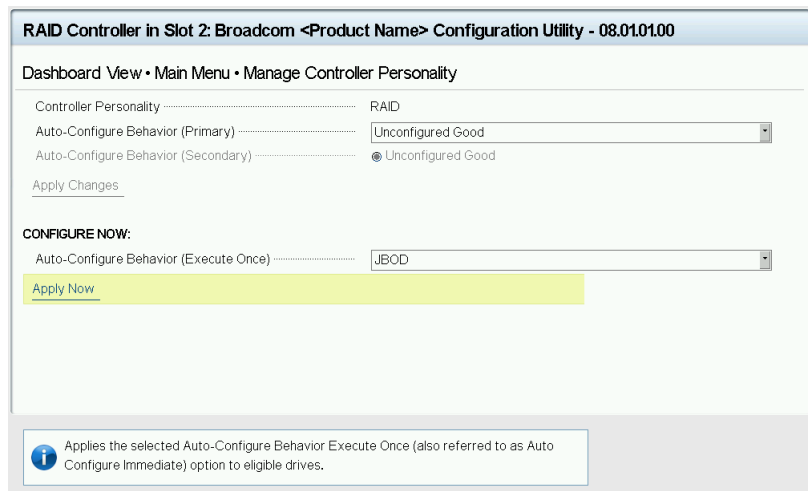
If your system is in personality mode (RAID), the firmware supports auto-configure options to allow the controller to function as appropriate for the user environment.

You can use the **Manage Controller Personality** setting to change the Auto-Configure Behavior (Primary) or Auto-Configure Behavior (Execute-Once).

1. In the **Advanced Controller Management** dialog, click **Manage Controller Personality**.

The **Manage Controller Personality** dialog appears.

Figure 37: Manage Controller Personality Dialog



2. Complete one of the following options to select the controller personality.
 - To select a primary auto-configuration behavior, select an option from the **Auto-Configure Behavior (Primary)** drop-down menu.
 - To select a one-time configuration, select an option from the **Auto-Configure Behavior (Execute-Once)** drop-down menu.
3. Click **Apply Now**.
A confirmation window appears.
4. Select the **Confirm** checkbox and then click **Yes**.

Saving or Clearing Persistent Events

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

NOTE

If the events log is empty, an error message appears.

Figure 38: Save Persistent Events Dialog

RAID Controller in Slot 5: Broadcom <> Configuration Utility - 00.FF.15.DA

Dashboard View • Main Menu • Save Persistent Events

Select File System JW NOLABEL0

Select Directory Current Dir

Enter Filename ctrlEvents.txt

Events to Save Events Since Last Shutdown

[Save Events](#)

Displays the options which will filter the persistent events being saved. (Press <F1> for more help)

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

IMPORTANT

If no file system is detected, then the following message appears.

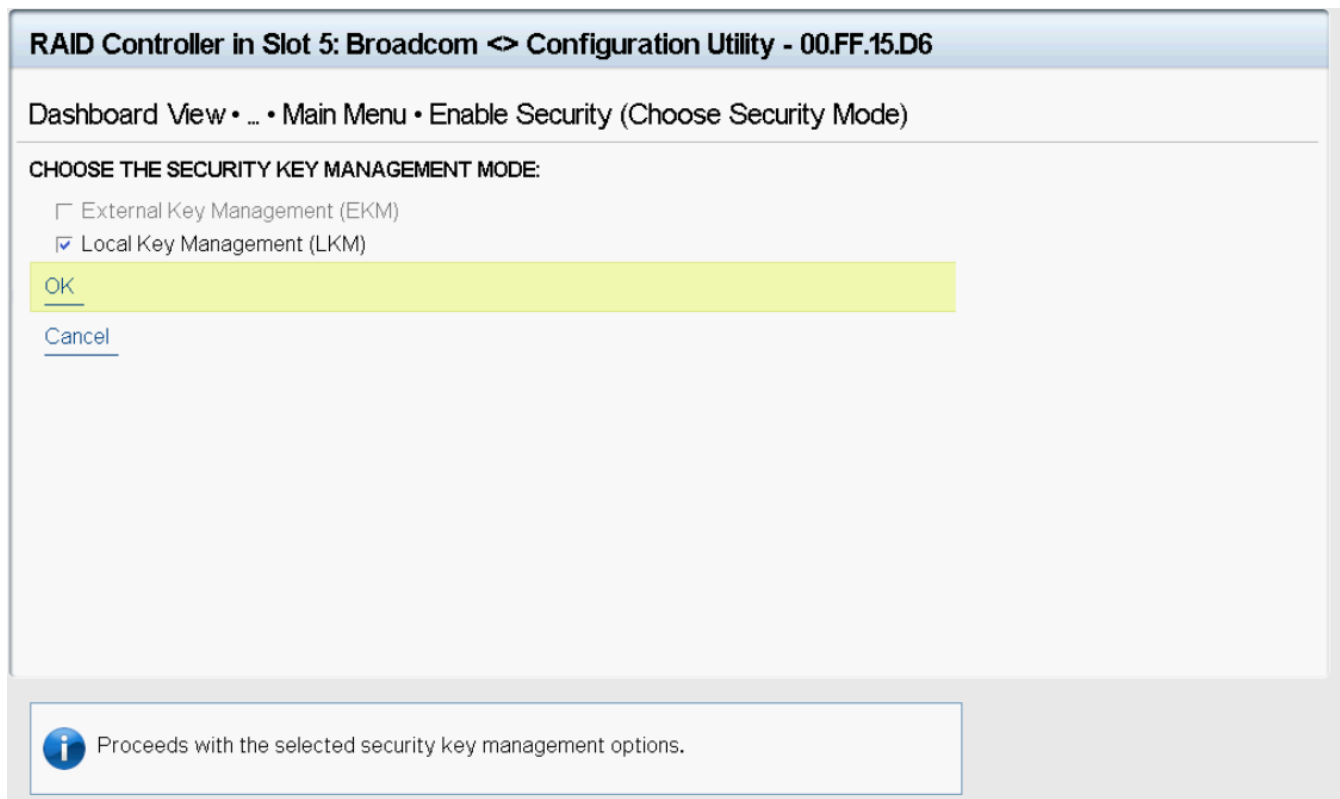
Persistent events cannot be saved because the file system is not available.

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**.
If there is no file system, an error message appears.
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 Persistent Events** in the **Advanced Controller Management** dialog. When the confirmation message appears, highlight **OK** and press **Enter**.

Enabling or Disabling Security

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

Figure 39: Enable Drive Security (Choose 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 the **Pause for password at boot** feature, which is part of the security feature, the system BIOS must support ECR 1085 and 1174; otherwise you cannot 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. Select the **Local Key Management (LKM)** field and then click **OK**.

The following dialog appears.

Figure 40: Enable Security Dialog

RAID Controller in Slot 5: Broadcom <Product Name> Configuration Utility - 00.FF.16.D1

Dashboard View • Main Menu • Enable Security

Security Key Identifier RAID_805e000004_00000000

[Suggest Security Key](#)

Caution: The security key that you enter will not be hidden.

Security Key
Confirm

PASSPHRASE:


Pause for Passphrase at Boot Time

Caution: The passphrase that you enter will not be hidden.

Passphrase
Confirm

I Recorded the Security Settings for Future Reference

[Enable Security](#)

 Displays the controller-suggested default security key identifier. Allows users to use the default string or enter their own identifier. (Press <F1> for more help)

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

- To request the controller to suggest a drive security key, select **Suggest Security Key**.
- To enter your own security key, type the new security key into the blank **Security Key** field.

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, > @ +).

- Re-enter the security key into the **Confirm** field to confirm it the new security key.
The security key must match exactly the characters you entered in the **Security Key** field.
- If you do not want the controller to require a password at boot time, deselect the **Pause for Password at Boot Time** option.

This option is selected by default.

- To enforce strong password restrictions, enter a passphrase in the **Passphrase** field.
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, > @ +).
- Re-enter the passphrase in the **Confirm** field to confirm the new passphrase.
The passphrase must match exactly the characters you entered in the **Passphrase** field.

- Record the drive security information and store it in a safe place.
- Select the **I Recorded The Security Settings** option.
- Select **Enable Security**.
- Select **Confirm** on the **Warning** dialog to confirm that you want to enable drive security, then select **Yes**.

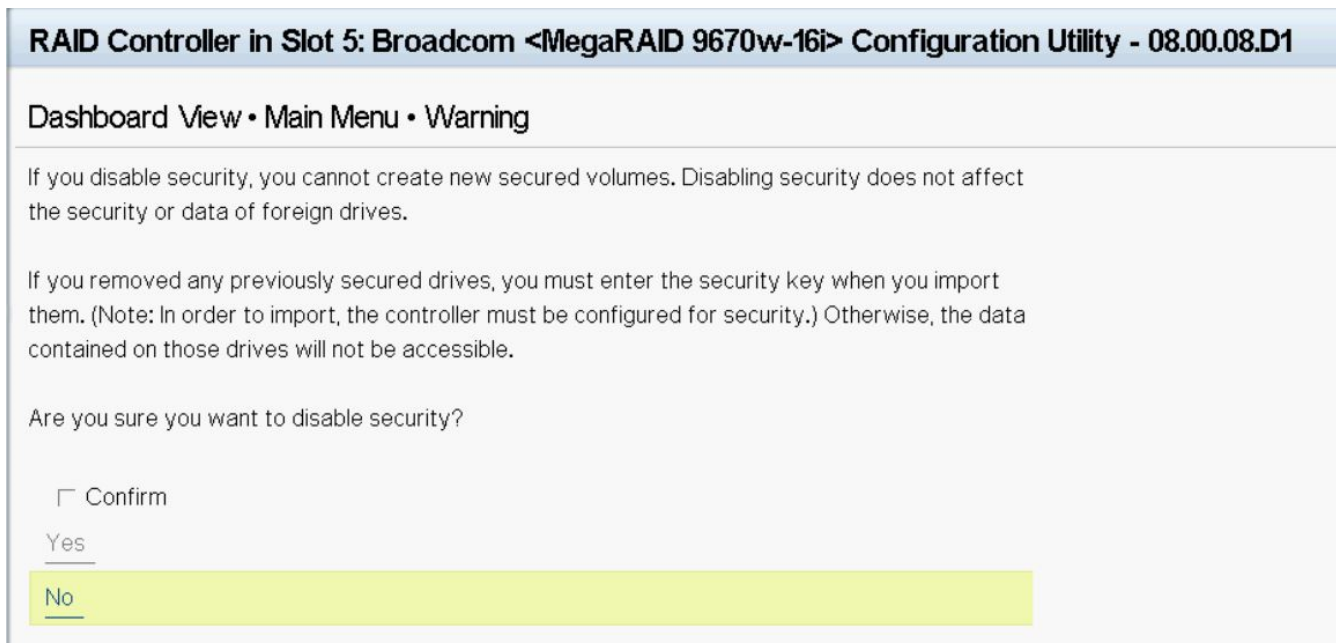
Figure 41: Enable Security Dialog



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

Follow these steps to disable LKM drive security:

- Select **Disable Security** from the **Advanced Controller Management** menu.
The following warning appears.

Figure 42: Disable Security Warning

2. Read the warning and be sure you understand what happens if you disable the drive security.
3. Select the **Confirm** checkbox, and then select **Yes**.
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 43: Change Security Settings Dialog

RAID Controller in Slot 5: Broadcom <Product Name> Configuration Utility - 00.FF.16.D1

Dashboard View • Main Menu • Change Security Settings

Enter a New Security Key Identifier

Use the Existing Security Key Identifier

Caution: The security key that you enter will not be hidden.

Enter Existing Security Key

Use the Existing Security Key

[Suggest Security Key](#)

Caution: The security key that you enter will not be hidden.

Enter a New Security Key


Confirm

PASSPHRASE:

Pause for Passphrase at Boot Time

Caution: The passphrase that you enter will not be hidden.

Passphrase

 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 ... (Press <F1> for more help)

By default, the same security key identifier is retained.

2. To change the security key identifier, deselect the **Use the Existing Security Key Identifier** option.
3. In the **Enter a New Security Key Identifier** field, enter the new security key identifier.
4. In the **Enter Existing Security Key** field, enter the current security key.

You are required to enter the security key to prevent unauthorized changes to the security settings.

5. Select the **Suggest Security Key** to have the system create a new security key.
6. To enter your own new security key, use the **Enter A New Security Key** field, 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, > @ +).

7. Confirm the new security key by reentering the key in the **Confirm** field.

The security key must match exactly the characters you entered in the **Enter a New Security Key** field.

8. If you do not want the controller to require a passphrase at boot time, deselect the **Pause for Passphrase at Boot Time** option.

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

This option is selected by default.

9. Enter the new boot time passphrase in the **Passphrase** field.
10. Highlight **Confirm** and reenter the new passphrase.
The passphrase must match exactly the characters you entered in the **Passphrase** field.
11. Record the drive security information and store it in a safe place.
12. Select the **I Recorded The Security Settings** field option.
13. Highlight **Save Security Settings** and press **Enter**.
14. 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 that are connected to this controller.

Perform Cryptographic Erase on Drives

The cryptographic erase operation erases all the security present on the drive.

If the controller firmware supports SafeStore and there is at least one PD where reprovision is allowed, then this menu option appears. A cryptographic erase allows the user to perform the operation on multiple drives in a single instance.

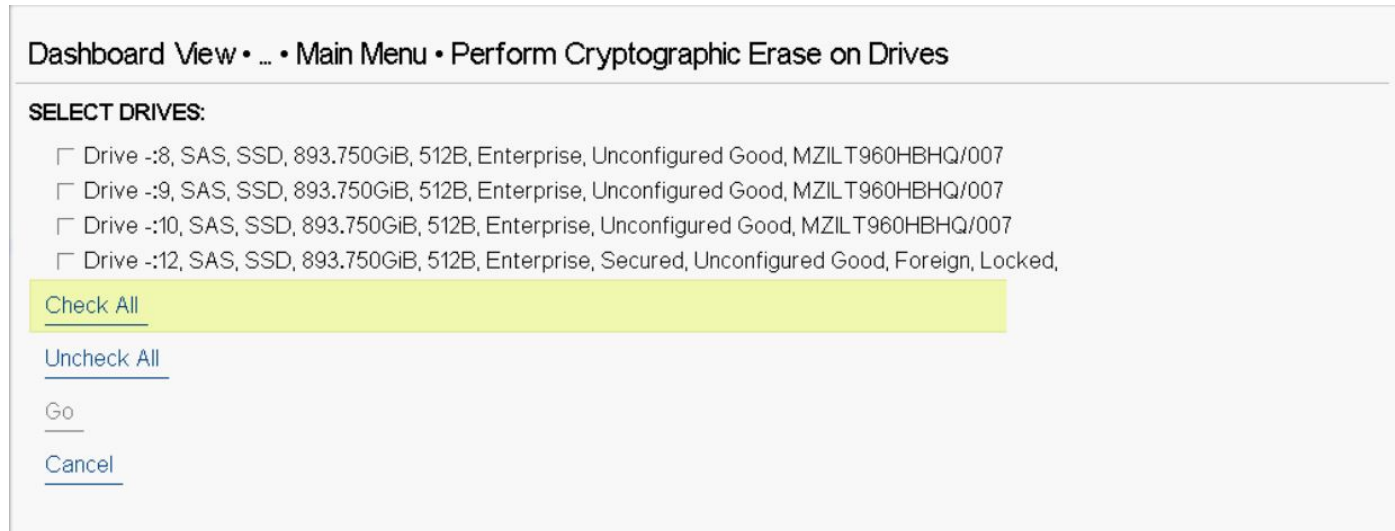
NOTE

You can perform the same operation at the drive level for individual drives.

1. Select **Perform Cryptographic Erase on Drives** from the **Advanced Controller Management** menu.
2. Select the drive that you wish to erase.

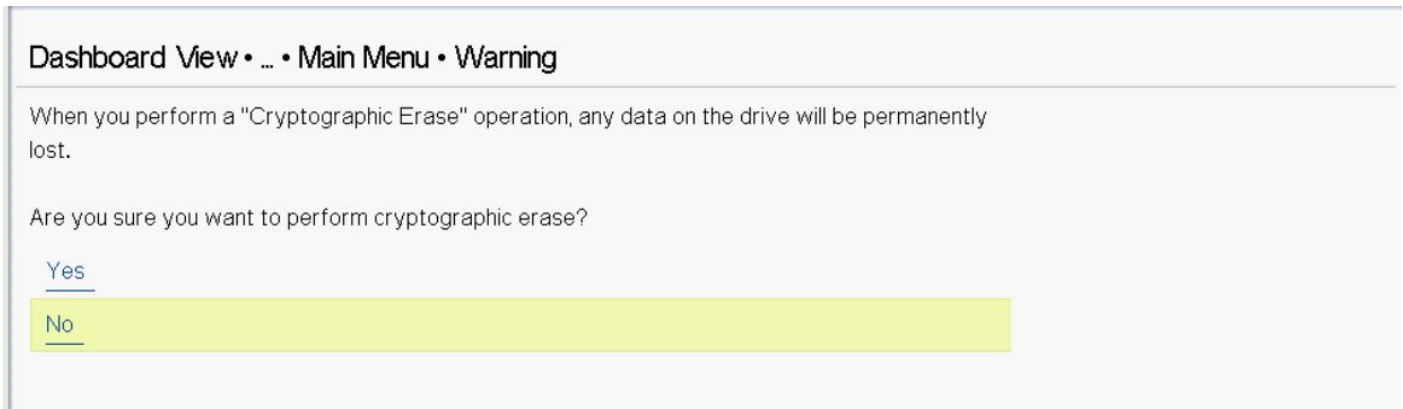
Alternatively, you can use the **Check All** and **Uncheck All** options at the bottom of the list of drives to either select all available drives or clear the selected drives.

Figure 44: Perform Cryptographic Erase on Drives Dialog



3. Click **Go**.

A warning message dialog appears indicating that a cryptographic erase permanently erases any data on the drive.

Figure 45: Perform Cryptographic Erase on Drives Warning Dialog

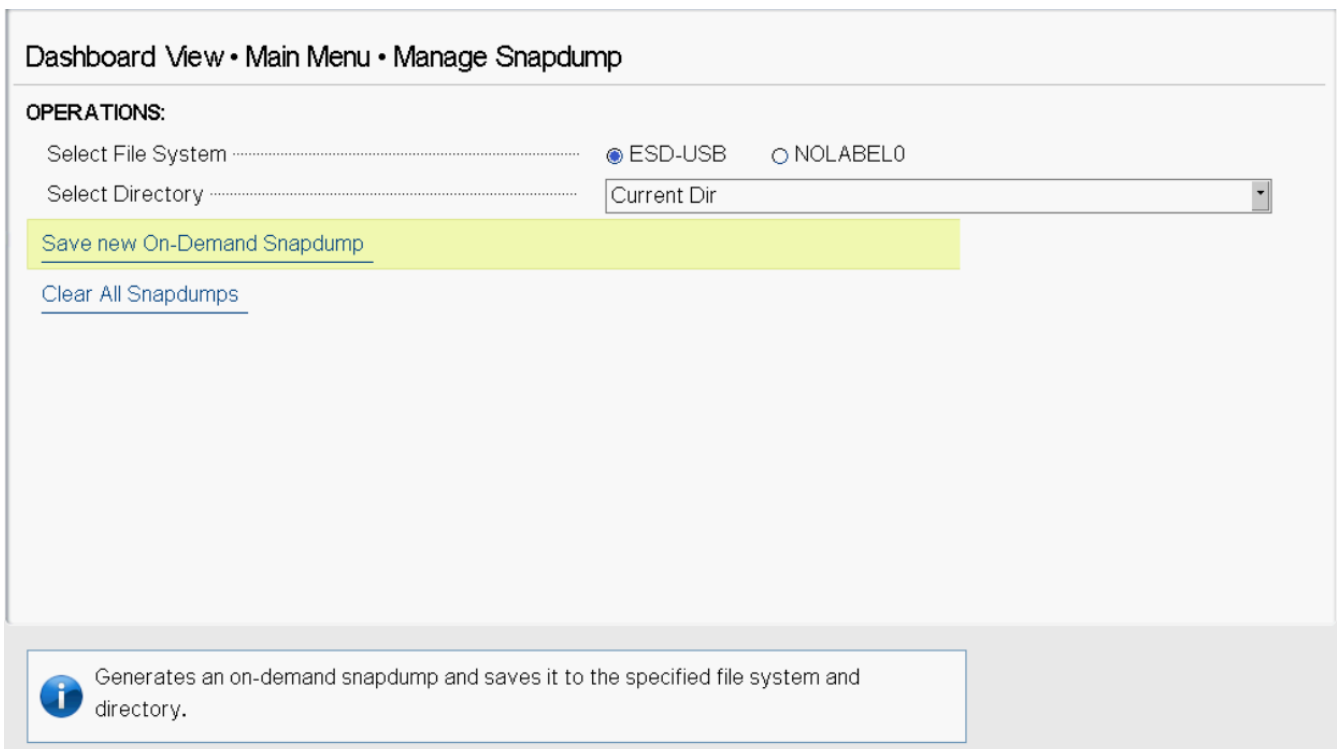
- To confirm the cryptographic erase, click **Yes**.

HII issues the command to the firmware for each drive and displays the result to the user.

Managing Snapdump

The Snapdump feature is a way to save a snapshot of the debug information at fault time. The intention is to collect all required information to be able to find a root cause of the defect at the first instance of defect detection. This ensures that multiple defect reproductions are not required for debugging.

The following dialog appears when you select **Manage Snapdump** on the **Advanced Controller Management** dialog.

Figure 46: Manage Snapdump Dialog

If a file system is not detected and if the firmware does not allow clearing Snapdumps, then the following message appears.

Snapdumps cannot be saved because the file system is not available.

Managing SAS Storage Link Speed

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.

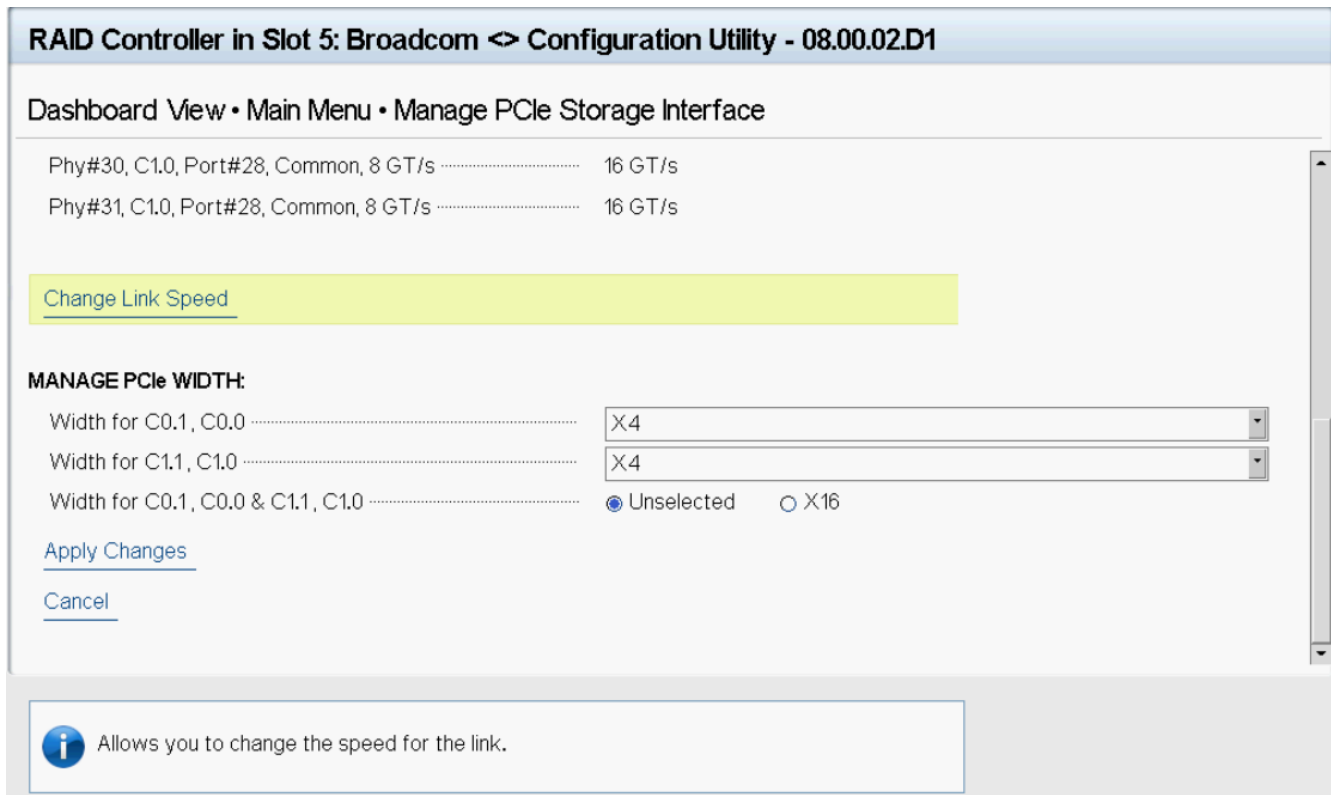
Follow these steps to change the link speed for one or more phys:

1. Use the radio buttons to select a link speed for each phy.
The link speed values are **6 GB/s**, **12 GB/s**, and **22.5 GB/s**.
2. Scroll to the bottom of the phy list, highlight **OK**, and press **Enter**.

Managing PCIe Storage Interface

The manage PCIe storage interface feature 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 47: Manage PCIe Storage Interface Dialog



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

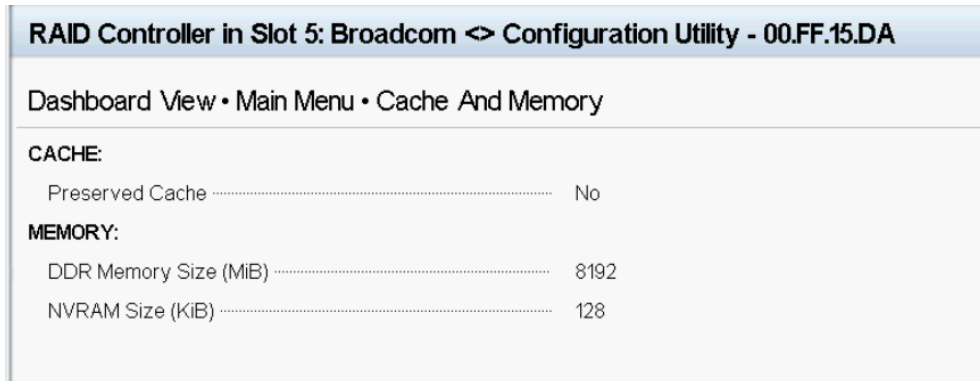
1. Highlight the field to the right of the phy number and press **Enter**.
2. Select an option from the popup menu.
The link speed values are **Unknown**, **2.5 GT/s**, **5 GT/s**, **8 GT/s**, and **16 GT/s**.

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

Setting Cache and Memory Properties

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

Figure 48: Cache and Memory Dialog



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.

Running a Patrol Read

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

Figure 49: Patrol Read Dialog



RAID Controller in Slot 5: Broadcom <Product Name> Configuration Utility - 08.00.04.00

Dashboard View • Main Menu • Patrol Read

[Start](#)


State Stopped

Iterations Completed 1

Rate

Corrects Unconfigured Areas Disabled Enabled

[Apply Changes](#)

 Starts patrol read for the controller.

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 the patrol read properties.

NOTE

You can only view the properties and options that are supported by your controller.

1. 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.
2. To select a patrol read setting for unconfigured space, highlight **Corrects Unconfigured Areas**, and press **Enter**.
Select either **Enabled** or **Disabled** and press **Enter**.
3. 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.

Setting Emergency Spare Properties

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

Figure 50: Spare Dialog

RAID Controller in Slot 5: Broadcom <> Configuration Utility - 00.FF.15.DA

Dashboard View • Main Menu • Spare

[Apply Changes](#)

Emergency Spare Unconfigured Good and Global Hot Spare

Emergency Spare for SMARTer Enabled Disabled

Replace Drive Enabled Disabled

Replace Drive on HDD SMART Error Enabled Disabled

Replace Drive on SSD SMART Error Enabled Disabled

[Apply Changes](#)

Submits the changes made to the entire form.

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

Follow these steps to set emergency spare properties:

- To specify whether it is acceptable to commission otherwise incompatible global hot spare drives, unconfigured good drives or both as emergency hot spare drives, use the **Emergency Spare** drop menu to select a mode. Select any of the following modes.
 - Global Hot Spare**
 - Unconfigured Good**
 - Unconfigured Good and Global Hot Spare**
 - None**
- Select **Enabled** or **Disabled** for each of the following properties.
 - Emergency for SMARTer** – specify whether it is acceptable to commission emergency hot spare drives for PFA events.
 - Replace Drive** – copy data back from a hot spare drive to a physical drive.
 - Replace Drive on HDD SMART Error** – if a Self-Monitoring Analysis and Report Technology (SMART) error is detected on a physical drive, start a Drive Replace operation.
 - Replace Drive on SDD SMART Error** – if a Self-Monitoring Analysis and Report Technology (SMART) error is detected on a physical drive, start a Drive Replace operation.
- Click **Apply Changes**.
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 51: Task Rates Dialog

RAID Controller in Slot 2: Broadcom <MegaRAID 9660-16i> Configuration Utility - 08.01.11.00

Dashboard View > Main Menu > Task Rates

Apply Changes

Background Initialization (BGI) Rate	30
Patrol Read Rate	30
Consistency Check Rate	30
Rebuild Rate	30

Apply Changes

Submits the changes made to the entire form.

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

Follow these steps to change the task rates.

NOTE

You can only view the properties and options that are 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 that are dedicated to running a background initialization of drives on this controller. You can configure the BGI rate between 1 percent and 100 percent. At 1 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 1 to 100 and press **Enter**.
The consistency check rate is the percentage of the compute cycles that are dedicated to running a consistency check on drives on this controller. You can configure the consistency check rate between 1 percent and 100 percent. At 1 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 1 to 100 and press **Enter**.
The patrol read rate is the percentage of the compute cycles that are dedicated to running a patrol read on drives on this controller. You can configure the patrol read rate between 1 percent and 100 percent. At 1 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 1 to 100 and press **Enter**.

The rebuild rate is the percentage of the compute cycles that are dedicated to rebuilding failed drives in virtual drives on this controller. You can configure the rebuild rate between 1 percent and 100 percent. At 1 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 an Online Capacity Expansion (OCE) on a virtual drive, highlight **Reconstruction Rate** and press **Enter**. Specify a number from 1 to 100 and press **Enter**.

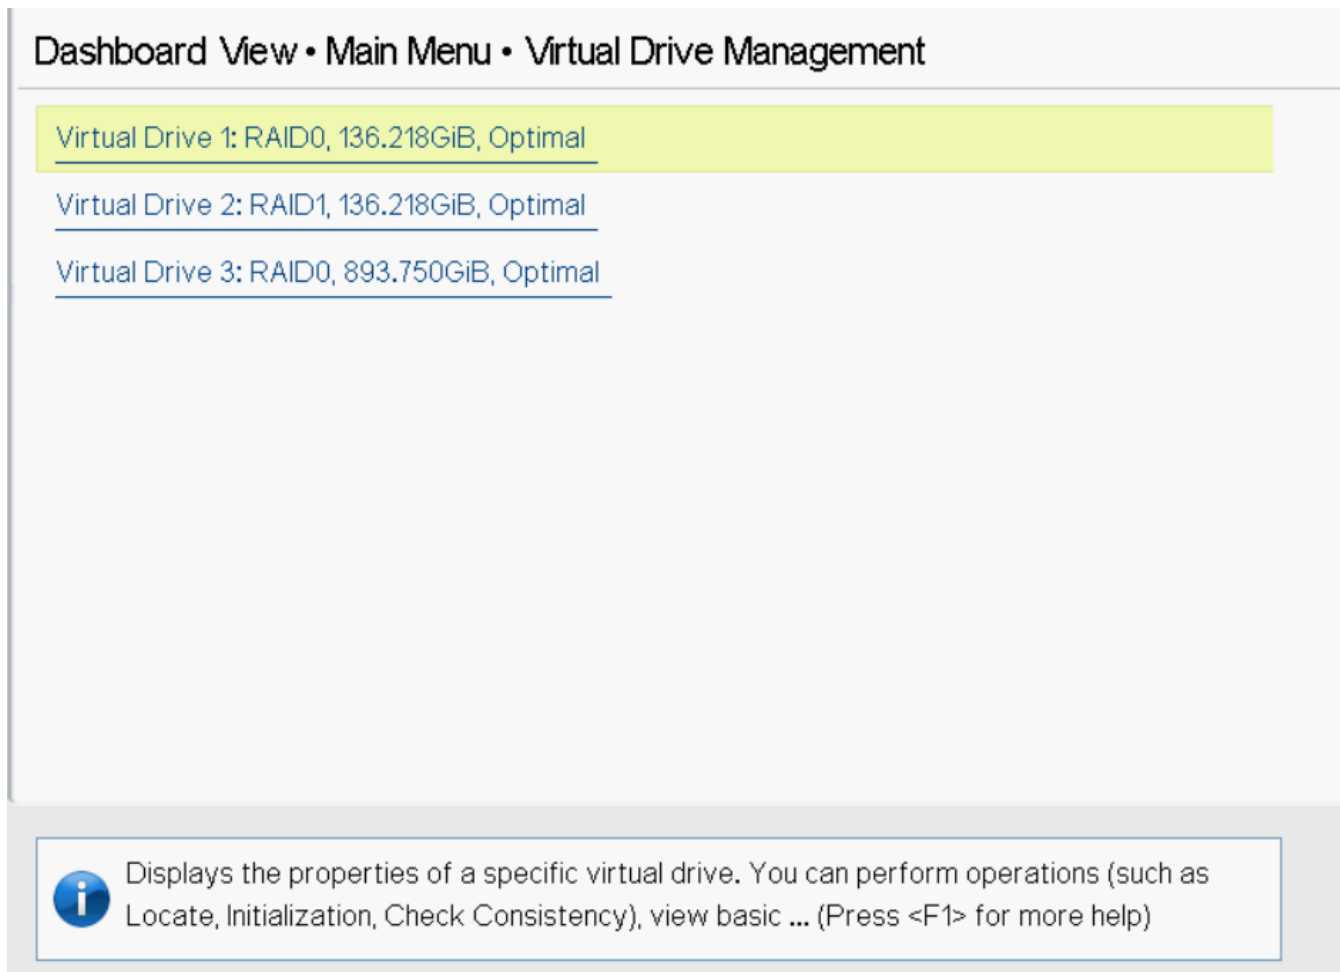
The reconstruction rate is the percentage of the compute cycles that are dedicated to reconstructing data on drives on this controller. You can configure the reconstruction rate between 1 percent and 100 percent. At 1 percent, the reconstruction operation runs only if the firmware is not doing anything 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.

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 52: Virtual Drive Management Dialog

The menu lists the virtual drives that currently exist on the controller. Highlight the virtual drive that you want to manage and press **Enter**. The following dialog appears.

Figure 53: Virtual Drive Management Dialog

Dashboard View • Main Menu • Virtual Drive 1: RAID0, 136.218GiB, Optimal

Operation

BASIC PROPERTIES:

Name

RAID Level RAID0


Status Optimal

Size 136.218 GiB

Drive Group Drive Group# 0

[View Associated Drives](#)

[Advanced...](#)

 Lists the operations that you can perform on a virtual drive.

This dialog lists the following basic virtual drive properties.

Table 25: Basic Virtual Drive Properties

Property	Description
Name	The name that is 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 MiB or GiB . 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.
Drive Group	The name of the drive group.

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 54: Virtual Drive Operations Popup Menu



Other options, such as **Secure Virtual Drive** and **Check Consistency**, might also appear, depending on the current configuration of the system.

Highlight the operation that 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.

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 **Full 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 can 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 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 RAID 0.

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.

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 that are 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](#).

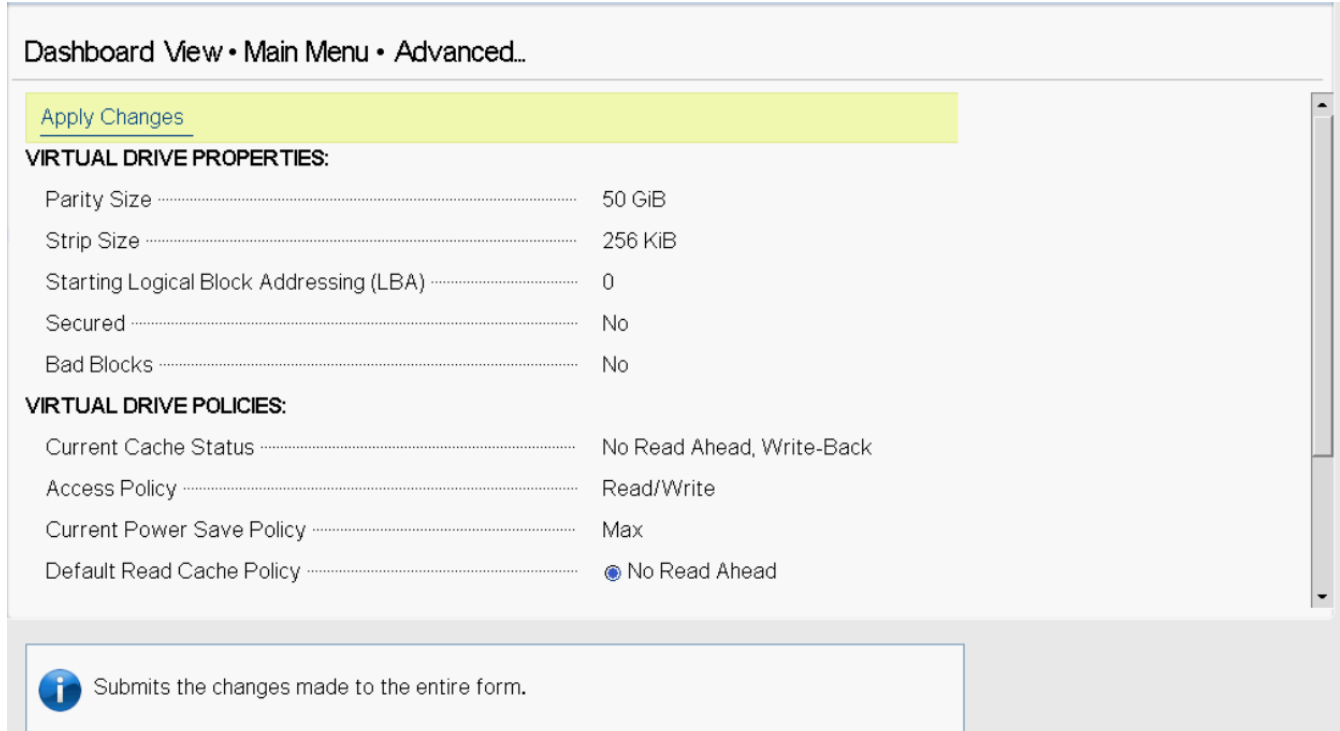
Viewing and Managing Virtual Drive Properties and Options

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 that are 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 55: Advanced Virtual Drive Properties Dialog



The following table describes the virtual drive properties that are listed in this dialog.

Table 26: Virtual Drive Properties

Property	Description
Parity Size	The size of the parity data on the virtual drive.
Strip Size	The size of the stripe that resides on the virtual drive.
Starting Logical Block Addressing (LBA)	The address of the first location of a block of data stored on the virtual drive.
Secured	Indicates whether the virtual drive is secured.
Bad Blocks	Indicates whether the virtual drive has bad blocks.

Following the virtual drive properties that are 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 27: Virtual Drive Policies

Property	Description
Current Cache Status	Displays the current cache policy. The possible values are as follows: <ul style="list-style-type: none"> • Write-Through 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 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. • Always Write Back
Access Policy	The access policy for the virtual drive. The options are Read/Write , Read Only , and Blocked .
Current Power Save Policy	Displays the current power save policy.
Default Read Cache Policy	Displays the read cache policy for the virtual drive. For any profile, if the drive is an SSD drive, the No Read Ahead and Read Ahead options are displayed. However, No Read Ahead is the <i>default</i> read policy. The possible options are as follows: <ul style="list-style-type: none"> • Default A virtual drive property that indicates whether the default read policy is Read Ahead or No Read Ahead. • 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.
Default Write Cache Policy	Displays the default write cache policy of the virtual drive.
Default Power Save Policy	If active, displays the default power save policy.

Property	Description
Drive Write Cache Policy	The drive write cache policy is the disk write cache policy for the individual drives in the virtual drive. The possible values are as follows: <ul style="list-style-type: none"> • Unchanged • Enable • 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.

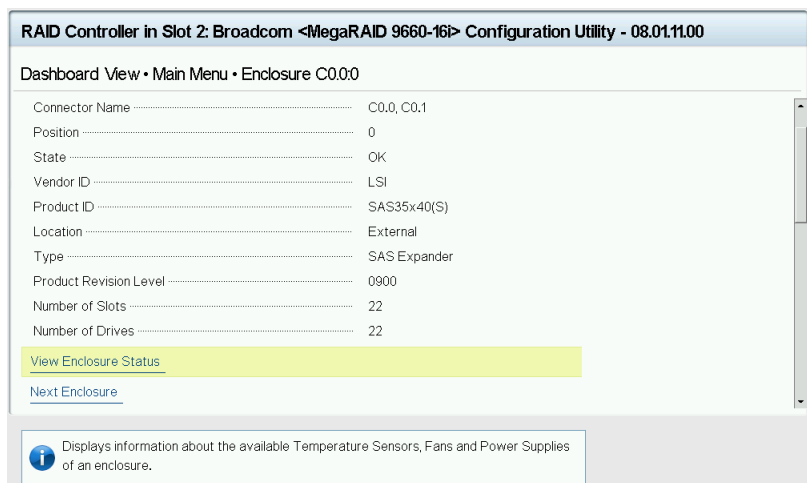
Managing Devices

When you select **Device Management** on the **Main Menu**, the **Device Management Selection** dialog appears.

To manage enclosures and view enclosure properties, select **Logical Enclosure** or **Enclosure** from the **Device Management** menu.

The **Logical Enclosure** dialog appears.

Figure 56: Enclosure Dialog



While on this page, if the enclosure or drives goes missing and you click the missing drive, a popup message appears stating `An unexpected error has occurred.`

The following table describes the enclosure details.

Table 28: Enclosure Properties

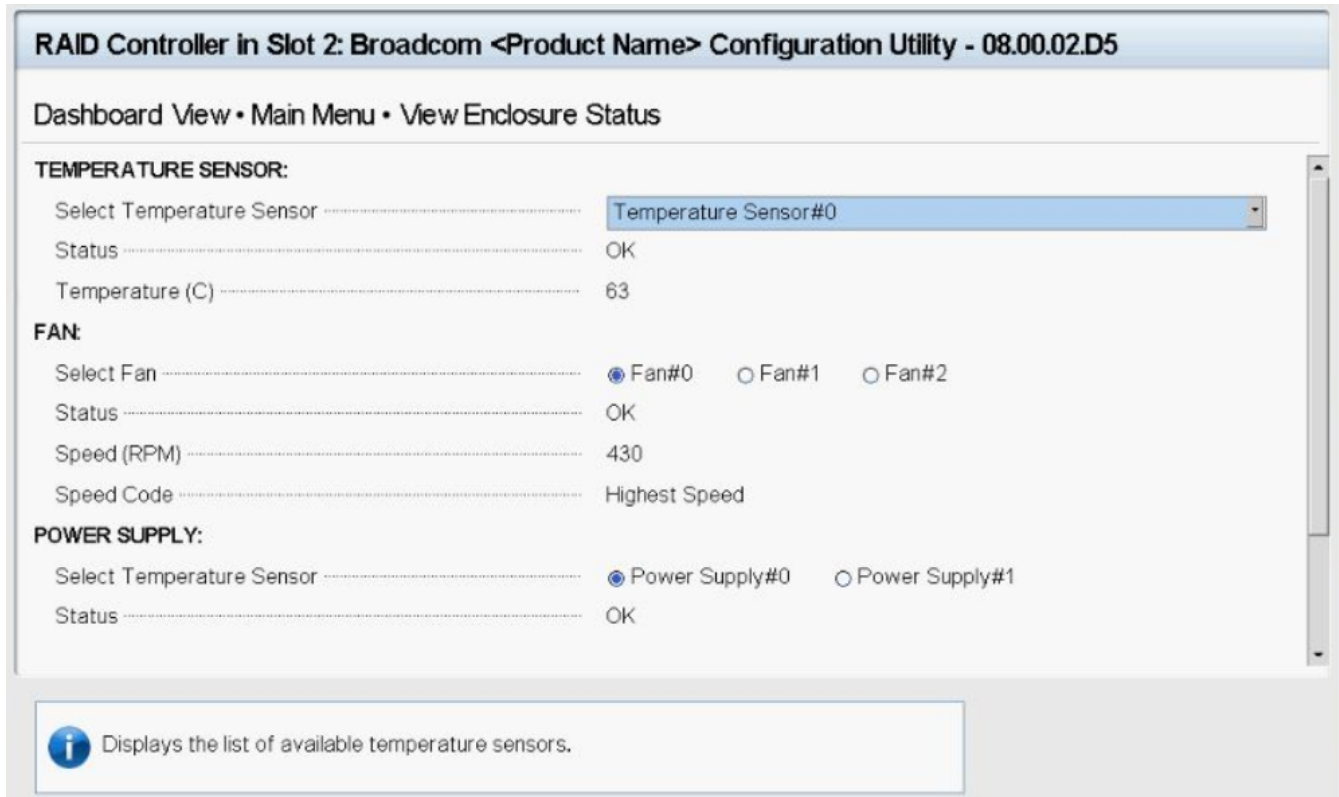
Property	Description
Connector Name	The name of the connector.
Position	The position of the enclosure.
State	The state of the enclosure.
Vendor ID	Manufacturing vendor.
Product ID	Vendor-assigned product ID.
Location	Drive location (internal or external).

Property	Description
Type	Type of enclosure (Virtual SES, SAS Expander).
Product Revision Level	The revision level of the enclosure.
Number of slots	The number of slots in the enclosure.
Number of drives	The number of drives in the enclosure.

To view more information about the enclosure status, select **View Enclosure Status**.

The **View Enclosure Status** dialog appears.

Figure 57: View Enclosure Status Dialog



The **View Enclosure Status** dialog shows information about the temperature sensors, fans, and power supplies installed in the selected enclosure.

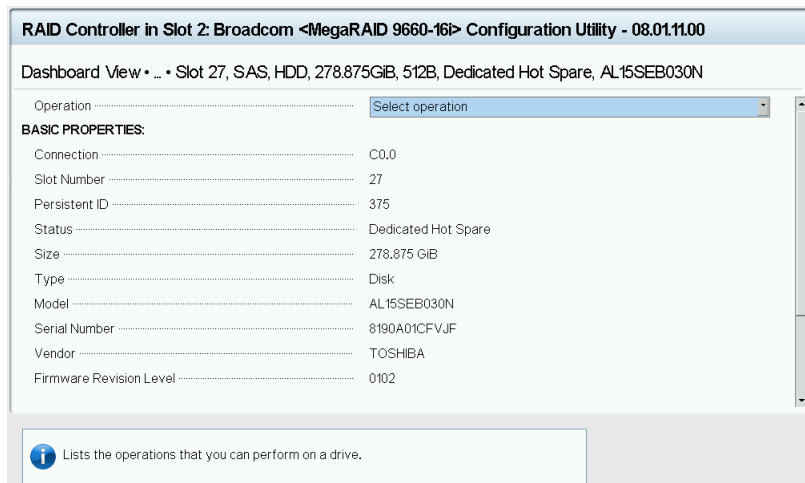
Select a different enclosure from the **ATTACHED DRIVES** list to view the information for a different enclosure.

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

Viewing Physical Drive Properties

To view physical drives properties, select **Logical Enclosure** or **Enclosure** from the **Device Management** menu.

The menu lists all the physical drives that are connected to the controller. Highlight the drive that you want to manage and press **Enter**. The following dialog appears.

Figure 58: Device Management Dialog**Table 29: Basic Physical Drive Properties**


Property	Description
Connection	The connection of the drive.
Slot Number	The slot number where the drive is located.
Persistent ID	The persistent ID.
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.
Serial Number	The serial number of the drive.
Vendor	The hardware vendor of the drive.
Firmware Revision Level	The current firmware revision.
Associated Virtual Drive	The virtual drive associated with the physical drive.

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 59: Select Drive Operations Menu

The 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 are 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 that you want to select, press **Enter** to select it, and return to this dialog.

While on this page, if the drive or enclosure where this drive is attached goes missing and you click the missing drive, a popup message appears stating `An unexpected error has occurred.`

When you perform operations such as Cryptographic Erase on ISE capable drives, when you return to this page after pressing **OK**, you may see fewer operations listed. This is due to the controller firmware (or drive) taking more time to complete the erase operation. To see all of the operations allowed for this drive, return to the previous form and select **Enter**.

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 can 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 **Device 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**.

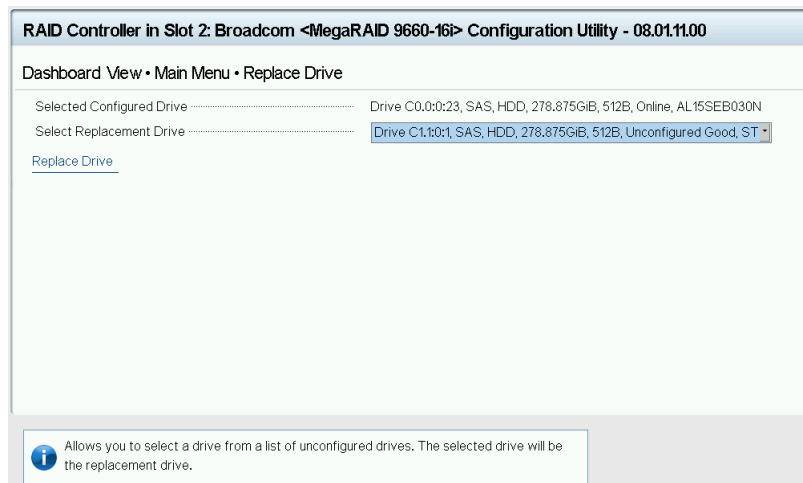
Replacing a Drive

You might want to replace a drive that is a member of a redundant virtual drive that is 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 60: Replace Drive Dialog

3. Select the drop-down menu to the right of **Select Replacement Drive**.

A list of available replacement drives appears.

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 **Device Management** menu. The status of the drive changes from **Online** to **Replace**. You can perform other tasks in the HII Configuration Utility while the replacement operation runs.

Make 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 that you force offline goes into the Unconfigured Good state.

1. Open the popup drive operations menu, highlight **Make Offline** and press **Enter**.

2. Highlight **Go**, which appears beneath **Operation** and press **Enter**.

The **Make 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.

Make Online

Perform these steps to force a selected member drive of a virtual drive online after it has been forced offline.

1. Open the popup drive operations menu, highlight **Make Online** and press **Enter**.

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

The **Make Online** warning appears.

ATTENTION

Do not force a drive that is part of a redundant array online.

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.

Mark 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 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

Global hot spare drives provide protection to redundant virtual drives on the controller. If you select an Unconfigured Good drive, you can 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** 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** and press **Enter**.

The status of the selected drive changes to hot spare.

NOTE

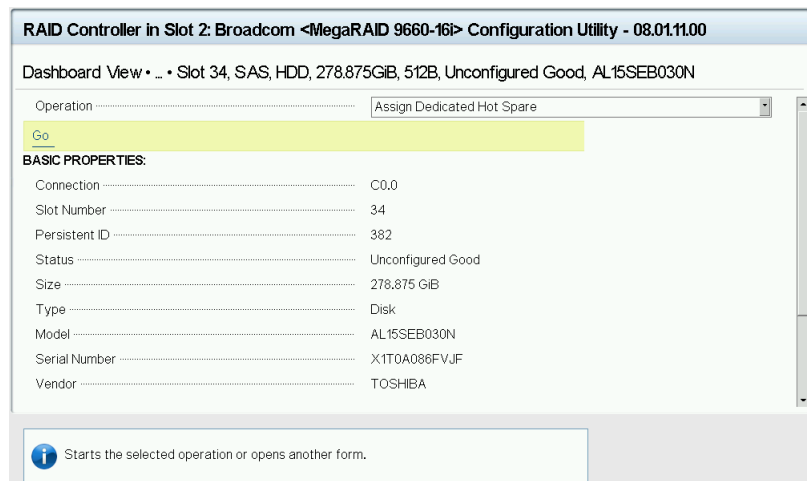
To refresh the status of the drive displayed in the dialog, exit back to the **Main Menu**, then re-enter the **Device Management** dialog.

Assigning a Dedicated Hot Spare

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 can assign it as a dedicated spare drive. Perform these steps to assign a dedicated hot spare.

1. From the **Operation** drop-down, select **Assign Dedicated Hot Spare**.
2. Select **Go**, which appears beneath **Operation**.

Figure 61: Assign Dedicated Hot Spare Dialog



3. Select the drive groups to which this hot spare drive is dedicated.
4. When your selection is complete, highlight **OK**.

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 **Device Management** dialog.

Unassigning a Hot Spare Drive

If the currently selected drive is a hot spare drive, you can unassign and return the drive to an 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 Hot Spare 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 or erase it. Back up any data that you want to keep before initializing or erasing a drive.

1. Open the popup drive operations menu, highlight **Drive Erase** 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 Drive** message appears. (The message is similar to that of erasing a drive.)

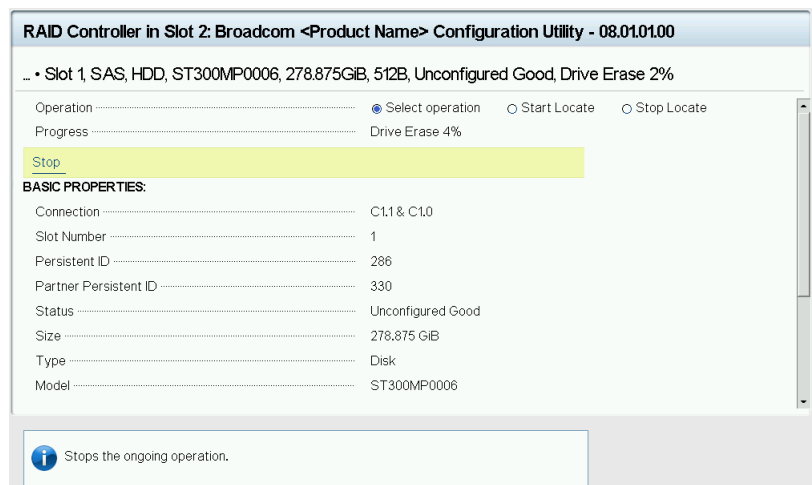
5. Highlight **Confirm** and press the spacebar to confirm the operation.

6. Highlight **Yes** and press **Enter**.

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

7. Highlight **Yes** and press **Enter** to return to the previous window.

This dialog displays a progress indicator that shows the percentage completion of the operation. The dialog also displays a `Stop` command, as shown in the following figure.

Figure 62: 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 Hot Spare**, 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. Back up any data that you want to keep before starting these operations.

Perform these steps to securely erase an FDE-capable drive:

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

A warning dialog 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 you have not created any JBODs, the Make JBOD option appears when you navigate to the **<Select operation>** under the **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.

Viewing Advanced Drive Properties

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

NOTE

Depending on your configuration, some advanced drive properties may not be available.

Figure 63: Advanced Drive Properties Dialog

Dashboard View • Main Menu • Advanced...

SMART Status	Enabled
SAS Address	0x5002538B011F07F2
Interface	SAS
Capable Speed	12.0Gb/s
Negotiated Speed	12.0Gb/s
Capable Link Width	x1
Negotiated Link Width	x1
Number of Connections	1
Cryptographic Erase Capable	Yes
SED Capable	Yes
Temperature (C)	31

The following table describes the entries that are listed on the **Advanced Drive Properties** dialog.

Table 30: Advanced Drive Properties

Property	Description
SMART Status	The SMART status of this device.
SAS Address	The SAS address of this device.
Interface	The interface type of this device.
Capable Speed	The capable speed of this device.
Negotiated Speed	The negotiated speed of this device.
Capable Link Width	The maximum phys/lanes supported by the drive.
Negotiated Link Width	The number of phys/lanes used by the drive.
Number of Connections	Indicates the connection of the drive.
Cryptographic Erase Capable	Indicates whether the drive is cryptographic erase capable.
SED Capable	Indicates whether the drive is encryption capable.
Temperature	The temperature of this device.

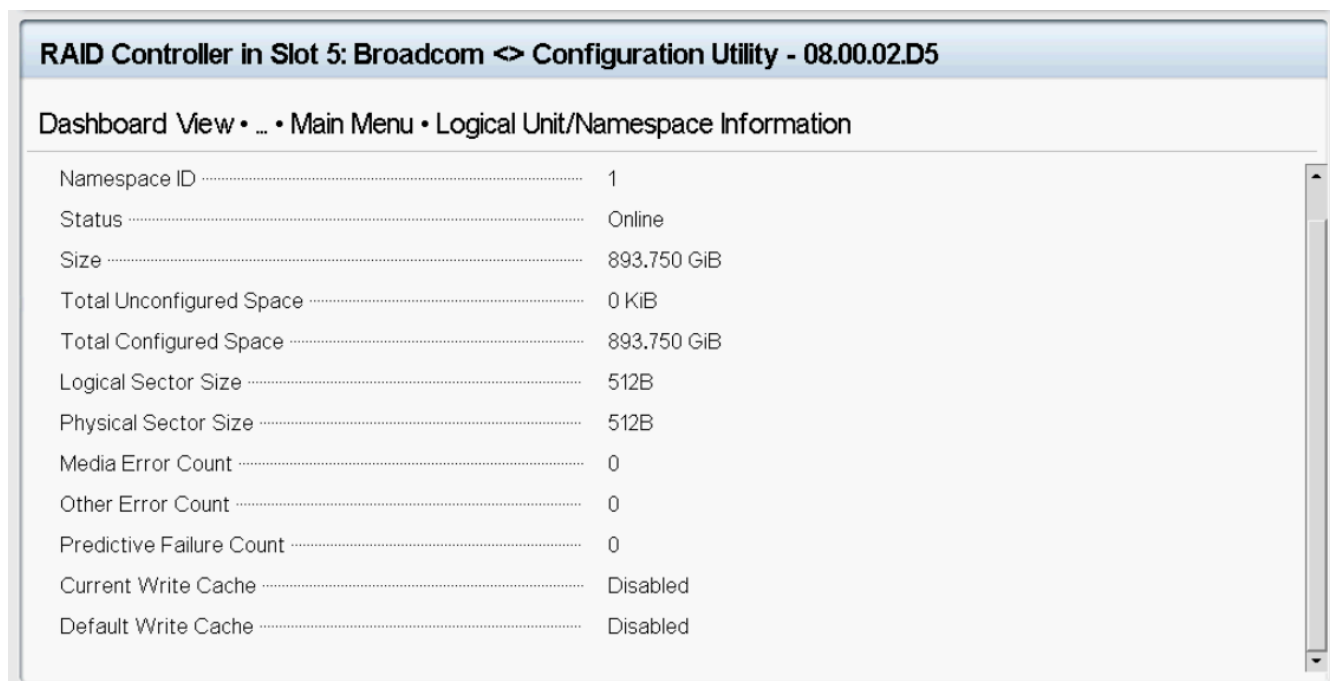
Logical Unit/Namespace Information

The following dialog appears when you select **Logical Unit/Namespace Information** on the **Device Management** menu. The property information in this dialog cannot be modified.

NOTE

Depending on your configuration, some properties may not be available.

Figure 64: Logical Unit/Namespaces Information Dialog



The following table describes the entries that are listed on the **Logical Unit/Namespaces Information** dialog.

Table 31: Logical Unit/Namespaces Information

Property	Description
Namespace ID	The namespace ID of this device (applicable to NVMe only).
Status	The status of this device.
Size	The size of this device.
Total Unconfigured Space	The total unconfigured space.
Total Configured Space	The total configured space.
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 .
Media Error Count	The number of errors that have been detected on the drive.
Other Error Count	The number of other errors that have been detected.
Predictive Failure Count	The predictive failure count.
BBM Error Count	The bad block management errors that have been detected on the disk media.
Firmware Managed Security	Indicates if the device security is managed by the controller firmware.
Secured	Indicates if the device is secured.
Locked	Indicates if the device is locked.
Current Write Cache	The current write cache.
Default Write Cache	The default write cache.

Managing Energy Packs

The following dialog appears when you select **Energy Pack Management** on the **Main Menu**.

Figure 65: Energy Pack Management Dialog

Dashboard View • Main Menu • Energy Pack Management	
Type	Supercap
Status	Optimal
Manufacturer	LSI
Date of Manufacture	06/19/2016
Module Version	B
Serial Number	23078
Design Capacity	6400 milliFarads
Temperature	23 degrees Celsius
Voltage	9728 mV

The following table describes the basic energy pack properties.

Table 32: Energy Pack Management Properties

Property	Description
Type	Type of the energy pack, such as Super Cap.
Status	The status of the energy pack, such as Optimal . The status field has six states. If operation is normal, the state is Optimal. <ul style="list-style-type: none"> • Optimal • Missing • Failed • Degraded • Degraded [Needs Attention] • Unknown
Manufacturer	Manufacturer of the energy pack.
Date of Manufacture	Manufacturing date of the energy pack.
Module Version	Module version of the energy pack.
Serial Number	Serial number of the energy pack.
Design Capacity	Theoretical capacity of the energy pack.
Temperature	Indicates the current temperature and if the current temperature is normal or high.
Voltage	Indicates the current voltage level in mV and if the current voltage is normal or low.

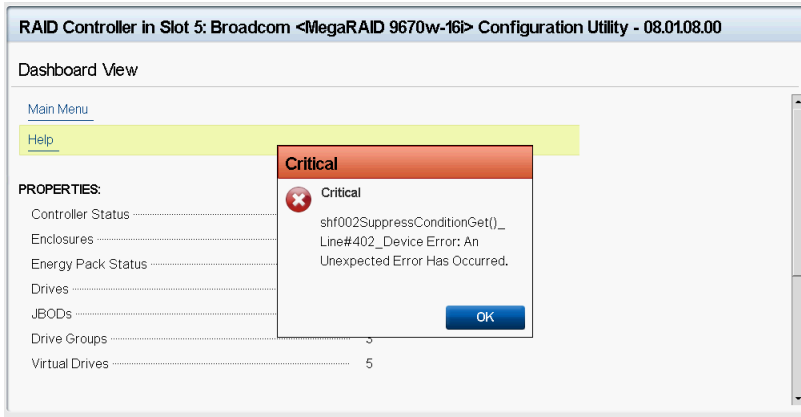
HII Popup Error Protocol

HII displays a popup critical error message when it encounters an error while sending a command to the controller or when a memory allocation fails.

For example, when HII asks for physical drive information the command fails because the drive no longer exists. This failure occurs because the drive was hot swapped or the drive was lost due to a power fluctuation.

The following dialog is an example of a popup error message in the HII Configuration Utility. The message format is:
FunctionName()_Line#_EfiStatus: An Unexpected Error Has Occurred.

Figure 66: HII Popup Error Example



StorCLI2 Utility

The Storage Command Line Tool2 (StorCLI2) is the command line management software designed for the MegaRAID product line.

- [Supported Controllers and Operating Systems](#)
- [Installing StorCLI2 on MegaRAID8](#)
- [StorCLI2 Commands](#)
- [Frequently Used Tasks](#)
- [SAS Address Assignment Rule](#)
- [StorCLI to StorCLI2 Command Conversion](#)

Supported Controllers and Operating Systems

The following topics provide information on supported controllers, operating systems, and default logging.

- [Supported Controllers](#)
- [Supported Operating Systems](#)
- [StorCLI2 Default Logging](#)

Supported Controllers

The StorCLI2 tool supports the following controllers:

- 9600 Family eHBA Adapters
- MegaRAID 9660 Family RAID Adapters
- MegaRAID 9670 Family RAID Adapters

Supported Operating Systems

The following table lists the supported operating systems.

Table 33: Supported Operating Systems

Supported Operating Systems	Version/Flavors
Microsoft	Microsoft Windows versions <ul style="list-style-type: none"> Windows 10 21H1 Microsoft Windows Client versions <ul style="list-style-type: none"> Windows 11 Client Windows 10 Client (RS5) Microsoft Windows Server versions <ul style="list-style-type: none"> Windows Server 2022 (LTSC) Windows Server 2019 (LTSC)
Linux	Red Hat <ul style="list-style-type: none"> Red Hat Enterprise Linux 8.3 Red Hat Enterprise Linux 8.4 Red Hat Enterprise Linux 8.5 Red Hat Enterprise Linux 8.6 Red Hat Enterprise Linux 9.0 SUSE <ul style="list-style-type: none"> SUSE Linux Enterprise Server 15 SP2 SUSE Linux Enterprise Server 15 SP3 Oracle Linux <ul style="list-style-type: none"> None Fedora <ul style="list-style-type: none"> None CentOS <ul style="list-style-type: none"> CentOS 8.3 CentOS 8.4 CentOS 8.5 CentOS 8.6 CentOS 9.0 Citrix <ul style="list-style-type: none"> None
VMware	<ul style="list-style-type: none"> VMware ESXi 7.0 (Update 2) VMware ESXi 7.0 (Update 3) VMware ESXi 6.7 (Update 3)
FreeBSD	<ul style="list-style-type: none"> FreeBSD 13.0 FreeBSD 12.2
Ubuntu	<ul style="list-style-type: none"> Ubuntu 20.04 LTS Ubuntu 22.04 LTS
Unified Extensible Firmware Interface	UEFI environment

StorCLI2 Default Logging

Default logging functionality has been enabled in StorCLI2. When a default log file is created, the file is saved as `storcli2.log`. Each time default logging occurs, the information is added to the `storcli2.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:

- storcli2.log
- storcli2.log.1
- storcli2.log.2
- storcli2.log.3

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

NOTE

StorCLI2 default logging requires a minimum of 20 MB of free space.

There are two conditions under which StorCLI2 logging occurs.

- When the `storcli2conf.ini` file is present in the same directory as the StorCLI2 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 StorCLI2 binary. In these situations, collect a StorCLI2 log file by placing the `storcli2conf.ini` file in the current running directory.
- When the `storcli2conf.ini` file is not present in the current running directory. Default logging occurs automatically.
- To disable default logging, set to `DEBUGLEVEL=0` in the `storcli2conf.ini`.

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

For example, include the `nolog` option in the `storcli2 /cx show nolog` command to prevent default logging.

Installing StorCLI2 on MegaRAID8

The following topics detail the steps that are required to install the StorCLI2 tool for MegaRAID8 controllers on various operating systems.

- [Installing the StorCLI2 Tool on Microsoft Windows Operating Systems](#)
- [Installing the StorCLI2 Tool on the UEFI Environment](#)
- [Installing the StorCLI2 Tool on Linux Operating Systems](#)
- [Installing the StorCLI2 Tool on VMware Operating Systems](#)

Installing the StorCLI2 Tool on Microsoft Windows Operating Systems

The Windows StorCLI2 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 StorCLI2, and run the tool.

Because Windows PowerShell is not fully supported by the StorCLI2 tool, use either one of the following techniques to run commands in the StorCLI2 tool in Windows PowerShell:

- Enclose commands in double quotation marks; for example,

```
storcli2 "/cx show"
```
- Launch the command prompt from within Windows PowerShell to run the StorCLI2 commands.

NOTE

The StorCLI2 tools must be run with the administrator privileges.

Installing the StorCLI2 Tool on the UEFI Environment

The UEFI StorCLI2 binary is provided in a binary format, and no separate installation is required.

1. Copy the binary file from the Broadcom website or from the CD provided to you onto 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 StorCLI2 commands.

Installing the StorCLI2 Tool on Linux Operating Systems

To install the StorCLI2 tool on Linux operating systems, perform the following steps:

1. Unzip the StorCLI2 tool package.
2. To install the StorCLI2 RPM feature, run the `rpm -ivh <StorCLI2- x.xx-x.x86_64.rpm >` command.
By default, the StorCLI2 tool will be installed in the `/opt/MegaRAID/storcli2` location.
3. To upgrade the StorCLI2 RPM feature, run the `rpm -Uvh <StorCLI2- x.xx-x.x86_64.rpm >` command.

Uninstalling the StorCLI2 Tool on Linux Operating Systems

To uninstall the StorCLI2 tool on Linux operating systems, perform the following steps:

Enter the `rpm -e <StorCLI2 -x.xx-x.x86_64.rpm >` command.

Installing the StorCLI2 Tool on VMware Operating Systems

To install the StorCLI2 tool on VMware operating systems, run the following from the command line:

```
esxcli software vib install -v=<path-to-vib-package>
```

Example:

```
esxcli software vib install -v=/vmfs/volumes/datastore1/StorCli2MN/vmware-esx-StorCli2-1.01.04.vib
```

Uninstalling the StorCLI2 Tool on VMware Operating Systems

Perform the following step to uninstall StorCLI2 on VMware OS:

Enter the `esxcli software vib remove -n =<StorCLI2 package name>` command.

StorCLI2 Commands

The section provides information on the commands that are used in StorCLI2.

- [StorCLI2 Tool Command Syntax](#)
- [System Commands](#)
- [Controller Help Commands](#)
- [Controller Commands](#)
- [Physical Drive Commands](#)
- [Virtual Drive Commands](#)
- [Foreign Configuration Commands](#)
- [Drive Group Commands](#)
- [Controller Power Savings Commands](#)
- [Enclosure Commands](#)
- [Controller Phy Commands](#)
- [Energy Pack Commands](#)
- [PCIe Storage Interface Commands](#)
- [Logging Commands](#)
- [Automated Physical Drive Configurations](#)

StorCLI2 Tool Command Syntax

This section describes the StorCLI2 command syntax and the valid values for each parameter in the general command syntax.

- In large configurations, do not run two instances of the StorCLI2 tool in parallel (at the same time).
- To get the output in JavaScript Object Notation (JSON) format, add `J` at the end of the command syntax. For example:

```
storcli2 /cx show <property1> J
```
- Background operations are blocked in the UEFI environment, and these operations are resumed in the operating system environment.

NOTE

Only the commands listed in the file `Schema_mapping_list.xlsx`, which is included in the `JSON_SCHEMA_FILES.zip`, support the JSON format output.

StorCLI2 is a command line utility tool and 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 may result in the last variable being executed or in a command failure.

The StorCLI2 tool syntax uses the following general format:

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

The StorCLI2 tool supports the object identifiers that are listed in the following table.

Here `x` can be any of the following numbers:

- Object identifier number
- A list of numbers
- A range of numbers

Table 34: Object Identifiers in the StorCLI2 Command Syntax

Object Identifier	Description
No object identifier specified	If no object identifier exists, the command is a system command.
<code>/cx /sasx</code>	This object identifier is for controller <code>x</code> or the controller with <code>sasaddress0x</code> .
<code>/call /sasall</code>	This object identifier is for sending the command to all controllers.

Object Identifier	Description
/cx/vx /sasx/vx	This object identifier is for a virtual drive <i>x</i> on controller <i>x</i> or the controller with <i>sasaddress0x</i> .
/cx/vall /sasx/vall	This object identifier is for all virtual drives on controller <i>x</i> or the controller with <i>sasaddress0x</i> .
/cx/ex /sasx/ex	This object identifier is for an enclosure <i>x</i> on controller <i>x</i> or the controller with <i>sasaddress0x</i> .
/cx/eall sasx/eall	This object identifier is for all enclosures on controller <i>x</i> or the controller with <i>sasaddress0x</i> .
/cx/fall /sasx/fall	This object identifier is for all foreign configurations on controller <i>x</i> or the controller with <i>sasaddress0x</i> .
/cx/ex/sx sasx/ex/sx	This object identifier for the drive is slot <i>x</i> on enclosure <i>x</i> on controller <i>x</i> or the controller with <i>sasaddress0x</i> .
/cx/ex/sall /sasx/ex/sall	This object identifier is for all the drives on enclosure <i>x</i> on controller <i>x</i> or the controller with <i>sasaddress0x</i> .
/cx/dx sasx/dx	This object identifier is for the drive group <i>x</i> on enclosure <i>x</i> on controller <i>x</i> or the controller with <i>sasaddress0x</i> .
/cx/dall sasx/dall	This object identifier is for the all drive groups on enclosure <i>x</i> on controller <i>x</i> or the controller with <i>sasaddress0x</i> .
/cx/ep /sasx/ep	The object identifier is for all the energy packs on controller <i>x</i> or the controller with <i>sasaddress0x</i> .

The StorCLI2 tool supports the following verbs.

Table 35: Verbs in the StorCLI2 Command Syntax

Verb	Description
add	This verb adds virtual drives, JBODs, and so on, to the object identifier.
delete	Deletes an element (for example, VD, or spare).
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.
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.
resume	This verb resumes paused operation.
reset	This verb resets 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.
start	This verb starts an operation.
stop	This verb stops an operation that is in progress. A stopped process cannot be resumed.
suspend	This verb suspends an ongoing operation.
undo	This verb is used to undo the requested operation.

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

System Commands

The StorCLI2 utility supports the system commands described in this section.

System Show Commands

StorCLI2 supports the following system `show` commands:

```
storcli2 show
storcli2 show all
storcli2 show file=<filepath>
storcli2 show ctrlcount
storcli2 get rttDump
```

The detailed description for each command follows.

storcli2 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.

storcli2 show all

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

storcli2 show file

This command shows the version information of the controller flash image file presented.

storcli2 show ctrlcount

This command shows the number of controllers that are connected to the system.

storcli2 get rttDump

This command gets windows rttDump of all the controllers and saves them to individual files.

Input example:

```
storcli2 get rttDump
```

NOTE

This command is only supported on a Windows operating system.

Controller Help Commands

To retrieve the list of commands supported by a specific controller, use the following command.

```
storcli2 /c0 ?", " ... /c1 ?
```

This command lists the supported commands on the controller index-0. Similarly, you can retrieve the command support with the proper index given.

Controller Commands

The StorCLI2 utility supports the controller commands described in this section.

Controller commands provide information and perform actions related to a specified controller.

Controller Show Commands

StorCLI2 supports the following `show` commands:

```
storcli2 /cx show
storcli2 /cx show all [logfile[=filename]]
storcli2 /cx start nvcacheerase
storcli2 /cx start diag [duration=<val>]
storcli2 /cx compare factory defaults
```

The detailed description for each command follows.

storcli2 /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 energy pack information.

Input example:

```
storcli2 /c0 show
```

storcli2 /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 energy pack information.

If you use the `logfile` option in the command syntax, the output is written to the specified file. If you do not specify the file name, then the output is written to the `storcli2.log` file. If you do not use the `logfile` option in the command syntax, the entire output is printed to the console.

Ensure that the file name does not contain a blank space.

Input example:

```
storcli2 /c0 show all logfile=log.txt
```

NOTE

The PCI information displayed as part of the `storcli2 /cx show` and `storcli2 /cx show all` commands is not applicable for the FreeBSD operating system. Hence, the PCI information fields are displayed as N/A.

storcli2 /cx start nvcacheerase

This command erases the NVCache flash on the controller.

Input example:

```
storcli2 /c0 start nvcacheerase
```

storcli2 /cx start diag [duration=<val>]

This command runs the self diagnostic for a specified duration (1 to 120 seconds). By default, the duration value is set to 20 seconds.

Input example:

```
storcli2 /c0 start diag duration=10
```

storcli2 /cx compare factory defaults

This command compares the default controller properties with the current property values.

Input example:

```
storcli2 /c0 compare factory defaults
```

Show and Set Controller Properties Commands

This section provides command information for the `show` and `set` controller properties.

NOTE

You cannot set multiple properties with a single command.

Table 36: Controller Commands Quick Reference

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

storcli2 /cx set nvcacherekey

This command rekeys the NVCache flash on the controller.

Input example:

```
storcli2 /cx set nvcacherekey
```

Show Controller Properties Commands

This section provides command information for the `show` controller properties.

NOTE

You cannot set multiple properties with a single command.

storcli2 /cx show <property>

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

General example output:

```
StorCLI2 /c0 show rebuildrate
CLI Version = 008.0000.0000.0074 Sep 2, 2020
Operating system = Windows Server 2016
Controller = 0
Status = Success
Description = None
```

Controller Properties :

=====

```
-----
Ctrl_Prop      Value
-----
RebuildRate(%)  40
-----
```

General example output using a SAS address:

```
storcli2 /sas0x500062b200000000 show rebuildrate
CLI Version = 008.0002.0000.0003 Sep 14, 2021
Operating system = Linux4.18.0-305.el8.x86_64
Controller = 0X500062B200000000
Status = Success
Description = None
```

Controller Properties :

=====

```
-----
Ctrl_Prop      Value
-----
Rebuild Rate(%)  100
-----
```

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

Table 37: Properties for Show Commands

Cmd	Property Name	Description
show	abortccconerror	Displays the abort consistency check on an error status.
show	alilog	Displays detailed information for components of the controller.
show	aso	Displays the advanced software options status, and provides the controller safeid and key information.
show	autoconfig	Displays the current autoconfiguration mode.
show	autorebuild	Displays the autorebuild status.

Cmd	Property Name	Description
show	BaseEnclLevel	Displays the value the base enclosure level.
show	bgirate	Displays the background initialization rate in percentage.
show	bootmode	Displays the bootmode status.
show	bootwithpreservedcache	Displays the bootwithpreservedcache status.
show	CacheOffloadEncType	Displays the cache offload encryption type status.
show	coercionmode	Displays the drive coercion mode status.
show	cc consistencycheck	Displays the cc consistencycheck (consistency check) information.
show	ccrate	Displays the consistency check rate in percentage.
show	drivewceforrebuild	Displays the drive write cache setting during rebuild.
show	eccbucketsize	Displays the size of the ECC single-bit-error bucket.
show	eccbucketleakrate	Displays the value of the leak rate of the single-bit bucket in minutes.
show	energypackwarning	Displays the energypackwarning status.
show	es	Displays the emergency global hot spare (GHS) and Emergency Unconfigured Good (UG) status.
show	esSMARTer	Displays the emergency SMARTer status.
show	events	Displays the events information. type=<blocking nonblocking sincereboot sinceshutdown includedeleted latest=x ccincondvd=<0, 1, ...> filter=<info warning critical fatal> file=<filepath>
show	eventseqinfo	Displays the event sequence information and status.
show	exposeencldevice	Displays the expose enclosure device status.
show	failonsmarterror	Displays the failonsmarterror status.
show	fwjbodsecurity	Displays the fwjbodsecurity status.
show	jbodsesmgmt	Displays the jbodsesmgmt status.
show	maintainjbodfailhistory	Displays the maintainjbodfailhistory status.
show	maintainpdfailhistory	Displays the maintainpdfailhistory status.
show	name	Displays the name of the controller.
show	ocerate	Displays the ocerate status.
show	ocr	Displays the online controller reset (ocr) status.
show	pdfaileventoptions	Displays the details of the PD predictive failures event information.
show	pci	Displays the PCI information.
show	pdtempoll	Displays the pdtempoll status.
show	personality	Displays the current, supported, and requested personalities. It also displays the current behavior and respective behavior parameters.
show	prcorrectunconfiguredareas	If set to On, corrects the media errors during a patrol read by writing 0s. If set to Off, unconfigured areas will be unchanged.
show	preservedcache	Displays a list of VDs that have pinned cache.
show	pr patrolRead	Displays the pr patrolRead status.

Cmd	Property Name	Description
show	prrate	Displays the patrol read rate of the virtual drives in percentage.
show	ps	Displays the Dimmer Switch information.
show	rebuildrate	Displays the rebuild rate of the drive in percentage.
show	replacedrive	Displays the <code>replacedrive</code> status.
show	security keyid	Displays the <code>security keyid</code> status.
show	snapdump	Displays the <code>snapdump</code> information.
show	sesvpdassociation	Displays the VPD association type for SES in a multipath configuration.
show	smartpoll	Displays whether SMART polling is on or off for PDs.
show	smartpollinterval	Displays the value of the SMART/temperature poll interval value for internal and external PDs in seconds.
show	spinupdelay	Displays the spin up delay in seconds.
show	spinupdrivecount	Displays the maximum number of drives to spin up at a time.
show	supportssdpatrolread	Displays the <code>supportssdpatrolread</code> status.
show	time	Displays the controller time.
show	unusabledriveinfo	Displays the list of unusable drives and their details.

Set Controller Properties Commands

This section provides command information for the `set` controller properties.

NOTE

You cannot set multiple properties with a single command.

`storcli2 /cx set <property> = <value>`

General example output:

```
storcli2 /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 `storcli2 /cx set <property>=<value>` command structure.

NOTE

Setting a property to `on` enables that feature, and setting a property to `off` disables that feature.

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

Table 38: Properties for Set Commands

Cmd	Property Name	Set Command Range	Description
set	abortcconererror	= [on off]	Aborts consistency check when it detects an inconsistency.
set	aso	deactivatetrailkey	Displays the enabled Advanced Software Options.
set	autoconfig factory	—	Sets the current autoconfiguration to factory default settings.
set	autoconfig primary option	= [UGOOD JBOD SecureJBOD R0 SecureR0 R0WB SecureR0WB]	Sets the primary behavior to UGOOD, JBOD, SecureJBOD, R0, SecureR0, R0WB or SecureR0WB.
set	autoconfig secondary option	= [UGOOD JBOD SecureJBOD R0 SecureR0 R0WB SecureR0WB]	Sets the secondary behavior to UGOOD, JBOD, SecureJBOD, R0, SecureR0, R0WB or SecureR0WB.
set	autoconfig immediate option	= [JBOD SecureJBOD R0 SecureR0 R0WB SecureR0WB] drives=all <e:s e:s-x e:s-x, y >	Immediately sets the autoconfig option chosen by the user for the selected drives.
set	autorebuild	= [on off]	Sets the autorebuild to on or off.
set	BaseEnclLevel	—	Sets the base enclosure level.
set	bgirate	= 1 to 100	Sets the background initialization rate in percentage.
set	bootmode	= [COE SMOE]	Sets the controller boot mode to continue on error or safe mode on error.
set	bootwithpreservedcache	= [on off]	Sets the boot with preserved cache to on or off.
set	CacheOffloadEncType	(0) None (1) 256-bit	Sets the cache offload encryption type.
set	cc consistencycheck	factory	This command schedules the consistency check operation to default values.
set	cc consistencycheck	=off	See Consistency Check .
set	cc consistencycheck	=on	This command enables the consistency check (CC) schedule operation.
set	cc consistencycheck	=on starttime=<yyyy/mm/dd hh> execfrequency hours days weeks=<value>	See Consistency Check .
set	cc consistencycheck	[starttime=<<yyyy/mm/dd hh>] [execfrequency hours days weeks=<value>] [maxvd=<value>] [excludevd=x-y, z none]	See Consistency Check .
set	ccrate	= 1 to 100	Sets the consistency check rate in percentage.

Cmd	Property Name	Set Command Range	Description
set	coercionmode	—	Sets the drive capacity in coercion mode. <ul style="list-style-type: none"> (0) – No coercion (1) – 128MiB (2) – 1GiB
set	dat LOSSwarning	= [on off]	Enables or disables the data loss warnings.
set	drive wceforrebuild	= [on off]	Sets the drive write cache setting during rebuild.
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	energy packwarning	= [on off]	Enables (on) or disables (off) energy pack warnings.
set	es	= [on off] ghs ug	Sets the use of a global hot spare or Unconfigured Good drive as an emergency drive to on or off .
set	esSMARTer	=on off	Sets the use of a SMARTer drive as an emergency drive to on or off .
set	exposeencldevice	= [on off]	Enables or disables the device drivers to expose the enclosure devices.
set	factory defaults	—	Restores the factory default settings.
set	failon smartererror	= [on off] pdtype=RAID JBOD	Enables (on) or disables (off) the <i>Fail on SMARTer</i> property.
set	fwjbodsecurity	=<on off>	Sets the firmware managed security on non-RAID physical drives.
set	hostjbodsecurity	=<on off>	Sets the host managed security on JBOD PDs.
set	jbodsesmgmt	=<on off>	Sets the SES management for the JBOD to on or off .
set	maintainjbodfailhistory	=<on off>	Sets the Maintain JBOD Physical Drive Fail History to on or off .
set	name	=<name>	Sets the name of the controller.
set	ocerate	= 1 to 100	Sets the virtual drive configuration OCE rate in percentage.
set	ocr	=<on off> type=<all auto>	type=<all> – Enables or disables the online controller reset feature. type=<auto> – Enables or disables the online controller reset feature for controller recovery and firmware update.
set	pr patrolread	factory	Sets the patrol read scheduling options to default values.
set	pr patrolread	=off	See Patrol Read .
set	pr patrolread	=on	Enables the patrol read scheduling on a controller.

Cmd	Property Name	Set Command Range	Description
set	pr patrolread	=on starttime=<yyyy/mm/dd hh> execfrequency hours days weeks=<value>	See Patrol Read .
set	pr patrolread	=[starttime=<yyyy/mm/dd hh>] [execfrequency hours days weeks=<value>] [maxconcurrentpd =<value>] [includessds=<on off>] [excludevd=x-y,z none]	See Patrol Read .
set	pdfaileventoptions	[detectionType=<val>] [correctiveaction=<val>] [errorThreshold=<val>]	See Predictive Failure Monitoring Commands .
set	pdtempoll	=[on off]	Sets the firmware poll setting for the physical drive temperature.
set	personality id	=<val> [force]	Sets the personality to RAID or eHBA . If you switch personalities, you must reboot the system for the changes to take effect.
set	prcorrectunconfigured areas	=[on off]	Correct media errors during patrol read by writing 0s to unconfigured areas of the disk.
set	prrate	= 1 to 100	Sets the patrol read rate of the virtual drives in percentage.
set	ps	=OFF type=UG HS all	See Controller Power Savings Commands .
set	ps	=ON type=UG HS properties	See Controller Power Savings Commands .
set	ps [properties]	—	See Controller Power Savings Commands .
set	rebuildrate	= 1 to 100	Sets the rebuild rate of the drive in percentage.
set	replacedrive	=<on off> type=ctrl smartssd smarthdd all	See Controller Replacedrive Commands .
set	security	securitykey=<xxxx> [passphrase=<xxxx>] [keyid=<xxxx>] [file=<filename>]	See Controller Security Commands .
set	security passphrase	<key>	See Controller Security Commands .
set	security rekey	oldsecuritykey=<xxxx> securitykey=<xxxx> [passphrase=<xxxx>] [keyid=<xxxx>] [file=<filename>]	See Controller Security Commands .
set	security rekey	securitykey=<xxxx> [passphrase=<xxxx>] [keyid=<xxxx>] [file=<filename>]	See Controller Security Commands .
set	security rekey useEKMS	—	See Controller Security Commands .

Cmd	Property Name	Set Command Range	Description
set	security rekey useekms	oldsecuritykey=<xxxx> [file=<filename>]	See Controller Security Commands .
set	security useekms		See Controller Security Commands .
set	sesvpdassociation	=<lun targetport>	Sets the VPD association type for SES in a multipath configuration. =lun – Association type based on the LUN. =targetport – Association type based on the target port.
set	smartpoll	=<on off> pdtype=RAID JBOD	Sets the SMART poll to On or Off for RAID and JBOD PDs.
set	smartpollinterval	=<value> pdtype=Internal External	Sets the value of SMART/Temperature Poll Interval for internal and external PDs in seconds. The minimum interval is 5 seconds and maximum interval is 30 minutes (1800 seconds).
set	spinupdrivecount	=<value>	Sets the maximum number of drives to spin up at a time.
set	supportssdpatrolread	=[on off]	Enables (on) or disables (off) patrol read for SSD drives.
set	time	=systemtime	Sets the controller time to the system time.

Controller Background Task Operation Commands

The StorCLI2 utility supports the controller commands background task operation commands described in this section.

Rebuild Rate

The StorCLI2 utility supports the following rebuild rate commands:

```
storcli2 /cx set rebuildrate=<value>
storcli2 /cx show rebuildrate
```

The detailed description for each command follows.

storcli2 /cx set rebuildrate=<value>

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

Input example:

```
storcli2 /c0 set rebuildrate=30
```

NOTE

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

storcli2 /cx show rebuildrate

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

Input example:

```
storcli2 /c0 show rebuildrate
```

Patrol Read

The StorCLI2 utility supports the following patrol read commands:

```
storcli2 /cx set pr|patrolread=off
storcli2 /cx set pr|patrolread factory
storcli2 /cx set pr|patrolread=on
storcli2 /cx set patrolread=on [starttime=<yyyy/mm/dd hh>] [ExecFrequency hours/days/dd hh=<value>]
storcli2 /cx set patrolread [starttime=<yyyy/mm/dd hh>] [ExecFrequencyhours/days/dd hh=<value>]
[maxconcurrentpd=<value>] [includessds=<on|off>] [excluded=x,y,z|none]
storcli2 /cx set prrate=<value>
storcli2 /cx show prrate
storcli2 /cx show pr|patrolread
storcli2 /cx start patrolread
storcli2 /cx stop patrolread
storcli2 /cx suspend pr|patrolread
storcli2 /cx resume patrolread
```

NOTE

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

The detailed description for each command follows.

storcli2 /cx set pr|patrolread factory

This command sets the patrol read scheduling options to the default values.

Input example:

```
storcli2 /c0 set patrolread factory
```

storcli2 /cx set pr|patrolread=on

This command enables the patrol read scheduling on a controller.

Input example:

```
storcli2 /c0 set patrolread factory
```

storcli2 /cx set patrolread=on [starttime=<yyyy/mm/dd hh>] [ExecFrequency hours/days/dd hh=<value>]

This command schedules a patrol read operation.

Input example:

```
storcli2 /c0 set patrolread starttime=2012/02/21 00
```

storcli2 /cx set pr|patrolread [starttime=<yyyy/mm/dd hh>] [execfrequency hours|days|weeks=<value>] [maxconcurrentpd =<value>] [includessds=<on|off>] [excludevd=x-y,z|none]

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

Table 39: Set Patrol Read Input Options

Option	Value Range	Description
starttime	A valid date and hour in 24 hours format	Sets the start time in a <i>yyyy/mm/dd hh</i> format.
maxconcurrentpd	Valid number of physical drives present	Sets the number of physical drives that can be patrol read at a single time.
includessds	—	Include SSDs in the patrol read operation.
excludevd	—	Excludes virtual drives from the patrolread. To exclude a particular virtual drives, provide list of the virtual drive IDs (x,y, z format) or the range of virtual drives. If this option is not specified in the command, no virtual drives are excluded. None : When specified the virtual drives are removed, if any were previously excluded.

NOTE

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

Input example:

```
storcli2 /c0 set patrolread starttime=2012/02/21 00
```

storcli2 /cx set patrolread [ExecFrequency=<value>]

This command delays the scheduled patrol read in hours.

Input example:

```
storcli2 /c0 set patrolread ExecFrequency=30
```

storcli2 /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 Frequency**, **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**, **Suspend**, **Aborted**, **Stopped**, or **Unknown**.

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

Input example:

```
storcli2 /c0 show patrolread
```

storcli2 /cx start patrolread

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

Input example:

```
storcli2 /c0 start patrolread
```

storcli2 /cx stop patrolread

This command stops a running patrol read operation.

Input example:

```
storcli2 /c0 stop patrolread
```

NOTE

You cannot resume a stopped patrol read.

storcli2 /cx suspend pr|patrolread

This command suspends a running patrol read operation.

Input example:

```
storcli2 /c0 suspend patrolread
```

storcli2 /cx resume patrolread

This command resumes a suspended patrol read operation.

Input example:

```
storcli2 /c0 resume patrolread
```

Consistency Check

The StorCLI2 utility supports the following commands to schedule, perform, and view the status of a consistency check (CC) operation:

```
storcli2 /cx set cc|consistencycheck [starttime=<yyyy/mm/dd hh>] [execfrequency hours|days|weeks=<value>]
[maxvd=<value>] [excludevd=x-y,z|none]
storcli2 /cx show cc
storcli2 /cx show ccrate
storcli2 /cx set cc|consistencycheck=off
storcli2 /cx set cc|consistencycheck=on starttime=<yyyy/mm/dd hh> execfrequency hours|days|weeks=<value>
storcli2 /cx set cc|consistencycheck=on
storcli2 /cx set cc|consistencycheck factory
```

The detailed description for each command follows.

**storcli2 /cx set consistencycheck[cc=[on|off]|sequential|concurrent][execfrequency=value][maxvd=<value>]
[starttime=<yyyy/mm/dd hh>] [excludevd=x-y,z|none]**

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

Table 40: Set CC Input Options

Option	Value Range	Description
cc	off or on	Sets CC to either sequential mode or concurrent mode, or turns off the CC. The concurrent mode slows I/O processing.
execfrequency	Hours, days, or weeks. These input options should be supported by the firmware.	Sets the execution frequency of a scheduled consistency check operation.
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 virtual drive IDs or None .	Excludes virtual drives from consistency checks. To exclude particular virtual drives, provide a list of virtual drive IDs (x, y, z format), or the range of virtual drives to exclude from a consistency check (x-y format). If this option is not specified in the command, no virtual drives are excluded. None : When specified the virtual drives are removed, if any were previously excluded.

Input example:

```
storcli2 /c0 set CC starttime=2022/02/21 00 excludevd=v0-v3
```

storcli2 /cx show cc

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

Input example:

```
storcli2 /c0 show cc
```

storcli2 /cx show cc

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

Input example:

```
storcli2 /c0 show cc
```

storcli2 /cx show ccrate

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

Input example:

```
storcli2 /c0 show ccrate
```

NOTE

A high CC rate slows I/O processing.

storcli2 /cx set cc|consistencycheck factory

This command resets the consistency check to factory settings.

Input example:

```
storcli2 /c0 set cc factory
```

Premium Feature Key Commands

The StorCLI2 utility supports the following commands for premium feature keys:

```
storcli2 /cx show aso
storcli2 /cx set aso key=<value> [trial][preview]
```

The detailed description for the command follows.

storcli2 /cx show aso

This command shows the advanced software options (ASO) for a controller.

Input example:

```
storcli2 /c0 show aso
```

storcli2 /cx set aso key=<value> [trial][preview]

This command is used to apply a key to enable Advanced Software Options (ASO) on a controller. You can use the following options with the advanced software options command.

Table 41: Set Advanced Software Options Input Options

Option	Description
key	Key to activate ASO on the controller. After they are activated, ASOs cannot be removed from the controller.
trial	Applies the trial activation key. This is for feature evaluation purposes only.
preview	Displays the preview of the advanced software options that gets enabled after applying this key on the controller.

Input example:

```
storcli2 /c0 set Aso key=LSI0000
```

Controller Security Commands

The StorCLI2 utility supports the following controller security commands:

```
storcli2 /cx set security { securitykey=xxxx [passphrase=xxxx] [keyid=xxxx] } | file=<filename>
storcli2 /cx set security rekey { oldsecuritykey=xxxx securitykey=xxxx [passphrase=xxxx] [keyid=xxxx] } |
  file=<filename>
storcli2 /cx set security rekey { securitykey=xxxx [passphrase=xxxx] [keyid=xxxx] } | file=<filename>
storcli2 /cx set security rekey useEKMS { oldsecuritykey=xxxx } | file=<filename>
storcli2 /cx set security useekms
storcli2 /cx set security rekey useekms
storcli2 /cx set security rekey passphrase=<key>
storcli2 /cx set security {passphrase=<key>}|file=<filename>
storcli2 /cx show fwjbodsecurity
storcli2 /cx set fwjbodsecurity=<on|off>
```

```
storcli2 /cx show hostjbodsecurity
storcli2 /cx set hostjbodsecurity=<on|off>
```

The detailed description for each command follows.

storcli2 /cx show securitykey keyid

This command shows the security key on the controller.

Input example:

```
storcli2 /c0 show securitykey keyid
```

storcli2 /cx set securitykey < keyid=xxxx | file=filename>

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

Input example:

```
storcli2 /c0 set securitykey=Lsi@12345 keyid=1
```

storcli2 /cx set securitykey < =xxxxxxxx [passphrase=xxxx] [keyid=xxxx] | file=filename > [useekms]

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

Table 42: Set Security Key Input Options

Option	Value Range	Description
passphrase	Should have a combination of numbers, uppercase letters, lowercase letters, and special characters. Minimum of 8 characters.	A string that is linked to the controller and is used in the next bootup to encrypt the lock key. If the <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:

```
storcli2 /c0 set securitykey=Lsi@12345 passphrase=Lsi@123456 keyid=1
```

storcli2 /cx set security {securitykey=xxxx [passphrase=xxxx] [keyid=xxxx | file=<filename>

This command changes the security key for the controller.

Input example:

```
storcli2 /c0 set securitykey=Lsi@12345 oldsecuritykey=pass123 passphrase=Lsi@123456 keyid=1
```

storcli2 /cx set security { passphrase=<key> } | file=filename

This command sets the passphrase of the controller.

Input example:

```
storcli2 /c0 set security passphrase=Lsi@123456
```


storcli2 /cx delete securitykey

This command deletes the security key of the controller.

Input example:

```
storcli2 /c0 delete security securitykey
```

storcli2 /cx show fwjbodsecurity

This command shows the status of the firmware managed security.

Input example:

```
storcli2 /c0 show fwjbodsecurity
```

storcli2 /cx set fwjbodsecurity=<on|off>

This command enables firmware managed security and sets the SafeStore on advanced host PDs.

Input example:

```
storcli2 /c0 set fwjbodsecurity=on
```

storcli2 /cx show hostjbodsecurity

This command shows the status of the host managed security.

Input example:

```
storcli2 /c0 show hostjbodsecurity
```

storcli2 /cx set hostjbodsecurity=<on|off>

This command enables and disables the host managed security on advanced host PDs. If this command is enabled, SED passthrough commands are supported on the JBOD drive.

Input example:

```
storcli2 /c0 set hostjbodsecurity=on
```

Flashing Controller Firmware

The StorCLI2 utility supports the following flashing controller firmware command:

```
storcli2 /cx download file=<filepath> [activationtype=online|offline] [noverchk]
storcli2 /cx reset
storcli2 /cx get activation status
storcli2 /cx delete activation offline
```

The detailed description for each command follows.

storcli2 /cx download file=<filepath> [activationtype=online|offline] [noverchk]

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 43: Flashing Controller Firmware Input Options

Option	Value Range	Description
file	filepath	The absolute file path.
offline	—	Offline activation will be issued post download.
noverchk	—	If a firmware downgrade is required, this option needs to be provided.

Input example:

```
storcli2 /c0 download file=c:\app.rom
```

storcli2 /cx reset

This command resets the controller firmware.

Input example:

```
storcli2 /c0 reset
```

storcli2 /cx get activation status

This command retrieves the the activation status.

Input example:

```
storcli2 /c0 get activation status
```

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.

Retrieving Snapdump Data Commands

The StorCLI2 utility supports the Snapdump commands that follow.

```
storcli2 /cx get snapdump all|id=<val> [norttdump]
storcli2 /cx get snapdump ondemand debugfile=<filename> [norttdump]
storcli2 /cx get snapdump ondemand [force] [norttdump]
storcli2 /cx get snapdump ondemand
storcli2 /cx show snapdump
```

Detailed descriptions for each command follow.

storcli2 /cx get snapdump all|id=<val>

To download a specific Snapdump ID, you must read the ID from the firmware. The StorCLI2 utility keeps writing the data to the file, truncating the file and adding new information.

Input example:

```
storcli2 /c0 get snapdump ID=<val>
```

Where:

- `val` – Specifies the Snapdump ID number.

storcli2 /cx get snapdump all

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

With this command, the file name 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:

```
storcli2 /c0 get snapdump all
```

storcli2 /cx get snapdump ondemand

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 file name 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:

```
storcli2 /c0 get snapdump ondemand
```

NOTE

An interval of 10 minutes is required between two consecutive on-demand Snapdump requests. Snapdump is a high resource operation and in certain cases can lead to I/O timeouts.

storcli2 /cx show snapdump

This command shows if any Snapdumps are present.

Input example:

```
storcli2 /c0 show snapdump
```

Clearing Snapdump Data Commands

The StorCLI2 utility is able to delete all Snapdump data from the firmware.

NOTE

Save all previous Snapdumps, as personality changes and flashing a new firmware package discards all Snapdumps on both DDR and flash.

```
storcli2 /cx delete snapdump [force]
```

```
storcli2 /cx delete snapdump enhanced
```

A detailed description for this command follows.

storcli2 /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 StorCLI2 utility 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:

```
storcli2 /c0 delete snapdump [force]
```

storcli2 /cx delete snapdump enhanced

This command clears the Snapdump data from the firmware.

Input example:

```
storcli2 /c0 delete snapdump enhanced
```

Predictive Failure Monitoring Commands

The StorCLI2 utility supports the following command for drive performance monitoring:

```
storcli2 /cx set pdfaileventoptions {[detectionType=<val>] [correctiveaction=<val>] [errorrthreshold=<val>]}
storcli2 /cx show pdfaileventoptions
```

The detailed description for each command follows.

storcli2 / cx set pdfaileventoptions {[detectionType=<val>] [correctiveaction=<val>] [errorrthreshold=<val>]}

This command provides the current settings of the `pdfaileventoptions` set on the controller and the various options to change these settings.

Input example 1:

```
storcli2 /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:

```
storcli2 /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:

```
storcli2 /c0 set pdfaileventoptions errorrthreshold=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.

storcli2 /cx show pdfaileventoptions

This command displays the current settings and checks the settings if the `set pdfaileventoptions` command is used to change settings.

Input example:

```
storcli2 /c0 show pdfaileventoptions
```

Controller Replacedrive Commands

The StorCLI2 utility supports the following commands for controller replacedrive:

```
storcli2 /cx show replacedrive
storcli2 /cx set replacedrive=<on|off> type=ctrl|smartssd|smarthdd|all
```

The detailed description for each command follows.

NOTE

In the replacedrive commands, `cx/ex/sx` indicates the source drive and `eid:sid` indicates the target drive.

NOTE

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

storcli2 /cx show replacedrive

This command displays the replacedrive status.

Input example:

```
storcli2 /c0 show replacedrive
```

storcli2 /cx set replacedrive=on type=ctrl

This command sets a control replacedrive operation.

Input example:

```
storcli2 /c0/e25/s4 set replacedrive type=ctrl
```

storcli2 /cx set replacedrive=on type=smarthdd

This command sets a smart HDD replacedrive operation.

Input example:

```
storcli2 /c0/e25/s4 show replacedrive type=smarthdd
```

storcli2 /cx set replacedrive=on type=smartssd

This command sets a smart SSD replacedrive operation.

Input example:

```
storcli2 /c0/e25/s4 show replacedrive type=smartssd
```

storcli2 /cx set replacedrive=on type=all

This command sets a replacedrive operation.

Input example:

```
storcli2 /c0/e25/s4 show replacedrive type=all
```

Drive Performance Monitoring Commands

The StorCLI2 utility supports the following commands for drive performance monitoring:

```
storcli2 /cx start DPM [delay=<value>] [maxconcurrentpd =<value>] drives=e:s|e:s-x|e:s-x,y
storcli2 /cx stop DPM
storcli2 /cx/ex/sx show DPM type = HIST | LCT | RA | EXT
storcli2 /cx delete DPM
storcli2 /cx get DPM status
storcli2 /cx get DPM config
storcli2 /cx set DPM [duration=<val>] [raFactor=<value>]
```

The detailed description for each command follows.

storcli2 /cx start DPM [delay=<value>] [maxconcurrentpd =<value>] drives=e:s|e:s-x|e:s-x,y

This command starts the drive performance monitoring on the requested drive.

Table 44: Drive Performance Monitoring Options

Option	Description
delay	The delay in seconds before the operation starts.
maxconcurrentpd	The number of PDs to run concurrently (0 for maximum).

NOTE

If the options are not included in the command, the default value is 0.

Input example:

```
storcli2 /c0 start DPM
```

storcli2 / cx stop DPM

This command stops the drive performance monitoring on the controller.

Input example:

```
storcli2 /c0 stop DPM
```

storcli2 /cx/ex/sx shows dpmstat type = HIST | LCT | RA | EXT | All

This command displays the requested drive performance monitoring data of the respective drive.

Where:

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

storcli2 /cx delete DPM

This command deletes the drive performance monitoring data on the controller.

Input example:

```
storcli2 /c0 delete DPM
```

storcli2 /cx get DPM status

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

Input example:

```
storcli2 /c0 get DPM status
```

storcli2 /cx get DPM config

This command displays the drive performance monitoring configuration on the controller.

Input example:

```
storcli2 /c0 get DPM config
```

storcli2 /cx set DPM [duration=<val>] [raFactor=<value>]

This command sets the drive performance monitoring configuration on the controller.

Table 45: Drive Performance Monitoring Configuration Options

Option	Description
duration	The total duration the PD list is running (in seconds).
raFactor	The divisor used when calculating the running average.

Input example:

```
storcli2 /c0 set DPM
```

SPDM Commands

StorCLI2 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.

```
storcli2 /cx show security spdm
storcli2 /cx show security spdm slotgroup=xx slot=yy
storcli2 /cx export security spdm slotgroup=xx slot=yy subject=subjectname file=filename
storcli2 /cx import security spdm slotgroup=xx slot=yy file=filename
storcli2 /cx set security spdm slotgroup=xx slot=yy invalidate [froce]
storcli2 /cx get security spdm slotgroup=xx slot=yy file=filename
```

storcli2 /cx show security spdm

This command shows the SPDM details for all the slots and slot groups.

storcli2 /cx show security spdm slotgroup=xx slot=yy

This command show the SPDM details for the specific slots and slotgroups.

storcli2 /cx export security spdm slotgroup=xx slot=yy subject=subjectname file=filename

This command exports the CSR for a particular slotgroup or slot combination.

storcli2 /cx import security spdm slotgroup=xx slot=yy file=filename

This command imports the SPDM certificate slot chain for a particular slotgroup or slot combination.

storcli2 /cx set security spdm slotgroup=xx slot=yy invalidate [froce]

This command invalidates the certificate chain storage slot.

storcli2 /cx get security spdm slotgroup=xx slot=yy file=filename

This commands gets the SPDM status.

Physical Drive Commands

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

Table 46: 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 .	Displays physical drive information.

Drive Show Commands

The StorCLI2 utility supports the following drive `show` commands:

```
storcli2 /cx/ex/sx show
storcli2 /cx/ex/sx| show all
storcli2 /cx/ex/sx show smart
storcli2 /cx/ex/sx show phyerrorcounters
storcli2 /cx/ex/sx reset phyerrorcounters
```

The detailed description for each command follows.

storcli2 /cx/ex/sx show

This command displays the summary of the physical drives specified.

Input example:


```
storcli2 /c0/e25/s4 show
```

storcli2 /cx/ex/sx|show all

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

Input examples:

```
storcli2 /c0/e25/s0-3 show all
storcli2 /c0/e25/sall show all
```

storcli2 /cx/ex/sx show smart

This command displays the SMART information of a physical drive.

Input example:

```
storcli2 /c0/e25/s4 show smart
```

storcli2 /cx/ex/sx show phyerrorcounters

This command retrieves the drive phyerrorcounters information.

Input example:

```
storcli2 /c0/e25/s4 show phyerrorcounters
```

storcli2 /cx/ex/sx reset phyerrorcounters

This command resets the drive phyerrorcounters.

Input example:

```
storcli2 /c0/e25/s4 reset phyerrorcounters
```

Missing Drives Commands

The StorCLI2 utility supports the following commands to mark and replace missing physical drives with the specified Unconfigured Good drive:

```
storcli2 /cx/ex/sx set missing
```

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`.

storcli2 /cx/ex/sx set missing

This command marks a drive as missing.

Input example:

```
storcli2 /c0/e5/s4 set missing
```

Set Drive State Commands

The StorCLI2 utility supports the following commands to set the status of physical drives:

```
storcli2 /cx/ex/sx set good [force]
storcli2 /cx/ex/sx set offline
storcli2 /cx/ex/sx set online [force]
storcli2 /cx/ex/sx set missing
storcli2 /cx/ex/sx reset phyerrorcounters
storcli2 /cx/ex/sx set JBOD [force]
storcli2 /cx/ex/sx set uconf [force]
storcli2 /cx/ex/sx set resume current
```

The detailed description for each command follows.

storcli2 /cx/ex/sx set good [force]

This command sets the drive state to good.

Input example:

```
storcli2 /c0/e25/s4 set good
```

storcli2 /cx/ex/sx set offline

This command changes the drive state to offline.

Input example:

```
storcli2 /c0/e25/s4 set offline
```

NOTE

Setting a drive to offline might trigger a hot spare to be commissioned. When this occurs, the offline drive transitions to Unconfigured Good. This transition makes the drive eligible for further use.

storcli2 /cx/ex/sx set online [force]

This command changes the drive state to online.

Input example:

```
storcli2 /c0/e25/s4 set online
```

storcli2 /cx/ex/sx set missing

This command marks a drive as missing.

Input example:

```
storcli2 /c0/e25/s4 set missing
```

storcli2 /cx/ex/sx set JBOD [force]

This command converts the drive to JBOD.

Input example:

```
storcli2 /c0/e25/s4 set jbod
```

storcli2 /cx/ex/sx set uconf [force]

This command converts the drive to an unconfigured drive.

NOTE

If the drive has an operating system or a file system on it, the StorCLI2 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:

```
storcli2 /c0/e25/s4 set uconf
```

storcli2 /cx/ex/sx set resume current

This command resumes an ongoing current process. You can run this command only when an operation is running on a drive.

Input example:

```
storcli2 /c0/e25/s4 set resume current
```

Drive Initialization Commands

When you initialize drives, all the data from the drives is cleared. The StorCLI2 utility supports the following commands to initialize drives:

```
storcli2 /cx/ex/sx show clear
storcli2 /cx/ex/sx start clear
storcli2 /cx/ex/sx stop clear
```

The detailed description for each command follows.

storcli2 /cx/ex/sx show clear

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

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

Input example:

```
storcli2 /c0/e25/s4 show clear
```

storcli2 /cx/ex/sx start clear

This command starts the initialization process on a drive.

Input example:

```
storcli2 /c0/e25/s4 start clear
```

storcli2 /cx/ex/sx stop clear

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

Input example:

```
storcli2 /c0/e25/s4 stop clear
```

Drive Firmware Download Commands

The StorCLI2 utility supports the following command to download the drive firmware:

```
storcli2 /cx/ex/sx download file=<filepath> mode=5|7|E|activatenow|F [chunksize=value]]
```

NOTE

StorCLI2 default logging should be disabled when an expander or drive firmware download is initiated in the presence of I/Os.

storcli2 /cx/ex/sx download file=<filepath> mode=5|7|E |activatenow] | F [chunksize=<value>]]

This command flashes the drive firmware with the specified file.

The `mode` options specify the SCSI write buffer mode. The description follows:

- `/cx` – Specifies the controller, where `x` is the index of the controller.
- `/ex` – Specifies the enclosure ID of the controller.
- `/sx` – Specifies the drive slot ID of the controller.
- `5` – The entire drive firmware file is downloaded at once.
- `7` – The drive firmware file is downloaded in chunks of 32 KB.
- Mode `E` – Downloads the microcode and defers the activation.
- Mode `F` – Activates the deferred microcode and allows you to issue this command to all devices in a safe manner. The default delay time is 60 seconds.

Input example:

```
storcli2 /c0/e25/s4 download file=c:\file.bin mode=7
```

NOTE

Prepare and Complete requests are only sent for Mode 5, 7, and F downloads.

Locate Drives Commands

The StorCLI2 utility supports the following commands to locate a drive and activate the physical disk activity LED:

```
storcli2 /cx/ex/sx start locate
storcli2 /cx/ex/sx stop locate
```

The detailed description for each command follows.

storcli2 /cx/ex/sx start locate

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

Input example:

```
storcli2 /c0/e25/s4 start locate
```

storcli2 /cx/ex/sx stop locate

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

Input example:

```
storcli2 /c0/e25/s4 stop locate
```

Prepare to Remove Drives Commands

The StorCLI2 utility supports the following commands to prepare the physical drive for removal:

```
storcli2 /cx/ex/sx start prepformv1
storcli2 /cx/ex/sx undo prepformv1
```

The detailed description for each command follows.

storcli2 /cx/ex/sx start prepformv1

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

Input example:

```
storcli2 /cx/e25/s4 start prepformv1
```

storcli2 /cx/ex/sx undo prepformv1

Use this command to undo the prepare drive for removal command.

Input example:

```
storcli2 /c0/e25/s4 undoprepformv1
```

Drive Security Command

The StorCLI2 utility supports the following drive security commands:

```
storcli2 /cx/ex/sx show securitykey keyid
```

storcli2 /cx/ex/sx show securitykey keyid

This command shows the security key for secured physical drives.

Input example:

```
storcli2 /c0/e25/s4 show SecurityKey keyid
```

storcli2 /cx/ex/sx set security = on

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

Input example:

```
storcli2 /c0/e25/s4 set security = on
```

Drive Secure Erase Commands

The StorCLI2 utility supports the following drive erase commands:

```
storcli2 /cx/ex/sx start erase type=reprovision [force]
storcli2 /cx/ex/sx start erase type=<simple | normal | thorough | crypto> [aue=<0|1>] [patternA=<val>]
[patternB=<val>] [force]
storcli2 /cx/ex/sx start erase type=sanitize mode=1|2|3|4 [aue=<0|1>] [patternA=<val>] [patternB=<val>]
[overwritecount=<val> invert=<val>] [force]
storcli2 /cx/ex/sx stop erase
```

```
storcli2 /cx/ex/sx show erase
```

The detailed description for each command follows.

storcli2 /cx/ex/sx start erase type=reprovision [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:

```
storcli2 /c0/e25/s4 start erase type=reprovision aue=0 force
```

NOTE

This command deletes data on the drive and the security configuration, and this data is no longer accessible.

This command is applicable for SED capable drives.

storcli2 /cx/ex/sx start erase type= <simple | normal | thorough | crypto> [aue=<0|1>] [patternA=<val>] [patternB=<val>] [force]

storcli2 /cx/ex/sx start erase type=sanitize mode=1|2|3|4 [aue=<0|1>] [patternA=<val>] [patternB=<val>] [overwritecount=<val> invert=<val>] [force]

These commands securely erase 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.

NOTE

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

This command is issued when this erase type is listed under the property **supported erase type** in the drive properties (`/cx/ex/sx show all`).

Table 47: Drive Erase Command Options

Options	Description
simple	Single pass, single pattern write.
normal	Three pass, three pattern write.
thorough	Nine pass, repeats the normal write three times.
standard	Performs standard erase for SSD drives.
threepass	Three pass verify: <ul style="list-style-type: none"> pass1 – Random pattern write. pass2,3 – Write zero and verify.
crypto	Cryptographic erase.
sanitize	Erase type sanitize.
Mode	Modes: <ul style="list-style-type: none"> 1 – Sanitize Overwrite for physical drives. 2 – Sanitize Block Erase for physical drives. 3 – Sanitize FREEZE LOCK for SATA physical drives. 4 – Sanitize. ANTIFREEZE LOCK for SATA physical drives.

Options	Description
patternA	8-bit binary value. Erases pattern A to overwrite the data.
patternB	8-bit binary value. Erases pattern B to overwrite the data.
ae	Allow uninterrupted erase.

Input example:

```
storcli2 /c0/e25/s4 start erase type=thorough aue=0 patternA=10010011 patternB=11110000
```

storcli2 /cx/ex/sx stop erase

This command stops the erase operation.

Input example:

```
storcli2 /c0/e25/s4 stop erase
```

storcli2 /cx/ex/sx show erase

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

Input example:

```
storcli2 /c0/e25/s4 show erase
```

Rebuild Drives Commands

The following commands rebuild drives in the StorCLI2 utility:

```
storcli2 /cx/ex/sx suspend rebuild
storcli2 /cx/ex/sx resume rebuild
storcli2 /cx/ex/sx show rebuild
storcli2 /cx/ex/sx start rebuild
storcli2 /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.

storcli2 /cx/ex/sx suspend rebuild

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

Input example:

```
storcli2 /c0/s4 suspend rebuild
```

storcli2 /cx/ex/sx resume rebuild

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

Input example:

```
storcli2 /c0/s4 resume rebuild
```

storcli2 /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:

```
storcli2 /c0/s4 show rebuild
```

storcli2 /cx/ex/sx start rebuild

This command starts a rebuild operation for a drive.

Input example:

```
storcli2 /c0/s4 start rebuild
```

storcli2 /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:

```
storcli2 /c0/s4/e4 stop rebuild
```

Hot Spare Drive Commands

The following commands create and delete hot spare drives:

```
storcli2 /cx/ex/sx add hotsparedrive {dgs=<n|0,1,2...>}[enclaffinity]
storcli2 /cx/ex/sx delete hotsparedrive
```

The detailed description for each command follows.

storcli2 /cx/ex/sx add hotspare [dgs=<n|0,1,2...>] [enclaffinity]

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

Table 48: 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 this option is not specified, there is no affinity. Affinity cannot be removed after it is set for a hot spare drive.

Input example:

```
storcli2 /c0/e25/s4,5 add hotspare
```

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

Input example:

```
storcli2 /c0/e25/s6,8 add hotspare dgs=0
```

This command sets /c0/e25/s6,8 as dedicated hot spare for disk group 0.

storcli2 /cx/ex/sx delete hotspare

This command deletes a hot spare drive.

Input example:

```
storcli2 /c0/e25/s4,5 delete hotspare
```

NVMe Drive Commands

The StorCLI utility supports the following NVMe drive commands.

```
storcli /cx show unusabledriveinfo
storcli /cx/ex/sx start recovery [force]
storcli /cx/ex/sx stop recovery
```

storcli /cx show unusabledriveinfo

This command displays the list of initialization failed NVMe drives and their details.

Input example:

```
storcli /c0 show unusabledriveinfo
```

storcli /cx/ex/sx start recovery [force]

This command starts the recovery operation on the specified drive.

Where:

force – Deletes all data present on the drive.

Input example:

```
storcli /c0/e25/s4 start recovery
```

storcli /cx/ex/sx stop recovery

This command stops the recovery operation on the specified drive.

Input example:

```
storcli storcli /c0/e25/s4 stop recovery
```

NOTE

If any unusable NVMe drives are detected, the controller state will move to `Need Attention`.

Replacedrive Commands

The StorCLI2 utility supports the following commands for replacedrive:

```
storcli2 /cx/ex/sx start replacedrive target=e:s
storcli2 /cx/ex/sx stop replacedrive
```

```
storcli2 /cx/ex/sx suspend replacedrive
storcli2 /cx/ex/sx resume replacedrive
storcli2 /cx/ex/sx show replacedrive
```

The detailed description for each command follows.

NOTE

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

storcli2 /cx/ex/sx start replacedrive target=e:s

This command starts a replacedrive operation for a drive.

Input example:

```
storcli2 /c0/e25/s4 start replacedrive target=25:8
```

storcli2 /cx/ex/sx stop replacedrive

This command stops a replacedrive operation.

Input example:

```
storcli2 /c0/e25/s4 stop replacedrive
```

storcli2 /cx/ex/sx suspend replacedrive

This command suspends a replacedrive operation.

Input example:

```
storcli2 /c0/e25/s4 suspend replacedrive
```

storcli2 /cx/ex/sx resume replacedrive

This command resumes a suspended replacedrive operation.

Input example:

```
storcli2 /c0/e25/s4 resume replacedrive
```

storcli2 /cx/ex/sx show replacedrive

This command shows the progress of the replacedrive operation in percentage. The estimated time (in minutes) left to complete the operation is also shown.

Input example:

```
storcli2 /c0/e25/s4 show replacedrive
```

Spinup Drive Commands

The StorCLI2 utility supports the following spinup drive command:

```
storcli2 /cx spinup drives=<e:s|e:s-x|e:s-x,y>
```

The detailed description for the command follows.

storcli2 /cx spinup drives= <e:s|e:s-x|e:s-x,y>

This command spins up the requested drive.

Input example:

```
storcli2 /c0 spinup drives=318:0,1
```

Virtual Drive Commands

The StorCLI2 utility supports the following virtual drive commands.

The following table describes frequently used virtual drive commands.

Table 49: Virtual Drives Commands Quick Reference Table

Commands	Value Range	Description
add	See Add RAID Configuration Input Options .	Creates virtual drives.
set	See Add RAID Configuration Input Options and Change Virtual Drive Properties Commands .	Sets virtual drive properties.
show	all : Shows all properties of the virtual drive. See Virtual Drive Show Commands .	Shows virtual drive information.

Add Virtual Drives Commands

The StorCLI2 utility supports the following commands to add virtual drives:

```
storcli2 /cx add vd r[0|1|5|6|10|50|60]
    [Size=<VD1_Sz>,<VD2_Sz>,...|remaining|all] [name=<VDNAME1>,...]
    drives=e:s|e:s-x|e:s-x,y,e:s-x,y,z [PDperArray=x] [secure]
    [pdcache=on|off|default] [ps=default|none]
    [WT|WB|AWB]
    [Strip=<64|256] [AfterVd=X]
    [hotspare = e:s|e:s-x|e:s-x,y]
    [cachebypass=None|All|64|128|256|512|1024] [init=0|1|2] [NoAutoBGI]
storcli2 /cx add vd each r0 [name=<VDNAME1>,...]
    drives=e:s|e:s-x|e:s-x,y [secure] [pdcache=on|off|default]
    [ps=default|none]
    [WT|WB|AWB] [Strip=<64|256>]
    [cachebypass=None|All|64|128|256|512|1024] [init=0|1|2] [NoAutoBGI]
storcli2 /cx/vx show BBMT
storcli2 /cx/vx delete BBMT
```

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

The detailed description for each command follows.

Table 50: Add RAID Configuration Input Options

Option	Value Range	Description
raid	[0 1 5 6 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.	
remaining	—	Considers the remaining space in the drive group.
drives	Valid enclosure number and valid slot numbers for the enclosure.	In <code>e:s e:s-x e:s-x,y</code> : <ul style="list-style-type: none"> <code>e</code> specifies the enclosure ID. <code>s</code> represents the slot in the enclosure. <code>e:s-x</code> is the range convention used to represent slots <code>s</code> to <code>x</code> in the enclosure <code>e</code> (250 characters maximum). 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 block sizes between the <code>[e:s]</code> pairs is not supported.
pdperarray	1–16.	Specifies the number of physical drives per array.
secure	—	Creates security-enabled drives.
pdcache	on off default.	Enables or disables PD cache.
pi	—	Enables protection information.
dimmerswitch	default: Logical device uses controller default power-saving policy. automatic (auto): Logical device power savings are managed by firmware. none: No power-saving policy.	Specifies the power-saving policy. Sets to default automatically.
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.
strip	64, 256	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.
hotspares	Number of spare physical drives present.	Specifies the physical drives that are to be assigned to a disk group for spares.
init	—	Initializes the virtual drive.
noautobgi	—	Prevents a background initialization operation.

Input example:

```
storcli2 /c0 add vd raid50 names=vdname drives=252:0-9 pdperarray=5
```

storcli2 /cx/vx show BBMT

This command shows the bad block management table details for the given VD.

Input example:

```
storcli2 /c0/v2 compare factory defaultsshows BBMT
```

storcli2 /cx/vx delete BBMT

This command deletes the bad block management table for the given VD.

Input example:

```
storcli2 /c0/v2 delete BBMT
```

Delete Virtual Drives Commands

The StorCLI2 utility supports the following virtual drive delete commands:

```
storcli2 /cx/vx|val del
storcli2 /cx/vx|vall del force
storcli2 /cx/vx del [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 uneraser 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.

storcli2 /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:

```
storcli2 /c0/v2 del
```

ATTENTION

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

storcli2 /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:

```
storcli2 /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.

storcli2 /cx/vx del [discardcache] [force]

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

Input example:

```
storcli2 /c0/v2 delete discardcache
```

Virtual Drive Show Commands

The StorCLI2 utility supports the following virtual drive show commands:

```
storcli2 /cx/vx show
storcli2 /cx/vx show all [logfile[=filename]]
```

The detailed description for each command follows.

storcli2 /cx/vx show

This command shows the summary of the virtual drive information.

Input example:

```
storcli2 /c0/v0 show
```

storcli2 /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:

```
storcli2 /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 StorCLI2 utility supports the following commands for preserved cache:

```
storcli2 /cx/vx delete preservedCache [force]
storcli2 /cx show preservedCache
```

The detailed description for each command follows.

storcli2 /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:

```
storcli2 /c0/v1 delete preservedcache
```

storcli2 /cx show preservedCache

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

Input example:

```
storcli2 /c0 show preservedCache
```

Change Virtual Drive Properties Commands

The StorCLI2 utility supports the following commands to change virtual drive properties:

```
storcli2 /cx/vx set name=<namestring>
storcli2 /cx/vx set pdcache=<on|off|default>
storcli2 /cx/vx set rdcache=<ra|nora>
storcli2 /cx/vx set wrcache=<wt|wb|awb>
storcli2 /cx/vx set ps=Default|None
storcli2 /cx/vx set autobgi=On|Off
storcli2 /cx/vx set Unmap=<On|Off>
storcli2 /cx/vx set wsUnmap=<On|Off>
storcli2 /cx/ex/sx set resume current
```

The detailed description for each command follows.

storcli2 /cx/vx set name=<namestring>

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

Input example:

```
storcli2 /c0/v0 set name=testdrive123
```

storcli2 /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:

```
storcli2 /c0/v0 set pdcache=on
```

storcli2 /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:

```
storcli2 /c0/v0 set rdcache=nora
```

storcli2 /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:

```
storcli2 /c0/v0 set wrcache=wt
```

storcli2 /cx/vx set Unmap=<On|Off>

This command unmaps the virtual drive.

storcli2 /cx/vx set wsUnmap=<On|Off>

This command unmaps the virtual drive.

storcli2 /cx/ex/sx set resume current

This command resumes an ongoing current process. You can run this command only when an operation is running on a drive.

Input example:

```
storcli2 /c0/e25/s4 set resume current
```

Virtual Drive Initialization Commands

The StorCLI2 utility supports the following commands to initialize virtual drives:

```
storcli2 /cx/vx show init
storcli2 /cx/vx start init [full] [Force]
storcli2 /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.

storcli2 /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:

```
storcli2 /c0/v2 show init
```

storcli2 /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:

```
storcli2 /cx/vx start init full
```

storcli2 /cx/vx stop init

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

Input example:

```
storcli2 /c0/v0 stop init
```

Virtual Drive Erase Commands

The StorCLI2 utility supports the following commands to erase virtual drives:

```
storcli2 /cx/vx start erase type=simple|normal|thorough[patternA=<val>] [patternB=<val>] [deletevd] [force]
storcli2 /cx/vx show erase
```

The detailed description for each command follows.

storcli2 /cx/vx start erase type=simple|normal|thorough[patternA=<val>] [patternB=<val>] [deletevd][force]

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

Input example:

```
storcli2 /c0/v0 start
```

storcli2 /cx/vx show erase

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

Input example:

```
storcli2 /c0/v0 show erase
```

Virtual Drive Consistency Check Commands

The StorCLI2 utility supports the following commands for virtual drive consistency checks:

```
storcli2 /cx/vx suspend cc
storcli2 /cx/vx resume cc
storcli2 /cx/vx show cc|consistencycheck
storcli2 /cx/vx start cc [force]
storcli2 /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.

storcli2 /cx/vx suspend cc

This command suspends 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:

```
storcli2 /c0/v4 suspend cc
```

storcli2 /cx/vx resume cc

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

Input example:

```
storcli2 /c0/v4 resume cc
```

storcli2 /cx/vx show cc|consistencycheck

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:

```
storcli2 /c0/v5 show cc|consistencycheck
```

storcli2 /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:

```
storcli2 /c0/v4 start cc
```

storcli2 /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:

```
storcli2 /c0/v4 stop cc
```

NOTE

You cannot resume a stopped consistency check process.

Background Initialization Commands

The StorCLI2 utility supports the following commands for background initialization:

```
storcli2 /cx/vx resume bgi
storcli2 /cx/vx set autobgi=<on|off>
storcli2 /cx/vx show bgi
storcli2 /cx/vx stop bgi
storcli2 /cx/vx suspend bgi
```

The detailed description for each command follows.

storcli2 /cx/vx resume bgi

This command resumes a suspended background initialization operation.

Input example:

```
storcli2 /c0/v0 resume bgi
```

storcli2 /cx/vx set autobgi=<on|off>

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

Input example:

```
storcli2 /c0/v0 set autobgi=on
```

storcli2 /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:

```
storcli2 /c0/v0 show bgi
```

storcli2 /cx/vx stop bgi

This command stops a background initialization operation.

Input example:

```
storcli2 /c0/v4 stop bgi
```

storcli2 /cx/vx suspend bgi

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

Input example:

```
storcli2 /c0/v4 suspend bgi
```

Virtual Drive Expansion Commands

The StorCLI2 utility supports the following commands for virtual drive expansion:

```
storcli2 /cx/dx expand
storcli2 /cx/vx expand drives=[e:]s|[e:]s-x|[e:]s-x,y
storcli2 /cx/vx expand [percent=<xx>][expandVdAndArray]
storcli2 /cx/vx show oce
storcli2 /cx/vx show ocedriveinfo
storcli2 /cx/vx show expansion
```

The detailed description for each command follows.

storcli2 /cx/dx expand

This command expands the virtual drive.

Input example:

```
storcli2 /c0/v25 expand
```

storcli2 /cx/vx expand drives=[e:]s|[e:]s-x|[e:]s-x,y

This command expands a virtual drive by adding new drives to the specified virtual drive.

```
storcli2 /c0/v25 expand drives=42:1
```

storcli2 /cx/vx expand [percent=<xx>][expandVdAndArray]

This command expands the virtual drive to the specified percent in the array, or to the fullest amount of free space available.

Input example:

```
storcli2 /c0/v25 expand percent=10
```

storcli2 /cx/vx show oce

This command displays expansion information on the virtual drive.

Input example:

```
storcli2 /c0/v25 show oceinfo
```

storcli2 /cx/vx show ocedriveinfo

This command displays information on the virtual drive.

Input example:

```
storcli2 /c0/v25 show ocedriveinfo
```

storcli2 /cx/vx show expansion

This command shows the expansion information.

Input example:

```
storcli2 /c0/v25 show expansion
```

Foreign Configuration Commands

The StorCLI2 utility supports the following commands to view, import, and delete foreign configurations:

```
storcli2 /cx/fall del|delete {[securitykey = xxx] [keyid = xxx]}|file=filename
storcli2 /cx/fall import [preview] {[securitykey = xxx] [keyid=xxx]}|file=filename
storcli2 /cx/fall show [all] {[securityKey = xxx] [keyid = xxx]}|file=filename
```

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.

storcli2 /cx/fall del|delete {[securityKey = xxx] [keyid = xxx]}|file=filename

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

Input example:

```
storcli2 /c0/fall delete
```

storcli2 /cx/fall import [preview] {[securityKey = xxx] [keyid = xxx]}|file=filename

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

Input example:

```
storcli2 /c0/fall import
```

storcli2 /cx/fall show [all] {[securityKey = xxx] [keyid = xxx]}|file=filename

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.

Input example:

```
storcli2 /c0/fall show all
```

Drive Group Commands

This section describes the drive group commands.

NOTE

The drive group index changes based on the topology changes because the index will be assigned based on VD persistent Id.

Drive Group Commands

The StorCLI2 utility supports the following drive group commands:

```
storcli2 /cx/dx show
storcli2 /cx/dx show all
storcli2 /cx/dx set security=on
storcli2 /cx/dx expand
```

storcli2 /cx/dx show

This command shows the topology information of the drive group.

Input example:

```
storcli2 /c0/d0 show
```

storcli2 /cx/dx show all

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

Input example:

```
storcli2 /c0/d0 show all
```

storcli2 /cx/dx set security=on

This command enables security on the specified drive group.

Input example:

```
storcli2 /c0/d0 set security=on
```

storcli2 /cx/dx expand

This command expands the drive group with the available free space.

Input example:

```
storcli2 /c0/d0 expand
```

Controller Power Savings Commands

The StorCLI2 utility supports the following commands to change the dimmer switch settings.

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

```
storcli2 /cx set ps=OFF type=UG|HS|all|
storcli2 /cx set ps=ON type=UG|HS [properties]
storcli2 /cx set ps [properties]
```

The following table describes the power-saving options.

Table 51: Power Savings Input Options

Option	Value Range	Description
ps	on off	Turns the power save option on.
type	UG : Unconfigured HS : Hot spare all : Unconfigured and hot spare drives	Specifies the type of drives for which the power save feature is applicable. Note: The power save is only activated for unconfigured drives and hot spare drives, but not for configured drives.
properties	SpinDownTime : 30-1440, in minutes SpinUpEnclDriveCount : Value 0 to 255 SpinUpEnclDelay : Valid time in seconds	The number of idle minutes before the device spins-down or stops. The maximum number of drives within an enclosure to spin-up (start) at one time. Specifies the delay of spin-up groups within an enclosure in seconds.

storcli2 /cx show power save (ps)

This command shows the current power save setting for the controller.

Input example:

```
storcli2 /c0 show ps
```

Enclosure Commands

The StorCLI2 utility supports the following enclosure commands:

```
storcli2 /cx/ex show
storcli2 /cx/ex show all
storcli2 /cx/ex show status
storcli2 /cx/ex show phyerrorcounters
storcli2 /cx/ex download file=filepath mode=5|7|E|F [activatenow][chunksize=<val>][buffer=<val>]
```

The detailed description for each command follows.

NOTE

StorCLI2 supports and can be used to qualify only Broadcom expanders and enclosures.

storcli2 /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:

```
storcli2 /c0/e25 show all
```

storcli2 /cx/ex show status

This command displays the status of the enclosure element modules.

Input example:

```
storcli2 /c0/e25 show status
```

storcli2 /cx/ex show phyerrorcounters

This command displays the enclosure phyerrorcounters information.

Input example:

```
storcli2 /c0/e25 show phyerrorcounters
```

NOTE

This command is only supported for expanders and managed PCI switch topologies.

storcli2 /cx/ex download file=<filepath> mode=5|7|E|F [activatenow] [chunksize=<val>] [bufferid=<val>]

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.

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 is downloaded in chunks of 32 KB.
- Mode E – Downloads the microcode and defers the activation.
- Mode F – Activates the deferred microcode and allows you to issue this command to all devices in a safe manner.

Input example:

```
storcli2 /c0/e25 download file=c:\file2.bin mode 7
```

Controller Phy Commands

The StorCLI2 utility supports the following controller phy commands:

```
storcli2 /cx/px show [phytype=sas|pcie]
storcli2 /cx/px show all [phytype=sas|pcie]
storcli2 /cx/px set saslinkspeed=6|12|22.5
storcli2 /cx/px set pcielinkspeed=2.5|5|8|16
storcli2 /cx/px set state=on|off [phytype=sas|pcie]
```

The detailed description for each command follows.

storcli2 /cx/px show [phytype=sas|pcie]

This command displays the SAS and PCIe phy information. If phytype is not provided, the system displays both SAS and PCIe phy information.

Input example:

```
storcli2 /c0/p0 show
```

storcli2 /cx/px|pall show all

This command displays the detailed SAS and PCIe phy information including SAS and PCIe phyerrorcounters information. If phytype is not provided, the system displays both SAS and PCIe phy information.

Input example:

```
storcli2 /c0/p0 show all
```

storcli2 /cx/px set saslinkspeed=6|12|22.5

This command sets the user provided linkspeed on a particular phy. The possible SAS speeds are 6 GB/s, 12 GB/s, and 22.5 GB/s.

Input example:

```
storcli2 /c0/p0 set saslinkspeed=6
```

storcli2 /cx/px set pcielinkspeed=2.5|5|8|16

This command sets the provided link speed on a particular PCIe link. All the phys belonging to a particular link should be provided. The possible PCIe speeds are 2.5 GT/s, 5 GT/s, 8 GT/s, and 16 GT/s.

Input example:

```
storcli2 /c0/p0 set pcielinkspeed=16
```

storcli2 /cx/px set state=on|off [phytype=sas|pcie]

This command sets the SAS and PCIe phy to on or off. If the phytype is not provided, change both the SAS and PCIe states.

Input example:

```
storcli2 /c0/p0 set state=on|off phytype=sas
```

Energy Pack Commands

The StorCLI2 utility supports the following energy pack commands:

```
storcli2 /cx/ep show [all]
storcli2 /cx/ep set autolearnmode=<val>
```

The detailed description for each command follows.

storcli2 /cx/ep show [all]

This command shows the information of the energy pack that is connected to the controller.

Input examples:

```
storcli2 /c0/ep show
storcli2 /c0/ep show all
```

storcli2 /cx/ep set autolearnmode=<val>

This command sets AutoLearnMode on an energy pack.

The possible values for autolearnmode include:

- 0 – Auto
- 1 – Disabled
- 2 – Warning

Input example:


```
storcli2 /c0/ep set autolearnmode=0
```

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.

Logging Commands

The StorCLI2 utility supports the following commands to generate and maintain log files:

```
storcli2 /cx show file=<absolute path>
storcli2 /cx show eventseqinfo
storcli2 /cx ack events seqnum=<sequence-number>
```

The detailed description for each command follows.

storcli2 /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:

```
storcli2 /c0 show events file=C:\Users\brohan\test\eventreports
```

NOTE

The command output for this command cannot be JSON formatted.

storcli2 /cx show eventseqinfo

This command shows the history of log files generated.

Input example:

```
storcli2 /c0 show eventseqinfo
```

NOTE

The command output for this command cannot be JSON formatted.

storcli2 /cx ack events seqnum=<sequence-number>

This command is used to acknowledge blocking events.

Input example:

```
storcli2 /c0 ack events seqnum=1
```

Automated Physical Drive Configurations

The StorCLI2 utility supports the following automated physical drive configuration commands:

```
storcli2 /cx show autoconfig
storcli2 /cx set autoconfig factory
storcli2 /cx set autoconfig primary option=UGood|JBOD|SecureJBOD|R0|SecureR0|ROWB|SecureROWB
storcli2 /cx set autoconfig secondary option=UGood|JBOD|SecureJBOD|R0|SecureR0|ROWB|SecureROWB
```

```
storcli2 /cx set autoconfig immediate option=JBOD|SecureJBOD|R0|SecureR0|ROWB|SecureROWB drives=<e:s|e:s-x|e:s-x,y>
```

The detailed description for each command follows.

storcli2 /cx show autoconfig

This command lets you view the automatic physical drive policy.

Input example:

```
storcli2 /c0 show autoconfig
```

storcli2 /cx set autoconfig factory

This command sets the current autoconfiguration to factory settings.

Input example:

```
storcli2 /c0 set autoconfig factory
```

storcli2 /cx set autoconfig primary option=UGood|JBOD|SecureJBOD|R0|SecureR0|ROWB|SecureROWB

This command sets the primary autoconfiguration of the controller.

Input example:

```
storcli2 /c0 set autoconfig primary
```

storcli2 /cx set autoconfig secondary option=UGood|JBOD|SecureJBOD|R0|SecureR0|ROWB|SecureROWB

This command sets the secondary autoconfiguration of the controller.

Input example:

```
storcli2 /c0 set autoconfig secondary
```

storcli2 /cx set autoconfig immediate option=JBOD|SecureJBOD|R0|SecureR0|ROWB|SecureROWB drives=<e:s|e:s-x|e:s-x,y>

This command sets the autoconfiguration of the controller immediately.

Input example:

```
storcli2 /c0 set autoconfig immediate
```

NOTE

This command applies the autoconfigure option that is specified in the arguments on the unconfigured drives present in the system. This command does not change the primary or secondary autoconfig option of the adapter.

Frequently Used Tasks

The following commands are frequently used in StorCLI2.

- [Displaying the Version of the StorCLI2 Utility](#)
- [Displaying the StorCLI2 Utility Help](#)
- [Displaying System Summary Information](#)
- [Displaying Free Space in a Controller](#)
- [Adding Virtual Drives](#)
- [Setting the Cache Policy in a Virtual Drive](#)
- [Displaying Virtual Drive Information](#)
- [Deleting Virtual Drives](#)
- [Flashing Controller Firmware](#)

Displaying the Version of the StorCLI2 Utility

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

```
storcli2 v
```

Displaying the StorCLI2 Utility Help

The following command displays the StorCLI2 utility help:

```
storcli2 h
```

Help appears for all the StorCLI2 tool commands.

NOTE

To retrieve the list of commands supported by a specific controller, use the command `storcli2 /c0 ?`, "`... /c1 ?`". The generic `... /cx ?` command lists all available commands in Storcli2.

Displaying System Summary Information

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

```
storcli2 -show [all]
```

Displaying Free Space in a Controller

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

```
storcli2 /cx show freespace
```

Adding Virtual Drives

The following command creates a virtual drive:

```
storcli2 /cx add vd r[0|1|5|6|10|50|60]
[Size=<VD1_Sz>,<VD2_Sz>,...|remaining|all] [name=<VDNAME1>,...]
drives=e:s|e:s-x|e:s-x,y,e:s-x,y,z [PDperArray=x] [secure]
[pdcache=on|off|default] [ps=default|none]
[WT|WB|AWB] [nora|ra] [unmap] [wsunmap]
[Strip=<64|256] [AfterVd=X]
[hotspare=e:s|e:s-x|e:s-x,y]
[cachebypass=None|All|64|128|256|512|1024] [init=0|1|2] [NoAutoBGI]
storcli2 /cx add vd each r0 [name=<VDNAME1>,...] [drives=e:s|e:s-x|e:s-x,y]
[secure] [pdcache=on|off|default] [ps=default|none]
```

```
[WT|WB|AWB] [nora|ra] [unmap] [wsunmap] [Strip=<64|256>]
[cachebypass=None|All|64|128|256|512|1024] [init]=0|1|2] [NoAutoBGI]
```

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 represent 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 *secure* option creates security-enabled drives.
- The *PDcache* option can be set to *on* or *off*.
- The *dimmerswitch* 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 *64* and *256*.

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.

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:

```
storcli2 /cx/v(x|all) set wrccache=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:

```
storcli2 /cx/v(x/all) show
```

Deleting Virtual Drives

The following command deletes virtual drives:

```
storcli2 /cx/v(x|all) del
```

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.

Flashing Controller Firmware

The following command is used to flash the controller firmware.

```
storcli2 /cx download file=<filepath> [noverchk] [activatenow]
```

NOTE

The command output for this command cannot be JSON formatted.

SAS Address Assignment Rule

This section provides information on how to calculate the PHY SAS address.

The PHY SAS address is calculated by incrementing the controller SAS address by one, based on the number of PHYs.

Suppose you are using 16 or 8 PHY cards and four connectors exist: C3, C2, C1, and C0. Each connector has four PHYs, and the autoport configuration is always enabled. Connector C3 has PHYs 0 through 3, Connector C2 has PHYs 4 through 7, Connector C1 has PHYs 8 through 11, and Connector C0 has PHYs 12 through 15.

- If you are connecting four different target devices and want to plug a cable into Connector 1, the SAS address for this port is 0x5000_0000_8000_0008 because the connector's first PHY is 8.
- Furthermore, when you plug a cable into Connector 0, the SAS address for this port is 0x5000_0000_8000_0009.
- Assuming nothing is connected to the HBA and you plug a cable into Connector 0, the SAS address that is assigned to this port is 0x5000_0000_8000_0008.
- Again, assuming nothing is connected to the HBA and you plug a cable into Connector 3, the SAS address that is assigned to this port is 0x5000_0000_8000_0000.
- Next, when a cable is plugged into Connector 2, the SAS address that is assigned to this port is 0x5000_0000_8000_0001.

This logic is also applicable for cards with eight PHYs.

Controllers have two SAS cores; each core can have a wide port, with at the most x8 connections. While connectors C0 and C1 can belong to one core, connectors C2 and C3 can belong to another core.

StorCLI to StorCLI2 Command Conversion

The following table provides a conversion of StorCLI to StorCLI2 commands.

Table 52: StorCLI to StorCLI2 Command Conversion

StorCLI Command	StorCLI2 Command
Help Commands	
storcli -h -help h help	storcli2 h help

StorCLI Command	StorCLI2 Command
storcli -v -version v version	storcli2 v version
System Commands	
storcli show	storcli2 show
storcli show all	storcli2 show all
storcli show file = <>	storcli2 show file=<>
storcli show ctrlcount	storcli2 show ctrlcount
—	storcli2 get rttDump
Enclosure Commands	
storcli2 /cx/ex show	storcli2 /cx/ex show
storcli2 /cx/ex show all	storcli2 /cx/ex show all
storcli /cx/ex download src=<filepath> [forceActivate] [mode=5 7] [bufferid=<val>]	storcli2 /cx/ex download file=<filepath> mode=5 7 E F [activatenow] [chunksize=<val>] [bufferid=<val>]
storcli /cx/ex download src=<filepath> mode=e [offline] [forceActivate] [delay=<val>] [bufferid=<val>]	
storcli /cx/ex download src=<filepath> mode=f [offline] [delay=<val>] [bufferid=<val>]	
storcli /cx/ex show status [extended]	storcli2 /cx/ex show status
storcli /cx/ex show phyerrorcounters	storcli2 /cx/ex show phyerrorcounters
Disk Group	
storcli /cx/dx set security=on	storcli2 /cx/dx set security=on
storcli /cx/dx show [all]	storcli2 /cx/dx show [all]
—	storcli2 /cx/dx expand
storcli /cx/dall show cachecade	—
storcli /cx/dall show mirror	—
storcli /cx/dall split mirror	—
storcli /cx/dx set hidden=on off	—
storcli /cx/dall add mirror src=<val> [force]	—
storcli /cx/dx set transport=on off [EDHSP=on off] [SDHSP=on off]	—
Foreign Configuration Level Commands	
storcli2 /cx/fall show [all] [securityKey = xxx]	storcli2 /cx/fall show [all] {[securityKey = xxx] [keyid = xxx]} file=filename
storcli2 /cx/fall del delete [securityKey = xxx]	storcli2 /cx/fall del delete {[securityKey = xxx] [keyid = xxx]} file=filename
storcli2 /cx/fall import [preview] [securityKey = xxx]	storcli2 /cx/fall import [preview] {[securityKey = xxx] [keyid = xxx]} file=filename
Controller Firmware Download and Restart Commands	
storcli /cx download file=<filepath> [noverchk] [noreset] [forcehcb]	storcli2 /cx download file=<filepath> [activationtype=online offline] [noverchk]
storcli /cx download file=<filepath> [fwtype=<val>] [ResetNow] [nosigchk] [noverchk] [force]	

StorCLI Command	StorCLI2 Command
storcli /cx download efibios file=<filepath>	
storcli /cx download cpld file=<filepath>	
storcli /cx download psoc file=<filepath>	
storcli /cx download bios file=<filepath>	
storcli /cx download fcode file=<filepath>	
storcli /cx erase nvram	—
storcli /cx erase fwbackup	—
storcli /cx erase bootservices	—
storcli /cx erase all [excludemfg] [file=filename]	—
storcli /cx erase perconpage	—
storcli /cx erase mpb	—
storcli /cx restart	storcli2 /cx reset
Physical Drive Commands	
storcli /cx[/ex]/sx show	storcli2 /cx[/ex]/sx show
storcli /cx[/ex]/sx show all	storcli2 /cx[/ex]/sx show all
storcli /cx[/ex]/sx start rebuild	storcli2 /cx[/ex]/sx start rebuild
storcli /cx[/ex]/sx stop rebuild	storcli2 /cx[/ex]/sx stop rebuild
storcli /cx[/ex]/sx suspend rebuild	storcli2 /cx[/ex]/sx suspend rebuild
storcli /cx[/ex]/sx resume rebuild	storcli2 /cx[/ex]/sx resume rebuild
storcli /cx[/ex]/sx show rebuild	storcli2 /cx[/ex]/sx show rebuild
storcli /cx[/ex]/sx show poh [ignoreselftest]	storcli2 /cx[/ex]/sx show poh [ignoreselftest]
storcli /cx[/ex]/sx show smart	storcli2 /cx[/ex]/sx show smart
storcli /cx[/ex]/sx start copyback target=e:s	storcli2 /cx/ex/sx start replacedrive target=e:s
storcli /cx[/ex]/sx stop copyback	storcli /cx[/ex]/sx stop replacedrive
storcli /cx[/ex]/sx suspend copyback	storcli /cx[/ex]/sx suspend replacedrive
storcli /cx[/ex]/sx resume copyback	storcli2 /cx/ex/sx resume replacedrive
storcli /cx[/ex]/sx reset phyerrorcounters	storcli2 /cx[/ex]/sx reset phyerrorcounters
storcli /cx[/ex]/sx reset errorcounters type = 1 2	—
storcli /cx[/ex]/sx show copyback	storcli2 /cx/ex/sx show replacedrive
storcli /cx[/ex]/sx show patrolread	storcli2 /cx[/ex]/sx show patrolread
storcli /cx[/ex]/sx show phyerrorcounters	storcli2 /cx[/ex]/sx show phyerrorcounters
storcli /cx[/ex]/sx show errorcounters	—
storcli /cx[/ex]/sx start initialization	storcli2 /cx/ex/sx start clear
storcli /cx[/ex]/sx stop initialization	storcli2 /cx/ex/sx stop clear
storcli /cx[/ex]/sx show initialization	storcli2 /cx/ex/sx show clear
storcli /cx[/ex]/sx start locate	storcli2 /cx[/ex]/sx start locate
storcli /cx[/ex]/sx stop locate	storcli2 /cx[/ex]/sx stop locate

StorCLI Command	StorCLI2 Command
storcli /cx[/ex]/sx show securitykey keyid	storcli2 /cx[/ex]/sx show securitykey keyid
storcli /cx/ex/sx add hotsparedrive [DGs=<N 0,1,2...>] [enclaffinity] [nonrevertible]	storcli2 /cx/ex/sx add hotspare [DGs=<N 0,1,2...>] [enclaffinity]
storcli /cx[/ex]/sx delete hotsparedrive	storcli2 /cx[/ex]/sx delete hotspare
storcli /cx/ex/sx spindown	storcli2 /cx/ex/sx start prepformv1
storcli /cx/ex/sx spinup	storcli2 /cx/ex/sx undo prepformv1
storcli /cx[/ex]/sx set online	storcli2 /cx[/ex]/sx set state=online
storcli /cx[/ex]/sx set offline	storcli2 /cx[/ex]/sx set state=offline
storcli /cx[/ex]/sx set missing	storcli2 /cx[/ex]/sx set state=missing
storcli /cx[/ex]/sx set jbod	storcli2 /cx/ex/sx set JBOD [force]
storcli /cx[/ex]/sx set security=on	storcli2 /cx[/ex]/sx set security=on
storcli /cx[/ex]/sx set good [force]	storcli2 /cx/ex/sx set good [force]
—	storcli2 /cx/ex/sx set bad [force]
—	storcli2 /cx/ex/sx set failed
storcli /cx[/ex]/sx insert dg=A array=B row=C	storcli2 /cx/ex/sx insert dg=A span=B row=C
storcli /cx[/ex]/sx download src=<filepath> [satabridge] [mode= 5 7] [parallel [force]]	storcli2 /cx/ex/sx download file=<filepath> mode=5 7 E [activatenow] F [chunksize=<val>]
storcli /cx[/ex]/sx download src=<filepath> mode= E [offline] [activatenow [delay=<val>]]	
storcli /cx[/ex]/sx downloadmode= F [offline] [delay=<val>]	
storcli /cx[/ex]/sx secureerase [force]	storcli2 /cx/ex/sx start erase type=reprovision [force]
storcli /cx/ex/sx start erase [simple normal thorough standard threepass crypto] [patternA=<val>] [patternB=<val>]	storcli2 /cx[/ex]/sx start erase [simple normal thorough standard threepass crypto] [patternA=<val>] [patternB=<val>]
storcli /cx[/ex]/sx stop erase	storcli2 /cx[/ex]/sx stop erase
storcli /cx[/ex]/sx show erase	storcli2 /cx[/ex]/sx show erase
storcli /cx[/ex]/sx show repair	storcli2 /cx show driveRecoveryInfo
storcli /cx[/ex]/sx start repair [force]	storcli2 /cx/ex/sx start recovery [force]
storcli /cx[/ex]/sx stop repair	storcli2 /cx/ex/sx stop recovery
storcli /cx[/ex]/sx show dpmstat type = HIST LCT RA EXT [logfile=<filename>]	storcli2 /cx/ex/sx show DPM type = HIST LCT RA EXT
storcli /cx delete dpmstat type = Hist LCT RA EXT All	storcli2 /cx delete DPM
storcli /cx/ex/sx show jbod	—
storcli /cx/ex/sx show jbod all	—
storcli /cx/ex/sx del jbod [force]	—
storcli /cx/ex/sx set bootdrive=<on off>	—
Virtual Drive Commands	

StorCLI Command	StorCLI2 Command
storcli /cx add vd r[0 1 5 6 00 10 50 60] [Size=<VD1_Sz>,<VD2_Sz>,... all] [name=<VDNAME1>,...]drives=e:s e:s-x e:s-x,y,e:s-x,y,z [Strip=<8 16 32 64 128 256 512 1024>][PDperArray=x][SED][pdcache=on off default][pi][DimmerSwitch(ds)=default automatic(auto) none maximum(max) MaximumWithoutCaching(maxnocache)][WT WB AWB][nora ra][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]	storcli2 add vd r[0 1 5 6 10 50 60] [Size=<VD1_Sz>,<VD2_Sz>,... remaining all] [name=<VDNAME1>,...] drives=e:s e:s-x e:s-x,y,e:s-x,y,z [PDperArray=x] [secure] [pdcache=on off default] [ps=default none] [WT WB AWB] [nora ra] [unmap] [wsunmap] [Strip=<64 256>] [AfterVd=X] [hotspare = e:s e:s-x e:s-x,y] [cachebypass=None All 64 128 256 512 1024] [init=0 1 2] [NoAutoBGI]
storcli /cx add vd each r0 [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 64 128 256 512 1024>] [ExclusiveAccess][Cbsize=0 1 2 Cbmode=0 1 2 3 4 7] [unmap]	storcli2 /cx add vd each r0 [name=<VDNAME1>,...] drives=e:s e:s-x e:s-x,y [secure] [pdcache=on off default] [ps=default none] [WT WB AWB] [nora ra] [unmap] [wsunmap] [Strip=<64 256>] [cachebypass=None All 64 128 256 512 1024] [init=0 1 2] [NoAutoBGI]
storcli /cx add VD cachecade r[0 1 10] drives = [e:]s [e:]s-x [e:]s-x,y [WT WB] [assignvds = 0,1,2]	—
storcli /cx/vx del [cachecade] [discardcache] [force]	storcli2 /cx/vx del [discardcache] [force]
storcli /cx/vx set ssdcaching=on off	—
storcli /cx/vx set hidden=on off	—
storcli /cx/vx show expansion	storcli2 /cx/vx show expansion
—	storcli2 /cx/vx show ocedriveinfo
storcli /cx/vx expand Size=<xx> [expandarray]	storcli2 /cx/vx expand [percent=<xx>] [expandVdAndArray] storcli2 /cx/vx expand drives=[e:]s [e:]s-x [e:]s-x,y
storcli /cx/vx set emulationType=0 1 2	—
storcli /cx/vx set cbsize=0 1 2 cbmode=0 1 2 3 4 7	storcli2 /cx/vx set cachebypass=None All 64 128 256 512 1024
storcli /cx/vx set wrcache=WT WB AWB	storcli2 /cx/vx set wrcache=WT WB AWB
storcli /cx/vx set rdcache=RA NoRA	storcli2 /cx/vx set rdcache=RA NoRA
storcli /cx/vx set iopolicy=Cached Direct	—
storcli /cx/vx set accesspolicy=RW RO Blocked RmvBlkd	—
storcli /cx/vx set pdcache=On Off Default	storcli2 /cx/vx set pdcache=On Off Default
storcli /cx/vx set name=<NameString>	storcli2 /cx/vx set name=<NameString>
storcli /cx/vx set HostAccess=ExclusiveAccess SharedAccess	—
storcli /cx/vx set ds=Default Auto None Max MaxNoCache	storcli2 /cx/vx set ps=Default None
storcli /cx/vx set autobgi=On Off	storcli2 /cx/vx set autobgi=On Off
storcli /cx/vx set pi=Off	—
storcli /cx/vx show	storcli2 /cx/vx show
storcli /cx/vx show all [logfile=filename]]	storcli2 /cx/vx show all [logfile=filename]]
storcli /cx/vx show init	storcli2 /cx/vx show init
storcli /cx/vx show cc	storcli2 /cx/vx show cc
storcli /cx/vx show erase	storcli2 /cx/vx show erase

StorCLI Command	StorCLI2 Command
storcli /cx/vx show migrate	—
storcli /cx/vx show bgi	storcli2 /cx/vx show bgi
storcli /cx/vx start init [Full] [Force]	storcli2 /cx/vx start init[Full][Force]
storcli /cx/vx start erase [simple normal thorough standard] [patternA=<val>] [patternB=<val>]	storcli2 /cx/vx start erase type=simple normal thorough [patternA=<val>] [patternB=<val>] [deletevd] [force]
storcli /cx/vx start cc [Force]	storcli2 /cx/vx start cc
storcli /cx/vx start migrate type=raidx [option=add remove drives=[e:]s [e:]s-x [e:]s-x,y] [Force]	—
storcli /cx/vx stop init	storcli2 /cx/vx stop init
storcli /cx/vx stop erase	storcli2 /cx/vx stop erase
storcli /cx/vx stop cc	storcli2 /cx/vx stop cc
storcli /cx/vx stop bgi	storcli2 /cx/vx stop bgi
storcli /cx/vx suspend cc	storcli2 /cx/vx suspend cc
storcli /cx/vx suspend bgi	storcli2 /cx/vx suspend bgi
storcli /cx/vx resume cc	storcli2 /cx/vx resume cc
storcli /cx/vx resume bgi	storcli2 /cx/vx resume bgi
—	storcli2 /cx/vx suspend init
—	storcli2 /cx/vx resume init
storcli /cx/vx delete preservedcache [force]	storcli2 /cx/vx delete preservedcache [force]
storcli /cx/vx del [discardcache] [force]	storcli2 /cx/vx del [discardcache] [force]
storcli /cx/vx set bootdrive=<on off>	—
storcli /cx/vx show BBMT	storcli2 /cx/vx show BBMT
storcli /cx/vx delete BBMT	storcli2 /cx/vx delete BBMT
—	storcli2 /cx/vx suspend current
—	storcli2 /cx/vx resume current
—	storcli2 /cx/vx suspend oce
—	storcli2 /cx/vx resume oce
—	storcli2 /cx/vx show oce
storcli /cx/vx set Unmap=<On Off>	storcli2 /cx/vx set Unmap=<On Off> [wsunmap=<On Off>]
—	storcli2 /cx/vx set wsunmap=<On Off>
Security and Premium Features Commands	
storcli /cx set securitykey < =xxxxxxx [passphrase=xxxx] [keyid=xxx] file=filename >	storcli2 /cx set security { securitykey = xxxx [passphrase = xxxx] [keyid = xxx] } file = filename
—	storcli2 /cx set security passphrase=<key>
storcli /cx set securitykey < keyid=xxx file=filename >	storcli2 /cx set security rekey { oldsecuritykey = xxxx securitykey = xxxx [passphrase = xxxx] [keyid = xxx] } file = filename
storcli /cx set securitykey < =xxxxxxx [passphrase=xxxx] [keyid=xxx] file=filename >	
storcli /cx show securitykey keyid	storcli2 /cx show security keyid

StorCLI Command	StorCLI2 Command
storcli /cx compare securitykey <=xxxxxxxxx file=filename>	—
storcli /cx delete securitykey	storcli2 /cx delete security securitykey
storcli /cx show aso	storcli2 /cx show aso
storcli /cx set aso key=<key value> preview	storcli2 /cx set aso key=<key value>preview
storcli /cx set aso key=<key value>	storcli2 /cx set aso key=<key value>
storcli /cx set aso transfertovault	—
storcli /cx set aso rehostcomplete	—
storcli /cx set aso deactivatetrialkey	storcli2 /cx set aso deactivatetrialkey
—	storcli2 /cx set security rekey { securitykey = xxxx [passphrase = xxxx] [keyid = xxxx] } file = filename
—	storcli2 /cx set security rekey useEKMS { oldsecuritykey=xxxx } file = filename
—	storcli2 /cx set security useekms
—	storcli2 /cx set security rekey useekms
—	storcli2 /cx show security spdm
storcli /cx show security spdm slotgroup=xx slot=yy	storcli2 /cx show security spdm slotgroup=xx slot=yy
storcli /cx export security spdm slotgroup=xx slot=yy subject=subjectfile file=filename	storcli2 /cx export security spdm slotgroup=xx slot=yy subject=subjectfile file=filename
storcli /cx import security spdm slotgroup=xx slot=yy file=filename [seal]	storcli2 /cx import security spdm slotgroup=xx slot=yy file=filename
storcli /cx set security spdm slotgroup=xx slot=yy invalidate [force]	storcli2 /cx set security spdm slotgroup=xx slot=yy invalidate [force]
storcli /cx get security spdm slotgroup=xx slot=yy file=filename	storcli2 /cx get security spdm slotgroup=xx slot=yy file=filename
Snapdump Commands	
storcli /cx show snapdump	storcli2 /cx show snapdump
storcli /cx set snapdump state=on off	—
storcli /cx set snapdump [savecount=<value> delayocr=<value>]	—
storcli /cx get snapdump [id=[all <value> file=<fileName>]] [norttdump]	storcli2 /cx get snapdump < all id= <dump-id> > [norttdump]
storcli /cx delete snapdump [force]	storcli2 /cx delete snapdump [force]
—	storcli2 /cx get snapdump ondemand [norttdump]
—	storcli2 /cx get snapdump ondemand debugfile=<filename> [norttdump]
—	storcli2 /cx delete snapdump enhanced
Auto Configuration Commands	
storcli /cx show autoconfig	storcli2 /cx show autoconfig
—	storcli2 /cx set autoconfig factory

StorCLI Command	StorCLI2 Command
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]]	storcli2 /cx set autoconfig primary option=UGood JBOD SecureJBOD R0 SecureR0 R0WB SecureR0WB
	storcli2 /cx set autoconfig secondary option=UGood JBOD SecureJBOD R0 SecureR0 R0WB SecureR0WB
	storcli2 /cx set autoconfig immediate option=JBOD SecureJBOD R0 SecureR0 R0WB SecureR0WB drives= all <e:s e:s-x e:s-x,y>
Controller Properties Display Commands	
storcli /cx show events [[type= <sincereboot sinceshutdown includedeleted latest=x ccinon vd=<0,1,...>] [filter=<[info], [warning],[critical],[fatal]>] [file=<filepath>] [logfile=<filename>]]	storcli2 /cx show events [[type= <blocking nonblocking sincereboot sinceshutdown includedeleted latest=x ccinon vd=<vd persistent id> >] filter=<info warning critical fatal> file=<filepath>]
storcli /cx show eventloginfo	storcli2 /cx show eventseqinfo
storcli /cx show termlog [type=config contents] [logfile=<filename>]	—
storcli /cx show sesmonitoring	—
storcli /cx show failpdonsmarterror	storcli2 /cx show failonsmarterror
storcli /cx show freespace	storcli2 /cx show freespace
storcli /cx show cc consistencycheck	storcli2 /cx show cc consistencycheck
storcli /cx show ocr	storcli2 /cx show ocr
storcli /cx show sesmultipathcfg	storcli2 /cx show sesmultipathcfg
storcli /cx show	storcli2 /cx show
storcli /cx show all [logfile=<filename>]	storcli2 /cx show all [logfile=<filename>]
storcli /cx show preservedcache	storcli2 /cx show preservedcache
storcli /cx show bootdrive	—
storcli /cx show bootwithpinnedcache	storcli2 /cx show bootwithpreservedcache
storcli /cx show activityforlocate	—
storcli /cx show copyback	storcli2 /cx show replacedrive
storcli /cx show jbod	—
storcli /cx show autorebuild	storcli2 /cx show autorebuild
storcli /cx show cachebypass	—
storcli /cx show usefdeonlyencrypt	—
storcli /cx show prcorrectunconfiguredareas	storcli2 /cx show prcorrectunconfiguredareas
storcli /cx show batterywarning	storcli /cx show energypackwarning
storcli /cx show abortconerror	storcli2 /cx show abortconerror
storcli /cx show ncq	—
storcli /cx show configautobalance	—
storcli /cx show maintainpdfailhistory	storcli2 /cx show maintainpdfailhistory
storcli /cx show restorehotspare	—
storcli /cx show bios	storcli2 /cx show bootmode
storcli /cx show alarm	—

StorCLI Command	StorCLI2 Command
storcli /cx show deviceorderbyfirmware	—
storcli /cx show foreignautoimport	—
storcli /cx show directpdmapping	—
storcli /cx show rebuildrate	storcli2 /cx show rebuildrate
storcli /cx show loadbalancemode	—
storcli /cx show eghs	storcli2 /cx show es storcli2 /cx show esSMARTER
storcli /cx show cacheflushint	—
storcli /cx show prrate	storcli2 /cx show prrate
storcli /cx show ccrate	storcli2 /cx show ccrate
storcli /cx show bgirate	storcli2 /cx show bgirate
storcli /cx show dpm	storcli2 /cx/ex/sx show DPM type = HIST LCT RA EXT
storcli /cx show sgpioforce	—
storcli /cx show reconrate	storcli2 /cx show ocerate
storcli /cx show spinupdrivecount	storcli2 /cx show spinupdrivecount
storcli /cx show wbsupport	—
storcli /cx show spinupdelay	storcli2 /cx show spinupdelay
storcli /cx show coercion	storcli2 /cx show coercionmode
storcli /cx show limitMaxRateSATA	—
storcli /cx show HDDThermalPollInterval	—
storcli /cx show SSDThermalPollInterval	—
storcli /cx show smartpollinterval	storcli2 /cx show smartpollinterval
storcli /cx show eccbucketsize	storcli2 /cx show eccbucketsize
storcli /cx show eccbucketleakrate	storcli2 /cx show eccbucketleakrate
storcli /cx show backplane	—
storcli /cx show perfmode	—
storcli /cx show perfmodevalues	—
storcli /cx show pi	—
storcli /cx show time	storcli2 /cx show time
storcli /cx show ds	storcli2 /cx show ps
storcli /cx show safeid	storcli2 /cx show ASO
storcli /cx show rehostinfo	—
storcli /cx show pci	storcli2 /cx show pci
storcli /cx show ASO	storcli2 /cx show ASO
storcli /cx show linkconfig	—
storcli /cx show securitykey keyid	storcli2 /cx show securitykey keyid
storcli /cx show patrolRead	storcli2 /cx show patrolRead

StorCLI Command	StorCLI2 Command
storcli /cx show powermonitoringinfo	—
storcli /cx show ldlimit	—
storcli /cx show badblocks	—
storcli /cx show dequeuelog file=<filepath>	—
storcli /cx show maintenance	—
storcli /cx show personality	storcli2 /cx show personality
storcli /cx show profile	—
storcli /cx show jbodwritecache	—
storcli /cx show immediateio	—
storcli /cx show driveactivityled	—
storcli /cx show unmap	storcli2 /cx show unmap
storcli /cx show pdfailevents [lastoneday] [lastseqnum=<val>] [file=<filepath>]	—
storcli /cx show pdfaileventoptions	storcli2 /cx show pdfaileventoptions
storcli get rttDump	storcli2 get rttDump
storcli /cx show AliLog [logfile=<filename>]	storcli /cx show AliLog [logfile=<filename>]
storcli /cx show flushwriteverify	—
storcli /cx show largeQD	—
storcli /cx show assemblynumber	—
storcli /cx show tracernumber	—
storcli /cx show boardname	—
storcli /cx show sasadd	—
storcli /cx show vpd	—
storcli /cx show htbparams	—
—	storcli2 /cx show bootmode
storcli /cx start dpm	storcli2 /cx start DPM [delay=<val>] [maxconcurrentpd =<value>] drives=e:s e:s-x e:s-x,y
storcli /cx stop dpm	storcli2 /cx stop DPM
—	storcli2 /cx show smartpollintervaljbod
storcli /cx show failedNvmeDevices	storcli2 /cx show driverecoveryinfo
—	storcli2 /cx show name
—	storcli2 /cx show CacheOffloadEncType
—	storcli2 /cx show maintainjbodfailhistory
—	storcli2 /cx show jbodesmgmt
—	storcli2 /cx show BaseEnclLevel
—	storcli2 /cx show pdtempPoll
—	storcli2 /cx show fwjbodsecurity
—	storcli2 /cx show hostjbodsecurity

StorCLI Command	StorCLI2 Command
—	storcli2 /cx show smartpoll
—	storcli2 /cx show drivewceforrebuild
—	storcli2 /cx show maintainpdfailhistory
—	storcli2 /cx show bootmode
—	storcli2 /cx show prcorrectunconfiguredareas
—	storcli2 /cx show smartpollinterval
—	storcli2 /cx show esSMARTer
—	storcli2 /cx show exposeencldevice
—	storcli2 /cx get DPM status
—	storcli2 /cx show datalosswarning
—	storcli2 /cx get DPM config
—	storcli2 /cx delete activation offline
—	storcli2 /cx get activation status
—	storcli2 /cx show unusabledriveinfo
Controller Properties Set Commands	
storcli /cx set termlog= on off offthisboot	—
storcli /cx set sesmonitoring= on off	—
storcli /cx set failpdonsmarterror= on off	storcli2 /cx set failpdonsmarterror= on off
storcli /cx set consistencycheck cc= off seq conc [delay= value] [starttime= yyyy/mm/dd hh] [excludevd= x-y,z none]	storcli2 /cx set cc consistencycheck= off storcli2 /cx set cc consistencycheck= on starttime= <yyyy/mm/dd hh> execfrequency hours days weeks= <value> storcli2 /cx set cc consistencycheck [starttime= <yyyy/mm/dd hh>] [execfrequency hours days weeks= <value>] [maxvd= <value>] [excludevd= x-y,z none]] storcli2 /cx set cc consistencycheck= on storcli2 /cx set cc consistencycheck factory
storcli /cx set ocr= <on off>	storcli2 /cx set ocr= <on off> type= <all auto>
storcli /cx set sesmultipathcfg= <on off>	storcli2 /cx set sesvpdassociation= <lun targetport>
storcli /cx set bootwithpinnedcache= <on off>	storcli2 /cx set bootwithpinnedcache= <on off>
storcli /cx set activityforlocate= <on off>	—
storcli /cx set copyback= <on off> type= ctrl smartssd smarthdd all	storcli2 /cx set replacedrive= <on off> type= ctrl smartssd smarthdd all
storcli /cx set jbod= <on off> [force]	—
storcli /cx set autorebuild= <on off>	storcli2 /cx set autorebuild= <on off>
storcli /cx set ldlimit= <default max>	—
storcli /cx set cachebypass= <on off>	—
storcli /cx set usefdeonlyencrypt= <on off>	—
storcli /cx set prcorrectunconfiguredareas= <on off>	storcli2 /cx set prcorrectunconfiguredareas= <on off>
storcli /cx set cachebypass= <on off>	storcli2 /cx set cachebypass= <on off>
storcli /cx set batterywarning= <on off>	storcli2 /cx set energypackwarning= <on off>

StorCLI Command	StorCLI2 Command
storcli /cx set abortconerror=<on off>	storcli2 /cx set abortconerror=<on off>
storcli /cx set ncq=<on off>	—
storcli /cx set configautobalance=<on off>	—
storcli /cx set maintainpdfailhistory=<on off>	storcli2 /cx set maintainpdfailhistory=<on off>
storcli /cx set restorehotspare=<on off>	—
storcli /cx set bios [state=<on off>] [Mode=<SOE PE IE SME>] [abs=<on off>] [DeviceExposure=<value>]	storcli2 /cx set bootmode=COE/SMOE
storcli /cx set alarm=<on off silence>	—
storcli /cx set deviceorderbyfirmware=<on off>	—
storcli /cx set foreignautoimport=<on off>	—
storcli /cx set directpdmapping=<on off>	—
storcli /cx set rebuildrate=<value>	storcli2 /cx set rebuildrate=<value>
storcli /cx set loadbalancemode=<on off>	—
storcli /cx set eghs [state=<on off>] [eug=<on off>] [smarter=<on off>]	storcli2 /cx set es=<on off> ghs ug storcli2 /cx set esSMARTer=on off
storcli /cx set cacheflushhint=<value>	—
storcli /cx set prrate=<value>	storcli2 /cx set prrate=<value>
storcli /cx set ccrate=<value>	storcli2 /cx set ccrate=<value>
storcli /cx set bgirate =<value>	storcli2 /cx set bgirate =<value>
storcli /cx set dpm =<on off>	—
storcli /cx set sgpioforce =<on off>	—
storcli /cx set supportssdpatrolread =<on off>	storcli2 /cx set supportssdpatrolread =<on off>
storcli /cx set reconrate=<value>	storcli2 /cx set ocerate=<value>
storcli /cx set spinupdrivecount=<value>	storcli2 /cx set spinupdrivecount=<value>
storcli /cx set spinupdelay=<value>	storcli2 /cx set spinupdelay=<value>
storcli /cx set coercion=<value>	storcli2 /cx set coercion=<value>
storcli /cx set limitMaxRateSATA=on off	—
storcli /cx set HDDThermalPollInterval=<value>	—
storcli /cx set SSDThermalPollInterval=<value>	—
storcli /cx set smartpollinterval=<value>	storcli2 /cx set smartpollinterval=<value> pdtype=Internal External
storcli /cx set eccbucketsize=<value>	storcli2 /cx set eccbucketsize=<value>
storcli /cx set eccbucketleakrate=<value>	storcli2 /cx set eccbucketleakrate=<value>
storcli /cx set backplane mode=<value> expose=<on off>	storcli2 /cx set exposeencldevice =<on off>
storcli /cx set perfmode=<value> [maxflushlines=<value> numiostoorder=<value>]	—
storcli /cx set pi [state=<on off>] [import=<on off>]	—
storcli /cx set time=<yyymmdd hh:mm:ss systemtime>	storcli2 /cx set time=systemtime
storcli /cx set ds=OFF type=1 2 3 4	storcli2 /cx set ps=OFF type= UG HS all

StorCLI Command	StorCLI2 Command
storcli /cx set ds=ON type=1 2 [properties]	storcli2 /cx set ps=ON type= UG HS [SpinDownTime=30-1440] [SpinUpEncDrvCnt=<val>] [SpinUpEncDelay=<val>]
storcli /cx set ds=ON type=3 4 DefaultLdType=<val> [properties]	—
storcli /cx set ds [properties]	storcli2 /cx set ps=ON type= UG HS [SpinDownTime=30-1440] [SpinUpEncDrvCnt=<val>] [SpinUpEncDelay=<val>]
storcli /cx set factory defaults	storcli2 /cx set factory defaults
storcli /cx set linkconfig [conname=cx,cy] configid=<val>	—
storcli /cx set patrolread [= [[on mode=<auto manual>]] off] [starttime=< yyyy/mm/dd hh>] [maxconcurrentpd =<value>] [includessds=<on onlymixed off>] [includessds=<on onlymixed off>] [includessds=<on onlymixed off>] [uncfgareas=on off] [delay = <value>] [excludevd=x-y,z none]	storcli2 /cx set pr patrolread=off storcli2 /cx set pr patrolread=on starttime=<yyyy/mm/dd hh> execfrequency hours days weeks=<value> storcli2 /cx set pr patrolread [starttime=<yyyy/mm/dd hh>] [execfrequency hours days weeks=<value>] [maxconcurrentpd =<value>] [includessds=<on off>] [excludevd=x-y,z none] storcli2 /cx set pr patrolread=on storcli2 /cx set pr patrolread factory
storcli /cx set maintenance mode=normal nodevices	—
storcli /cx set personality=RAID HBA JBOD	storcli2 /cx set personality id=<personality-id> [force]
storcli /cx set profile profileid=<id>	—
storcli /cx set jbodwritecache=on off default	—
storcli /cx set immediateio=<on off>	—
storcli /cx set largeiosupport=<on off>	—
storcli /cx set unmap=<on off>	—
storcli /cx set driveactivityled=<on off>	—
—	storcli2 /cx ack events seqnum=<sequence-number>
storcli /cx set pdfaileventoptions [detectionType=<val>] [correctiveaction=<val>] [errorThreshold=<val>]	storcli2 /cx set pdfaileventoptions [detectionType=<val>] [correctiveaction=<val>] [errorThreshold=<val>]
storcli /cx set assemblynumber= xxxx	—
storcli /cx set config file=<fileName>	—
storcli /cx set flushwriteverify=<on off>	—
storcli /cx set largeQD=<on off>	—
storcli /cx set debug type=<value> option=<value> [level=<value in hex>]	—
storcli /cx set debug reset all	—
storcli /cx set tracernumber= xxxx	—
storcli /cx set sasadd = xxxx [devicename] [methodport]	—
storcli /cx set sasaddhi = xxxxsasaddlow = xxxxx [devicename] [methodport]	—
storcli /cx set updatevpd file=<filepath>	—
storcli /cx set htbpams=off	—
storcli /cx set htbpams [= on] maxsize=<value> minsize=<value> decrementsize=<value>	—

StorCLI Command	StorCLI2 Command
—	storcli2 /cx set bootmode=COE SMOE
storcli /cx start diag [duration=<val>]	storcli2 /cx start diag [duration=<val>]
storcli /cx start dpm	—
storcli /cx start patrolread	storcli2 /cx start patrolread
storcli /cx stop patrolread	storcli2 /cx stop patrolread
—	storcli2 /cx set nvccachekey
storcli /cx flasherase	storcli2 /cx start nvccacheerase
storcli /cx stop dpm	storcli2 /cx stop DPM
storcli /cx set sasaddress=xxxxxxx	—
storcli /cx compare bios ver =<bios version>	—
storcli /cx compare fwprodid ver =<fw product id version>	—
storcli /cx compare ssid ver =<ssid version>	—
storcli /cx compare firmware ver =<firmware version>	—
storcli /cx erase nvram	—
storcli /cx erase fwbackup	—
storcli /cx erase bootservices	—
storcli /cx erase all [excludemfg] [file=filename]	—
storcli /cx erase perconpage	—
storcli /cx erase mpb	—
—	storcli2 /cx set smartpollinterval=<value> pdtype=Internal External
—	storcli2 /cx set drivewceforrebuild=<on off>
—	storcli2 /cx set BaseEnclLevel=<val>
—	storcli2 /cx set CacheOffloadEncType=<val>
—	storcli2 /cx set name=<name>
—	storcli2 /cx set DPM [duration=<val>] [raFactor=<value>]
—	storcli2 /cx set maintainjbodfailhistory=<on off>
—	storcli2 /cx set jbodsesmgt=<on off>
—	storcli2 /cx set pdtempoll=<on off>
—	storcli2 /cx set fwjbodsecurity=<on off>
—	storcli2 /cx set hostjbodsecurity=<on off>
—	storcli2 /cx set smartpoll=<on off> pdtype=RAID JBOD
—	storcli2 /cx spinup drives= <e:s e:s-x e:s-x,y>
—	storcli2 /cx set datalosswarning=<on off>
Image Upload Commands	
storcli /cx get vpd file=<fileName>	—
storcli /cx get config file=<fileName>	—

StorCLI Command	StorCLI2 Command
storcli /cx get bios file=<filename>	storcli2 /c0 get component type=<val> [location=<val>] file=<filepath>
storcli /cx get firmware file=<filename>	
storcli /cx get mpb file=<filename>	
storcli /cx get fwbackup file=<filename>	
storcli /cx get nvdata file=<filename>	
storcli /cx get flash file=<filename>	
PHY Commands	
storcli /cx/px show	storcli2 /cx/px show [phytype=sas pcie]
storcli /cx/px show phyerrorcounters	storcli2 /cx/px show all [phytype=sas pcie]
storcli /cx/px show all	storcli2 /cx/px set saslinkspeed=6 12 22.5
storcli /cx/px set linkspeed=0 1.5 3 6 12	storcli2 /cx/px set pcielinkspeed=2.5 5 8 16 storcli2 /cx/px set pcielinkspeed=2.5 5 8 16
storcli /cx/px set state=on off	storcli2 /cx/px set state=on off [phytype=sas pcie]
storcli /cx/px reset [hard]	—
storcli /cx/px compare linkspeed=<speed>	—
Energy Pack Commands	
storcli /cx/bbu show	storcli2 /cx/ep show
storcli /cx/bbu show all	storcli2 /cx/ep show all
storcli /cx/bbu show status	—
storcli /cx/bbu show properties	—
storcli /cx/bbu show learn	—
storcli /cx/bbu show gasgauge Offset=xxxx Numbytes=n	—
storcli /cx/bbu start learn	—
storcli /cx/bbu show modes	—
storcli /cx/bbu set [learnDelayInterval=<val> bbuMode=<val> learnStartTime=[DDD HH off] autolearnmode=<val> powermode=sleep writeaccess=sealed]	storcli2 /cx/ep set autolearnmode=<val>
storcli /cx/cv show	storcli2 /cx/ep show
storcli /cx/cv show all	storcli2 /cx/ep show all
storcli /cx/cv show status	—
storcli /cx/cv show learn	—
storcli /cx/cv start learn	—

Events, Messages, and Behaviors

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

The software monitors the activity and performance of all controllers in the workstation and the devices that are 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 main menu window. The messages are also logged in the Windows Application log (Event Viewer).

- [Error Levels](#)
- [Event Messages](#)

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 53: Event Error Levels

Severity Level	Meaning
Progress	Progress message. No user action is necessary.
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.
Fault	The I/O Unit faulted due to a catastrophic error.

Event Messages

The following table lists the event messages. The event message descriptions include placeholders for specific values that are determined when the event is generated. For example, in message 0x0006 in the Event Messages table, “%s” is replaced by the connector number.

Table 54: Event Messages

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x0000	Information	Dummy event to know the start of boot events. This event is never logged.	Boot event start - unused, just a marker.
0x0001	Information	Memory/energy pack problems were detected. An unexpected power loss occurred. The adapter has been recovered, but the controller cache was lost. The volumes may be inconsistent. Check-Consistency is required to ensure redundancy.	Logged when Memory/energy pack problems are detected. The adapter has recovered, but the controller cache was discarded.
0x0002	Information	A foreign configuration found on adapter.	Logged when a foreign configuration is found on an adapter.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x0003	Information	An enclosure contains both SAS and SATA drives, but this controller does not allow mixed drive types in a single enclosure. Correct the problem and restart the system.	Logged when an enclosure is found containing both SAS and SATA drives, but the controller does not allow mixed drive types in a single enclosure.
0x0004	Information	SAS drives were detected, but this controller does not support SAS drives. Remove the SAS drives and restart the system.	Logged when SAS drives are detected, but the controller does not support SAS drives.
0x0005	Information	SATA drives were detected, but this controller does not support SATA drives. Remove the SATA drives and restart the system.	Logged when SATA drives are detected, but this controller does not support SATA drives.
0x0006	Critical	The total number of enclosures that are connected to connector %s has exceeded the maximum allowable limit. Remove the extra enclosures and restart the system.	Logged when the total number of enclosures that are connected to a connector has exceeded the maximum allowable limit of enclosures.
0x0007	Information	Invalid SAS topology detected: %s. Check the cable configurations, repair the problem, and restart the system.	Logged when an invalid SAS topology is detected.
0x0008	Information	The energy pack hardware is missing or malfunctioning, the energy pack is unplugged, or the energy pack could be fully discharged. If you continue to boot the system, the energy pack-backed cache does not function. If the energy pack is connected and has charged for 30 minutes and this message continues to appear, contact technical support for assistance.	Logged when the energy pack hardware is missing or malfunctioning, the energy pack is unplugged, or fully discharged.
0x0009	Information	An invalid SAS address is detected. Program a valid SAS address and restart the system.	Logged when an invalid SAS address is present.
0x000A	Information	Some configured disks have been removed from the system or are no longer accessible. Check the cables and ensure that the all disks are present.	Logged when configured disks have been removed from the system or the disks are no longer accessible.
0x000B	Information	The following VDs have missing disks: %s. If you proceed (or load the configuration utility), these VDs are marked OFFLINE and may be inaccessible.	Logged when the VDs have missing disks.
0x000C	Information	The following VDs are missing: %s. If you proceed (or load the configuration utility), these VDs are removed from the configuration. If you wish to use them later, they have to be imported. If you believe these VDs should be present, power off the system and check the cables to ensure that all the disks are present.	Logged when the VDs are missing.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x000D	Information	The following VD's are missing complete spans: %s. If you proceed (or load the configuration utility), these VD's are removed from the configuration, the remaining drives marked as foreign. If you wish to use them later, restore the missing spans and use the foreign import option to recover the VD's. If you believe these VD's should be present, power off the system and check the cables to ensure that all the disks are present.	Logged when the VD's are missing complete spans.
0x000E	Information	All the disks from the previous configuration are gone. If this message is unexpected, then power off the system and check the cables to ensure all that disks are present.	Logged when all the disks from the previous configuration are gone.
0x000F	Information	The write-back VD's are temporarily running in write-through mode. This issue occurs because the energy pack or super capacitor being charged is missing or bad. If you are using an energy pack, allow the energy pack to charge for 24 hours before evaluating the energy pack for replacement. You can evaluate the health of the energy pack or super capacitor by using the appropriate utility within the operating system or within POST.	Logged when the VD's that are configured for write-back are temporarily running in write-through mode.
0x0010	Information	The controller cache was discarded due to an unexpected power-off or a reboot during a write operation, but the adapter has recovered. This issue could be due to memory problems, a bad energy pack, or you may not have an energy pack installed.	Logged when the controller cache is discarded due to an unexpected power-off or a reboot during a write operation, but the adapter has recovered.
0x0011	Information	Multi-bit ECC errors were detected on the RAID controller. The DIMM on the controller needs replacement. Contact technical support to resolve this issue.	Logged when multi-bit ECC errors are detected on the RAID controller and the DIMM on the controller must be replaced.
0x0012	Information	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.	Logged when single-bit ECC errors were detected during the previous boot of the RAID controller and the DIMM on the controller needs replacement.
0x0013	Information	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.	Logged when single-bit overflow ECC errors were detected during the previous boot of the RAID controller and the DIMM on the controller needs replacement.
0x0014	Information	Number of disks exceeded the maximum supported count of %lu disks. Remove the extra drives and reboot the system to avoid losing data.	Logged when the number of disks exceeded the maximum supported count of disks.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x0015	Information	A discovery error has occurred: %s, power cycle the system and all the enclosures that are attached to this system.	Logged to post a topology error.
0x0016	Information	Drive security is enabled on this controller and a passphrase is required. Enter the passphrase.	Logged during boot for a boot time password.
0x0017	Information	Invalid passphrase. Enter the passphrase.	Logged when an invalid passphrase is entered.
0x0018	Information	A drive security key error occurred. All secure drives are locked and marked as foreign.	Logged when a controller key is invalid.
0x0019	Information	Invalid passphrase. If you continue, there is a drive security error and all the secured configurations are locked and marked as foreign. Reboot the machine to retry the passphrase.	Logged when a user entered an invalid passphrase at boot. If you continue, all the secure configurations are marked as foreign.
0x001A	Information	Unable to communicate with the EKM server. If you continue, a drive security key error occurs and all the secured configurations are marked as foreign. Check the connection with the EKM server, and reboot the machine to retry switching to EKM.	Logged during boot when the server failed to receive the EKM key.
0x001B	Information	Unable to change security to EKMS as not able to communicate to EKMS. If you continue, the drive security remains in the existing security mode. Check the connection with the EKMS, reboot the machine to retry the EKMS.	Logged during boot when the EKM rekey failed.
0x001C	Information	There are Offline or missing virtual drives with preserved cache. Check the cables and ensure that all drives are present.	Logged during boot when pinned cache is present.
0x001D	Information	There are offline or missing virtual drives with preserved cache. Check the cables and ensure that all drives are present.	Logged during boot when pinned cache is present.
0x001E	Information	The native configuration is not supported by the controller. Check the controller, ibutton, key-vault, or the Feature on-demand key. If you continue, the configuration is marked foreign.	Logged during DDF configuration and the configuration is incompatible.
0x001F	Information	The most recent configuration command could not be committed and must be retried.	Logged during DDF configuration read when the configuration is lost.
0x0020	Information	The firmware could not sync up configuration/property changes for some of the VD/PDs.	Logged when the firmware cannot sync up config/prop changes for some of the VD/PDs.
0x0021	Information	The Foreign configuration import did not import any drives.	Logged when the user tries to import invalid or incomplete foreign configurations.
0x0022	Information	Cannot communicate with iButton to retrieve the premium features. The loss of communication occurs because of extreme temperatures. The system has halted!	Logged when unable to communicate with iButton to retrieve the premium features. The loss of communication occurs because of extreme temperatures. The system is halted.
0x0023	Information	Waiting for the energy pack to get fully charged.	Logged when waiting for the energy pack to be fully charged.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x0024	Information	The energy pack charging is completed.	Logged when the energy pack charging is completed.
0x0025	Information	The energy pack charging is not completed.	Logged when the energy pack charging is completed.
0x0026	Information	Nonvolatile Cache capacity is too less to support data backup. The write-back VDs are converted to write-through.	Logged when Boot MSG ONFI backup capacity is over.
0x0027	Information	Data backup capacity of the Nonvolatile Cache device is degraded, consider a replacement.	Logged when the NVCache data backup capacity has decreased. Consider a replacement.
0x0028	Information	Nonvolatile Cache has gone bad.	Logged when an NVCache device failed and cannot support data retention.
0x0029	Information	Waiting for the energy pack learning to complete.	Logged when waiting for an energy pack learn to complete.
0x002A	Information	The energy pack learning is completed.	Logged when an energy pack learn is completed.
0x002B	Information	Error restoring offloaded cache. Data cache is lost.	Logged when an error restoring offloaded cache occurs. The data cache is lost.
0x002C	Information	Cache offload failed. Data cache is lost.	Logged when failed to initiate a cache offload. The data cache is lost.
0x002D	Information	OCM data restore failed upon reset. Data cache is lost.	Logged when OCM data restore failed upon reset. The data cache is lost.
0x002E	Critical	NVDATA Validation failed. Flash new firmware with the correct NVDATA.	Logged when NVDATA Validation failed. Flash new firmware with correct NVDATA.
0x002F	Information	MPB read failed for personality/OemId/SubOemId, switching to SafeMode.	Logged when a read failed for personality/OemId/SubOemId, or Switching to SafeMode.
0x0030	Information	Mismatch in the PNP ID during the premium feature initialization. Contact technical support for assistance.	Logged when a mismatch in the PNP ID during the premium feature initialization. Contact technical support for assistance.
0x0031	Information	NVMe PRP drives were detected, but this controller does not support NVMe PRP drives. Remove the NVMe PRP drives and restart the system.	Logged when NVMe PRP drives are detected, but the controller does not support NVMe PRP drives.
0x0032	Information	NVMe SGL drives were detected, but this controller does not support NVMe SGL drives. Remove the NVMe SGL drives and restart the system.	Logged when NVMe SGL drives are detected, but the controller does not support NVMe SGL drives.
0x0033	Information	SSD drives were detected, but this controller does not support SSD drives. Remove the SSD drives and restart the system.	Logged when SSD drives are detected, but the controller does not support SSD drives.
0x0034	Information	HDD drives were detected, but this controller does not support HDD drives. Remove the HDD drives and restart the system.	Logged when HDD drives are detected, but the controller does not support HDD drives.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x0035	Information	NVRAM of the controller is erased.	Whenever there is an NVRAM layout change between firmware upgrades/downgrades, NVRAM migration is done. For example, NVRAM contents are moved from one layout to another. If there is a power loss or reset while doing migration, NVRAM data is in an inconsistent state. The only way to recover is to erase the NVRAM and use it.
0x0036	Information	The controller successfully received an EKM key at boot time.	Logged during boot when the controller successfully received an EKM key at boot time.
0x0037	Information	Cache offload failed; however, no data cache was lost, potentially only write journals were lost.	Logged when a cache offload failed, but no data cache was lost. Potentially, only the write journals were lost.
0x0038	Information	The controller booted to safe mode due to critical errors.	Logged when the controller booted to safe mode due to critical errors.
0x0039	Information	The controller has exited the safe mode.	Logged when the controller has exited the safe mode.
0x0100	Information	Firmware initialization started.	Logged when the firmware initialization started.
0x0101	Information	Configuration cleared.	Logged when the configuration is cleared.
0x0102	Information	The Background initialization rate changed to %d%%	Logged when the BGI rate is changed with the DCMD.
0x0103	Fatal	The controller cache was discarded due to memory or energy pack problems.	Logged when, on boot, the controller cache is not dirty as per DDR, but the NVSRAM reflects that the cache should be dirty. This fault indicates a cache offload.
0x0104	Information	Cache data was recovered successfully.	Logged when the cache data is recovered successfully.
0x0105	Information	The consistency check rate changed to %d%%	Logged when the CC rate is changed with the DCMD.
0x0106	Fault	Fatal firmware error: %s	Logged when a fatal firmware error occurs.
0x0107	Information	The factory defaults are restored.	Logged during an end-user factory reset, not done for a manufacturing factory reset.
0x0108	Critical	Flash erase error.	Logged as part of a flash write error during the firmware download.
0x0109	Critical	Flash programming error.	Logged during IOP code is logging this event.
0x010A	Information	Shutdown activity completed.	Logged in the shutdown path. This event happens when the driver unloads.
0x010B	Information	The event log was cleared.	Logged when the PEL logs are cleared.
0x010C	Information	The event log was wrapped.	Logged when the PEL logs wrap around.
0x010D	Fault	Not enough controller memory.	Logged when not enough controller memory is available.
0x010E	Information	The patrol read is completed.	Logged when a patrol read cycle completes on the controller.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x010F	Information	The patrol read is paused.	Logged when a patrol read is paused.
0x0110	Information	The patrol read rate changed to %d%%	Logged when the patrol read rate is changed from the DCMD.
0x0111	Information	The patrol read resumed.	Logged when a patrol read is resumed from DCMD or from a power cycle.
0x0112	Information	The patrol read started.	Logged when a patrol read is started from DCMD.
0x0113	Information	The rebuild rate changed to %d%%	Logged when the rebuild rate is changed from DCMD.
0x0114	Information	The OCE rate changed to %d%%	Logged when the OCE rate is changed from DCMD.
0x0115	Information	The driver shutdown command was received from the host.	Logged when a driver shutdown request is received from the host.
0x0116	Information	Test event: %s	Logged for testing purpose.
0x0117	Warning	Background initialization aborted on VD 0x%x	Logged whenever the BGI operation is aborted.
0x0118	Information	Background initialization corrected medium error (VD 0x%x at 0x%llx, PD 0x%02x(e0x%02x/s%d) Lun 0x%llx at 0x%llx count 0x%x)	Logged when a medium error is corrected because of the BGI.
0x0119	Information	Background initialization completed on VD 0x%x	Logged when the BGI completes successfully.
0x011A	Fatal	Background initialization completed with uncorrectable errors on VD 0x%x	Logged when the BGI completes with uncorrectable errors.
0x011B	Fatal	Background initialization detected uncorrectable multiple medium errors (VD 0x%x at 0x%llx on PD 0x%02x(e0x%02x/s%d) Lun 0x%llx at 0x%llx count 0x%x)	Logged when the BGI encounters a double medium error.
0x011C	Critical	Background initialization failed on VD 0x%x	Logged when BGI completes with fatal errors.
0x011D	Progress	Background initialization is in progress on VD 0x%x is %s	Logged when a Background Initialization is in progress on an LD.
0x011E	Information	Background initialization started on VD 0x%x	Logged when BGI starts on an LD.
0x011F	Information	Cache policy change on VD 0x%x to [cp=%02x dc=%02x dbgi=%02x] from [cp=%02x dc=%02x dbgi=%02x]	Logged when the LD cache policy is changed.
0x0120	Information	Consistency check aborted on VD 0x%x	Logged when a CC operation is aborted.
0x0121	Information	Consistency check corrected medium error (VD 0x%x at 0x%llx, PD 0x%02x(e0x%02x/s%d) Lun 0x%llx at 0x%llx count 0x%x)	Logged when a medium error is corrected because of a CC.
0x0122	Information	Consistency check done on VD 0x%x	Logged when a CC operation completes.
0x0123	Information	Consistency check done with corrections on VD 0x%x, (corrections=%d)	Logged when a CC operation completes with inconsistencies found.
0x0124	Fatal	Consistency check detected uncorrectable multiple medium errors (VD 0x%x at 0x%llx on PD 0x%02x(e0x%02x/s%d) Lun 0x%llx at 0x%llx count 0x%x)	Logged when a CC encounters double medium error.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x0125	Critical	Consistency check failed on VD 0x%x	Logged when a CC completes with fatal errors.
0x0126	Fatal	Consistency check completed with uncorrectable errors on VD 0x%x	Logged when a CC completes with uncorrectable errors.
0x0127	Information	Consistency check found inconsistent parity on VD 0x%x at strip 0x%llx	Logged when a CC found inconsistent parity on LD at strip.
0x0128	Warning	Consistency check inconsistency logging disabled on VD 0x%x (too many inconsistencies)	Logged when the inconsistencies found is more than the logging limit.
0x0129	Progress	Consistency check progress on VD 0x%x is %s	Logged to report a CC progress.
0x012A	Information	Consistency check started on VD 0x%x	Logged when a CC operation is started.
0x012B	Information	Initialization aborted on VD 0x%x	Logged when an INIT operation is aborted.
0x012C	Critical	Initialization failed on VD 0x%x	Logged when an INIT operation fails.
0x012D	Progress	Initialization progress on VD 0x%x is %s	Logged to report the INIT progress.
0x012E	Information	Fast initialization started on VD 0x%x	Logged when a FAST INIT is started from DCMD.
0x012F	Information	Full initialization started on VD 0x%x	Logged when a FULL INIT is started from DCMD.
0x0130	Information	Initialization completed on VD 0x%x	Logged when an INIT completes successfully.
0x0131	Information	VD 0x%x properties updated to [cp=%02x,dc=%02x,dbgj=%02x]	Logged when an LD property changes.
0x0132	Information	OCE completed on VD 0x%x	Logged when the OCE completes successfully.
0x0133	Fatal	OCE of VD 0x%x stopped due to unrecoverable errors	Logged when the OCE completes with errors.
0x0134	Fatal	OCE detected uncorrectable multiple medium errors (VD 0x%x at 0x%llx on PD 0x%02x(e0x%02x/s%d) Lun 0x%llx at 0x%llx count 0x%x)	Logged when the OCE detected uncorrectable multiple medium errors at an LBA.
0x0135	Progress	OCE is in progress on VD 0x%x is %s	Logged to report the OCE progress.
0x0136	Information	OCE resumed on VD 0x%x	Logged when the OCE resumes from a power cycle or when OCE use resumes.
0x0137	Fatal	OCE resume of VD 0x%x failed due to a configuration mismatch.	Logged when the OCE fails to resume from a power cycle.
0x0138	Information	OCE started on VD 0x%x	Logged when the OCE is started.
0x0139	Information	LD state changed on VD 0x%x from %s(%x) to %s(%x)	Logged when an LD state changes.
0x013A	Information	Clear aborted on PD 0x%02x(e0x%02x/s%d)	Logged when the PD clear operation is aborted.
0x013B	Critical	Clear failed on PD 0x%02x(e0x%02x/s%d) Path 0x%llx (Error 0x%02x)	Logged when the PD clear operation fails.
0x013C	Progress	Clear is in progress on PD 0x%02x(e0x%02x/s%d) Lun 0x%llx is %s	Logged to update the progress of the PD clear operation.
0x013D	Information	Clear started on PD 0x%02x(e0x%02x/s%d)	Logged when the PD clear operation is started from the DCMD.
0x013E	Information	Clear completed on PD 0x%02x(e0x%02x/s%d)	Logged when the PD clear operation completes with success.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x013F	Information	Error on PD 0x%02x(e0x%02x/s%d) Path 0x%llx (Error 0x%02x)	Logged when an error occurs on the PD.
0x0140	Warning	PD 0x%02x(e0x%02x/s%d) is not supported	Logged when a PD is not supported.
0x0141	Information	Patrol read corrected a medium error on PD 0x%02x(e0x%02x/s%d) Lun 0x%llx at 0x%llx	Logged when a patrol read corrects an LBA.
0x0142	Progress	Patrol read is in progress on PD 0x%02x(e0x%02x/s%d) Lun 0x%llx is %s	Logged while a patrol read is in progress for every 10% completion.
0x0143	Fatal	A patrol read found an uncorrectable medium error on PD 0x%02x(e0x%02x/s%d) Lun 0x%llx at 0x%llx	Logged when a patrol read recovery fails for an LBA, as an uncorrectable medium error.
0x0144	Warning	Predictive failure: PD 0x%02x(e0x%02x/s%d)	Logged when a predictive failure is detected for a PD.
0x0145	Fatal	Puncturing bad block on PD 0x%02x(e0x%02x/s%d) Lun 0x%llx at 0x%llx	Logged when puncturing bad block at an LBA.
0x0146	Information	Rebuild aborted by user on PD 0x%02x(e0x%02x/s%d)	Logged when a rebuild operation is aborted by a user.
0x0147	Information	Rebuild completed on VD 0x%x	Logged when a rebuild operation completes on an LD.
0x0148	Information	Rebuild completed on PD 0x%02x(e0x%02x/s%d)	Logged when a rebuild operation is completed on a PD.
0x0149	Critical	Rebuild failed on PD 0x%02x(e0x%02x/s%d) due to a source drive error	Logged when a rebuild operation fails on a PD because of a source drive error.
0x014A	Critical	Rebuild failed on PD 0x%02x(e0x%02x/s%d) due to a target drive error	Logged when a rebuild operation fails on a PD because of a medium error.
0x014B	Progress	Rebuild is in progress on PD 0x%02x(e0x%02x/s%d) Lun 0x%llx is %s	Logged when a rebuild operation is in progress on a PD.
0x014C	Information	Rebuild resumed on PD 0x%02x(e0x%02x/s%d)	Logged when a rebuild operation is resumed on a PD.
0x014D	Information	Rebuild started on PD 0x%02x(e0x%02x/s%d)	Logged when a rebuild operation is started on a PD.
0x014E	Information	Rebuild automatically started on PD 0x%02x(e0x%02x/s%d)	Logged when a rebuild operation is automatically started on a PD.
0x014F	Fatal	Reassign-write operation fails on PD 0x%02x(e0x%02x/s%d) Lun 0x%llx at 0x%llx	Logged when a reassign write operation fails on a PD at the LBA.
0x0150	Fatal	Unrecoverable medium error during Rebuild on PD 0x%02x(e0x%02x/s%d) Lun 0x%llx at 0x%llx	Logged when an unrecoverable medium error is found during a rebuild operation at an LBA for a PD.
0x0151	Information	Corrected medium error during recovery on PD 0x%02x(e0x%02x/s%d) Lun 0x%llx at 0x%llx	Logged when a corrected medium error occurs during a recovery operation at an LBA for a PD.
0x0152	Fatal	Unrecoverable medium error during recovery on PD 0x%02x(e0x%02x/s%d) Lun 0x%llx at 0x%llx	Logged when an unrecoverable medium error occurs during a recovery operation at an LBA for a PD.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x0153	Information	Unexpected sense: PD 0x%02x(e0x%02x/s%d) Path 0x%llx, CDB:%s, Sense:%01x/%02x/%02x	Logged when an unexpected sense is received for a CDB on a PD.
0x0154	Information	State is changed on PD 0x%02x(e0x%02x/s%d) from %s(%x) to %s(%x)	Logged when the PD state is changed.
0x0155	Fatal	Unable to access device PD 0x%02x(e0x%02x/s%d)	Logged when unable to access the PD.
0x0156	Information	Dedicated hot spare created on PD 0x%02x(e0x%02x/s%d) (%s)	Logged when a dedicated hot spare is created.
0x0157	Information	Dedicated hot spare PD 0x%02x(e0x%02x/s%d) (%s) is disabled	Logged when a dedicated hot spare is removed.
0x0158	Warning	Dedicated hot spare PD 0x%02x(e0x%02x/s%d) is no longer useful for all arrays	Logged when a dedicated hot spare is no longer useful for all arrays.
0x0159	Information	Global hot spare created on PD 0x%02x(e0x%02x/s%d) (%s)	Logged when a global hot spare is created.
0x015A	Information	Global hot spare PD 0x%02x(e0x%02x/s%d) (%s) is disabled	Logged when a global hot spare is removed.
0x015B	Warning	Global hot spare PD 0x%02x(e0x%02x/s%d) does not cover all arrays	Logged when a global hot spare does not cover all arrays.
0x015C	Information	Created VD 0x%x	Logged when an LD is created.
0x015D	Information	Deleted VD 0x%x	Logged when an LD is deleted
0x015E	Information	The energy pack is present.	Logged when an energy pack is present.
0x015F	Warning	The energy pack is not present.	Logged when an energy pack is not present.
0x0160	Information	A new energy pack was detected.	Logged when a new energy pack is detected.
0x0161	Information	The energy pack has been replaced.	Logged when an energy pack has been replaced.
0x0162	Warning	The energy pack temperature is high.	Logged when the energy pack temperature is high.
0x0163	Warning	The energy pack voltage is low.	Logged when the energy pack voltage is low.
0x0164	Information	The energy pack started charging.	Logged when the energy pack starts charging.
0x0165	Information	The energy pack is discharging.	Logged when the energy pack is discharging.
0x0166	Information	The energy pack temperature is normal.	Logged when the energy pack temperature is normal.
0x0167	Fatal	The energy pack has failed and cannot support data retention. Replace the energy pack.	Logged when an energy pack has failed and cannot support data retention. This event indicates that the user must replace the energy pack.
0x0168	Information	The energy pack relearn operation started.	Logged when an energy pack relearn operation starts.
0x0169	Information	The energy pack relearn operation is in progress.	Logged when an energy pack relearn operation is in progress.
0x016A	Information	The energy pack relearn operation is completed.	Logged when an energy pack relearn operation is completed.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x016B	Warning	The energy pack relearn operation timed out.	Logged when an energy pack relearn operation has timed out.
0x016C	Information	The energy pack relearn operation is pending: Energy pack is under charge.	Logged when an energy pack relearn operation is pending and the energy pack is being charged.
0x016D	Information	The energy pack relearn operation is postponed.	Logged when an energy pack relearn is postponed
0x016E	Warning	The energy pack is removed.	Logged when an energy pack is removed.
0x016F	Warning	The current capacity of the energy pack is below the threshold.	Logged when the current capacity of the energy pack is below the threshold.
0x0170	Information	The current capacity of the energy pack is above the threshold.	Logged when the current capacity of the energy pack is above the threshold.
0x0171	Information	Enclosure (SES) discovered on PD 0x%02x(e0x%02x/s%d)	Logged when an enclosure (SES) is discovered.
0x0172	Critical	Enclosure PD 0x%02x communication lost.	Logged when an enclosure has lost communication.
0x0173	Information	Enclosure PD 0x%02x communication restored.	Logged when the enclosure's communication is restored.
0x0174	Critical	Enclosure PD 0x%02x(ELI-0x%llx/p%d) fan %d failed.	Logged when the enclosure's fan fails.
0x0175	Information	Enclosure PD 0x%02x(ELI-0x%llx/p%d) fan %d inserted.	Logged when the enclosure's fan is inserted.
0x0176	Warning	Enclosure PD 0x%02x(ELI-0x%llx/p%d) fan %d removed.	Logged when the enclosure's fan is removed.
0x0177	Information	Enclosure PD 0x%02x(ELI-0x%llx/p%d) power supply %d inserted.	Logged when the enclosure's power supply is inserted.
0x0178	Warning	Enclosure PD 0x%02x(ELI-0x%llx/p%d) power supply %d removed.	Logged when the enclosure's power supply is removed.
0x0179	Critical	Enclosure PD 0x%02x(ELI-0x%llx/p%d) EMM %d failed.	Logged when the enclosure's EMM is failed.
0x017A	Information	Enclosure PD 0x%02x(ELI-0x%llx/p%d) EMM %d inserted.	Logged when the enclosure's EMM is inserted.
0x017B	Critical	Enclosure PD 0x%02x(ELI-0x%llx/p%d) EMM %d removed.	Logged when the enclosure's EMM is removed.
0x017C	Warning	Enclosure PD 0x%02x(ELI-0x%llx/p%d) temperature sensor %d is below the warning threshold.	Logged when the enclosure's temperature sensor is below the warning threshold.
0x017D	Critical	Enclosure PD 0x%02x(ELI-0x%llx/p%d) temperature sensor %d is below the error threshold.	Logged when the enclosure's temperature sensor is below the warning threshold.
0x017E	Warning	Enclosure PD 0x%02x(ELI-0x%llx/p%d) temperature sensor %d is above the warning threshold.	Logged when the enclosure's temperature sensor is above the warning threshold.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x017F	Critical	Enclosure PD 0x%02x(ELI-0x%llx/p%d) temperature sensor %d is above the error threshold.	Logged when the enclosure's temperature sensor is above the error threshold.
0x0180	Critical	Enclosure PD 0x%02x shutdown.	Logged when the enclosure shutdowns.
0x0181	Warning	Enclosure PD 0x%02x is not supported; too many enclosures that connected to the port.	Logged when the enclosure is not supported due to too many enclosures that are connected to the port.
0x0182	Critical	Enclosure PD 0x%02x(ELI-0x%llx/p%d) firmware mismatch (EMM %d)	Logged when the enclosure firmware revisions are inconsistent between SIMs in the same enclosure.
0x0183	Warning	Enclosure PD 0x%02x(ELI-0x%llx/p%d) sensor %d is bad.	Logged when the enclosure's sensor becomes bad.
0x0184	Critical	Enclosure PD 0x%02x(ELI-0x%llx/p%d) phy is bad for the slot %d	Logged when the enclosure has phy bad for slot.
0x0185	Critical	Enclosure PD 0x%02x is unstable.	Logged when the enclosure is unstable.
0x0186	Critical	Enclosure PD 0x%02x has a hardware error.	Logged when the enclosure has a hardware error.
0x0187	Critical	Enclosure PD 0x%02x is not responding.	Logged when the enclosure is not responding.
0x0188	Information	PD 0x%02x(e0x%02x/s%d) is too small to be used for auto-rebuild.	Logged when the PD too small to be used for an auto-rebuild on operation.
0x0189	Information	BEM enabled; changing write-through virtual disks to write-back.	Logged when the BEM is enabled.
0x018A	Warning	BEM disabled; changing write-back virtual disks to write-through, Forced write-back VDs are not affected.	Logged when the BEM is disabled.
0x018B	Critical	Energy pack/charger problems detected; SOH is Bad.	Logged when an energy pack or charger problems are detected. The SOH is bad.
0x018C	Warning	Enclosure PD 0x%02x(ELI-0x%llx/p%d) Power supply %d switched off.	Logged when the enclosure power supply is switched off.
0x018D	Information	Enclosure PD 0x%02x(ELI-0x%llx/p%d) Power supply %d switched on.	Logged when the enclosure power supply is switched on.
0x018E	Warning	Enclosure PD 0x%02x(ELI-0x%llx/p%d) Power supply %d cable removed.	Logged when the enclosure power supply cable is removed.
0x018F	Information	Enclosure PD 0x%02x(ELI-0x%llx/p%d) Power supply %d cable inserted.	Logged when an enclosure power supply cable is inserted.
0x0190	Information	Enclosure PD 0x%02x(ELI-0x%llx/p%d) Fan %d returned to normal.	Logged when the enclosure fan returns to normal operation.
0x0191	Information	BEM retention test was initiated on a previous boot.	Logged when the BEM retention test was initiated on a previous boot.
0x0192	Information	BEM retention test passed.	Logged when the BEM retention test passed.
0x0193	Critical	BEM retention test failed!	Logged when the BEM retention test failed.
0x0194	Information	NVRAM retention test was initiated on a previous boot	Logged when the NVRAM retention test was initiated on a previous boot.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x0195	Information	NVRAM retention test passed.	Logged when the NVRAM retention test passed.
0x0196	Critical	NVRAM retention test failed!	Logged when the NVRAM retention test failed.
0x0197	Information	Diagnostic %s test completed %d passes successfully.	Logged when individual tests under the diagnostic test completed the required number of passes successfully.
0x0198	Critical	Diagnostic %s test FAILED on %d pass. Fail data: errorOffset=0x%x goodData=0x%x badData=0x%	Logged when individual tests under the diagnostic test failed on a specific pass.
0x0199	Information	Self-check diagnostics completed.	Logged when the self-check diagnostics completes.
0x019A	Information	Foreign configuration detected.	Logged when a foreign configuration is detected.
0x019B	Information	Foreign configuration imported.	Logged when a foreign configuration is imported.
0x019C	Information	Foreign configuration cleared.	Logged when a foreign configuration is cleared.
0x019D	Warning	Enclosure PD 0x%02x(ELI-0x%llx/p%d) temperature sensor %d differential detected.	Logged when the enclosure temperature sensor differential is detected.
0x019E	Information	Diagnostic disk test cannot start. No qualifying disks found. ^a	Logged when the disk test cannot start. No qualifying disks are found.
0x019F	Information	Time duration that is provided by host is not sufficient for self-check diagnostics.	Logged when the time duration that is provided by host is not sufficient for a self-check.
0x01A0	Information	PD 0x%02x(e0x%02x/s%d) marked as missing on the array %d row %d	Logged when a PD is marked as missing in an array on a row.
0x01A1	Information	Replaced missing for PD 0x%02x(e0x%02x/s%d) on the array %d row %d	Logged when a PD is replaced from a missing array on a row.
0x01A2	Information	Enclosure PD 0x%02x(ELI-0x%llx/p%d) temperature sensor %d returned to normal.	Logged when the enclosure temperature sensor has returned to normal.
0x01A3	Information	Enclosure PD 0x%02x Firmware download is in progress.	Logged when the enclosure firmware download is in progress.

a. If the test is run without any drives or enclosures present, the test result is printed as `Unsupported` and the link test results prints `Resource Unavailable`.

Detailed Status :

Type Name StartTime(LocalTime yyyy/mm/dd hh:mm:sec) CompletionTime(LocalTime yyyy/mm/dd hh:mm:sec)

PassCount Status

DMA DMA 2023/02/20 12:04:02 2023/02/20 12:04:22 72411 Success

XOR XOR 2023/02/20 12:04:02 2023/02/20 12:04:22 24425 Success

DISK DISK 2023/02/20 12:04:02 2023/02/20 12:04:02 0 Unsupported

Memory MEMORY 2023/02/20 12:04:02 2023/02/20 12:04:22 2127 Success

NVRAM NVRAM 2023/02/20 12:04:02 2023/02/20 12:04:22 34033 Success

LINK LINK 2023/02/20 12:04:02 2023/02/20 12:04:02 1 Resource Unavailable

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x01A4	Warning	Enclosure PD 0x%02x firmware download failed.	Logged when the enclosure firmware download fails.
0x01A5	Warning	PD 0x%02x(e0x%02x/s%d) is not a certified drive.	Logged when the PD is not a certified drive.
0x01A6	Information	PD missing: %s	Logged when the PDs are missing from the configuration at boot.
0x01A7	Warning	VDs missing drives and will go offline at boot: %s	Logged when the VDs are missing drives and go offline at boot.
0x01A8	Warning	VDs missing at boot: %s	Logged when the VDs are missing at boot.
0x01A9	Information	Energy pack charging is completed.	Logged when the energy pack charge is complete.
0x01AA	Information	The Dedicated spare PD 0x%02x(e0x%02x/s%d) imported as global due to missing arrays.	Logged when the dedicated spare is imported as global due to missing arrays.
0x01AB	Information	PD 0x%02x(e0x%02x/s%d) rebuild is not possible as a SAS/SATA mix is not supported in the array.	Logged when a rebuild operation is not possible because SAS/SATA is not supported in an array.
0x01AC	Information	VD 0x%x is now OPTIMAL.	Logged when the LD is optimal.
0x01AD	Warning	VD 0x%x is now PARTIALLY DEGRADED.	Logged when the LD is partially degraded.
0x01AE	Critical	VD 0x%x is now DEGRADED.	Logged when the LD is degraded.
0x01AF	Fatal	VD 0x%x is now OFFLINE.	Logged when the LD is offline.
0x01B0	Critical	Enclosure PD 0x%02x(ELI-0x%llx/p%d) EMM %d is not installed.	Logged when the enclosure's EMM is not installed.
0x01B1	Warning	Global affinity hot spare PD 0x%02x(e0x%02x/s%d) commissioned in a different enclosure.	Logged when the global affinity hot spare is commissioned in a different enclosure.
0x01B2	Warning	Foreign configuration table overflow.	Logged when the foreign configuration table overflows.
0x01B3	Warning	Partial foreign configuration imported, PDs not imported:%s	Logged when the partial foreign configuration is imported, but the PDs are not imported.
0x01B4	Warning	Command timeout on PD 0x%02x(e0x%02x/s%d) Path 0x%llx, CDB:%s	Logged when a command timeout occurs on the PD for a CDB.
0x01B5	Warning	PD 0x%02x(e0x%02x/s%d) Path 0x%llx reset (Type 0x%02x)	Logged when the TM is performed on a PD with a specific task type.
0x01B6	Warning	The bad block table on VD 0x%x is 80%% full.	Logged when the VD bad block table is 80% full.
0x01B7	Fatal	The bad block table on VD 0x%x is full; unable to log the block 0x%llx (on PD 0x%02x(e0x%02x/s%d) Lun 0x%llx at 0x%llx count 0x%x)	Logged when the VD bad block table is full.
0x01B8	Fatal	An uncorrectable medium error was logged for VD 0x%x at 0x%llx (on PD 0x%02x(e0x%02x/s%d) Lun 0x%llx at 0x%llx count 0x%x)	Logged when an uncorrectable medium error is logged for the PD at an LBA.
0x01B9	Information	A medium error was corrected on VD 0x%x at 0x%llx	Logged when a VD medium error is corrected on the VD at an LBA.
0x01BA	Warning	The bad block table on VD 0x%x is 100%% full.	Logged when the VD bad block table is 100% full.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x01BB	Information	Replace started on PD 0x%02x(e0x%02x/s%d) from PD 0x%02x(e0x%02x/s%d)	Logged when a replacedrive operation is started.
0x01BC	Information	Replace aborted on PD 0x%02x(e0x%02x/s%d) and the source is PD 0x%02x(e0x%02x/s%d)	Logged when a replacedrive operation is aborted.
0x01BD	Information	Replace completed on PD 0x%02x(e0x%02x/s%d) from PD 0x%02x(e0x%02x/s%d)	Logged when a replacedrive operation is completed.
0x01BE	Progress	Replace is in progress on PD 0x%02x(e0x%02x/s%d) Lun 0x%llx is %s	Logged when a replacedrive operation is in progress.
0x01BF	Information	Replace resumed on PD 0x%02x(e0x%02x/s%d) from PD 0x%02x(e0x%02x/s%d)	Logged when a replacedrive operation is resumed.
0x01C0	Information	Replace automatically started on PD 0x%02x(e0x%02x/s%d) from PD 0x%02x(e0x%02x/s%d)	Logged when a replacedrive operation is automatically started.
0x01C1	Critical	Replace failed on PD 0x%02x(e0x%02x/s%d) due to a source PD 0x%02x(e0x%02x/s%d) error.	Logged when a replacedrive operation fails.
0x01C2	Warning	Foreign import results in a backward incompatible upgrade of configuration metadata.	Logged when a foreign import results in a backward incompatible upgrade of the configuration metadata.
0x01C3	Information	Redundant path restored for PD 0x%02x(e0x%02x/s%d) Path (%x) 0x%llx	Logged when a redundant path is restored for a PD.
0x01C4	Warning	Redundant path broken for PD 0x%02x(e0x%02x/s%d) Path (%x) 0x%llx	Logged when a redundant path is broken for a PD.
0x01C5	Information	Redundant enclosure EMM Encl PD 0x%02x inserted for EMM Encl PD 0x%02x	Logged when a redundant enclosure EMM is inserted.
0x01C6	Warning	Redundant enclosure EMM Encl PD 0x%02x removed for EMM Encl PD 0x%02x	Logged when a redundant enclosure EMM is removed.
0x01C7	Information	Patrol read cannot 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 a patrol read cannot be started because the PDs are not online, in a VD with an active process, or are in an excluded VD.
0x01C8	Information	Replace aborted by user on PD 0x%02x(e0x%02x/s%d) and source is PD 0x%02x(e0x%02x/s%d)	Logged when a replacedrive operation is aborted by a user on a destination PD.
0x01C9	Critical	Replace aborted on hot spare PD 0x%02x(e0x%02x/s%d) from PD 0x%02x(e0x%02x/s%d), as a hot spare is needed for the rebuild.	Logged when a replacedrive operation is aborted on a hot spare because the hot spare is needed for a rebuild operation.
0x01CA	Warning	Replace aborted on PD 0x%02x(e0x%02x/s%d) from PD 0x%02x(e0x%02x/s%d), as Rebuild is required in an array	Logged when a replacedrive operation is aborted on the destination from source because a rebuild operation is required in the array.
0x01CB	Fatal	Controller cache is discarded for deleted, missing, or offline VD 0x%x	Logged when a controller cache is discarded for being deleted, missing, or offline.
0x01CC	Information	Replace cannot be started as PD 0x%02x(e0x%02x/s%d) is too small for the source PD 0x%02x(e0x%02x/s%d)	Logged when a replacedrive operation cannot be started because the destination PD is too small for the source PD.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x01CD	Information	Replace cannot be started on PD 0x%02x(e0x%02x/s%d) from PD 0x%02x(e0x%02x/s%d), as a SAS/SATA mix is not supported in the array.	Logged when a replacedrive operation cannot be started because SAS/SATA is not supported in an array.
0x01CE	Warning	Microcode update timeout on PD 0x%02x(e0x%02x/s%d)	Logged when the microcode update times out on a device.
0x01CF	Information	Controller properties changed.	Logged when the controller properties are changed.
0x01D0	Information	Patrol read properties changed.	Logged when the patrol read properties are changed.
0x01D1	Information	Consistency check schedule properties changed.	Logged when the CC properties are changed.
0x01D2	Information	Energy pack properties changed.	Logged when the energy pack properties are changed.
0x01D3	Information	Drive security key created.	Logged when a new key is created.
0x01D4	Information	Drive security key changed.	Logged when a key is re-keyed.
0x01D5	Warning	Drive security key re-key operation failed.	Logged when a rekey is progress or a rekey fails.
0x01D6	Warning	Drive security key is invalid.	Logged when the drive security key is invalid.
0x01D7	Information	Drive security key destroyed.	Logged when a key is destroyed.
0x01D8	Warning	Drive security key from ESCROW is invalid.	Logged when the drive unlock fails.
0x01D9	Information	VD 0x%x is now secured.	Logged when an LD is secured.
0x01DA	Warning	VD 0x%x is partially secured.	Logged when at least one drive is not secured in an LD.
0x01DB	Information	PD 0x%02x(e0x%02x/s%d) security activated.	Logged when a drive is secured with a lock key.
0x01DC	Information	PD 0x%02x(e0x%02x/s%d) security disabled.	Logged when a drive is unsecured.
0x01DD	Information	PD 0x%02x(e0x%02x/s%d) is reprovisioned.	Logged when a drive is reprovisioned.
0x01DE	Information	PD 0x%02x(e0x%02x/s%d) security key changed.	Logged when a drive security key is rekeyed.
0x01DF	Fatal	Security subsystem problems detected for PD 0x%02x(e0x%02x/s%d)	Logged when a drive security key fails.
0x01E0	Fatal	The controller cache was preserved for missing or offline VD 0x%x	Logged when an LD is pinned.
0x01E1	Information	The controller cache was discarded by user for VDs: %s	Logged when a user clears the pinned cache.
0x01E2	Information	The controller cache was de-staged for VD 0x%x	Logged when the pinned cached is recovered.
0x01E3	Warning	Consistency check started on an inconsistent VD 0x%x	Logged when a CC operation is started on inconsistent VDs.
0x01E4	Warning	Drive security key failure, cannot access the secured configuration.	Logged when an invalid key is detected.
0x01E5	Warning	Drive security passphrase from the user is invalid.	Logged when the drive security passphrase from the user is invalid.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x01E6	Information	PD 0x%02x(e0x%02x/s%d) rebuild is not possible as an HDD/SSD mix is not supported in the array.	Logged during a rebuild operation when a mix violation is detected.
0x01E7	Information	Replace cannot be started on PD 0x%02x(e0x%02x/s%d) from PD 0x%02x(e0x%02x/s%d), as an HDD/SSD mix is not supported in the array.	Logged during a replacedrive operation when a mix violation is detected.
0x01E8	Information	Bad block table on VD 0x%x is cleared.	Logged when an LD BBM table is cleared.
0x01E9	Information	Cluster of medium errors corrected for VD 0x%x at 0x%llx (on PD 0x%02x(e0x%02x/s%d) Lun 0x%llx at 0x%llx count 0x%x)	Logged when a medium error is corrected.
0x01EA	Information	The controller requests a host bus rescan.	Logged when a drive spin up is completed during the configuration creation.
0x01EB	Information	Controller manufacturing factory defaults restored.	Logged when a user does a factory reset.
0x01EC	Information	Drive security is in external key management mode.	Logged when a key binding type is set to the EKM.
0x01ED	Warning	Drive security failed to communicate with an external key manager.	Logged when failed to receive a key from the EKM.
0x01EE	Information	PD 0x%02x(e0x%02x/s%d) needs key to be %s %s	Logged when a key is required from the EKM
0x01EF	Critical	The configuration command could not be committed to the disk, retry.	Logged when a configuration command could not be committed to a disk.
0x01F0	Information	COD on PD 0x%02x(e0x%02x/s%d) was updated as it was stale.	Logged during boot when the configuration on a disk is not updated.
0x01F1	Warning	VD 0x%x is not available.	Logged when the LUN is not ready.
0x01F2	Information	VD 0x%x is available.	Logged when the LUN is ready.
0x01F3	Information	Advanced Software Options Serial number %s	Logged when a controller serial number is generated.
0x01F4	Information	Host driver is loaded and operational.	Logged when the driver is loaded.
0x01F5	Warning	The foreign configuration auto-import did not import any drives.	Logged during boot if the auto import, imported none.
0x01F6	Warning	Cache-vault microcode update required.	Logged when a battery microcode update is required.
0x01F7	Warning	LD (0x%x) protection information lost.	Logged during the LD import with protection disabled.
0x01F8	Information	Diagnostics passed for PD 0x%02x(e0x%02x/s%d)	Logged when the diagnostic test completes.
0x01F9	Critical	Diagnostics failed for PD 0x%02x(e0x%02x/s%d)	Logged when the diagnostic test fails.
0x01FA	Information	Drive cache settings are enabled during rebuild for PD 0x%02x(e0x%02x/s%d)	Logged during a rebuild operation when the drive is spun up.
0x01FB	Information	Drive cache settings are restored after rebuild for PD 0x%02x(e0x%02x/s%d)	Logged when a rebuild operation is completed.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x01FC	Information	Drive PD 0x%02x(e0x%02x/s%d) commissioned as an emergency spare.	Logged when the PD is commissioned as a spare.
0x01FD	Warning	Reminder: Potential non-optimal configuration due to drive PD 0x%02x(e0x%02x/s%d) commissioned as an emergency spare	Logged periodically when the PD is a spare.
0x01FE	Information	Consistency check suspended on VD 0x%x	Logged when a user suspends a CC.
0x01FF	Information	Consistency check resumed on VD 0x%x	Logged when a CC is resumed.
0x0200	Information	Background initialization suspended on VD 0x%x	Logged when a BGI operation is suspended.
0x0201	Information	Background initialization resumed on VD 0x%x	Logged when a BGI operation is resumed.
0x0202	Information	OCE suspended on VD 0x%x	Logged when an OCE is suspended.
0x0203	Information	Rebuild suspended on PD 0x%02x(e0x%02x/s%d)	Logged when a user suspends the rebuild operation.
0x0204	Information	Replace suspended on PD 0x%02x(e0x%02x/s%d)	Logged when a user suspends the replacedrive operation.
0x0205	Information	Reminder: Consistency check suspended on VD 0x%x	Logged periodically when the CC is suspended.
0x0206	Information	Reminder: Background initialization suspended on VD 0x%x	Logged periodically when the BGI is suspended.
0x0207	Information	Reminder: Rebuild suspended on PD 0x%02x(e0x%02x/s%d)	Logged periodically if the rebuild operation is suspended.
0x0208	Information	Reminder: Replace suspended on PD 0x%02x(e0x%02x/s%d)	Logged periodically if the replacedrive operation is suspended.
0x0209	Information	Reminder: Patrol read suspended	Logged periodically if the patrol read is suspended.
0x020A	Information	Erase aborted on PD 0x%02x(e0x%02x/s%d)	Logged when a PD erase is aborted.
0x020B	Critical	Erase failed on PD 0x%02x(e0x%02x/s%d) Path 0x%llx (Error 0x%02x)	Logged when a PD erase is failed.
0x020C	Progress	Erase is in progress on PD 0x%02x(e0x%02x/s%d) Lun 0x%llx is %s	Logged when a PD erase is in progress.
0x020D	Information	Erase started on PD 0x%02x(e0x%02x/s%d)	Logged when a PD erase is started.
0x020E	Information	Erase completed on PD 0x%02x(e0x%02x/s%d)	Logged when a PD erase is successful.
0x020F	Information	Erase aborted on VD 0x%x	Logged when an LD erase operation is aborted.
0x0210	Critical	Erase failed on VD 0x%x	Logged when an LD erase operation fails.
0x0211	Progress	Erase is in progress on VD 0x%x is %s	Logged when an LD erase operation is in progress.
0x0212	Information	Erase started on VD 0x%x	Logged when an LD erase operation is started.
0x0213	Information	Erase completed on VD 0x%x	Logged when an LD erase operation is successful.
0x0214	Warning	Potential leakage during erase on the VD 0x%x	Logged when an erase operation that is tried on a non-optimal LD or PD erase is not fully erased.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x0215	Warning	Energy pack charging was suspended due to a high energy pack temperature.	Logged when the battery reaches threshold values.
0x0216	Warning	Patrol read aborted on PD 0x%02x(e0x%02x/s%d)	Logged when the user aborts a patrol read.
0x0217	Information	The controller soft reset was completed.	Logged when the controller reset that is requested by the host completes.
0x0218	Critical	Nonvolatile cache capacity is not enough to support the data backup. Write-back VDs are converted to write-through.	Logged when the NVCache capacity is not enough to support the data backup.
0x0219	Critical	Nonvolatile cache device failed, cannot support data retention.	Logged when the NVCache device fails.
0x021A	Information	The controller operating temperature is within normal range, full operation restored.	Logged when the controller operating temperature is within normal range.
0x021B	Warning	Controller temperature threshold exceeded. This warning may indicate inadequate system cooling. Switching to a low-performance mode.	Logged when the controller temperature threshold is exceeded.
0x021C	Information	Configuration automatically created by %s	Logged when the configuration is automatically created.
0x021D	Critical	Initialization aborted on VD 0x%x due to a controller reset.	Logged when the initialization is aborted due to a controller reset.
0x021E	Critical	MegaRAID Solution is forced to shut down due to the maximum temperature threshold exception. This error may indicate inadequate system cooling.	Logged when the MegaRAID solution is forced to shut down due to a maximum temperature threshold exception.
0x021F	Information	Locate LED started on PD 0x%02x(e0x%02x/s%d)	Logged when the locate LED has started.
0x0220	Information	Locate LED stopped on PD 0x%02x(e0x%02x/s%d)	Logged when the locate LED has stopped.
0x0221	Information	Patrol read aborted on PD 0x%02x(e0x%02x/s%d) due to conflict with other background operations.	Logged when a patrol read is aborted due to conflict with other background operations.
0x0222	Warning	System reset is required.	Logged when a system reset is required.
0x0223	Information	Auto-configuration parameters changed.	Logged when the auto configuration parameters change.
0x0224	Information	Inserted: PD 0x%02x(e0x%02x/s%d)	Logged when a PD is inserted.
0x0225	Information	Removed: PD 0x%02x(e0x%02x/s%d)	Logged when a PD is removed.
0x0226	Information	VD 0x%x cannot be secured in the future due to a non-SED drive.	Logged when a specific LD cannot be secured in the future due to a non-SED drive.
0x0227	Information	PD 0x%02x(e0x%02x/s%d) security key unlocked.	Logged when a security key is unlocked.
0x0228	Warning	SAS/SATA mixing is not supported in the enclosure; PD 0x%02x(e0x%02x/s%d) is disabled.	Logged when SAS/SATA mixing is not supported in an enclosure. %s is disabled.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x0229	Information	ASO %s key - %s	Logged when a trial key feature is disabled.
0x022A	Information	VD 0x%x unmap support cannot be enabled.	Logged when an unmap support operation cannot be enabled.
0x022B	Information	VD 0x%x WriteSame unmap support cannot be enabled.	Logged when a WriteSame unmap support operation cannot be enabled.
0x022C	Critical	Rebuild is not possible as the firmware did not find a suitable unmap-capable drive.	Logged when a rebuild operation is not possible because the firmware did not find a suitable unmap-capable drive.
0x022D	Critical	Replace is not possible as the firmware did not find a suitable unmap-capable drive.	Logged when a replacedrive is not possible because the firmware did not find a suitable unmap-capable drive.
0x022E	Warning	The dedicated hot spare PD 0x%02x(e0x%02x/s%d) is not unmap-capable and is no longer useful for one or more arrays.	Logged when a dedicated hot spare is not unmap-capable and no longer useful for one or more arrays.
0x022F	Information	The Snapdump Available Id is %lu	Logged when a Snapdump availability with an ID.
0x0230	Information	The Snapdump %s is deleted.	Logged when a Snapdump is overwritten with an ID.
0x0231	Critical	SAS topology error: %s	Logged when a SAS topology error occurs.
0x0232	Critical	SAS topology error: Multiple ports to the same SAS address.	Logged when the following SAS topology error occurs: Multiple ports to the same SAS address.
0x0233	Critical	SAS topology error: Table to table.	Logged when the following SAS topology error occurs: Table to table.
0x0234	Critical	SAS topology error: Multiple subtractive.	Logged when the following SAS topology error occurs: Multiple subtractive.
0x0235	Critical	SAS topology error: SMP CRC error.	Logged when the following SAS topology error occurs: SMP CRC error.
0x0236	Critical	SAS topology error: SMP function failed.	Logged when the following SAS topology error occurs: SMP function failed.
0x0237	Critical	SAS topology error: SMP timeout.	Logged when the following SAS topology error occurs: SMP timeout.
0x0238	Critical	SAS topology error: Device with an invalid SAS address.	Logged when the following SAS topology error occurs: Device with invalid SAS address.
0x0239	Critical	SAS topology error: Loop detected.	Logged when the following SAS topology error occurs: Loop detected.
0x023A	Information	PD 0x%02x(e0x%02x/s%d) rebuild not possible as SAS/SATA and NVMe mix is not supported in the array.	Logged when a rebuild operation is not possible because SAS/SATA and NVMe mixing is not supported in an array.
0x023B	Information	Replace cannot be started on PD 0x%02x(e0x%02x/s%d) from PD 0x%02x(e0x%02x/s%d), as SAS/SATA and NVMe mix is not supported in the array.	Logged when a replacedrive operation cannot be started because mixing is not supported in an array.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x023C	Information	PD 0x%02x(e0x%02x/s%d) rebuild not possible as the drive is not supported in the array.	Logged when a replacedrive operation cannot be started because the drive is not supported in the array.
0x023D	Information	Replace cannot be started on PD 0x%02x(e0x%02x/s%d) from PD 0x%02x(e0x%02x/s%d), as the drive is not supported in the array.	Logged when a rebuild operation is not possible because the drive is not supported in the array.
0x023E	Information	Auto configuration option is set to - %s	Logged when the auto configuration option is set.
0x023F	Information	The Enclosure PD 0x%02x(ELI-0x%llx/p%d) power supply %d failed.	Logged when an enclosure power supply fails.
0x0240	Information	Failing PD 0x%02x(e0x%02x/s%d) as WR_UNCOR is not supported.	Logged when a failing PD, such as WR_UNCOR is not supported.
0x0241	Information	The property of the Snapdump module is changed	Logged when a Snapdump module property is changed.
0x0242	Warning	DDF configuration clear failed on PD 0x%02x(e0x%02x/s%d)	Logged when a DDF Config clear failed on a PD.
0x0243	Critical	PCIe hot reset failed on PD 0x%02x(e0x%02x/s%d)	Logged when a PCIe hot reset fails.
0x0244	Warning	One or more images in flash do not match the package manifest.	Logged when one or more images in the flash do not match the package manifest.
0x0245	Warning	SAS-wide port %d lost link on PHY %d	Logged when a SAS-wide port lost a link on a PHY.
0x0246	Information	SAS-wide port %d link restored on PHY %d	Logged when a SAS-wide port link is restored on a PHY.
0x0247	Information	The energy pack is bad or missing. Incomplete writes during power loss may cause data integrity issues on parity VD 0x%x	Logged when incomplete writes are present during power loss because the energy pack is bad or missing.
0x0248	Information	NVMe recover started on PD 0x%02x(e0x%02x/s%d)	Logged when an NVMe recover is started on a PD.
0x0249	Information	NVMe recover successfully completed on PD 0x%02x(e0x%02x/s%d)	Logged when an NVMe recover successfully completes on a PD.
0x024A	Warning	NVMe recover failed on PD 0x%02x(e0x%02x/s%d)	Logged when an NVMe recovery fails on a PD.
0x024B	Warning	NVMe recover aborted on PD 0x%02x(e0x%02x/s%d)	Logged when an NVMe recovery is aborted on a PD.
0x024C	Warning	PD 0x%02x(e0x%02x/s%d) has 0x%04x bad media events	Logged when a PD has bad media events.
0x024D	Warning	PD 0x%02x(e0x%02x/s%d) has bad perf, %s	Logged when a PD has bad performance.
0x024E	Information	Power state change occurred on PD 0x%02x(e0x%02x/s%d) from %s(%x) to %s(%x)	Logged when a PD power state change occurs.
0x024F	Warning	Power state change failed on PD 0x%02x(e0x%02x/s%d) (from %s(%x) to %s(%x))	Logged when a PD power state change fails.
0x0250	Information	Link speed changed on SAS port %d and PHY %d	Logged when a link speed changed on a SAS port and PHY.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x0251	Information	Write journal throttling enabled.	Logged when write journal throttling is enabled.
0x0252	Information	Write journal throttling disabled.	Logged when write journal throttling is disabled.
0x0253	Information	The dedicated hot spare PD 0x%02x(e0x%02x/s %d) is no longer useful due to the deleted array.	Logged when a dedicated hot spare PD is no longer useful due to a deleted array.
0x0254	Warning	The enclosure PD 0x%02x(ELI-0x%llx/p%d) slot %d is critical.	Logged when an enclosure element indicates a critical condition.
0x0255	Critical	The controller booted to safe mode due to critical errors.	Logged when a controller is booted to safe mode due to critical errors.
0x0256	Warning	A validation error occurred during the firmware update (%s)	Logged when a validation error occurs during a firmware update.
0x0257	Information	An update to a new eFUSE key was successfully completed.	Logged when an update to a new eFUSE key is successfully completed.
0x0258	Warning	A programming error occurred during the firmware update (%s)	Logged when a programming error occurs during a firmware update.
0x0259	Warning	A pinned cache or a write journal was found when attempting firmware activation.	Logged when a pinned cache or write journal was found when attempting a firmware activation.
0x025A	Information	Offline activation is pending (%s)	Logged when an offline activation is starting or pending.
0x025B	Information	A repair of images was completed successfully (%s)	Logged when a repair of images is completes successfully.
0x025C	Warning	Online activation request was converted to an offline activation.	Logged when an online activation request was converted to an offline activation.
0x025D	Warning	Online activation failed during the controller preparation.	Logged when an online activation failed during the controller preparation.
0x025E	Information	Firmware activation completed successfully.	Logged when a firmware activation completes successfully.
0x025F	Information	Escrow keyId %s cleaned up after timeout.	Logged when an escrow key cleaned up after a timeout.
0x0260	Warning	Pending rekey operation discarded due to unavailability of an EKM security key.	Logged when a pending rekey operation is discarded due to the unavailability of an EKM security key.
0x0261	Information	The controller successfully received the EKM key.	Logged when the controller successfully received the EKM key.
0x0262	Information	The time on the controller has been set.	Logged when the time has been set on the controller.
0x0263	Critical	An error has been detected in the L2 cache.	Logged when an A15 L2 cache error is detected.
0x0264	Information	The driver on the host must be updated.	Logged when the host driver must be updated to support the controller.
0x0265	Warning	A transient error has occurred while accessing the flash.	Logged when a transient error with the flash is detected.
0x0266	Critical	The flash device has failed.	Logged when a fatal error with the flash is detected.

Number	Severity Level	Event Text	Generic conditions when each event occurs
0x0267	Critical	A multi-bit ECC error has been detected in the OCM. This error is fatal.	Logged when a fatal multi-bit ECC error is detected in the OCM.
0x0268	Warning	A single-bit ECC error has been detected in the OCM. This error is non-fatal.	Logged when a non-fatal single-bit ECC error is detected in the OCM.
0x0269	Warning	The number of single-bit ECC errors in the OCM has crossed the warning threshold.	Logged when the number of single-bit ECC errors in the OCM has crossed the warning threshold.
0x026A	Critical	The number of single-bit ECC errors in the OCM has crossed the critical threshold.	Logged when the number of single-bit ECC errors in the OCM has crossed the critical threshold.
0x026B	Warning	The detection of single-bit ECC errors in the OCM has been disabled.	Logged when the detection of single-bit ECC errors in the OCM is disabled.

Glossary

This glossary defines the terms that are used in this document.

A

Absolute state of charge	Predicted remaining battery capacity that is expressed as a percentage of Design Capacity. 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 that is 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	The 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 if there is 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 the subsequent access to the same data. When data is read from or written to the main memory, a copy is also saved in the 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 the cache memory. If it is already in the cache memory (a cache hit), it is read from the 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 the main memory and saved in the 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, and 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	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). <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 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 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 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 the data in a transaction.
Design capacity	Designed charge capacity of the supercapacitors, which is 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, which are 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 that are attached to a RAID controller on which one or more virtual drives can be created. All virtual drives in the drive group use 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 StorCLI2 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 StorCLI2 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 StorCLI2 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 StorCLI2 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 StorCLI2 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 StorCLI2 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 StorCLI2 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 StorCLI2 commands, Offline is displayed as OfLn.
Drive state drive subsystem		<p>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.</p>
Drive type		<p>A drive property indicating the characteristics of the drive.</p>
	E	
EKM		<p>External Key Management</p>
Estimated time to recharge		<p>Estimated time necessary to complete recharge of the supercapacitors at the current charge rate.</p>
Expected margin of error		<p>Indicates how accurate the reported supercapacitors capacity is in terms of percentage.</p>
	F	
Fast initialization		<p>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.</p>
Fault tolerance		<p>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 provide 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.</p>

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 a drive or from a network and then passes control to the operating system.
Foreign configuration	A RAID configuration that 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 that are controlled by the controller. When a drive fails, LSI Storage Authority software automatically uses a hot spare to replace it and then rebuilds the data from the failed drive to the hot spare. Hot spares can be used in RAID 1, 5, 6, 10, 50, and 60 storage configurations.
I	
Initialization	The process of making a redundant virtual drive consistent. Foreground initialization writes zeros to the data fields, erasing all existing data. 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.
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 the cache memory. In Direct I/O mode, reads are not buffered in the cache memory. Data is transferred to cache and the host concurrently. If the same data block is read again, it comes from the cache memory. The IO Policy applies to reads on a specific virtual drive and 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 saves 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 that are 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 the Max Error to 100 percent on a full reset. The gas gauge sets the 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 the Max Error to 8 percent unless the Max Error was already below 8 percent. In this case, the Max Error does not change. The gas gauge increments the 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 the 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 if there is a drive failure.
NVMe	Acronym for nonvolatile memory express. NVMe is a logical device interface specification for accessing NVM storage media that is attached by 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 that are 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 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, 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 the 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. The redundant data can also be a complete mirrored copy of the data that is 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 uses 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, 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 that is consumed by a stripe is the same on each drive. Striping by itself does not provide data redundancy. Striping in combination with parity provides 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, which is measured in Celsius.
	U	
Uncorrectable error count		A controller property that lists the number of uncorrectable errors that are detected on drives that are connected to the controller. If the error count reaches a certain level, a drive is 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 that is 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 if there is 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 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	<i>See 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 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 1.3, April 28, 2023

- Updated [RAID 1 Drive Groups](#).
- Updated [Drive Autoconfiguration](#).
- Updated [Enable Security](#).
- Added [Managing Unconfigured Secure Drives](#).
- Updated [Viewing and Importing a Foreign Configuration](#).
- Updated [Managing SAS Storage Link Speed](#).
- Updated [System Show Commands](#).
- Added [Controller Help Commands](#).
- Added [Controller Show Commands](#).
- Updated [Show Controller Properties Commands](#).
- Updated [Set Controller Properties Commands](#).
- Updated [Patrol Read](#).
- Updated [Consistency Check](#).
- Updated [Controller Security Commands](#).
- Updated [Retrieving Snapdump Data Commands](#).
- Added [Drive Performance Monitoring Commands](#).
- Added [SPDM Commands](#).
- Updated [Drive Firmware Download Commands](#).
- Updated [Drive Secure Erase Commands](#).
- Updated [Set Drive State Commands](#).
- Updated [Rebuild Drives Commands](#).
- Updated [Hot Spare Drive Commands](#).
- Added [Replacedrive Commands](#).
- Added [Spinup Drive Commands](#).
- Updated [Add Virtual Drives Commands](#).
- Updated [Change Virtual Drive Properties Commands](#).
- Updated [Background Initialization Commands](#).
- Updated [Virtual Drive Expansion Commands](#).
- Updated [Foreign Configuration Commands](#).
- Updated [Drive Group Commands](#).
- Updated [Controller Power Savings Commands](#).
- Updated [Enclosure Commands](#).
- Added [Controller Phy Commands](#).
- Added [Energy Pack Commands](#).
- Added [Logging Commands](#).
- Updated [Automated Physical Drive Configurations](#).
- Updated [StorCLI to StorCLI2 Command Conversion](#).
- Removed Supported Commands on Initiator-Target Controllers.
- Updated [Event Messages](#).
- Minor rewrites for clarity and consistency.

Version 1.2, January 26, 2023

- Updated [Set Drive State Commands](#).
- Updated [Drive Firmware Download Commands](#).
- Updated [Displaying the StorCLI2 Utility Help](#).
- Minor rewrites for clarity and consistency.

Version 1.1, October 11, 2022

- Updated [Tri-Mode Technology](#).
- Updated [Serial ATA III Features](#).
- Updated [Background Initialization](#).
- Updated [Summary of RAID Levels](#).
- Updated [RAID 6 Drive Groups](#).
- Updated [RAID 10 Drive Groups](#).
- Updated [RAID 50 Drive Groups](#).
- Updated [RAID 60 Drive Groups](#).
- Updated [SafeStore Disk Encryption](#).
- Updated [Workflow](#).
- Updated [MegaRAID ADVANCED SOFTWARE OPTIONS](#).
- Updated [Manually Creating a Virtual Drive](#).
- Updated [Viewing Advanced Controller Management Options](#).
- Updated [Viewing Advanced Controller Properties](#).
- Updated [Managing MegaRAID Advanced Software Options](#).
- Updated [Displaying the Controller Personality](#) .
- Updated [Changing Task Rates](#).
- Added [Perform Cryptographic Erase on Drives](#).
- Updated [Viewing and Managing Virtual Drive Properties and Options](#).
- Updated [Securely Erasing a Drive](#).
- Added [Logical Unit/Namespace Information](#).
- Updated [Supported Operating Systems](#).
- Updated [StorCLI2 Tool Command Syntax](#).
- Updated [Show and Set Controller Properties Commands](#).
- Updated [Patrol Read](#).
- Updated [Consistency Check](#).
- Updated [Premium Feature Key Commands](#).
- Updated [Controller Security Commands](#).
- Updated [Retrieving Snapdump Data Commands](#).
- Updated [Rebuild Drives Commands](#).
- Removed [Drive Copyback Commands](#).
- Updated [Hot Spare Drive Commands](#).
- Updated [Add Virtual Drives Commands](#).
- Updated [Delete Virtual Drives Commands](#).
- Updated [Change Virtual Drive Properties Commands](#).
- Updated [Virtual Drive Erase Commands](#).
- Updated [Background Initialization Commands](#).
- Updated [Virtual Drive Power Settings Commands](#).
- Updated [StorCLI to StorCLI2 Command Conversion](#).
- Minor rewrites for clarity and consistency.

Version 1.0, April 20, 2022

- Updated [Blocking Boot Events](#).
- Updated [ACTIONS](#).
- Updated [BACKGROUND OPERATIONS](#).
- Updated [Manually Creating a Virtual Drive](#).
- Updated [Viewing and Importing a Foreign Configuration](#).
- Updated [Clearing a Foreign Configuration](#).
- Added [HII Popup Error Protocol](#).
- Minor rewrites for clarity and consistency.

Advance, Version 0.2, November 30, 2021

- Updated [Overview](#).
- Updated [Broadcom 9600 Series Features](#).
- Updated [Configuration Scenarios](#).
- Updated [Introduction to RAID](#).
- Updated [Background Initialization](#).
- Updated [Drive States](#).
- Updated [Workflow](#).
- Added [Blocking Boot Events](#).
- Updated [HII Dashboard View](#).
- Updated [Managing Configurations](#).
- Updated [Managing Controllers](#).
- Updated [Managing Virtual Drives](#).
- Updated [Managing Devices](#).
- Updated [Managing Energy Packs](#).
- Updated [Show and Set Controller Properties Commands](#).
- Added [Controller Replacedrive Commands](#).
- Updated [Prepare to Remove Drives Commands](#).
- Added [NVMe Drive Commands](#).
- Updated [Supported Commands on Initiator-Target Controllers](#).
- Added [System Commands](#).
- Added [Download Commands](#).
- Added [Drive Commands](#).
- Added [Get Commands](#).
- Added [Other Commands](#).
- Added [SAS Address Assignment Rule](#).
- Added [Event Messages](#).
- Added [StorCLI to StorCLI2 Command Conversion](#).
- Updated graphics.
- Minor rewrites for clarity and consistency.

Advance, Version 0.1, April 16, 2021

Initial document release.

Documentation Legal Notice

Copyright © 2021-2023 Broadcom. All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. For more information, go to www.broadcom.com. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

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.

