LSI Storage Authority Software
User Guide

Note

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

First Edition (June 2017)

Contents

Chapter 1: How is This Guide Organized

The *LSI Storage Authority Software User Guide* contains the following sections:

Section	Description
LSI Storage Authority Overview	Provides an overview of the LSI Storage Authority Software including monitoring and maintaining storage devices and the required hardware and software to run the application.
LSI Storage Authority Feature Comparison Matrix	Outlines the LSI Storage Authority feature differences for MegaRAID®, iMegaRAID, , and Initiator-Target (IT) controllers.
Types of Installation	Provides information on LSI Storage Authority Installers and steps to install and uninstall the LSI Storage Authority software.
Performing Initial Setup	Provides certain initial setups that you need to perform.
Server Dashboard	Provides information about the Server Dashboard.
Controller Dashboard	Provides information about the Controller Dashboard.
Configuration	Provides information on how to create and modify storage configurations on systems with Lenovo controllers.
Background Operations Support	Provides information on Background Operations Support, such as Pause, Resume, Abort, and so on.
Managing Controllers	Provides information on how to monitor the activity of all the controllers present in the system and the devices attached to them.
MegaRAID Advanced Software	Provides information on certain premium features that the LSI Storage Authority software supports on ThinkSystem 12Gb/s RAID controllers.
Managing Drive Groups	Provides information on how to monitor the status of the drive groups and spanned drive groups.
Managing Virtual Drives	Provides information on how to perform various operations on the virtual drives.
Managing Physical Drives	Provides information on how to manage physical drives that are connected to the controller.
Managing Hardware Components	Provides information on managing hardware components.
Viewing Event Logs	Provides information on how to view event logs.
Customizing the Theme of the LSI Storage Authority Software	Provides information on customizing the theme of the LSI Storage Authority software, such as adding your company logo or change the default colors.

Chapter 2: LSI Storage Authority Overview

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

- **Monitoring and Configuring**: LSA enables you to monitor the controllers and configure the drives on the controller. It displays the status of the controller cards, virtual drives, and physical drives on the controller. The device status icons are displayed on their respective pages to notify you in case of drive failures and other events that require your immediate attention. Real-time email notifications on the status of the server are sent based on your alert settings. The system errors and events are recorded and displayed in an event log file. Additionally, you can also import or clear foreign configurations.

  NOTE MegaRAID Storage Manager is not supported on the ThinkSystem controllers. Simple Network Management Protocol (SNMP) Management Information Base (MIB) is not used in LSA anymore. So MegaRAID Storage Manager SNMP MIB will be ended with the ThinkServer and System X controllers.

- **Maintaining**: Using LSA, you can perform system maintenance tasks, such as updating the controller firmware.
- **Troubleshooting**: LSA displays information related to drive failures, device failures, and so on. It also provides recommendations and displays contextual links, helping you to easily locate the drives/devices that have issues and troubleshoot them. You can also download a complete report of the all the devices and their configurations, properties, and settings and send it to the Support Teams for further troubleshooting.
- In addition to SAS and SATA drives, starting with LSA 2.4 and later, non-volatile memory express (NVMe) drives are also supported.

2.1 Support Matrix

The following table provides the support requirements for the LSI Storage Authority software.

Table 1 Hardware and Software Support Matrix

Operating System	Version/Flavors
Supported Controllers	■ ThinkSystem 12Gb/s SAS RAID controllers ■ Integrated ThinkSystem (iMR) 12Gb/s SAS RAID controllers ■ ThinkSystem 6Gb/s SAS RAID controllers ■ ThinkSystem host bus adapters (HBA)
Supported operating systems	To check for the latest list of supported operating systems and to download the device drivers for those operating systems, go to http://support.lenovo.com
Supported web browsers	■ Windows Internet Explorer 9.0 and later ■ Mozilla Firefox version 9.0 and later ■ Google Chrome version 16.0 and later
Supported networks	■ Internet Protocol versions 4 and 6 ■ Network Address Translation ■ Lightweight Directory Access Protocol (LDAP) ■ Domain ■ HTTP, HTTPS

2.2 Technical Support

For assistance with running or configuring the LSI Storage Authority Software, contact a Lenovo Technical Support representative. Click the following link to send an email or call a Technical Support representative, or submit a new service request and view its status.

Contact support:

http://support.lenovo.com

Chapter 3: LSI Storage Authority Feature Comparison Matrix

The following tables outline the LSI Storage Authority feature differences for MegaRAID, iMegaRAID, , and Initiator-Target controllers with respect to software features and firmware features. The tables also indicate the supported and unsupported features for a specific controller.

Some of the features might not be supported on all the controllers. Refer to these feature comparison matrices for information on the features that are supported on your controller.

Table 2 Firmware Feature Comparison Matrix

Feature Name	MegaRAID	iMegaRAID	Initiator-Target
RAID Level	RAID 0, RAID 1,RAID 5, RAID 6, RAID 00, RAID 10,RAID 50, RAID 60,RAID 1E, andSpanned RAID 1E (PRL-11)	RAID 0, RAID 1,RAID 5, RAID 10, RAID 50, andRAID 1E	No
Maximum Physical Drives	64	58	1024
Maximum ConfigurablePhysical Drives	240	iMegaRAID 32.Rest of the drives can be used as JBODs.	1024
Maximum Spans	8	8	No
Maximum Virtual Drives	64	32	No
Dimmer Switch	DS-I and DS-II	No	No
Maximum Media Errors	256	102	No
Drive-mixing Support	Yes	Yes	No
Strip Size Support	64 KB, 128 KB, 256 KB, 512 KB, and 1024 KB	64 KB	No
Maximum VDs per Drive Group	64	16	No
Multipath	Yes	Yes	Yes
Controller Reset Support	Yes	Yes	No

Table 3 Software Feature Comparison Matrix

Feature Name	MR	iMR	IT
LDAP Authentication	Yes	Yes	Yes
Server Discovery and ManagingServers	Yes	Yes	Yes
Server Dashboard	Yes	Yes	Yes
Controller Dashboard	Yes	Yes	Yes
Simple Configuration	Yes	Yes	No
Advance Configuration	Yes	Yes	No
CacheCade - SSD CachingConfiguration	Yes	No	No
Foreign Configuration (Import/Clear)	Yes	Yes	No
Clear Configuration	Yes	Yes	No
Update Firmware	Yes	Yes	Yes
Online Firmware Update	Yes	Yes	No
Controller Operations			
Setting Consistency Check Properties	Yes	Yes	No
Scheduling Consistency Check	Yes	Yes	No
Setting Patrol Read Properties	Yes	Yes	No

Table 3 Software Feature Comparison Matrix (Continued)

Feature Name	MR	iMR	IT
Starting Patrol Read	Yes	Yes	No
Stopping Patrol Read	Yes	Yes	No
Managing Link Speed	Yes	Yes	No
Setting Adjustable Task Rates	Yes	Yes	No
Enable/Disable Alarm	Yes	Yes	No
Silence Alarm	Yes	Yes	No
Manage Power-save Settings	Yes	Yes	No
Enable and Disable SSD Guard	Yes	Yes	No
Enable and Disable Security	Yes	Yes	No
Change Drive Security	Yes	Yes	No
Discarding Preserved Cache	Yes	No[a]	No
Downloading TTY Log	Yes	Yes	No
Background Operations	Yes	Yes	No
Advanced Software Features			
Fast Path	Yes	Yes	No
CacheCade SSD	Yes	Yes[b]	No
CacheCade Pro	Yes	Yes[c]	No
SafeStore	Yes	Yes	No
RAID 5 and RAID 6	Yes	Yes[d]	No
Modify Drive Group	Yes	Yes	No
Secure Using FDE	Yes	Yes	No
Disable Data Protection	Yes	Yes	No
Hide/Unhide	Yes	Yes	No
Virtual Drive Operations			
Virtual Drive Settings/ModifyingVirtual Drive Properties	Yes	Yes[e]	No
Start and Stop Locating a Virtual Drive	Yes	Yes	No
Erasing a Virtual Drive	Yes	Yes	No
Initializing a Virtual Drive	Yes	Yes	No
Starting Consistency Check on aVirtual Drive	Yes	Yes	No
Expanding the Online Capacity of a Virtual Drive	Yes	Yes	No
Deleting a Virtual Drive	Yes	Yes	No
Hide/Unhide	Yes	Yes	No
Physical Drive Operations			
Assign Global Hot Spare	Yes	Yes	No
Remove Global Hot Spare	Yes	Yes	No
Assign Dedicated Hot Spare	Yes	Yes	No
Remove Dedicated Hot Spare	Yes	Yes	No
Start and Stop Locating Drive	Yes	Yes	Yes
Making a Drive online and Offline	Yes	Yes	No
Replacing a Drive	Yes	Yes	No

Table 3 Software Feature Comparison Matrix (Continued)

Feature Name	MR	iMR	IT
Rebuilding a Drive	Yes	Yes	No
Prepare for Removal	Yes	Yes	No
Erasing a Drive	Yes	Yes	No
Instant Secure Erase	Yes	Yes	No
Converting Unconfigured Bad Drive to Unconfigured Good Drive	Yes	Yes	No
Make Unconfigured Good Drive	Yes	Yes	No
Make JBOD/ Delete JBOD	Yes	Yes	No
Energy Pack Operations			
Learn Cycle	Yes	No	No
Event Logs			
Viewing Event Logs	Yes	Yes	Yes

a. Energy Pack is not supported
b. Not supported.
c. Not supported
d. RAID 6 not supported.
e. Cached IO and Write Back options are not supported.

Chapter 4: LSI Storage Authority Preinstallation, Postinstallation, and Upgrade Requirements

This section describes the tasks that you must complete before you install/after the installation is complete/while upgrading the LSI Storage Authority Software.

4.1 Preinstallation Requirements

The following are the tasks that you must complete before you install LSI Storage Authority Software:

- If you want to access any other server with a particular gateway, ensure that you have installed the same version of LSA on both the systems.

4.1.1 OpenSLP

SLP or Service Location Protocol is a process by which nodes on a network and select services/resources can be discovered. By nature, this process is dynamic and requires little or no static configuration. OpenSLP is just an open source implementation of SLP, suitable for commercial and non-commercial applications.

From an LSA perspective, OpenSLP requires multicasting functionality to discover the servers that are connected over a subnet. For the **Remote Server Discovery** page to display all the registered servers, ensure that the servers are connected to a network configuration that supports multicasting. For more information, refer to Network Settings.

If OpenSLP is Not Installed

No action is required. OpenSLP is bundled with LSA 2.2 and later. While installing LSA, ensure that you select the option to install OpenSLP, and LSA seamlessly installs the required version of OpenSLP.

If OpenSLP is Already Installed Before Installing LSA

If an instance of OpenSLP is already installed, LSA 2.2 and later packages skip installing OpenSLP.

An Instance of OpenSLP was Already Installed, But LSA is Unable to Display All the Registered Servers From the Remote Discovery page

- Restart the SLPD Services and LSA Service if LSA is not discovering the Servers from the Remote Discovery Page.

4.2 Upgrade Requirements

The following are the tasks that you must complete while upgrading the LSI Storage Authority Software.

4.2.1 Browser Cache

If you are upgrading from a previous version of LSA, clear the browser cache on the client on which you are using LSA.

Chapter 5: Types of Installation

The following are the different types of LSI Storage Authority installers:

- Gateway
- StandAlone
- DirectAgent
- Lightweight Monitor (LWM)

The following table provides more information on each of these installers and their associated advantages.

Table 4 Types of Installers and Their Advantages

Feature	Gateway Installer	StandAlone Installer	DirectAgent Installer	Lightweight Monitor
Permits discovery of other servers that run the LSI Storage Authority software	Yes	No	No	No
Permits self-registration using OpenSLP and has interface for server discovery detection from the network	Yes	Yes **NOTE** No interface for server discovery	No	No
Allows to manage the servers from the list of discovered servers through the user interface (UI).	Yes	No	No	No
Provides capability to configure LDAP information	No	Yes	No	No
Provides server monitoring capabilities and helps to monitor the health of theserver and alerts the end-user of any issues with event logs and email notifications.	Yes	No	No	Yes

5.1 Gateway Installer

The Gateway installer has the following components:

- A back-end with local agent and remote agent management capabilities.
- A monitor with remote monitoring capability.
- A client with remote and managed server capabilities.

The Gateway installer has the following features:

- Permits discovery of other servers that run the LSI Storage Authority software.
- Permits self registration using OpenSLP and has interface for server discovery detection from the network.
- Allows you to manage the servers from the list of discovered servers through the user interface (UI).

5.2 StandAlone Installer

The standAlone installer has the following components:

- A back-end with local agent (without remote agent management capability).
- A monitor (without remote monitoring capability).
- A client (without remote and managed server capabilities).

The standAlone installer has the following features and limitations:

- Does not permit the discovery of other hosts that are running the LSI Storage Authority software.
- Permits self registration of the current host using OpenSLP, but will not have any interface for server discovery detection from the network.
- Provides capability to configure LDAP information.
- Does not permit to add managed servers through the user interface (UI).

5.3 DirectAgent Installer

The following are the types of DirectAgent installations:

- Indirect agent (MegaRAID SMI-S provider)
- DirectAgent

The DirectAgent installer has the following components:

- A back-end with local agent and a monitor component.
- A thin agent, which supports discovery (using SLP), authentication, and DCMD tunneling.

The indirect agent installer has the following components:

- OpenSLP
- SMI-S

5.4 Indirect Agent Installer

The Indirect Agent installer has the following components:

- A back-end with local agent and a monitor component.
- A thin agent, which supports discovery (using SLP), authentication, and DCMD tunneling.
- OpenSLP.
- MegaRAID SMI-S provider.

Chapter 6: Installing the LSI Storage Authority Software on the Microsoft Windows Operating System

Perform the following steps to install the LSI Storage Authority software.

1. Run the LSI Storage Authority `setup.exe` file.
   The **InstallShield Wizard** dialog appears.

Figure 1 InstallShield Wizard Dialog

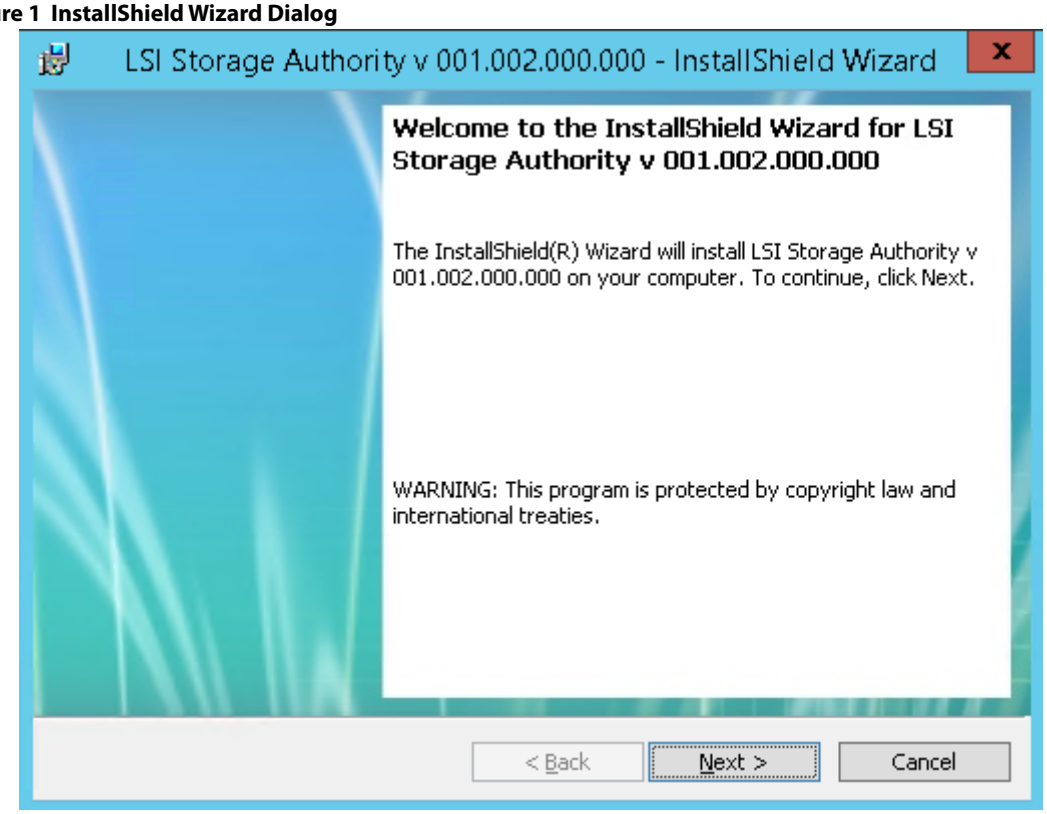

2. Click **Next**.
   The **License Agreement** dialog appears.
3. Read the agreement and choose the **I accept the terms in the license agreement** radio button, and click **Next**.
   The **Customer Information** dialog appears.
4. Enter your user name and the organization name, and click **Next**.
   The **Port Configuration Settings** dialog appears.

Figure 2 Port Configuration Settings Dialog

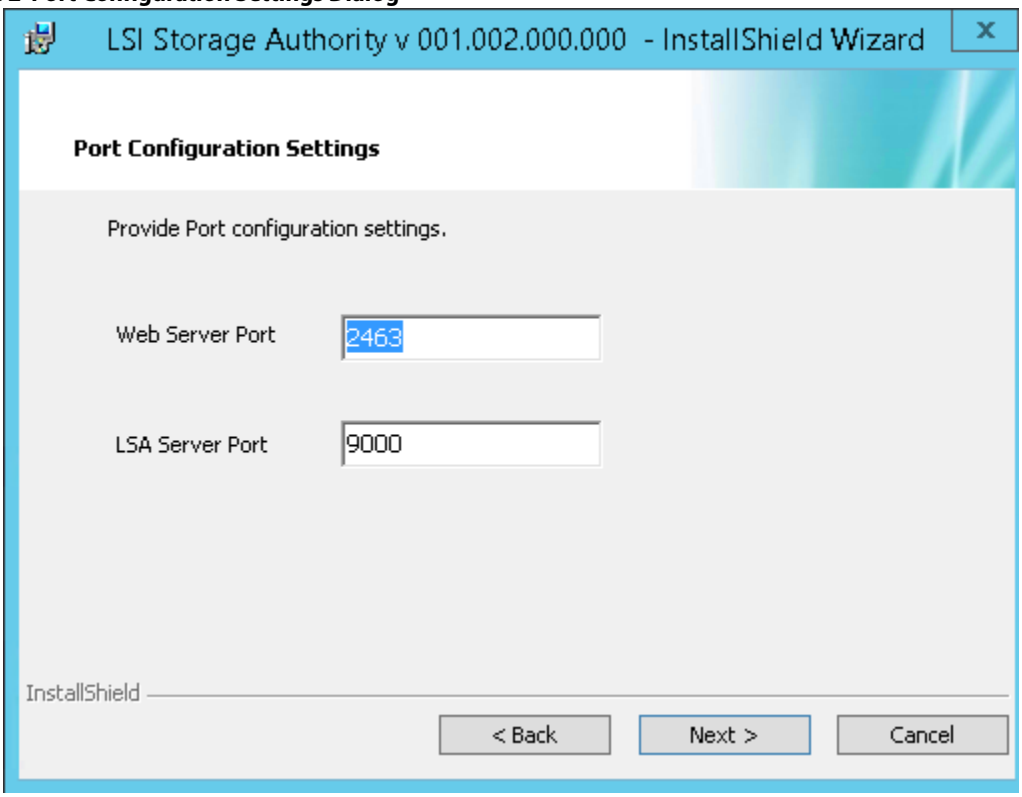

By default, LSA communicates on **Web Server Port 2463** and **LSA Server Port 9000**. Ensure that these ports are available to be used by LSA. Depending on your environment, if these ports are not available, specify the port details here. You can edit this port details after installation also. See Changing the LSI Storage Authority Application Port Number and Changing the nginx Web Server Port Number.

5. Click **Next** to proceed.

   The **Destination Folder** dialog appears with the default file path.

Figure 3 Destination Folder Dialog

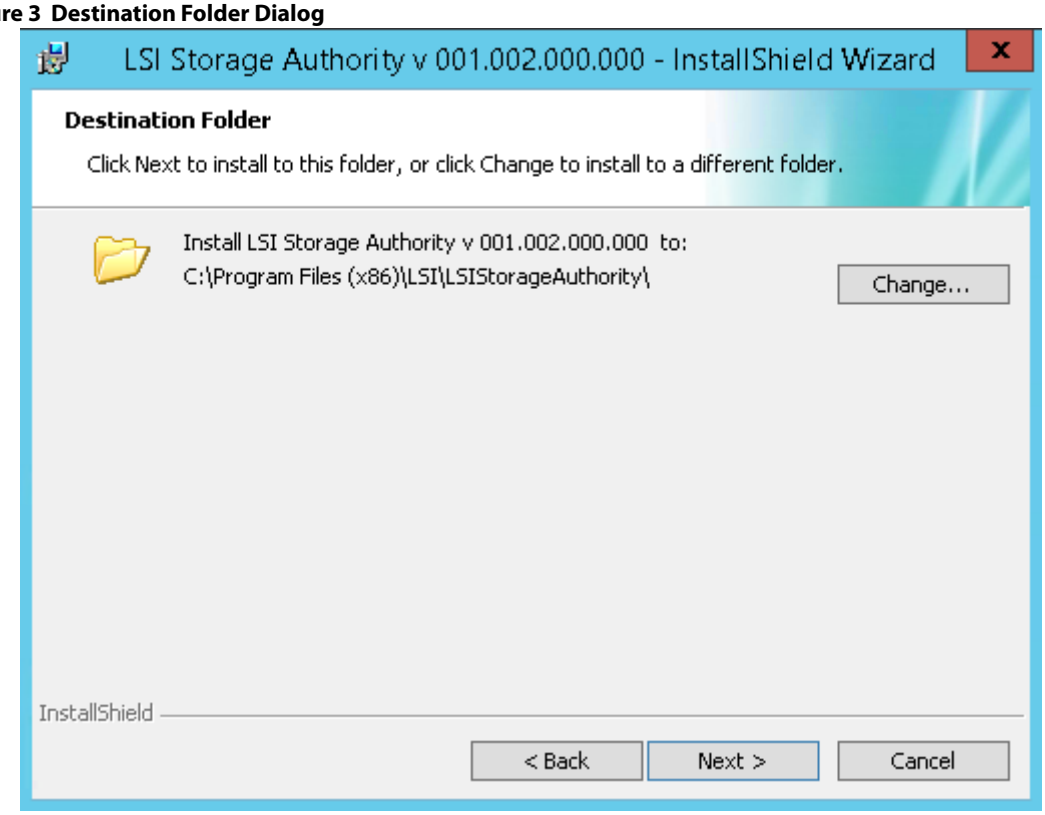

6. (Optional) Click **Change** to select a different destination folder for the installation files.
7. Click **Next**.

   The **Configure Range of Events to Generate Alert Notifications** dialog appears. You can configure alert notifications to get early notification of application or service issues/problem occurrences.

Figure 4 Configure Range of Events to Generate Alert Notifications

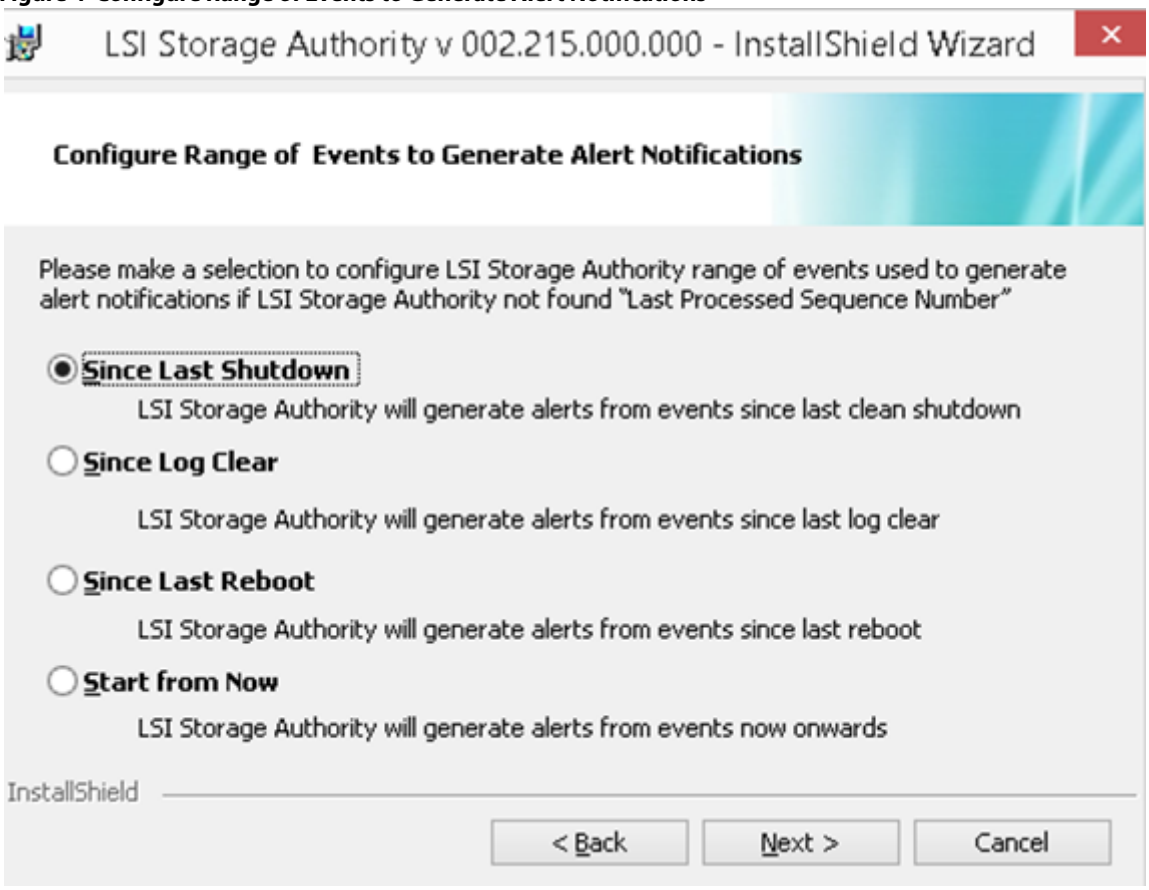

The following configuration options are available:

- **Since Last Shutdown**: Select this option to retrieve events from the last clean shutdown. By default, you can only retrieve the last 30 events. If there are any progress events as part of the last 30 events, then those progress events will not be part of the event history. If the sequence numbers are less than the last log that was cleared (*Since Log Clear* option), then LSA always retrieves events from the *Since Log Clear* option.
- **Since Log Clear**: Select this option to retrieve events from the last log that was cleared. By default, you can only retrieve the last 30 events. If there are any progress events as part of the last 30 events, then those progress events will not be part of the event history.
- **Since Last Reboot**: Select this option to retrieve events from the last time the system was restarted. By default, you can only retrieve the last 30 events. If there are any progress events as part of the last 30 events, then those progress events will not be part of the event history. If the sequence numbers are less than the last log that was cleared (*Since Log Clear* option), then LSA always retrieves events from the *Since Log Clear* option.
- **Start From Now**: Select this option to retrieve events from now.

You can also change these configuration options as per your requirement at any point in time by editing the `lsa.conf` file in the `LSI Storage Authority/conf` directory and choosing the required parameter. For example, if you have selected **Since Last Shutdown** as a configuration option to retrieve events during the time of installation and you want to change it to **Since Last Reboot**, through the `lsa.conf` file, go to `# Retrieve range of events used to generate alert notification, if LSA not found LastProcessedSeqNum` section in the lsa.conf file, change the `retrieve_range_of_events_since =` to 2 (`retrieve_range_of_events_since = 2`).

NOTE You must restart the LSI Storage Authority service for the configuration changes to take effect.

8. Click **Next**.
   The **Setup Type** dialog appears.

Figure 5 Setup Type Dialog

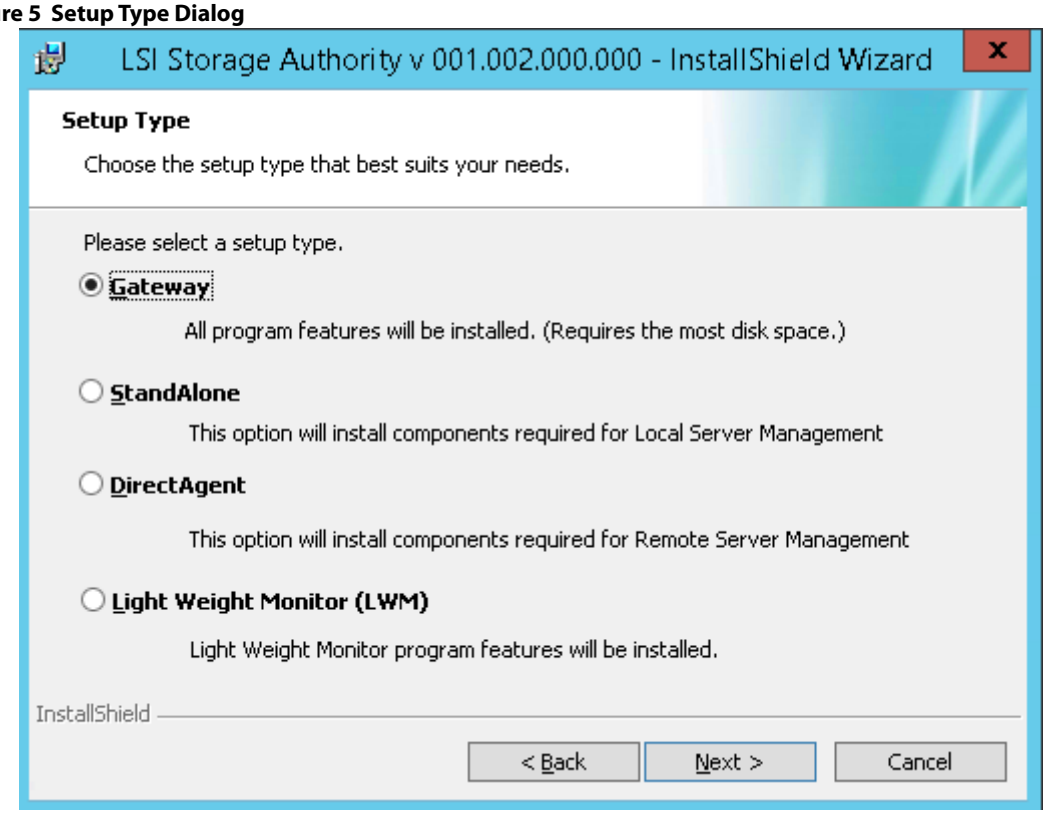

9. Select a setup type that suits your needs. The following options are available:
   For more information on each of these installers and their associated advantages, refer to Table 5, Types of InstallationTable 1, Types of Installers and Their Advantages.
   - **Gateway**
   - **StandAlone**
   - **DirectAgent**
   - **Lightweight Monitor (LWM)**
10. Click **Next**. The **Ready to Install the Program** windows appears. Click **Next**.
    Depending on the setup type you have selected, the **InstallShield Wizard Completed** dialog appears.
11. (optional) Select the **Show the Windows Installer log** checkbox to view the windows installer log file.
    The log file (`LSA_install.log`) is created in the same folder from where the `setup.exe` is installed.
12. Turn off the Windows Firewall.
    The Windows Firewall settings are located under **Control Panel > Windows Firewall**.
13. Click **Finish**.
    The **LSI Storage Authority** homepage appears.

Figure 6 LSI Storage Authority Homepage

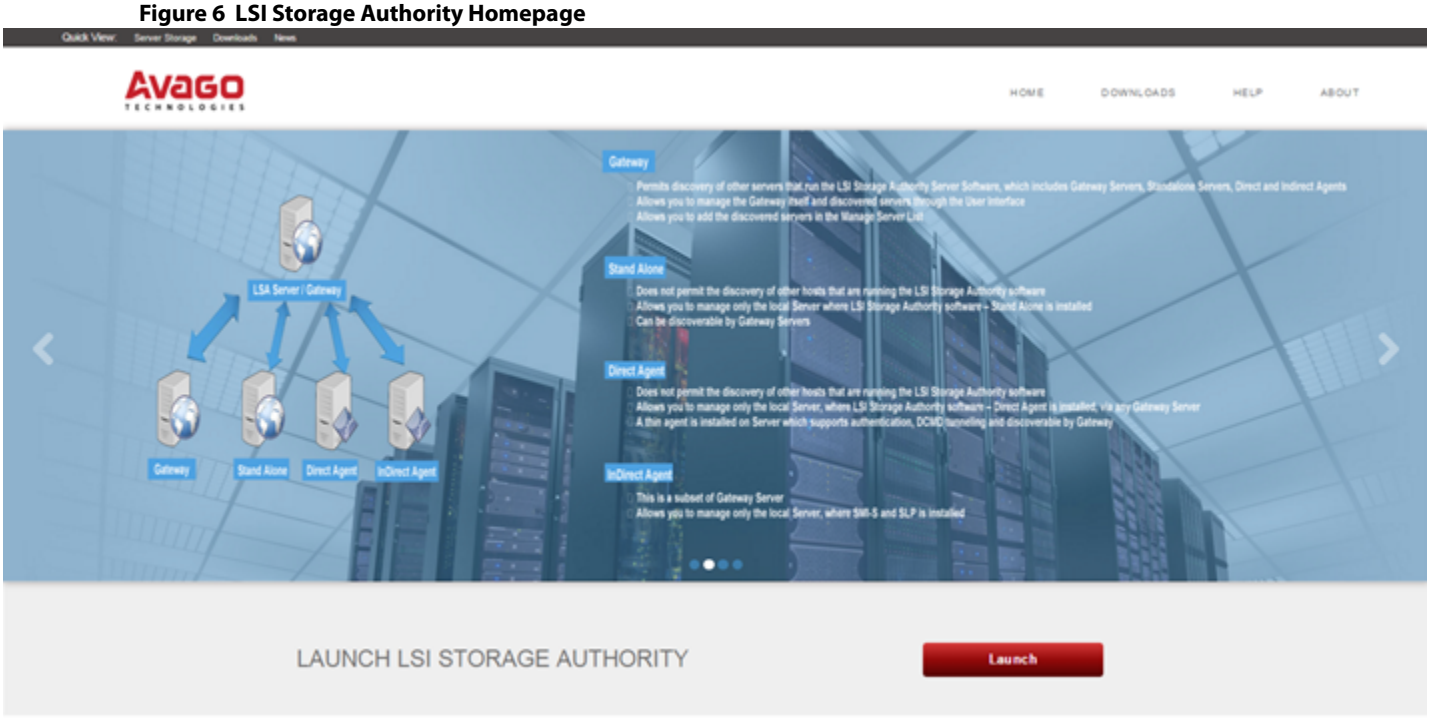

14. Click **Launch LSI Storage Authority**.

    The **Remote Server Discovery** page appears.

Figure 7 LSI Storage Authority Remote Server Discovery Page

On the Remote Server Discovery Page, you can:

- View a list of managed servers with their health status.
- View all the registered servers in the network with their associated server information.
- Manually refresh the list of managed servers.
- Add and remove the managed servers from the list. For more information, see Managing Servers from the Remote Server Discovery Page
- Re-discover servers or go back to the Manage Servers page.
- Using the Check Boxes, you can select only those servers that you require to be displayed.

The Remote Server Discovery page appears for the **Complete** set up option only. The Remote Server Discovery page will not be displayed for a standalone server.

NOTE The OpenSLP requires multicasting functionality to discover the servers that are connected over a subnet. Therefore, for the Remote Discovery page to display all the registered servers. make sure that the servers are connected to the network configuration that supports multicasting.

You can also launch LSA by selecting **Start > All Programs > LSI > LSIStorageAuthority > Launch LSA** or by double-clicking the **Launch LSA** shortcut icon on the desktop.

6.1 Installing LSI Storage Authority Software on the Microsoft Windows Operating System (Non-interactive Mode)

You must log in to the system with root privileges. You can also open the command prompt as root and run the installer through the command line.

Perform the following steps to install the LSI Storage Authority software in a non-interactive mode:

1. From the command line, run the `vcredist_x86.exe /Q` command to install the *Microsoft Visual C++ 2008 Redistributable Package for x86* if it is not already installed.

   The Microsoft Visual C++ 2008 Redistributable Package for x86 (vcredist_x86.exe) is available under the directory `<Package_Dir>\ISSetupPrerequisites\{0BE9572E-8558-404f-B0A5-8C347D145655}\VC Redist 2008 Installation`.
2. Depending on the type of installation required, run the `setup.exe /s /v/qn ADDLOCAL=` command. The types of installation available are:
   - `Gateway–INSTALLATIONCHOICES=0`
   - `StandAlone–INSTALLATIONCHOICES=1`
   - `DirectAgent–INSTALLATIONCHOICES=2`
   - `LightWeightMonitor–INSTALLATIONCHOICES=129`

   Example: If you require the LightWeightMonitor to be installed, you need to run the `setup.exe /s /v/qn ADDLOCAL=LightWeightMonitor INSTALLATIONCHOICES=129 INSTALLDIR=CustomDirecotryLocation` command.

6.2 Uninstalling the LSI Storage Authority Software on the Microsoft Windows Operating System

You can uninstall the LSI Storage Authority Software either through the **Control Panel** or the application shortcut in the **Start** menu.

Uninstalling the LSI Storage Authority Software through the Application Shortcut in the Start Menu

1. Select **Start > All Programs > LSI > LSI Storage Authority > Uninstall LSI Storage Authority**.

Uninstalling the LSI Storage Authority Software through the Control Panel

1. If you are using the Microsoft Windows Server 2008 or the Microsoft Windows Server 2012 operating systems, select **Add/Remove Programs** from the **Control Panel**. If you are using the Microsoft Windows 7 and Microsoft Windows 8 operating systems, select **Programs and Features** from the **Control Panel**.
2. Select the LSI Storage Authority software from the list and click **Uninstall**.

6.3 Uninstalling LSI Storage Authority Software on the Microsoft Windows Operating System (Non-interactive Mode)

You must log in to the system with root privileges. You can also open the command prompt as root and run the installer through the command line.

Perform the following steps to install the LSI Storage Authority software in a non-interactive mode:

1. From the command line, run the `SET /p LSA_ProductCode=<LSA_HOME_PATH\productcode.txt msiexec.exe /x %LSA_ProductCode% /qn` command to uninstall LSA

   Where *LSA_HOME_PATH* is the location where the LSA is installed and *LSA_ProductCode* is a unique product code associated with each LSA installation.

Chapter 7: Installing the LSI Storage Authority Software on the Linux Operating System

The LSI Storage Authority software supports both the Interactive and the Non-interactive modes of Linux installation.

7.1 Installing in the Interactive Mode

You must log in to the system with root privileges. You can also open the command prompt as root and run the installer through the command line.

Perform the following steps to install the LSI Storage Authority software in the interactive mode.

1. Run the `./install.csh` command from the installation disk.
2. Read the license agreements for the software package. If you agree to the terms of the entire license agreements, press Y. Otherwise, press N to exit the installation.
3. Select a setup type that suits your needs. The following options are available:
   - **Gateway** - press 1. Selecting this option installs all the program features.
   - **StandAlone**- press 2. Selecting this option installs components that are required for Local Server Management.
   - **DirectAgent**- press 3. Selecting this option installs components that are required for Remote Server Management.
   - **Lightweight Monitor**- press 4. Selecting this option installs the Lightweight Monitor program features.
4. The **Configure Range of Events to Generate Alert Notifications** dialog appears. You can configure alert notifications to get early notification of application or service issues/problem occurrences.

   The following configuration options are available:
   - **Since Last Shutdown**

     : Select this option to retrieve events from the last clean shutdown.
   - **Since Log Clear**: Select this option to retrieve events from the last log that was cleared.
   - **Since Last Reboot**: Select this option to retrieve events from the last time the system was restarted.
   - **Start From Now**: Select this option to retrieve events from now.

   You can also change these configuration options as per your requirement at any point in time by editing the `lsa.conf` file in the `LSI Storage Authority/conf` directory and choosing the required parameter. For example, if you have selected **Since Last Shutdown** as a configuration option to retrieve events during the time of installation and you want to change it to **Since Last Reboot**, through the `lsa.conf` file, go to `# Retrieve range of events used to generate alert notification, if LSA not found LastProcessedSeqNum` section in the lsa.conf file, change the `retrieve_range_of_events_since =` to 2 (`retrieve_range_of_events_since = 2`).

   You must restart the LSI Storage Authority service for the configuration changes to take effect.
5. Enter the nginx server port number. The port range is from1 to 65535. The default port number is 2463.

6. Enter the LSI Storage Authority Application port numbers. The port range is from 1 to 65535. The default port number is 9000.

   Ensure that the nginx_port number and the LSA_port number are in the between the range, 1-65535 and /or different. If the nginx_port number and the LSA_port number are not specified in the command line, the default values are used.

   By default, LSA communicates on Web Server Port 2463 and LSA Server Port 9000. Ensure that these ports are available to be used by LSA. Depending on your environment, if these ports are not available, specify the port details here. You can edit this port details after installation also.
7. Turn off the Linux Firewall.
8. Extract the contents of the zip file and install the appropriate package on the 32-bit Linux operating systems or the 64-bit Linux operating systems. The `LSA_Linux.zip` file contents are as follows:
   - `x86` - Contains files for 32-bit platforms.
   - `x64` - Contains files for 64-bit platforms.

NOTE Ensure that **Connect automatically** checkbox is selected. This is available under **Network Connections**.

7.2 Installing in the Noninteractive Mode

You must log in to the system with root privileges. You can also open the command prompt as root and run the installer through the command line.

Perform the following steps to install the LSI Storage Authority software in the noninteractive mode.

1. Run the `./install.csh [-options] [nginx_port] [LSA_port]` command from the installation disk.

   Where:

   — `Options`: `c` for complete setup and `m` for monitor setup.
   — `nginx_port`: The nginx server port number.
   — `LSA_port`: The LSI Storage Authority Application port numbers.

   Ensure that the nginx_port number and the LSA_port number are in the between the range, 1-65535 and are different. If the nginx_port number and the LSA_port number are not specified in the command line, the default values (nginx default port 2463 and LSA default) are used.

   Command Usage Examples:

   — Gateway Installation with default ports: `./install.csh -g`
   — StandAlone Installation with default ports: `./install.csh -s`
   — DirectAgent Installation with default ports: `./install.csh -d`
   — Light Weight Monitor Installation with default ports: `./install.csh -l`
   — Gateway installation with different ports: `./install.csh -g 1234 8000`
   — StandAlone installation with different ports: `./install.csh -s 4321 7000`
   — DirectAgent installation with different ports: `./install.csh -d 1254 8800`
   — Light Weight Monitor installation with different ports: `./install.csh -l 4388 9900`
2. Extract the contents of the zip file and install the appropriate package on the 32-bit Linux operating systems or the 64 bit Linux operating systems. The `LSA_Linux.zip` file contents are as follows:

   — `x86` - Contains files for 32-bit platforms.
   — `x64` - Contains files for 64-bit platforms.

7.3 Uninstalling the LSI Storage Authority Software on the Linux Operating System

Perform the following step to uninstall the Linux operating system.

1. Run the `uninstaller.sh script` (`/opt/lsi/LSIStorageAuhority/uninstaller.sh`). Alternatively, you can run the `rpm -e <rpm_name>` command to uninstall the RPM's from the target system.

   Command Usage Example: `rpm -e LSIStorageAuhority-1.00xx.xxxx-xxxx`

Chapter 8: LSA Support on the VMware ESXi Operating System

This section outlines the pre-installation/post-installation requirements needed to support the VMware ESXi Operating System.

LSA cannot be installed directly on the VMware ESXi operating system. Management of LSA is performed through the LSA installed on a Linux/Windows machine in the same subnet.

1. **SMI-S Provider Details**

   For VMware ESXi 4.x, 5.x, and 6.x to work with LSA, depending on your VMware ESXi environment, the following SMI-S Provider must be installed:

Table 5 SMI-S Provider Versions for VMware ESXi

VMWare ESXi Version	SMI-S Provider Version
VMware ESX 4.x	VMware-ESX4-Provider.zip
VMware ESXi 5.4 and below	VMware-ESXi-Provider.zip
VMware ESXi 5.5 and above	VMware-ESXi5.5-Provider.zip

   To deploy the MegaRAID SMI-S provider on an ESXi machine by using the VIB file provided by Lenovo, copy the `VIB` file or the `offline-bundle.zip` file to the ESXi machine. Use the `esxcli` file to load the MegaRAID SMI-S provider, and run the following command:

   ```
   ESXi# esxcli software vib install -v <local_file_path or file:///URL> --force
   ```

   NOTE A reboot is required after installing the SMI-S provider on VMware ESXi environment.

2. **Firewall Details**

   On every reboot, the firewall gets enabled, ensure that you run the following command after every reboot to disable the firewall:

   ```
   esxcli network firewall unload
   ```

   On a VMware ESXi environment, to check whether the firewall is enabled, execute the following command:

   ```
   esxcli network firewall get
   ```

3. **Provider Services**

   Ensure that the Provider Services are up and running before performing the VMware ESXi Discovery from the LSA Client (Windows and Linux).

   Run the following commands to make sure that the provider services are up and running on VMware ESXi:

   ```
   /etc/init.d/slpd status
   /etc/init.d/sfcb-watchdog status
   ```

4. **Configuration Change Details**

   If there is any configuration change, ensure that you perform the below actions:

   — `/etc/init.d/sfcb-watchdog stop`
   — `/etc/init.d/slpd stop`
   — `/etc/init.d/slpd start`
   — `/etc/init.d/sfcb-watchdog start`

5. **Storage Controller**

   Ensure that the storage controller on VMware ESXi has the right configuration (Firmware/Driver) and the storage controller is working as expected before connecting through LSA. The following command helps you verify whether the controller is getting detected or not.

```
enum_instances cim_system lsi/lsimr13
```

6. **Network Communication Details**

   Network communication is a key element for a proper communication between the VMware ESXi CIM provider and the Broadcom management software. Make sure that the network settings are correct by making the following changes:

   — Provide a proper host name and an IP address while performing the initial configurations for the VMware ESXi host. Refer to How to Configure Networking on VMware ESXi? for more information.

   — For networks that do not have DNS configured, the "hosts" file in the machine on which LSA is installed must be edited as follows:

   a. Add an entry to map the VMware host's IP address with the host name. This is for the discovery process to happen correctly. In the absence of this entry, the VMware host would be discovered as 0.0.0.0.

   b. Add an entry to map the actual IP address of the localhost with its hostname (an entry for the loopback address would be present by default in the hosts file and it should not be removed). This is to ensure that the Asynchronous Event Notifications (AENs) are delivered correctly. For example, if 135.24.228.136 is the IP address of your VMware host and 135.24.228.137 is the IP address of your Linux host, the following entries must be added in the hosts file:

```
135.24.228.136 dhcp-135-24-228-136.lsi.com dhcp-135-24-228-136 #VMware
135.24.228.137 dhcp-135-24-228-137.lsi.com dhcp-135-24-228-137 #Linux
```

   NOTE Ensure to disable the firewall in the client system as well (for both Windows and Linux).

Remote management of VMware ESXi is supported only in a Gateway installation of LSA on the following operating systems:

- Microsoft Windows Server
- RHEL
- SuSE Linux

Windows/Linux Steps

The following steps are required to be performed on the Windows/Linux Client:

NOTE Both the Client and the Server should be in the same subnet.

1. Stop or disable the firewall on the client machine.
2. Install the latest LSA Client in a Gateway installation mode.
3. Launch LSA.
4. Ensure that the LSA Service is up and running.
5. Ensure that other LSA server(s) in the network are being discovered.
6. Ensure that VMware IP is being discovered as part of the Remote Server Discovery page.
7. Login with your VMware credentials to monitor and manage the storage controller through the LSA Client.

Configuring the Network on VMware ESXi Environment

- By default, during the VMware ESXi Operating System installation, the IP and host name should be configured appropriately.
- If an already installed VMware ESXi Operating System is moved from one network to the other, and if the host name mapping is not correct, follow the steps mentioned in the below link to configure the network and host name:

  http://www.unixarena.com/2015/05/how-to-configure-the-network-on-vmware-esxi-6-0.html

Multi-subnet Configuration

- When a gateway is part of multiple subnet, and the discovered VMware is part of one of these subnets, you need to configure both the LSA gateway and the VMware gateway under the same subnet.

  You can add an irrelevant IP address to the LSA configuration file (`conf\LSA.conf -> private_ip_range=*`) to avoid registering the wrong IP to the VMware gateway.
- CURL error in CIMOM server results in a blocked AEN to the upper layer (CIMProvider-->LSA). This happens if the servers are in different subnets or if there is any incorrect/incomplete AEN subscriptions. To avoid this error, you are required to have both the client and the server in the same subnet. Any incomplete AEN subscriptions must be removed via CIMClient.

  To view the existing subscriptions, enter:

  `host-ind -s`

  To remove an existing subscription, enter:

  `host-ind -d -k <handler name>`

  For example, `host-ind -d -k dhcp-x.y.z.k.dhcp.company.net_LSA_127.0.0.1`

It is recommended to either restart the `sfcb` service or reboot the server after making any changes to the VMware Server.

8.1 Increasing the Memory Limit of Host Hardware RAID Controller (HHRC)

To increase the memory limit of HHRC, perform the following steps:

1. Edit the `/etc/sfcb/sfcb.cfg` file.
2. Insert `provMemOveride: hhrc=100`into the file.
3. Reboot the system.
4. Verify to ensure that the changes you made have taken effect by running the below command:

   ```
   memstats -r group-stats -u mb -s name:min:max:memsize:memsizepeak | grep -E "hhrc|memSizePeak|--"
   ```

8.2 Configuring the Firewall on Various LSA Installers

The following topics provide information on how to configure firewall on various LSA installers on different operating systems:

- Gateway/StandAlone installer configuration on Windows Operating System.
- DirectAgent installer configuration on VMware Operating System.
- Gateway/StandAlone installer configuration on Linux Operating System.

8.2.1 Configuring the Firewall on Gateway/StandAlone Installer

You can configure the Gateway/StandAlone installer firewall on Windows Operating System.

A firewall profile is a way of grouping settings, such as firewall rules, connection security rules, etc., which are applied to the system depending on where the system is connected.

Windows Operating System has three profiles, **Public**, **Private**, or **Domain**. You need to enable one of these profiles appropriately based on your connection type.

Once the firewall is enabled, inbound settings must have LSA as an exception. By default, inbound settings block all incoming connections unless specified as a rule or as an exception. You must add `nginx.exe\port number` as an exception.

Perform the following steps to add `nginx.exe\port number` as an exception:

1. Go to **Control Panel > Windows Firewall > Allow a program or feature through Windows Firewall**.
2. Select **Allow Another Program**.
3. Browse to the folder where LSIStorageAuthority is installed.
4. From the LSIStorageAuthority installation folder, select, **nginx.exe**.

   Usually **nginx.exe** is installed under the `C:\Program Files (x86)\LSI\LSIStorageAuthority\server`location.

   Nginx.exe is the webserver used by LSA as an interface with remote systems.
5. Click **Add** or select the check box to allow `nginx.exe` as an exception.
6. Click **OK**.

   Alternatively, you can also create a new rule and set the profile type to unblock LSA from the firewall.

8.2.2 Configuring the Firewall on Gateway/StandAlone Installer (Linux)

You can configure the Gateway/StandAlone installer firewall on Linux Operating System. In most Linux systems, by default, all the inbound requests are accepted. You can also check the `iptable` entries to verify the configuration rules.

To verify to configuration rules set in any Linux system, execute the following command:

```
iptable -L
```

If the configuration rules are blocking access to LSA, run the following command to allow input connection to LSA:

```
iptables -I INPUT -p tcp --dport <webserver port> -j ACCEPT
```

For example, if your Nginx port number is 2463 for LSA, you must run `iptables -I INPUT -p tcp -dport 2463 -j ACCEPT` command.

8.2.3 Configuring the Firewall on DirectAgent Installer

LSA uses SMI-S provider to discover and manage its storage controllers on VMware environment. To discover the ESXi servers where SMI-S providers are available, LSA uses the SLP as a discovery mechanism. CMISLP advertises and allows remote systems to discover CIM servers. CIMHttpServer is a CIM server and is required to interact with the SMI-S provider.

You can configure the DirectAgent installer firewall in VMware environment:

To check whether the firewall is enabled, execute the following command:

```
esxcli network firewall get
```

If the firewall is enabled, CIMSLP services should be enabled in the rule set. To check whether CIMSLP services are enabled in the rule set, execute the following command:

```
esxcli network firewall ruleset list
```

If the rule ID of CIMSLP is set to false, CIMSLP gets disabled, and LSA will be unable to discover the CIM service. To enable the CIM service, execute the following command:

```
esxcli network firewall ruleset set --enabled true --ruleset-id=CIMSLP
```

Similarly, CIMHttpServer should be enabled, so that it can interact with the SMI-S provider.

8.3 Collecting LSA Logs (Windows/Linux)

On successfully installing LSA, you can recreate the issue to collect the required log files for Windows/Linux.

Perform the following steps:

1. Stop `LSAService`
2. Browse to `LSA_HOME/Conf` directory and open the `LSA.conf` file.
3. In the `LSA.conf` file, search for `log_level` field.
4. Modify the existing or default value in the log level field to `32`.
5. In the `LSA.conf` file, search for `log_cache_mode` field.
6. Modify the existing or default value in the log_cache_mode field to `0`.
7. Browse to `LSA_HOME/logs` directory and delete the `logs.txt` file.
8. Restart the `LSAService`.
9. Once you see the issue, share the `logs.txt` file from the `LSA_HOME/logs` directory.

8.4 Collecting LSA Logs (VMware)

On successfully installing LSA, you can recreate the issue to collect the required log files for VMware.

Perform the following steps:

1. Stop the provider server by running the `/etc/init.d/sfcbd-watchdog stop` command.
2. Copy the `providerTraceLog.properties` file from `/etc/cim/lsi` to a temporary directory.
3. Modify the `providerTraceLog.properties` and set the debug level by uncommenting the below line:

   ```
   #LEVEL=ERROR" and change to "LEVEL=DEBUG
   ```

4. Restart the provider service by running the `/etc/init.d/sfcbd-watchdog start` command.
5. Recreate the issue.
6. Run the following command to collect the complete LSA logs:

   ```
   vm-support
   ```

7. Share the location of the log file and where the log file was generated.

8.5 Logout and Reboot Requirements on VMware

Some features/functionalities such as flashing the firmware, managing SAS storage link speed, managing PCIe lane speed, etc., may require server reboot for the changes to take effect.

In case of VMware, for the above-mentioned functionalities, you need to follow these instructions:

- Log out from the LSA client.
- Reboot the VMware Server.
- Once the VMware Server comes up, login again to check whether the changes have taken effect.

NOTE In cases where logout and reboot are required for certain functionalities as mentioned above, you will be notified through message - `Please Logout and Re-login to the server once Reboot is complete`.

8.6 Behavior of Event History

LSI Storage Authority supports retrieving the list of previous events, also known as event history. LSA maintains the event history in three different log levels:

- `0` – always read from firmware.
  - Firmware maintains a separate space to save event history, which can be retrieved using the `LSA.conf` file.
  - **MegaRAID Behavior** – The source of an event is firmware. All events are persisted in firmware and can be retrieved at any point in time.
  - **Non-MegaRAID Behavior**– The source of an event is Firmware, Driver, and StoreLIb. Firmware generated events can only be persisted and can be retrieved.
- `1` – by default, persists only non-MegaRAID events in the log file.
  - Non-MegaRAID events include IR/HBA.
- `255` – persists all events in the log file.
  - LSA monitors the log location of the individual controller under `$LSA_HOME/conf/monitor/logs`
  - Maximum number of events persisted in a file is 300. This can be configured through the `LSA.conf` file.
  - Maximum number of events that can be retrieved on a single page is 30. This can be configured through the `LSA.conf` file.
  - LSA maintains a set of events in the log file.
  - LSA starts maintaining the event history for Gateway, StandAlone, and DirectAgent installers from the time LSA is started.
  - For Indirect Agent installers, LSA maintains the event history from the time you login and as long as LSA Service is running.

8.7 Behavior of Event Monitoring on Non-ESXi Versus ESXi Server

The following table details the behavior of event monitoring on non-ESXi Versus ESXi Servers.

Table 6 Event Monitoring on Non-ESXi Versus ESXi Servers

Event	Non-ESXi Server (Gateway, StandAlone, and DirectAgent)	Non-ESXi (Indirect Agent)
SysLog/e-mail	Immediately after the LSA service is started. This includes new installations as well as restart of LSA services.	User should login at least once to the ESXi Server after the Gateway service starts.
Syslog Location	On Windows: Event viewer. On Linux: `/var/log/messages`	Events are logged in Gateway server's Syslog location.
Differentiating syslog	No special attribute is added as part of the event description.	Corresponding server ID is added as part of event description, which helps the user to identify any particular event that belongs to a particular ESXi server.
Time Sync	Handled	Not handled
Alert Configuration	Follows `config-current.json`	Follows Gateway's alert configuration.
SMTP server communication	LSA Server communicates with SMTP server for e-mail communication.	Instead of ESXi, Gateway server communicates with SMTP server for e-mail communication.

8.8 Limitations of Installation and Configuration

The following are the limitations of this installation and configuration.

- No status information exists for the controller.
- Events are collected as long as LSA runs on the client.
- LSA on VMware responds slower as compared to the response of the LSA on the Windows or Linux operating systems.

  Events are collected from the time the client logs in to a VMware ESXi machine for the first time. Events continue to be collected as long as the LSA Service is running.

8.8.1 Upgrading and Downgrading on IR/IT Controllers

Due to some issues with VMware ESXi 5.5, VMware ESXi 6.0, and VMware ESXi 6.5, you cannot upgrade or downgrade IR/IT firmware through LSA following the usual firmware flashing procedure. You need to follow the below suggested workaround:

1. Edit the `/etc/sfcb/sfcb.cfg`
2. Add the `httpMaxContentLength: 4194304` parameter.
3. Restart the LSA Service, `/etc/init.d/sfcbd-watchdog restart`
4. Now flash the firmware.

8.9 Differences in LSA for VMware ESXi

The following are some of the differences in LSA when you manage a VMware server.

- The following limitations apply to the system information exposed through the application:
  - Only the host name appears.
  - No support exists for the controller health information.
- **Authentication support**:
  - The LSI Storage Authority Software allows CIMOM server authentication with the user ID and the password for VMware.
- **Event logging**:

  Event logging support is available for the VMware ESXi operating system, but it works differently than the normal LSI Storage Authority Software framework mode. The event logging feature for LSA client connected to a VMware ESXi system behaves as follows:

  - The system logs are logged in the remote server instead of logging in the VMware ESXi Server.
  - The "View Log" option allows you to view the logs saved in a text file on the Event Logger dialog.
  - Refreshing the LSA GUI after any updates on the firmware is slower for a client connected to VMware ESXi hosts, compared to one that is connected to a Windows/Linux/Solaris host.
- VMware ESXi is supported only on **Gateway** installation. StandAlone, DirectAgent,, and Lightweight Monitor (LWM) installation modes are not supported.
- VMware ESXi is supported on following operating systems:
  - Microsoft Windows Server
  - RHEL
  - SuSE Linux

Chapter 9: Performing Initial Configuration

After successfully installing the LSI Storage Authority Software, you need to set up these initial configurations.

9.1 Using LDAP Authentication

To access the LDAP service, the LSI Storage Authority server must know some information about the LDAP server settings. Apart from the user name and password details for the LDAP authentication, the LSA back-end must know some parameters to enable authentication. Perform the following steps to configure these parameters in the `lsa.conf` file in the `LSI Storage Authority/conf` directory.

1. Open the `lsa.conf` file in the `LSI Storage Authority/conf` directory.
2. Enter a value for the `ldap_mode` field. If you set is as `0`, the LDAP authentication using the LSI Storage Authority software is disabled. If you set it as `1`, the LDAP authentication using the LSI Storage Authority software is enabled.

   Example:

```
LDAP Login
ldap_mode = 1
```

3. Enter the hostname of the LDAP server in the `ldap_server` field. This value is used to connect to the specific LDAP server for the user authentication.

   Example:

```
# LDAP Server
ldap_server = <Hostname of the LDAP server>
```

4. Optional step - Enter the LDAP protocol version in the `ldap_protocol_version` field. This value is used to define the protocol that is used to create an LDAP session.

   Example:

```
# LDAP Protocol version
ldap_protocol_version = v3
```

   The default value is `v3`.

5. Enter the LDAP authentication mode in the `ldap_binding` field. In LDAP, the authentication is supplied through the Bind operation. LDAP supports three types of authentication modes:
   - Anonymous – When an LDAP session is created, that is, when an LDAP client connects to the server, the authentication state of the session is set to the anonymous mode.
   - BASIC (default) – The simplest form of client authentication is to bind to the server using a clear-text password. This mechanism has security problems because the password can be read from the network.
   - SECURE – A more secured method is to use an Simple Authentication and Security Layer (SASL) authentication mechanisms, such as DIGEST-MD5[4]. This is based on an encryption known to both the client and the server, allowing for a simple challenge-response scheme. The SASL authentication mechanism is also capable of negotiating data encryption to protect subsequent operations.

   Example:

```
# LDAP_BINDING
ldap_binding = BASIC
```

6. Optional step - Enter the LDAP server port number in the `ldap_port_number` field.

   If you do not specify a port number, the standard LDAP port 389 is used for the BASIC authentication mode. For the SECURE authentication mode, the Port 636 is used

   Example:

```
# LDAP Port Number

ldap_port_number = 389
```

7. Enter the DN (distinguished name) details in the `dn_details` field. The format is as follows:

   Example:

```
# LDAP_DN_DETAILSdn_details
={"DN":[{"key":"DC","values":["ldapdomain"]},{"key":"DC","values":["com"]},{"key"
:"ou","values":["TEST"]}]}
```

   Where:

   — `DC` – This attribute contains the Domain Component type.
   — `ou` – This attribute contains the name of an organizational unit.

8. Optional step - Enter the LDAP user access privilege details in the `readOnly` field. The values follow:
   — `1` (default) – Read only access.
   — `0` – Full access
9. Restart the nginx Service and the LSI Storage Authority Service for the changes to take affect.

9.2 Accessing LSA Over Network Address Translation (NAT)

Network Address Translation (NAT) enables private IP networks that use unregistered IP addresses to connect to the Internet. NAT operates on a router, usually connecting two networks together, and translates the private addresses in the internal network into legal addresses.

To access the LSI Storage Authority (LSA) application over a NAT environment, the LSA server must know some information about the NAT server settings.

Perform the following steps to configure the parameters in the `lsa.conf` file in the `conf` directory.

1. Open the `lsa.conf` file in the `LSI Storage Authority/conf` directory.
2. Specify the public IP of `nat_ipv4_ipv6`

   For example, if the public NAT IP address configured is as `135.24.227.198`, you need to specify `nat_ipv4_ipv6 = 135.24.227.198`.
3. Restart the nginx service and the LSA Service for the changes to take effect.

   If you have multiple public NATs (for example, `135.24.227.198`, `135.24.227.199`, `fe80::dc8d:e156:41e1:b06`), you need to specify them as `nat_ipv4_ipv6 = 135.24.227.198`, `135.24.227.199`, `fe80::dc8d:e156:41e1:b06`

9.3 Changing the LSI Storage Authority Application Port Number

Perform the following steps to change the LSI Storage Authority Application port numbers.

1. Open the `lsa.conf` file in the `LSI Storage Authority/conf` directory.

2. Enter the new port number in the `listening_port` field.

   Prior to assigning the port number, ensure that the port is available for usage.
3. Save the `lsa.conf` file.
4. Open the `nginx.conf` file in the `LSI Storage Authority/server/conf` directory.
5. Replace all of the `fastcgi_pass 127.0.0.1:9000` instances with `fastcgi_pass 127.0.0.1:<new port number>`.
6. Save the `nginx.conf` file.
7. Open the `portconfig.properties` file in the `LSI Storage Authority` directory.
8. Enter the new port number in the `<Client Port> new port number </Client Port>` field.
9. Save the `portconfig.properties` file.
10. Restart the nginx Service and the LSI Storage Authority Service.

9.4 Hiding an Empty Backplane

By default, LSA displays all the empty backplanes connected to the controller in the **Other Hardware** tab. However, if you need to hide the empty backplanes from showing up in the **Other Hardware** tab, perform the following steps:

1. Stop the LSA`Service`.
2. Open the `lsa.conf` file in the `LSI Storage Authority/conf` directory.
3. In the LSA`.conf` file, search for `empty_SGPIO_display` field.

   Modify the existing or default value in the `empty_SGPIO_display` field to `0`.
   - `0` –Disables the empty backplanes from showing up in the **Other Hardware** tab.
   - `1` – Enables the empty backplanes from showing up in the **Other Hardware** tab.

   By default, the `empty_SGPIO_display` field is set to `1`.
4. Start the LSA`Service` once again.

9.5 Changing the nginx Web Server Port Number

Perform the following steps to change the nginx web server port numbers.

1. Open the `nginx.conf` file in the `LSI Storage Authority/server/conf` directory.
2. Replace all of the `listen 2463 default_server ssl instances` with `listen <new port> default_server ssl`.
3. Save the `nginx.conf` file.
4. Restart the nginx service and the LSI Storage Authority service.

9.6 Blacklisting Private IPs on the Remote Server Discovery Page

If you need to discover the sever using the Remote Server Discovery page, it is recommended not to use the default private IP addresses starting with the first octet (for example, `Default=10,169,172,192`).

If you are still using these default private IP addresses, they will be blacklisted. You need to remove the IP address range starting from with the first octet.

Chapter 10: Performing Initial Setup

After you successfully log on to the LSI Storage Authority software, it is suggested that you perform certain initial setup tasks before proceeding.

10.1 Managing Servers from the Remote Server Discovery Page

The LSI Storage Authority software allows you to set up a list of servers to monitor and manage. Perform the following steps to manage the servers:

1. On **Remote Server Discovery** page, click the **Go To - Manage Server Page** hyperlink.
   The **Gateway - Authenticate** dialog opens.

Figure 8 Gateway Authenticate Dialog

2. Enter the administrator credentials for the Gateway server.
   a. Select either **DOMAIN** or **HOST**as the option from the drop-down list.
   b. Specify the user name and the password in their respective fields.

   The gateway server persists the login credentials in an encrypted file.
3. Click **Sign In**.
   The **Remote Server Discovery** page switches to the **Managing Servers** page.

Figure 9 Managing Server Mode

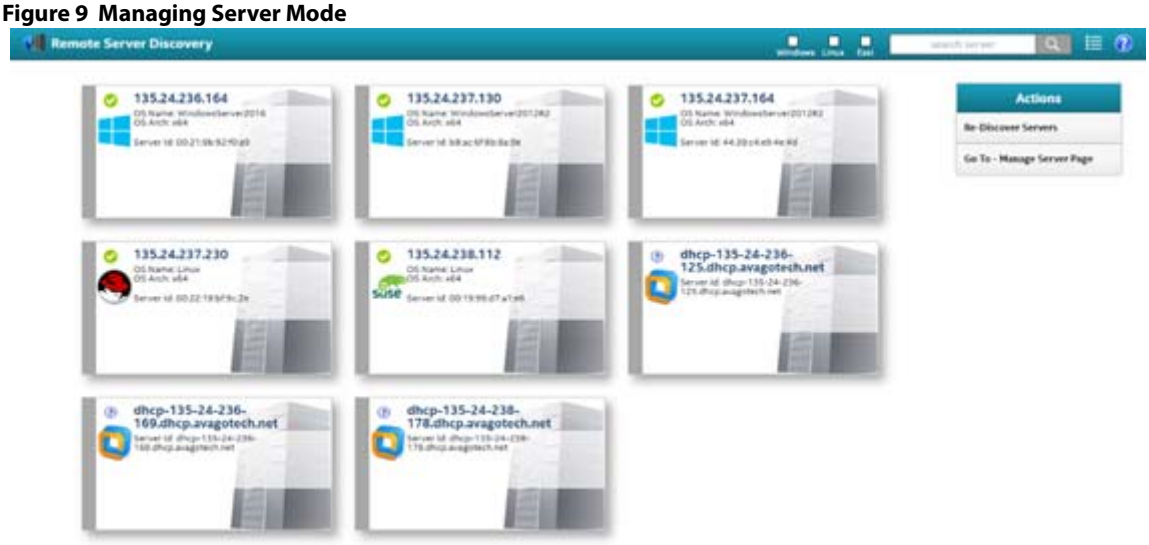

On the **Remote Server Discovery** or **Manage Servers** page, you can:

- See the list of managed servers with their health status.
- Add and remove the managed servers from the list. For more information, see Adding Managed Servers and Removing Managed Servers.
- Rediscover the servers or go back to the **Remote Server Discovery** page.
- Using the Check Boxes, you can select only those servers that you require to be displayed.

10.2 Displaying or Blocking a Private IP Address

This topic outlines the strategy that LSA is following for displaying or blocking a private IP address in a corresponding sub-net.

Private IP address – A private IP address is a non-internet facing IP address on an internal network. Private IP addresses are provided by network devices, such as routers, using network address translation (NAT).

Virtual IP address – A virtual IP address (VIPA) is an IP address assigned to multiple domain names or servers that share an IP address based on a single network interface card (NIC). VIPAs are allocated to virtual private servers, websites, or any other application residing on a single server. The host server for these applications has a network IP address assigned by a network administrator, whereas the different server applications have VIPAs. VIPAs enhance network load balancing and redundancy.

Automatic Private IP Addressing – Automatic Private IP Addressing (APIPA) is a feature of Windows-based operating systems that enable a computer to automatically assign itself an IP address when there is no Dynamic Host Configuration Protocol (DHCP) server available to perform that function. APIPA serves as a DHCP server failover mechanism and makes it easier to configure and support small local area networks.

Private IP Address Range– The following is the IP address range which falls under either of private (or) Virtual (or) APIPA category:

- **NAT**– `10.0.0.0 - 10.255.255.255`
- **Private (or) Virtual**– `172.16.0.0 - 172.31.255.255` or `192.168.0.0 - 192.168.255.255`
- **APIPA**– `169.254.0.0` to `169.254.255.255`

Behavior or LSA: During the discovery of the servers, LSA checks if there are any changes in the network configuration. If LSA finds any changes in the network configuration, LSA deregisters the initial IP address with OpenSLP and

reregisters with the latest network IP. In this way, LSA ensures that any network changes that have happened in a real-time are addressed without even restarting the LSA services.

The following use cases provide details on how LSA behaves in various situations:

Table 7 Use case#1: Without blocking the Private IP

Use Case	Standalone / Client	Gateway / Client	Remarks
No NIC CARD (Windows)	Loopback (or) 127.0.0.1	Cannot access	As the server is not in network, LSA Gateway cannot access the Standalone server.
No NIC CARD (Linux)	Loopback (or) 127.0.0.1	Cannot access	As the server is not in network, LSA Gateway cannot access the Standalone server.
Static IP	Using Static IP	Using Static IP	
DHCP IP	Using the DHCP IP	Using the DHCP IP	
Private IP	Using the Private IP	Can display the Private IP irrespective of whether it is under the Ethernet IP (or) Private IP.	In a more secured environment, Private IP address cannot be accessed outside the server.

Table 8 Use case#2: After blocking the Private IP

Use Case	Standalone / Client	Gateway / Client	Remarks
No NIC CARD (Windows)	Loopback (or) 127.0.0.1	Cannot access	As the server is not in network, LSA Gateway cannot access the Standalone server.
No NIC CARD (Linux)	Loopback (or) 127.0.0.1	Cannot access	As the server is not in network, LSA Gateway cannot access the Standalone server.
Static IP	Using Static IP	Using Static IP	
DHCP IP	Using the DHCP IP	Using the DHCP IP	
Private IP	If there is a valid IP, it will be. displayed. If there is no valid IP, Loopback (or) 127.0.0.1 will be displayed.	Can display the Private IP irrespective of whether it is under the Ethernet IP (or) Private IP.	In a more secured environment, as Private IP address cannot be accessed outside the server, LSA does not populate a Private IP Address.

Why LSA blocks certain IP Addresses: In an enterprise world, when a computer is assigned a private IP address, the local devices see this computer through its private IP address. However, the devices residing outside of your local network cannot directly communicate through the private IP address, but uses your router's public IP address to communicate. To allow direct access to a local device which is assigned a private IP address, a Network Address Translator (NAT) must be used.

In a more secured environment, although LSA is able to discover and display the Private IP address through the Gateway server, when the request(s) is made through the Gateway server, the Private IP will not be accessible. As LSA cannot access the Private IP, LSA is unable to service the requests which are meant for the Private IP.

Due to the above-mentioned reason, when the LSA installation is Gateway, the corresponding Gateway server will not be able to communicate with the Private IP Address which in turn will be an issue. LSA will work if the private IP addresses are behind the NAT router which is the most preferable option in an enterprise world.

The following illustration explains how a private IP address should be accessed in enterprise networks and the problems with the Private IP address:

Figure 10 Private IP Address Access

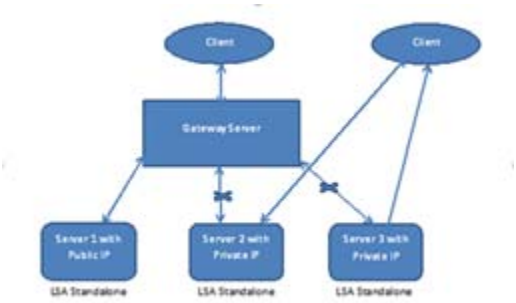

10.3 Adding Managed Servers

Perform the following steps from the **Manage Servers** page to add the managed servers.

1. Select a server that you want to add from the list of discovered servers, and click the  icon.
   The **Remote - Authenticate** dialog appears.

Figure 11 Remote - Authenticate Dialog

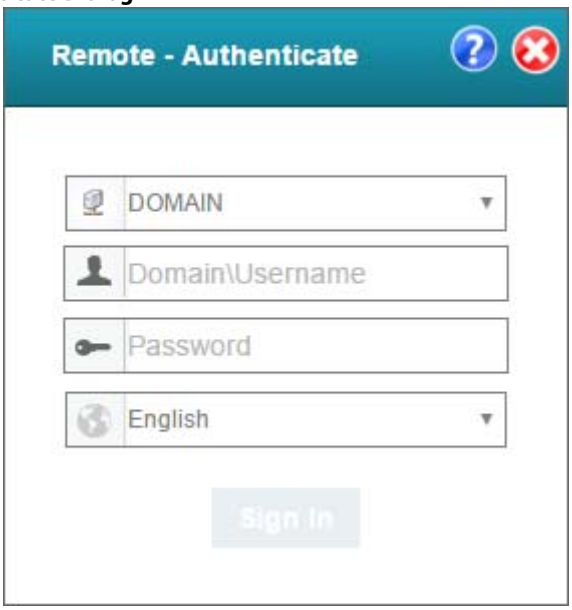

2. Enter the user credentials for the server you want to add.
   a. Select either **DOMAIN** or **HOST** as the option from the drop-down list.
   b. Specify the user name and the password in their respective fields.
3. Click **Sign In**.
   The server is added to the list of managed servers. The  icon changes to  icon.
4. Click the server that you have added to the managed server list.
   The Server dashboard page for the server appears. See Server Dashboard.

10.4 Removing Managed Servers

Perform the following step from the **Manage Servers** page to remove the managed servers.

1. Click the  icon.
   The host is removed from the list of managed servers. The  icon changes to the  icon.

10.5 Alert Settings

The **Alert Settings** tab lets you perform the following actions:

- Change the alert delivery method for different severity levels.
- Specify different alert delivery methods for inside and outside the application.
- Revert back to the default alert delivery methods and the default severity level of an individual event.
- Save the alert settings on the server.

Based on the severity level (Information, Warning, Critical, and Fatal), the default alert delivery methods change. By default, each severity level has one or more alert delivery methods configured for it. The different alert delivery methods are as follows:

- **System Log** – By default, all of the severity events are logged in the local syslog. In the Windows operating system (OS), the system log is logged in **Event Viewer > Application**. In the Linux OS, the system log is logged in **var > log**.
- **Event Log** – By default, all the severity events appear in the event log. Click **View Event Log** to view the event log. Each message that appears in this log has a severity level that indicates the importance of the event (severity), an event ID, a brief description, and a date and timestamp (when it occurred).
- **System Messages** – By default, fatal and critical events are displayed as system messages. System messages are displayed in a yellow bar at the top of the Server dashboard and the controller dashboard. System messages let you view multiple events in a single location.
- **Email** – By default, fatal events are displayed as email notifications. Based on your configuration, the email notifications are delivered to your inbox. In the email notification, besides the event's description, the email also contains system information and the controller's image details. Using this additional information, you can determine the system and the controller on which the fatal error occurred.

To change the alert delivery method for each severity level, perform these steps:

1. Click **Username > Settings** in the Server dashboard.
   The **Alert Settings** window appears, which the default alert delivery methods for each severity level.

Figure 12 Alert Settings Window

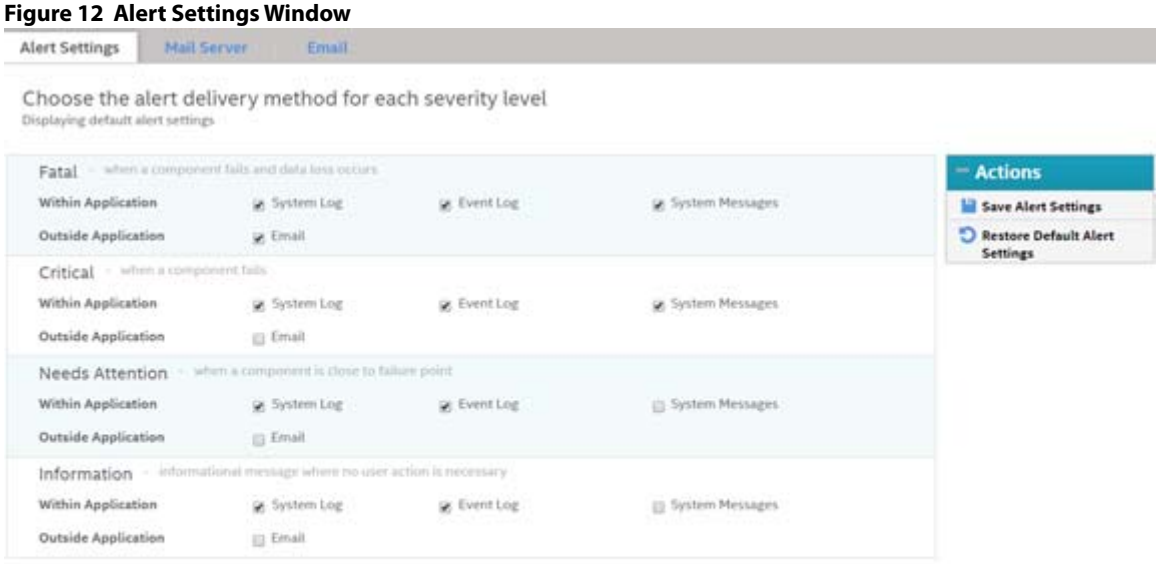

2. Select the desired alert delivery method for each severity level by clicking the required check box.
3. Click **Save Alert Settings** to save the settings on the server.

Click **Restore Default Alert Settings** to revert back to the default alert delivery method settings.

10.6 Setting Up the Email Server

Perform the following steps to enter or edit the mail and the SMTP server settings.

1. In the **Settings** window, click the **Mail Server** tab.
   The **Mail Server** tab appears and displays the current mail server settings.

Figure 13 Mail Server Window

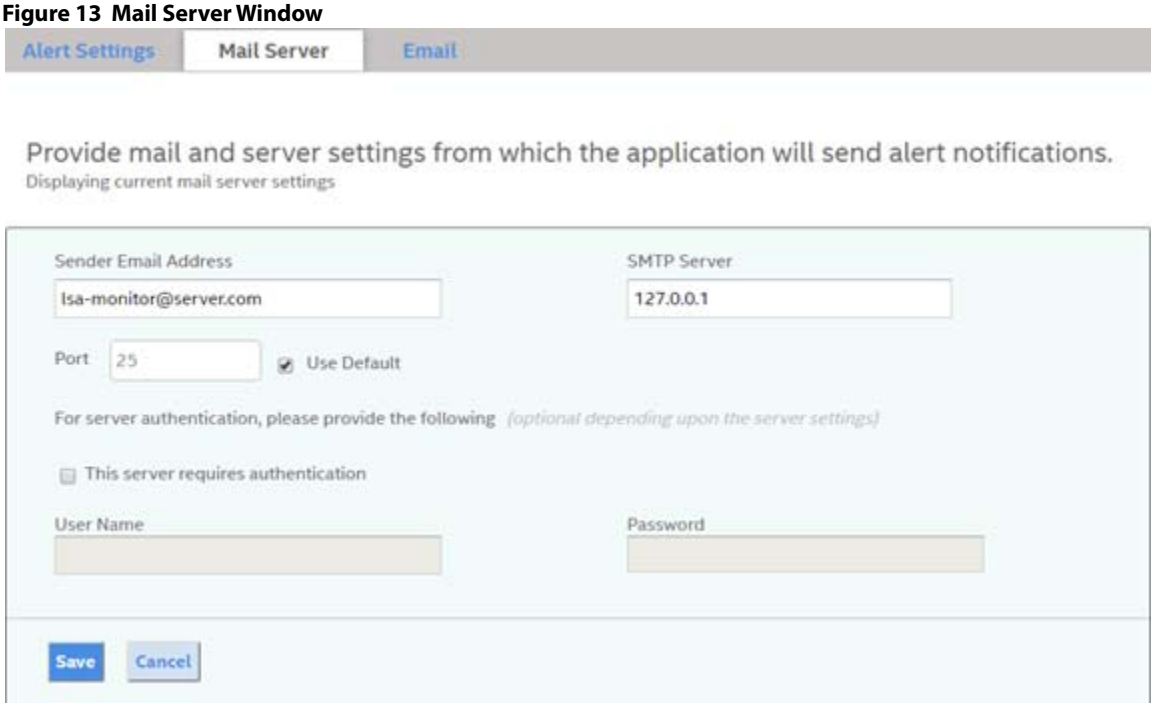

2. Enter a sender's email address in the **Sender Email Address** field, or edit the existing sender email address.
3. Enter your SMTP server name/IP address in the **SMTP Server** field, or edit the existing details.
4. Clear the **Use Default** check box to enter the desired port number in the **Port** field.
5. On your SMTP server, if the Auth Login feature is enabled and if you want to enable this feature on the LSI Storage Authority software, select the **This server requires authentication** check box and specify the authentication details in the **User Name** and **Password** fields.
6. Click **Save**.

10.7 Adding Email Addresses of Recipients of Alert Notifications

Perform the following steps to add email addresses of recipients of the alert notifications.

1. In the **Setting** window, click the **Email** tab.
   The **Email** tab appears and displays the current email settings.

Figure 14 Email Window

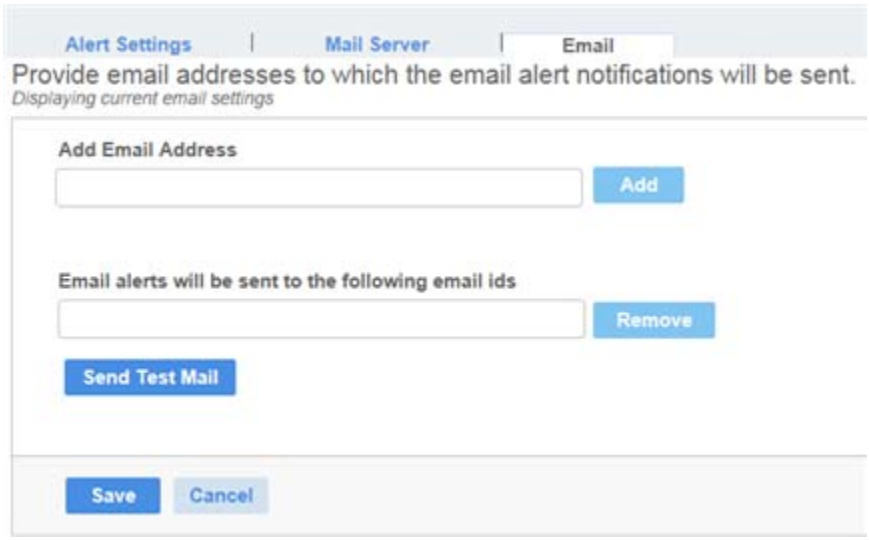

2. Enter the email address you want to add in the **Add Email Address** field.
3. Click **Add**.

   The new email address appears in the **Email alerts will be sent to the following email ids** field.

   You can click **Remove** to delete the email addresses that are added.
4. Click **Send Test Email** to send a test message to the email addresses that you added for the recipients of alert notifications.

   A pop-up message indicates if the test message was successfully sent to the email address.
5. Click **Save** to save the email settings.

Chapter 11: Server Dashboard

The Server dashboard is the default landing page in the LSI Storage Authority software. The Server dashboard displays the overall summary of the server and the devices attached to it. You can troubleshoot, configure, maintain, and monitor the controllers from the Server dashboard. The following figure and table describe this page.

Figure 15 Server Dashboard

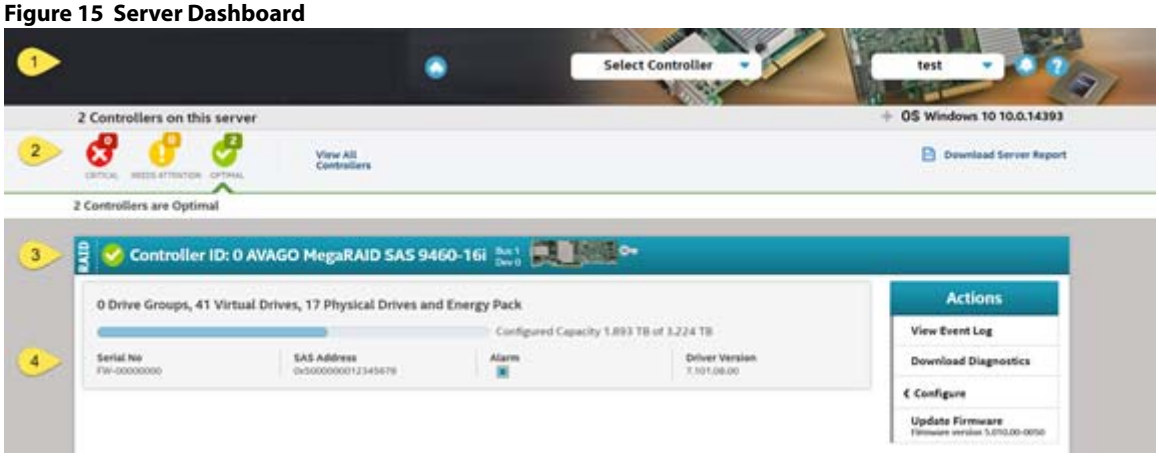

Chapter 12: Controller Dashboard

You can perform controller related actions and view all the information pertaining to a controller from the Controller dashboard. The following figure and table describe this page.

Figure 16 Controller Dashboard

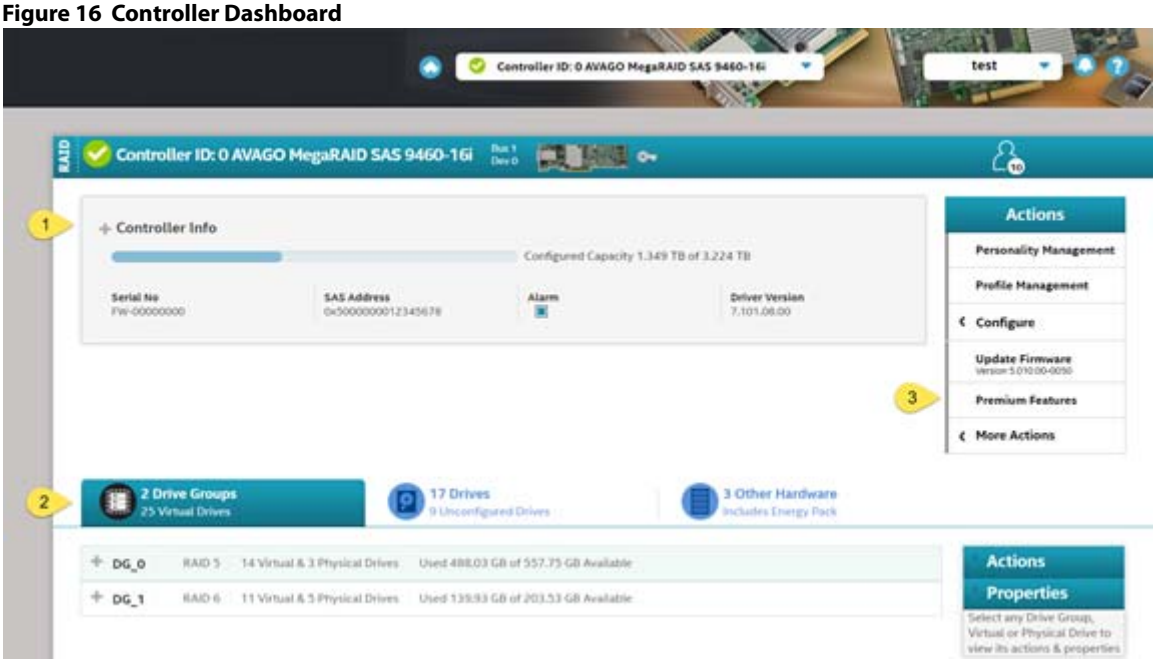

Table 9 Controller Dashboard Description

Callout	Description
1	**Controller Summary** - Displays the name of the ServeRAID controller card. The color-coded icons indicate the status of the controller card. Displays the basic controller properties, such as the controller serial number, vendor ID, SAS address, driver version, device ID, host interface, and so on. Click the + icon to view the advanced properties of the controller, such as the NVRAM details, data protection information properties, BIOS version, firmware properties, drive security properties, emergency spare properties, and so on.
2	**Controller Views** - Displays all of the configured drive groups, virtual drives, and physical drives associated with the selected controller card. It also displays the hardware, such as enclosures, backplanes, and the supercapacitor associated with the controller. All these views are displayed as tabs. Click the + icon to view to view detailed information about the device. For example, click a drive group to view the associated virtual drives and physical drives. Select any device from the expanded view to perform relevant actions and view device properties.
3	**Controller Actions** - Lets you perform the following actions: ■ Create configuration ■ Clear configuration ■ Enable or disable an alarm ■ Update the controller firmware ■ Import or clear foreign configurations ■ View Premium features ■ View event log

Chapter 13: Configuration

You can use the LSI Storage Authority software to create and modify storage configurations on systems with Lenovo controllers.

You can create RAID 0, RAID 1, RAID 5, RAID 6, RAID 00, RAID 10, RAID 50, RAID 60, RAID 1E, and Spanned R1E (PRL-11) storage configurations.

The supported RAID levels differ or might not be supported for some controllers. For more information, see LSI Storage Authority Feature Comparison Matrix.

You can create the following types of configurations:

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

13.1 Creating a New Storage Configuration Using The Simple Configuration Option

Simple configuration is the quickest and easiest way to create a new storage configuration. When you select simple configuration mode, the system creates the best configuration possible using the available drives.

NOTE When a physical drive is in the **Prepare for Removal** state, you cannot create a virtual drive using that physical drive. To create a virtual drive when the physical drive is in the **Prepare for Removal** state, you must manually undo the operation by navigating to the **Undo Removal** option.

Perform the following steps to create a simple storage configuration:

1. On the Server dashboard or on the Controller dashboard, select **Configure > Simple Configuration**.
   The **Simple Configuration** window opens.

Figure 17 Simple Configuration Window

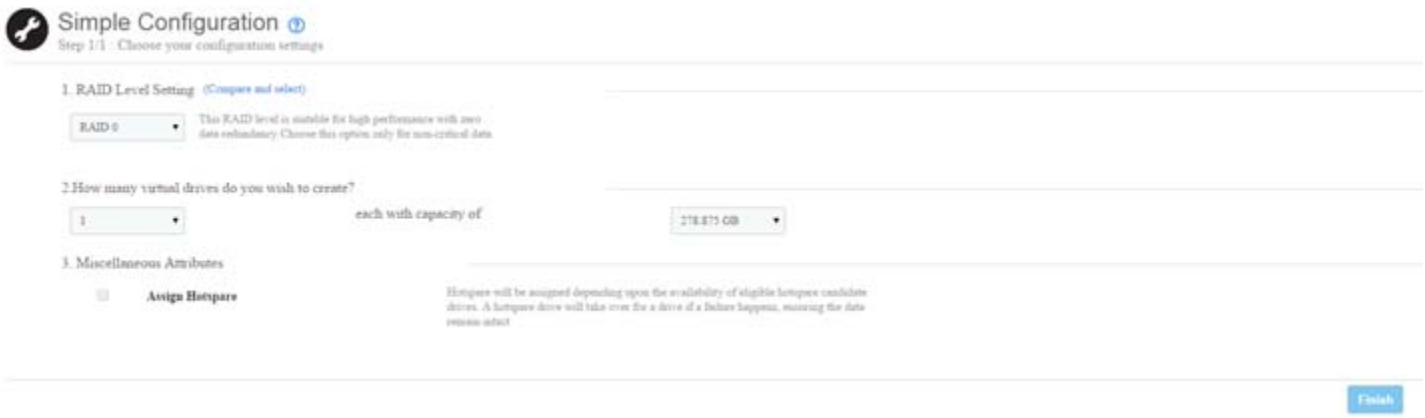

2. Select a RAID level for the drive group from the drop-down box.

3. (Optional) click **Compare and Select** to view the detailed information on each RAID level.

   When you use simple configuration, the RAID controller supports RAID levels 0, 1, 5, 6, and PRL-11 (RAID-1E). The window text gives a brief description of the RAID level that you select. The RAID levels that you can choose depend on the number of drives available.
4. Select the number of virtual drives you want to create.
5. Select the capacity of the virtual drives. Each virtual drive has the same capacity.
6. Select the **Assign Hotspare** check box if you want to assign a dedicated hot spare to the new virtual drive.

   If an unconfigured good drive is available, that drive is assigned as a hot pare. Hot spares are drives that are available to replace failed drives automatically in a redundant virtual drive (RAID 1, RAID 5, RAID 6, or RAID-1E (PRL-11)).
7. Click **Finish**.

   A message appears stating that the configuration is successfully created.

13.2 Creating a New Storage Configuration Using The Advanced Configuration Option

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

NOTE When a physical drive is in the **Prepare for Removal** state, you cannot create a virtual drive using that physical drive.

Perform the following steps to create an advanced storage configuration.

1. On the Server dashboard or the Controller dashboard, select **Configure > Advanced Configuration**.

   The **Advanced Configuration** window opens.

Figure 18 Advance Configuration Window

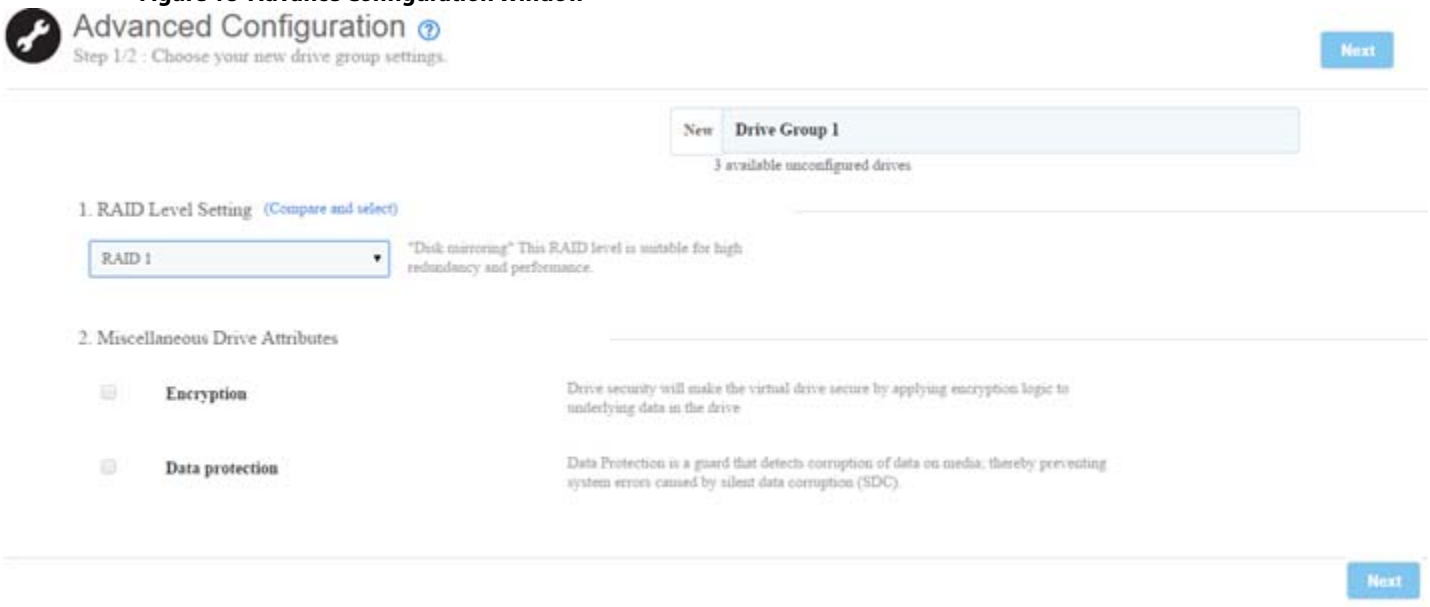

2. Select a RAID level for the drive group from the drop-down box.

3. (Optional) click **Compare and Select** to view the detailed information on each RAID level.

   When you use advanced configuration, the RAID controller supports RAID levels 00, 10, 50, 60, RAID 1E, and Spanned R1E (PRL-11). The window text gives a brief description of the RAID level that you select. The RAID levels that you can choose depend on the number of drives available.

4. (Optional) Select the **Encryption** check box if you want to apply the encryption logic to secure the data in the virtual drive.

   You can add a hot spare to all of the RAID levels except RAID 0. Also, you can create a secured virtual drive only when the security capable drives are present. This check box is disabled when there are no secured drives.

5. (Optional) Select the **Data Protection** check box to detect data corruption on media and prevent system errors caused by silent data corruption (SDC).

   This check box is disabled when there are no secured drives.

6. Click **Next**.

7. Click **Add Physical Drives** to add physical drives to the drive group.

   The **Available Unconfigured Drive** window appears.

Figure 19 Available Unconfigured Drive Window

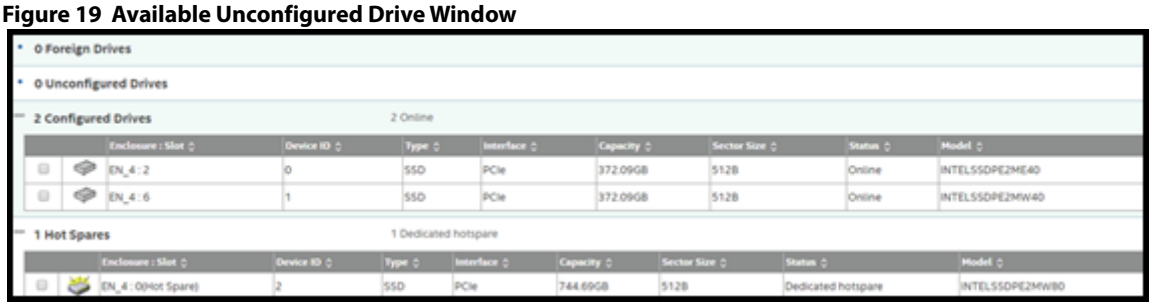

   For information on adding the unconfigured drives to the drive group, see Selecting Available Unconfigured Drive.

8. Select the span depth using the slider bar.

9. Click **Add Virtual Drives** to add virtual drives to the drive group.

   The **Virtual Drive Settings** window appears.

Figure 20 Virtual Drive Settings Window

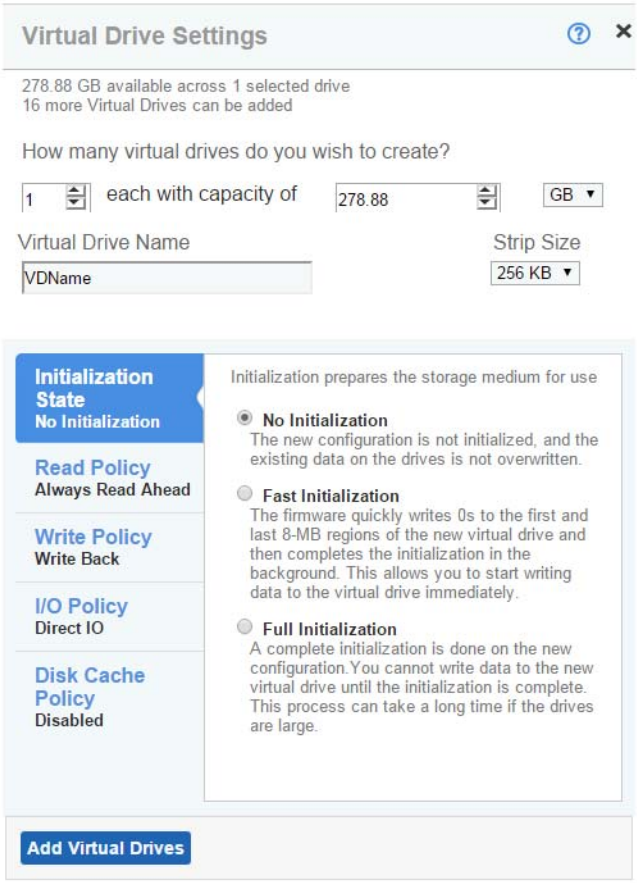

For information on configuring virtual drives, see Selecting Virtual Drive Settings.

10. Click **Finish**.
    A message appears stating that the configuration is complete.

13.2.1 Selecting Available Unconfigured Drive

The **Available Unconfigured Drive** window lets you add physical drives and hot spares to the drive group.

Perform the following steps to add physical drives and hot spares to the drive group.

1. In the **Available Unconfigured Drives** window, select the physical drives and click **Add Physical Drives**.
   The selected physical drives appear in the **Advanced Configuration** window.
   You can click the ✕ icon to remove the physical drives that you have already added.
2. Click **Add Hot Spares** to add dedicated hot spare drives to the drive group.
   The **Available Unconfigured Drives** window appears.
3. Select the drives you want to add as hot spares and click **Add Hot Spares**.
   The selected hot spares appear in the **Advanced Configuration** window.

13.2.2 Selecting Virtual Drive Settings

The **Virtual Drive Settings** window enables you to configure the virtual drives. Detailed descriptions for all of the parameters are present in the **Virtual Drive Settings** window.

The virtual drive settings differ or might not be supported for some controllers. For more information, see LSI Storage Authority Feature Comparison Matrix.

Perform the following steps to configure a virtual drive:

1. Specify the number of virtual drives you want to create.
2. Specify the size of the virtual drives you want to create.

   Each virtual drive has the same capacity. If you specify the capacity first and then the number of virtual drives, the virtual drive capacity is adjusted with the available capacity.
3. Enter a name for the virtual drive in the **Virtual Drive Name** field.

   The virtual drive name can have a maximum of 15 characters.
4. Select a strip size from the **Strip Size** drop-down list.

   Strip sizes of 64 KB, 128 KB, 256 KB, 512 KB, and 1024 KB are supported.
5. Specify the initialization status. The options follow:
   - **Fast Initialization**
   - **Full Initialization**
   - **No Initialization**
6. Specify the read policy for the virtual drive. The options follow:
   - **No Read Ahead**
   - **Always Read Ahead**
7. Specify the write policy for the virtual drive. The options follow:
   - **Write Through**
   - **Write Back**
   - **Always Write Back**

   The write policy depends on the status of the Energy Pack. If the Energy Pack is not present, is low, is failed, or is being charged, the current write policy switches to Write Through.
8. Specify the I/O policy for the virtual drive. The options follow:
   - **Cached IO**
   - **Direct IO**
9. Specify a disk cache setting for the virtual drive. The options follow:
   - **Unchanged**
   - **Disabled**
   - **Enabled**
10. Click **Add Virtual Drives**.

    The newly created virtual drive appears in the **Advanced Configuration** window just below the **Virtual Drives** section.

    NOTE You will lose some drive capacity if you choose drives with uneven and large capacity while creating a virtual drive.

    If you want to modify the virtual drive settings before finishing the configuration, click the  icon.

    The **Virtual Drive Settings** window opens.

    You can modify the settings and click **Modify Virtual Drive**.

13.3 Clearing the Configuration

You can clear all existing configuration on a selected controller.

Perform the following steps to clear the existing configurations on a controller.

1. Navigate to the Controller dashboard whose configurations you want to clear.
2. Click **Configure** and then click **Clear Configuration**.
   A confirmation message appears.
3. Select **Confirm** and click **Yes, Clear configuration** to clear all the existing configurations on the controller.
   Operating system drives cannot be cleared.

13.4 Importing or Clearing the Foreign Configurations

A foreign configuration is a RAID configuration that already exists on a replacement set of drives that you install in a computer system. You can use the LSI Storage Authority software to import the foreign configuration to the controller or clear the foreign configuration so that you can create a new configuration using these drives.

Perform the following steps to import or clear foreign configurations.

1. Navigate to the Controller dashboard.
2. Click **Configure** and then click **Foreign Configuration**.
   The **Foreign Configuration** window appears, which lists all of the foreign configurations.
3. Click one of the following options:
   - **Import All**: Import the foreign configurations from all the foreign drives.
   - **Clear All**: Remove the configurations from all the foreign drives.
4. Click **Re-Scan** to refresh the window.
   You can import or clear the foreign configuration on security enabled drives. See Importing or Clearing a Foreign Configuration - Security Enabled Drives.

Chapter 14: Background Operations Support

The LSI Storage Authority software provides a background Pause, Resume, Abort, Pause All, Resume All, and Abort All features that enhance the functionality where in the background operations running on a physical drive or a virtual drive can be paused for some time, and resumed later.

The background operations, including **Consistency Check**, **Rebuild**, **Replace**, and **Initialization** are supported by an Abort operation. If any operation is stopped before completion, it is considered to be aborted. An aborted operation cannot be resumed from the place where it was stopped.

To perform **Pause**, **Resume** and **Abort** operations, go to the **Background Processes in Progress** window in the Server dashboard or the Controller dashboard, and perform the following steps. The **Background Processes in Progress** window is as shown in the following figure.

Figure 21 Background Processes in Progress Window

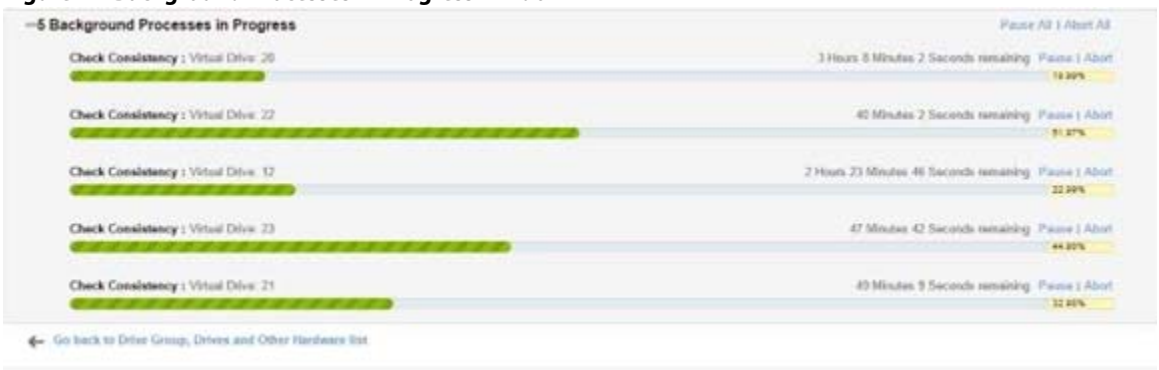

- **Pause** – Click **Pause** to suspend the background operation taking place at that particular point of time. When the operations gets paused, the **Resume** option appears instead of the **Pause** option.
- **Resume** – Click **Resume** to resume the operation from the point where it was suspended last.
- **Abort** – Click **Abort** to abort the ongoing active operation.
- **Pause All** – Click **Pause All** to suspend all the active operations. This option is enabled only if one or more background operations are in active state.
- **Resume All** – Click **Resume All** to resume all the paused operations from the point they were paused. This option is disabled if no operations are paused.
- **Abort All** – Click **Abort All** to abort all the active operations.

NOTE In case the Copyback progress bar is not displaying the progress of the Copyback operation for small-size volumes automatically, set the maximum event grouping time gap to `0` in the LSA`.conf` file.

Chapter 15: Managing Controllers

The LSI Storage Authority software enables you to monitor the activity of all the controllers present in the system and the devices attached to them.

15.1 Viewing Controller Properties

The Controller dashboard displays basic controller properties. Click the + icon to see the advanced properties of the controller.

Click the **Click to download all the controller properties** link to download the properties in the in the . JSON format.

Figure 22 Basic and Advanced Controller Properties

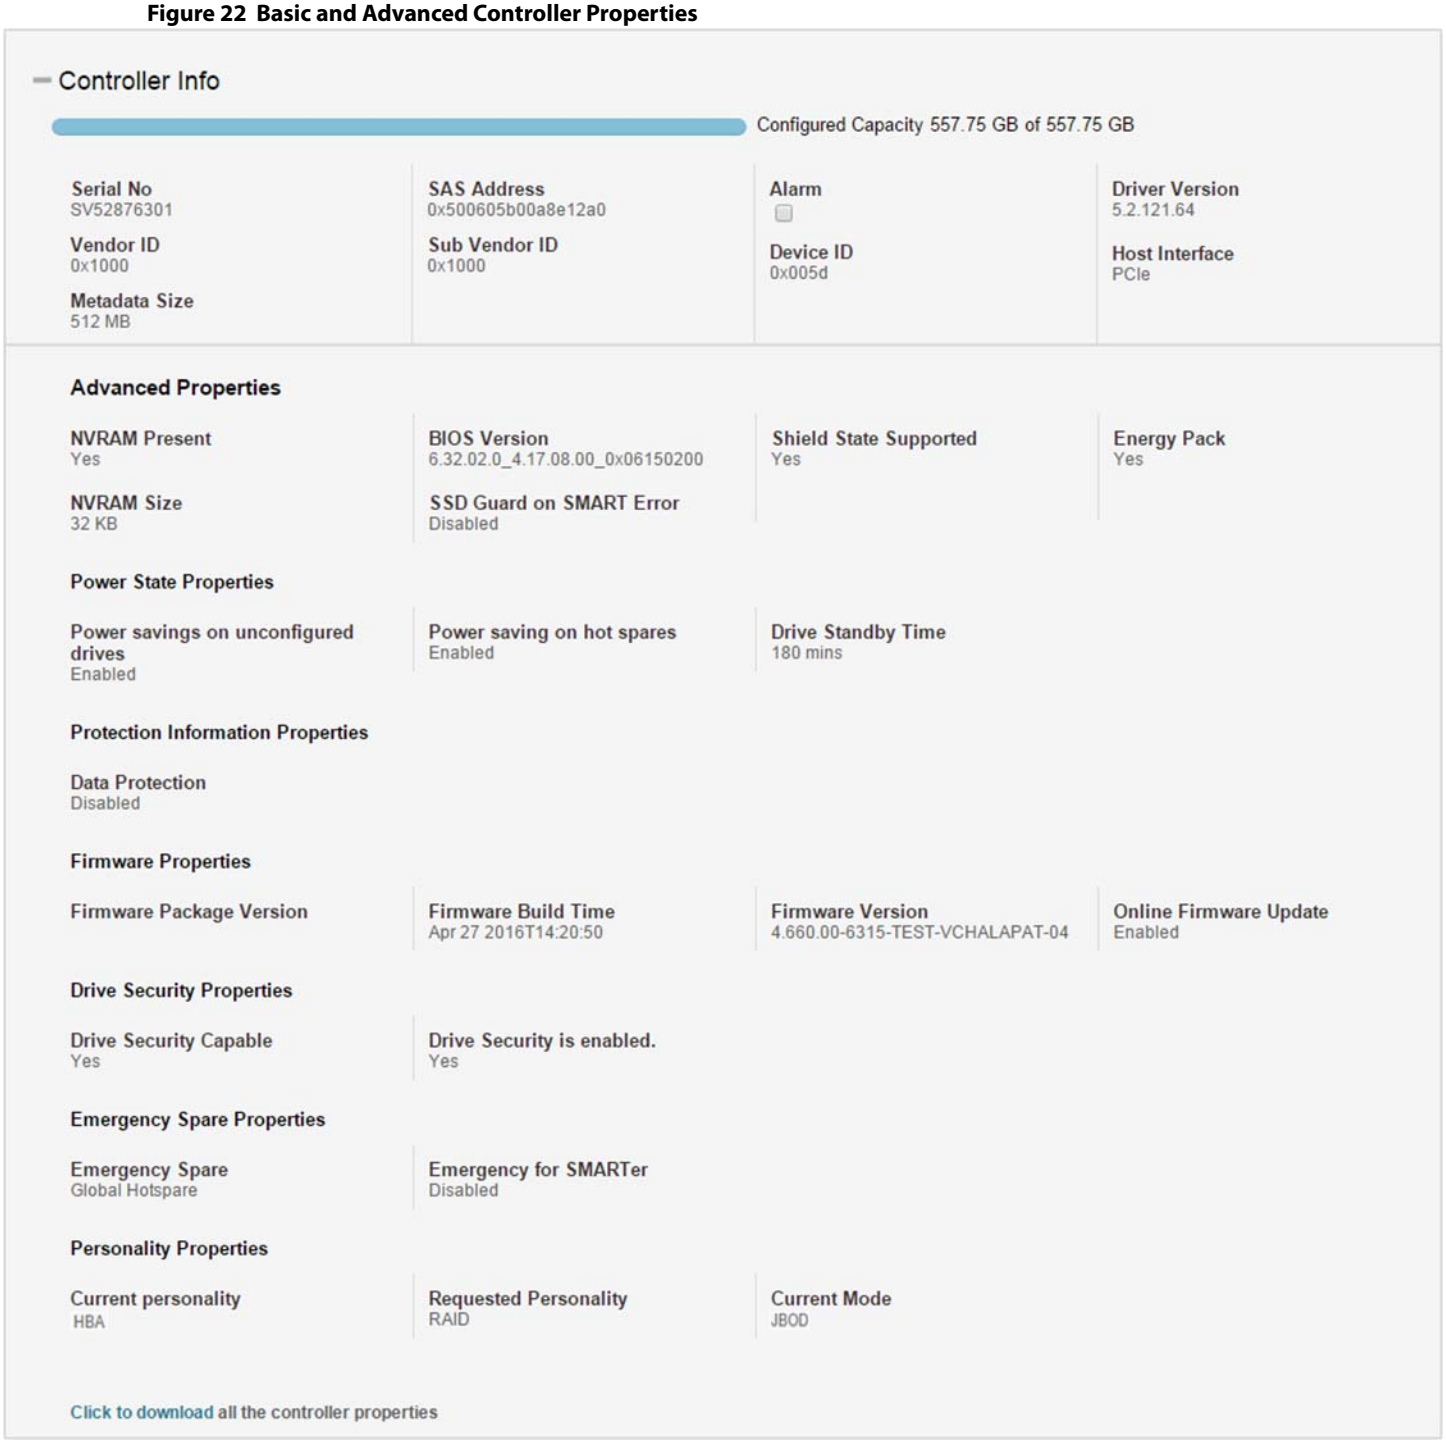

Table 10 Basic and Advanced Controller Properties

Property	Description	MegaRAID	iMegaRAID	Initiator-Target
Serial Number	The serial number of the controller.	Yes	Yes	Yes
SAS Address	The SAS address of the controller.	Yes	Yes	Yes
Alarm	Enables or disables the alarm.	Yes	Yes	No
Driver Version	The driver version of the controller.	Yes	Yes	Yes

Table 10 Basic and Advanced Controller Properties (Continued)

Property	Description	MegaRAID	iMegaRAID	Initiator-Target
Vendor ID	A unique controller ID assigned to a specific vendor.	Yes	Yes	Yes
Sub Vendor ID	Additional vendor ID information about the controller.	Yes	Yes	Yes
Device ID	The device ID that is assigned by the manufacturer.	Yes	Yes	Yes
Host Interface	The type of interface used by the computer host system.	Yes	Yes	Yes
Meta Data Size	The total space used for metadata. The following are displayed as size units: ■ If the size is less than 1 MB (1024 KB), the size is displayed in KB. ■ If the size is greater than or equal to 1 MB but less than 1 GB (1024 MB), the size is displayed in MB. ■ If the size is greater than or equal to 1 GB but less than 1 TB (1024 GB), the size is displayed in GB.	Yes	Yes	Yes
NVRAM Present	Indicates if a nonvolatile random access memory (NVRAM) is present on the controller.	Yes	Yes	Yes
NVRAM Size	Indicates the capacity of the controller's NVRAM.	Yes	Yes	No
BIOS Version	The BIOS version of the controller.	Yes	Yes	No
SSD Guard on SMART Error	Indicates if the SSD Guard feature is enabled on the controller.	Yes	Yes	No
Shield State Supported	Indicates whether t controller supports the shield state.	Yes	Yes	No
Energy Pack	Indicates if the energy pack is present.	Yes	Yes	No
Power State Properties				
Power savings on unconfigured drives	Indicates if the power savings on the unconfigured drives is enabled.	Yes	Yes	No
Power saving on hot spares	Indicates if the power savings on the hot spares is enabled or not	Yes	Yes	No
Drive Standby Time	Shows the drive standby time in minutes	Yes	Yes	No
Power Information Properties				
Data Protection	Indicates if data protection is enabled	Yes	Yes	Yes
Firmware Properties				
Firmware Package Version	The firmware package version of the controller	Yes	Yes	Yes
Firmware Build Time	The last firmware build time.	Yes	Yes	No
Online Firmware Update	Indicates if the Online Firmware Update Feature is enabled in the firmware.	Yes	Yes	No
Firmware Version	The firmware version of the controller.	Yes	Yes	Yes
Drive Security Properties				
Drive Security Capable	Indicates the drive security (encryption) feature status on the controller	Yes	No	Yes

Table 10 Basic and Advanced Controller Properties (Continued)

Property	Description	MegaRAID	iMegaRAID	Initiator-Target
Drive Security Enabled	Indicates whether the drive security is enabled	Yes	No	Yes
Emergency Spare Properties				
Emergency Spare	Indicates the Emergency Spare controller properties. It can be set to **Unconfigured Good** or**Unconfigured Good and Global Hotspare**.	Yes	No	Yes
Emergency for SMARTer	Indicates if emergency hot spare drives are commissioned for predictive analysis	Yes	No	Yes
CacheCade Properties				
CacheCade SSD Caching	Indicates if SSD Caching feature is enabled.	Yes	No	Yes
Write Cache Capable	Indicates if write cache feature is enabled	Yes	No	Yes[a]
Total Cache Size	Total available cache size	Yes	No	Yes
Maximum Cache Size	Maximum available cache size.	Yes	No	Yes

a. Write Cache of the LUN is not supported by all target enclosures. Check your enclosure's documentation to determine whether Write Cache is supported on your device.

15.2 Running Consistency Check

Consistency check operation verifies the correctness of the data in virtual drives that use RAID levels 1, 5, 6, 10, 50, 60, 1E, and Spanned RAID 1E configurations. 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. You should periodically run a consistency check on fault-tolerant virtual drives.

Because RAID 0 does not provide data redundancy, you cannot run a consistency check on RAID 0 volumes.

To run a consistency check, you must first set the consistency check properties, and then you can either schedule a consistency check to be run at a defined interval chosen by you or you can start the consistency check operation immediately.

15.2.1 Setting Consistency Check Properties

Perform the following steps to set the properties for a consistency check.

1. In the Controller dashboard, select **More Actions > Set Consistency Check Properties**.

   The **Set Consistency Check Properties** dialog appears.
2. Choose one of the two options:
   - **Continue Consistency Check and Fix Error** - The RAID controller continues the consistency check, and if any errors are found, fixes them.
   - **Stop Consistency Check On Error** - The RAID controller stops the consistency check operation if it finds any errors.
3. Click **Save**.

15.2.2 Scheduling Consistency Check

Perform the following steps to schedule a consistency check:

1. In the Controller dashboard, select **More Actions > Schedule Consistency Check**.
   The **Schedule Consistency Check** page appears.
2. Set the **Consistency Check Mode**. The available options are:
   - **Concurrent** - Run consistency check concurrently on all virtual drives.
   - **Sequential** - Run consistency check on one virtual drive at a time.
   - **Disable** - Disables consistency check.
3. Set the desired interval at which you want to run the consistency checks. The available options are:
   - **Hourly**, **Daily**, **Weekly**, **Monthly**, and **Continuously**.
   - Select an appropriate date and time range.
4. Click **Next**.
   The Schedule Consistency Check page appears which allows you to add virtual drives on which you want to perform consistency check.
5. Click **Add Virtual Drives**.
   The Available Virtual Drive dialog appears which lists all the virtual drives present in the selected drive group.
6. Select the Virtual Drive (s) on which you want to run the consistency check.
7. Click **Save**.
   Now consistency check will run based on the frequency/interval chosen by you. You can also monitor the progress of the consistency check operation. See Background Operations Support
8. (Optional) If you want to perform a consistency check operation immediately, from the Controller View section, select the Virtual Drive on which you want to perform a consistency check operation, go to **More Actions**> **Start Consistency Check**.
   If you attempt to run a consistency check on a virtual drive that has not been initialized, a confirmation dialog appears, asking for your confirmation.

15.3 Running Patrol Read

A patrol read periodically verifies all sectors of the drives connected to a controller, including the system reserved area in the RAID configured drives. You can run a patrol read for all RAID levels and for all hot spare drives. A patrol read is initiated only when the controller is idle for a defined period and has no other background activities. You can set the patrol read properties and start the patrol read operation, or you can start the patrol read without changing the properties.

15.3.1 Setting Patrol Read Properties

Perform the following steps to set the patrol read properties.

1. In the Controller dashboard, select **More Actions > Set Patrol Read Properties**.
   The **Available Virtual Drives** dialog appears.
2. Select the virtual drives for which you want to set the patrol read properties and click **Add Virtual Drives**.
   The **Set Patrol Read Properties** dialog appears.

3. Click **Select Virtual Drives**.
   Click the ✕ icon to remove the virtual drives that you have already added.
4. Click **Next**.
5. Perform the following steps to set the properties:
   a. Select an operation mode for patrol read from the **Set Patrol Read Mode** drop-down list. The options follow:
      - **Automatic** – Patrol read runs automatically at the time interval you specify.
      - **Manual** – Patrol read runs only when you manually start it, by selecting Start Patrol Read from the Controller dashboard.
      - **Disabled** – Patrol read does not run.
   b. (Optional) Specify a maximum count of drives to include in the patrol read concurrently. The count must be a number from 1 to 255.
   c. Select the frequency at which the patrol read runs from the drop-down list. The default frequency is weekly (168 hours), which is suitable for most configurations. The other options are hourly, daily, and monthly.
   d. Select the month, day, and year on which to start the patrol read.
   e. Select the time of day to start the patrol read.
   f. (Optional) Select the **Start Patrol Read Now** check box.
   g. (Optional) Select the **Run Patrol Read Non-Stop** check box.
6. Click **Finish**.
   You can monitor the progress of the patrol read operation. See Background Operations Support.

15.3.2 Starting a Patrol Read

Perform the following steps to start a patrol read.

1. In the Controller dashboard, select **More Actions > Start Patrol Read**.
   A warning message appears.
2. Click **Start Patrol Read** to start a patrol read.
   You can monitor the progress of the patrol read operation. See Background Operations Support.

15.3.3 Stopping Patrol Read

Perform the following step to stop a patrol read.

1. In the Controller dashboard, select **More Actions > Stop Patrol Read**.

15.4 Managing SAS Storage Link Speed

The Managing SAS Storage Link Speed feature allows 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. All phys in a SAS port can have different link speeds or can have the same link speed. You can select a link speed setting. However, if phys in a SAS port have different link speed settings and if a phy is connected to a drive or an expander, the firmware overrides the link speed setting you have selected and instead uses the common maximum link speed among all the phys.

Perform the following steps to change the link speed.

1. In the Controller dashboard, select **More Actions > Manage SAS Storage Link Speed**.
   The **Manage SAS Storage Link Speed** dialog appears.

Figure 23 Manage SAS Storage Link Speed

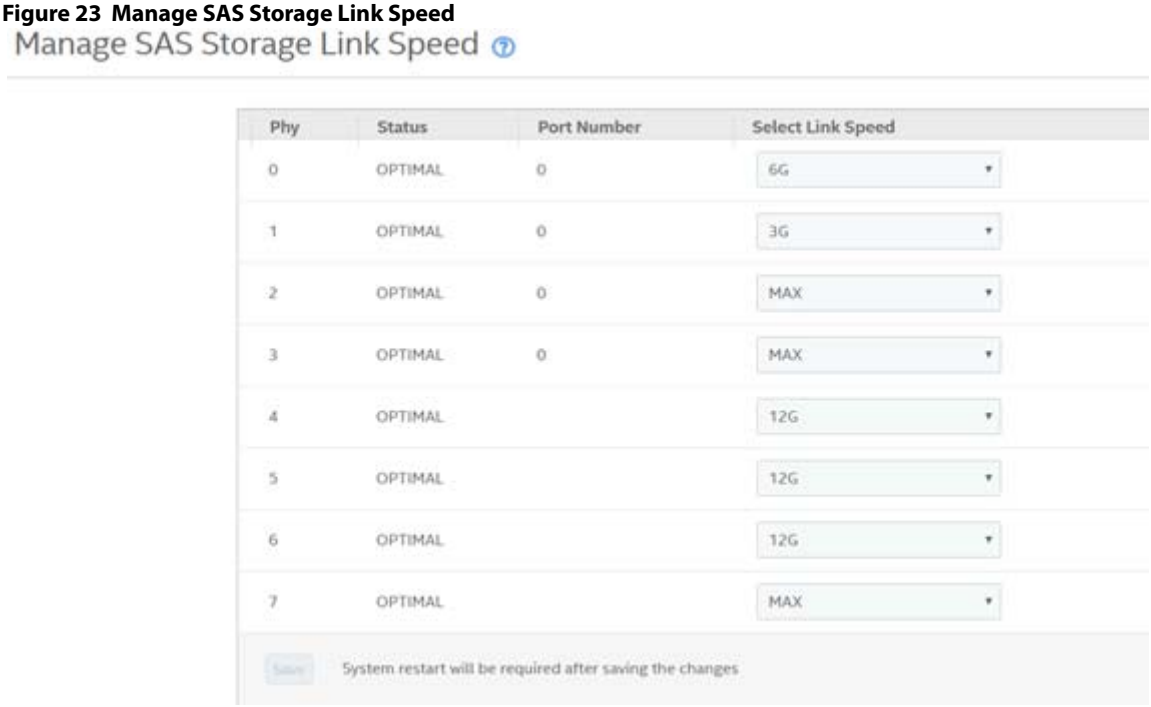

- The **Phy** column displays the system-supported phy link values. The phy link values are from 0 through 7.
- The **Status** column displays the status of the link speed.
- The **Port Number** column displays the port numbers.
- The **Select Link Speed** column displays the phy link speeds.

2. Select the desired link speed from the **Select Link Speed** field using the drop-down selector. The link speed values are **MAX**,**1.5G**, **3G**, **6G**, or **12G**.

   By default, the link speed in the controller is **MAX** or the value last saved by you. The 12G link speed is supported for some SAS-3 expanders.

3. Click **Save**.

   The link speed value is now reset. The change takes place after you restart the system.

15.5 Managing 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 Managing PCIe Storage Interface feature allows you to change the lane speed between a controller and an expander or between the controller and a drive that is directly connected to the controller. LSA 2.4 and later versions support both SAS/SATA topologies as well as PCIe topologies using the same device phys to manage the lane speed.

Perform the following steps to change the lane speed.

1. In the Controller dashboard, select **More Actions > Manage PCIe Storage Interface**.

   The **Manage PCIe Storage Interface** dialog appears.

Figure 24 Manage PCIe Storage Interface

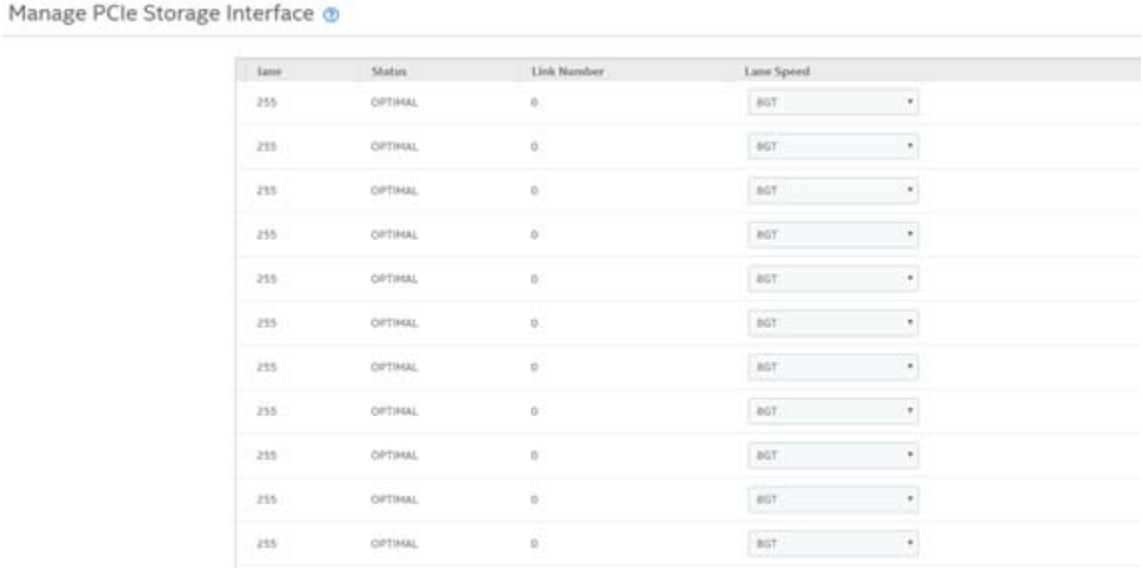

- The **Lane** column displays the system-supported lane values.
- The **Status** column displays the status of the lane.
- The **Link Number** column displays the link numbers.
- The **Lane Speed** column displays the lane speed.

2. Select the desired lane speed from the **Lane Speed** field using the drop-down selector. The lane speed values are **Unknown**,**2.5GT/s**, **5GT/s**, and **8GT/s**.

   By default, the lane speed in the controller is **8 GT** or the value last saved by you.

3. Click **Save**.

   The lane speed value is now reset. The change takes place after you restart the system.

15.6 Setting Adjustable Task Rates

Perform the following steps to set the adjustable task rates.

1. In the Controller dashboard, select **More Actions > Set Adjustable Task Rate**.

   The **Set Adjustable Task Rates** dialog appears.

Figure 25 Set Adjustable Task Rate Dialog

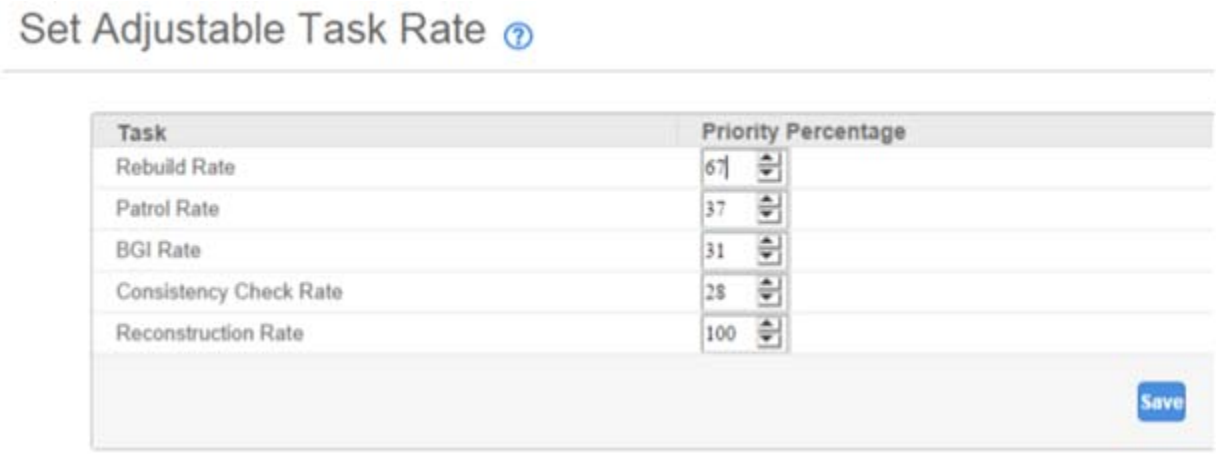

2. Enter changes, as needed, in the following task rates:
   - **Rebuild Rate** – Enter a number from 0 to 100 to control the rate at which a rebuild is performed on a drive when it is necessary. The higher the number, the faster the rebuild will occur (and the system I/O rate might be slower as a result).
   - **Patrol Rate** – Enter a number from 0 to 100 to control the rate at which patrol reads is performed. Patrol read monitors drives to find and resolve potential problems that might cause drive failure. The higher the number, the faster the patrol read will occur (and the system I/O rate might be slower as a result).
   - **Background Initialization (BGI) Rate** – Enter a number from 0 to 100 to control the rate at which virtual drives are initialized in the background. Background initialization establishes mirroring or parity for a RAID virtual drive while allowing full host access to the virtual drive. The higher the number, the faster the initialization will occur (and the system I/O rate might be slower as a result).
   - **Check Consistency Rate** – Enter a number from 0 to 100 to control the rate at which a consistency check is done. A consistency check scans the consistency data on a fault tolerant virtual drive to determine if the data has become corrupted. The higher the number, the faster the consistency check is performed (and the system I/O rate might be slower as a result).
   - **Reconstruction Rate**. Enter a number from 0 to 100 to control the rate at which reconstruction of a virtual drive occurs. The higher the number, the faster the reconstruction occurs (and the system I/O rate might be slower as a result).
3. Click **Save** to set the new task rates.

15.7 Managing Power-Save Settings

Dimmer Switch Technology

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

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

Dimmer Switch 2 – Spin down hot spares. This feature is configurable and can be disabled.

The RAID controller includes Dimmer Switch technology that conserves energy by placing certain unused drives into Power-Save mode. In Power-Save mode, the drives use less energy, and the fan and the enclosure require less energy to cool and house the drives, respectively. Also, this technology helps avoid application timeouts caused by spin-up delays and drive wear caused by excessive spin-up/down cycles.

Perform the following steps to manage the power-save settings.

1. In the Controller dashboard, select **More Actions > Manage Power Save Settings**.
   The **Manage Power Save Settings** dialog appears.

Figure 26 Manage Power Save Settings Dialog

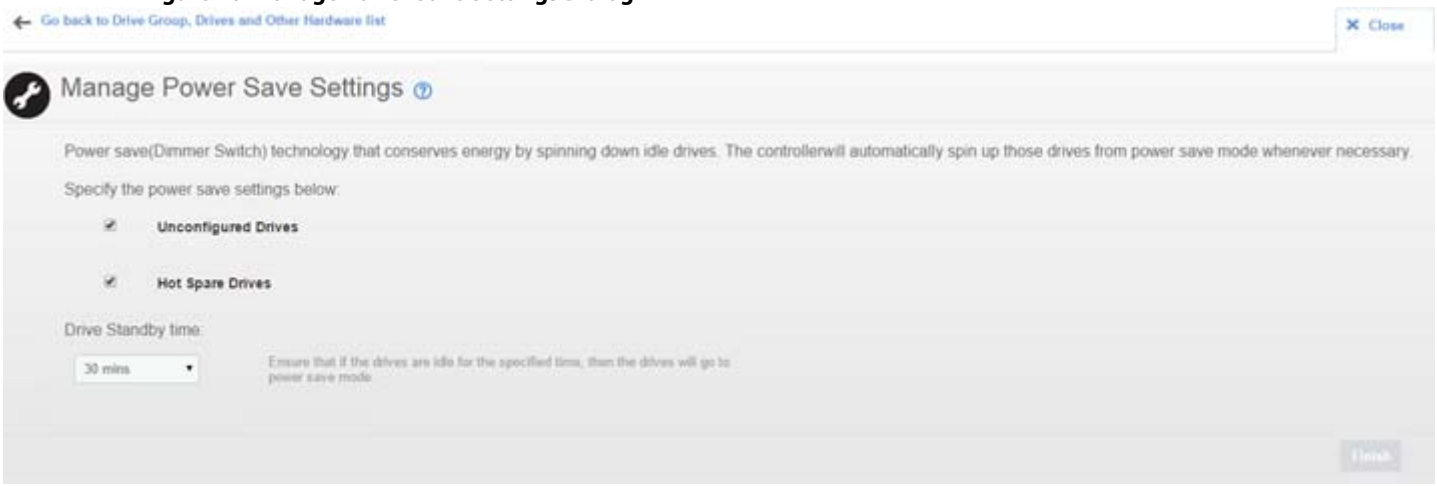

2. Select the **Unconfigured Drives** check box to let the controller enable the unconfigured drives to enter the Power-Save mode.
3. Select the **Hot Spare Drives** check box to let the controller enable the Hot spare drives to enter the Power-Save mode.
4. Select the drive standby time using the drop-down list from the **Drive standby time:** field.
   The **Drive Standby time:** drop-down list is enabled only if any of the check boxes above it are checked. The drive standby time can be 30 minutes, 1 hour, 1.30 hours, or 2 hours through 12 hours.
5. Click **Finish** to save the settings.
   A confirmation message appears.

15.8 Enabling and Disabling SSD Guard

SSDs are known for their reliability and performance. The SSD Guard technology, that is unique to MegaRAID controller cards, 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. A predictive failure event notification, or S.M.A.R.T command, automatically initiates this rebuild to help preserve the data on an SSD whose health or performance falls below par. For RAID volumes that are using CacheCade software, SSD Guard technology can help ensure that the health and performance of SSDs being used for second tier cache are being monitored in the background.

1. In the Controller dashboard, select **More Actions > Enable SSD Guard** to enable the SSD Guard feature.
2. To disable the SSD Guard feature, select **More Actions > Disable SSD Guard**.

15.9 Discarding Pinned Cache

If the controller loses access to one or more virtual drives, the controller preserves the data from the virtual drive. This preserved cache is called as pinned cache. This cache is preserved until you import the virtual drive or discard the cache. As long as there is pinned cache, you cannot perform certain operations on the virtual drive.

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

Perform the following steps to discard the pinned cache.

1. In the Controller dashboard, select **More Actions > Discard Preserved Cache**.

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

   A message appears, prompting you to confirm your choice.
2. Select **Confirm** and click **Yes, Discard**.

15.10 Downloading TTY Log

You can download TTY log file, which contains the firmware terminal log entries for the controller. The log information is shown as total number of entries available on the firmware side. Perform the following steps to download the TTY log file.

1. In the Controller dashboard, select **More Actions > Download TTY Log**.

   The `tty.log` file is downloaded.

15.11 Updating the Controller Firmware

The LSI Storage Authority software enables you to update the controller firmware.

Perform the following steps to update the controller firmware.

1. Navigate to the Controller dashboard.
2. Click **Update Firmware**.

   The **Update Firmware** window appears. It displays the current controller firmware version.
3. Click **Browse** to locate and open the:
   - `.rom` update file for MR controllers.
   - `.bin` update file for IR3/IT controllers.
4. Click **Update**.

   The **Update Firmware** windows displays the current controller firmware version and the controller firmware version that is selected to be flashed.

Figure 27 Update Firmware Window-Selected Version

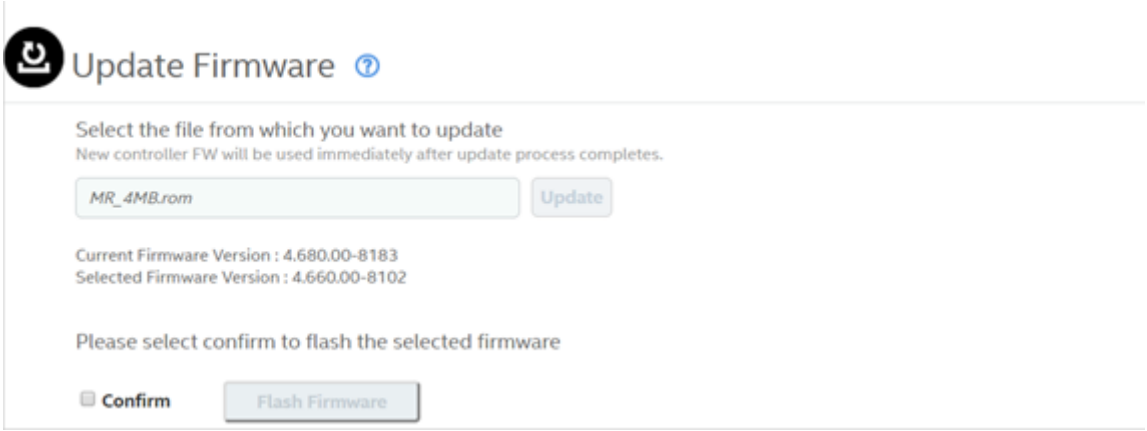

5. Select the **Confirm** check box and click **Flash Firmware**.

   After the update is complete, a message displays that confirms the success of the update and displays the new version of the controller firmware.

Chapter 16: MegaRAID Advanced Software

The MegaRAID advanced software (Premium) are features that the LSI Storage Authority software supports on certain 12Gb/s RAID controllers.

The MegaRAID advanced software includes the following features:

- MegaRAID FastPath
- MegaRAID CacheCade SSD Read Caching software
- MegaRAID CacheCade Pro 2.0 SSD Read/Write Caching software
- MegaRAID SafeStore

The MegaRAID software licensing authorizes you to enable the MegaRAID advanced software features. You have to obtain the activation key to enable, and use the advanced software features present in the controller.

16.1 Activating MegaRAID Advanced Software

The **Premium Features** window allows you to use the advanced software features.

Perform the following steps to enable the activation key to use the advanced controller features:

1. In the Controller dashboard, select **Actions > Premium Features**.
   The **Premium Features** window opens.

Figure 28 Premium Features Window

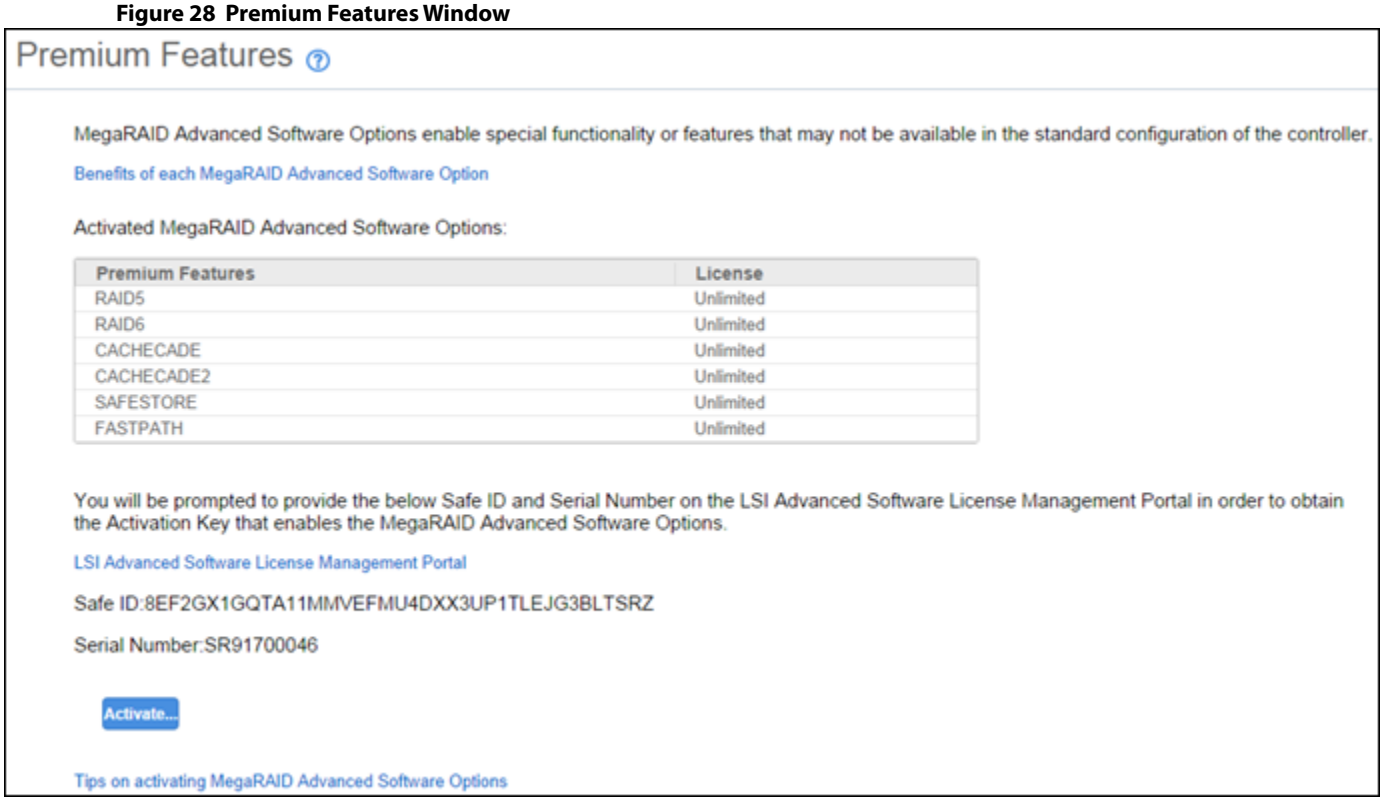

The **Activated MegaRAID Advanced Software Options:** table consists of the **Premium Features** and the **License** columns.

- The **Premium Features** column displays the list of advanced software options present in the controller.
- The **License** column displays the license details for the list of advanced software options present in the **Advanced Software Option** column. The license details validate if the software is under a trial period, or if it can be used without any trial period (Unlimited).

For more information on the benefits of these features, click the **Benefits of each MegaRAID Advanced Software** link.

2. Click the **LSI Advanced Software License Management Portal** link to obtain the license authorization code and the activation key.

   Both the **Safe ID** field and the **Serial Number** field consists of a pre-defined value generated by the controller.

   For more information on activating the advanced software options, click the **Tips on activating MegaRAID Advanced Software Options** link.

3. Click **Activate**.

   The **Activate Features** window appears.

Figure 29 Activate Features Window

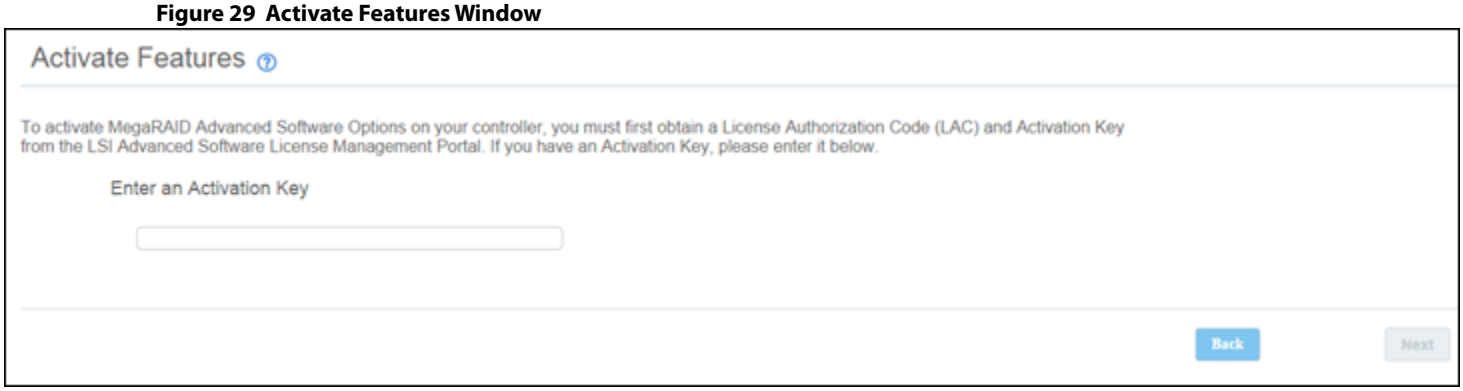

4. Enter the activation key in the text box provided.
5. Click **Next**.

   After you click **Next**, one of the following two scenarios occurs:

   — Depending on whether you are activating an unlimited key or a trial key, the relevant **Activate Features – Summary** dialog appears. See Advanced MegaRAID Software Status Summary.

   — If you have entered an invalid key or if there is a key mismatch, relevant error messages are shown. See Application Scenarios and Messages.

16.1.1 Advanced MegaRAID Software Status Summary

After you enter the activation key and click **Next**, the **Activate Features** window appears as shown in the following figure. It displays the list of the advanced software features along with their *former status* and *new status* in the controller.

Figure 30 Activate Features – Summary

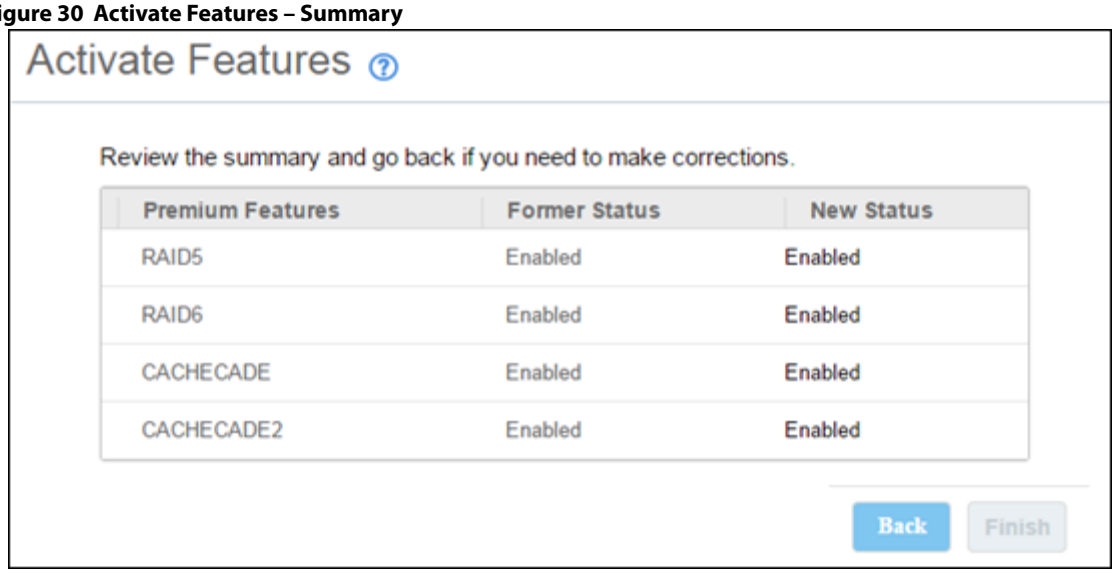

- The **Premium Features** column displays the currently available software in the controller.
- The **Former Status** column displays the status of the available advanced software before entering the activation key.
- The **New Status** column displays the status of the available advanced software, after entering the activation key.

1. Click **Finish**.

   The status of the advanced software is enabled, and the advanced features are secured in the Key Vault.
2. Click **Back** to return to the previous window to change any selections.

16.1.1.1 Activating a Trial Key

When you activate a trial key, a message `This trial software expires in 30 days.` appears.

Figure 31 Activating a Trial Software

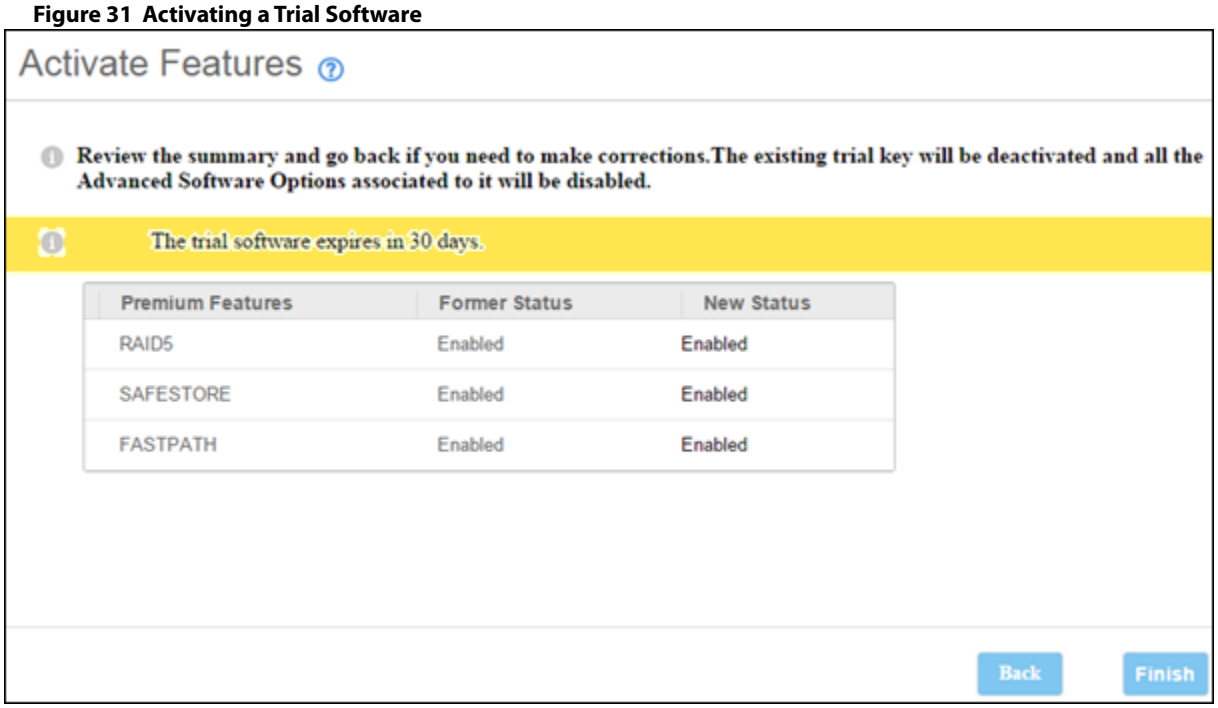

16.1.1.2 Activating an Unlimited Key over a Trial Key

When you activate an unlimited key over a trial key, a message, `The existing trial key will be deactivated and all the Advanced Software Options associated to it will be disabled.` appears.

Figure 32 Activating an Unlimited Key over a Trial Key

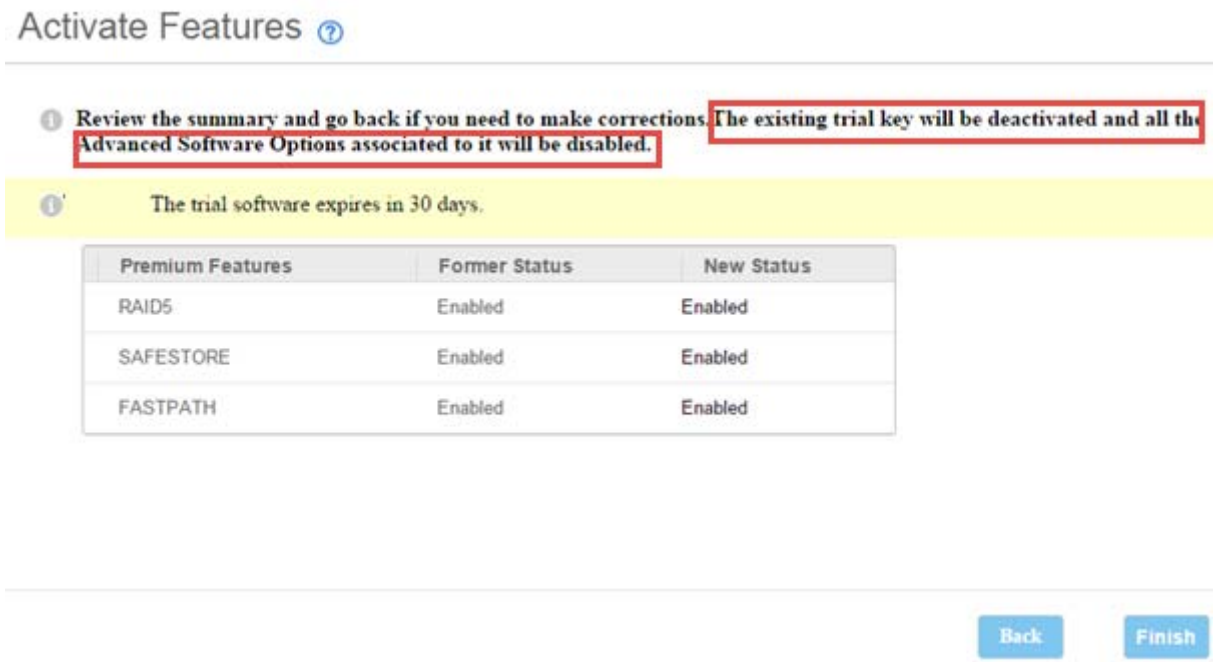

16.1.1.3 Reusing the Activation Key

If you are using an existing activated key, the features are transferred to the key vault, and a message appears.

Figure 33 Reusing the Activation Key

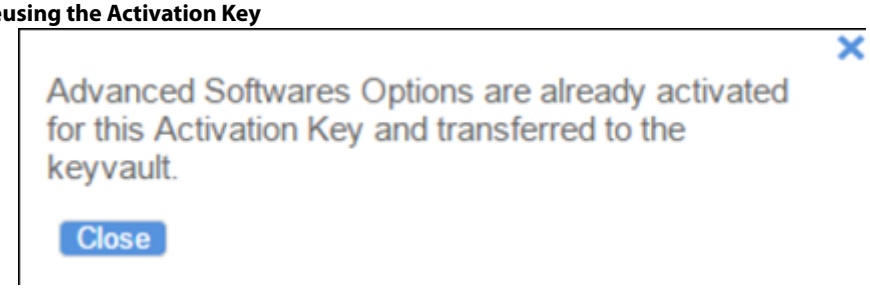

16.1.1.4 Application Scenarios and Messages

Scenario # 1

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

Figure 34 Invalid Activation Key Message

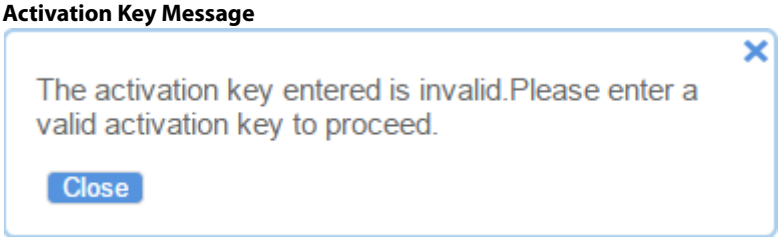

Scenario # 2

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

Figure 35 Activation Key Mismatch Message

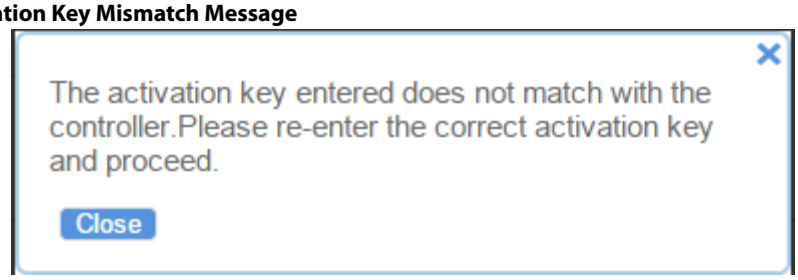

16.2 Securing Advanced MegaRAID Software

You can transfer the advanced software from the controller to the key vault. This feature is conditional, and appears only when the key vault and the unsecured keys exist.

Perform the following steps to secure the advanced MegaRAID software.

1. In the **Premium Features** window, click **Configure Key Vault**.
   The **Activate Features** window opens.

Figure 36 Activate Features– Secure Key Vault Option

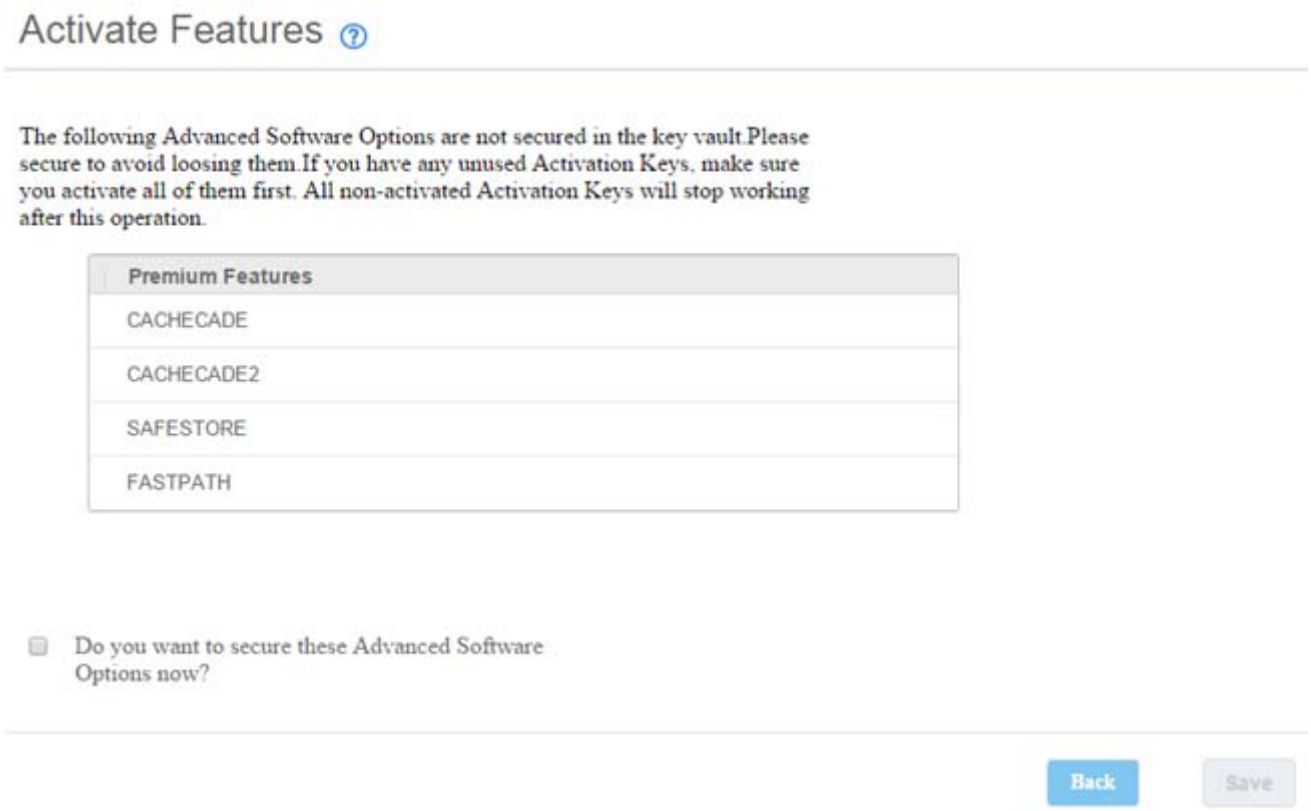

2. Select the **Do you want to secure these Advanced Software Options now?** check box, if you want to secure the advanced software.

   After you select the check box, the **Save** button is enabled. This situation implies that the advanced software is secured in the key vault.

16.3 Configuring Key Vault (Re-hosting Process)

Re-hosting is a process of transferring the advanced software features from one controller to another. This feature is conditional and appears only if the re-hosting process is necessary, and when both the key vault and the unsecured keys are present at the same time.To implement the re-hosting process, perform the following steps.

1. In the **Premium Features** window, click **Configure Key Vault**.

   The following window appears.

Figure 37 Premium Features – Configure Key Vault

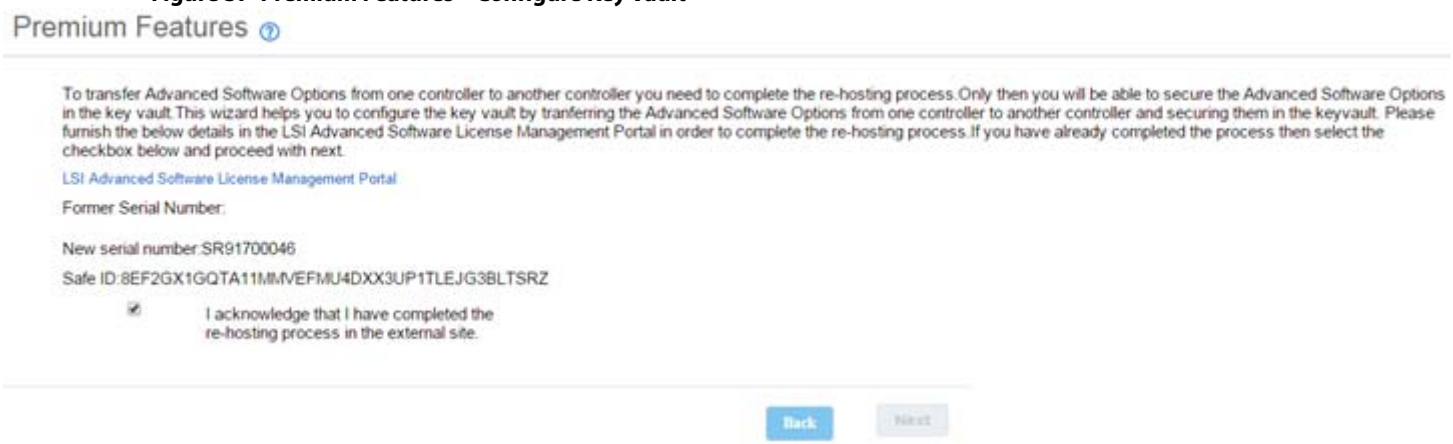

2. Select the **I acknowledge that I have completed the re-hosting process in the external site.** check box.
3. Click **Next**.
   The **Next** button in the screen is enabled only if you select the check box.
   The **Activate Features** window appears.

Figure 38 Activate Features – Configure Key Vault Window

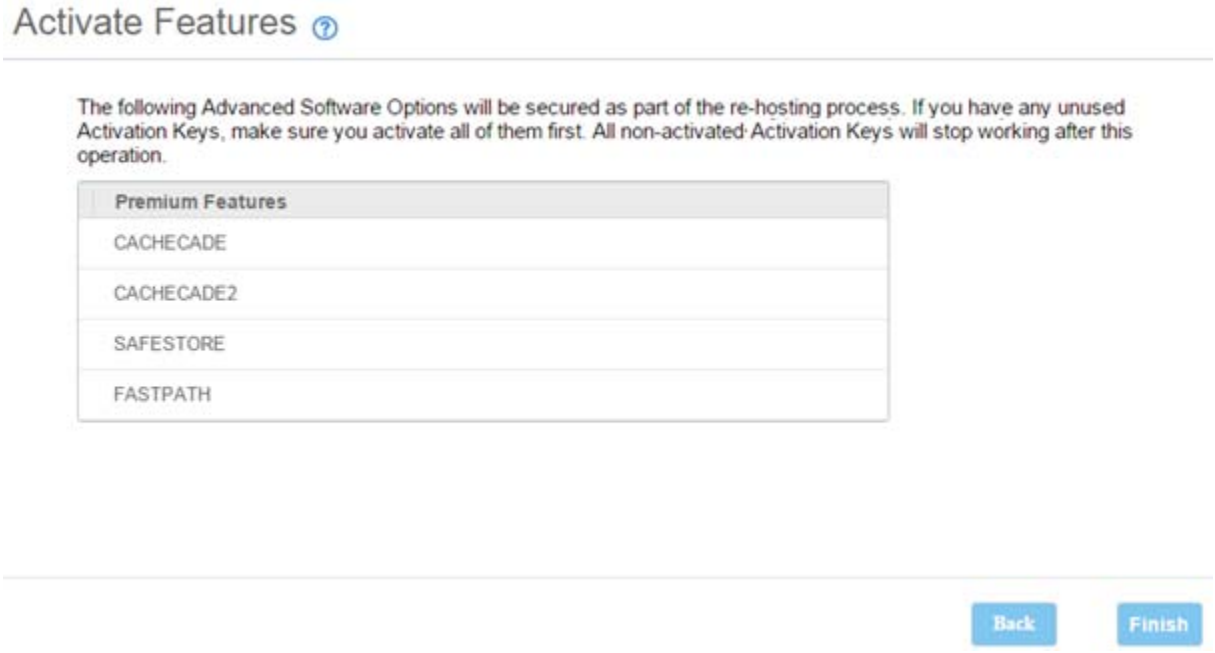

4. Click **Finish**, and the advanced software options are secured in the key vault.

16.4 Re-hosting Complete

If you want to transfer the advanced software options from one controller to another, use the re-hosting process. The re-hosting process makes sure that these options are secured in the Key Vault. You have to configure the Key Vault to complete the re-hosting process. To implement the re-hosting process, perform the following steps.

1. In the **Premium Features** window, click **Configure Key Vault**.
   The following window appears.

Figure 39 Premium Features Window – Re-hosting Complete

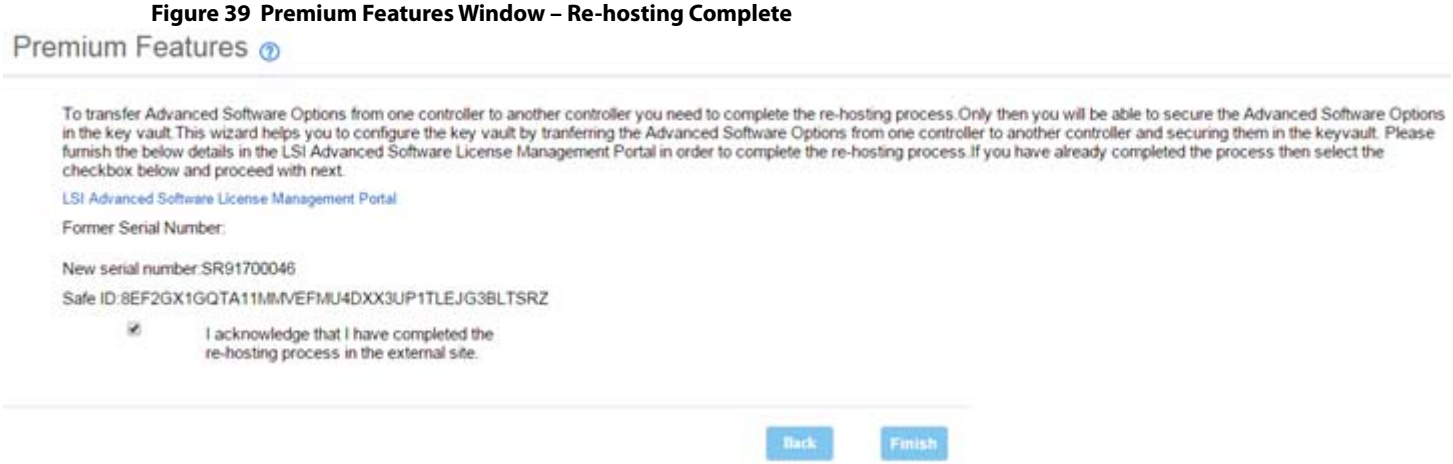

2. Select the **I acknowledge that I have completed the re-hosting process in the external site.** check box.
   This setting makes sure that the advanced software features are transferred to the controller.
3. Click **Finish** and the advanced software options are secured in the key vault.
   Click **Cancel** if you do not want to activate the re-hosting process.

16.5 Deactivating Trial Software

When you want to deactivate a trial software, use the **Deactivate All Trial Software** wizard.

Perform the following steps to enable the deactivate trial software button:

1. Click **Deactivate All Trial Software** in the **Premium Features** window.
   A confirmation dialog appears.

Figure 40 Deactivate All Trial Software - Confirmation Dialog

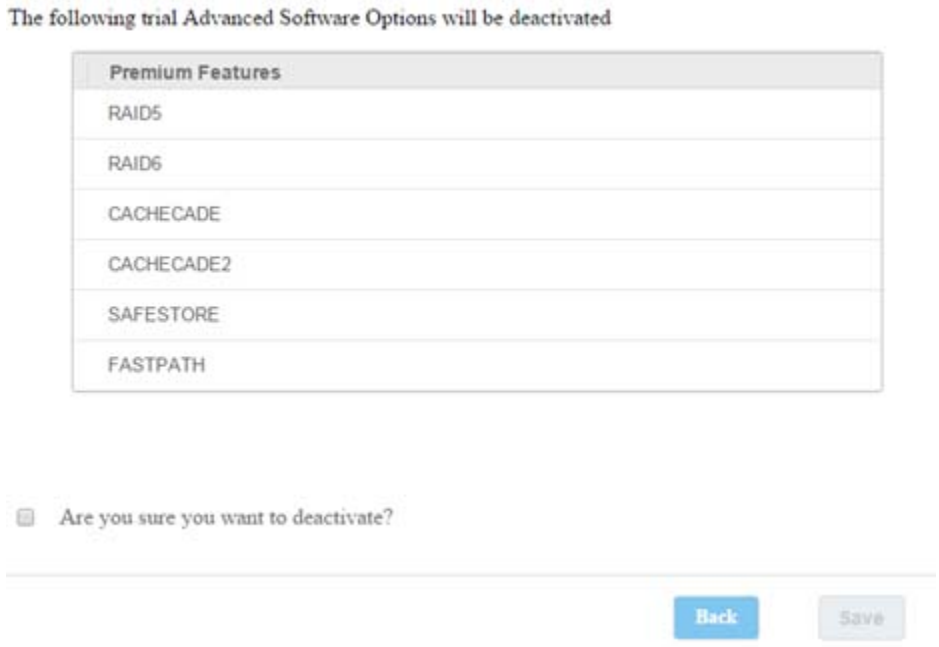

2. Select the **Are you sure you want to deactivate?** check box, if you want to deactivate the software applications, that are used with a trial key.
3. Click **Save**.

   The trial software is deactivated.

16.6 Using the MegaRAID CacheCade Pro 2.0 Feature

The MegaRAID CacheCade Pro 2.0 read and write software eliminates the need for manually configured hybrid arrays by intelligently and dynamically managing frequently-accessed data and copying it from HDD volumes to a higher performance layer of SSD cache. Copying the most accessed data (hot spot) to flash cache relieves the primary HDD array from time-consuming transactions, which allows for more efficient hard disk operation, reduced latency, and accelerated read and write speeds. CacheCade Pro 2.0 software is the industry's first software solution that offers both read and write controller-based caching on SSDs, dramatically enhancing the performance gains achieved by the previous generation CacheCade software. With the addition of write caching support, read/write-intensive workloads such as Exchange server, high performance computing (HPC) applications, Web 2.0 and other IO-intensive OLTP database system workloads, experience dramatic performance improvements.

16.6.1 Creating a CacheCade Virtual Drive

Perform the following steps to create a CacheCade virtual drive.

1. In the Server dashboard or the Controller dashboard, select **Configure > CacheCade - SSD Caching Configuration**.

   The **CacheCade - SSD Caching Configuration** window opens.

Figure 41 CacheCade - SSD Caching Configuration Window

2. Select a RAID level for the drive group. For example, select **RAID 0**.
3. (Optional) click **Compare and Select** to view the detailed information on each RAID level.
4. Select the **Encryption** check box if you want to apply the encryption logic to secure the data in the virtual drive.
5. Click **Next**.
6. Click **Add SSD Physical Drives** to add SSD drives to the drive group.

   The **Available Unconfigured SSD Drives** window appears.

Figure 42 Available Unconfigured SSD Drive Window

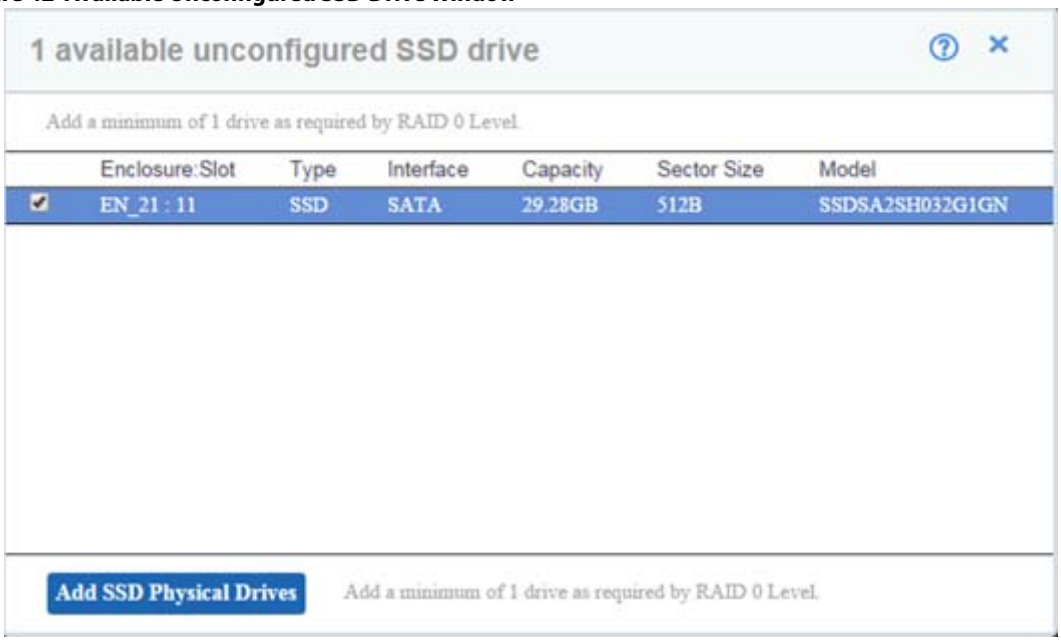

7. Select the SSD physical drives and click **Add SSD Physical Drives**.
8. Click **Add CacheCade - SSD Caching Virtual Drives** to add CacheCade virtual drives to the drive group. The **CacheCade - SSD Caching Virtual Drive Settings** window appears.

Figure 43 CacheCade - SSD Caching Virtual Drive Settings Window

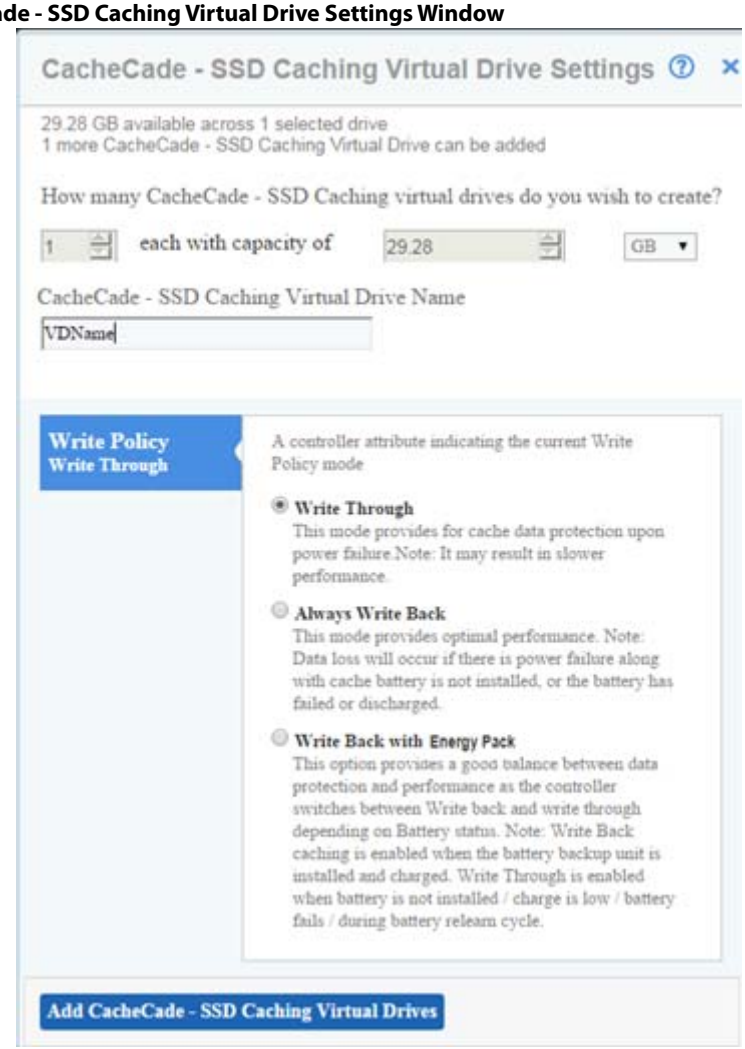

NOTE You can create only one CacheCade-SSD Caching virtual drive as the full capacity of the virtual drive is used for the creation of CacheCade-SSD Caching virtual drive.

9. Enter a name for the CacheCade - SSD Caching virtual drive in the **CacheCade - SSD Caching Virtual Drive Name** field.

   The virtual drive name can have a maximum of 15 characters.

10. Specify the write policy for the CacheCade - SSD Caching virtual drives. The write policy depends on the status of the Energy Pack. If the Energy Pack is not present, is low, is failed, or is being charged, the current write policy switches to write through.

    The options follow:

    — **Always Write Back**
    — **Write Back with Energy Pack**
    — **Write Through**

11. Click **Add CacheCade - SSD Caching Virtual Drives**.

    The newly created CacheCade - SSD Caching virtual drives appears in the **CacheCade - SSD Caching Configuration** window just below the **Add CacheCade - SSD Caching Virtual Drives** section.

    If you want to modify the CacheCade - SSD Caching virtual drives settings before finishing the configuration, click the icon.

12. Click **Finish**.

    A message appears stating that the configuration is complete.

16.6.2 Modifying CacheCade Virtual Drive Properties

You can modify the name and the write policy of a CacheCade - SSD Caching Virtual drive any time after a CacheCade - SSD Caching Virtual drive is created. Perform the following steps to change the virtual drive properties:

1. Navigate to the Controller dashboard, click a drive group name (for example, **DG_1**). Click the icon corresponding to a drive group to display its contents.

   The virtual drives and physical drives associated with the selected drive group appear.

2. Click a CacheCade - SSD Caching Virtual drive whose settings you want to change.
3. Select **Actions > More Actions > Modify Properties**.

   The **Modify Virtual Drive: <Virtual Drive Name> Properties** dialog appears.

Figure 44 SSD Caching Virtual Drive - VDName Properties Dialog

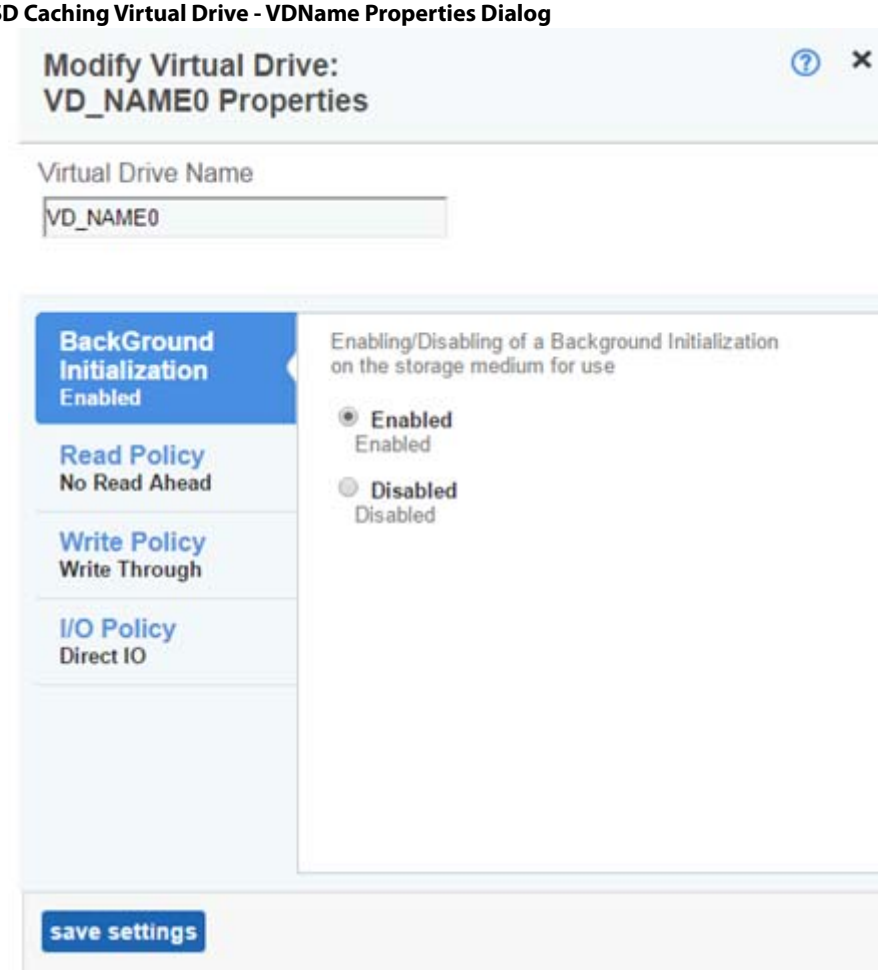

4. Change the **CacheCade - SSD Caching Virtual Drive Name** and the **Write Policy** properties as needed.
5. Click **Save Settings**.

16.6.3 Enabling SSD Caching on a Virtual Drive

You can enable SSD caching on a virtual drive. When you enable SSD caching on a virtual drive, that virtual drive becomes associated with an existing or with a future CacheCade - SSD Caching virtual drive. This option is only available when the virtual drive's caching is currently disabled.

1. Navigate to the Controller dashboard, click a drive group name (for example, **DG_1**). Click the + icon corresponding to a drive group to display its contents.

   The virtual drives and physical drives associated with the selected drive group appear.
2. Click the virtual drive on which you want to enable SSD caching.
3. Select **Actions > Enable SSD Caching**.

   The following dialog appears.

Figure 45 Enable SSD Caching

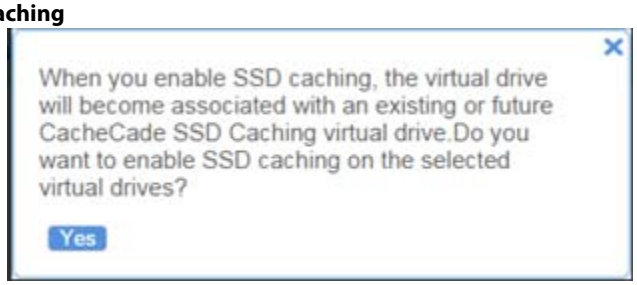

4. Click **Yes**.

   A confirmation message appears.

16.6.4 Disabling SSD Caching on a Virtual Drive

You can disable caching on a virtual drive. When you disable SSD caching on a virtual drive, any associations that the selected virtual drive has with a CacheCade - SSD Caching virtual drive is removed. This option is only available when the virtual drive's caching is currently enabled.

1. Navigate to the Controller dashboard, click a drive group name (for example, **DG_1**). Click the + icon corresponding to a drive group to display its contents.

   The virtual drives and physical drives associated with the selected drive group appear.
2. Click the virtual drive on which you want to disable SSD caching.
3. Select **Actions > Disable SSD Caching**.

   The following dialog appears.

Figure 46 Disable SSD Caching

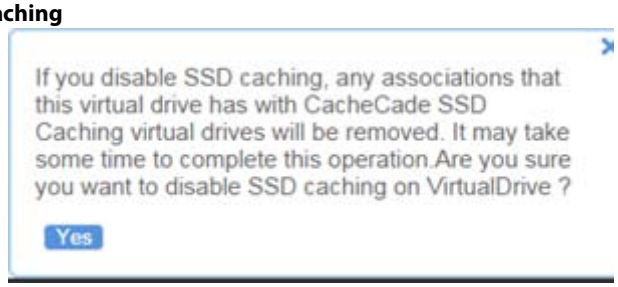

4. Click **Yes**.

   A confirmation message appears.

16.6.5 Clearing Configuration on Controllers that Have CacheCade Virtual Drives

You can clear all existing configurations on a selected controller that has CacheCade Pro 2.0 virtual drives.

1. Navigate to the Controller dashboard whose configurations you want to clear.
2. Click **Configure** and then click **Clear Configuration**.

   The following confirmation message appears.

Figure 47 Clear Configuration - CacheCade - SSD Caching

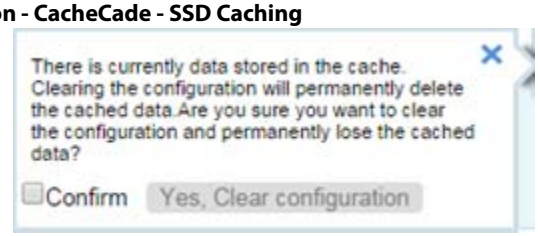

3. Select **Confirm** and click **Yes, Clear configuration** to clear all the existing configurations on the controller.
   However, the operating system drives cannot be cleared.

16.6.6 Deleting a CacheCade - SSD Caching Virtual Drive

Perform the following steps to delete a CacheCade - SSD Caching virtual drive.

1. Navigate to the Controller dashboard, click a drive group name (for example, **DG_1**). Click the + icon corresponding to a drive group to display its contents.
   The virtual drives and physical drives associated with the selected drive group appear.
2. Click the CacheCade - SSD Caching virtual drive that you want to delete.
3. Select **Actions > Delete**.
   The following confirmation message appears.

Figure 48 CacheCade - SSD Caching Virtual Drive - Delete Confirmation

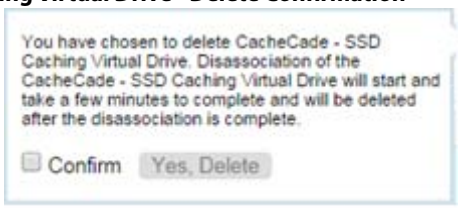

4. Select **Confirm** and click **Yes, Delete** to proceed with the delete operation.
   A message appears confirming that the CacheCade - SSD Caching virtual drive is deleted successfully.

16.7 MegaRAID Fast Path Advanced Software

The MegaRAID FastPath software is a high-performance I/O accelerator for Solid State Drive (SSD) arrays connected to a MegaRAID controller card. This advanced software is an optimized version of Broadcom MegaRAID technology that can dramatically boost storage subsystem and overall application performance. Particularly those that demonstrate high random read/write operation workloads – when deployed with a Broadcom MegaRAID SATA+SAS controller connected to SSDs.

16.8 MegaRAID SafeStore Encryption Services

The MegaRAID SafeStore software, together with self-encrypting drives (SEDs), secures a drive's data from unauthorized access or modification resulting from theft, loss, or repurposing of drives. If you remove a self-encrypting drive from its storage system or the server in which it resides, the data on that drive is encrypted, and becomes useless to anyone who attempts to access it without the appropriate security authorization.

Auto Lock with Local Key Management locks the SED using an authentication key. When secured in this manner, the drive's data encryption key is locked whenever the drive is powered down. In other words, the moment the SED is switched off or unplugged, it automatically locks down the drive's data. When the drive is powered back on, it requires authentication before being able to unlock its encryption key and read any data on the drive. This action protects against any type of insider or external theft of drives or systems.

Instant Secure Erase feature allows you to instantly and securely render data on SED drives unreadable, saving businesses time and money by simplifying decommissioning of drives and preserving hardware value for returns and repurposing.

You can enable, change, and disable the drive security feature. You can also import a foreign configuration using the SafeStore Encryption Services advanced software.

16.8.1 Enabling Drive Security

Ensure that MFC settings related to security are enabled in the firmware.

Perform the following steps to enable security on the drives.

1. In the Controller dashboard, select **More Actions > Enable Drive Security**.

   The **Enable Drive Security** dialog appears.

Figure 49 Enable Drive Security Dialog

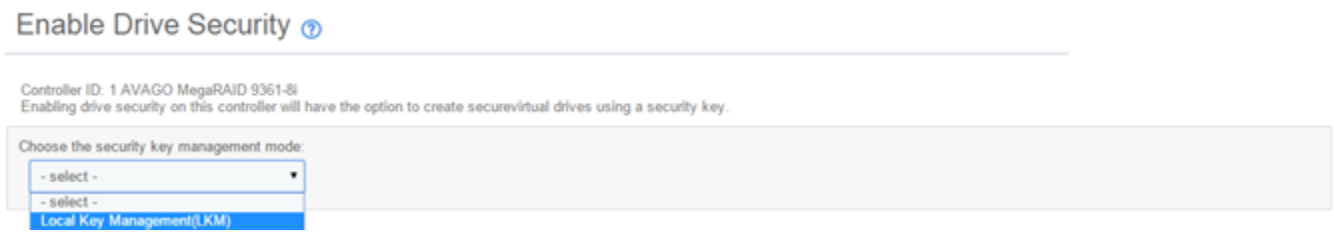

2. Select the **Local Key Management (LKM)** option from the **Choose the security key management mode** drop-down list.

   The **Enable Drive Security** dialog appears with the following options that lets you enable the drive security.

Figure 50 Enable Drive Security

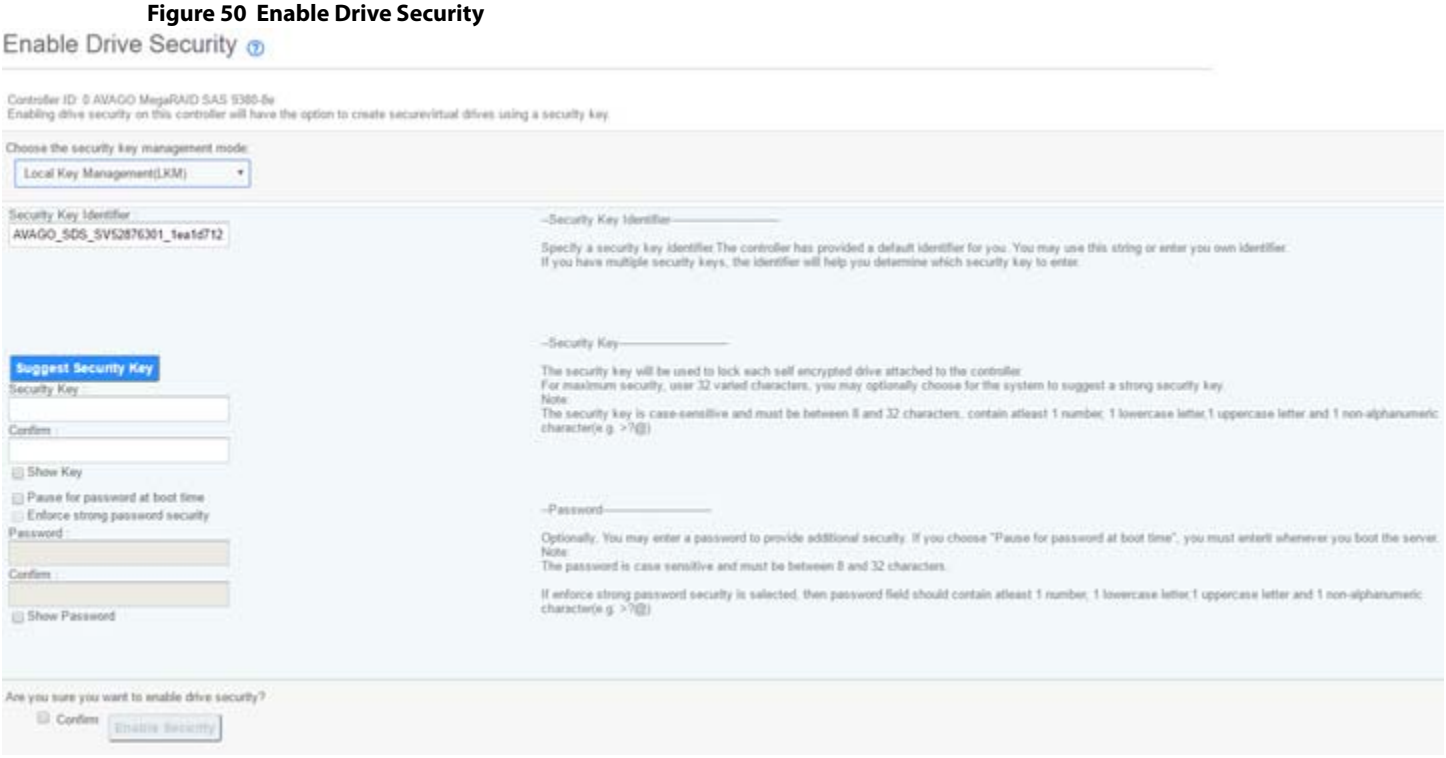

To enable drive security, the following details must be specified:

- **Security key identifier** - The controller, by default, assigns a security key identifier. However, you can change this security key identifier as per your requirement. If you have more than one security key identifier, the controller helps you to determine which security key identifier to enter.
- **Security key** - Provides you with an option to create secure virtual drives by specifying the security key. The security key provided by you is used to lock each self-encrypted drive attached to the controller.
- **Suggest Security Key** - Alternatively, you can click this option to have the system create a security key for you.
- **Password** - You can also specify a password to provide additional drive security.
- **Pause for password at boot time** and **Enforce strong password security** - If you select the **Pause for password at boot time**, you are prompted to provide the password each time you restart your server. If you select **Enforce strong password security**, the system enforces you to specify a strong password.
- **Show Key and Show Password** - You can either select or clear the **Show Key** and **Show Password** check boxes. By default, they are unchecked.

To enable drive security, perform the following steps:

3. Either use the default security key identifier provided by the controller or specify a new security key identifier.

   NOTE If you create more than one security key, ensure that you change the security key identifier. Otherwise, you cannot differentiate between the security keys.

4. Either click **Suggest Security Key** to have the system create a security key for you or enter a new security key in the **Security Key:** field and confirm.
5. (Optional) Select the **Show Key** check box.

   If you choose this option, the security key that you specify or the security key that is created by the system if you have clicked on Suggest Security Key, will be visible to you. If you do not select this option, the security key will not be visible to you.

NOTE **Ensure that you note down this security key somewhere for future reference. If you are unable to provide the security key when it is required by the system, you will lose access to your data.**

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

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

6. (Optional) Select the **Pause for password at boot time** check box.

   If you choose this option, you are prompted to provide the password each time you restart your server.

7. (Optional) Select the **Enforce strong password security** check box.

   If you choose this option, make sure the password is between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g. < > @ +). The space character is not permitted. The password is case-sensitive.

8. (Optional) Enter a password in the **Password** field and confirm the same password once again in the **Confirm** field.

9. (Optional) Select the **Show Password** check box.

   If you choose this option, the password that you specify will be visible to you. If you do not select this option, the password will not be visible to you.

   Warning messages appear if there is a mismatch between the characters entered in the **Password** field and the **Confirm** field, or if you have entered an invalid character.

CAUTION **Ensure that you note down this password somewhere for future reference. If you are unable to provide the password when it is required by the system, you will lose access to your data.**

10. Select the **Confirm** check box, then click **Enable Security** to confirm that you want to enable drive security on this controller.

16.8.2 Changing Security Settings

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

1. In the Controller dashboard, select **More Actions > Change Drive Security**.

   The **Change Drive Security** dialog appears.

Figure 51 Change Drive Security Dialog

2. Select the **Change current security settings** radio button from the **Current drive security mode is LKM** field.

   The following options appear. The options list the actions you can perform, which include editing the security key identifier, security key, and the password.

Figure 52 Change Drive Security Options

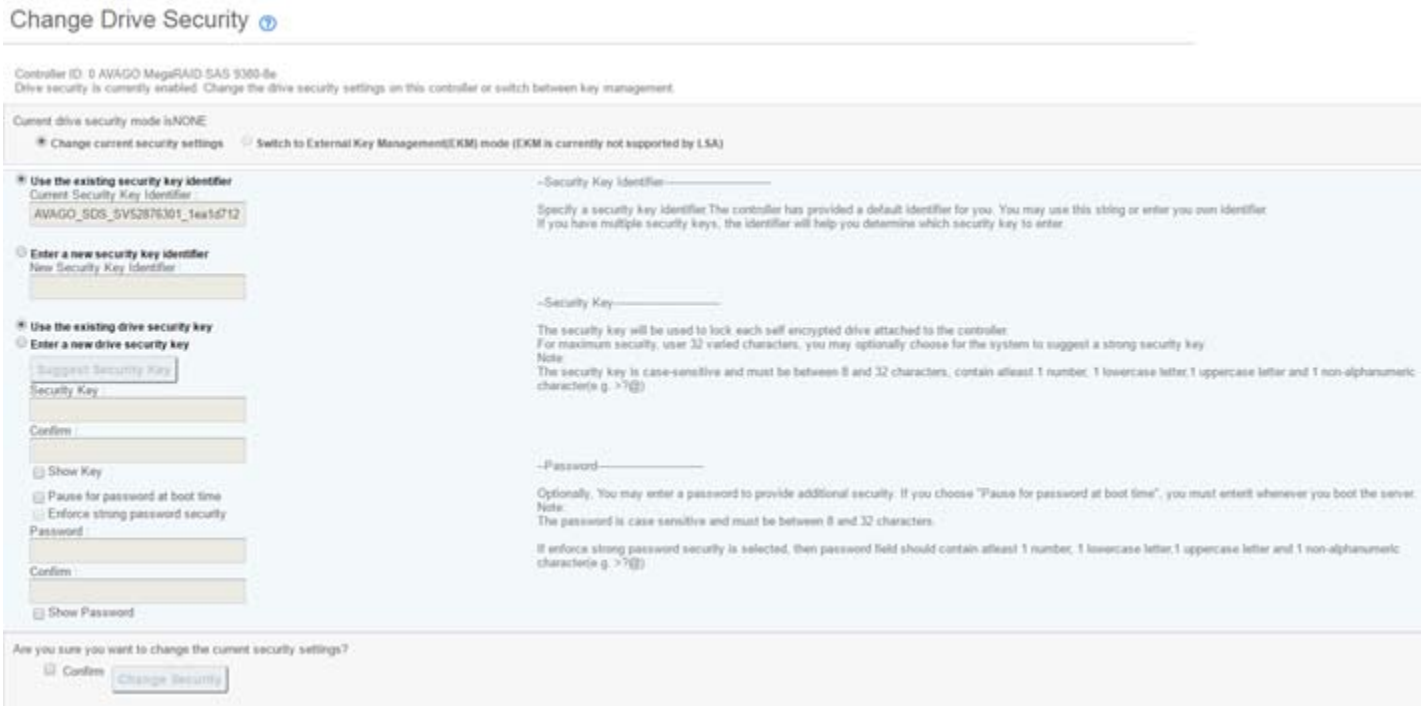

3. Either you can use the existing security key identifier assigned by the controller or specify a new security key identifier.

   If you change the security key, you need to change the security key identifier. Otherwise, you cannot differentiate between the security keys.
4. Either select the **Use the existing drive security key** option or select the **Enter a new drive security key** to specify a new security key and confirm once again.
5. Either click **Suggest Security Key** to have the system create a security key, or you can enter a new security key in the **Security Key:** text field.
6. (Optional) Select the **Show Key** checkbox.

   NOTE The security key is case-sensitive. It must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (for example, < > @ +).
7. (Optional) Select the **Pause for password** at boot time check box.

   If you choose this option, you are prompted to provide the password each time you restart your server.
8. (Optional) Select the **Enforce strong password security** check box.

   If you choose this option, make sure the password is between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (example, < > @ +). The space character is not permitted. The password is case-sensitive
9. If you chose to use a password, either enter the existing password or enter a new password, and confirm once again.
10. (Optional) Select the **Show Password** checkbox.

    If you choose this option, the password that you specify will be visible to you. If you do not select this option, the password will not be visible to you.

11. Select the **Confirm** checkbox and click **Change Security** to change the security settings.

    The **Authenticate Drive Security Settings** dialog appears. Your authentication is required for the changes to take effect. Enter the new security key that you just specified in the Security Key field

12. Enter the new security key that you just specified and click **Authenticate** to authenticate the changes.

    The existing configuration on the controller is updated to use the new security settings.

16.8.3 Disabling Drive Security

ATTENTION If you disable drive security, your existing data is not secure and you cannot create any new secure virtual drives. Disabling drive security does not affect the security of data on foreign drives. If you have removed any drives that were previously secured, you still need to enter the password when you import them. Otherwise, you cannot access the data on those drives. If there are any secure drive groups on the controller, you cannot disable drive security. A warning dialog appears if you attempt to do so. To disable drive security, you must first delete the virtual drives on all of the secure drive groups.

Perform the following steps to disable drive security:

1. In the Controller dashboard, select **More Actions > Disable Drive Security**.

   A warning message appears asking for your confirmation.

2. Select **Confirm** and click **Yes, Disable Drive Security**.

   The software disables drive security.

16.8.4 Importing or Clearing a Foreign Configuration - Security Enabled Drives

Perform the following steps to import or clear foreign configuration for security enabled drives.

1. Enable drive security to allow importation of security enabled foreign drives.
2. After you create a security key, navigate to the Controller dashboard and click **Configure** and then click **Foreign Configuration**.

   If locked drives (security is enabled) exist, the **Unlock Foreign Drives** dialog appears.

3. Enter the security key to unlock the configuration.

   The **Foreign Configuration** window appears, which lists all of the foreign configurations.

4. Click one of the following options:
   - **Import All**: Import the foreign configurations from all the foreign drives.
   - **Clear All**: Remove the configurations from all the foreign drives.

5. Click **Re-Scan** to refresh the window.
6. Repeat the import process for any remaining drives because locked drives can use different security key, and you must verify whether there are any remaining drives to be imported.

Chapter 17: Managing Drive Groups

The LSI Storage Authority software allows you to monitor the status of the drive groups and spanned drive groups.

17.1 Viewing Drive Group Properties

Select a drive group in the Controller dashboard to view its properties.

The following figure and table describe the Drive Group properties.

Figure 53 Drive Group Properties

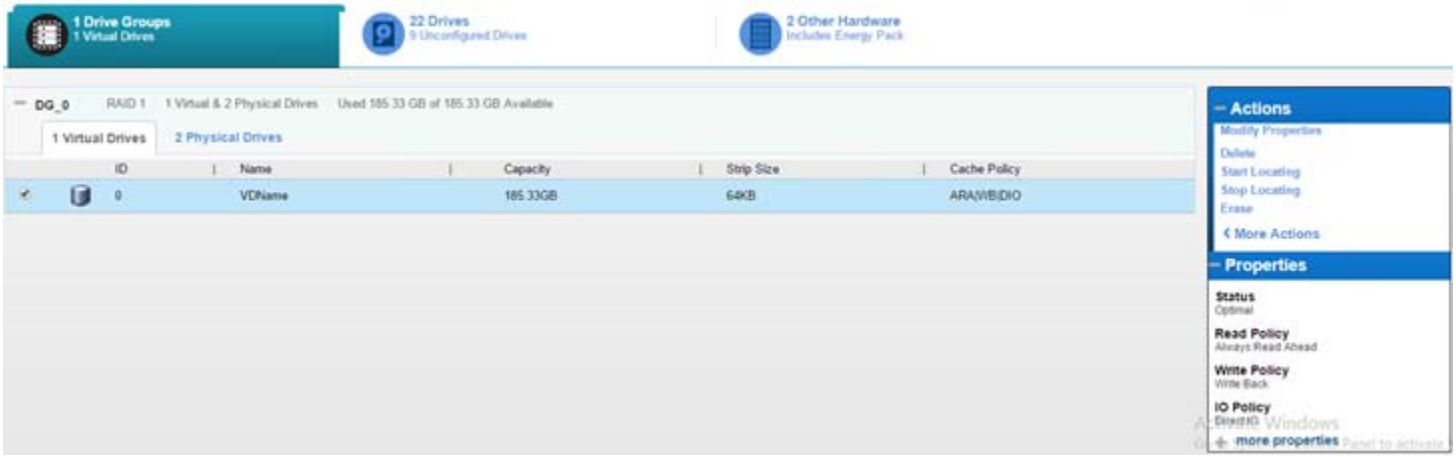

If you have selected multiple virtual drives or multiple physical drives, you have to click on the (Expand button) to perform actions such as starting a a consistency check and so on. This is applicable for all the scenarios where you have selected multiple virtual drives or multiple physical drives and performing certain actions through the **Actions** dialog.

Table 11 Drive Group Properties

Property	Description	MegaRAID	iMegaRAID	Initiator-Target
Data Protection	Indicates if the data protection feature is enabled for the drive group.	Yes	Yes	Yes
Free Capacity	Indicates the free space available in the drive group.	Yes	Yes	Yes
Secured	Indicates if the drive group is secured.	Yes	Yes	Yes
Drive Security Method	Indicates if drive security is enabled.	Yes	Yes	Yes

17.2 Adding a Virtual Drive to a Drive Group

You can add virtual drives to an existing drive group provided there is sufficient storage space in the existing virtual drives of the drive group.

Perform the following steps to add a virtual drive to an existing drive group:

1. Navigate to the Controller dashboard and click a drive group name (for example, **DG_1**).
   In the right pane, under **Actions**, the **Add Virtual Drives** option appears.

2. Click **Add Virtual Drives**.
   The **Virtual Drive Settings** window appears.
3. Specify the settings you want for the virtual drives you want to create.
   See Selecting Virtual Drive Settings for details on creating virtual drives.
4. Click **Add Virtual Drives**.
   The newly created virtual drive gets added to the selected drive group.

17.3 RAID Level Migration

RAID level migration is the process of converting one RAID configuration to another. You can perform RAID level migration at the drive group level. The following table describes the valid RAID level migration matrix.

Table 12 Drive Group – RAID Level Migration

Initial RAID Level	Migrated RAID Level
RAID 0	RAID 1
RAID 0	RAID 5
RAID 0	RAID 6
RAID 1	RAID 0
RAID 1	RAID 5
RAID 1	RAID 6
RAID 5	RAID 0
RAID 5	RAID 6
RAID 6	RAID 0
RAID 6	RAID 5

17.3.1 Migrating the RAID Level of a Drive Group

Perform the following steps to migrate the RAID level of a drive group.

1. Navigate to the Controller dashboard and click a drive group name (for example, **DG_1**).
   In the right pane, under **Actions**, the **Modify Drive Group** option appears.
2. Click **Modify Drive Group**.
   The **Modify Drive Group** window appears.

Figure 54 Modify Drive Group Window

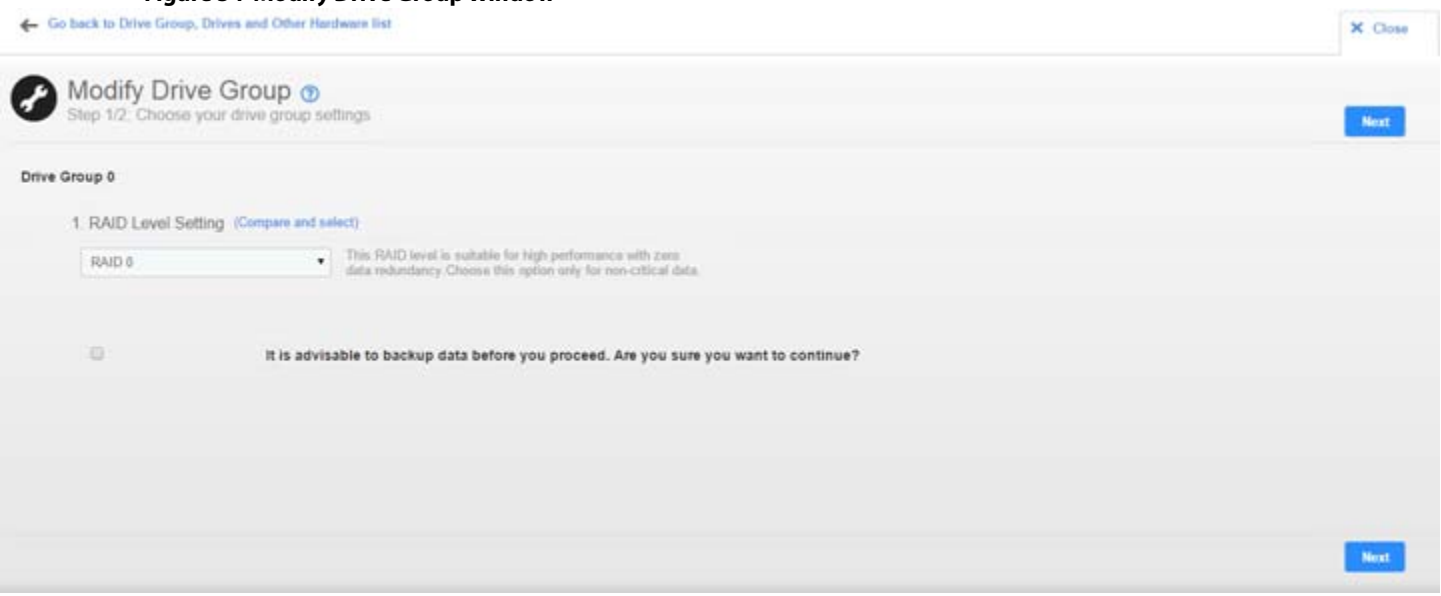

3. In **RAID Level Setting**, select the RAID level to which you want to migrate the drive group.
   - Select the **Auto Back-up** check box to back up the data before you change the RAID level.
4. Click **Next**.

   The **Modify Drive Group** window appears and provides you an option to add, remove, or directly change the RAID level. Depending on the source and the target RAID levels, you can also add drives directly without having to choose an option.

Figure 55 Modify Drive Group Settings

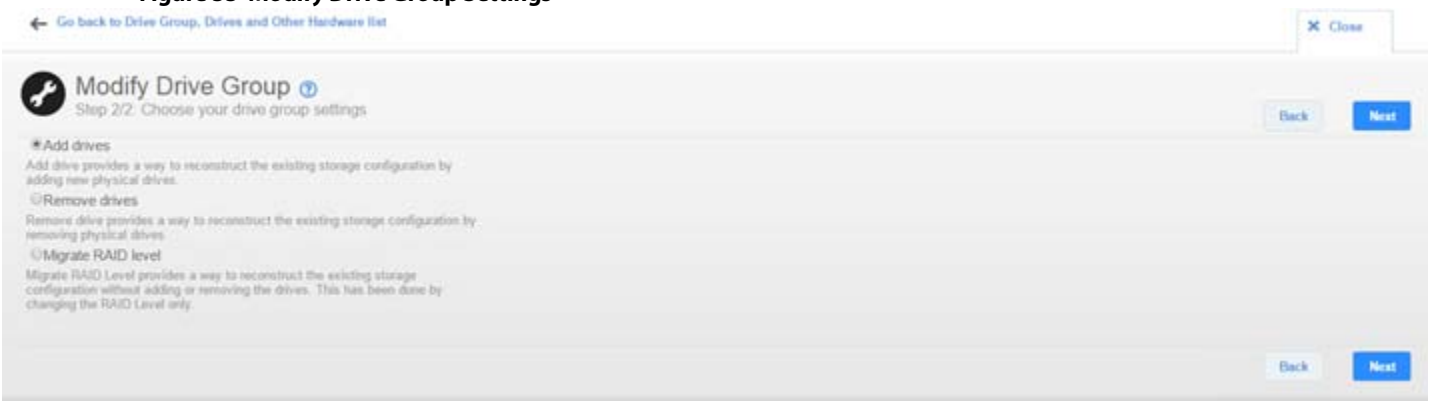

17.3.1.1 Adding Physical Drives to a Configuration

For example, if you are migrating the RAID level of a drive group from RAID 0 to RAID 5, the **Modify Drive Group** wizard allows you to add unconfigured physical drives to the existing configuration to enable the RAID level migration.

1. In the **Modify Drive Group** window, click **Add Physical Drives**.

   The drives you add must have the same capacity as or greater capacity than the drives already in the drive group, or you cannot change the RAID level.

   The **Available Unconfigured Drive** window appears. It lists the drives you can add, and it states whether you have to add a minimum number of drives to change the RAID level from the current level to the new RAID level.

Figure 56 Available Unconfigured Drive Window

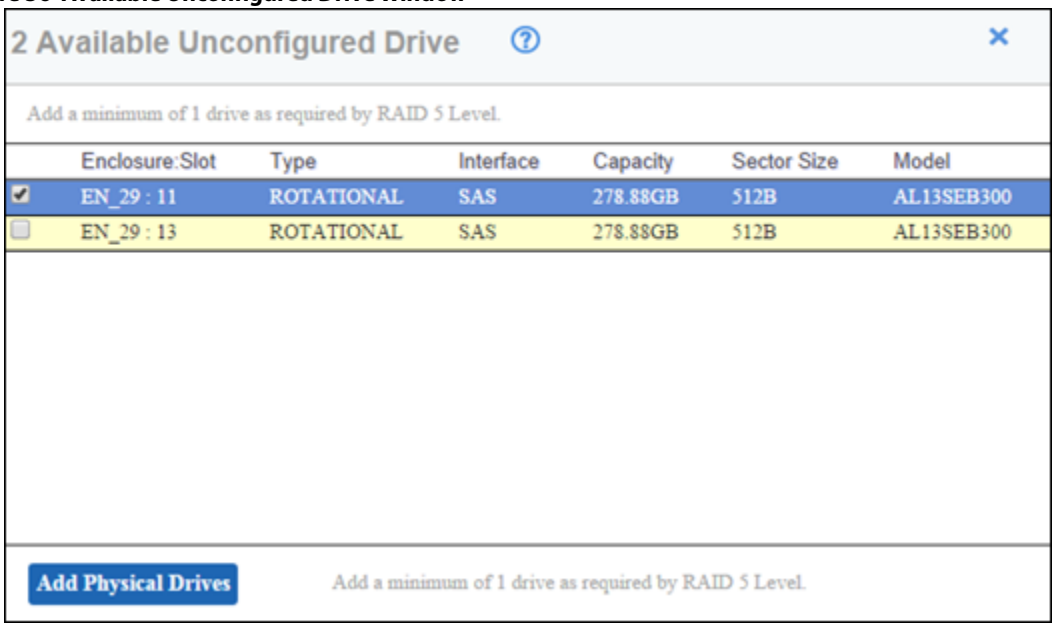

2. Select the available unconfigured drives and click **Add Physical Drives**.
3. Click **Finish**.

   The RAID level is migrated. A confirmation message appears. You can monitor the progress of the reconstruction. See Background Operations Support.

17.3.1.2 Removing Drives From a Configuration

For example, if you are migrating the RAID level of a drive group from RAID 5 to RAID 0, the **Modify Drive Group** wizard allows you to remove physical drives from the existing configuration to enable the RAID level migration.

1. In the **Modify Drive Group** window, select **Remove drives** and click **Next**.

   The **Modify Drive Group** window appears and it states the number of physical drives that you have to remove to change the RAID level from the current level to a new RAID level and the maximum number of physical drives that can be removed.
2. Click on the X mark to remove the drives.
3. Click **Finish**.

   The RAID Level is migrated. A confirmation message appears. You can monitor the progress of the reconstruction. See Background Operations Support.

17.3.1.3 Migrating the RAID Level Without Adding or Removing Drives

For example, if you are migrating the RAID level of your drive group from RAID 5 to RAID 0, the **Modify Drive Group** wizard allows you to migrate the RAID level without adding or removing the drives.

1. In the **Modify Drive Group**, select **Migrate RAID level** and click **Next**.

   The RAID level is migrated. A confirmation message appears. You can monitor the progress of the reconstruction. See Background Operations Support.

Chapter 18: Managing Virtual Drives

The LSI Storage Authority software enables you to perform various operations on the virtual drives.

18.1 Viewing Virtual Drive Properties

Select a virtual drive from a drive group in the controller dashboard to view its properties.

Figure 57 Virtual Drive Properties

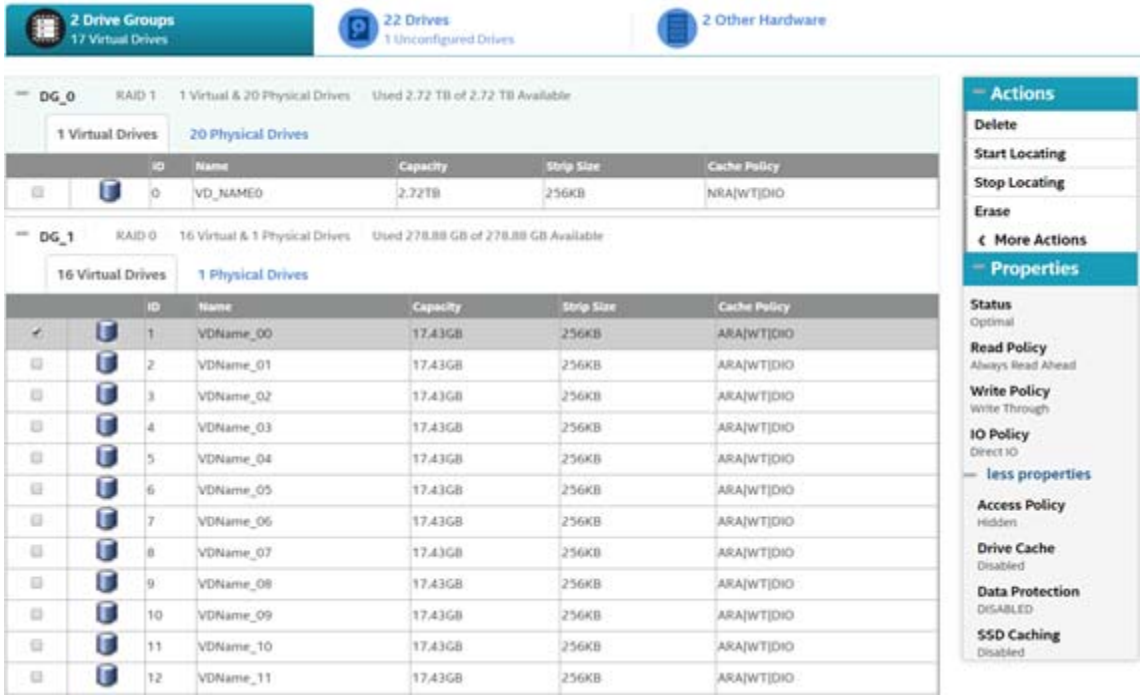

Table 13 Virtual Drive Properties

Property	Description	MegaRAID	iMegaRAID	Initiator-Target
Status	The current status of the virtual drive. The following options are available: ■ **Optimal** ■ **Partially Degraded** ■ **Degraded** ■ **Offline**	**Yes**	**Yes**	**No**
Read Policy	The read cache policy for the virtual drive. The following options are available: ■ **Read Ahead** ■ **No Read Ahead**	**Yes**	**Yes**	**No**
Write Policy	The write cache policy for the virtual drive. The following options are available: ■ **Write Back** ■ **Write Through** ■ **Always Write Back**	**Yes**	**Yes**	**No**[a]
IO Policy	The input/output policy for the virtual drive. The following options are available: ■ **Direct IO** ■ **Cached IO**	**Yes**	**Yes**	**No**
Host Access Policy **NOTE** This property appears only if the controller supports High Availability DAS.	Indicates whether or not the virtual drive is shared between the servers in a cluster. The values for this property are **Shared**, **Exclusive**, and **Exclusive to Peer Controller**.	**No**	**No**	**No**
Peer Has No Access **NOTE** This property appears only if the controller supports High Availability DAS.	Indicates whether the peer controller has access to the shared virtual drive. This property appears only if the virtual drive is shared.	**No**	**No**	**No**
Access Policy	The access policy for the virtual drive. The following options are available: ■ **Read Write** ■ **Read Only** ■ **Hidden** The Hidden policy is applicable for only hidden VDs. No other access policies will be applicable once you select Hidden as the access policy.	**Yes**	**Yes**	**No**
Drive Cache	The virtual drive cache setting. The following possible options are available: ■ **Unchanged** ■ **Enable** ■ **Disable**	**Yes**	**Yes**	**No**
Data Protection	Indicates if data protection feature is enabled for the virtual drive.	**Yes**	**Yes**	**No**
SSD Caching	Indicates if SSD caching is enabled.	**Yes**	**Yes**	**No**

a. Write Cache of the LUN is not supported by all target enclosures. Check your enclosure's documentation to determine whether Write Cache is supported on your device.

18.2 Modifying Virtual Drive Properties

You can change the read policy, write policy, and other virtual drive properties at any time after a virtual drive is created. Perform the following steps to modify the virtual drive settings.

1. Navigate to the Controller dashboard, click a drive group name (for example, **DG_1**). Click the + icon corresponding to a drive group to display its contents.

   The virtual drives and physical drives associated with the selected drive group appear.
2. Click the virtual drive whose settings you want to change.
3. Select **Actions > Modify Properties**.

   The **Modify <Virtual Drive Name>** dialog appears.

Figure 58 Modify Virtual Drive Dialog

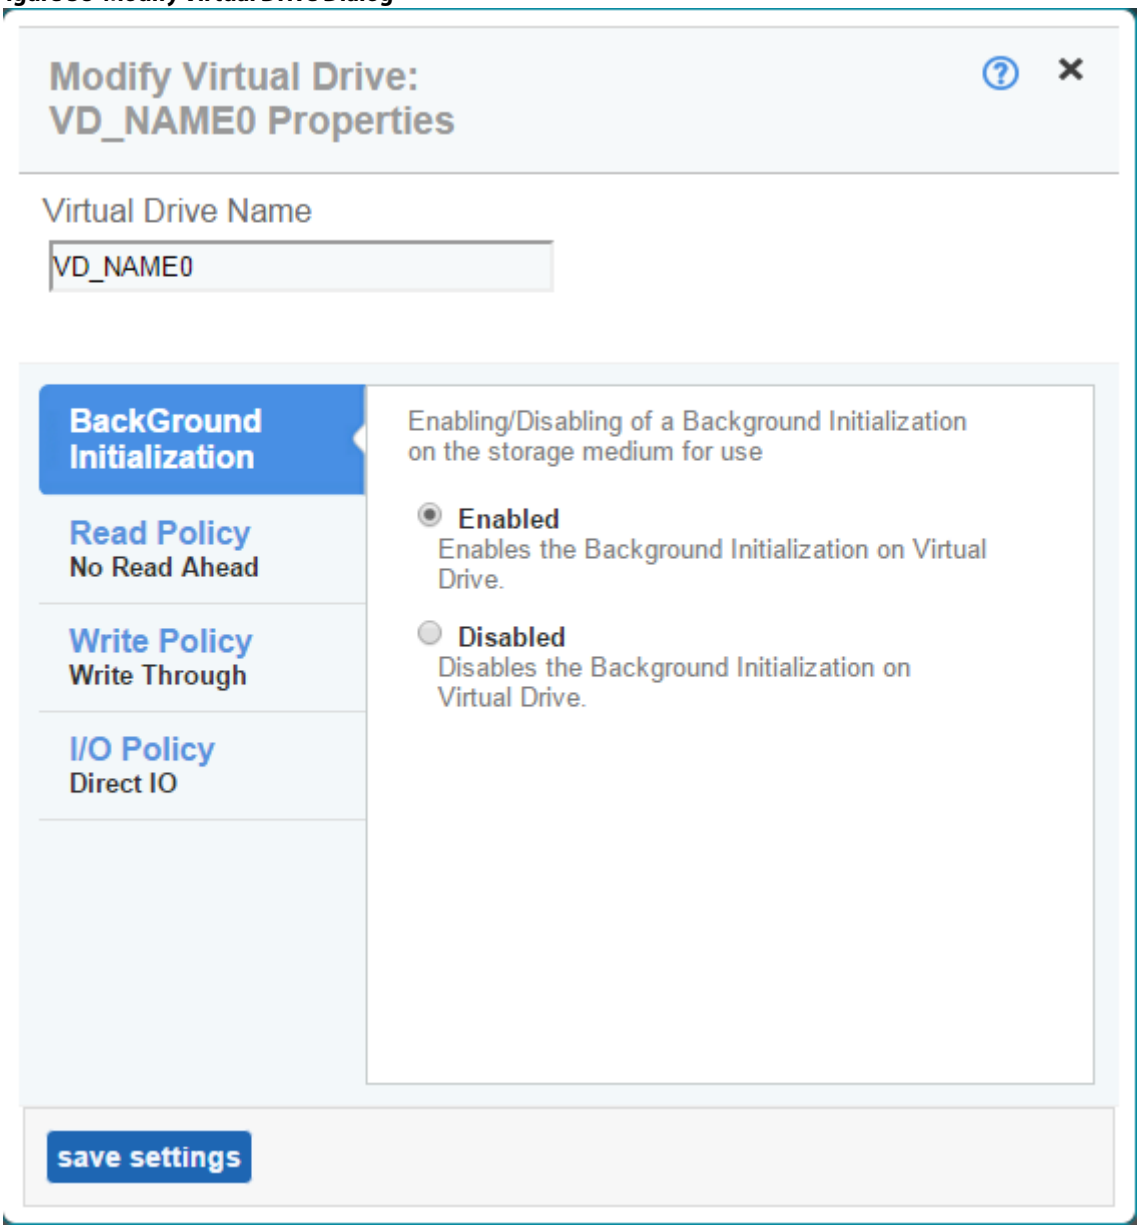

4. Change the virtual drive properties as needed. For information about these properties, see Selecting Virtual Drive Settings.
5. Click **Save Settings**.

18.3 Start and Stop Locating a Virtual Drive

If the drives in the virtual drives are in a disk enclosure, you can identify them by making their LEDs blink. Perform the following steps to identify the virtual drives:

1. Navigate to the Controller dashboard, click a drive group name (for example, **DG_1**). Click the + icon corresponding to a drive group to display its contents.

   The virtual drives and physical drives associated with the selected drive group appear.
2. Click the virtual drive that you want to locate in the disk enclosure.
3. Select **Actions > Start Locate**.

   The LEDs on the drives in the virtual drive start blinking.
4. To stop the LEDs from blinking, select **Actions > Stop Locate**.

18.4 Erasing a Virtual Drive

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

Perform the following steps to erase a virtual drive.

1. Navigate to the Controller dashboard, click a drive group name (for example, **DG_1**). Click the + icon corresponding to a drive group to display its contents.

   The virtual drives and physical drives associated with the selected drive group appear.
2. Click the virtual drive whose content you want to erase.
3. Select **Actions > Erase**.

   The **Virtual Drive Erase** dialog appears.

Figure 59 Virtual Drive Erase Dialog

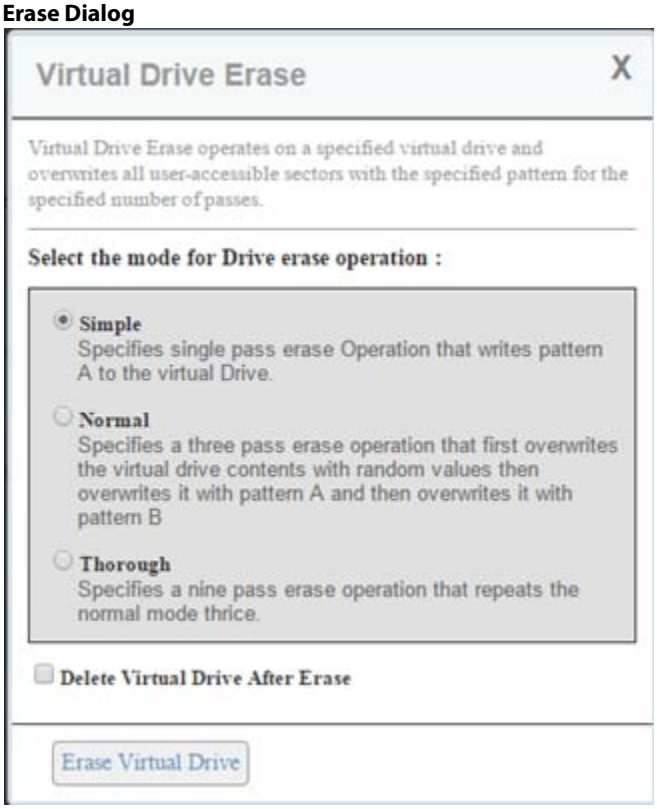

The dialog shows the following modes:

- **Simple**
- **Normal**
- **Thorough**

4. Select a mode and click **Erase Virtual Drive**.

   A warning message appears asking for your confirmation.

5. Click **Yes, Erase Drive**.

   After the virtual drive erase operation has started, the **Stop Erase** option is enabled in the **Actions** menu. You can monitor the progress of the erase operation. See Background Operations Support.

To delete the virtual drive after the erase operation has been completed, select the **Delete Virtual Drive After Erase** check box.

18.5 Initializing a Virtual Drive

When you create a new virtual drive with the **Advanced Configuration** wizard, you can select the **Fast Initialization** or **Full Initialization** option to initialize the drive immediately. However, you can select **No Initialization** if you want to initialize the virtual drive later.

Perform the following steps to initialize a virtual drive after completing the configuration process.

1. Navigate to the Controller dashboard, click a drive group name (for example, **DG_1**). Click the + icon corresponding to a drive group to display its contents.

   The virtual drives and physical drives associated with the selected drive group appear.

2. Click the virtual drive that you want to initialize.

3. Select **Actions > Start Initialize**.

   A warning message appears.

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

4. Select the **Fast Initialization** check box if you want to use this option. If you leave the check box unselected, the software runs a Full Initialization on the virtual drive.
5. Click **Yes, Start Initialization** to begin the initialization.

   You can monitor the progress of the initialization. See Background Operations Support.

18.6 Starting Consistency Check on a Virtual Drive

Perform the following steps to start consistency check on a virtual drive. For more information of consistency check, see Running Consistency Check.

1. Navigate to the Controller dashboard, click a drive group name (for example, **DG_1**). Click the + icon corresponding to a drive group to display its contents.

   The virtual drives and physical drives associated with the selected drive group appear.
2. Click the virtual drive on which you want to start consistency check.
3. Select **Actions > Start Consistency Check**.

   The consistency check operation starts. You can see the progress of this operation in the **Background Processes in Progress** section. After the consistency check operation has started, the **Stop Consistency Check** option is enabled in the **Actions** menu.

18.7 Expanding the Online Capacity of a Virtual Drive

Online Capacity Expansion (OCE) allows the capacity of a virtual disk to be expanded by adding new physical disks or making use of unused space on existing disks, without requiring a reboot. Perform the following steps to expand the capacity of a virtual drive.

ATTENTION Make sure to back up the data on the virtual drive before you proceed with the online capacity expansion.

1. Navigate to the Controller dashboard, click a drive group name (for example, **DG_1**). Click the + icon corresponding to a drive group to display its contents.

   The virtual drives and physical drives associated with the selected drive group appear.
2. Click the virtual drive whose capacity you want to expand.
3. Select **Actions > Expand**.

   The **Expand Virtual Drive** dialog appears.

Figure 60 Expand Virtual Drive Dialog

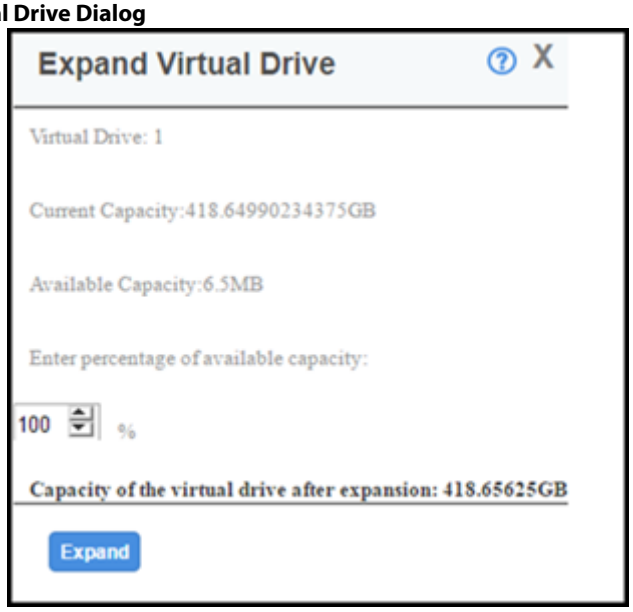

4. Select the percentage of the available capacity that you want the virtual drive to use.
5. Click **Expand**.

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

18.8 Deleting a Virtual Drive

You can delete virtual drives on a controller to reuse that space for new virtual drives.

CAUTION All data on a virtual drive is lost when you delete it. Make sure to back up the data before you delete a virtual drive.

Perform the following steps to delete a virtual drive.

1. Navigate to the Controller dashboard, click a drive group name (for example, **DG_1**). Click the + icon corresponding to a drive group to display its contents.

   The virtual drives and physical drives associated with the selected drive group appear.
2. Click the virtual drive that you want to delete.
3. Select **Actions > Delete**.

   A confirmation dialog appears.
4. Select **Confirm** and click **Yes, Delete** to proceed with the delete operation.

   A message appears confirming that the virtual drive is deleted successfully.

   Operating system drives cannot be deleted. If you try to do so, an error message appears.

18.9 Hiding and Unhiding a Virtual Drive or a Drive Group

You can hide or unhide either a virtual drive or a drive group on the controller.

18.9.1 Hiding a Virtual Drive

You can hide a virtual drive on the controller.

Perform the following steps to hide a virtual drive:

1. Navigate to the Controller dashboard and click **Drive Groups** (for example, **DG_1**).
2. Click the + icon corresponding to a drive group to display its contents.
   The virtual drives associated with the selected drive group appear.
3. Select a virtual drive that you want to hide.
4. Select **Actions > More Actions > Hide**.
   A message box appears, asking you to confirm the operation.
5. Click **Yes** to confirm and hide the virtual drive.

18.9.2 Unhiding a Virtual Drive

You can unhide a virtual drive on the controller.

Perform the following steps to unhide a virtual drive:

1. Navigate to the Controller dashboard and click **Drive Groups** (for example, **DG_1**).
2. Select the entire virtual drive group that you want to unhide.
3. Select **Actions > Un Hide**.
   A message box appears, asking you to confirm the operation.
4. Click **Yes** to unhide the virtual drive.

18.9.3 Hiding a Drive Group

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

Perform the following steps to hide a drive group:

1. Navigate to the Controller dashboard and click **Drive Groups** (for example, **DG_1**).
2. Select a drive group that you want to hide.
3. Navigate to **Actions > Hide All Virtual Drives**.
   A message box appears, asking you to confirm the operation.
4. Select the **Confirm** checkbox and click **Yes** to hide the drive group.

18.9.4 Unhiding a Drive Group

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

Perform the following steps to unhide a drive group:

1. Navigate to the Controller dashboard and click **Drive Groups** (for example, **DG_1**).
2. Select a drive group that you want to unhide.
3. Navigate to **Actions > Un Hide All Virtual Drives**.
   A message box appears, asking you to confirm the operation.
4. Select the **Confirm** checkbox and click **Yes** to unhide the drive group.

Chapter 19: Managing Physical Drives

The LSI Storage Authority software allows you to manage all of the physical drives that are connected to the controller.

19.1 Viewing Physical Drive Properties

Select a physical drive from a drive group in the Controller dashboard to view its properties. The following figure and table describe the physical drive properties.

Figure 61 Physical Drive Properties

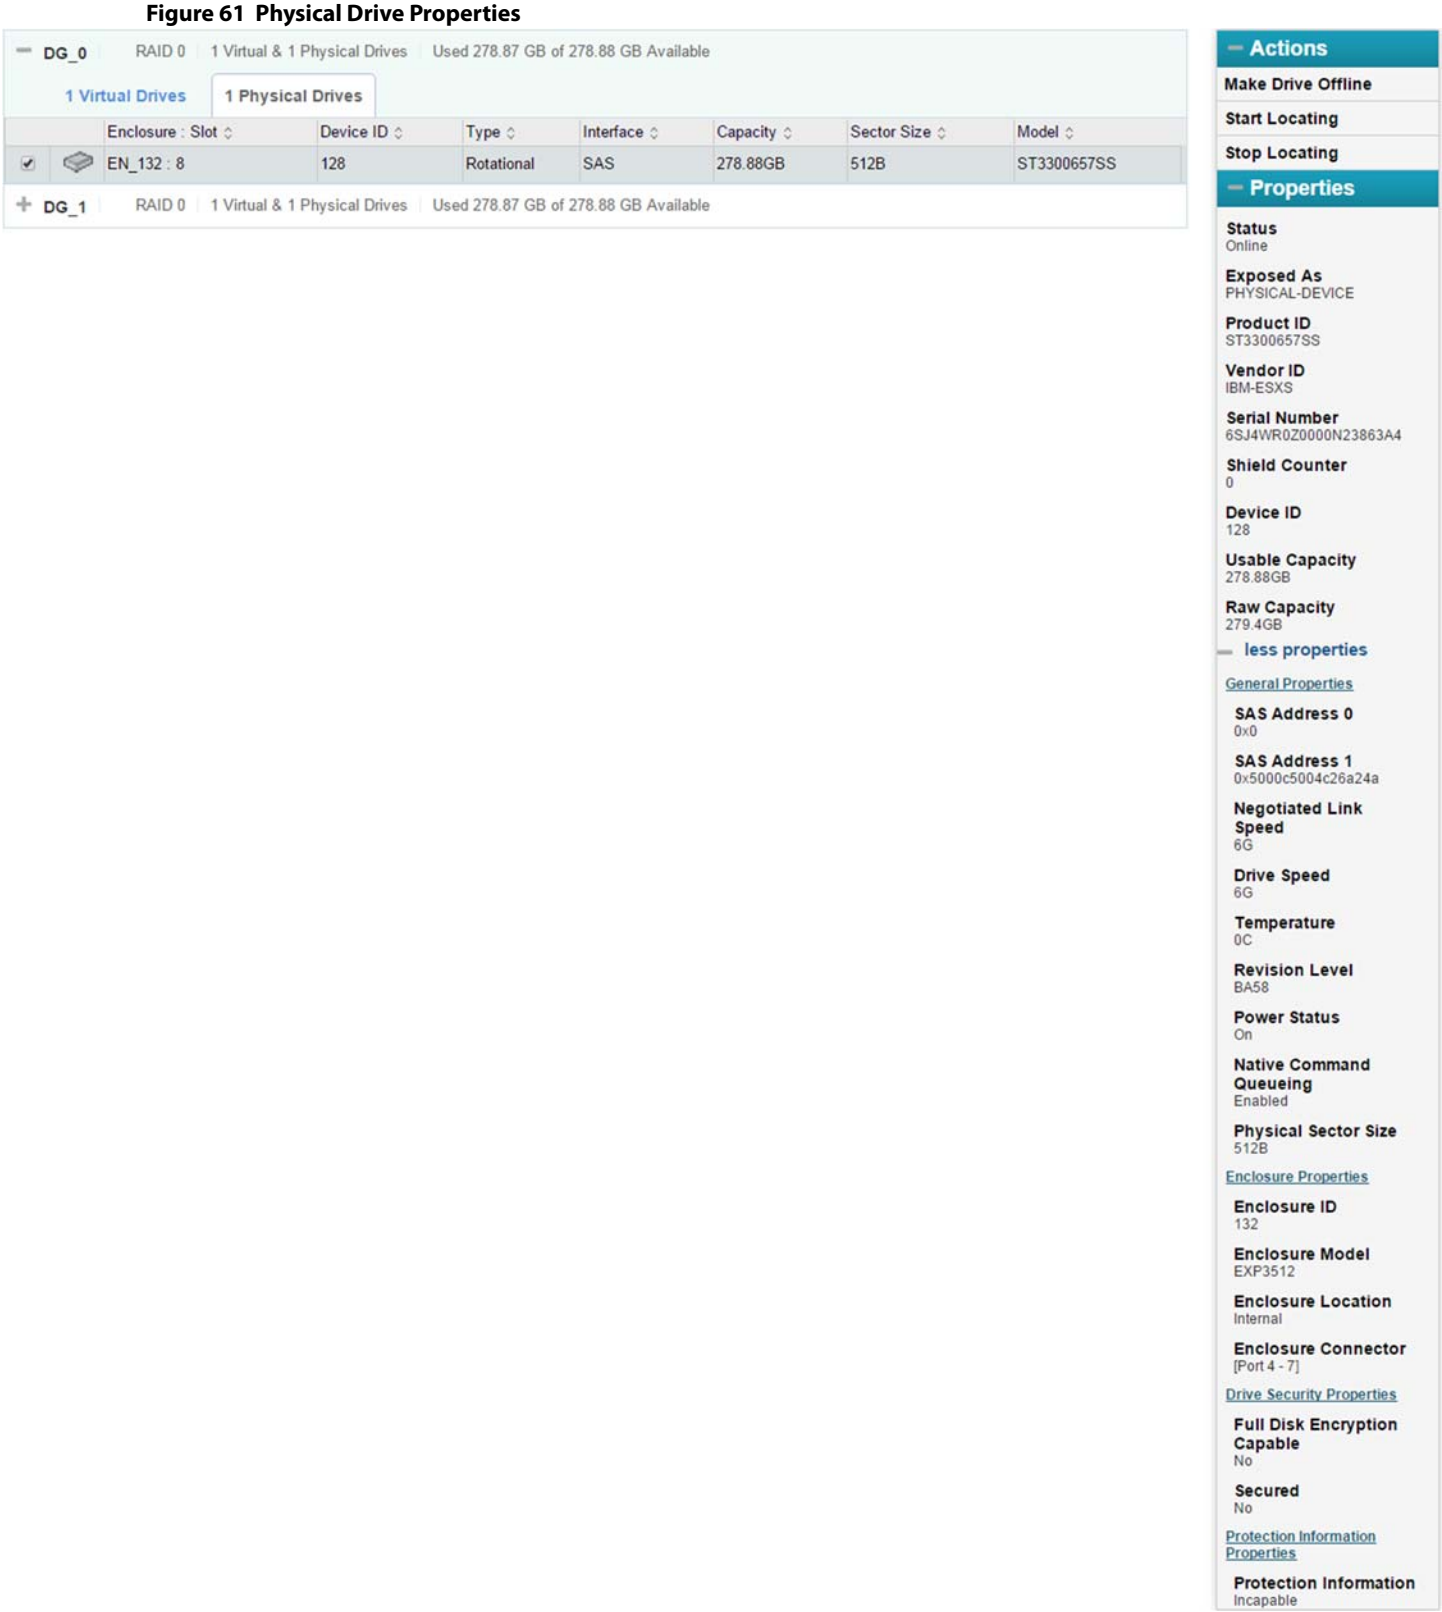

Table 14 Physical Drive Properties

Property	Description	MegaRAID	iMegaRAID	Initiator-Target
Status	The current status of the physical drive.	Yes	Yes	Yes
Exposed As	To differentiate the physical drives, the drives are exposed as one of the following drive: ■ JBOD ■ PHYSICAL-DEVICE	Yes	Yes	Yes
Product ID	The product ID of the physical drive.	Yes	Yes	Yes
Vendor ID	The ID assigned to the physical drive by the vendor.	Yes	Yes	Yes
Serial Number	The serial number of the physical drive.	Yes	Yes	Yes
Shield Counter	The shield counter value.	Yes	Yes	Yes
Device ID	The device ID of the physical drive that is assigned by the manufacturer.	Yes	Yes	Yes
Usable Capacity	The usable storage capacity, based on the RAID level used.	Yes	Yes	Yes
Raw Capacity	The actual full capacity of the drive before any coercion mode is applied to reduce the capacity.	Yes	Yes	Yes
General Properties				
SAS Address 0	The World Wide Name (WWN) for the physical drive.	Yes	Yes	Yes
SAS Address 1	The World Wide Name (WWN) for the physical drive.	Yes	Yes	Yes
Negotiated Link Speed	The negotiated link speed for data transfer to and from the physical drive.	Yes	Yes	Yes
Drive Speed	The speed of the physical drive.	Yes	Yes	Yes
Temperature	The temperature of the physical drive.	Yes	Yes	Yes
Revision Level	The revision level of the physical drive's firmware.	Yes	Yes	Yes
Power Status	The Power Status displays the following status: ■ **On**- when a physical drive is spun up. ■ **Powersave**- when a physical drive is spun down.	Yes	Yes	Yes
Native Command Queueing	Indicates if the Native Command Queueing function is enabled. Native Command Queueing enables the physical drive to queue the I/O requests and reorder them for efficiency.	Yes	Yes	Yes
Physical Sector Size	The size of the physical sector of the drive. The possible options are 4 KB or 512 KB.	Yes	Yes	Yes
Enclosure Properties				
Enclosure ID	The ID of the enclosure in which the physical drive is located.	Yes	Yes	Yes[a]
Enclosure Model	The type of enclosure in which the physical drive is located.	Yes	Yes	Yes
Enclosure Location	The port number of the enclosure to which the physical drive is connected.	Yes	Yes	Yes
Enclosure Connector	Indicates the connector name, position, and size of the wide port. **Wide ports**: `[Port 0-3 x4]` - Single X4 wide port `[Port 0-3 x4] & [Port 4-7 x4]`- two X4 wide ports `[Port 0-3 & Port 4-7 x8]`- Single X8 wide port **Single Drive** `[Port 0-3 x1]`	Yes	Yes	Yes[b]

Table 14 Physical Drive Properties (Continued)

Property	Description	MegaRAID	iMegaRAID	Initiator-Target
Drive Security Properties				
Full Disk EncryptionCapable	Indicates if disk encryption is enabled for the physical drive.	Yes	Yes	Yes
Secured	Indicates if the drive is secured.	Yes	Yes	Yes
Protection Information Properties				
Protection Information	■ Indicates if the SCSI Protection Information type is active for the drive.	Yes	Yes	Yes

a. The Locate feature is not supported by all target enclosures. Check your enclosure's documentation to determine whether the Locate feature is supported on your device.

b. The external wide port address of the SAS HBA is not a fixed address. The auto-configuration feature dynamically assigns the SAS address based on the port linkup sequence and port number. As the address assignment is based on port link sequence, when you do a DC/AC power cycle or move a cable, the port's SAS address may change resulting in the original LUN mapping getting changed.

19.2 Start and Stop Locating a Drive

If the physical drives are in a disk enclosure, you can identify them by making their LEDs blink. Perform the following steps to identify the physical drives:

1. On the Controller dashboard, navigate to the physical drive and select the drive that you want to identify such as, Unconfigured Good drive, Online physical drive, Configured drive and so on.
2. Select **Actions > Start Locating**.

   The corresponding LED on the physical drive starts blinking.
3. To stop the LED from blinking, select **Actions > Stop Locating**.

19.3 Making a Drive Offline

Perform the following steps to make a drive offline.

ATTENTION After you perform this procedure, all of the data on the drive will be lost.

1. Navigate to the Controller dashboard, click a drive group name (for example, **DG_1**). Click the + icon corresponding to a drive group to display its contents.

   The virtual drives and physical drives associated with the selected drive group appear.
2. Click the **Physical Drive** tab, and select a drive that you want to make offline.
3. Select **Actions > Make Drive Offline**.

   The drive status changes to Offline.

19.4 Making a Drive Online

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

1. Navigate to the Controller dashboard, click a drive group name (for example, **DG_1**). Click the + icon corresponding to a drive group to display its contents.

   The virtual drives and physical drives associated with the selected drive group appear.
2. Click the **Physical Drive** tab, and select the offline drive that you want to make online.
3. Select **Actions > Make Drive Online**.

   The drive status changes to Online.

19.5 Replacing a Drive

You might want to replace a drive 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. Perform the following steps to replace a drive.

ATTENTION Make sure to back up the data on the drive before you replace it.

1. Navigate to the Controller dashboard, click a drive group name (for example, **DG_1**). Click the + icon corresponding to a drive group to display its contents.

   The virtual drives and physical drives associated with the selected drive group appear.
2. Click the **Physical Drive** tab, and select a drive which you want to replace.
3. Select **Actions > Replace Drive**.

   The **Replace Drive** dialog appears.

Figure 62 Replace Drive

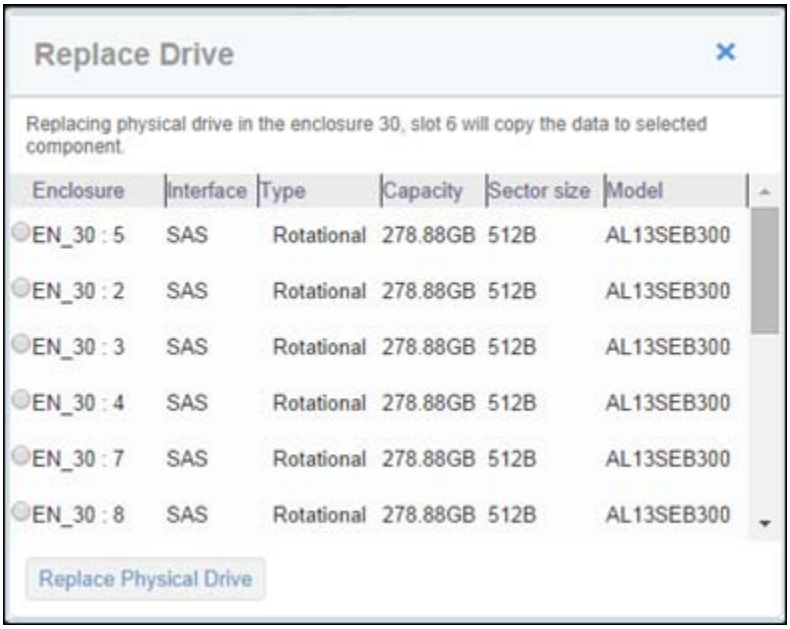

4. Select a replacement drive and click **Replace Physical Drive**.

   A confirmation message appears.

5. Select **Confirm** and click **Yes, Replace Drive** to proceed with the replace operation.
   The drive is replaced and the data is copied to the selected component.

19.6 Assigning Global Hot Spares

A global hot spare replaces a failed physical drive in any redundant array, as long as the capacity of the global hot spare is equal to or larger than the coerced capacity of the failed physical drive. Perform the following steps to assign global hot spares.

1. Navigate to the Controller dashboard and click the **Drives** tab.
   All of the associated drives appear.
2. Expand **Unconfigured Drives** and select an unconfigured good drive.
3. Select **Actions > Assign Global Hotspare**.
   The unconfigured good drive is changed to a global hot spare. The status of the unconfigured good drive appears as a global hot spare in the **Hot Spares** section.

19.7 Removing Global Hot Spares

Perform the following steps to remove a hot spare.

1. Navigate to the Controller dashboard and click the **Drives** tab.
   All of the associated drives appear.
2. Expand **Hot Spares** and select a hot spare that you want to remove.
3. Select **Actions > Remove Global Hotspare**.
   The hot spare drive is removed and is listed in the **Unconfigured Drives** section as an unconfigured good drive.

19.8 Assigning Dedicated Hot Spares

Dedicated hot spare drives provide protection to one or more specified drive groups on the controller. If you select an Unconfigured Good drive, you have the option of assigning it as a dedicated hot spare drive. Perform these steps to assign a dedicated hot spare.

1. Navigate to the Controller dashboard and click the **Drives** tab.
   All of the associated drives appear.
2. Expand **Unconfigured Drives** and select an unconfigured good drive.
3. Select **Actions** > **More Actions** > **Assign Dedicated Hotspare**.
   The **Drive Groups** dialog appears.

Figure 63 Drive Groups Dialog

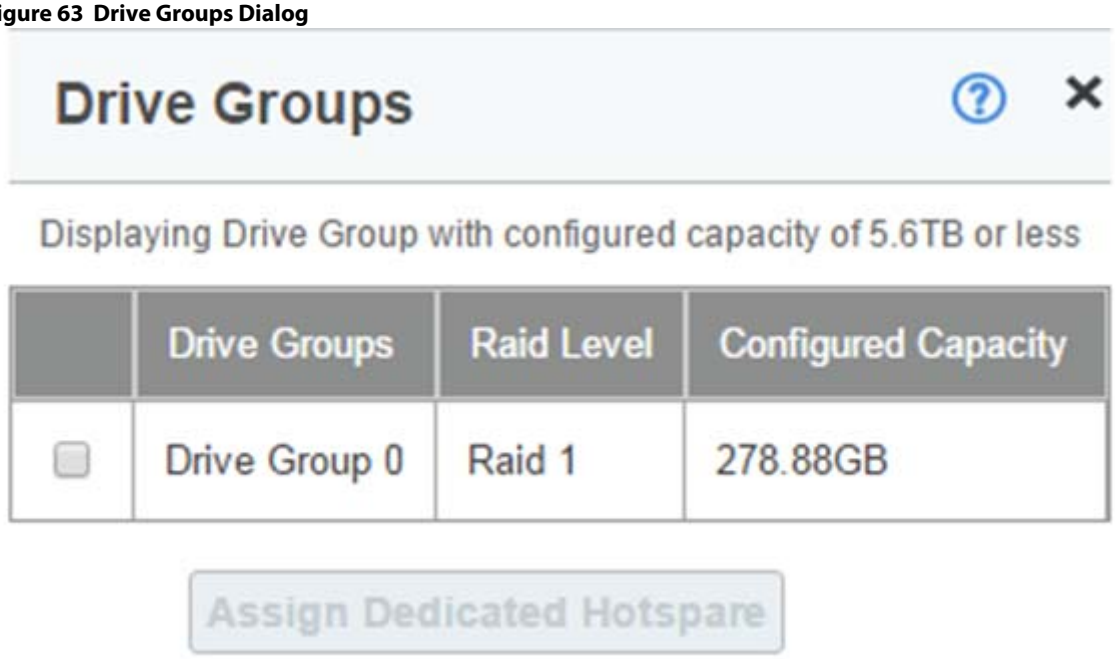

4. Select a drive group and click **Add Dedicated Hotspare**.

   A confirmation message appears.
5. Click **Done**.

   The unconfigured good drive is changed to a dedicated hot spare. The status of the unconfigured good drive appears as a dedicated hot spare in the **Hot Spares** section.

19.9 Rebuilding a Drive

If a drive, which is configured as RAID 1, 5, 6, 10, 50, or 60 fails, the LSI Storage Authority software automatically rebuilds the data on a hot spare drive to prevent data loss. The rebuild is a fully automatic process. You can monitor the progress of drive rebuilds in the **Background Processes in Progress** window. See Background Operations Support.

19.10 Converting Unconfigured Bad Drive to Unconfigured Good Drive

Perform the following steps to convert an unconfigured bad drive to an unconfigured good drive.

1. Navigate to the Controller dashboard and click the **Drives** tab.

   All of the associated drives appear.
2. Expand **Unconfigured Drives** and select an unconfigured bad drive.
3. Select **Actions > Make Unconfigured Good**.

   A confirmation message appears.
4. Select **Confirm** and click **Yes, Make Unconfigured Good** to proceed with the operation.

   The unconfigured bad drive is changed to unconfigured good drive. The status of the unconfigured bad drive appears as unconfigured good in the **Unconfigured Drives** section.

19.11 Removing a Drive

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

1. Navigate to the Controller dashboard and click the **Drives** tab.
   All of the associated drives appear.
2. Expand **Unconfigured Drives** and select a drive that you want to remove.
3. Select **Actions > Prepare for Removal**.
   The drive is in the power save mode and is ready for removal.
4. Wait until the drive spins down and then remove it.
   If you do not want to remove the drive, select **Actions > Undo Prepare for Removal**.

19.12 Make Unconfigured Good and Make JBOD

When you power down a controller and insert a new physical drive and if the inserted drive does not contain valid DDF metadata, the drive status is listed as JBOD (Just a Bunch of Drives) when you power the system again. When you power down a controller and insert a new physical drive and if the drive contains valid DDF metadata, its drive state is Unconfigured Good. A new drive in the JBOD drive state is exposed to the host operating system as a stand-alone drive. You cannot use JBOD drives to create a RAID configuration, because they do not have valid DDF records. Therefore, you must convert JBOD drives to unconfigured good drives.

If the controller supports JBOD drives, the LSI Storage Authority includes options for converting JBOD drives to an unconfigured good drive, or vice versa.

19.12.1 Making Unconfigured Good Drives

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

1. Navigate to the Controller dashboard and click the **Drives** tab.
   All of the associated drives appear.
2. Expand **JBOD** and select a JBOD drive.
3. Select **Actions > Make Unconfigured Good**.
   A confirmation message appears.
4. Select **Confirm** and click **Yes, Make Unconfigured Good** to proceed with the operation.
   The JBOD drive is changed to an unconfigured good drive.

19.12.2 Making JBOD

Perform these steps to change the status of unconfigured good drives to JBOD.

1. Navigate to the Controller dashboard and click the **Drives** tab.
   All of the associated drives appear.
2. Expand **Unconfigured Drives** and select an unconfigured good drive.
3. Select **Actions > Make JBOD**.
   The unconfigured good drive is changed to a JBOD drive.

19.13 Erasing a Drive

You can erase data on Non SEDs (normal HDDs) by using the **Drive Erase** option. For Non–SEDs, the erase operation consists of a series of write operations to a drive that overwrites every user-accessible sector of the drive with specified patterns. It can be repeated in multiple passes using different data patterns for enhanced security. The erase operation is performed as a background task. Perform the following steps to erase a drive.

1. Navigate to the Controller dashboard and click the **Drives** tab.

   All of the associated drives appear.
2. Expand **Unconfigured Drives** and select an unconfigured good drive.
3. Select **Actions > More Actions > Drive Erase**.

   The **Physical Drive Erase** dialog appears.

Figure 64 Physical Drive Erase Dialog

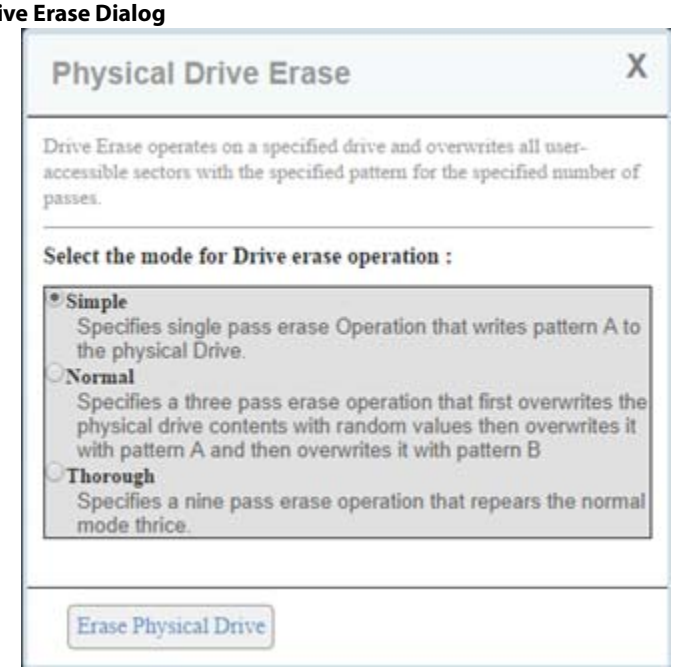

   The dialog shows the following modes:
   - **Simple**
   - **Normal**
   - **Thorough**
4. Select a mode and click **Erase Physical Drive**.

   A warning message appears asking for your confirmation.
5. Click **Yes, Erase Drive**.

   After the drive erase operation has started, the **Stop Erase** option is enabled in the **Actions** menu. You can monitor the progress of the erase operation. See Background Operations Support.

19.14 Erasing a Drive Securely

The Instant Secure Erase erases data from encrypted drives.

ATTENTION All data on the drive is lost when you erase it. Before starting this operation, back up any data that you want to keep.

1. Navigate to the Controller dashboard and click the **Drives** tab.

   All of the associated drives appear.
2. Expand **Unconfigured Drives** and select an unconfigured good drive.
3. Select **Actions > Instant Secure Erase**.

   A confirmation message appears.
4. Select **Confirm** and click **Yes, Securely Erase Drive** to proceed with the operation.

   After the secure erase operation has started, the **Stop Erase** option is enabled in the **Actions** menu. You can monitor the progress of the erase operation. See Background Operations Support.

Chapter 20: Managing Hardware Components

When you select the **Other Hardware** tab from the Controller dashboard, the hardware components appear as shown in the following figure.

Figure 65 Other Hardware

20.1 Monitoring Energy Packs

When the LSI Storage Authority software is running, you can monitor the status of all of the energy packs connected to the controllers in the server.

Learn Cycle

Learn cycle is an energy pack calibration operation that is performed by the controller periodically to determine the condition of the energy pack. You can start the learn cycles manually or automatically. To choose automatic learn cycles, enable the automatic learn cycles feature. If you enable automatic learn cycles, you can delay the start of the learn cycles for up to 168 hours (7 days).

20.1.1 Viewing Energy Pack Properties

Select an energy pack from the **Other Hardware** tab in the Controller dashboard to view its properties.

The following figure and table describe the energy pack basic and advanced properties.

Figure 66 Energy Pack Properties

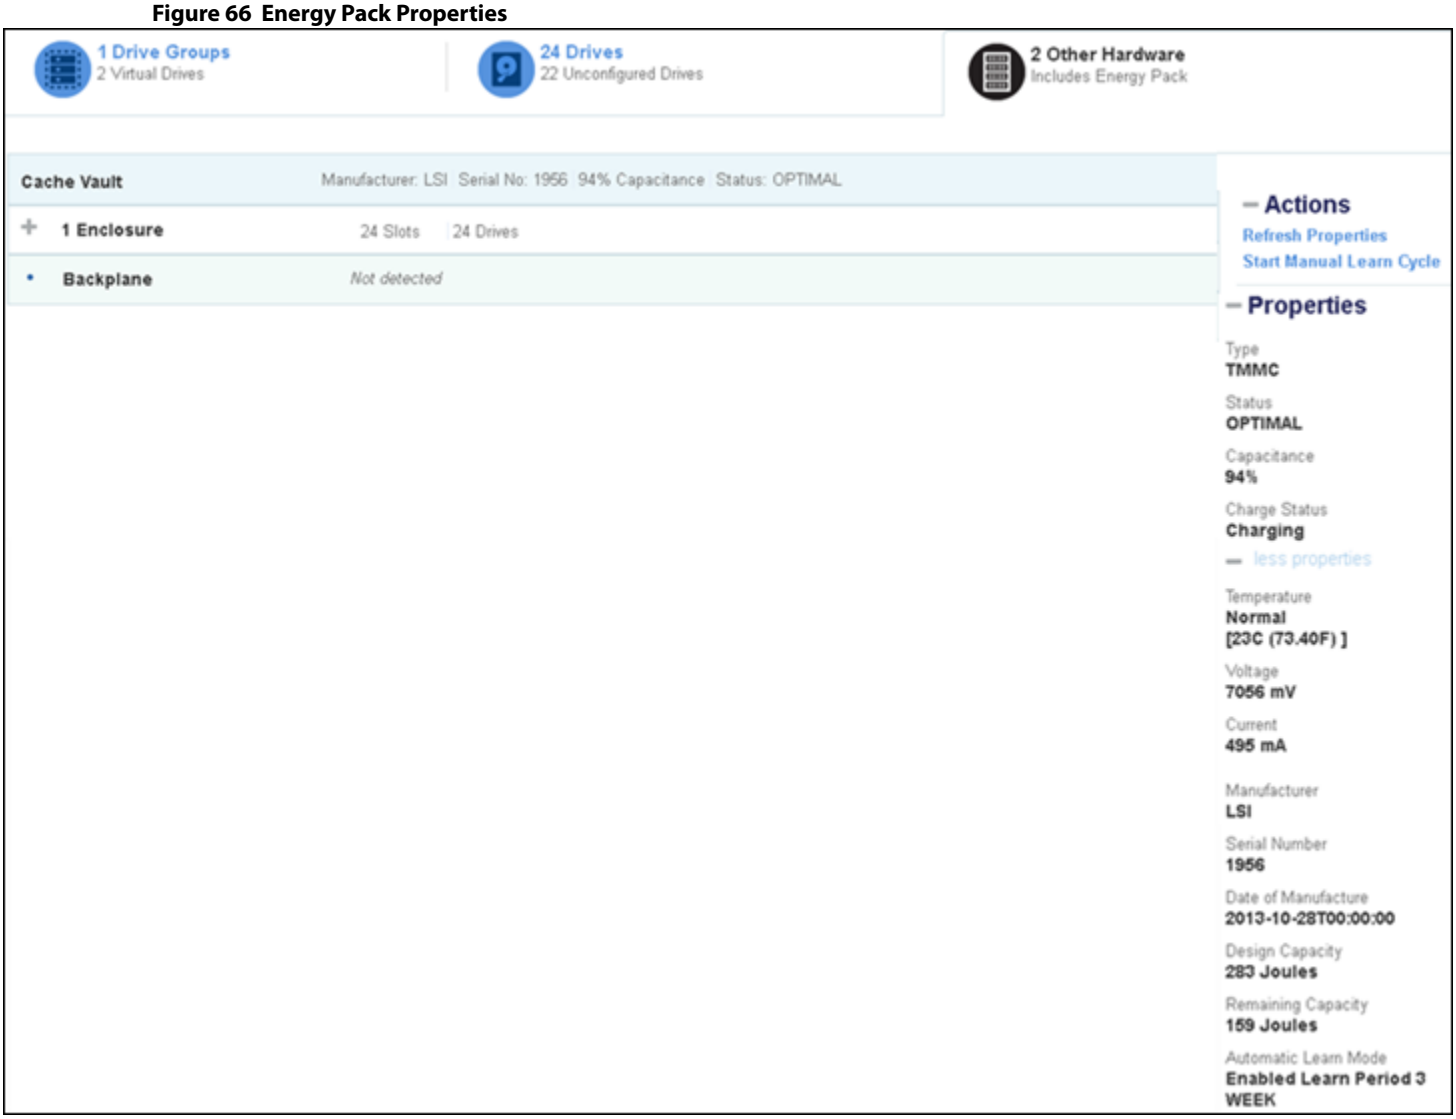

Table 15 Energy Pack Properties

Property	Description	MegaRAID	iMegaRAID	[illegible]grg_rpp R_pecr
Type	Type of the battery. For example, TTMC.	Yes	No	No
Status	Current status of the battery. The battery status field has the following states: ■ Optimal ■ Missing ■ Failed ■ Degraded ■ Degraded [Needs Attention] ■ Unknown	Yes	No	No
Capacitance	Available capacitance of the battery, stated as a percentage.	Yes	No	No
Charge Status	Indicates the charge status	Yes	No	No

Table 15 Energy Pack Properties (Continued)

Property	Description	MegaRAID	iMegaRAID	Initiator-Target
Temperature	Indicates the current temperature of the battery. Also indicates whether the current temperature of the battery is normal or high.	Yes	No	No
Voltage	Voltage level of the battery, in mV. Also indicates if the current battery voltage is normal or low.	Yes	No	No
Current	Current of the battery, in mA.	Yes	No	No
Manufacturer	Manufacturer of the battery.	Yes	No	No
Serial Number	Serial number of the battery.	Yes	No	No
Date of Manufacture	Manufacturing date of the battery.	Yes	No	No
Design Capacity	Theoretical capacity of the battery.	Yes	No	No
Remaining Capacity	Remaining capacity of the battery.	Yes	No	No
Automatic Learn Mode	Indicates whether automatic learn mode is enabled or disabled. A learn cycle is a battery calibration operation that the controller performs periodically to determine the battery condition. This operation cannot be disabled.	Yes	No	No
Next Learn Cycle	Date and hour of the next scheduled learn cycle.			

20.1.2 Refresh Properties

Some of the properties, such as temperature, voltage in the **Properties** section do not refresh automatically. You need to manually refresh the **Properties** section to view the latest data. Perform the following steps to refresh the data.

1. Navigate to the Controller dashboard and click the **Other Hardware** tab.
   All of the associated hardware connected to the controller appear.
2. Expand **Energy Pack** and select a energy pack.
3. Select **Actions > Refresh Properties**.
   The properties are updated.

20.1.3 Setting Learn Cycle Properties

Perform the following steps to set automatic learn cycle properties.

1. Navigate to the Controller dashboard and click the **Other Hardware** tab.
   All of the associated hardware connected to the controller appear.
2. Expand **Energy Pack** and select an energy pack.
3. Select **Actions > Set Learn Cycle Properties**.
   The **Set Learn Cycle Properties** dialog appears.

4. In the **Learn Cycle** drop-down list, select the **Enable** option. The other two options are **Disable** and **Warn Via Event**.
   - If you select **Disable**, the automatic learn cycle is disabled. The **Start On** and **Delay next learn cycle by** fields are also disabled.
   - If you select **Warn Via Event**, an event is generated notifying you when to start a learn cycle manually.
   - If a learn cycle is disabled or not scheduled, the value **None** appears in the **Next learn cycle time** field.
   - If a learn cycle is already scheduled, the day of the week, date, and time of the next learn cycle appears in the **Next learn cycle time** field.

   NOTE After selecting **Disable**, if you select **Enable**, the controller firmware resets the energy pack module properties to initiate an immediate learn cycle. The **Next Learn cycle** field is updated only after the energy pack relearn is completed. Once the relearning cycle is completed, the value in the **Next Learn cycle** field displays the new date and the time of the next learn cycle.
5. In the **Start On** field, specify a day and time to start the automatic learn cycle.
6. You can delay the start of the next learn cycle up to 7 days (168 hours) by specifying the day and hours in the **Delay next learn cycle by** field.
7. Click **Save**.

20.1.4 Starting a Learn Cycle Manually

Perform the following steps to start the learn cycle properties manually.

1. Navigate to the Controller dashboard and click the **Other Hardware** tab.
   All of the associated hardware connected to the controller appear.
2. Expand **Energy Pack** and select an energy pack.
3. Select **Actions > Start Manual Learn Cycle**.
   A confirmation message appears.
4. Select **Confirm** and click **Yes, start manual learn cycle** to proceed with the operation.
   The learn cycle operation starts.

20.2 Monitoring Enclosures

When the LSI Storage Authority software is running, you can monitor the status of all of the enclosures connected to the controllers in the server.

20.2.1 Viewing Enclosure Properties

From the **Other Hardware** tab, under **Enclosures**, select an enclosure to view its properties.

The following figure and table describe the enclosure basic and advanced properties.

Figure 67 Enclosure Properties

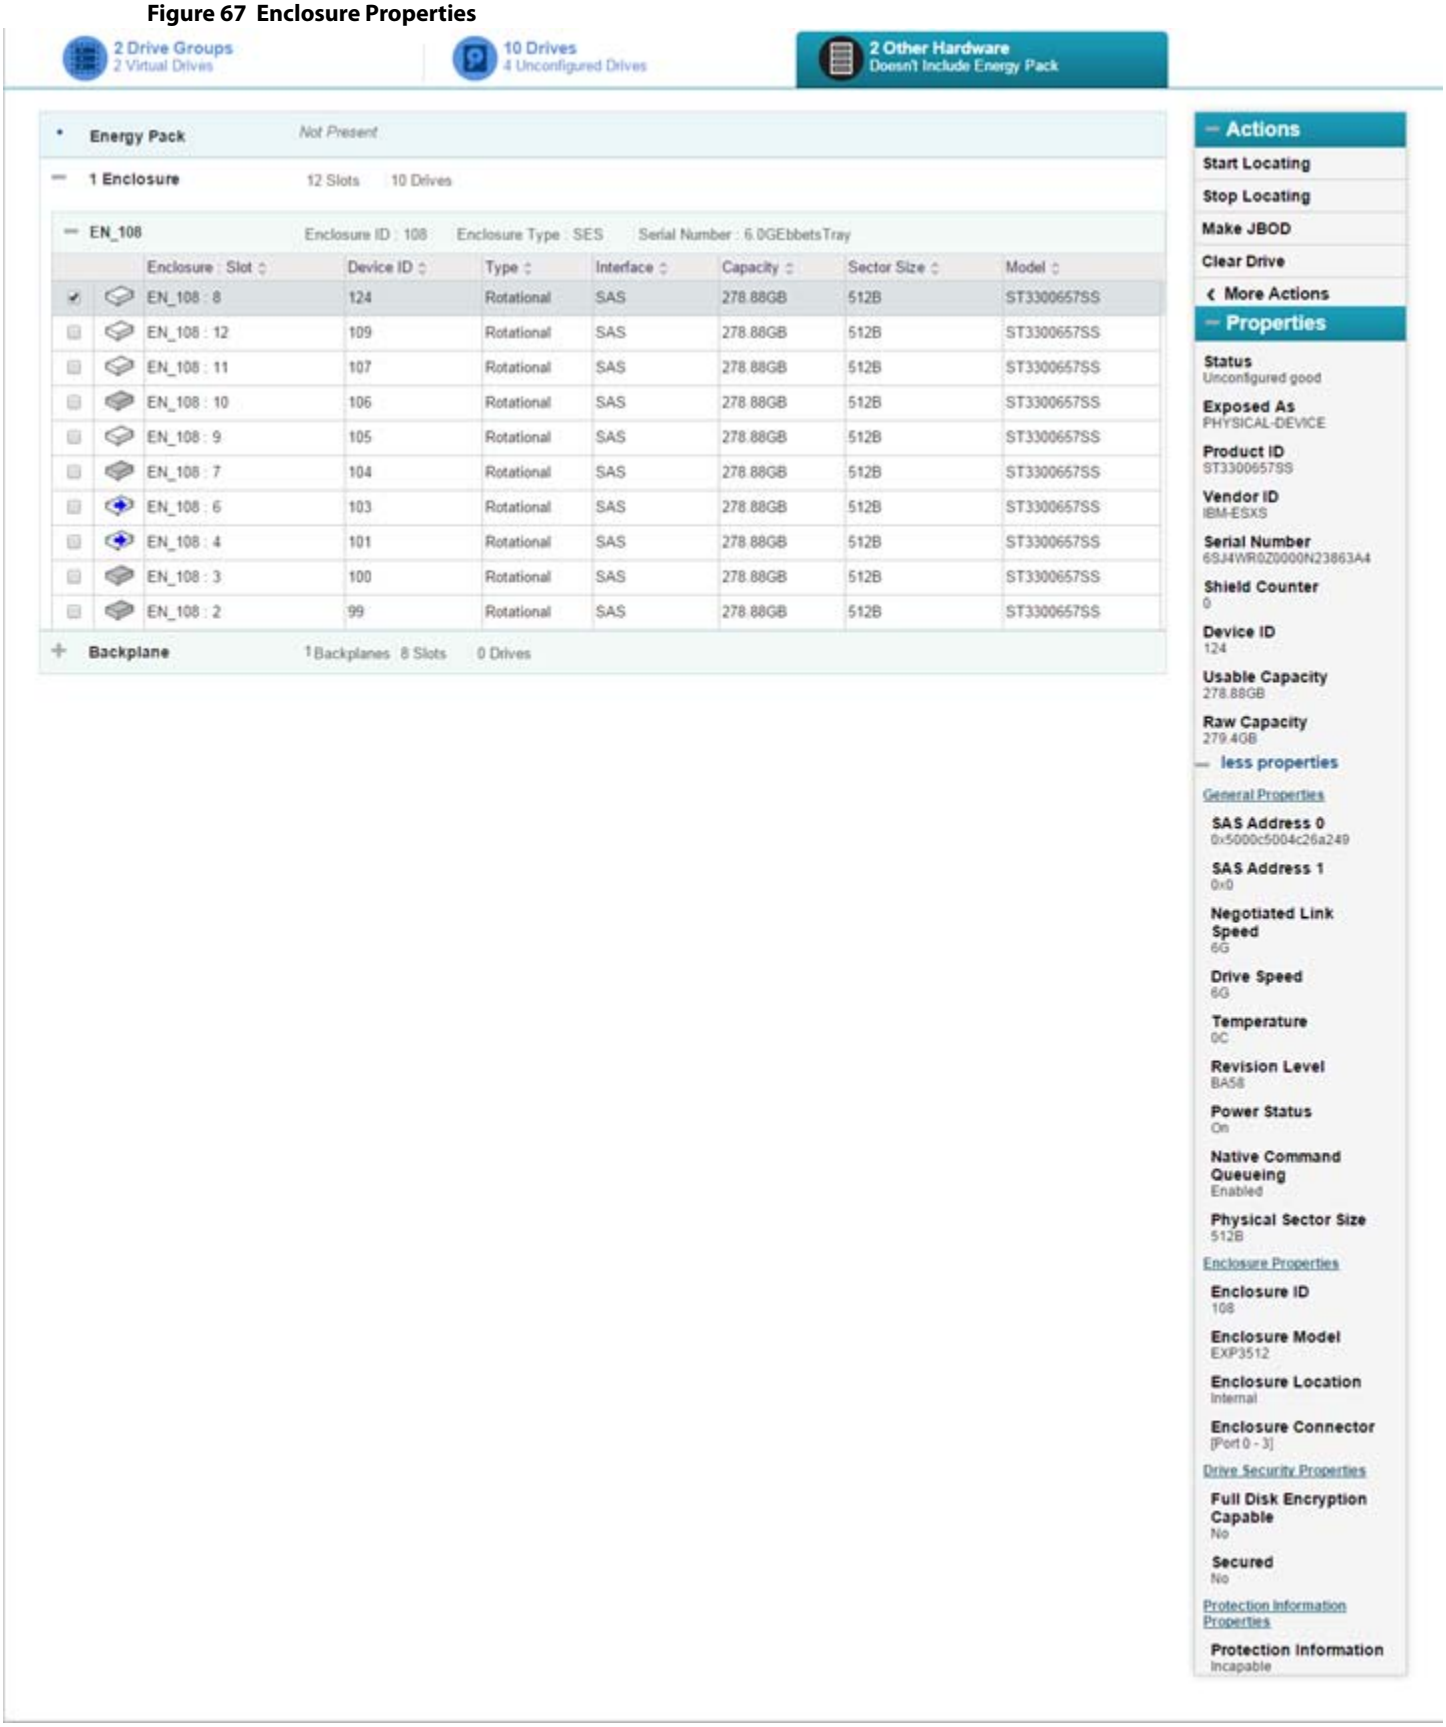

If you have selected multiple virtual drives or multiple physical drives, you have to click on the + (Expand button) to perform actions such as starting a a consistency check and so on. This is applicable for all the scenarios where you have selected multiple virtual drives or multiple physical drives and performing certain actions through the **Actions** dialog.

Table 16 Enclosure Properties

Property	Description	MegaRAID	iMegaRAID	Initiator-Target
Vendor ID	The vendor-assigned ID number of the enclosure.	**Yes**	**Yes**	**Yes**
Enclosure ID	The ID of the enclosure in which the drive is located.	**Yes**	**Yes**	**Yes**[a]
Enclosure Type	Type of the enclosure.	**Yes**	**Yes**	**Yes**
Serial Number	The serial number of the enclosure.	**Yes**	**Yes**	**Yes**
Enclosure Model	The enclosure model.	**Yes**	**Yes**	**Yes**
Enclosure Location	Indicates whether the drive is attached to an internal connector or an external connector of the enclosure.	**Yes**	**Yes**	**Yes**
Enclosure Connector	Indicates the connector name and size of wide port. **Single Path**: `[Port 0-3 ×4]` - Single ×4 **Multipath**: `[Port 0-3 ×4] & [Port 4-7 ×4]`- two ×4 **Wide ports**: `[Port 0-3 & Port 4-7 ×8]`- Single ×8 wide port **Single Drive** `[Port 0-3 x1]`	**Yes**	**Yes**	**No**[b]
Revision Level	The revision level of the enclosure's firmware.	**Yes**	**Yes**	**Yes**
No of Slots	Total number of available slots.	**Yes**	**Yes**	**Yes**
No of Fans	Total number of fans that are connected.	**Yes**	**Yes**	**Yes**
No of Temperature Sensors	Total number of temperature sensors that are connected.	**Yes**	**Yes**	**Yes**
No of Power Supplies	Total number of power supplies that are connected.	**Yes**	**Yes**	**Yes**
No of Voltage Sensors	Total number of voltage sensors that are connected.	**Yes**	**Yes**	**Yes**

a. The Locate feature is not supported by all target enclosures. Check your enclosure's documentation to determine whether the Locate feature is supported on your device.

b. The external wide port address of the SAS HBA is not a fixed address. The auto-configuration feature dynamically assigns the SAS address based on the port linkup sequence and port number. As the address assignment is based on port link sequence, when you do a DC/AC power cycle or move a cable, the port's SAS address may change resulting in the original LUN mapping getting changed.

Chapter 21: Viewing Event Logs

The LSI Storage Authority software monitors the activity and performance of the server and all of the controllers cards attached to it. Perform the following steps to view the event logs.

1. In the Server dashboard or the Controller dashboard, select **Actions > View Event Log**.

   The **View Event Log** window appears that displays a list of events. Each entry has an event ID, a severity level that indicates the severity of the event, a date and time entry, and a brief description of the event. The event logs are sorted by date and time in the chronological order.

Figure 68 View Event Log Window

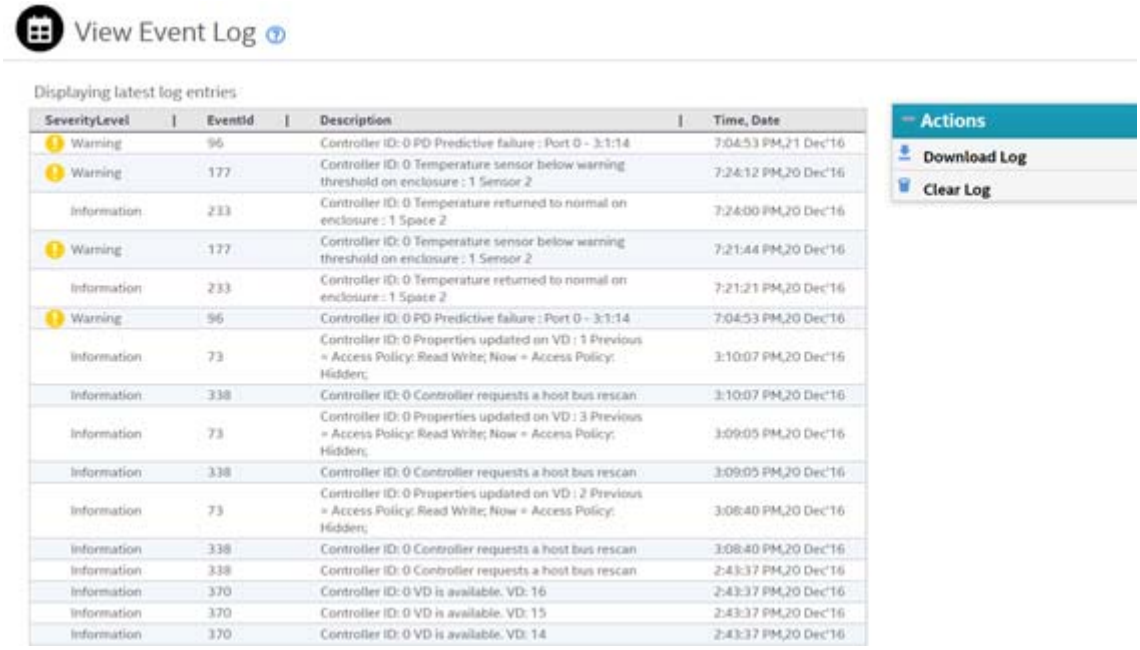

2. (Optional) click **Load More** to view more events in the same page.

21.1 Downloading Logs

To download the event logs, navigate to the **View Event Log** window, then click **Download Log** to download the event log file.

21.2 Clearing the Event Logs

Perform the following steps to clear the event logs.

1. In the **View Event Log** window, click **Clear Log**.

   A confirmation dialog appears.
2. Select **Confirm**, and click **Yes, Clear Log**.

   The event logs are cleared.

Chapter 22: Customizing the Theme of the LSI Storage Authority Software

You can customize the theme of the LSI Storage Authority software to create a uniform look and feel that matches your organization's brand. For example, you can add a company logo or change the default colors. The theme colors are applied globally throughout the software. You can make changes to the following themes:

- Company logo
- Header or banner background color

22.1 Default Theme Settings

The following table lists the default logo, color themes, and their associated values for the UI elements used in the LSA software.

Table 17 Default Theme Settings

Theme	Default	Default File Name/Property Name

Logo	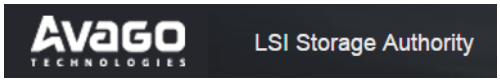	`mainlogo.png`
		`<root>\LSI\LSIStorageAuthority\server\html\ui\images`
		Dimensions
		■ Width - 1172 pixels
		■ Height - 125 pixels
		■ Bit depth - 32
		LSI Storage Authority is present in `<root>\LSI\LSIStorageAuthority\server\html\js\message_en.jsi` n the form of `<Key>:<Value>` format. This value string can be customized.

Header		`headbackground.png`
		`<root>\LSI\LSIStorageAuthority\server\html\ui\images`
		Dimensions
		■ Width - 1172 pixels
		■ Height - 125 pixels
		■ Bit Depth - 32

22.2 Customizing the Logo

Prerequisites

- The new logo must be in the `.png` format.
- Before you begin, make sure that the image already looks the way you want it to appear on the web page.
- Make sure the image has the right size (dimensions 372 x120 pixels)
  - Width - 372 pixels
  - Height -120 pixels

— Bit depth - 32

The logo appears in the header or banner of the software and is visible in all the pages you navigate in the software.

Perform the following steps to change the company logo.

1. Navigate to the Images directory: `<root>\LSI\LSIStorageAuthority\server\html\ui\images`
2. Remove the default logo image file (`mainlogo.png`).
3. Copy the new logo image file.

   NOTE Do not change the file name. Retain the same name, that is `mainlogo.png`.

4. Refresh the browser for the changes to take effect.

22.3 Customizing the Header Background Image

Prerequisites

- The new logo must be in the `.png` format.
- Before you begin, make sure that the image already looks the way you want it to appear on the web page.
- Make sure the image has the right size
  - — Width - 1172 pixels
  - — Height -125 pixels
  - — Bit depth - 32
  - — Dimensions 372x120 pixels

The logo appears in the header or banner of the software and is visible in all the pages you navigate in the software.

Perform the following steps to change the company logo.

1. Navigate to the Images directory: `<root>\LSI\LSIStorageAuthority\server\html\ui\images`
2. Remove the default logo image file (`headbackground.png`).
3. Copy the new logo image file.

   NOTE Do not change the file name. Retain the same name, that is `headbackground.png`.

4. Refresh the browser for the changes to take effect.

Appendix A: Introduction to RAID

Redundant Array of Independent Disks (RAID) is an array, or group of multiple independent physical drives, that provide high performance and fault tolerance because it improves I/O performance and reliability. The RAID drive group appears to the host computer as a single storage unit or as multiple virtual units. I/O 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 that results from a drive failure can be prevented by reconstructing missing data from the remaining drives. RAID improves I/O performance and increases storage subsystem reliability.

RAID Functions

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

Your drives must be organized into virtual drives in a drive group, and they must be able to support the RAID level that you choose. 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
- Enabling Copy back
- Erasing drives
- Performing patrol read
- Updating controller firmware
- Verifying that the redundancy data in virtual drives on RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, RAID 60, PRL-11, or Spanned PRL-11is correct
- Reconstructing virtual drives after changing RAID levels or adding or removing drives to the same drive group
- Selecting a host controller on which to work

A.1 RAID Components and Features

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

A.1.1 Drive Group

A drive group is a group of physical drives. These drives are managed in partitions known as virtual drives. You can create one or more virtual drives on a group of drives attached to a controller card. However, this is based on the support of sliced VD and RAID level of the controller.

A.1.2 Physical Drive States

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

Table 18 Drive States

State	Description
Online	The physical drive is working normally and is a part of a configured logical drive.
Unconfigured Good	A drive that is functioning normally but is not configured as a part of a virtual drive or as a hotspare.
Hotspare	A drive that is powered up and ready for use as a spare in case an online drive fails.
Failed	A fault has occurred in the physical drive, placing it out of service.
Rebuild	A drive to which data is being written to restore full redundancy for a virtual drive.
Unconfigured Bad	A drive on which firmware detects some unrecoverable error; the drive was Unconfigured Good or the drive could not be initialized.
Missing	A drive that was Online but has been removed from its location.
Offline	A drive that is part of a virtual drive but has invalid controller configuration data.

A.1.3 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
- A part of a drive group
- A combination of any two of these conditions

A.1.4 Virtual Drive States

The virtual drive states are described in the following table.

Table 19 Virtual Drive States

State	Description
Optimal	The virtual drive operating condition is good. All configured drives are online.
Degraded	The virtual drive operating condition is not optimal. One of the configured drives has failed or is offline.
Partial Degraded	The operating condition in a RAID 6 and a RAID 60 virtual drive is not optimal. One of the configured drives has failed or is offline. If two drives fail in a RAID 6 drive group or from a single span RAID 60 drive group, the drives become degraded.
Failed	If one drive gets failed from a degraded virtual drive, the virtual drive is failed.
Offline	The virtual drive is not available to the controller card.

A.1.5 Fault Tolerance

Fault tolerance is the capability of the subsystem to undergo a drive failure or failures without compromising data integrity, and processing capability. The MegaRAID controller provides this support through redundant drive groups in RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, RAID 60, PRL-11 and Spanned PRL-11 levels. The system can still work correctly even with a drive failure in a drive group, though performance might be degraded to some extent.

In a span of RAID 1 drive groups, each RAID 1 drive group has two drives and can tolerate one drive failure. RAID 1 drive groups can contain up to 2 drives. 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 in each RAID 6 drive group.

Each span support single drive fault tolerance. A RAID 50 virtual drive can tolerate eight drive failures, as long as each failure is in a separate drive group. RAID 60 drive groups can tolerate up to 16 drive failures in each drive group.

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

Fault tolerance is often associated with system availability because it lets the system 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 that, in case of a disk failure in a redundant RAID drive group, can rebuild the data and re-establish redundancy. 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.

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

A.1.5.1 Multipathing

Firmware supports detecting and using multiple paths from the controller cards to the SAS devices that are in enclosures. Devices connected to enclosures have multiple paths to them. With redundant paths to the same port of a device, if one path fails, another path can communicate between the controller and the device. Using multiple paths with load balancing, instead of a single path, can increase reliability through redundancy.

Multipathing provides the following features:

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

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

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

A.1.5.2 Wide Port

The term "port" is used to identify a single connection point between devices, while the term "wide port" defines a group of individual phys used as a single connection point between SAS initiators, expanders, and/or targets.

A.1.6 Consistency Check

Consistency check verifies the accuracy of the data in virtual drives that use RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, RAID 60, PRL-11, and Spanned PRL-11. RAID 0 does not provide data redundancy. For example, in a system with parity, checking consistency means computing the data on one drive and comparing the results to the contents of the parity drive.

It is recommended to perform a consistency check at least once a month.

A.1.7 Copyback

Copyback lets you copy data from a source drive to a destination drive that is not a part of the virtual drive. Copyback 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 Copyback automatically or manually.

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

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

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

Order of Precedence

In the following scenarios, rebuild takes precedence over Copyback:

- If Copyback is already taking place to a hot spare drive, and any virtual drive on the controller degrades, Copyback aborts, and a rebuild starts. Rebuild changes the virtual drive to the Optimal state.
- The Rebuild takes precedence over Copyback when the conditions exist to start both operations. Consider the following examples:
  - A hotspare drive 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 Copyback, Copyback ends abruptly, and rebuild starts on the hotspare drive.

A.1.8 Background Initialization

Background initialization checks for media errors (soft and hard) on the drives when you create a virtual drive. It is an automatic operation that starts five minutes after you create the virtual drive. This automatic feature might not be supported for all the customers. This check makes sure 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 only a background initialization is forced on new virtual drives.

The new RAID 5 virtual drives and RAID 6 virtual drives require a minimum number of drives for a background initialization to start. If there are fewer drives than the minimum required, the background initialization does not start. The following number of drives are required. However, it is customer-specific:

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

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

A.1.9 Patrol Read

Patrol read reviews your system for possible drive errors that could lead to drive failure and then performs 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 upon 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. It can continue to run during heavy I/O processes.

When Patrol Read starts, the progress bar takes some time to display the actual progress. To inform the user that Patrol Read is started, the progress bar displays the progress status as Unknown. The progress bar displays the actual progress once the actual progress status is available.

A.1.10 Disk Striping

Disk striping lets you write data across multiple drives instead of just one drive. Disk striping partitions each drive storage space into stripes that can vary in size from 8 KB to 1024 KB. These stripes are interleaved in a repeated sequential manner. The combined storage space contains stripes from each drive. You should keep stripe sizes the same across RAID drive groups.

For example, in a four-disk system that uses only disk striping (used in RAID 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 it does not provide data redundancy.

The following figure shows an example of disk striping.

Figure 69 Example of Disk Striping (RAID 0)

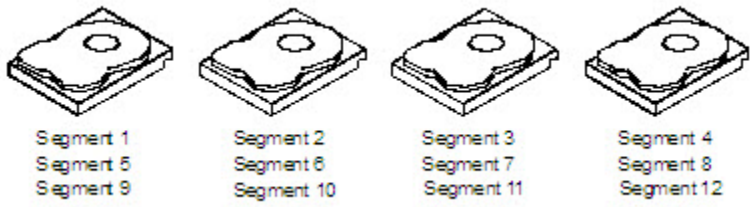

Stripe Width

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

Stripe Size

The stripe size is the length of the interleaved data segments that the controller writes across multiple drives, excluding parity drives. For example, consider a stripe that contains 64 KB of disk space and has 16 KB of data residing on each disk in the stripe. In this case, the stripe size is 64 KB, and the strip size is 16 KB.

Strip Size

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

A.1.11 Disk Mirroring

With disk mirroring (used in RAID 1, RAID 10, PRL-11 and spanned PRL-11), data written to one drive is simultaneously written to another drive. The primary advantage of disk mirroring is that it provides 100-percent data redundancy. Because the contents of the disk are completely written to a second disk, data is not lost if one disk fails. In addition,

both drives contain the same data at all times, so either disk can act as the operational disk. If one disk fails, the contents of the other disk can run the system and reconstruct the failed disk.

The following figure shows an example of disk mirroring.

Figure 70 Example of Disk Mirroring (RAID 1)

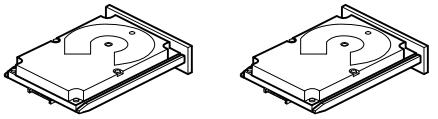

A.1.12 Parity

Parity generates a set of redundancy data from two or more parent data sets. The redundancy data can reconstruct one of the parent data sets in the event of a drive failure. Parity data does not fully duplicate the parent data sets, but parity generation can slow the write process. In RAID, this method is applied to entire drives or stripes across all of the drives in a drive group. The following table describes the types of parity.

Table 20 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.

RAID 5 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. RAID 5 uses parity to provide redundancy for one drive failure without duplicating the contents of entire drives. RAID 6 also uses distributed parity and disk striping, but it adds a second set of parity data so that the drive can survive up to two drive failures.

Figure 71 Example of Distributed Parity (RAID 5)

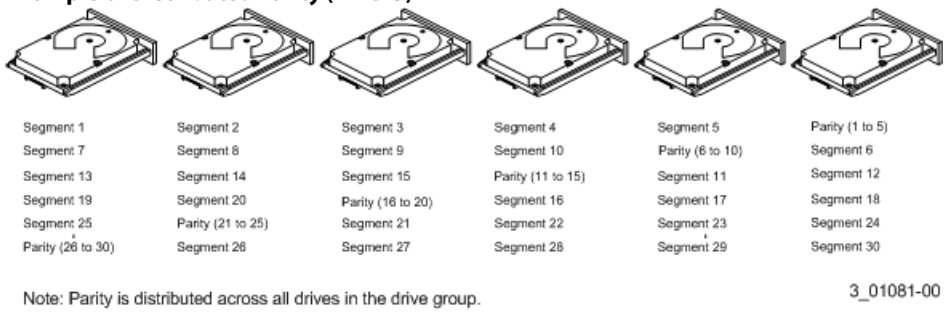

A.1.13 Disk Spanning

Disk spanning lets multiple drives 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, you can combine four 20-GB drives 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.

ATTENTION Even if one span fails, the entire virtual drives will go off line and data will be lost.

Figure 72 Example of Disk Spanning

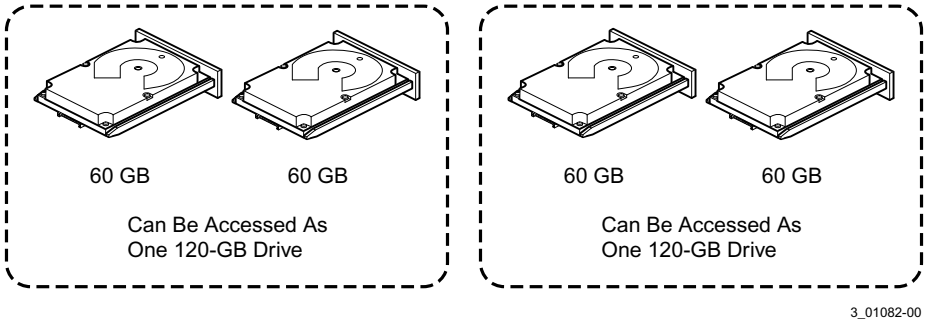

Spanning two contiguous RAID 0 virtual drives does not produce a new RAID level or add fault tolerance. It increases the capacity of the virtual drive and improves performance by doubling the number of spindles.

Spanning for RAID 10, RAID 50, RAID 60, and Spanned PRL-11

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

Table 21 Spanning for RAID 00, RAID 10, RAID 50, and RAID 60

Level	Description
00	Configure RAID 00 by spanning two contiguous RAID 0 virtual drives, up to the maximum number of supported devices for the controller.
10	Configure RAID 10 by spanning two contiguous RAID 1 virtual drives, up to the maximum number of supported devices for the controller. RAID 10 supports a maximum of 16 drives (8 spans X 2). 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 RAID 50 by spanning two contiguous RAID 5 virtual drives. The RAID 5 virtual drives must have the same stripe size.
60	Configure RAID 60 by spanning two contiguous RAID 6 virtual drives. The RAID 6 virtual drives must have the same stripe size.

NOTE In a spanned virtual drive (RAID 00, RAID 10, RAID 50, RAID 60, and Spanned PRL-11), the span numbering starts from Span 0, Span 1, Span 2, and so on.

A.1.14 Hot Spares

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

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

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

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

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

The hot spares are of two types:

- **Global Hot Spare**
- **Dedicated Hot Spare**

Observe the following parameters when using hot spares:

- Hot spares are used only in drive groups with redundancy, which include RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, RAID 60, PRL-11, and Spanned PRL-11drive groups.
- You must assign the hot spare to one or more drives through the controller BIOS or use drive group management software to place it in the hot spare pool.

A.1.15 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 recreates the data using the data stored on the other drives in the drive group. Rebuilding can be done only in drive groups with data redundancy, which include RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, RAID 60, PRL-11, and Spanned PRL-11drive 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 starts automatically when a drive fails. If a hot spare is not available, you must replace the failed drive 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 to a hot spare starts. If the system goes down during a rebuild, the RAID controller automatically resumes the rebuild after the system reboots.

When the rebuild to a hot spare starts, the failed drive often is removed from the virtual drive before management applications detect the failed drive. When the rebuild occurs, the event logs show the drive rebuilding to the hot spare without showing the failed drive. The formerly failed drive is marked as *ready* after a rebuild starts to a hot spare. If a source drive fails during a rebuild to a hot spare, the rebuild 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 fails, because of a source drive failure, the dedicated hot spare is still dedicated and assigned to the correct drive group, and the global hot spare is still global.

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

A.1.16 Rebuild Rate

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

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

A.1.17 Hot Swap

A hot swap manually replaces a defective drive unit when the computer is still running. When a new drive is installed, a rebuild occurs automatically if these situations occur:

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

You can configure the controller to detect the new drives and automatically rebuild the contents of the drive.

A.1.18 Enclosure Management

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

A.2 RAID Levels

The subsequent sections describe the RAID levels in detail.

A.2.1 Summary of RAID Levels

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

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

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

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

RAID 10, a combination of RAID 0 and RAID 1, contains 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. RAID 10 allows a maximum of eight 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. RAID 10 provides high data throughput and complete data redundancy, but it uses a larger number of spans.

RAID 50, a combination of RAID 0 and RAID 5, 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.

NOTE Having virtual drives of different RAID levels, such as RAID 0 and RAID 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 must be RAID 5 only.

RAID 60, a combination of RAID 0 and RAID 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.

NOTE RAID 50 and RAID 60 work best with data that requires high reliability, high request rates, high data transfers, and medium-to-large capacity.

A.2.2 Selecting a RAID Level

To make sure of the best performance, you must choose the optimal RAID level when you create a system drive. The optimal RAID level for your drive group depends on a number of factors:

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

A.2.3 RAID 0

RAID 0 provides disk striping across all drives in the RAID drive group. RAID 0 does not provide any data redundancy, but RAID 0 offers the best performance of any RAID level. RAID 0 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. RAID 0 offers high bandwidth.

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

RAID 0 does not perform parity calculations to complicate the write operation. This situation makes RAID 0 ideal for applications that require high bandwidth but do not require fault tolerance. The following table provides an overview of RAID 0. The following figure shows an example of a RAID 0 drive group.

Table 22 RAID 0 Overview

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

Figure 73 RAID 0 Drive Group Example with Two Drives

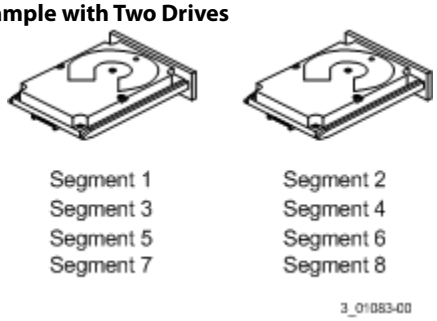

A.2.4 RAID 1

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

Table 23 RAID 1 Overview

Uses	Use RAID 1 for small databases or any other environment that requires fault tolerance but small capacity.
Strong points	Provides complete data redundancy. RAID 1 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 74 RAID 1 Drive Group

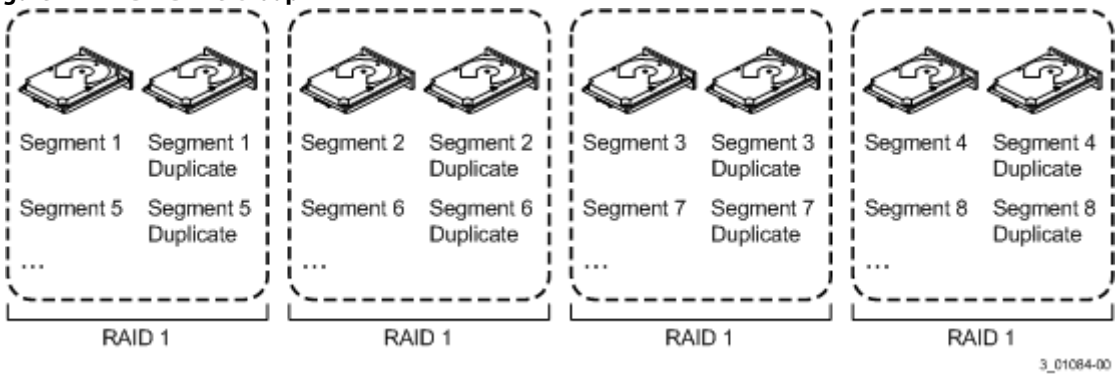

A.2.5 RAID 5

RAID 5 includes disk striping at the block level and parity. Parity is the data's property of being odd or even, and parity checking detects errors in the data. In RAID 5, the parity information is written to all drives. RAID 5 is best suited for networks that perform many small I/O transactions simultaneously.

RAID 5 addresses the bottleneck issue for random I/O operations. Because each drive contains both data and parity, numerous writes can take place concurrently.

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

Table 24 RAID 5 Overview

Uses	Provides high data throughput, especially for large files. Use RAID 5 for transaction-processing applications because each drive can read and write independently. If a drive fails, the controller card uses the parity drive to re-create all missing information. Also use it for office automation and online customer service that requires fault tolerance. Use it for any application that has high read request rates but low write request rates.
Strong points	Provides data redundancy, high read rates, and good performance in most environments. Provides redundancy with the lowest loss of capacity.
Drives	3 through 32.

Figure 75 RAID 5 Drive Group with Six Drives

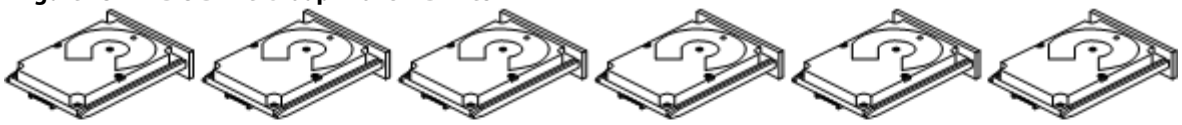

Segment 1	Segment 2	Segment 3	Segment 4	Segment 5	Parity (1–5)
Segment 7	Segment 8	Segment 9	Segment 10	Parity (6–10)	Segment 6
Segment 13	Segment 14	Segment 15	Parity (11–15)	Segment 11	Segment 12
Segment 19	Segment 20	Parity (16–20)	Segment 16	Segment 17	Segment 18
Segment 25	Parity (21–25)	Segment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

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

3_01085-00

A.2.6 RAID 6

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

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

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

Table 25 RAID 6 Overview

Uses	RAID 6 for office automation and online customer service that requires fault tolerance. Use it for any application that has high read request rates but low write request rates.
Strong points	Provides data redundancy, high read rates, and good performance in most environments. Can survive the loss of two drives or the loss of a drive while another drive is being rebuilt. Provides the highest level of protection against drive failures of all of the RAID levels. The read performance is similar to that of RAID 5.
Weak points	Not well-suited to tasks that require many writes. 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 writes. 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. RAID 6 costs more because of the extra capacity required by using two parity blocks per stripe.
Drives	3 through 32.

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

Figure 76 Distributed Parity across Two Blocks in a Stripe (RAID 6)

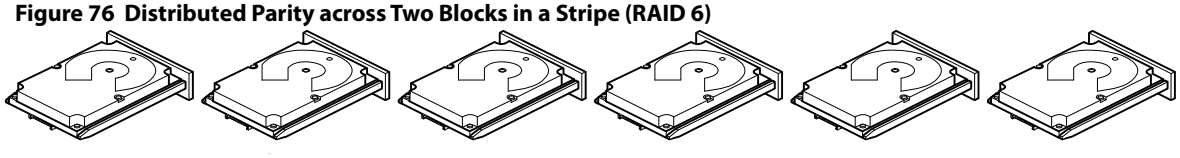

Segment 1	Segment 2	Segment 3	Segment 4	Parity (P1–P4)	Parity (Q1–Q4)
Segment 6	Segment 7	Segment 8	Parity (P5–P8)	Parity (Q5–Q8)	Segment 5
Segment 11	Segment 12	Parity (P9–P12)	Parity (Q9–Q12)	Segment 9	Segment 10
Segment 16	Parity (P13–P16)	Parity (Q13–Q16)	Segment 13	Segment 14	Segment 15
Parity (P17–P20)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

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

3_01086-00

A.2.7 RAID 00

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

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

By breaking up a large file into smaller segments, the controller can use both SAS drives and SATA drives to read or write the file faster. RAID 00 does not perform parity calculations to complicate the write operation. This situation makes RAID 00 ideal for applications that require high bandwidth but do not require fault tolerance. The following table provides an overview of RAID 00. The following figure provides a graphic example of a RAID 00 drive group.

Table 26 RAID 00 Overview

Uses	Provides high data throughput, especially for large files. Use it for any environment that does not require fault tolerance.
Strong points	Provides increased data throughput for large files. Does not have capacity loss penalty for parity.
Weak points	Does not provide fault tolerance or high bandwidth. All data lost if any drive fails.
Drives	2 through 240.

Figure 77 RAID 00 Drive Group Example with Two Drives

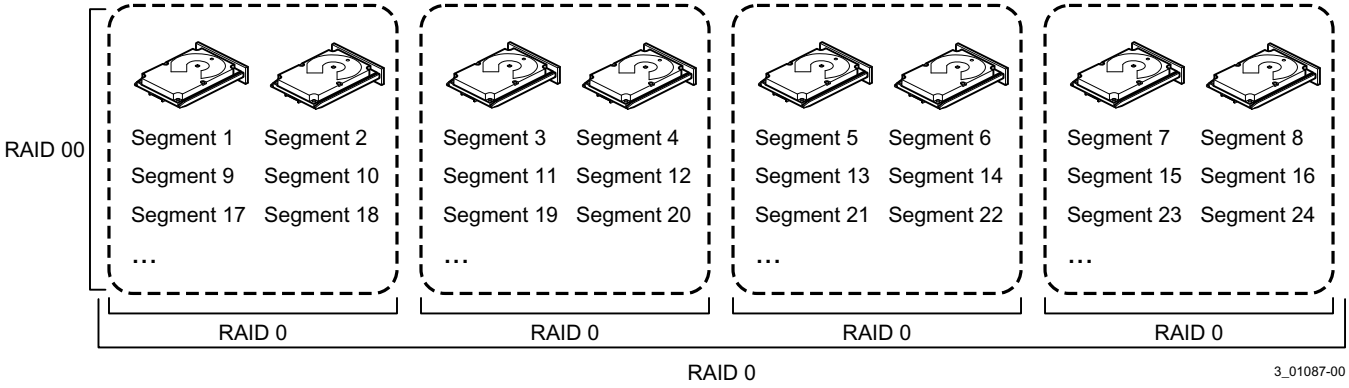

3_01087-00

A.2.8 RAID 10

RAID 10 is a combination of RAID 0 and RAID 1, and it consists of stripes across mirrored drives. RAID 10 breaks up data into smaller blocks and 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 stripe size parameter determines the size of each block, which is set during the creation of the RAID set. The RAID 1 virtual drives must have the same stripe size.

Spanning is used because one virtual drive is defined across more than one drive group. Virtual drives defined across multiple RAID 1 level drive groups are referred to as RAID 10 (RAID 1+RAID 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 a single drive failure. If drive failures occur, less than the total drive capacity is available.

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

Factors, such as the type of controller, can restrict the number of drives supported by RAID 10 virtual drives. A maximum of16 drives are supported on MR/iMR controllers. For IR3/SWR controllers, less than 16 drives are supported, depending on the controller and the number of physical drives the controller supports.

The following table provides an overview of RAID 10.

Table 27 RAID 10 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). RAID 10 works well for medium-sized databases or any environment that requires a higher degree of fault tolerance and moderate-to-medium capacity.
Strong Points	Provides both high data transfer rates and complete data redundancy.
Drives	4 to 32 in multiples of 4 – The maximum number of drives supported by the controller (using an even number of drives in each RAID 10 virtual drive in the span). The MegaRAID/iMegaRAID controller supports 16 drives.

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

Figure 78 RAID 10 Level Virtual Drive

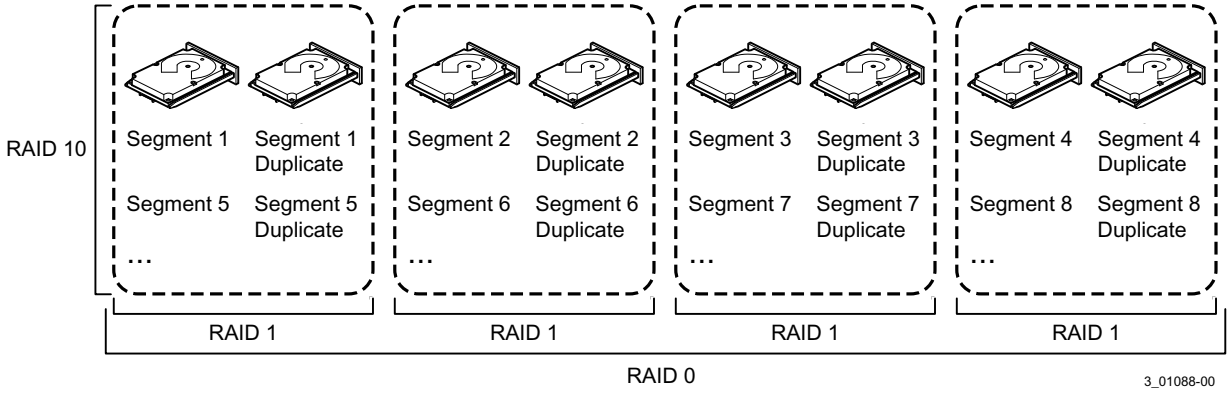

A.2.9 RAID 50

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

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

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

The following table provides an overview of RAID 50.

Table 28 RAID 50 Overview

Uses	Appropriate when used with data that requires high reliability, high request rates, high data transfer, and medium-to-large capacity.
Strong points	Provides high data throughput, data redundancy, and very good performance.
Weak points	Requires two times to eight times as many parity drives as RAID 5.
Drives	Eight spans of RAID 5 drive groups containing 3 to 32 drives per span. However, you can use 256 drives (32x8). The MegaRAID controller supports a total number of 240 drives.

Figure 79 RAID 50 Level Virtual Drive

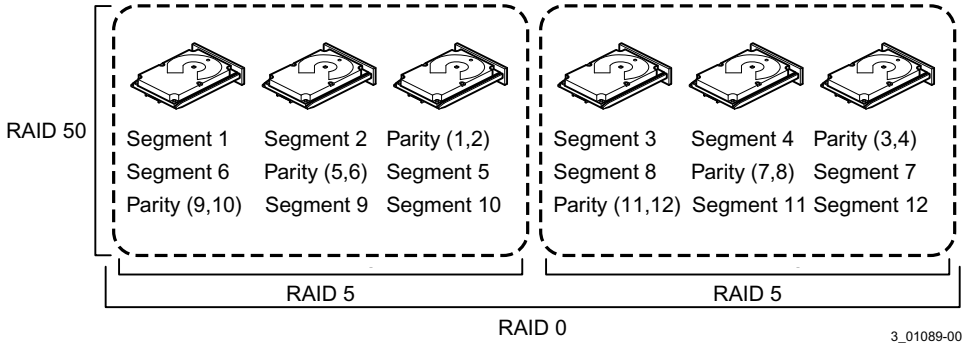

A.2.10 RAID 60

RAID 60 provides the features of both RAID 0 and RAID 6 and includes both distributed parity and drive striping across multiple drive groups. RAID 6 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 sets without losing data. RAID 60 is best implemented on two RAID 6 drive groups with data striped across both drive groups.

RAID 6 breaks up data into smaller blocks, calculates parity by performing an exclusive OR operation on the blocks, and then writes the blocks of data and parity to each drive in the drive group. The e stripe size parameter determines the size of each block, which is set during the creation of the RAID set.

RAID 60 supports up to eight spans and tolerates up to 16 drive failures, though less than the total drive capacity is available. Each RAID 6 level drive group can tolerate two drive failures.

Table 29 RAID 60 Overview

Uses	Provides a high level of data protection through the use of a second parity block in each stripe. Use RAID 60 for data that requires a very high level of protection from loss. In the case of a failure of one drive or two drives in a RAID set in a virtual drive, the controller card uses the parity blocks to re-create all of the missing information. If two drives in a RAID 6 set in a RAID 60 virtual drive fail, two drive rebuilds are required, one for each drive. These rebuilds can occur at the same time. Use RAID 60 for office automation and online customer service that require fault tolerance. Use it for any application that has high read request rates but low write request rates.
Strong points	Provides data redundancy, high read rates, and good performance in most environments. Each RAID 60 set can survive the loss of two drives or the loss of a drive while another drive is being rebuilt. Provides the highest level of protection against drive failures of all of the RAID levels. Read performance is similar to that of RAID 50, though random reads in RAID 60 might be slightly faster because data is spread across at least one more disk in each RAID 60 set.
Weak points	Not well-suited to tasks requiring lot of writes. A RAID 60 virtual drive must generate two sets of parity data for each write operation, which results in a significant decrease in performance during writes. 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. RAID 60 costs more because of the extra capacity required by using two parity blocks per stripe.
Drives	A minimum of 8 drives and maximum of 240 drives.

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 60 parity scheme.

Figure 80 RAID 60 Level Virtual Drive

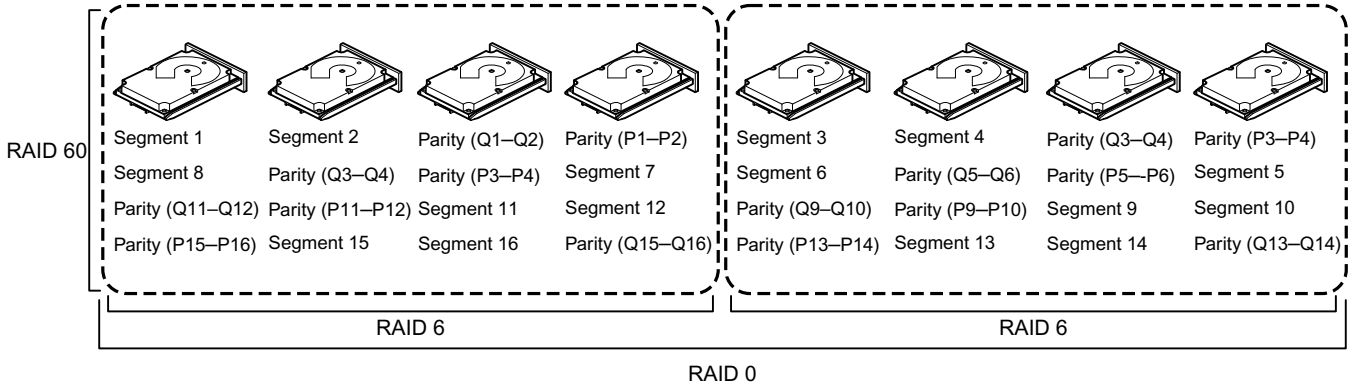

Note: Parity is distributed across all drives in the drive group.

3_01090-00

A.3 RAID Configuration Strategies

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

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

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

A.3.1 Maximizing Fault Tolerance

Fault tolerance is achieved through the ability to perform automatic and transparent rebuilds using hot spare drives and hot swaps. A hot spare drive is an unused online available drive that the controller card 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 occurs.

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

Table 30 RAID Levels and Fault Tolerance

RAID Level	Fault Tolerance
0	Does not provide fault tolerance. All data is lost if any drive fails. Disk striping writes data across multiple drives instead of just one drive. It involves partitioning each drive storage space into stripes that can vary in size. RAID 0 is ideal for applications that require high performance but do not require fault tolerance.
1	Provides complete data redundancy. If one drive fails, the contents of the other drive in the drive group can be used to run the system and reconstruct the failed drive. The primary advantage of disk mirroring is that it provides 100 percent data redundancy. Because the contents of the drive are completely written to a second drive, no data is lost if one of the drives fails. Both drives contain the same data at all times. RAID 1 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 controller card uses the parity data to reconstruct all missing information. In RAID 5, this method is applied to entire drives or stripes across all drives in a drive group. Using distributed parity, RAID 5 offers fault tolerance with limited overhead.
00	Does not provide fault tolerance. All data in a virtual drive is lost if any drive in that virtual drive fails. Disk striping writes data across multiple drives instead of just one drive. It involves partitioning each drive storage space into stripes that can vary in size. RAID 00 is ideal for applications that require high bandwidth but do not require fault tolerance.
6	Combines distributed parity with disk striping. RAID 6 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 controller card uses the parity data to reconstruct all missing information. In RAID 6, this method is applied to entire drives or stripes across all of the drives in a drive group. Using distributed parity, RAID 6 offers fault tolerance with limited overhead.
10	Provides complete data redundancy using striping across spanned RAID 1 drive groups. RAID 10 works well for any environment that requires the 100 percent redundancy offered by mirrored drive groups. RAID 10 can sustain a drive failure in each mirrored drive group and maintain data integrity.
50	Provides data redundancy using distributed parity across spanned RAID 5 drive groups. RAID 50 includes both parity and disk striping across multiple drives. If a drive fails, the controller card uses the parity data to re-create all missing information. RAID 50 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. RAID 60 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. RAID 60 includes both parity and disk striping across multiple drives. If a drive fails, the controller card uses the parity data to re-create all missing information.

A.3.2 Maximizing Performance

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

Table 31 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. RAID 0 partitions each drive 's storage space into stripes that can vary in size from 8 KB to 1024 KB. These stripes are interleaved in a repeated sequential manner. Disk striping enhances performance because multiple drives are accessed simultaneously.
1	With RAID 1 (mirroring), each drive in the system must be duplicated, which requires more time and resources than striping. Performance is impaired during drive rebuilds.
5	RAID 5 provides high data throughput, especially for large files. Use this RAID level for any application that requires high read request rates, but low write request rates, such as transaction processing applications, because each drive can read and write independently. Because each drive contains both data and parity, numerous writes can take place concurrently. In addition, robust caching algorithms and hardware-based exclusive-or assist make RAID 5 performance exceptional in many different environments.Parity generation can slow the write process, which makes write performance significantly lower for RAID 5 than for RAID 0 or RAID 1. Drive performance is reduced when a drive is being rebuilt. Clustering also can reduce drive performance. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.
6	RAID 6 works best when used with data that requires high reliability, high request rates, and high data transfer. It provides high data throughput, data redundancy, and very good performance. However, RAID 6 is not well-suited to that requires many writes. 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 writes. Drive performance is reduced during a drive rebuild. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.
00	RAID 00 (striping in a spanned drive group) offers excellent performance. RAID 00 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. Striping partitions each drive's storage space into stripes that can vary in size from 8 KB to 1024 KB. These stripes are interleaved in a repeated sequential manner. Disk striping enhances performance because multiple drives are accessed simultaneously.
10	RAID 10 works best for data storage that needs the enhanced I/O performance of RAID 0 (striped drive groups), which provides high data transfer rates. Spanning increases the capacity of the virtual drive and improves performance by doubling the number of spindles. The system performance improves as the number of spans increases. (The maximum number of spans is 8.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans, and RAID performance degrades to that of a RAID 1 or RAID 5 drive group.
50	RAID 50 works best when used with data that requires high reliability, high request rates, and high data transfer. It provides high data throughput, data redundancy, and very good performance. Spanning increases the capacity of the virtual drive and improves performance by doubling the number of spindles. The system performance improves as the number of spans increases. (The maximum number of spans is eight.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans, and RAID performance degrades to that of a RAID 1 or RAID 5 drive group.
60	RAID 60 works best when used with data that requires high reliability, high request rates, and high data transfer. It provides high data throughput, data redundancy, and very good performance. Spanning increases the capacity of the virtual drive and improves performance by doubling the number of spindles. The system performance improves as the number of spans increases. (The maximum number of spans is 8.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans, and RAID performance degrades to that of a RAID 1 or RAID 6 drive group.RAID 60 is not well suited to tasks that requires many 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 writes. 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.

A.3.3 Maximizing Storage Capacity

Storage capacity is an important factor when selecting a RAID level. Consider several variables to consider. Striping alone (RAID 0) requires less storage space than mirrored data (RAID 1) or distributed parity RAID 5. The following table explains the effects of the RAID level on storage capacity.

Table 32 RAID Levels and Capacity

RAID Level	Capacity
0	RAID 0 (striping) partitions each drive's storage space into stripes that can vary in size. The combined storage space is composed of stripes from each drive.

A.4 RAID Availability

A.4.1 RAID Availability Concepts

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 are associated with failed servers. The RAID technology 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 drives or defective drives in a drive group. A replacement drive must be at least as large as the drive it replaces. Spare drives include hot swaps, hot spares, and cold swaps.

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

Hot spare drives are drives that power up along with the RAID drives and operate in a Standby state. If a drive used in a RAID virtual drive fails, a hot spare automatically takes its place, and the data on the failed drive is rebuilt on the hot spare. Hot spares can be used for RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, and RAID 60, PRL-11, and Spanned PRL-11.

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, RAID 5, RAID 6, RAID 10, RAID 50, RAID 60, PRL-11, and Spanned PRL-11virtual drive, you can recover the lost data by rebuilding the drive. If you have configured hot spares, the controller card automatically tries to use them to rebuild failed drives. 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.

A.5 Configuration Planning

The factors to consider when planning a configuration are the number of drives the that the controller card can support, the purpose of the drive group, and the availability of spare drives.

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

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

Appendix B: Events and Messages

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

The LSI Storage Authority software monitors the activity and performance of all of the controllers in the workstation and the devices attached to them. When an event occurs, such as the start of an initialization, an event message appears in the log at the bottom of the Server dashboard or Controller dashboard. The messages are also logged in the Windows Application log (Event Viewer).

B.1 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 33 Event Error Levels

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

B.2 Event Messages

The following table lists all of the event messages. The event message descriptions include placeholders for specific values that are determined when the event is generated. For example, in message No. 1 in the Event Messages table, "%s" is replaced by the firmware version, which is read from the firmware when the event is generated.

Table 34 Event Messages

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

Table 34 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0009	Information	Background initialization rate changed to %d%%	Logged to display background initialization progress indication in percentage.
0x000a	Fatal	Controller cache discarded due to memory/battery problems	Logged on cache discard due to hardware problems.
0x000b	Fatal	Unable to recover cache data due to configuration mismatch	Currently not logged.
0x000c	Information	Cache data recovered successfully	Logged when cache data is successfully recovered after reboot.
0x000d	Fatal	Controller cache discarded due to firmware version incompatibility	Logged when cache data discarded because of firmware version mismatch.
0x000e	Information	Consistency Check rate changed to %d%%	Logged to display Consistency check progress indication percentage.
0x000f	Fatal	Fatal firmware error: %s	Logged in case of fatal errors and also while entering debug monitor.
0x0010	Information	Factory defaults restored	Logged while controller is reset to factory defaults.
0x0011	Information	Flash downloaded image corrupt	Logged to inform downloaded flash image is corrupt.
0x0012	Critical	Flash erase error	Logged in case of flash erase failure, generally after flash update.
0x0013	Critical	Flash timeout during erase	Logged to indicate flash erase operation timed out.
0x0014	Critical	Flash error	Generic unknown internal error during flash update flash.
0x0015	Information	Flashing image: %s	Logged to display flash image name string before getting updated to controller.
0x0016	Information	Flash of new firmware images complete	Logged to inform successful update of flash image(s).
0x0017	Critical	Flash programming error	Logged to notify, write failure during flash update, not being allowed usually due to internal controller settings.
0x0018	Critical	Flash timeout during programming	Logged to indicate flash write operation timed out.
0x0019	Critical	Flash chip type unknown	Logged during flash update tried with unsupported flash chip type.
0x001a	Critical	Flash command set unknown	Logged while unsupported flash command set detected, most likely because of unsupported flash chip.
0x001b	Critical	Flash verify failure	Logged when compare operation fails between written flash data and original data.
0x001c	Information	Flush rate changed to %d seconds	Logged to notify modified cache flush frequency in seconds.
0x001d	Information	Hibernate command received from host	Logged to inform about reception of hibernation command from host to controller, generally during host shutdown.
0x001e	Information	Event log cleared	Logged when controller log has been cleared.
0x001f	Information	Event log wrapped	Logged when controller log has been wrapped around, when the maximum logs are written.
0x0020	Fatal	Multi-bit ECC error: ECAR=%x, ELOG=%x, (%s)	Logged to notify ECC multi bit error in memory, ELOG: ecc info (source, type, syndrome), ECAR:ecc address.
0x0021	Warning	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s)	Logged to notify ECC single bit error in memory, ELOG: ecc info (source, type, syndrome), ECAR:ecc address.
0x0022	Fatal	Not enough controller memory	Logged to notify fatal controller condition, when you run out of memory to allocate.
0x0023	Information	Patrol Read complete	Logged when patrol read completes.

Table 34 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0024	Information	Patrol Read paused	Logged when patrol read is paused.
0x0025	Information	Patrol Read Rate changed to %d%%	Logged to indicate progress of patrol read in percentage.
0x0026	Information	Patrol Read resumed	Logged when patrol read is resumed.
0x0027	Information	Patrol Read started	Logged when patrol read is started.
0x0028	Information	Reconstruction rate changed to %d%%"	Logged to indicate progress of reconstruction in percentage.
0x0029	Information	Drive group modification rate changed to %d%%	Logged to indicate the change in Drive group modification frequency.
0x002a	Information	Shutdown command received from host	Logged when shutdown command is received from host to controller.
0x002b	Information	Test event: %s	General controller event, with a generic string.
0x002c	Information	Time established as %s; (%d seconds since power on)	Logged when controller time was set from host, also displaying time since power on in seconds.
0x002d	Information	User entered firmware debugger	Logged when user enters controller debug shell.
0x002e	Warning	Background Initialization aborted on %s	Logged to inform about user aborted background initialization on displayed LD number.
0x002f	Warning	Background Initialization corrected medium error (%s at %lx	logged to inform about corrected medium error on displayed LD number, LBALBA number, PD number and PDLBA number in that order.
0x0030	Information	Background Initialization completed on %s	Logged to inform Background Initialization completion on displayed LD.
0x0031	Fatal	Background Initialization completed with uncorrectable errors on %s	Logged to inform Background Initialization completion with error on displayed LD.
0x0032	Fatal	Background Initialization detected uncorrectable double medium errors (%s at %lx on %s)	Logged to inform Background Initialization completion with double medium error on displayed PD, PDLBA and LD in that order.
0x0033	Critical	Background Initialization failed on %s	Logged to inform Background Initialization failure on displayed LD.
0x0034	Progress	Background Initialization progress on %s is %s	Logged to inform Background Initialization progress in percentage of displayed LD.
0x0035	Information	Background Initialization started on %s	Logged to inform Background Initialization started for displayed LD.
0x0036	Information	Policy change on %s from %s to %s	Logged to inform the changed policy for displayed LD with old and new policies.
0x0038	Warning	Consistency Check aborted on %s	Logged to inform aborted Consistency check for displayed LD.
0x0039	Warning	Consistency Check corrected medium error (%s at %lx	Logged when Consistency check corrected medium error.
0x003a	Information	Consistency Check done on %s	Logged when Consistency check has completed successfully on the LD.
0x003b	Information	Consistency Check done with corrections on %s	Logged when Consistency check completed and inconsistency was found during check and was corrected.
0x003c	Fatal	Consistency Check detected uncorrectable double medium errors (%s at %lx on %s)	Logged when uncorrectable double medium error are detected while consistency check.
0x003d	Critical	Consistency Check failed on %s	Logged when Consistency check failed as fatal error was found.

Table 34 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x003e	Fatal	Consistency Check completed with uncorrectable data on %s	Logged when Uncorrectable error occurred during consistency check.
0x003f	Warning	Consistency Check found inconsistent parity on %s at strip %lx	Logged when consistency check finds inconsistency parity on a strip.
0x0040	Warning	Consistency Check inconsistency logging disabled on %s (too many inconsistencies)	Logged when consistency check finds too many inconsistent parity (greater than 10) and the inconsistency parity logging is disabled.
0x0041	Progress	Consistency Check progress on %s is %s	Logs Consistency Check progress, the progress is logged only if the progress is greater than 1% at an interval of every 15 seconds.
0x0042	Information	Consistency Check started on %s	Logged when consistency check has started
0x0043	Warning	Initialization aborted on %s	Logged when consistency check is aborted by you or for some other reason.
0x0044	Critical	Initialization failed on %s	Logged when initialization has failed.
0x0045	Progress	Initialization progress on %s is %s	Logs initialization progress, the progress is logged only if the progress is greater than 1% at an interval of every 15 seconds.
0x0046	Information	Fast initialization started on %s	Logged when quick initialization has started on a LD. The parameter to decide Quick init or Full init is passed by you.
0x0047	Information	Full initialization started on %s	Logged when full initialization has started.
0x0048	Information	Initialization complete on %s	Logged when initialization has completed successfully.
0x0049	Information	LD Properties updated to %s (from %s)	Logged when LD properties has been changed.
0x004a	Information	Reconstruction complete on %s	Logged when reconstruction has completed successfully.
0x004b	Fatal	Reconstruction of %s stopped due to unrecoverable errors	Logged when reconstruction has finished because of failure (unrecoverable errors).
0x004c	Fatal	Reconstruct detected uncorrectable double medium errors (%s at %lx on %s at %lx)	Logged while reconstructing if an unrecoverable double medium error is encountered.
0x004d	Progress	Reconstruction progress on %s is %s	Logs reconstruction progress, the progress is logged only if the progress is greater than 1% at an interval of every 15 seconds.
0x004e	Information	Reconstruction resumed on %s	Logged when reconstruction resumes after a power cycle.
0x004f	Fatal	Reconstruction resume of %s failed due to configuration mismatch	Logged when reconstruction resume failed due to configuration mismatch.
0x0050	Information	Reconstruction started on %s	Logged on start of reconstruction on a LD.
0x0051	Information	State change on %s from %s to %s	Logged when there is change in LD state. The event gives the new and old state. The state could be one of the following, LDS_OFFLINE, LDS_PARTIALLY_DEGRADED, LDS_DEGRADED, LDS_OPTIMAL.
0x0052	Information	Drive Clear aborted on %s	Logged when PD clear is aborted.
0x0053	Critical	Drive Clear failed on %s (Error %02x)	Logged when drive clear is failed and the even is logged along with error code.
0x0054	Progress	Drive Clear progress on %s is %s	Logs drive clear progress, the progress is logged only if the progress is greater than 1% at an interval of every 15 seconds.
0x0055	Information	Drive Clear started on %s	Logged when drive clear started on a PD.
0x0056	Information	Drive Clear completed on %s	Logged when PD clear task is completed successfully on a PD.

Table 34 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0057	Warning	Error on %s (Error %02x)	Logged if Read returns with Uncorrectable error or same errors on both the drives or write long returns with an error (ie. puncture operation could failed).
0x0058	Information	Format complete on %s	Logged when Format has completed.
0x0059	Information	Format started on %s	Logged when format unit is started on a PD.
0x005a	Critical	Hot Spare SMART polling failed on %s (Error %02x)	Currently not logged.
0x005b	Information	Drive inserted: %s	Logged when drive is inserted and slot/enclosure fields of PD are updated.
0x005c	Warning	Drive %s is not supported	Logged when the drive is not supported; reason could be the number of drive has exceeded the MAX supported drives or an unsupported drive is inserted like a SATA drive in SAS only enclosure or could be a unsupported drive type.
0x005d	Warning	Patrol Read corrected medium error on %s at %lx	Logged when Patrol read has successfully completed recovery read and recovered data.
0x005e	Progress	Patrol Read progress on %s is %s	Logs patrol read progress, the progress is logged only if the progress is greater than 1% at an interval of every 15 seconds.
0x005f	Fatal	Patrol Read found an uncorrectable medium error on %s at %lx	Logged when Patrol read is unable to recover data.
0x0060	Critical	Predictive failure: CDB: %s	Logged when a failure is found during smart (predictive failure) poll.
0x0061	Fatal	Patrol Read puncturing bad block on %s at %lx	Logged when patrol read punctures a block due to unrecoverable medium error.
0x0062	Information	Rebuild aborted by user on %s	Logged when the user aborts a rebuild operation.
0x0063	Information	Rebuild complete on %s	Logged when the rebuild operation on a logical drive on a physical drive (which can have multiple LDs) is completed.
0x0064	Information	Rebuild complete on %s	Logged when rebuild operation is completed for all logical drives on a given physical drive.
0x0065	Critical	Rebuild failed on %s due to source drive error	Logged if one of the source drives for the rebuild operation fails or is removed.
0x0066	Critical	Rebuild failed on %s due to target drive error	Logged if the target rebuild drive (on which rebuild operation is going on) fails or is removed from the controller.
0x0067	Progress	Rebuild progress on %s is %s	Logged to indicate the progress (in percentage) of the rebuild operation on a given physical drive.
0x0068	Information	Rebuild resumed on %s	Logged when the rebuild operation on a physical drive resumes.
0x0069	Information	Rebuild started on %s	Logged when the rebuild operation is started on a physical drive.
0x006a	Information	Rebuild automatically started on %s	Logged when the rebuild operation kicks in on a spare.
0x006b	Critical	Rebuild stopped on %s due to loss of cluster ownership	Logged when the rebuild operation is stopped due to loss of ownership.
0x006c	Fatal	Reassign write operation failed on %s at %lx	Logged when a check condition or medium error is encountered for a reassigned write.
0x006d	Fatal	Unrecoverable medium error during rebuild on %s at %lx	Logged when the rebuild I/O encounters an unrecoverable medium error.
0x006e	Information	Corrected medium error during recovery on %s at %lx	Logged when recovery completed successfully and fixed a medium error.

Table 34 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x006f	Fatal	Unrecoverable medium error during recovery on %s at %lx	Logged when the recovery for a failed I/O encounters a medium error.
0x0070	Information	Drive removed: %s	Logged when a drive is removed from the controller.
0x0071	Warning	Unexpected sense: %s, CDB%s, Sense: %s	Logged when an I/O fails due to unexpected reasons and sense data needs to be logged.
0x0072	Information	State change on %s from %s to %s	Logged when the state of a drive is changed by the firmware or by you.
0x0073	Information	State change by user on %s from %s to %s	Not logged by the firmware.
0x0074	Warning	Redundant path to %s broken	Not logged by the firmware.
0x0075	Information	Redundant path to %s restored	Not logged by the firmware
0x0076	Information	Dedicated Hot Spare Drive %s no longer useful due to deleted drive group	Not logged by the firmware.
0x0077	Critical	SAS topology error: Loop detected	Logged when device discovery fails for a SAS device as a loop was detected.
0x0078	Critical	SAS topology error: Unaddressable device	Logged when device discovery fails for a SAS device as an unaddressable device was found.
0x0079	Critical	SAS topology error: Multiple ports to the same SAS address	Logged when device discovery fails for a SAS device multiple ports with same SAS address were detected.
0x007a	Critical	SAS topology error: Expander error	Not logged by the firmware.
0x007b	Critical	SAS topology error: SMP timeout	Logged when device discovery fails for a SAS device due to SMP timeout.
0x007c	Critical	SAS topology error: Out of route entries	Logged when device discovery fails for a SAS device as expander route table is out of entries.
0x007d	Critical	SAS topology error: Index not found	Logged when device discovery fails for a SAS device as expander route table out of entries.
0x007e	Critical	SAS topology error: SMP function failed	Logged when device discovery fails for a SAS device due to SMP function failure.
0x007f	Critical	SAS topology error: SMP CRC error	Logged when device discovery fails for a SAS device due to SMP CRC error.
0x0080	Critical	SAS topology error: Multiple subtractive	Logged when device discovery fails for a SAS device as a subtractive-to-subtractive link was detected.
0x0081	Critical	SAS topology error: Table to table	Logged when device discovery fails for a SAS device as table-to-table link was detected.
0x0082	Critical	SAS topology error: Multiple paths	Not logged by the firmware.
0x0083	Fatal	Unable to access device %s	Logged when the inserted drive is bad and unusable.
0x0084	Information	Dedicated Hot Spare created on %s (%s)	Logged when a drive is configured as a dedicated spare.
0x0085	Information	Dedicated Hot Spare %s disabled	Logged when a drive is removes as a dedicated spare.
0x0086	Critical	Dedicated Hot Spare %s no longer useful for all drive groups	Logged when an array with a dedicated spare is resized. The hot spare (dedicated to this array and possibly others) will not be applicable to other arrays.
0x0087	Information	Global Hot Spare created on %s (%s)	Logged when a drive is configured as a global hot spare.
0x0088	Information	Global Hot Spare %s disabled	Logged when a drive configured as global host spare fails or is unconfigured by you.

Table 34 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0089	Critical	Global Hot Spare does not cover all drive groups	Logged when the global hotspare is too small (or doesn't meet the SAS/SATA restricitons) to cover certain arrays.
0x008a	Information	Created %s}	Logged as soon as the new logical drive created is added to the firmware configuration.
0x008b	Information	Deleted %s}	Logged when the firmware removes an LD from its configuration upon a user request from the applications.
0x008c	Information	Marking LD %s inconsistent due to active writes at shutdown	Logged when we have active writes on one of the target disks of a Raid 5 LD at the time of shutdown.
0x008d	Information	Battery Present	Logged during firmware initialization when we check if there is a battery present and the check turns out true. This event is also logged when a battery is inserted or replaced with a new one and the battery present check returns true.
0x008e	Warning	Battery Not Present	Logged if the user has not disabled "Battery Not Present" warning at the boot time or if a battery has been removed.
0x008f	Information	New Battery Detected	Logged when we have a subsequent boot after a new battery has been inserted.
0x0090	Information	Battery has been replaced	Logged when a new battery has been replaced with an old battery.
0x0091	Critical	Battery temperature is high	Logged when we detect that the battery temperature is high during the periodic battery status check.
0x0092	Warning	Battery voltage low	Not logged by the firmware.
0x0093	Information	Battery started charging	Logged as part of monitoring the battery status when the battery is getting charged.
0x0094	Information	Battery is discharging	Logged as part of monitoring the battery status when the battery is getting discharged.
0x0095	Information	Battery temperature is normal	Logged as part of monitoring the battery status when the temperature of the battery is normal.
0x0096	Fatal	Battery has failed and cannot support data retention. Please replace the battery.	Logged when there is not enough capacity left in battery for expected data retention time. Battery has to be replaced.
0x0097	Information	Battery relearn started	logged when the battery relearn started, initiated either by the user or automatically.
0x0098	Information	Battery relearn in progress	Logged as part of monitoring the battery status when the battery relearn is in progress.
0x0099	Information	Battery relearn completed	Logged as part of monitoring the battery status when the battery relearn is complete.
0x009a	Critical	Battery relearn timed out	Not logged by the firmware.
0x009b	Information	Battery relearn pending: Battery is under charge	Logged as part of monitoring the battery status when the battery relearn is requested but yet to start.
0x009c	Information	Battery relearn postponed	Logged as part of monitoring the battery status when the battery relearn is requested but postponed as there is valid pinned cache present. This event can also be logged when learn delay interval has been explicitly set.
0x009d	Information	Battery relearn will start in 4 days	Logged as part of providing battery learn cycle information when auto learn is enabled.
0x009e	Information	Battery relearn will start in 2 day	Logged as part of providing battery learn cycle information when auto learn is enabled.
0x009f	Information	Battery relearn will start in 1 day	Logged as part of providing battery learn cycle information when auto learn is enabled.

Table 34 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x00a0	Information	Battery relearn will start in 5 hours	Logged as part of providing battery learn cycle information when auto learn is enabled.
0x00a1	Information	Battery removed	Logged as part of periodic monitoring of the battery status when a battery has been removed.
0x00a2	Information	Current capacity of the battery is below threshold	Logged as part of monitoring the battery status when the capacity of the battery is below threshold.
0x00a3	Information	Current capacity of the battery is above threshold	Logged as part of monitoring the battery status when the capacity of the battery is above threshold.
0x00a4	Information	Enclosure (SES) discovered on %s	Logged when an Enclosure (SES) is discovered for the first time.
0x00a5	Information	Enclosure (SAFTE) discovered on %s	Not logged by the firmware.
0x00a6	Critical	Enclosure %s communication lost	Logged when the communication with an enclosure has been lost.
0x00a7	Information	Enclosure %s communication restored	Logged when the communication with an enclosure has been restored
0x00a8	Critical	Enclosure %s fan %d failed	Logged when an enclosure fan has failed.
0x00a9	Information	Enclosure %s fan %d inserted	Logged when an enclosure fan has been inserted newly.
0x00aa	Critical	Enclosure %s fan %d removed	Logged when an enclosure fan has been removed.
0x00ab	Critical	Enclosure %s power supply %d failed	Not logged by the firmware.
0x00ac	Information	Enclosure %s power supply %d inserted	Logged when power supply has been inserted to an enclosure.
0x00ad	Critical	Enclosure %s power supply %d removed	Logged when power supply has been removed from an enclosure.
0x00ae	Critical	Enclosure %s SIM %d failed	Logged when the enclosure SIM has failed.
0x00af	Information	Enclosure %s SIM %d inserted	Logged when an enclosure SIM has been inserted.
0x00b0	Critical	Enclosure %s SIM %d removed	Logged when an enclosure initialization was completed but later the SIM was removed.
0x00b1	Warning	Enclosure %s temperature sensor %d below warning threshold	Logged when the enclosure services process has detected a temperature lower than a normal operating temperature or lower than the value indicated by the LOW WARNING THRESHOLD field in the Threshold In diagnostic page.
0x00b2	Critical	Enclosure %s temperature sensor %d below error threshold	Logged when the enclosure services process has detected a temperature lower than a safe operating temperature or lower than the value indicated by the LOW CRITICAL THRESHOLD field in the Threshold In diagnostic page.
0x00b3	Warning	Enclosure %s temperature sensor %d above warning threshold	Logged when the enclosure services process has detected a temperature higher than a normal operating temperature or higher than the value indicated by the HIGH WARNING THRESHOLD field in the Threshold In diagnostic page.
0x00b4	Critical	Enclosure %s temperature sensor %d above error threshold	Logged when the enclosure services process has detected a temperature higher than a safe operating temperature or higher than the value indicated by the HIGH CRITICAL THRESHOLD field in the Threshold In diagnostic page.
0x00b5	Critical	Enclosure %s shutdown	Logged when an unrecoverable condition is detected in the enclosure.
0x00b6	Warning	Enclosure %s not supported; too many enclosures connected to port	Logged when the maximum allowed enclosures per port is exceeded.
0x00b7	Critical	Enclosure %s firmware mismatch	Logged when two ESMs have different firmware versions.

Table 34 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x00b8	Warning	Enclosure %s sensor %d bad	Logged when the device is present on the phy, but the status does not indicate its presence.
0x00b9	Critical	Enclosure %s phy %d bad	Logged when the status indicates a device presence, but there is no corresponding SAS address is associated with the device.
0x00ba	Critical	Enclosure %s is unstable	Logged when the enclosure services process reports the sense errors.
0x00bb	Critical	Enclosure %s hardware error	Logged when a critical or an unrecoverable enclosure failure has been detected by the enclosure services process.
0x00bc	Critical	Enclosure %s not responding	Logged when there is no response from the enclosure.
0x00bd	Information	SAS/SATA mixing not supported in enclosure; Drive %s disabled	Logged when the SAS/SATA mixing in an enclosure is being violated.
0x00be	Information	Enclosure (SES) hotplug on %s was detected, but is not supported	Not reported to the user.
0x00bf	Information	Clustering enabled	Logged when the clustering is enabled in the controller properties.
0x00c0	Information	Clustering disabled	Logged when the clustering is disabled in the controller properties.
0x00c1	Information	Drive too small to be used for auto-rebuild on %s	Logged when the size of the drive is not sufficient for auto-rebuild.
0x00c2	Information	BBU enabled; changing WT virtual drives to WB	Logged when changing WT virtual drives to WB and the BBU status is good.
0x00c3	Warning	BBU disabled; changing WB virtual drives to WT	Logged when changing WB virtual drives to WT and the BBU status is bad.
0x00c4	Warning	Bad block table on drive %s is 80% full	Logged when the Bad block table on a drive is 80% full.
0x00c5	Fatal	Bad block table on drive %s is full; unable to log block %lx	Logged when the Bad block table on a drive is full and not able to add the bad block in the Bad block table.
0x00c6	Information	Consistency Check Aborted due to ownership loss on %s	Logged when the Consistency Check is aborted due to ownership is lost.
0x00c7	Information	Background Initialization (BGI) Aborted Due to Ownership Loss on %s	Logged when the Background Initialization (BGI) is aborted due to ownership loss.
0x00c8	Critical	Battery/charger problems detected; SOH Bad	Logged when the battery is not presented or removed and SOH is bad.
0x00c9	Warning	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); warning threshold exceeded	Logged when the Single-bit ECC errors exceeded the warning threshold.
0x00ca	Critical	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); critical threshold exceeded	Logged when the Single-bit ECC errors exceeded the critical threshold.
0x00cb	Critical	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); further reporting disabled	Logged when the Single-bit ECC errors exceeded all the thresholds and disable further logging.
0x00cc	Critical	Enclosure %s Power supply %d switched off	Logged when the enclosure services process has detected that the Enclosure Power supply is switched off and it was switched on earlier.

Table 34 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x00cd	Information	Enclosure %s Power supply %d switched on	Logged when the enclosure services process has detected that the Enclosure Power supply is switched on and it was switched off earlier.
0x00ce	Critical	Enclosure %s Power supply %d cable removed	Logged when the enclosure services process has detected that the Enclosure Power supply cable is removed and it was inserted earlier.
0x00cf	Information	Enclosure %s Power supply %d cable inserted	Logged when the enclosure services process has detected that the Enclosure Power supply cable is inserted and it was removed earlier.
0x00d0	Information	Enclosure %s Fan %d returned to normal	Logged when the enclosure services process has detected that the current status of a fan is good and it was failed earlier.
0x00d1	Information	BBU Retention test was initiated on previous boot	Logged when the Battery Retention test was initiated on previous boot.
0x00d2	Information	BBU Retention test passed	Logged when the Battery Retention test passed successfully.
0x00d3	Critical	BBU Retention test failed!	Logged when the Battery Retention test failed.
0x00d4	Information	NVRAM Retention test was initiated on previous boot	Logged when the NVRAM Retention test was initiated on previous boot.
0x00d5	Information	NVRAM Retention test passed	Logged when the NVRAM Retention test passed successfully.
0x00d6	Critical	NVRAM Retention test failed!	Logged when the NVRAM Retention test failed.
0x00d7	Information	%s test completed %d passes successfully	Logged when the controller diagnostics test passes successfully.
0x00d8	Critical	%s test FAILED on %d pass. Fail data: errorOffset=%x goodData=%x badData=%x	Logged when the controller diagnostics test fails.
0x00d9	Information	Self check diagnostics completed	Logged when Self check diagnostics is completed.
0x00da	Information	Foreign Configuration detected	Logged when Foreign Configuration is detected.
0x00db	Information	Foreign Configuration imported	Logged when Foreign Configuration is imported.
0x00dc	Information	Foreign Configuration cleared	Logged when Foreign Configuration is cleared.
0x00dd	Warning	NVRAM is corrupt; reinitializing	Logged when NVRAM is corrupt and re-initialized.
0x00de	Warning	NVRAM mismatch occurred	Logged when NVRAM mismatch occurs.
0x00df	Warning	SAS wide port %d lost link on PHY %d	Logged when SAS wide port lost link on a PHY.
0x00e0	Information	SAS wide port %d restored link on PHY %d	Logged when a SAS wide port restored link on a PHY.
0x00e1	Warning	SAS port %d, PHY %d has exceeded the allowed error rate	Logged when a SAS PHY on port has exceeded the allowed error rate.
0x00e2	Warning	Bad block reassigned on %s at %lx to %lx	Logged when a Bad block is reassigned on a drive from a error sector to a new sector.
0x00e3	Information	Controller Hot Plug detected	Logged when a Controller Hot Plug is detected.
0x00e4	Warning	Enclosure %s temperature sensor %d differential detected	Logged when an Enclosure temperature sensor differential is detected.
0x00e5	Information	Drive test cannot start. No qualifying drives found	Logged when Disk test cannot start. No qualifying disks found.
0x00e6	Information	Time duration provided by host is not sufficient for self check	Logged when Time duration provided by the host is not sufficient for self check.

Table 34 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x00e7	Information	Marked Missing for %s on drive group %d row %d	Logged when a physical drive is Marked Missing on an array at a particular row.
0x00e8	Information	Replaced Missing as %s on drive group %d row %d	Logged when a physical drive is Replaced Missing on an array at a particular row.
0x00e9	Information	Enclosure %s Temperature %d returned to normal	Logged when an Enclosure temperature returns to normal.
0x00ea	Information	Enclosure %s Firmware download in progress	Logged when Enclosure a Firmware download is in progress.
0x00eb	Warning	Enclosure %s Firmware download failed	Logged when Enclosure a Firmware download failed.
0x00ec	Warning	%s is not a certified drive	Logged if the drive is not certified.
0x00ed	Information	Dirty cache data discarded by user	Logged when Dirty cache data is discarded by the user.
0x00ee	Information	Drives missing from configuration at boot	Logged when physical drives are missing from configuration at boot.
0x00ef	Information	Virtual drives (VDs) missing drives and will go offline at boot: %s	Logged when virtual drives missing drives and will go offline at boot.
0x00f0	Information	VDs missing at boot: %s	Logged when virtual drives missing at boot.
0x00f1	Information	Previous configuration completely missing at boot	Logged when Previous configuration completely missing at boot.
0x00f2	Information	Battery charge complete	Logged when Battery charge is completed.
0x00f3	Information	Enclosure %s fan %d speed changed	Logged when an Enclosure fan speed changed.
0x00f4	Information	Dedicated spare %s imported as global due to missing arrays	Logged when a Dedicated spare is imported as global due to missing arrays.
0x00f5	Information	%s rebuild not possible as SAS/SATA is not supported in an array	Logged when a rebuild is not possible as SAS/SATA is not supported in an array.
0x00f6	Information	SEP %s has been rebooted as a part of enclosure firmware download. SEP will be unavailable until this process completes.	Logged when SEP has been rebooted as part of enclosure firmware download. It will be unavailable until reboot completes.
0x00f7	Information	Inserted PD: %s Info: %s	Logged when a physical drive is inserted.
0x00f8	Information	Removed PD: %s Info: %s	Logged when a physical drive is removed.
0x00f9	Information	VD %s is now OPTIMAL	Logged when a logical drive state changes to OPTIMAL.
0x00fa	Warning	VD %s is now PARTIALLY DEGRADED	Logged when a logical drive state changes to a partially degraded state.
0x00fb	Critical	VD %s is now DEGRADED	Logged when a logical drive state changes to degraded state.
0x00fc	Fatal	VD %s is now OFFLINE	Logged when a logical drive state changes to offline state.
0x00fd	Warning	Battery requires reconditioning; please initiate a LEARN cycle	Logged when a Battery requires reconditioning; please initiate a LEARN cycle.
0x00fe	Warning	VD %s disabled because RAID-5 is not supported by this RAID key	Logged when a virtual drive is disabled because RAID-5 is not supported by this RAID key.
0x00ff	Warning	VD %s disabled because RAID-6 is not supported by this controller	Logged when a virtual drive is disabled because RAID-6 is not supported by this controller.
0x0100	Warning	VD %s disabled because SAS drives are not supported by this RAID key	Logged when a virtual drive is disabled because SAS drives are not supported by this RAID key.
0x0101	Warning	PD missing: %s	Logged to provide information about the missing drive during boot.

Table 34 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0102	Warning	Puncturing of LBAs enabled	Currently not logged in the firmware.
0x0103	Warning	Puncturing of LBAs disabled	Currently not logged in the firmware.
0x0104	Critical	Enclosure %s EMM %d not installed	Logged when Enclosure SIM is not installed.
0x0105	Information	Package version %s	Prints the Package version number.
0x0106	Warning	Global affinity Hot Spare %s commissioned in a different enclosure	Logged when a hot spare that is a part of an enclosure is commissioned in a different enclosure.
0x0107	Warning	Foreign configuration table overflow	Logged when the number of GUIDs to import exceeds the total supported by the firmware.
0x0108	Warning	Partial foreign configuration imported, PDs not imported:%s	Logged when all the foreign configuration drives could not be imported.
0x0109	Information	Connector %s is active	Logged during initial boot when a SAS MUX connector is found for the controller.
0x010a	Information	Board Revision %s	Logged during boot.
0x010b	Warning	Command timeout on PD %s, CDB:%s	Logged when command to a PD Timesout.
0x010c	Warning	PD %s reset (Type %02x)	Logged when PD is reset.
0x010d	Warning	VD bad block table on %s is 80% full	Logged when number of Bad Blocks entries is at 80 % of what can be supported in the firmware.
0x010e	Fatal	VD bad block table on %s is full; unable to log block %lx (on %s at %lx)	Logged when number of Bad Blocks exceed what can be supported in the firmware.
0x010f	Fatal	Uncorrectable medium error logged for %s at %lx (on %s at %lx)	Logged when an uncorrectable medium error is detected.
0x0110	Information	VD medium error corrected on %s at %lx	Logged on the corrected medium error.
0x0111	Warning	Bad block table on PD %s is 100% full	Logged when Bad block table is 100 % Full. Any more media errors on this physical drive will not be logged in the bad block table.
0x0112	Warning	VD bad block table on PD %s is 100% full	Logged when Bad block table is 100 % Full. Any more media errors on this logical drive will not be logged in the bad block table.
0x0113	Fatal	Controller needs replacement, IOP is faulty	Currently not logged in the firmware.
0x0114	Information	Replace Drive started on PD %s from PD %s	Logged when Replace is started.
0x0115	Information	Replace Drive aborted on PD %s and src is PD %s	Logged when Replace is aborted.
0x0116	Information	Replace Drive complete on PD %s from PD %s	Logged when Replace is completed.
0x0117	Progress	Replace Drive progress on PD %s is %s	Logged to provide the progress of Replace.
0x0118	Information	Replace Drive resumed on PD %s from %s	Logged when Replace operation is resumed.
0x0119	Information	Replace Drive automatically started on PD %s from %s	Logged on automatic start of Replace.

Table 34 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x011a	Critical	Replace Drive failed on PD %s due to source %s error	Logged when the source physical drive of a Replace fails. The Replace stops and rebuild starts on the destination physical drive.
0x011b	Warning	Early Power off warning was unsuccessful	Currently not logged in the firmware.
0x011c	Information	BBU FRU is %s	Logged only for IBM.
0x011d	Information	%s FRU is %s	Logged if FRU data is present. Logged only for IBM.
0x011e	Information	Controller hardware revision ID %s	Currently not used in the firmware.
0x011f	Warning	Foreign import shall result in a backward incompatible upgrade of configuration metadata	Currently not used in the firmware.
0x0120	Information	Redundant path restored for PD %s	Logged when new path is added for the physical drives.
0x0121	Warning	Redundant path broken for PD %s	Logged when one path is removed.
0x0122	Information	Redundant enclosure EMM %s inserted for EMM %s	Logged when an enclosure is added.
0x0123	Information	Redundant enclosure EMM %s removed for EMM %s	Logged when an enclosure is removed
0x0124	Warning	Patrol Read can't be started, as PDs are either not ONLINE, or are in a VD with an active process, or are in an excluded VD	Logged when none of the disks can start PR.
0x0125	Information	Replace Drive aborted by user on PD %s and src is PD %s	Logged when Replace is aborted by the user.
0x0126	Critical	Replace Drive aborted on hot spare %s from %s, as hot spare needed for rebuild	Logged when Replace is aborted on a Hotspare.
0x0127	Warning	Replace Drive aborted on PD %s from PD %s, as rebuild required in the array	Logged when Replace is stopped for a higher priority rebuild operation on a drive.
0x0128	Fatal	Controller cache discarded for missing or offline VD %s When a VD with cached data goes offline or missing during runtime, the cache for the VD is discarded. Because the VD is offline, the cache cannot be saved.	Logged when pinned cache lines are discarded for a LD.
0x0129	Information	Replace Drive cannot be started as PD %s is too small for src PD %s	Logged when destination PD is too small for Replace.
0x012a	Information	Replace Drive cannot be started on PD %s from PD %s, as SAS/SATA is not supported in an array	Logged when there is a SAS/SATA mixing violation for the destination PD.
0x012b	Information	Microcode update started on PD %s	Logged when PD Firmware download starts.
0x012c	Information	Microcode update completed on PD %s	Logged when PD Firmware download completes.
0x012d	Warning	Microcode update timeout on PD %s	Logged when PD Firmware download does not complete and times out.
0x012e	Warning	Microcode update failed on PD %s	Logged when PD Firmware download fails.
0x012f	Information	Controller properties changed	Logged when any of the controller properties has changed.

Table 34 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0130	Information	Patrol Read properties changed	Currently not logged in the firmware.
0x0131	Information	CC Schedule properties changed	Logged when consistency check scheduling property has changed.
0x0132	Information	Battery properties changed	Logged when any of the BBU properties has changed.
0x0133	Warning	Periodic Battery Relearn is pending. Please initiate manual learn cycle as Automatic learn is not enabled	Logged when BBU periodic relearn is pending.
0x0134	Information	Drive security key created	Logged when controller lock key is created.
0x0135	Information	Drive security key backed up	Logged when controller lock key is backed up.
0x0136	Information	Drive security key from escrow, verified	Logged when controller lock key is verified from escrow.
0x0137	Information	Drive security key changed	Logged when controller lock key is re-keyed.
0x0138	Warning	Drive security key, re-key operation failed	Logged when controller lock re-key operation failed.
0x0139	Warning	Drive security key is invalid	Logged when the controller lock is not valid.
0x013a	Information	Drive security key destroyed	Logged when the controller lock key is destroyed.
0x013b	Warning	Drive security key from escrow is invalid	Logged when the controller escrow key is not valid. This escrow key can not unlock any drive.
0x013c	Information	VD %s is now secured	Logged when secure LD is created.
0x013d	Warning	VD %s is partially secured	Logged when all the drives in the array are not secure.
0x013e	Information	PD %s security activated	Logged when PD security key is set.
0x013f	Information	PD %s security disabled	Logged when security key is removed from an FDE drive.
0x0140	Information	PD %s is reprovisioned	Logged when PD security is cleared.
0x0141	Information	PD %s security key changed	Logged when PD lock key is re-keyed.
0x0142	Fatal	Security subsystem problems detected for PD %s	Logged when PD security can not be set.
0x0143	Fatal	Controller cache pinned for missing or offline VD %s	Logged when LD cache is pinned.
0x0144	Fatal	Controller cache pinned for missing or offline VDs: %s	Logged when pinned cache is found during OCR.
0x0145	Information	Controller cache discarded by user for VDs: %s	Logged when LD pinned cache is discarded by the user.
0x0146	Information	Controller cache destaged for VD %s	Logged when LD pinned cache is recovered.
0x0147	Warning	Consistency Check started on an inconsistent VD %s	Logged when consistency check is started on an inconsistent LD.
0x0148	Warning	Drive security key failure, cannot access secured configuration	Logged when an invalid lock key is detected.
0x0149	Warning	Drive security password from user is invalid	Not logged.
0x014a	Warning	Detected error with the remote battery connector cable	Not logged.
0x014b	Information	Power state change on PD %s from %s to %s	Logged when PD power state (spun up, spun down, in-transition) changes.
0x014c	Information	Enclosure %s element (SES code 0x%x) status changed	Not logged.

Table 34 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x014d	Information	PD %s rebuild not possible as HDD/CacheCade software mix is not supported in a drive group	Logged when mixing violation occurs due to HDD/SSD mismatch.
0x014e	Information	Replace Drive cannot be started on PD %s from %s, as HDD/CacheCade software mix is not supported in a drive group	Logged when Replace could not be started on a PD because HDD/CacheCade software mix was not supported in a drive group.
0x014f	Information	VD bad block table on %s is cleared	Logged when a VD bad block table was cleared.
0x0150	Caution	SAS topology error: 0x%lx	Logged when a SAS topology error occurred.
0x0151	Information	VD cluster of medium errors corrected for %s at %lx (on %s at %lx)	Logged when medium errors were corrected for a PD for a LD.
0x0152	Information	Controller requests a host bus rescan	Logged when controller requested a host bus rescan.
0x0153	Information	Controller repurposed and factory defaults restored	Logged when controller repurposed and factory defaults were restored.
0x0154	Information	Drive security key binding updated	Logged when drive security key binding was updated.
0x0159	Critical	Controller encountered a fatal error and was reset	Logged when a controller encountered a fatal error and was reset.
0x015a	Information	Snapshots enabled on %s (Repository %s)	Logged when snapshot was enabled on a LD.
0x015b	Information	Snapshots disabled on %s (Repository %s) by the user	Logged when snapshot was disabled on a LD by the user.
0x015c	Critical	Snapshots disabled on %s (Repository %s), due to a fatal error	Logged when snapshot was disabled on a LD due to a fatal error.
0x015d	Information	Snapshot created on %s at %s	Logged when snapshot was created on a LD.
0x015e	Information	Snapshot deleted on %s at %s	Logged when snapshot was deleted on a LD.
0x015f	Information	View created at %s to a snapshot at %s for %s	Logged when view was created at a LD.
0x0160	Information	View at %s is deleted, to snapshot at %s for %s	Logged when View at a LD was deleted
0x0161	Information	Snapshot rollback started on %s from snapshot at %s	Logged when snapshot rollback was started on a LD.
0x0162	Fatal	Snapshot rollback on %s internally aborted for snapshot at %s	Logged when snapshot rollback was internally aborted.
0x0163	Information	Snapshot rollback on %s completed for snapshot at %s	Logged when snapshot rollback on a LD was completed.
0x0164	Information	Snapshot rollback progress for snapshot at %s, on %s is %s	Logged to report snapshot rollback progress on a LD.
0x0165	Warning	Snapshot space for %s in snapshot repository %s, is 80%% full	Logged when snapshot space for a LD in a snapshot repository was 80% full.
0x0166	Critical	Snapshot space for %s in snapshot repository %s, is full	Logged when snapshot space for a LD in a snapshot repository was full.
0x0167	Warning	View at %s to snapshot at %s, is 80%% full on snapshot repository %s	Logged when view at a LD to a snapshot was 80% full on a snapshot repository.
0x0168	Critical	View at %s to snapshot at %s, is full on snapshot repository %s	Logged when view at a LD to a snapshot was full on a snapshot repository.
0x0169	Critical	Snapshot repository lost for %s	Logged when snapshot repository was lost for a LD.
0x016a	Warning	Snapshot repository restored for %s	Logged when snapshot repository was restored for a LD.

Table 34 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x016b	Critical	Snapshot encountered an unexpected internal error: 0x%lx	Logged when snapshot encountered an unexpected internal error.
0x016c	Information	Auto Snapshot enabled on %s (snapshot repository %s)	Logged when auto snapshot was enabled.
0x016d	Information	Auto Snapshot disabled on %s (snapshot repository %s)	Logged when auto Snapshot was disabled.
0x016e	Critical	Configuration command could not be committed to disk, please retry	Logged when configuration command could not be committed to disk and was asked to retry.
0x016f	Information	COD on %s updated as it was stale	Logged when COD in DDF is updated due to various reasons.
0x0170	Warning	Power state change failed on %s (from %s to %s)	Logged when power state change failed on a PD.
0x0171	Warning	%s is not available	Logged when a LD was not available.
0x0172	Information	%s is available	Logged when a LD was available.
0x0173	Information	%s is used for CacheCade with capacity 0x%lx logical blocks	Logged when a LD was used for CacheCade with the indicated capacity in logical blocks.
0x0174	Information	%s is using CacheCade %s	Logged when a LD was using CacheCade.
0x0175	Information	%s is no longer using CacheCade %s	Logged when a LD was no longer using CacheCade.
0x0176	Critical	Snapshot deleted due to resource constraints for %s in snapshot repository %s	Logged when the snapshot is deleted due to resource constraints in snapshot repository.
0x0177	Warning	Auto Snapshot failed for %s in snapshot repository %s	Logged when the Auto Snapshot is failed for a VD in snapshot repository.
0x0178	Warning	Controller reset on-board expander	Logged when the chip reset issued to on-board expander.
0x0179	Warning	CacheCade (%s) capacity changed and is now 0x%lx logical blocks	Logged when the CacheCade capacity is changed along with the current capacity.
0x017a	Warning	Battery cannot initiate transparent learn cycles	Logged when the Battery cannot initiate transparent learn cycles.
0x017b	Information	Premium feature %s key was applied for - %s	Logged when the Premium feature key was applied.
0x017c	Information	Snapshot schedule properties changed on %s	Logged when the Snapshot schedule properties changed.
0x017d	Information	Snapshot scheduled action is due on %s	Logged when the Snapshot scheduled action is due.
0x017e	Information	Performance Metrics: collection command 0x%lx	Logged during the Performance Metrics collection.
0x017f	Information	Premium feature %s key was transferred - %s	Logged when the Premium feature key was transferred.
0x0180	Information	Premium feature serial number %s	Logged when displaying the Premium feature serial number.
0x0181	Warning	Premium feature serial number mismatched. Key-vault serial num - %s	Logged when Premium feature serial number mismatched.
0x0182	Warning	Battery cannot support data retention for more than %d hours. Please replace the battery	Logged during the Battery monitoring and it displays the remaining data retention time of the battery.
0x0183	Information	%s power policy changed to %s (from %s)	Logged when the power policy of an LD is changed.

Table 34 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x0184	Warning	%s cannot transition to max power savings	Logged when LD cannot transition to max power savings.
0x0185	Information	Host driver is loaded and operational	This event is not reported to the user.
0x0186	Information	%s mirror broken	Logged when the mirror is broken for an LD.
0x0187	Information	%s mirror joined	Logged when joining the LD with its broken mirror.
0x0188	Warning	%s link %d failure in wide port	This event is not reported to the user.
0x0189	Information	%s link %d restored in wide port	This event is not reported to the user.
0x018a	Information	Memory module FRU is %s	This event is not reported to the user.
0x018b	Warning	Cache-vault power pack is sub-optimal. Please replace the pack	This event is not reported to the user.
0x018c	Warning	Foreign configuration auto-import did not import any drives	Logged when the Foreign configuration auto-import did not import any drives.
0x018d	Warning	Cache-vault microcode update required	Logged when the BMU is not in Normal mode and Cache-vault microcode update required.
0x018e	Warning	CacheCade (%s) capacity exceeds maximum allowed size, extra capacity is not used	Logged when CacheCade capacity exceeds maximum allowed size, extra capacity is not used.
0x018f	Warning	LD (%s) protection information lost	Logged when the protection information is lost for an LD.
0x0190	Information	Diagnostics passed for %s	Logged when the SHIELD Diagnostics passed for a PD.
0x0191	Critical	Diagnostics failed for %s	Logged when the SHIELD Diagnostics failed for a PD.
0x0192	Information	Server Power capability Diagnostic Test Started	Logged when the Server Power capability Diagnostic Test starts.
0x0193	Information	Drive Cache settings enabled during rebuild for %s	Logged when the Drive Cache settings enabled during rebuild for a PD.
0x0194	Information	Drive Cache settings restored after rebuild for %s	Logged when the Drive Cache settings restored after rebuild for a PD.
0x0195	Information	Drive %s commissioned as Emergency spare	Logged when the Drive commissioned as Emergency spare.
0x0196	Warning	Reminder: Potential non-optimal configuration due to drive %s commissioned as emergency spare	Logged when the PD being imported is an Emergency Spare.
0x0197	Information	Consistency Check suspended on %s	Logged when the Consistency Check is suspended on an LD.
0x0198	Information	Consistency Check resumed on %s	Logged when the Consistency Check is resumed on an LD.
0x0199	Information	Background Initialization suspended on %s	Logged when the Background Initialization is suspended on an LD.
0x019a	Information	Background Initialization resumed on %	Logged when the Background Initialization is resumed on an LD.
0x019b	Information	Reconstruction suspended on %s	Logged when the Reconstruction is suspended on an LD.
0x019c	Information	Rebuild suspended on %	Logged when the Rebuild is suspended on a PD.
0x019d	Information	Replace Drive suspended on %s	Logged when the Replace is suspended on a PD.
0x019e	Information	Reminder: Consistency Check suspended on %	Logged as a reminder when the Consistency Check is suspended on an LD.
0x019f	Information	Reminder: Background Initialization suspended on %s	Logged as a reminder when the Background Initialization is suspended on an LD.
0x01a0	Information	Reminder: Reconstruction suspended on %s	Logged as a reminder when the Reconstruction is suspended on an LD.

Table 34 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x01a1	Information	Reminder: Rebuild suspended on %s	Logged as a reminder when the Rebuild is suspended on a PD.
0x01a2	Information	Reminder: Replace Drive suspended on %s	Logged as a reminder when Replace is suspended on a PD.
0x01a3	Information	Reminder: Patrol Read suspended	Logged as a reminder when the Patrol Read is suspended.
0x01a4	Information	Erase aborted on %s	Logged when the Erase is aborted on a PD.
0x01a5	Critical	Erase failed on %s (Error %02x)	Logged when the Erase is failed on a PD along with the error.
0x01a6	Progress	Erase progress on %s is %s	Logged to display the Erase progress on a PD along with its current progress.
0x01a7	Information	Erase started on %s	Logged when Erase is started on a PD.
0x01a8	Information	Erase completed on %s	Logged when the Erase is completed on a PD.
0x01a9	Information	Erase aborted on %s	Logged when the Erase is aborted on an LD.
0x01aa	Critical	Erase failed on %s	Logged when the Erase is failed on an LD.
0x01ab	Progress	Erase progress on %s is %s	Logged to display the Erase progress on an LD along with its current progress.
0x01ac	Information	Erase started on %s	Logged when the Erase is started on an LD.
0x01ad	Information	Erase complete on %s	Logged when the Erase is complete on an LD.
0x01ae	Warning	Potential leakage during erase on %s	Logged to inform the Potential leakage during erase on an LD.
0x01af	Warning	Battery charging was suspended due to high battery temperature	Logged when the Battery charging was suspended due to high battery temperature.
0x01b0	Information	NVCache firmware update was successful	This event is not reported to the user.
0x01b1	Warning	NVCache firmware update failed	This event is not reported to the user.
0x01b2	Fatal	%s access blocked as cached data in CacheCade is unavailable	This event is not reported to the user.
0x01b3	Information	CacheCade disassociate started on %s	This event is not reported to the user.
0x01b4	Information	CacheCade disassociate completed on %s	This event is not reported to the user.
0x01b5	Critical	CacheCade disassociate failed on %s	This event is not reported to the user.
0x01b6	Progress	CacheCade disassociate progress on %s is %s	This event is not reported to the user.
0x01b7	Information	CacheCade disassociate aborted by user on %s	This event is not reported to the user.
0x01b8	Information	Link speed changed on SAS port %d and PHY %d	Logged when the Link speed changed on SAS port and PHY.
0x01b9	Warning	Advanced Software Options was deactivated for - %s	This event is not reported to the user.
0x01ba	Information	%s is now accessible	This event is not reported to the user.
0x01bb	Information	%s is using CacheCade	This event is not reported to the user.
0x01bc	Information	%s is no longer using CacheCade	This event is not reported to the user.
0x01bd	Warning	Patrol Read aborted on %s	Logged when the Patrol Read is aborted on a PD.

Table 34 Event Messages (Continued)

Number	Severity Level	Event Text	Generic Conditions when Each Event Occurs
0x01c2	Information	Periodic Battery Relearn was missed, and rescheduled to %s	Logged if Battery Relearn was missed at the scheduled time due to a system power off then the controller will reschedule automatically when you power on the system.
0x01c3	Information	Controller reset requested by host	Logged when the Controller Reset process started on the corresponding controller.
0x01c4	Information	Controller reset requested by host, completed	Logged when the Controller Reset process completed on the corresponding controller.
0x01c7	Warning	Controller booted in headless mode with errors	Logged when the Controller is booted to safe mode due to warning errors.
0x01c8	Critical	Controller booted to safe mode due to critical errors	Logged when the Controller is booted to safe mode due to critical errors.
0x01c9	Warning	Warning Error during boot - %s	Logged when a warning error occurs during booting the controller to safe mode.
0x01ca	Critical	Critical Error during boot - %s	Logged when a critical error occurs during booting the controller to safe mode
0x01cb	Fatal	Fatal Error during boot - %s	Logged when a fatal error occurs during booting the controller to safe mode

Appendix C: HTTP Status codes and description

HTTP status codes notify you about the status of the request made.This section describes the meaning of the HTTP status codes.

Table 35 HTTP Status Codes

Code	Description	Example

200
OK
	The request was successfully completed and includes a representation in its body (if applicable). In most cases, this is the code the client hopes to see. It indicates that the server successfully carried out whatever action the client requested. Also use for partial success e.g. few of the controller fields are updated (by checking read-modify-write-read) process.	`GET /servers/{id}/controllers`
		`PUT/servers/{id}/controllers/1`
		`PUT/servers/{id}/controllers/1/virtualdrives/0`

201 Created	A request that created a new resource completed successfully.	`POST/servers/{id}/controllers/1/virtualdrives`
301 Moved Permanently	The requested resource resides under a different URI	TBD
302 Found	The requested resource resides temporarily under a different URI.	TBD

400
Bad Request
	The request could not be processed because it contains missing or invalid information (such as validation error on an input field, a missing required value, and so on). In general either the request body is not valid against schema or semantical error in request	`PUT/servers/{id}/controllers/1`
		`with {"nonExistantAttribute":"0"}`
		Try to run consistency-check on R0 virtual drive. The status code will be MFI_STAT if coming from FW.

401 Unauthorized	?The authentication credentials are missing or invalid.	`GET/servers/{id}/controllers?` Before authentication
403 Forbidden	The server recognized the credentials in the request, but those credentials do not possess authorization to perform this request. E.g. a read-only user is trying to create a configuration.	`POST /servers/{id}/controllers/1/virtualdrives` For a read-only user

404
Not Found
	The request specified a URI of a resource that does not exist. The server has not found anything matching the Request-URI. No indication is given of whether the condition is temporary or permanent.	`GET /servers/{id}/controlers/0`
		`GET /servers/{id}/controllers/99`
		WHEN there is not controller with id 99

405 Method Not Allowed	The URI is valid but the HTTP verb specified in the request (e.g. DELETE, GET, POST, PUT) is not supported for this request URI.	`POST /servers/{id}/controllers`
410 Gone	The requested resource is no longer available at the server and no forwarding address is known. This condition is expected to be considered permanent.	Client wants to download a log (e.g. ttylog) but the file is being deleted and no longer available (and will never be available).

Table 35 HTTP Status Codes (Continued)

Code	Description	Example
422 Unprocessable Entity	Semantical error (see status code 400 for details)	Try to run consistency-check on R0 virtual drive.
500 Internal Server Error	The server encountered an unexpected condition that prevented it from fulfilling the request.	Most of the cases it is a defect/error in backend itself e.g. memory allocation error or initialization error etc.
501 Not Implemented	Implies future availability. The current server lacks the ability to fulfill the request.	None so far.

Appendix D: Glossary

This glossary defines the terms used in this document.

A

Access policy	A virtual drive property indicating what kind of access is allowed for a particular virtual drive. The possible values are *Read/Write*, *Read Only*, or *Blocked*.

B

BIOS	Basic Input/Output System. The computer BIOS is stored on a flash memory chip. The BIOS controls communications between the microprocessor and peripheral devices, such as the keyboard and the video controller, and miscellaneous functions, such as system messages.

C

Cache	Fast memory that holds recently accessed data. Use of cache memory speeds subsequent access to the same data. When data is read from or written to main memory, a copy is also saved in cache memory with the associated main memory address. The cache memory software monitors the addresses of subsequent reads to see if the required data is already stored in cache memory. If it is already in cache memory (a cache hit), it is read from cache memory immediately and the main memory read is aborted (or not started). If the data is not cached (a cache miss), it is fetched from main memory and saved in cache memory.
Caching	The process of using a high speed memory buffer to speed up a computer system's overall read/write performance. The cache can be accessed at a higher speed than a drive subsystem. To improve read performance, the cache usually contains the most recently accessed data, as well as data from adjacent drive sectors. To improve write performance, the cache can temporarily store data in accordance with its write back policies.
Capacity	A property that indicates the amount of storage space on a drive or virtual drive.
Coerced capacity	A drive property indicating the capacity to which a drive has been coerced (forced) to make it compatible with other drives that are nominally the same capacity. For example, a 4-GB drive from one manufacturer might be 4,196 MB, and a 4-GB from another manufacturer might be 4,128 MB. These drives could be coerced to a usable capacity of 4,088 MB each for use in a drive group in a storage configuration.
Coercion mode	A controller property indicating the capacity to which drives of nominally identical capacity are coerced (forced) to make them usable in a storage configuration.
Consistency check	An operation that verifies that all stripes in a virtual drive with a redundant RAID level are consistent and that automatically fixes any errors. For RAID 1 drive groups, this operation verifies correct mirrored data for each stripe.
Consistency check rate	The rate at which consistency check operations are run on a computer system.
Controller	A chip that controls the transfer of data between the microprocessor and memory or between the microprocessor and a peripheral device such as a drive. RAID controllers perform RAID functions such as striping and mirroring to provide data protection.
Copyback	The procedure used to copy data from a source drive of a virtual drive to a destination drive that is not a part of the virtual drive. The copyback operation is often used to create or restore a specific physical configuration for a drive group (for example, a specific arrangement of drive group members on the device I/O buses). The copyback operation can be run automatically or manually. Typically, a drive fails or is expected to fail, and the data is rebuilt on a hot spare. The failed drive is replaced with a new drive. Then the data is copied from the hot spare to the new drive, and the hot spare reverts from a rebuild drive to its original hot spare status. The copyback operation runs as a background activity, and the virtual drive is still available online to the host.
Current	Measure of the current flowing to (+) or from (-) the battery, reported in milliamperes.

Current write policy

A virtual drive property that indicates whether the virtual drive currently supports Write Back mode or Write Through mode.

- In Write Back mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction.
- In Write Through mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction.

D

Device ID

A controller or drive property indicating the manufacturer-assigned device ID.

Drive group

A group of drives attached to a RAID controller on which one or more virtual drives can be created. All virtual drives in the drive group use all of the drives in the drive group.

Drive state

A physical drive or a virtual drive property indicating the status of the appropriate drive.

Physical Drive State

A physical drive can be in any one of the following states:

- **Unconfigured Good** – A drive accessible to the RAID controller but not configured as a part of a virtual drive or as a hot spare.
- **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.
- **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.
- **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.
- **None** – A drive with an unsupported flag set. An Unconfigured Good or Offline drive that has completed the prepare for removal operation.

Virtual Drive State

A virtual drive can be in any one of the following states:

- **Optimal** – A virtual drive whose members are all online.
- **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.
- **Degraded** – A virtual drive with a redundant RAID level with one or more member failures and can no longer sustain a subsequent drive failure.
- **Offline** - A virtual drive with one or more member failures that make the data inaccessible.

Drive type

A drive property indicating the characteristics of the drive.

E

Energy Pack

Refers to a battery backup unit or a CacheVault.

F

Fast initialization

A mode of initialization that quickly writes zeroes to the first and last sectors of the virtual drive. This allows you to immediately start writing data to the virtual drive while the initialization is running in the background.

Fault tolerance

The capability of the drive subsystem to undergo a single drive failure per drive group without compromising data integrity and processing capability. Lenovo SAS RAID controllers provides fault tolerance through redundant drive groups in RAID levels 1, 5, 6, 10, 50, and 60. They also support hot spare drives and the auto-rebuild feature.

Firmware	Software stored in read-only memory (ROM) or programmable ROM (PROM). Firmware is often responsible for the behavior of a system when it is first turned on. A typical example would be a monitor program in a system that loads the full operating system from drive or from a network and then passes control to the operating system.
Foreign configuration	A RAID configuration that already exists on a replacement set of drives that you install in a computer system. MegaRAID Storage Manager 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.
G	
GUI	Graphical User Interface.
GT/s	Giga Transfers per second.
H	
Hot spare	A standby drive that can automatically replace a failed drive in a virtual drive and prevent data from being lost. A hot spare can be dedicated to a single redundant drive group or it can be part of the global hot spare pool for all drive groups controlled by the controller. When a drive fails, MegaRAID Storage Manager software automatically uses a hot spare to replace it and then rebuilds the data from the failed drive to the hot spare. Hot spares can be used in RAID 1, 5, 6, 10, 50, and 60 storage configurations.
I	
Initialization	The process of writing zeros to the data fields of a virtual drive and, in fault-tolerant RAID levels, generating the corresponding parity to put the virtual drive in a Ready state. Initialization erases all previous data on the drives. Drive groups will work without initializing, but they can fail a consistency check because the parity fields have not been generated.
IO policy	A virtual drive property indicating whether Cached I/O or Direct I/O is being used. In Cached I/O mode, all reads are buffered in cache memory. In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to cache and the host concurrently. If the same data block is read again, it comes from cache memory. (The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.)
L	
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
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.
M	
Manufacturing date	Date on which the battery pack assembly was manufactured.
Manufacturing name	Device code that indicates the manufacturer of the components used to make the battery assembly.
Migration	The process of moving virtual drives and hot spare drives from one controller to another by disconnecting the drives from one controller and attaching them to another one. The firmware on the new controller will detect and retain the virtual drive information on the drives.
Mirroring	The process of providing complete data redundancy with two drives by maintaining an exact copy of one drive's data on the second drive. If one drive fails, the contents of the other drive can be used to maintain the integrity of the system and to rebuild the failed drive.

Multipathing	The firmware provides support for detecting and using multiple paths from the RAID controllers to the SAS devices that are in enclosures. Devices connected to enclosures have multiple paths to them. With redundant paths to the same port of a device, if one path fails, another path can be used to communicate between the controller and the device. Using multiple paths with load balancing, instead of a single path, can increase reliability through redundancy.
N	
NVMe	Non-volatile memory. NVMe is a logical device interface specification for accessing non-volatile storage media attached via PCI Express (PCIe) bus.
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.
R	
RAID	A group of multiple, independent drives that provide high performance by increasing the number of drives used for saving and accessing data. A RAID drive group improves input/output (I/O) performance and data availability. The group of drives appears to the host system as a single storage unit or as multiple virtual drives. Data throughput improves because several drives can be accessed simultaneously. RAID configurations also improve data storage availability and fault tolerance. Redundant RAID levels (RAID levels 1, 5, 6, 10, 50, and 60) provide data protection.
RAID 0	Uses data striping on two or more drives to provide high data throughput, especially for large files in an environment that requires no data redundancy.
RAID 00	Uses data striping on two or more drives in a spanned drive group to provide high data throughput, especially for large files in an environment that requires no data redundancy.
RAID 1	Uses data mirroring on pairs of drives so that data written to one drive is simultaneously written to the other drive. RAID 1 works well for small databases or other small applications that require complete data redundancy.
RAID 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. Lenovo 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. In Always Read Ahead mode, the controller reads sequentially ahead of requested data and stores the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data. In No Read Ahead mode (known as Normal mode in WebBIOS), read ahead capability is disabled.
Rebuild	The regeneration of all data to a replacement drive in a redundant virtual drive after a drive failure. A drive rebuild normally occurs without interrupting normal operations on the affected virtual drive, though some degradation of performance of the drive subsystem can occur.
Rebuild rate	The percentage of central processing unit (CPU) resources devoted to rebuilding data onto a new drive after a drive in a storage configuration has failed.
Reclaim virtual drive	A method of undoing the configuration of a new virtual drive. If you highlight the virtual drive in the Configuration Wizard and click Reclaim, the individual drives are removed from the virtual drive configuration.
Reconstruction rate	The user-defined rate at which a drive group modification operation is carried out.
Redundancy	A property of a storage configuration that prevents data from being lost when one drive fails in the configuration.
Redundant configuration	A virtual drive that has redundant data on drives in the drive group that can be used to rebuild a failed drive. The redundant data can be parity data striped across multiple drives in a drive group, or it can be a complete mirrored copy of the data stored on a second drive. A redundant configuration protects the data in case a drive fails in the configuration.
S	
SAS	Acronym for Serial-Attached SCSI. SAS is a serial, point-to-point, enterprise-level device interface that leverages the Small Computer System Interface (SCSI) protocol set. The SAS interface provides improved performance, simplified cabling, smaller connectors, lower pin count, and lower power requirements when compared to parallel SCSI.
SATA	Acronym for Serial Advanced Technology Attachment. A physical storage interface standard. SATA is a serial link that provides point-to-point connections between devices. The thinner serial cables allow for better airflow within the system and permit smaller chassis designs.
SCSI device type	A drive property indicating the type of the device, such as drive.
Serial no.	A controller property indicating the manufacturer-assigned serial number.
Stripe size	A virtual drive property indicating the length of the interleaved data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. The user can select the stripe size.
Striping	A technique used to write data across all drives in a virtual drive. Each stripe consists of consecutive virtual drive data addresses that are mapped in fixed-size units to each drive in the virtual drive using a sequential pattern. For example, if the virtual drive includes five drives, the stripe writes data to drives one through five without repeating any of the drives. The amount of space consumed by a stripe is the same on each drive. Striping by itself does not provide data redundancy. Striping in combination with parity does provide data redundancy.
Strip size	The portion of a stripe that resides on a single drive in the drive group.
Subvendor ID	A controller property that lists additional vendor ID information about the controller.

T

Temperature	Temperature of the battery pack, measured in Celsius.

V

Vendor ID	A controller property indicating the vendor-assigned ID number of the controller.
Vendor info	A drive property listing the name of the vendor of the drive.
Virtual drive	A storage unit created by a RAID controller from one or more drives. Although a virtual drive can be created from several drives, it is seen by the operating system as a single drive. Depending on the RAID level used, the virtual drive can retain redundant data in case of a drive failure.
Virtual drive state	A virtual drive property indicating the condition of the virtual drive. Examples include Optimal and Degraded.

W

Write-back	In Write-Back Caching mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a drive write transaction. Data is written to the drive subsystem in accordance with policies set up by the controller. These policies include the amount of dirty/clean cache lines, the number of cache lines available, and elapsed time from the last cache flush.
Write policy	See *Default Write Policy*.
Write-through	In Write-Through Caching mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data and has completed the write transaction to the drive.