



LSA LSI[®] Storage Authority Software

**User Guide
Version 2.2**

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

Copyright © 2014–2021 Broadcom. All Rights Reserved.

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

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

Table of Contents

How This Guide Is Organized.....	9
LSI Storage Authority Overview.....	10
Support Matrix.....	10
Technical Support.....	11
LSI Storage Authority Features.....	12
LSI Storage Authority Preinstallation, Postinstallation, and Upgrade Requirements.....	15
Preinstallation Requirements.....	15
OpenSLP.....	15
Upgrade Requirements.....	15
Browser Cache.....	15
Types of Installation.....	16
Gateway Installer.....	16
StandAlone Installer.....	17
DirectAgent Installer.....	17
Light Weight Monitor.....	17
Light Weight Agent Installer.....	17
Configuring HTTPS.....	18
Resetting Encryption Keys.....	19
Installing the LSI Storage Authority Software on the Microsoft Windows Operating System.....	20
Installing in Noninteractive Mode.....	24
Uninstalling in Interactive Mode.....	25
Uninstalling in Noninteractive Mode.....	25
Installing the LSI Storage Authority Software on the Linux Operating System.....	26
Installing in the Interactive Mode.....	26
Installing in the Noninteractive Mode.....	27
Uninstalling the LSI Storage Authority Software on the Linux Operating System.....	27
Adding a Digital Signature and Enabling ASLR and DEP for LSI Utilities.....	28
Adding a Digital Signature in a Windows Environment.....	28
Adding a Digital Signature in a Linux Environment.....	28
Managing Light Weight Agent.....	29
Installing LWA on VMware ESXi.....	29
Uninstalling LWA.....	29
Managing LSA and Web Services.....	29
Configuring the Firewall on Various LSA Installers.....	30

Configuring the Firewall on Gateway/StandAlone Installer (Windows).....	30
Configuring the Firewall on Gateway/StandAlone Installer (Linux).....	30
Configuring the Firewall on DirectAgent Installer.....	31
Collecting LSA Logs on Windows and Linux Operating Systems.....	32
Behavior of Event History.....	32
Behavior of Event Monitoring.....	33
Limitations of Installation and Configuration.....	34
Performing Initial Configuration.....	35
Domain Authentication Behavior.....	35
Adding Translated Names in Windows.....	35
Adding Translated Names in Linux.....	35
Using LDAP Authentication.....	36
Accessing LSA over Network Address Translation.....	37
Changing the LSI Storage Authority Application Port Number.....	37
Hiding an Empty Backplane.....	38
Changing the nginx Web Server Port Numbers.....	38
Changing the nginx Read Timeout.....	38
Performing Initial Setup.....	40
Managing Servers from the Remote Server Discovery Page.....	40
Adding Managed Servers.....	41
Removing Managed Servers.....	42
Manually Discovering Servers.....	42
Alert Settings.....	44
Setting Up the Email Server.....	45
Adding Email Addresses of Recipients of Alert Notifications.....	46
Server Dashboard.....	49
Controller Dashboard.....	52
Configuration.....	54
Creating a New Storage Configuration Using the Simple Configuration Option.....	54
Creating a New Storage Configuration Using the Advanced Configuration Option.....	55
Adding Physical Drives.....	59
Adding Hot Spares to the Existing Drive Group.....	59
Adding Virtual Drives.....	60
Clearing the Configuration.....	61
Importing or Clearing the Foreign Configurations.....	62
Personality Management.....	62
Changing Personality Modes.....	63
Changing Behavior Modes.....	63

Background Operations Support.....	65
Managing Controllers.....	66
Viewing Controller Properties.....	66
Running the Consistency Check.....	68
Setting the Consistency Check Properties.....	68
Scheduling a Consistency Check.....	68
Running Patrol Read.....	70
Scheduling a Patrol Read.....	70
Starting a Patrol Read Operation.....	71
Stopping a Patrol Read Operation.....	71
Managing SAS Storage Link Speed.....	71
Managing PCIe Storage Lane Speed.....	73
Setting the Adjustable Task Rates.....	74
Managing Power-Save Settings.....	75
Enabling and Disabling SSD Guard.....	76
Discarding Pinned Cache.....	76
Updating the Controller Firmware.....	76
Advanced Software.....	78
Activating Advanced Software.....	78
Advanced Software Status Summary.....	79
Activating a Trial Key.....	79
Activating an Unlimited Key over a Trial Key.....	80
Reusing the Activation Key.....	80
Application Scenarios and Messages.....	81
Securing Advanced Software.....	81
Configuring the Key Vault (Re-Hosting Process).....	82
Implementing the Re-Hosting Process.....	83
Snapdump Feature.....	83
Snapdump Feature Support.....	83
Retrieving the Snapdump Output.....	84
Deactivating Trial Software.....	85
Fast Path Advanced Software.....	86
SafeStore Encryption Services.....	86
Enabling Drive Security.....	86
Changing Drive Security Settings.....	88
Disabling Drive Security.....	90
Importing or Clearing a Foreign Configuration – Security-Enabled Drives.....	90
Managing Drive Groups.....	91
Viewing Drive Group Properties.....	91

Adding a Virtual Drive to a Drive Group	91
Managing Virtual Drives	92
Viewing Virtual Drive Properties	92
Modifying Virtual Drive Properties	93
Start and Stop Locating a Virtual Drive	94
Erasing a Virtual Drive	94
Initializing a Virtual Drive	96
Starting Consistency Check on a Virtual Drive	96
Deleting a Virtual Drive	96
Behavior of Virtual Drive Operations on VMware	97
Enabling the Schedule Panel Feature.....	97
Disabling the Schedule Panel Feature.....	98
Modifying the Time Limit of the Schedule Panel.....	98
Managing Physical Drives	99
Viewing Physical Drive Properties	99
Locating Tape Drives	100
Start and Stop Locating a Drive	101
Making a Drive Offline	101
Making a Drive Online	101
Replacing a Drive	101
Marking a Drive as a Missing Drive	102
Replacing a Missing Drive	104
Assigning Global Hot Spares	104
Removing Global Hot Spares	104
Assigning Dedicated Hot Spares	104
Rebuilding a Drive	105
Converting an Unconfigured Bad Drive to an Unconfigured Good Drive	105
Removing a Drive	106
Making Different Types of Drives	106
Making an Unconfigured Drive.....	106
Making a Good Drive.....	106
Making a JBOD Drive.....	107
Erasing a Drive	107
Erasing a Drive Securely	108
Managing Hardware Components	110
Monitoring Energy Packs	110
Viewing Energy Pack Properties.....	110
Refreshing Properties.....	111
Monitoring Enclosures	112

Viewing Enclosure Properties.....	112
Show Events.....	114
Downloading Events.....	115
Clearing Events.....	115
Customizing the Theme of the LSI Storage Authority Software.....	116
Viewing Default Theme Settings.....	116
Customizing the Logo.....	116
Customizing the Header Background Image.....	117
Introduction to the Light Weight Monitor System.....	118
Overview.....	118
Supported Operating Systems.....	118
Alert Delivery Methods Based on Severity Levels.....	118
System Log.....	119
Email Notification.....	119
Time Synchronization.....	120
Installing the Light Weight Monitor System.....	121
Installing the LSI Storage Authority Software on the Windows Operating System through the LSA Master Setup.....	121
Uninstalling Light Weight Monitor Software on the Windows Operating System.....	123
Installing the Light Weight Monitor on the Linux Operating System.....	123
Uninstalling the Light Weight Monitor on the Linux Operating System.....	124
Configuring the Light Weight Monitor System.....	125
Setting Up the Email Server.....	125
Adding the Email Addresses of Alert Notification Recipients.....	127
Configuring Alert Settings.....	128
Changing the Default Alert Delivery Method for Each Severity Level.....	128
Changing the Alert Delivery Method for a Specific Event.....	128
Changing the Severity Level for a Specific Event.....	129
Introduction to RAID.....	131
RAID Components and Features.....	131
Drive Group.....	131
Physical Drive States.....	131
Virtual Drive.....	132
Virtual Drive States.....	132
Fault Tolerance.....	132
Multipathing.....	133
Wide Port.....	133
Consistency Check.....	133
Replace.....	134

Background Initialization.....	134
Patrol Read.....	135
Disk Striping.....	135
Disk Mirroring.....	135
Parity.....	136
Disk Spanning.....	137
Hot Spares.....	138
Disk Rebuilds.....	138
Rebuild Rate.....	139
Hot Swap.....	139
Enclosure Management.....	139
RAID Levels.....	139
Summary of RAID Levels.....	139
Selecting a RAID Level.....	140
RAID 0 Drive Groups.....	140
RAID 1 Drive Groups.....	141
RAID 5 Drive Groups.....	142
RAID 6 Drive Groups.....	142
RAID 10 Drive Groups.....	143
RAID 50 Drive Groups.....	144
RAID 60 Drive Groups.....	145
RAID Configuration Strategies.....	146
Maximizing Fault Tolerance.....	147
Maximizing Performance.....	148
Maximizing Storage Capacity.....	149
RAID Availability.....	149
RAID Availability Concepts.....	149
Configuration Planning.....	150
Possible Raid Levels.....	150
Events and Messages.....	152
Error Levels.....	152
HTTP Status Codes and Description.....	153
SAS Address Assignment Rule.....	155
Multi-Selection Threshold for Virtual and Physical Drives.....	156
Known Issues and Workarounds.....	157
Glossary.....	159
Revision History.....	164

How This Guide Is 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, including monitoring and maintaining storage devices and the required hardware and software to run the application.
LSI Storage Authority Features	Outlines the LSI Storage Authority feature differences for MegaRAID®.
Types of Installation	Provides information on LSI Storage Authority installers and steps to install and uninstall the LSI Storage Authority.
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 Broadcom® 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.
Advanced Software	Provides information on certain premium features that the LSI Storage Authority supports on MegaRAID SAS 24Gb/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.
Show Events	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, such as adding your company logo or changing the default colors.

LSI Storage Authority Overview

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

- **Monitoring and Configuring:** LSA lets you 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.
- **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 and 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.
- **Personality Management:** The Personality Management solutions for MegaRAID controllers support personality modes in the form of a new JBOD (Just a Bunch of Drives) personality. This new JBOD personality is a unique personality derived from the standard RAID personality with additional features. The additional features that the JBOD personality provides are:
 - Switching between behavior modes; you can switch between **JBOD** and **None**.
 - Auto-configuration option.
The following additional features are supported in the new JBOD personality mode: RAID 0, RAID 1, RAID 10, PRL 11, and spanned PRL 11 are supported.
 - SEP is exposed to the operating system.
You can switch to JBOD personality mode from the standard RAID personality mode to use these additional features.
 - In addition to SAS and SATA drives, starting with LSA 2.4 and later, nonvolatile memory express (NVMe) drives are also supported.

Support Matrix

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

Table 1: Hardware and Software Support Matrix

Supported Items	Version
Supported controllers	<ul style="list-style-type: none"> • 9600 Family eHBA Adapters • 9620 Family eHBA Adapters • 9660 Family RAID Adapters • 9670 Family RAID Adapters
Supported operating systems	<p>Microsoft</p> <p>It is recommended you keep the Microsoft Windows operating systems up to date.</p> <ul style="list-style-type: none"> • Windows Server 2022 • Windows Server 2019 • Windows 10 (RS5) <p>Linux</p> <ul style="list-style-type: none"> • Red Hat Enterprise Linux 8.4 • Red Hat Enterprise Linux 8.3 • Red Hat Enterprise Linux 8.2 • SUSE Linux Enterprise Server 15 SP3 • SUSE Linux Enterprise Server 15 SP2 <p>VMware</p> <ul style="list-style-type: none"> • VMware vSphere 7.0 U2 • VMware vSphere 7.0 U1 • VMware vSphere 7.0 <p>Ubuntu</p> <ul style="list-style-type: none"> • Ubuntu 20.04 LTS
Supported web browsers	<ul style="list-style-type: none"> • Microsoft Edge 94.0 and above • Mozilla Firefox version 89.0 • Google Chrome Version 90
Supported networks	<ul style="list-style-type: none"> • Internet Protocol versions 4 and 6 • Network Address Translation • Lightweight Directory Access Protocol (LDAP) • Domain • HTTP, HTTPS
Supported languages	<ul style="list-style-type: none"> • English See Domain Authentication Behavior for non English server support.

Technical Support

For assistance with running or configuring the LSI Storage Authority, contact a Broadcom 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 Technical Support: <https://www.broadcom.com/support/call-us>.

LSI Storage Authority Features

The following tables outline the LSA features for MegaRAID controllers with respect to software features and firmware features.

Table 2: Firmware Features

Feature Name	MegaRAID
RAID level	RAID 0, RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, and RAID 60
Maximum configurable physical drives	64
Maximum spans	8
Maximum virtual drives	64
Dimmer switch	DS-I and DS-II
Maximum media errors	256
Strip size support	64 KiB and 256 KiB
Maximum VD's per drive group	16
Multipath	Yes
Controller reset support	Yes

Table 3: Software Features

Feature Name	Description
LDAP authentication	Supported network authentication protocol.
Server discovery and managing servers	Ability to set up a list of servers to monitor and manage.
Server dashboard	Displays the overall summary of the server and devices attached to it.
Controller dashboard	Interface in which you can perform controller-related actions and view all the information pertaining to a controller.
Simple configuration	Specifies a limited number of settings and has the system select drives for you.
Advanced configuration	Allows you to choose additional settings and customize virtual drive creation.
Foreign configuration (import/clear)	A RAID configuration that already exists on a replacement set of drives that you install in a computer system.
Clear configuration	Clears all existing configurations on a selected controller.
Update firmware	Update the controller firmware.
Online firmware update	Update the firmware while the system is online.
Controller Operations	
Setting consistency check properties	Set the consistency check properties.
Scheduling consistency check	Schedule a consistency check.

Feature Name	Description
Setting patrol read properties	Set the patrol read properties.
Starting patrol read	Start a patrol read.
Stopping patrol read	Stop a patrol read.
Managing link speed	Change the link speed between the controller and an expander or between the controller and a drive that is directly connected to the controller
Setting adjustable task rates	Set the adjustable task rates.
Manage power-save settings	Manage the power-save settings to reduce power consumption of the devices connected to a controller.
Enable and disable SSD Guard™	Enable or disable the SSD Guard feature.
Enable and disable security	Enable or disable controller security.
Change drive security	Change the drive security settings. Drive security settings cannot be changed when EKM is enabled.
Discarding preserved cache	Discard the controller data preserved from the virtual drive. This option is only available if pinned cache is present on the controller.
Background operations support	Pause, Resume, Abort, Pause All, Resume All, and Abort All features that enhance the functionality where the background operations running on a physical drive or a virtual drive can be paused for some time, and resumed later.
Personality management	Allows you to switch between different personality modes that are supported by your firmware.
Advanced Software Features	
Fast Path	A high-performance I/O accelerator for SSD arrays connected to a MegaRAID controller card.
SafeStore™	Used to secure drive data from unauthorized access or modification resulting from theft, loss, or repurposing of drives.
RAID 5 and RAID 6	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.
Modify drive group	Allows you to add, remove, or directly change the RAID level.
Secure using FDE	Secure a drive using Full-disk encryption (FDE).
Virtual Drive Operations	
Virtual drive settings and modifying virtual drive properties	Allows you to set and modify virtual drive properties.
Start and stop locating a virtual drive	Identify a virtual drive within a disk enclosure.
Erasing a virtual drive	Overwrite user-accessible locations or deletes a virtual drive.
Initializing a virtual drive	Initialize a virtual drive after completing the configuration process.

Feature Name	Description
Starting consistency check on a virtual drive	Start a consistency check on a virtual drive.
Deleting a virtual drive	Delete virtual drives on a controller to reuse that space for new virtual drives.
Physical Drive Operations	
Assign global hot spare	Assign a global hot spare to replace a failed physical drive in a redundant array.
Remove a global hot spare	Remove a global hot spare and lists the drive in the Unconfigured Drives section.
Assign a dedicated hot spare	Assign a dedicated hot spare drives to provide protection to one or more specified drive groups on the controller
Remove a dedicated hot spare	Remove a dedicated hot space and lists the drive in the Unconfigured Drives section.
Start and stop a locating drive	Identifies a physical drive by making their LEDs blink.
Making a drive online and offline	Changes the state of a physical drive.
Replacing a drive	Replaces a failing or failed drive.
Rebuilding a drive	Automatically rebuilds the data onto a hot spare to prevent data loss if a drive fails.
Prepare for removal	Prepares a physical drive for removal by spinning the drive into a power save mode.
Erasing a drive	Erases data on non-SEDs (normal HDDs).
Instant secure erase	Erases data from encrypted drives.
Converting an unconfigured bad drive to an unconfigured good drive	Convert an unconfigured bad drive to an unconfigured good drive.
Make an unconfigured good drive	Changes the status of a JBOD drive to an Unconfigured Good drive.
Make a JBOD and delete a JBOD	Changes the status of an Unconfigured Good drive to a JBOD drive.
Energy Pack Operations	
Auto Learn	Calibration operation that the controller performs to determine the condition of the energy pack.
Event Logs	
Viewing event logs	Views event logs, which contain information about the activity and performance of the server and all of the controller cards attached to the server.

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, and while upgrading the LSI Storage Authority.

Preinstallation Requirements

You must complete this task before you install LSI Storage Authority (LSA):

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

NOTE

You must stop LSA services if configuring a SAS 3.5 controller pass-through using DDA in a Windows environment.

OpenSLP

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

If OpenSLP Is Not Installed

No action is required. With the latest versions of LSA, OpenSLP is bundled within the LSA package itself. 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 does not discover the servers from the **Remote Discovery** page.

Uninstalling or Cancelling an LSA Installation

OpenSLP will remain on the system if LSA is uninstalled or the LSA installation is cancelled.

Upgrade Requirements

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

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.

Types of Installation

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

- Gateway
- StandAlone
- DirectAgent
- Light Weight Monitor (LWM)
- Light Weight Agent (LWA) Installer

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

Table 4: Types of Installers and Associated Support

Feature	Gateway Installer	StandAlone Installer	DirectAgent Installer	Light Weight Agent (LWA) Installer – ESXi Server Support
Permits discovery of other servers that run the LSI Storage Authority software	Yes	No	No	No
Permits self-registration using OpenSLP and has an interface for server discovery detection from the network	Yes	Yes Note: No interface exists for remote server discovery.	No	No
Allows the management of servers from the list of discovered servers through the user interface (UI)	Yes	No	No	No
Provides capability to configure LDAP information	Yes	Yes	No	No
Provides server monitoring capabilities and helps to monitor the health of the server and alerts the end-user of any issues with event logs and email notifications	Yes	Yes	Yes	Yes

NOTE

In a new LSA installation, LSA processes only the last 30 events and performs the corresponding alert delivery.

Gateway Installer

The Gateway installer has the following components:

- A backend 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.
- Permits self-registration using OpenSLP and has an interface for server discovery detection from the network.
- Allows you to manage the servers from the list of discovered servers through the user interface.

StandAlone Installer

The StandAlone installer has the following components:

- A backend with a 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.
- Permits self-registration of the current host using OpenSLP but does 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 UI.

DirectAgent Installer

The DirectAgent installer has the following components:

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

Light Weight Monitor

The Light Weight Monitor has the following benefits:

- Provides server monitoring capabilities.
- Light Weight Monitor (LWM) monitors the status of the controller cards, virtual drives, drives, and other devices on the server.
- Alerts you of any issues that require immediate attention with system logs and real-time email notifications (based on the alert settings).

Light Weight Agent Installer

The Light Weight Agent installer has the following components:

- A backend with a local agent (without remote agent management capability) to support ESXi 7.x.
- A monitor (without remote monitoring capability).
- A client (without remote and managed server capabilities).

The Light Weight Agent installer has the following features and limitations:

- Does not permit the discovery of other hosts that are running the LSI Storage Authority.
- Does not permit self-registration of the current host using OpenSLP and does not have any interface for server discovery detection from the network.
- Does not provide capability to configure LDAP information.
- Does not permit to add managed servers through the UI.

Configuring HTTPS

To configure HTTPS, perform the following steps:

1. Stop **LSAService** and **NginxService** in Windows, and **/etc/init.d/LsiSASH stop** in Linux.
2. Replace the `nginx.conf` file with `LSA_HOME\server\conf`.
3. Change the `nginx_default` value in `nginx.conf` file to the `Web Server Port` value user selected during the LSA installation.
If necessary, change the same value in multiple places in the same file.
4. Change the `LSA_Default` value in `nginx.conf` file to the `LSA Server Port` value user selected during the LSA installation.
If necessary, change the same value in multiple places in the same file.
5. Place "ssl.crt" and "ssl.key" into "LSA_HOME\server\conf"
6. Go to `LSA_HOME\conf` directory and change the `protocol_type` value to 1 (HTTPS) in `LSA.conf`.
7. Start **NginxService** and **LSAService** in Windows and **/etc/init.d/LsiSASH start** in Linux.
8. From the `<LSA_HOME>` directory, open the file `startupLSAUI.bat`, and start `http` to start `https`.

NOTE

When the certificate and key files are created, rename them to `ssl.crt` and `ssl.key`.

To create your own SSL self-signed certificate, see http://www.akadia.com/services/ssh_test_certificate.html.

Resetting Encryption Keys

LSA maintains sensitive data in an encrypted format. To encrypt data, LSA uses an AES algorithm with a key pair.

NOTE

Ensure the user has administrative privileges before attempting to reset the key pair.

To change the key pair, use the following procedure:

1. Stop the LSA service.
 - a) Launch the command prompt in Windows.
 - b) To stop the LSA service, use the `-> sc stop LSAService` command.

NOTE

To stop the LSA service in Linux, use the `LSA.sh` command.

2. Enter the `>$LSA_HOME/LSA.exe -rekey` command.
LSA will automatically delete the old keys and encrypt the sensitive data with a new key pair.
3. Start the LSA service.
 - a) Launch the command prompt in Windows.
 - b) To start the LSA service, use the `-> sc start LSAService` command.

Installing the LSI Storage Authority Software on the Microsoft Windows Operating System

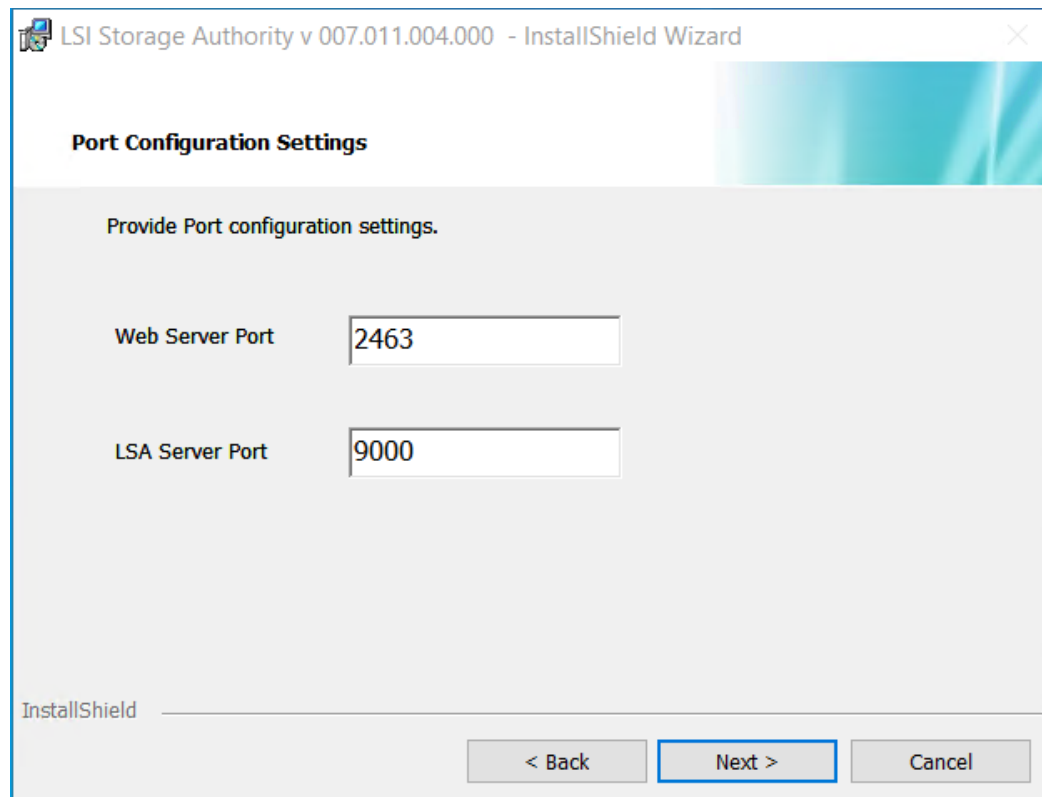
Perform the following steps to install the LSI Storage Authority.

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 select 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

LSI Storage Authority v 007.011.004.000 - InstallShield Wizard

Port Configuration Settings

Provide Port configuration settings.

Web Server Port

LSA Server Port


InstallShield

< Back Next > Cancel

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 also edit this port's details after installation. See [Changing the LSI Storage Authority Application Port Number](#) and [Changing the nginx Web Server Port Numbers](#).

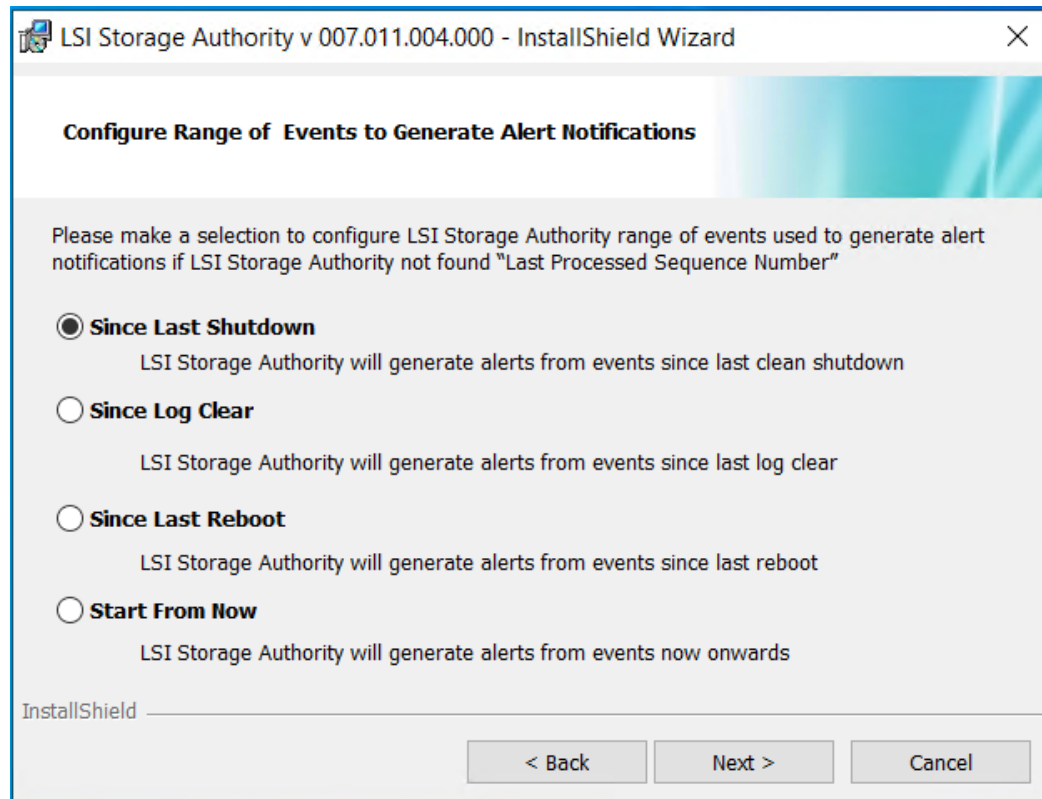
5. Click **Next** to proceed.

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 and 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, 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), 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, 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, 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), 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`).

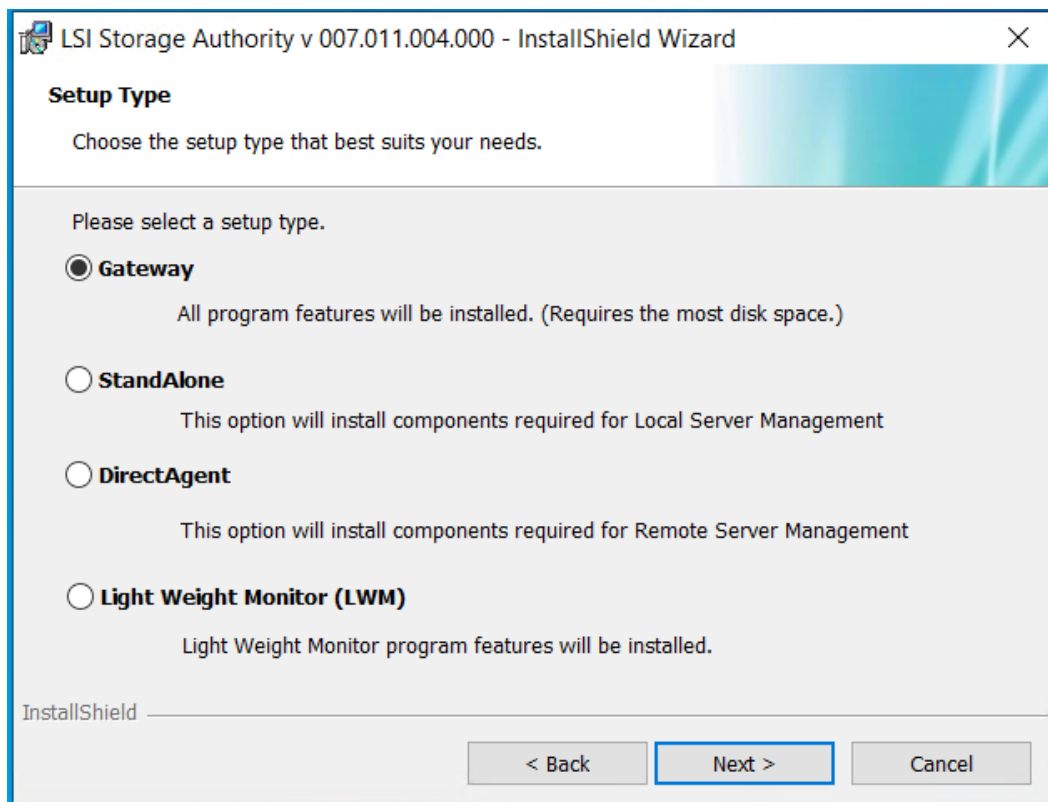
NOTE

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, see [Table 1](#), .

- **Gateway**
- **StandAlone**
- **DirectAgent**
- **Light Weight 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** check box 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. Ensure that port 2463 is not blocked by a firewall.

The Windows Firewall settings are located under **Control Panel > Windows Firewall**.

13. Click **Finish**.

Installing in 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 noninteractive mode:

1. From the command line, run the `vc_redist_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 (`vc_redist_x86.exe`) is available under the directory `<Package_Dir>\ISSetupPrerequisites\{270b0954-35ca-4324-bbc6-ba5db9072dad}\VC Redist 2008 Installation`.

OpenSLP is bundled with LSA 2.2 and later. While installing LSA, if OpenSLP is not installed ensure that you select the option to install OpenSLP, and LSA seamlessly installs the required version of OpenSLP. See [OpenSLP](#) for more information.

2. Depending on the type of installation required, run the `setup.exe /s /v/qn ADDLOCAL=` command. The types of installation and their associated alert notifications available are as follows:

Type of Installation	Type of Event Notification	Event Notification Choice
Gateway	Since Last Shutdown	0
StandAlone	Since Log Clear	1
DirectAgent	Since Last Reboot	2
Light Weight Monitor	Start From Now	3

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

Uninstalling in Interactive Mode

You can uninstall the LSI Storage Authority 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 2012 operating systems, select **Add/Remove Programs** from the **Control Panel**. If you are using the 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**.

Uninstalling in 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 step to install the LSI Storage Authority software in noninteractive mode:

From the command line, run the `msiexec.exe /x <productcode>/qn` command to uninstall LSA.

Where `<productcode>` is a unique product code associated with each LSA installation and `<LSA_HOME_PATH>` is the location where the LSA is installed.

Example: `msiexec.exe /x {20660CCB-7C70-4D61-8D18-FB7FA3C476C9}/qn`

Before you begin to uninstall LSA, if any file has been manually copied to the `<LSA_HOME>` directory by you other than the standard installation package contents, make sure you delete those files. If you fail to remove the files that have been manually copied to the `<LSA_HOME>` directory, the uninstallation process may fail.

Installing the LSI Storage Authority Software on the Linux Operating System

The LSI Storage Authority software supports both the interactive and the noninteractive modes of Linux installation.

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 license agreements, press **Y**. Otherwise, press **N** to exit the installation.
3. 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.

4. 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.
 - **Light Weight Monitor** – Press **4**. Selecting this option installs the Light Weight Monitor program features.
5. Enter the nginx server port number. The port range is from 1 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 to 65535 and are 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 also edit this port details after installation.

7. Ensure that port 2463 is not blocked by a firewall.
8. Extract the contents of the zip file and install the 64-bit Linux operating systems. The `LSA_Linux.zip` file contains files for 64-bit platforms.

NOTE

LSA only supports 64-bit operating systems.

NOTE

Ensure that **Connect automatically** check box is selected, which is available under **Network Connections**.

9. To launch the LSI Storage Authority, navigate to your browser and enter your IP Address followed by:2463. For example, `http://135.24.237.36:2463`.

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.

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 64-bit Linux operating systems. The `LSA_Linux.zip` file contains the files for the 64-bit platform.

Uninstalling the LSI Storage Authority Software on the Linux Operating System

Perform the following step to uninstall the Linux operating system.

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

Command usage example: `rpm -e LSIStorageAuthority-1.00xx.xxxx-xxxx`

Adding a Digital Signature and Enabling ASLR and DEP for LSI Utilities

The LSI Storage Authority software supports both address space layout randomization (ASLR) and data execution prevention (DEP).

Adding a Digital Signature in a Windows Environment

To add a digital signature in a Windows environment:

1. Ensure that the `Setup.exe` has been signed.
2. If necessary, install the signed certificate.
The signed certificate will not be installed during the installation.
3. Run the `signtool.exe verify /v /pa <LSA executable name>` command to complete the signature verification.

NOTE

The `signtool.exe` is required to validate the signature.

Adding a Digital Signature in a Linux Environment

To add a digital signature in a 64-bit environment:

1. Ensure that only the LSA RPM has been signed.
During the installation of LSA, the signed certificate will be imported into the servers.
The certificate will not be removed during the uninstallation or upgrade of LSA.
2. Verify the LSA RPM signature.
 - a) Import the public key to RPM DB using the following command: `rpm --import pubKey.asc`.
 - b) Verify the RPM signature using the following command: `rpm -Kv<LSIStorageAuthority-xxx.xxx.xxx.xxx.rpm>`.
3. Upon successful installation, the following message is displayed:

```
md5 gpg OK
```

Managing Light Weight Agent

Previously LSA support for VMware ESXi 6.x was achieved using an SMI-S provider. The Broadcom Light Weight Agent (LWA) software provides server monitoring and management capabilities. LWA is a stand-alone installation on an ESXi server without SLP or LDAP support. LWA has a smaller footprint and provides the same set of features as the StandAlone installation.

See [Light Weight Agent Installer](#) for information on features and limitations.

Installing LWA on VMware ESXi

To deploy the Light Weight Agent (LWA) on a VMware ESXi machine, use the VIB file provided by Broadcom.

1. Download the `LSA_ESXi.zip` file.
2. Extract the `vmware-lsa-xxxx-xxxx-xx.vib` (x86_64 arch only) file from `LSA_ESXi.zip`.
3. Copy `vmware-lsa-xxxx-xxxx-xx.vib` to the host VMware ESXi machine.
4. Run the following `esxcli` command to install the VIB.

```
esxcli software vib install --viburl=/vmware-lsa-xxxxxxxx-0x.vib --no-sig-check --force
```

This step installs LSA to the `/opt/lsa` directory and starts the service. The service can be managed using `/etc/init.d/lsad {start|stop|restart}`.

LSA is running along with the `lighttpd` server.

Uninstalling LWA

To uninstall LWA, run the following command:

```
esxcli software vib remove -n <name of LSA package>
```

Managing LSA and Web Services

Use the following commands to manage LSA and Web Service using LWA.

- To start LSA and `lighttpd`, run `/etc/init.d/lsad start`.
- To stop LSA and `lighttpd`, run `/etc/init.d/lsad stop`.
- To restart LSA and `lighttpd`, run `/etc/init.d/lsad restart`.
- To check status of LSA and `lighttpd`, run `/etc/init.d/lsad status`.

Configuring the Firewall on Various LSA Installers

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

- Gateway/StandAlone installer configuration on the Windows operating system.
- DirectAgent installer configuration on the VMware operating system.
- Gateway/StandAlone installer configuration on the Linux operating system.

NOTE

Ensure that port 2463 is not blocked by a firewall.

Configuring the Firewall on Gateway/StandAlone Installer (Windows)

You can configure the Gateway/StandAlone installer firewall on the Windows operating system.

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

The Windows operating system has three profiles, **Public**, **Private**, or **Domain**. You must enable one of these profiles appropriately based on your connection type.

After 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 LSI Storage Authority 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 web server 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.

Configuring the Firewall on Gateway/StandAlone Installer (Linux)

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

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

```
iptables -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 the `iptables -I INPUT -p tcp -dport 2463 -j ACCEPT` command.

Configuring the Firewall on DirectAgent Installer

LSA uses the SMI-S Provider to discover and manage its storage controllers on the VMware environment. To discover the ESXi servers where SMI-S Providers are available, LSA uses the SLP as a discovery mechanism. CIMSLP 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 the 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 is 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.

Collecting LSA Logs on Windows and Linux Operating Systems

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

Perform the following steps:

1. Stop the 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. After you see the issue, share the logs.txt file from the LSA_HOME/logs directory.

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:

Figure 6: Behavior of Event History

The screenshot shows a web interface for viewing event logs. At the top, there is a navigation bar with a back arrow and the text "Go back to Drive Group, Drives and Other Hardware list", and a "Close" button on the right. Below this is a "Show Events" section with a calendar icon. The main content area displays "Displaying latest log entries" and a table of log entries. To the right of the table is an "Actions" menu with "Download Events" and "Clear Events" options.

Severity Level	Event ID	Locale	Description	Time, Date
Information	590	Physical Device	Power state change on PD 0x102(e0x11a/s1) from ON(0) to POWERSAVE(1)	1:15:23 AM,6 Oct'21
Information	590	Physical Device	Power state change on PD 0x101(e0x11a/s0) from ON(0) to POWERSAVE(1)	1:15:23 AM,6 Oct'21
Information	340	Physical Device	State change on PD 0x102(e0x11a/s1) from ONLINE(18) to UNCONFIGURED_GOOD(0)	12:44:33 AM,6 Oct'21
Information	340	Physical Device	State change on PD 0x101(e0x11a/s0) from ONLINE(18) to UNCONFIGURED_GOOD(0)	12:44:33 AM,6 Oct'21
Information	349	Virtual Drive Configuration	Deleted VD 0x1	12:44:33 AM,6 Oct'21
Information	287	Virtual Drive	Policy change on VD 0x1 to [ID=01 cp=00 dc=2] from [ID=01 cp=01 dc=2]	12:44:07 AM,6 Oct'21
Information	287	Virtual Drive	Policy change on VD 0x1 to [ID=01 cp=00 dc=2] from [ID=01 cp=01 dc=2]	12:44:07 AM,6 Oct'21
Information	498	Virtual Drive	VD 0x1 is available	12:44:06 AM,6 Oct'21
Information	348	Virtual Drive Configuration	Created VD 0x1	12:44:06 AM,6 Oct'21
Information	428	Virtual Drive	VD 0x1 is now OPTIMAL	12:44:06 AM,6 Oct'21
Information	340	Physical Device	State change on PD 0x102(e0x11a/s1) from UNCONFIGURED_GOOD(0) to ONLINE(18)	12:44:06 AM,6 Oct'21
Information	340	Physical Device	State change on PD 0x101(e0x11a/s0) from UNCONFIGURED_GOOD(0) to ONLINE(18)	12:44:06 AM,6 Oct'21
Information	340	Physical Device	State change on PD 0x104(e0x11a/s3) from HOT SPARE(2) to UNCONFIGURED_GOOD(0)	12:43:53 AM,6 Oct'21

- 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`.
 - The maximum number of events persisted in a file is 300. This can be configured through the `LSA.conf` file.
 - The 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.

Behavior of Event Monitoring

The following table details the behavior of event monitoring.

Table 5: Event Monitoring Behavior

Event	Description
SysLog or email	Immediately after the LSA service is started. This includes new installations as well as restart of LSA services.
Syslog location	On Windows: Event viewer. On Linux: <code>/var/log/messages</code>
Differentiating syslog	No special attribute is added as part of the event description.
LSA boot events	Boot events are handled for local controllers.
Time sync	Handled
Alert configuration	Follows <code>config-current.json</code>
SMTP server communication	LSA server communicates with SMTP server for email communication.

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.

IMPORTANT

All security related browser settings must be set to the standard security settings. LSA is web-based application, and any browser setting changes will impact LSA.

Performing Initial Configuration

After successfully installing the LSI Storage Authority, you must set up these initial configurations.

Domain Authentication Behavior

The default installation of LSA contains entries in a properties file for administrators and authenticated users. The properties file contains two access groups for LSA.

- `full_access_groups` – Administrators and non-administrators
- `readonly_access_groups` – Authenticated users

LSA supports multiple comma (,) separated group names for both administrators and authenticated users. Users that are part of both groups are assigned administrative access to LSA. Server administrators can edit entries or add multiple groups to the individual entries.

Users who are part of the `full_access_group` will have full access to LSA even if they are not an administrator. If the administrator group value is removed from the `full_access_group`, an error message will display *Invalid Credentials*, and you will not be able to access LSA.

Users who are part of the `readonly_access_group` will be given read-only access by default. If the authenticated users group value is removed from the `readonly_access_group`, an error message will display *Invalid Credentials*, and you will not be able to access LSA even in read-only mode.

LSA will prompt users to provide the domain name, user name, and password of the server they are trying to access. Once the credentials are verified, users will given access based on their assigned access level.

Adding Translated Names in Windows

Complete the following steps for non English versions of Windows.

1. Stop the `LSAService`.
2. Browse to `LSA_HOME/Conf` directory and open the `LSA.conf` file.
3. Add the translated names of the Administrators group to the `full_access_groups` in the `LSA.conf` file.
This will allows users to have full access privileges in the LSA post login.
4. Add the translated names of the Authenticated Users group to the `readonly_access_groups` in the `LSA.conf` file.
This will allows users to have read access privileges in the LSA post login.
5. Restart the `LSAService` and proceed with the login.

Adding Translated Names in Linux

Complete the following steps for non English versions of Windows.

1. Stop the `LSAService` by using the `/etc/init.d/LsiSASH stop` command.
2. Navigate to `LSA_HOME/Conf` directory and open the `LSA.conf` file.
3. Add the translated names of the root group to the `full_access_groups` in the `LSA.conf` file.
This will allows users to have full access privileges in the LSA post login.
4. Add the translated names of the users group to the `readonly_access_groups` in the `LSA.conf` file.
This will allows users to have read access privileges in the LSA post login.

5. Restart the `LSAService` using the `/etc/init.d/LsiSASH start` command, and proceed with the login.

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 backend must know some parameters to enable authentication. Perform the following steps to configure these parameters in the `lsa.conf` file in the `LSIStorageAuthority/conf` directory.

1. Open the `lsa.conf` file in the `LSIStorageAuthority/conf` directory.
2. Enter a value for the `ldap_mode` field. If you set it 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 host name of the LDAP server in the `ldap_server` field. This value connects the specific LDAP server for the user authentication.

Example:

```
# LDAP Server
ldap_server = <IP Address of the LDAP server>
```

4. (Optional) Enter the LDAP protocol version in the `ldap_protocol_version` field. This value defines the protocol 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 method 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) Enter the LDAP server port number in the `ldap_port_number` field.

Example:

```
# LDAP Port Number = 636
ldap_port_number = 389
```

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

Example:

```
# LDAP_DN_DETAILS
dn_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) 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 effect.

Accessing LSA over Network Address Translation

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 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 `LSIStorageAuthority/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 must 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 must specify them as `nat_ipv4_ipv6 = 135.24.227.198, 135.24.227.199, fe80::dc8d:e156:41e1:b06`

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 `LSIStorageAuthority/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 `LSIStorageAuthority` 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.

Hiding an Empty Backplane

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

1. Stop the `LSAService`.
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 `LSAService` once again.

Changing the nginx Web Server Port Numbers

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

1. Open the `nginx.conf` file in the `LSIStorageAuthority/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.

Changing the nginx Read Timeout

On VMware, when you request process-intensive operations such as creating **64 Virtual Drives**, **Drive Initialization**, **Consistency Check**, **Drive Erase**, and so on, the VMware ESXi Server may time out, resulting in a delay of the operation that is being performed.

To avoid the VMware ESXi Server getting timed out, perform the following steps to change the nginx FCGI Read Timeout.

1. Open the `nginx.conf` file in the `LSIStorageAuthority/server/conf` directory.
2. In the `nginx.conf` file, search for the `fastcgi_read_timeout` field.
3. Modify or increment the value present in the `fastcgi_read_timeout` to anywhere between 900 to 2000 depending on your requirement.
4. Save the `nginx.conf` file.
5. Restart the nginx service as well as the LSI Storage Authority Service.

Performing Initial Setup

After you successfully log in to LSI Storage Authority, you must perform some initial setup tasks before proceeding.

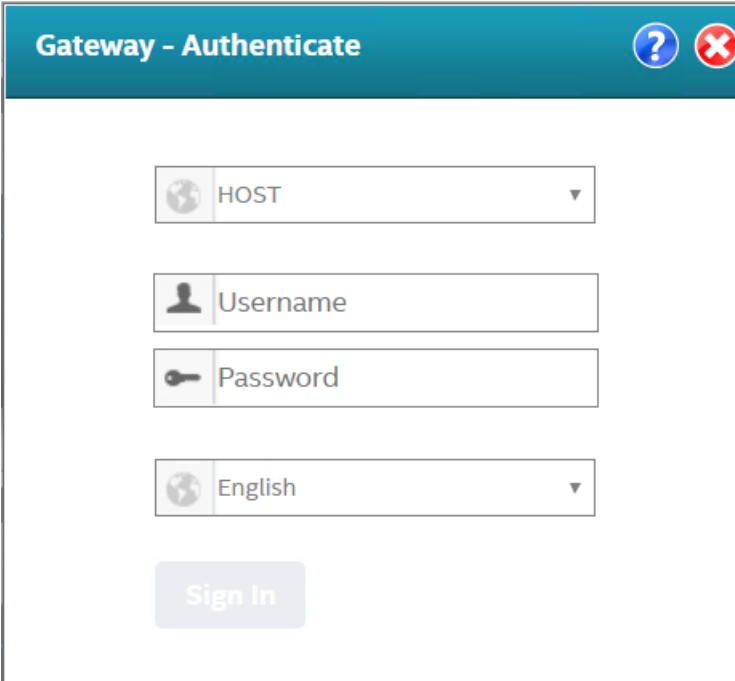
Managing Servers from the Remote Server Discovery Page

The LSI Storage Authority lets you 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 7: Gateway – Authenticate Dialog



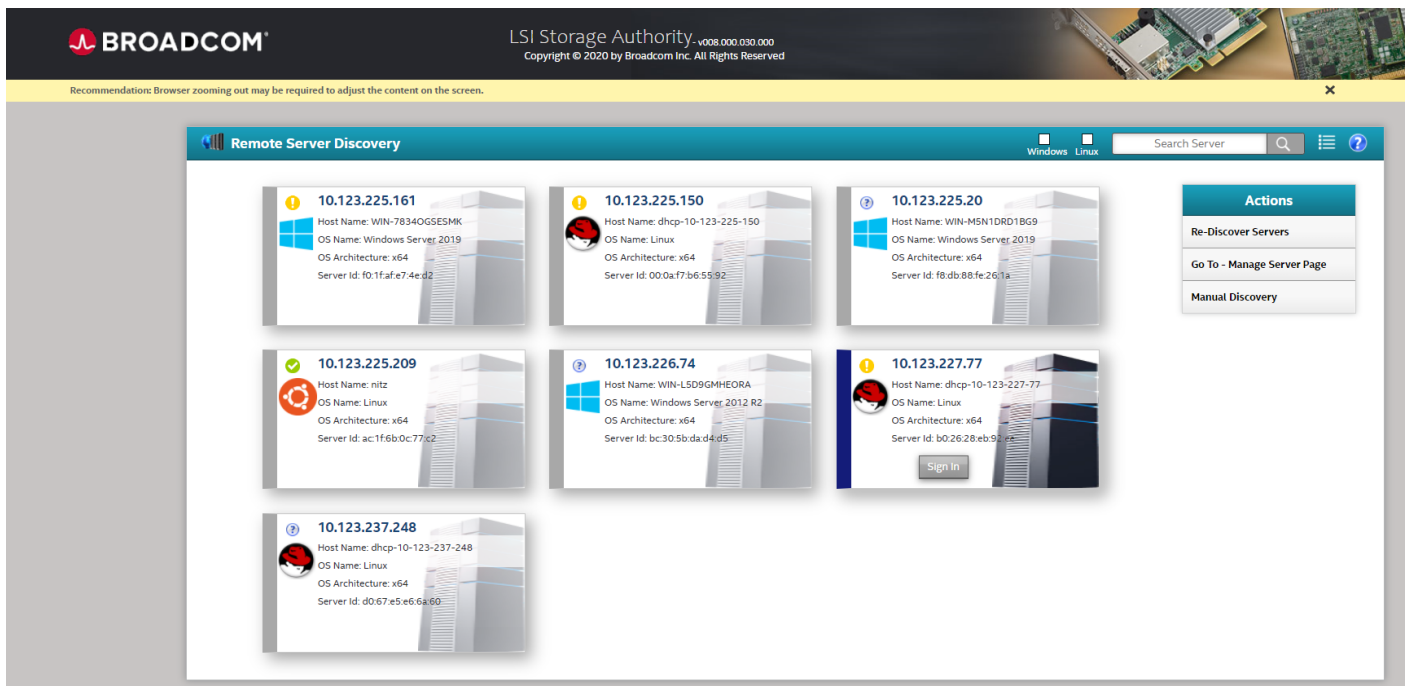
The screenshot shows a dialog box titled "Gateway - Authenticate". It contains the following elements from top to bottom:

- A teal header bar with the text "Gateway - Authenticate" and a blue question mark icon on the right.
- A dropdown menu with a globe icon and the text "HOST".
- A text input field with a person icon and the text "Username".
- A text input field with a key icon and the text "Password".
- A dropdown menu with a globe icon and the text "English".
- A light blue button with the text "Sign In".

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 **Manage Servers** page.

Figure 8: Remote Server Discovery Page



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

- View 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.
- Manually discover the servers using either the IP address or the host name.
See [Manually Discovering Servers](#).
- Using the check boxes of different operating system types, you can select only those servers that are associated with that particular operating system, to be displayed.
- Specify the IP address of the server in the search box to display that particular server.
- You can toggle between **List** and **Grid** view by clicking the *List/Grid* icon.
By default, LSA displays all the servers in Grid view.

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 9: Remote – Authenticate Dialog

The dialog box is titled "Remote - Authenticate" and features a blue header bar with a question mark icon and a red close button. Below the header, there are four input fields: a dropdown menu with a globe icon and the text "HOST", a text field with a person icon and the text "Username", a text field with a key icon and the text "Password", and a dropdown menu with a globe icon and the text "English". At the bottom center, there is a light blue button labeled "Sign In".

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 the  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](#).

Once the administrator user logs into the Manage Server page, the local Gateway server will be added to the Manage Server list.



There is not an option to add the self server because the user is already logged in with full access.

Users will not see the Verify option for the self server or local gateway server because the details are not stored locally.

Removing Managed Servers

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

Click the  icon.

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

Manually Discovering Servers

LSI Storage Authority discovers servers that are in the same subnet automatically. However, using the **Manual Discovery** option that is available on the **Remote Server Discovery** page, you can manually discover the servers using either the IP address or the host name of the server. **Manual Discovery** option enables you to discover and manage servers that are outside the subnet.

Manual Discovery server health will not get updated unless the server is part of **Auto Discovered** list. If the **Manual Discovery** server is part of the **Auto Discovered** list, then the **Manual Discovery** server will be removed and will be displayed as Auto Discovered Server.

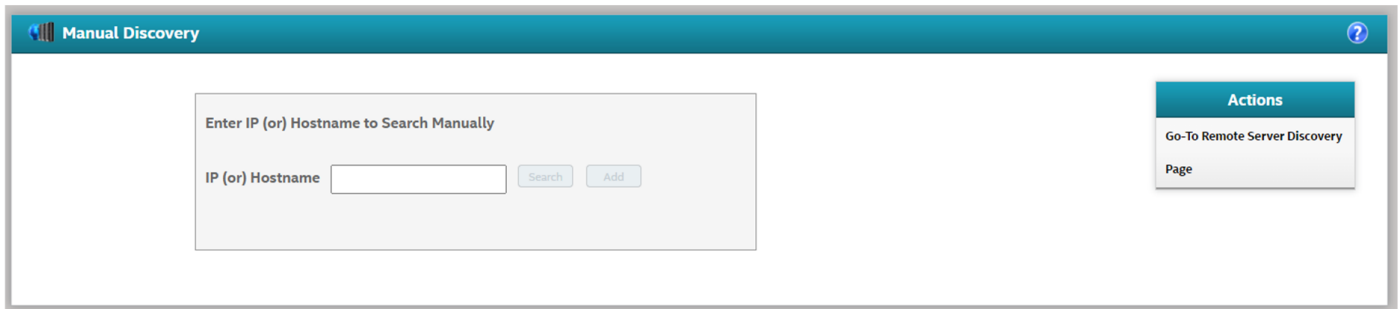
Adding the same server to **Manual Discovery** in different forms (different IP/Hostname) is not recommended. Adding a server with the same name will cause the same server to appear multiple times in the **Manual Server** or **Remote Discovery** pages.

Perform the following steps from the **Remote Server Discovery** page to manually discover servers that are outside the subnet.

1. On the **Remote Server Discovery** page, click **Manual Discovery**

The Manual Discovery dialog appears.

Figure 10: Manual Discovery Dialog



The screenshot shows a web-based dialog titled "Manual Discovery". At the top left is a logo and the title "Manual Discovery". At the top right is a help icon. The main content area has a heading "Enter IP (or) Hostname to Search Manually". Below this is a text input field labeled "IP (or) Hostname" with "Search" and "Add" buttons. On the right side, there is an "Actions" panel with a button labeled "Go-To Remote Server Discovery Page".

2. Specify either the IP address or the host name of the server that is outside the LSA subnet, which you wish to discover.

3. Click **Search**

Based on the IP address or the host name of the server specified by you, the corresponding server is displayed.


4. Click **Add** to add the newly discovered server to the list of Managed Servers.

a) (Optional) If you do not need this server to be listed in the Managed Servers list, click the **Remove Manually Added Server** icon to remove the server.

5. Click **Sign In** to login to the server

6. Specify the administrator credentials for the server and click **Sign In**.

The manually added server is now displayed or listed on the Remote Server Discovery page.

Figure 11: Remote Server Discovery – Manually Added Server

You can also choose to remove this server from the Remote Server Discovery list by clicking on the **Remove Manually Added Server** icon.

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 system log (syslog). In the Windows operating system, the system log is logged in **Event Viewer > Application**. In the Linux operating system, the system log is logged in `var/log/messages`.
- **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. Select **Username > Settings** in the Server dashboard.
The **Alert Settings** window appears with 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 selecting 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 settings.

NOTE

Restore Default Alert Settings will not save the changes by default. To restore the default alert settings, click **Save Settings**.

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

Alert Delivery **Mail Server** **Email**

Provide mail and server settings from which the application will send alert notifications.
Displaying current mail server settings

Sender Email Address: isa-monitor@server.com

SMTP Server: 127.0.0.1

Port: 25 Use Default

Security Protocol: NONE SSL SSL/TLS

For server authentication, please provide the following *(optional depending upon the server settings)*

This server requires authentication

User Name: _____

Password: _____

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. If the SMTP server requires authentication, select the **This server requires authentication** check box and specify the authentication details in the **User Name** and **Password** fields.
6. Click **Save**.

Adding Email Addresses of Recipients of Alert Notifications

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

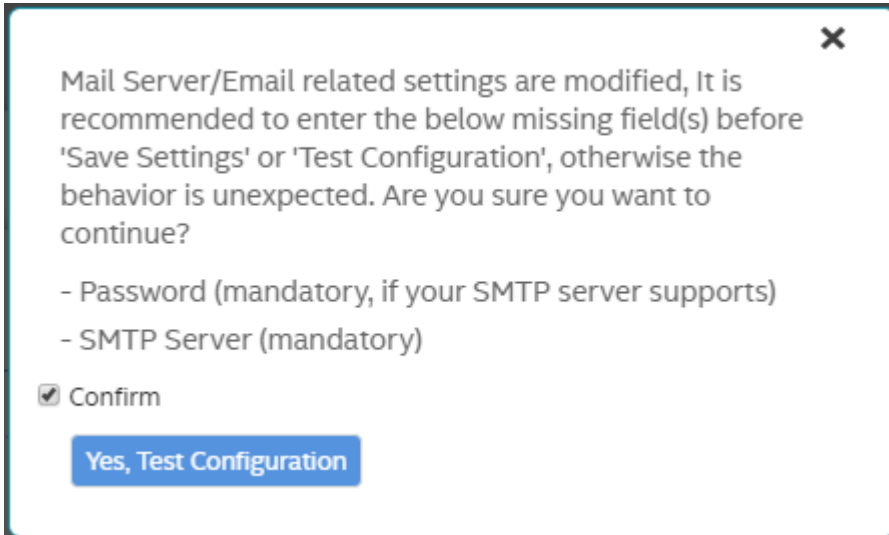
1. In the **Setting** window, click the **Email** tab.

The **Email** tab appears and displays the current email settings.

Figure 14: Email Window

The screenshot shows the 'Email' configuration window. At the top, there are three tabs: 'Alert Delivery', 'Mail Server', and 'Email'. Below the tabs, the main heading reads 'Provide email addresses to which the email alert notifications will be sent.' followed by the subtitle 'Displaying current email settings'. The main content area is a light blue box containing the following elements: an 'Add Email Address' section with an empty text input field and an 'Add' button; an 'Email alerts will be sent to the following email ids' section with a list item 'root@localhost' (preceded by a checkbox) and a 'Remove' button; and a blue 'Test Configuration' button at the bottom.

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 **Test Configuration** to send an email to the addresses that you added for the recipients of alert notifications.
A test configuration verification popup appears.



5. Verify the **Mail Server** and **Email** settings are correct.
6. If the STMP server supports authentication, ensure the password field in **Mail Server** tab is populated.
7. Click **Yes, Test Configuration**.
A popup message indicates if the test message was successfully sent to the email address.
8. Click **Save** to save the email settings.

Server Dashboard

After you log in, 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

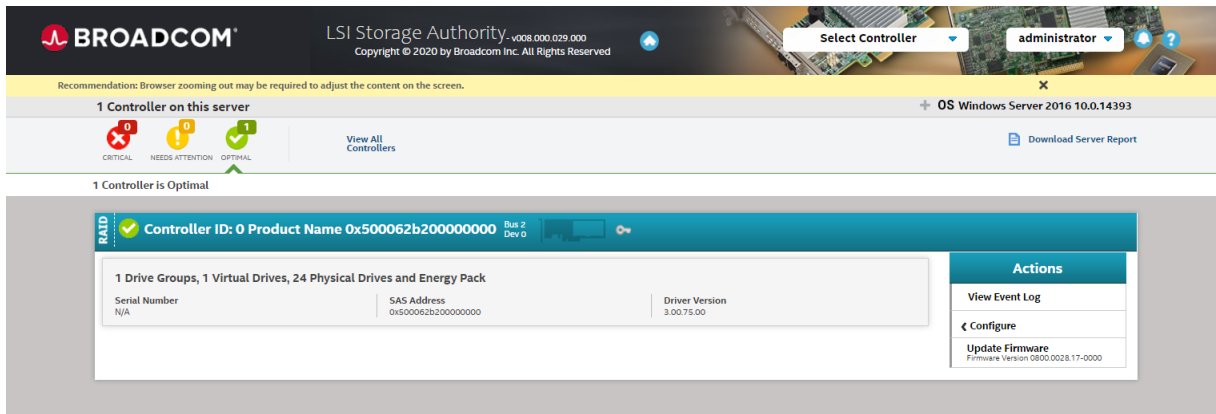


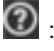


Table 6: Server Dashboard Description

Description
<p>Main Navigation – The main navigation window helps you to traverse among the various views. This navigation is available across all of the pages in the software. The description follows:</p> <ul style="list-style-type: none"> • : Helps you to navigate to the Server dashboard from any page in the software. • Select Controllers: Lists the controllers that you are monitoring. The color-coded controller status icons (red, amber, and green) indicate the health status of all the controllers based on their criticality. Click a controller to navigate to its dashboard. • Username: Displays the name of the user. <ul style="list-style-type: none"> – Click Settings to perform initial settings. – Click Send Feedback to email your feedback to the Broadcom Technical Support using your Gmail or Microsoft Outlook accounts. – Click View Server Profile and expand the + button to view the server configuration, such as the server IP, server name, OS Name, OS version, OS architecture, and the version of the LSI Storage Authority software that is installed. You can also view the controller information, such as controller hardware, enclosure of the controller, and information about the physical drives and virtual drives associated with the controller. – Click Logout to exit from the software. • : Lets you enable or disable system messages. • : Displays the LSI Storage Authority software context-sensitive help.
<p>User's Menu – Provides a method to back out of the Server dashboard and log into the Remote Server Discovery Page or back to the Login Page.</p> <p>DownloadServer Report – Enables you to download the server report, which contains the following information. Downloading the server report is available only for admin users.</p> <ul style="list-style-type: none"> • LSA Server Report • Information on each individual controller: <ul style="list-style-type: none"> – Diagnosis report – Events report – TTY log – Corpus of individual entity – DEQUEUE log – If Snapdump is not supported – PL log – If Snapdump is not supported – Collection of Snapdump files, including all supporting indexes • LSA developer debug logs

Description
<ul style="list-style-type: none"> • Controller Information: Displays information about the controller. • Controller Personality: Displays the current personality of the controller. • Controller Status: When multiple controllers are connected, the controllers are sorted based on the Bus device function. The controllers are indexed with numbers 0, 1, 2, and so on. • Controller summary • Controller properties • Controller issues • Controller event logs • Allows you to perform the following tasks: <ul style="list-style-type: none"> – Configure the controllers. See Configuration. – Download diagnostics. – Update the controller firmware. – View, download, and clear event logs. – Perform various operations on the controller. See Managing Controllers – Navigate to any of the controllers to see its specific view by clicking on the appropriate controller.

Figure 16: View Server Profile Window

— Server Profile

Server IP: 135.15.165.237
Host Name: LEONSMCX10
OS Name: Windows Server 2016
OS Version: Windows Server 2016 10.0.14393
OS Architecture: x64
Application Version: 008.000.047.000
Storelib8 Version: 08.0000.0098

Controller ID: 0 Product Name 0x500062b200000000

— Hardware(s) Information

Controller ID: 0 Product Name 0x500062b200000000
Status: CRITICAL
Firmware Version: 8.0.0.100-00000-00001
Driver Version: 8.00.02.00
Enclosure(s): 1
Drive(s): 26
Logical Enclosure(s): 1
Energy Pack: Yes

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 17: Controller Dashboard

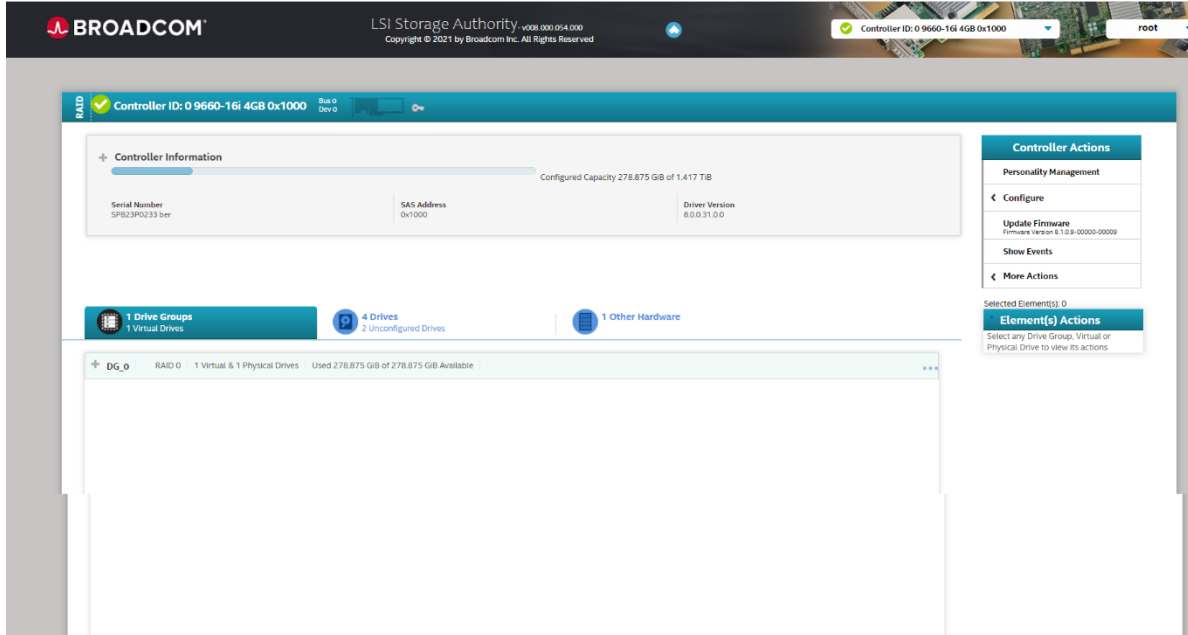




Table 7: Controller Dashboard Description

Callout	Description
1	<p>Controller Summary – Displays the name of the controller. 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.</p> <p>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.</p>
2	<p>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 CacheVault™ module associated with the controller. All these views are displayed as tabs.</p> <p>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.</p>

Callout	Description
3	<p>Controller Actions – Lets you perform some of the following actions:</p> <ul style="list-style-type: none">• Create a configuration.• Clear a configuration.• Manage the Profile.• Manage link speed.• Update the controller firmware.• Enable or disable SSD Guard.• View Premium features.• Show events.• Dump a snapshot when a fault occurs.• Enable or disable drive security.• Import or clear foreign configurations.

Configuration

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

You can create RAID 0, RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, RAID 60, 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 Features](#).

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.

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 18: 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, and 6). 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 spare. Hot spares are drives that are available to replace failed drives automatically in a redundant virtual drive (RAID 1, RAID 5, or RAID 6).
7. Click **Finish**.
A message appears stating that the configuration is successfully created.

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 19: Advanced 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.
Also see [RAID Levels](#) for more information.
When you use advanced configuration, the RAID controller supports RAID levels 0, 1, 5, 6, 10, 50, and 60). 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 **Secure** 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. Click **Next**.
6. Click **Add Physical Drives** to add physical drives to the drive group.
The **Available Unconfigured Drive** window appears.

Figure 20: Available Unconfigured Drive(s) Window

<input type="checkbox"/>	▲ Enclosure:Slot	Device/Persistent ID	Media	Interface	Capacity	Logical Sector Size	Model	NS/LU Count
<input type="checkbox"/>	EN_262 : 0	257	SSD	NVMe	893.75GiB	512B	SAMSUNG MZQLW960HJMJP-00003	1
<input type="checkbox"/>	EN_262 : 1	258	SSD	NVMe	893.75GiB	512B	SAMSUNG MZQLW960HJMJP-00003	1
<input type="checkbox"/>	EN_368 : 2	283	HDD	SAS	278.937GiB	4KiB	HUC156030CS4204	1
<input type="checkbox"/>	EN_368 : 3	284	HDD	SAS	558.406GiB	4KiB	HUC101860CS4200	1
<input type="checkbox"/>	EN_368 : 4	285	HDD	SAS	558.406GiB	4KiB	HUC101860CS4200	1
<input type="checkbox"/>	EN_368 : 5	286	HDD	SAS	278.937GiB	4KiB	HUC156030CS4204	1
<input type="checkbox"/>	EN_368 : 6	287	HDD	SAS	278.875GiB	512B	ST300MM0006	1
<input type="checkbox"/>	EN_368 : 7	288	HDD	SAS	278.875GiB	512B	ST300MM0006	1
<input type="checkbox"/>	EN_368 : 8	289	HDD	SAS	278.937GiB	512B	HUC156030CS5200	1
<input type="checkbox"/>	EN_368 : 9	290	HDD	SAS	278.937GiB	512B	HUC156030CS5200	1

Add Physical Drives 0 drives selected.

Refer the following sections for more information:

- For information on how to add physical drives, see [Adding Physical Drives](#).
- For information on adding hot spare drives to the drive group, see [Adding Hot Spares to the Existing Drive Group](#).

7. Select the span depth using the slider bar.
8. Click **Add Virtual Drives** to add virtual drives to the drive group.
The **Virtual Drive Settings** window appears.

Figure 21: Virtual Drive Settings Window

Virtual Drive Settings ? ×

558.406 GiB available across 1 selected drive
16 Virtual Drives can be added for a drive group

How many virtual drives do you wish to create?

↑ ↓ each with capacity of ↑ ↓ ▼

Virtual Drive Name Strip Size ▼

Initialization State
No Initialization

Read Policy
No Read Ahead

Write Policy
Write Back

Drive Write Cache Policy
Disabled

Initialization prepares the storage medium for use

- No Initialization**
The new configuration is not initialized, and the existing data on the drives is not overwritten.
- Fast Initialization**
The Firmware erases the first and last 8 MB/MiB of the data area of the virtual drive by writing 0x00 to wipe out any remains of Master boot record (MBR) or partition tables. This operation is extremely fast, so the virtual drive is almost instantly accessible to the user.
- Full Initialization**
A complete initialization is done on the new configuration. You cannot write data to the new virtual drive until the initialization is complete. This process can take a long time if the drives are large.

Add Virtual Drives

For information on configuring virtual drives, see [Adding Virtual Drives](#).

9. Click **Add Virtual Drives**.

10. Click **Finish**.

A confirmation message appears if the storage configuration is successfully completed.

Adding Physical Drives

The **Available Unconfigured Drive** window lists all the available unconfigured physical drives. You can either select each physical drive individually or select all the physical drives based on their respective drive type.

Perform the following steps to add physical drives to the drive group.

1. On the **Available Unconfigured Drives** window, click on the corresponding check box of a particular physical drive to add it to the drive group.

The selected physical drives appear in the **Advanced Configuration** window.

You can click the **X** icon to remove the physical drives that you have already added.

For adding physical drives to the drive group based on their drive type:

2. Click the **ALL** drop-down list.

The physical drives are categorized by the system based on their drive type. For example, *4K drives*, *SAS/SATA drives*, *512K drives*, and so on).

Figure 22: Available Unconfigured Drives – Drive Type Window

24 Available Unconfigured Drive(s)

Add a minimum of 1 drive as required by RAID 0 Level.

	Enclosure/Slot	Device/Persistent ID	Media	Interface	Capacity	Logical Sector Size	Model	NSID	Type
<input type="checkbox"/>	EN_262 : 0	257	SSD	NVMe	893.75GiB	512B	SAMSUNG MZQLW960HMJP-00003	1	HDD-4KiB
<input type="checkbox"/>	EN_262 : 1	258	SSD	NVMe	893.75GiB	512B	SAMSUNG MZQLW960HMJP-00003	1	HDD-512B
<input type="checkbox"/>	EN_368 : 2	283	HDD	SAS	278.937GiB	4KiB	HUC156030CS4204	1	SSD-512B
<input type="checkbox"/>	EN_368 : 3	284	HDD	SAS	558.406GiB	4KiB	HUC101860CS4200	1	NVME-512B
<input type="checkbox"/>	EN_368 : 4	285	HDD	SAS	558.406GiB	4KiB	HUC101860CS4200	1	ALL
<input type="checkbox"/>	EN_368 : 5	286	HDD	SAS	278.937GiB	4KiB	HUC156030CS4204	1	
<input type="checkbox"/>	EN_368 : 6	287	HDD	SAS	278.875GiB	512B	ST300MM0006	1	
<input type="checkbox"/>	EN_368 : 7	288	HDD	SAS	278.875GiB	512B	ST300MM0006	1	
<input type="checkbox"/>	EN_368 : 8	289	HDD	SAS	278.937GiB	512B	HUC156030CS5200	1	
<input type="checkbox"/>	EN_368 : 9	290	HDD	SAS	278.937GiB	512B	HUC156030CS5200	1	

Add Physical Drives 0 drives selected.

- a) From the available drop-down list, select the required drive type.
- b) Click on the check box that is located on the very first row to select the all the physical drives based on the drive type.

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

Adding Hot Spares to the Existing Drive Group

To add hot spares to the drive group, you should have already added the required physical drives to the drive group.

Perform the following steps to add hot spares to the drive group.

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

The **Advanced Configuration** window opens.

2. Click **Add Hot Spares** to add dedicated hot spare drives to the drive group.

The **Available Unconfigured Drives** window appears which lists the drives that can be added as hot spares.

You can either select each hot spare individually or select all the hot spares based on type of physical drive that you have selected to add to the drive group.

3. Perform one of the following actions:
 - a) Click on the corresponding check box of a hot spare drive to add it to the drive group.
 - b) Click on the check box that is located on the very first row to select the all the hot spare drives to add them to the drive group.

Figure 23: Available Unconfigured Drives – Hot Spares

The screenshot shows a window titled "15 Available Unconfigured Drive(s)" with a close button and a help icon. Below the title bar, there is a note: "Add a minimum of 1 drive as required by RAID 0 Level." To the right of this note are two dropdown menus: "Type" and "ALL". Below these is a table with the following columns: Enclosure:Slot, Device ID, Media, Interface, Capacity, Logical Sector Size, Model, and NS/LU Count. The table contains 15 rows of drive information. At the bottom of the window, there is a button labeled "Add Physical Drives" and a status message that says "0 drives selected."

<input type="checkbox"/>	▲ Enclosure:Slot	Device ID	Media	Interface	Capacity	Logical Sector Size	Model	NS/LU Count
<input type="checkbox"/>	EN_318 : 0	275	HDD	SAS	278.937GiB	4KiB	HUC156030CS4204	1
<input type="checkbox"/>	EN_318 : 5	280	HDD	SAS	465.25GiB	512B	ST9500431SS	1
<input type="checkbox"/>	EN_318 : 9	284	HDD	SAS	278.937GiB	4KiB	HUC156030CS4204	1
<input type="checkbox"/>	EN_318 : 10	285	HDD	SAS	278.937GiB	4KiB	HUC156030CS4204	1
<input type="checkbox"/>	EN_318 : 11	286	HDD	SAS	278.937GiB	4KiB	HUC156030CS4204	1
<input type="checkbox"/>	EN_318 : 12	287	HDD	SAS	558.406GiB	512B	ST600MM0099	1
<input type="checkbox"/>	EN_318 : 14	289	HDD	SAS	558.406GiB	512B	ST600MM0099	1
<input type="checkbox"/>	EN_318 : 15	290	HDD	SAS	558.406GiB	512B	ST600MM0099	1
<input type="checkbox"/>	EN_318 : 17	292	HDD	SAS	558.406GiB	512B	ST600MM0099	1
<input type="checkbox"/>	EN_318 : 18	293	HDD	SAS	558.406GiB	512B	ST600MM0099	1

4. Click **Add Hot Spares** to add dedicated hot spare drives to the drive group.

Adding Virtual Drives

The **Virtual Drive Settings** window enables you to configure and add virtual drives to the drive group. 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 Features](#).

Perform the following steps to configure and add virtual drives:

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 KiB** and **256 KiB** are supported.
5. Specify the initialization policy for the virtual drive. The options follow:
 - **No Initialization**
 - **Fast Initialization**
 - **Full Initialization**

6. Specify the read policy for the virtual drive as **No 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 a drive cache setting for the virtual drive. The options follow:
 - **Default**
 - **Enabled**
 - **Disabled**
9. 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 capacities 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**.

10. Click **Finish** to successfully add the virtual drive to the drive group.

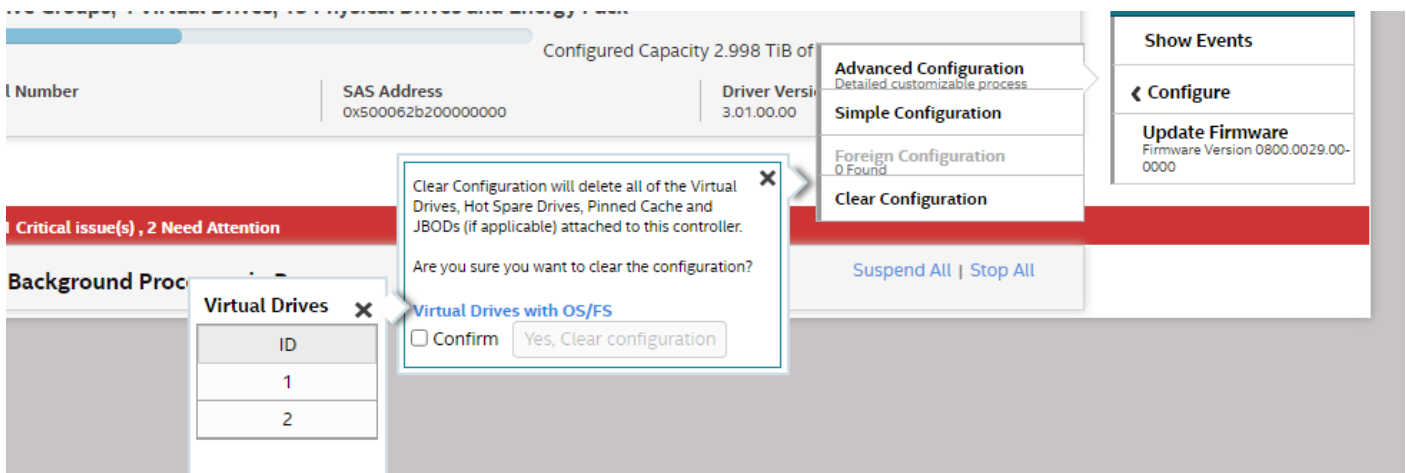
Clearing the Configuration

You can clear all existing configurations on a selected 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.

Figure 24: Clear Configuration Dialog



3. Select **Confirm** and click **Yes, Clear configuration** to clear all the existing configurations on the controller.

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](#).

Personality Management

The Personality Management feature allows you to switch between different personality and behavior modes that are supported by your firmware. Using the Personality Management feature, you can switch between the following personality and behavior modes.

- Personality modes:
 - RAID
 - HBA
- Behavior modes
 - JBOD
 - Secure JBOD
 - Unconfigured Good
 - Single Drive RAID0
 - Single Drive RAID 0 WB
 - Secure Single Drive RAID 0
 - Secure Single Drive RAID 0 WB

For more information on behavior modes and parameters, see [Changing Behavior Modes](#).

Figure 25: Personality Properties

Personality Properties		
Controller Personality RAID	Controller Primary Auto-configure behavior Unconfigured Good	Controller Secondary Auto-configure behavior Unconfigured Good

Changing Personality Modes

Perform the following steps to switch between different personality modes. Currently, you can switch between RAID and JBOD personality modes.

1. On the Controller dashboard, select **Actions > Personality Management**.

The **Change Personality** page appears.

Figure 26: Change Personality

2. Select the required personality mode from the **Select personality** drop-down list, and then click **Change**.
3. Reboot the system for the changes to take effect.

Changing Behavior Modes

Perform the following steps to change the behavior mode and parameters:

1. In the **Change Personality** page, select the **Change Auto-configure behavior** radio button.
2. From the **Change Primary Auto-configure behavior** drop-down list, select an appropriate behavior mode.

The available behavior modes are based on the current firmware support.

- **Unconfigured Good** – If a user selects this option, then all the configured drives are converted to Unconfigured Good Drive.
- **JBOD** – If all the physical drives are in unconfigured good state, once you select the JBOD mode, all the physical drives are automatically converted to JBOD physical drives.
- **Single Drive RAID 0** – If a user selects this option, then all the physical drives are automatically converted to Single Drive RAID 0.
- **Single Drive RAID 0 WB** – If a user selects this option, then all the physical drives are automatically converted to Single Drive RAID 0 With Write Back.
- **Secure JBOD** – If all the physical drives are in unconfigured good state, once you select the JBOD mode, all the physical drives are automatically converted to secure JBOD physical drives.
- **Secure Single Drive RAID 0** – This option converts all the UGOOD physical drives to Secure Single Drive Raid0 on secure physical drives and Single Drive Raid0 on normal physical drives.
- **Secure Single Drive RAID 0 WB** – This option auto-secures the SED drive to Single Drive RAID 0 Write Back.

The **Change Primary Auto-configure behavior** and **Change Secondary Auto-configure behavior** drop-down list selection is disabled if a change is not allowed.

- (Optional) To immediately change the behavior mode, select the **Configure Now** radio button from the **Change Personality** page.

Figure 27: Configure Now Dialog

The available behavior modes are based on the current firmware support.

- **Unconfigured Good** – If a user selects this option, then all the configured drives are converted to Unconfigured Good Drive.
- **JBOD** – If all the physical drives are in unconfigured good state, once you select the JBOD mode, all the physical drives are automatically converted to JBOD physical drives.
- **Single Drive RAID 0** – If a user selects this option, then all the physical drives are automatically converted to Single Drive RAID 0.
- **Single Drive RAID 0 WB** – If a user selects this option, then all the physical drives are automatically converted to Single Drive RAID 0 With Write Back.
- **Secure JBOD** – If all the physical drives are in unconfigured good state, once you select the JBOD mode, all the physical drives are automatically converted to secure JBOD physical drives.
- **Secure Single Drive RAID 0** – This option converts all the UGOOD physical drives to Secure Single Drive Raid0 on secure physical drives and Single Drive Raid0 on normal physical drives.
- **Secure Single Drive RAID 0 WB** – This option auto-secures the SED drive to Single Drive RAID 0 Write Back.

- Select the appropriate behavior mode and click **Change**.

Background Operations Support

The LSI Storage Authority software provides background Suspend, Resume, Stop, Suspend All, Resume All, and Stop All features that enhance the functionality where 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 a **Stop** 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 **Suspend**, **Resume**, and **Stop** operations, go to the **Background Processes in Progress** window in the Server dashboard or the Controller dashboard, and perform the steps that follow. The **Background Processes in Progress** window is shown in the figure that follows.

Figure 28: Background Processes in Progress Window

The screenshot displays the 'Background Processes in Progress' window. At the top, it shows '1 Background Processes in Progress' with a progress bar for a 'Rebuild' operation on a physical drive. Summary statistics include: 1 Drive Groups (1 Virtual Drive), 19 Drives (13 Unconfigured), and 2 Other Hardware (including Energy Pack). Below this is a table for 'DG_0' RAID 1, showing 1 Virtual Drive and 2 Physical Drives. The table lists two physical drives with their respective IDs, media types, interfaces, capacities, and models. A 'More Actions' menu on the right lists options like 'Make Offline', 'Start Locate', 'Stop Locate', 'Make Failed', 'Resume Rebuild', and 'Stop Rebuild'.

Enclosure:Slot	Device/Persistent ID	Media	Interface	Capacity	Logical Sector Size	Model	NS/LU Count
EN_318:15	290	HDD	SAS	558.406GiB	512B	ST600MM0099	1
EN_318:16	291	HDD	SAS	558.406GiB	512B	ST600MM0099	1

- **Suspend** – Click **Suspend** to suspend the background operation taking place at that particular point of time. When the operation is paused, the **Resume** option appears instead of the **Suspend** option.
- **Resume** – Click **Resume** to resume the operation from the point of its last suspension.
- **top** – Click **top** to stop the ongoing active operation.
- **Suspend All** – Click **Suspend All** to suspend all 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.
- **Stop All** – Click **Stop All** to stop all the active operations.


While refreshing the LSA Client or the LSA GUI, if there is any delay such as the **Replace** progress bar not displaying the progress of the **Replace** operation or the **Replace** progress bar itself not getting displayed for small-size volumes, set the maximum event grouping time gap to 0 in the `LSA.conf` file to see the real-time events. For more information on setting up LSA retrieve real-time events, see [Configuring LSA on Gateway Servers to Get Real-Time Events](#).

Note that setting real-time events has an impact on the performance of LSA.

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.

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 29: Basic and Advanced Controller Properties

Controller Information			
Configured Capacity 0 KiB of 10,199 TiB			
Serial Number N/A	SAS Address 0x500062b200000000	Driver Version 8.00.01.00	PCI Address 00:02:00:0
PCI_Vendor_ID 0x1000	PCI Subsystem Vendor ID 0x1000	PCI Device ID 0x00a5	PCI Subsystem ID 0x4600
Chip Temperature 60 C (140 F)			
Advanced Properties			
NVRAM Present Yes	Shield State Supported Yes	Energy Pack Yes	NVRAM Size 128 KiB
SSD Guard on SMART Error Enabled			
Power State Properties			
Power savings on unconfigured drives Enabled	Power saving on hot spares Enabled	Spin Down Time 30 mins	
Firmware Properties			
Package Version MANIFEST PKG VER: 8.0.30.00028-0000-00000	Firmware Version FW Version:8.030.28-00000	All Online Controller Reset Enabled	
Security Properties			
Security Capable Yes	Security is enabled. No		
Emergency Spare Properties			
Emergency Spare Unconfigured Good & Global Hot Spare	Emergency spare for SMARTer Enabled		
Personality Properties			
Controller Personality RAID	Controller Primary Auto-configure behavior Unconfigured Good	Controller Secondary Auto-configure behavior Unconfigured Good	
Snapdump Properties			
Save Count through reset 4			

Table 8: Basic and Advanced Controller Properties

Property	Description
Basic Controller Properties	
Serial Number	The serial number of the controller.
SAS Address	The SAS address of the controller.

Property	Description
Driver Version	The driver version of the controller.
PCI Address	A unique address assigned to the controller.
PCI_Vendor_ID	A unique controller ID assigned to a specific vendor.
PCI Subsystem Vendor ID	Additional vendor ID information about the controller.
PCI Device ID	The device ID that is assigned by the manufacturer.
PCI Subsystem ID	Additional device ID that is assigned by the manufacturer.
Chip Temperature	Indicates the temperature of the controller.
Advanced Properties	
NVRAM Present	Indicates if a nonvolatile random access memory (NVRAM) is present on the controller.
Shield State Supported	Indicates whether the controller supports the shield state.
Energy Pack	Indicates whether the controller supports the associated energy pack status management.
NVRAM Size	Indicates the capacity of the controller's NVRAM.
SSD Guard on SMART Error	Indicates if the SSD Guard feature is enabled on the controller.
Power State Properties	
Power savings on unconfigured drives	Indicates if the power savings on the unconfigured drives is enabled.
Power saving on hot spares	Indicates if the power savings on the hot spares is enabled.
Spin Down Time	Shows the drive spin down time in minutes.
Firmware Properties	
Package Version	The firmware package version of the controller.
Firmware Version	The firmware version of the controller.
All Online Firmware Update	Indicates if the online firmware update feature is enabled in the firmware.
Security Properties	
Security Capable	Indicates the security (encryption) feature status on the controller.
Security is enabled	Indicates whether the security is enabled.
Emergency Spare Properties	
Emergency Spare	Indicates the emergency spare controller properties. It can be set to Unconfigured Good or Unconfigured Good and Global Hotspare .
Emergency spare for SMARTer	Indicates if emergency hot spare drives are commissioned for predictive analysis.
Personality Properties	
Controller Personality	Shows the current personality mode.
Controller Primary Auto-configure behavior	Shows the primary auto-configure behavior.
Controller Secondary Auto-configure behavior	Shows the secondary auto-configure behavior.
Snap Dump Properties	
Save Count through reset	

Running the Consistency Check

The consistency check operation verifies the correctness of the data in virtual drives that use RAID levels 1, 5, 6, 10, 50, and 60 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 these RAID 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.

Setting the 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**.

Scheduling a 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.

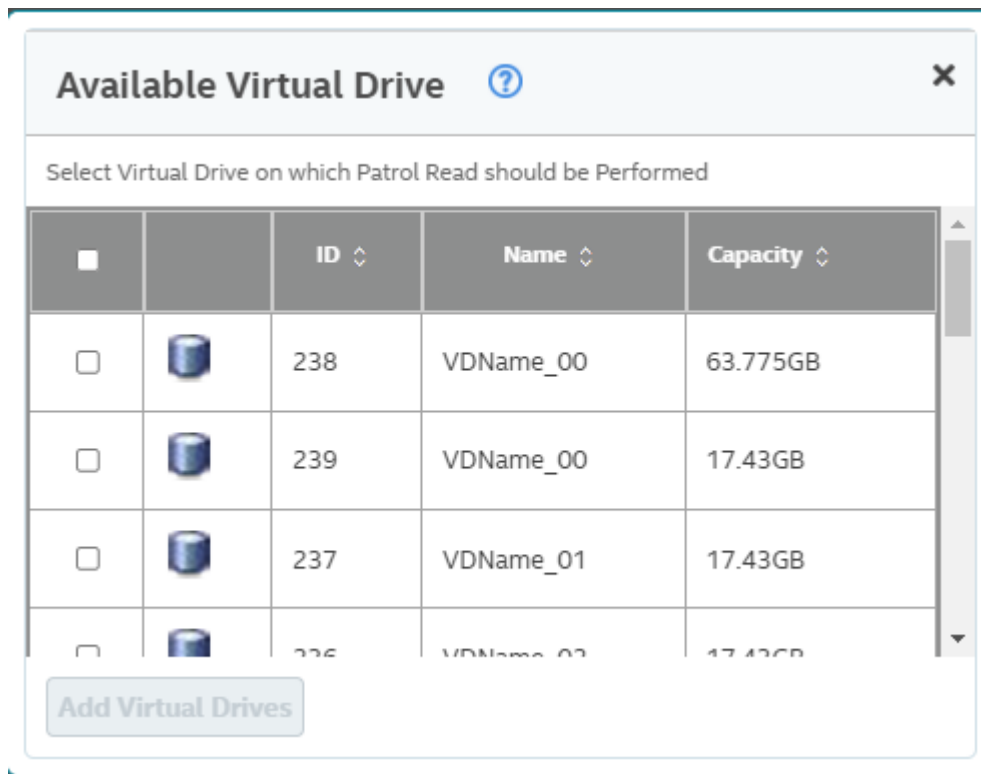
Figure 30: Schedule Consistency Check Dialog

2. In the **Set Mode** dialog, select an available mode. The available options are:
 - **Enable** – Run a consistency check concurrently on all virtual drives or one drive at a time.
 - **Disable** – Disables the consistency check.
3. Set the maximum number of virtual drives.
4. Set the desired interval at which you want to run the consistency check.
 - a) Select an appropriate date and time range.
5. Click **Next**.

The **Schedule Consistency Check** page appears, which allows you to add virtual drives on which you want to perform a consistency check.

6. Click **Add Virtual Drives**.

The **Available Virtual Drive** dialog appears, which lists all the virtual drives present in the selected drive group.

Figure 31: Available Virtual Drives Dialog

7. You can either select each virtual drive individually by selecting the corresponding check box of a particular virtual drive or select all the virtual drives by selecting the check box that is located on the very first row.
8. Click **Add Virtual Drives**.
The consistency check runs based on the frequency or interval chosen by you. You can also monitor the progress of the consistency check operation. See [Background Operations Support](#).
9. Click **Save**.
10. (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, select **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, which asks for your confirmation.

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.

Scheduling a Patrol Read

Perform the following steps to schedule a patrol read.

1. In the Controller dashboard, select **More Actions > Schedule Patrol Read**.
The **Schedule Patrol Read** dialog appears.

2. Perform the following steps to set the properties:
 - a) Select an operation mode for patrol read from the **Set Mode** drop-down list. The options follow:
 - **Enable** – The patrol read operation runs automatically at the time interval you specify or when you manually start it by selecting **Start Patrol Read** from the Controller dashboard.
 - **Disable** – The patrol read operation does not run.
 - b) (Optional) Specify a maximum number of physical drives to include in the patrol read concurrently. The count must be a number from 1 to the maximum number of configurable physical devices.
 - c) Select the frequency at which the patrol read operation runs from the drop-down list.
 - 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.
3. Click **Next**.
4. Select the virtual drives for which you want to set the patrol read properties.
5. Click **Finish**.

You can monitor the progress of the Patrol Read operation. See [Background Operations Support](#).

Starting a Patrol Read Operation

Perform the following steps to start a patrol read operation.

1. In the Controller dashboard, select **Element(s) Actions > Start Patrol Read**.

A warning message appears.
2. Click **Start Patrol Read** to start a patrol read operation.

You can monitor the progress of the patrol read operation. See [Background Operations Support](#).

Stopping a Patrol Read Operation

Perform the following step to stop a patrol read operation.

In the Controller dashboard, select **Element(s) Actions > Stop Patrol Read**.

Managing SAS Storage Link Speed

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

Figure 32: Manage SAS Storage Link Speed DialogManage SAS Storage Link Speed [?](#)

SAS/SATA Phy Information

Phy	Status	Receptacle Name	Port Number	Select Link Speed	Negotiated Speed	Connected Device
0	DISABLED	N/A	N/A		N/A	
1	DISABLED	N/A	N/A		N/A	
2	DISABLED	N/A	N/A		N/A	
3	DISABLED	N/A	N/A		N/A	
4	DISABLED	N/A	N/A		N/A	
5	DISABLED	N/A	N/A		N/A	
6	DISABLED	N/A	N/A		N/A	
7	DISABLED	N/A	N/A		N/A	
8	DISABLED	N/A	N/A		N/A	
9	DISABLED	N/A	N/A		N/A	
10	DISABLED	N/A	N/A		N/A	
11	DISABLED	N/A	N/A		N/A	
12	DISABLED	N/A	N/A		N/A	
13	DISABLED	N/A	N/A		N/A	
14	DISABLED	N/A	N/A		N/A	
15	DISABLED	N/A	N/A		N/A	
16	OPTIMAL	C0.1	16	6.0Gb/s	N/A	
17	OPTIMAL	C0.1	16	6.0Gb/s	N/A	
18	OPTIMAL	C0.1	16	6.0Gb/s	N/A	
19	OPTIMAL	C0.1	16	6.0Gb/s	N/A	
20	OPTIMAL	C0.0	16	6.0Gb/s	N/A	

- The **Phy** column displays the system-supported phy link values.
 - 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.
 - The **Negotiated Speed** column displays the current negotiated speed of the phy.
 - The **Connected Device** column displays the device connected to the phy.
2. Select the desired link speed from the **Select Link Speed** field using the drop-down selector. The link speed values are **MAX**, **6.0Gb/s**, **12.0Gb/s**, or **24.0Gb/s**.

By default, the link speed in the controller is **MAX** or the value last saved by you. The **12.0Gb/s** 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.

Managing PCIe Storage Lane Speed

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 Lane Speed 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 Lane Speed**.

The **Manage PCIe Storage Lane Speed Dialog** appears.

Figure 33: Manage PCIe Storage Lane Speed Dialog

Manage PCIe Storage Lane Speed ⓘ

View PCIe Information

lane	Status	Receptacle Name	Port Number	Current Clock Mode	Current Speed	Connected Device
16	OPTIMAL	C0.1	16	0	N/A	
17	OPTIMAL	C0.1	16	0	N/A	
18	OPTIMAL	C0.1	16	0	N/A	
19	OPTIMAL	C0.1	16	0	N/A	
20	OPTIMAL	C0.0	20	0	N/A	
21	OPTIMAL	C0.0	20	0	N/A	
22	OPTIMAL	C0.0	20	0	N/A	
23	OPTIMAL	C0.0	20	0	N/A	
24	OPTIMAL	C1.1	24	0	N/A	
25	OPTIMAL	C1.1	24	0	N/A	
26	OPTIMAL	C1.1	24	0	N/A	
27	OPTIMAL	C1.1	24	0	N/A	
28	OPTIMAL	C1.0	28	0	N/A	
29	OPTIMAL	C1.0	28	0	N/A	
30	OPTIMAL	C1.0	28	0	N/A	
31	OPTIMAL	C1.0	28	0	N/A	

Change Link Speed

Port Number	Lane Speed
16	16GT/s ▼
20	16GT/s ▼
24	16GT/s ▼
28	16GT/s ▼

System restart will be required after saving the changes

- The **Lane** column displays the system-supported lane values.
- The **Status** column displays the status of the lane.
- The **Port Number** column displays the port 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**, **8GT/s**, and **16GT/s**.

By default, the lane speed in the controller is **8GT/s** 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.

Setting the 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 34: Set Adjustable Task Rate Dialog

Task	Priority Percentage
Rebuild Rate	30
Patrol Rate	30
BGI Rate	30
Consistency Check Rate	30
OCE Rate	30

Save

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 the 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 are performed.
The 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).
- **BGI Rate (Background Initialization 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 performed.

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).

- **OCE Rate** – Enter a number from 0 to 100 to control the rate at which online capacity expansion occurs.

3. Click **Save** to set the new task rates.

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, which helps to share resources more efficiently and lowers the cost.

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

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 35: 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, and 2 hours through 24 hours.

5. Click **Finish** to save the settings.

A confirmation message appears.

Enabling and Disabling SSD Guard

Solid state drives (SSDs) are known for their reliability and performance. The SSD Guard technology, which 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.

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**.

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 is displayed 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**.

Updating the Controller Firmware

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

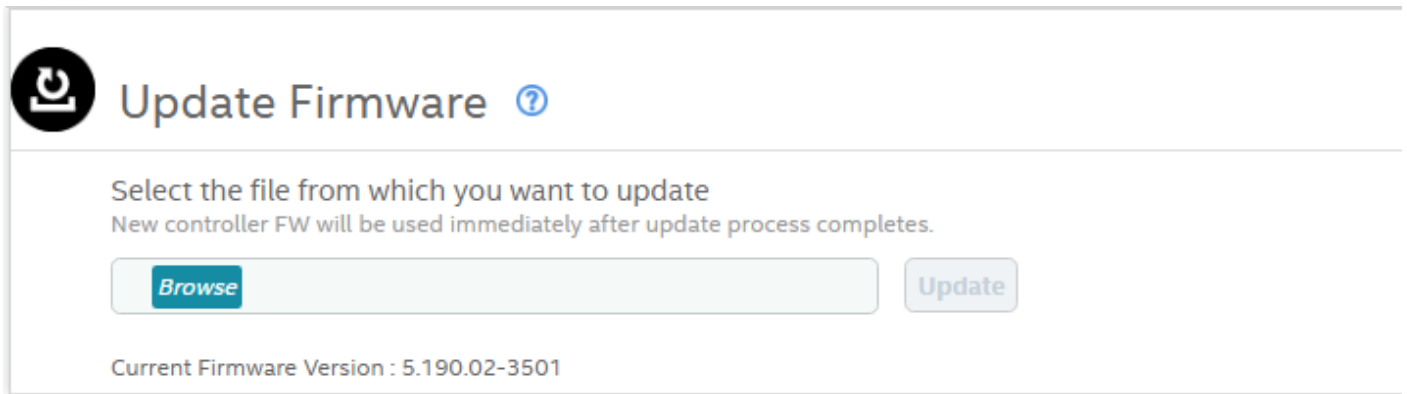
NOTE

You can flash both the BIOS and EFI using any of the BIOS or EFI radio buttons.

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` file.
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 36: Update Firmware

Update Firmware ?

Select the file from which you want to update
New controller FW will be used immediately after update process completes.

[Browse](#) [Update](#)

Current Firmware Version : 5.190.02-3501

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

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

Advanced Software

The advanced software (Premium) are features that the LSI Storage Authority software supports on certain RAID controllers.

The advanced software includes the following features:

- Fast Path
- SafeStore

The software licensing authorizes you to enable the advanced software features. You must obtain the activation key, which will enable you to use the advanced software features present in the controller.

Activating 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 > Advanced Software Options**.

The **Advanced Software Options** window opens.

Figure 37: Advanced Software Options Window

The screenshot shows the 'Advanced Software Options' window. At the top, there are fields for SERIAL NUMBER (SP82418319), SAS ADDRESS (0x500062b203f74f80), and FIRMWARE VERSION (7.705.08.00). Below this, there is a section for '2 Background Processes in Progress' with a 'Suspend Patrol' button. The main content area is titled 'Advanced Software Options' and contains the following text:

Advanced Software Options enable special functionality or features that may not be available in the standard configuration of the controller.

[Benefits of each Advanced Software Option](#)

Activated Advanced Software Options:

Software Options	Time Remaining	Location
SafeStore	Unlimited	ELECTRONIC-KEY
FastPath	Unlimited	FACTORY

You will be prompted to provide the below Safe ID and Serial Number on the LSI Advanced Software License Management Portal in order to obtain the Activation Key that enables the Advanced Software Options.

[LSI Advanced Software License Management Portal](#)

Safe ID: FAGVMVW48N66SLUQRP2T4H4A1PE95ELD762FLGRZ

Serial Number: SP82418319

[Activate](#)

[Tips on activating Advanced Software Options](#)

On the right side of the window, there is a 'More Actions' menu with the following options: Schedule Consistency Check, Set Consistency Check Properties, Schedule Patrol Read, Stop Patrol Read, Manage SAS Storage Link Speed, Advanced Software Options, Set Adjustable Task Rate, Disable SSD Guard, Enable SCSI Unmap, Download TTY Log, Snapdump Actions, Disable Security, Change Security, and Manage Power Save Setting.

The Activated Advanced Software Options table consists of the Software Options, Time Remaining, and Location columns.

- The Software Options column displays the list of advanced software options present in the controller.
- The Time Remaining column displays the time remaining for the software option to deactivate.
- The Location column displays the key location from which the option is activated.

(Optional) – For more information on the benefits of these features, click the **Benefits of each Advanced Software Option** link.

- 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 **KeyId** field consist of a predefined value generated by the controller.

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

- Click **Activate**.
- Enter the activation key in the text box provided.
- Click **Next**.

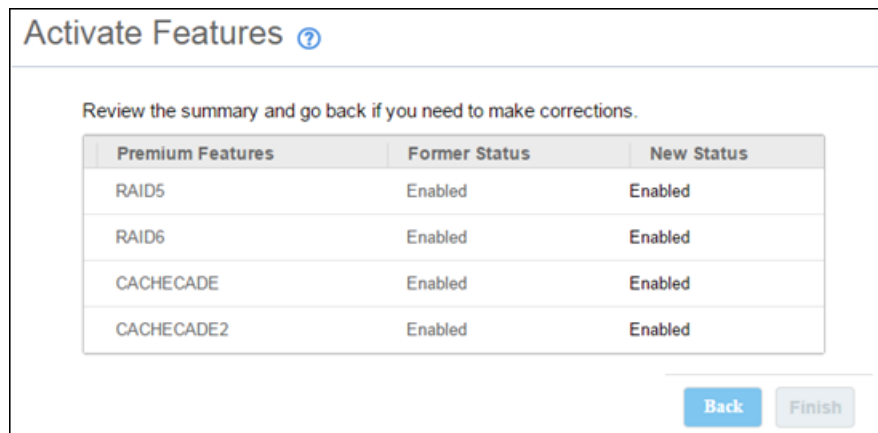
After you click **Next**, one of the following two scenarios occurs:

- Depending on whether you are activating an unlimited key or a trial key, the relevant **Activate Features – Summary** dialog appears. See [Advanced 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](#).

Advanced 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 38: Activate Features Window – Summary



The screenshot shows a window titled "Activate Features" with a help icon. Below the title is a message: "Review the summary and go back if you need to make corrections." Below this is a table with three columns: "Premium Features", "Former Status", and "New Status". The table contains four rows of data. At the bottom right of the window are two buttons: "Back" and "Finish".

Premium Features	Former Status	New Status
RAID5	Enabled	Enabled
RAID6	Enabled	Enabled
CACHECADE	Enabled	Enabled
CACHECADE2	Enabled	Enabled

- 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.

To activate the premium features, perform the following steps.

- Click **Finish**.

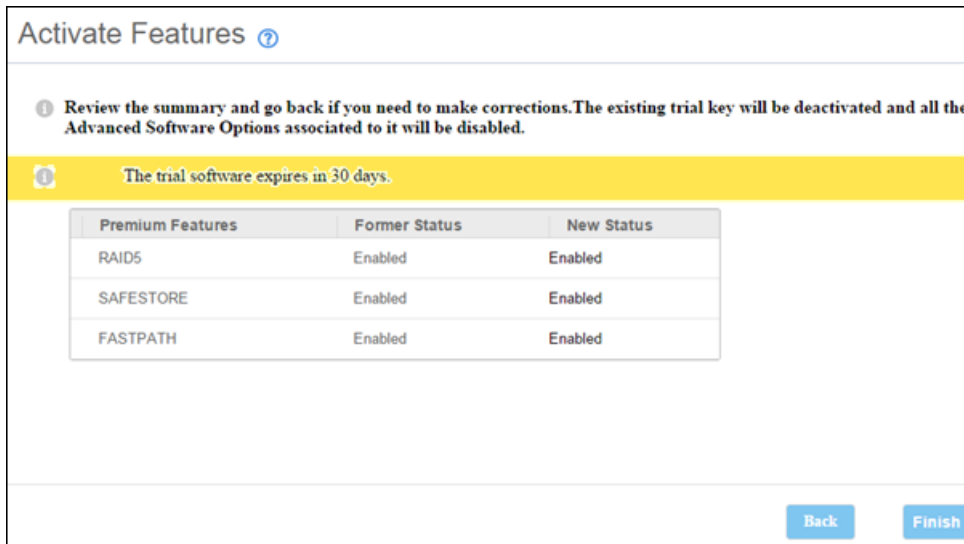
The status of the advanced software is enabled, and the advanced features are secured in the key vault.

- Click **Back** to return to the previous window to change any selections.

Activating a Trial Key

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

Figure 39: Activating Trial Software

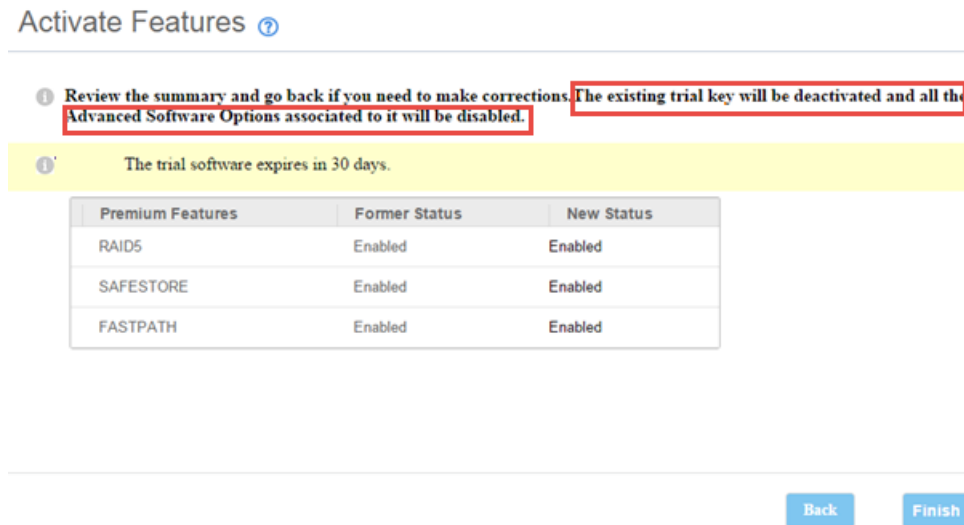


Activating an Unlimited Key over a Trial Key

When you activate an unlimited key over a trial key, the following message appears:


The existing trial key will be deactivated and all the Advanced Software Options associated to it will be disabled.

Figure 40: Activating an Unlimited Key over a Trial Key



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 41: Reusing the Activation Key

Application Scenarios and Messages

Scenario 1

If you enter an *invalid* activation key, an error message is displayed.

Scenario 2

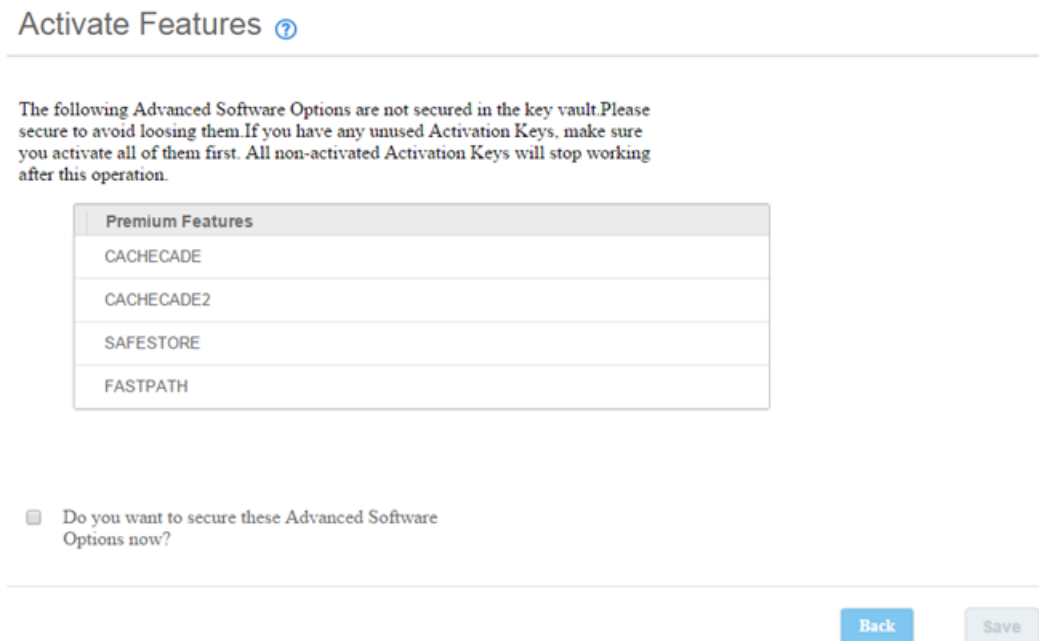
If you enter an *incorrect* activation key, and if a mismatch exists between the activation key and the controller, an error message is displayed.

Securing Advanced 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 software.

1. In the **Advanced Software Options** window, click **Configure Key Vault**.
The **Activate Features** window opens.

Figure 42: Activate Features Window – 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.

Configuring the 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 43: Premium Features Window – Configure Key Vault

Premium Features

To transfer Advanced Software Options from one controller to another controller you need to complete the re-hosting process. Only then you will be able to secure the Advanced Software Options in the key vault. This wizard helps you to configure the key vault by transferring the Advanced Software Options from one controller to another controller and securing them in the key vault. Please furnish the below details in the LSI Advanced Software License Management Portal in order to complete the re-hosting process. If you have already completed the process then select the checkbox below and proceed with next.

LSI Advanced Software License Management Portal

Former Serial Number:

New serial number: SR91700046

Site ID: 8EF20X1GGTA118BAVEFMU4DXXUP1TLEJ03BLTSRZ

I acknowledge that I have completed the re-hosting process in the external site.

Back Next

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 selected the check box.

The **Activate Features** window appears.

Figure 44: Activate Features Window – Configure Key Vault Window

Activate Features

The following Advanced Software Options will be secured as part of the re-hosting process. If you have any unused Activation Keys, make sure you activate all of them first. All non-activated Activation Keys will stop working after this operation.

Premium Features
CACHECADE
CACHECADE2
SAFESTORE
FASTPATH

Back Finish

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

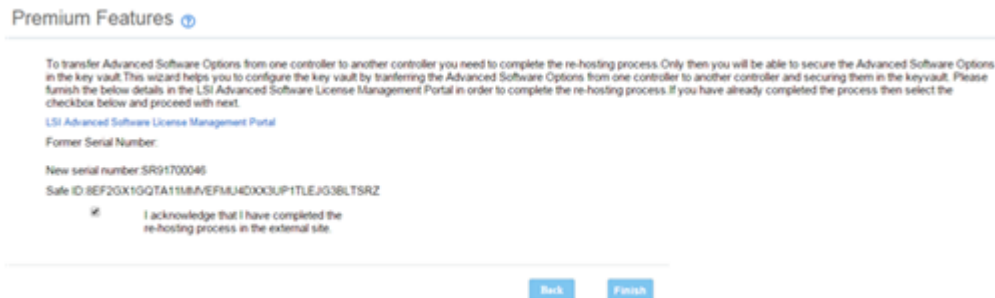
Implementing the Re-Hosting Process

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 45: 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. Either click **Finish**, and the advanced software options are secured in the key vault, or click **Cancel** if you do not want to activate the re-hosting process

Snapdump Feature


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

This section describes the LSA design `get` and `set` Snapdump properties and how to download the respective dump files from the firmware.

Snapdump Feature Support

Using the Snapdump feature support, the user can enable or disable the Snapdump feature. If the Snapdump feature is enabled, the user can set the `CacheOffloadLimit` properties of the Snapdump. If the user chooses to disable the Snapdump feature, the `CacheoffloadLimit` modification is not applicable and is disabled at the client level.

Figure 46: Snapdump Properties Menu



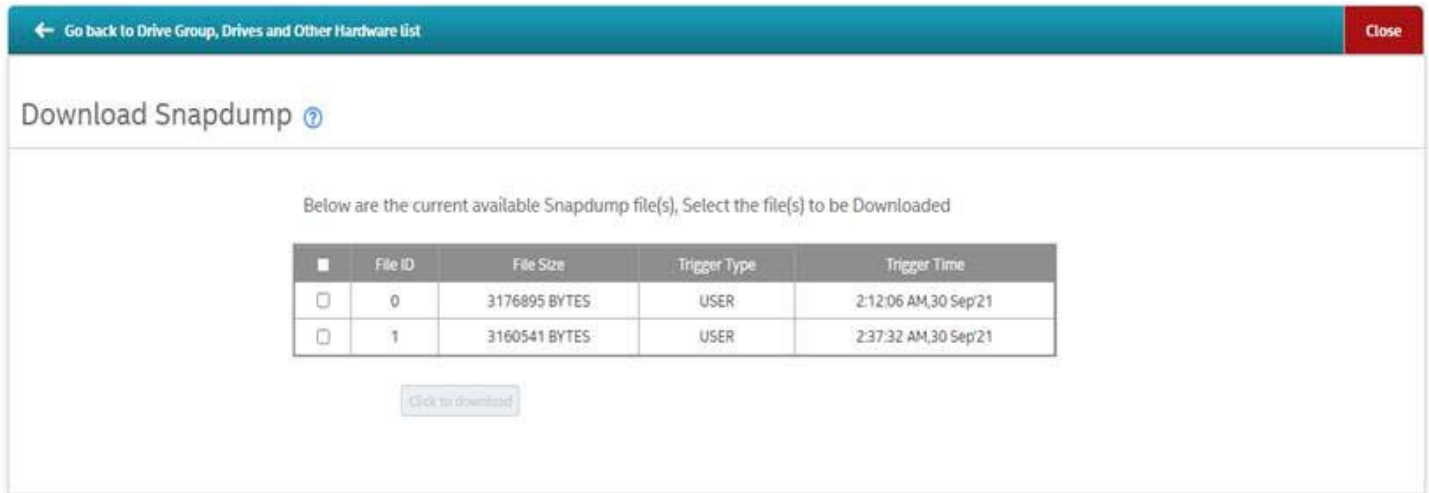
The user can set the Snapdump properties from the **Snapdump Actions** menu of the controller dashboard. Snapdump properties can be edited when using advanced software.

Retrieving the Snapdump Output

The LSI Storage Authority software has an option to read the Snapdump logs from the firmware. In certain cases, firmware might hold multiple Snapdump logs. The LSA displays the list of available Snapdump logs with file ID and file size details.

Snapdump logs are generated when Snapdump is triggered by either the firmware or by a user. When a Snapdump log is generated, the user can select the which Snapdump log to downloaded.

Figure 47: Download Snapdump Window



NOTE

The user can only download a Snapdump log when the Snapdump feature is enabled.

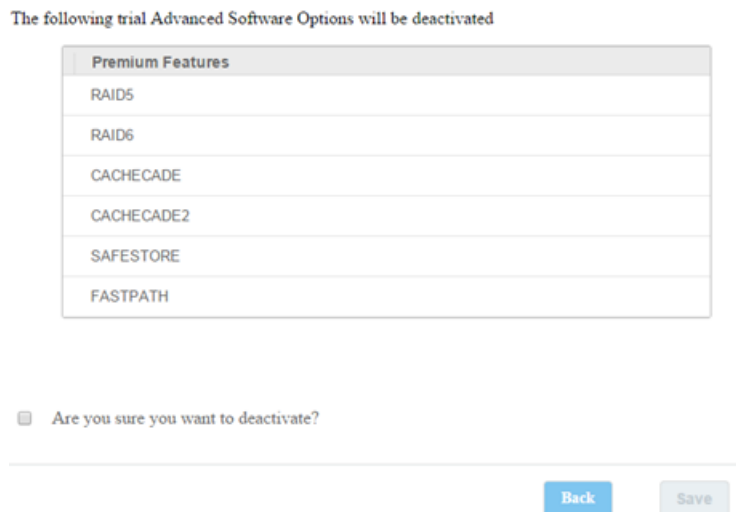
Deactivating Trial Software

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 48: Deactivate All Trial Software – Confirmation Window



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.

Fast Path Advanced Software

The Fast Path software is a high-performance I/O accelerator for SSD arrays connected to a MegaRAID controller card. This advanced software is an optimized version of Broadcom MegaRAID technology, particularly those that demonstrate high random read/write operation workloads, when deployed with a Broadcom MegaRAID SATA+SAS controller connected to SSDs.

SafeStore Encryption Services

The 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 an SED drive from its storage system or the server in which it resides, the data on that drive is encrypted, and it 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.

The instant Secure Erase feature allows you to instantly and securely render data on SED drives unreadable, saving businesses time and money by simplifying the 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.

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 Security**.

The **Enable Security** dialog appears.

Figure 49: Enable Security Dialog

← Go back to Drive Group, Drives and Other Hardware list

Enable Security [?](#)

Controller ID: 0 Product Name 0x500062b200000000
Enabling security on this controller will have the option to create secure virtual drives using a security key.

Choose the security key management mode:
Local Key Management(LKM) ▼

Security Key Identifier: --Security Key Identifier-----
Specify a security key identifier.The controller has provided a default identifier for you. You may use this string or a new identifier.
If you have multiple security keys, the identifier will help you determine which security key to enter.

Suggest Security Key
Security Key: --Security Key-----
The security key will be used to lock each self-encrypted drive attached to the controller.
For maximum security, user 32 varied characters, you may optionally choose for the system to suggest a strong security key.
Note:
The security key is case-sensitive and must be between 8 and 32 characters, contain atleast 1 number, 1 lowercase letter 1 uppercase letter and 1 non-alphanumeric character(e.g. >?@)

Show Key
 Pause for Passphrase
 Enforce strong password security
Password: --Password-----
Optionally, You may enter a password to provide additional security. If you choose "Pause for password at boot time", you must enterit whenever you boot the server.
Note:
The password is case sensitive and must be between 8 and 256 characters.
If enforce strong password security is selected, then password field should contain atleast 1 number, 1 lowercase letter,1 uppercase letter and 1 non-alphanumeric character(e.g. >?@)

Confirm:

Show Password

Are you sure you want to enable security?
 Confirm

More Actions:

- Schedule Consistency Check (Starts on Monday : 12 AM)
- Set Consistency Check Properties
- Schedule Patrol Read (Starts on Wednesday : 5 PM)
- Start Patrol Read
- Advanced Software Options
- Set Adjustable Task Rate
- Disable SSD Guard
- Enable SCSI Unmap
- Download RTT Dump
- Snap Dump Actions
- Enable Security
- Manage Power Save Setting

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

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

- Security Key Identifier** – The controller, by default, assigns a security key identifier.
- Security Key** – Provides you with an option to create secure virtual drives by specifying the security key.
- 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 not selected.

To enable drive security, perform the following steps:

- 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 **Suggest Security Key**, will be visible to you. If you do not select this option, the security key will not be visible to you.

IMPORTANT

Ensure that you write down this security key 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 nonalphanumeric 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 (for example, < > @ +). 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.

IMPORTANT

Ensure you write down the password 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, and then click **Enable Security** to confirm that you want to enable drive security on this controller.

Changing Drive Security Settings

NOTE

Drive security settings cannot be changed when EKM is enabled. Changes to drive security settings for EKM will fail from LSA.

When EKMS is enabled from another application, LSA:

- Does not allow security settings to be disabled or changed.
- Does not support changes to EKM or ESKM.
- Allows users to create secured VD.
- Allows users to secure PD.

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 Security**.

The **Change Security** dialog appears.

Figure 50: Change Security Dialog

The screenshot shows the 'Change Security' dialog box. At the top, it says 'Change Security' with a help icon. Below that, it displays the Controller ID and Product Name. The current security mode is 'FW'. There are two radio buttons: 'Change current security settings' (selected) and 'Switch to External Key Management(EKM) mode (EKM is currently not supported by LSA)'. Under 'Change current security settings', there are three main sections: 'Use the existing security key identifier', 'Enter a new security key identifier', and 'Use the existing drive security key'. The 'Use the existing security key identifier' section has a text field for the 'Current Security Key Identifier' containing 'RAID_805e00014_6077fc11'. The 'Use the existing drive security key' section has a 'Suggest Security Key' button and a 'Security Key' text field. There are also checkboxes for 'Show Key', 'Pause for Passphrase', and 'Enforce strong password security'. A 'Password' text field and a 'Confirm' text field are also present. At the bottom, there is a confirmation question: 'Are you sure you want to change the current security settings?' with 'Confirm' and 'Change Security' buttons. On the right side, there is a 'More Actions' menu with options like 'Schedule Consistency Check', 'Set Consistency Check Properties', 'Schedule Patrol Read', 'Start Patrol Read', 'Advanced Software Options', 'Set Adjustable Task Rate', 'Disable SSD Guard', 'Enable SCSI Unmap', 'Download RTT Dump', 'Snap Dump Actions', 'Disable Security', 'Change Security', and 'Manage Power Save Setting'.

2. Select the **Change current security settings** radio button from the **Current drive security mode is FW** field. When LKMS is enabled, LSA will show the current drive security mode as FW instead of LKM.
3. Either you can use the existing security key identifier assigned by the controller, or you can 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 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** check box.

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 (for 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** 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.
11. Select the **Confirm** check box 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.

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 Security**.
A warning message appears asking for your confirmation.
2. Select **Confirm** and click **Yes, Disable Security**.
The software disables drive security.

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**, 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**: Import the foreign configuration from all the foreign drives.
 - **Clear**: Remove the configuration 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.

Managing Drive Groups

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

Viewing Drive Group Properties

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

The following figure and table describes the drive group properties.

Figure 51: Drive Group Properties Window

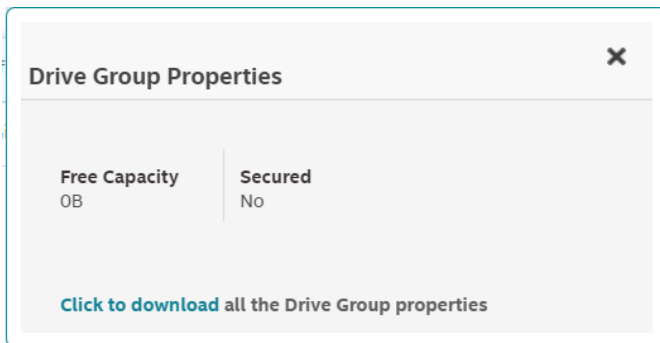



Table 9: Drive Group Properties

Property	Description
Free Capacity	Indicates the free space available in the drive group.
Secured	Indicates whether the drive group is secured.

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

Adding a Virtual Drive to a Drive Group

You can add virtual drives to an existing drive group provided sufficient storage space exists in the current 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 [Adding Virtual Drives](#) for details on creating virtual drives.
4. Click **Add Virtual Drives**.
The newly created virtual drive gets added to the selected drive group.

Managing Virtual Drives

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

Viewing Virtual Drive Properties

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

Figure 52: Virtual Drive Properties Window

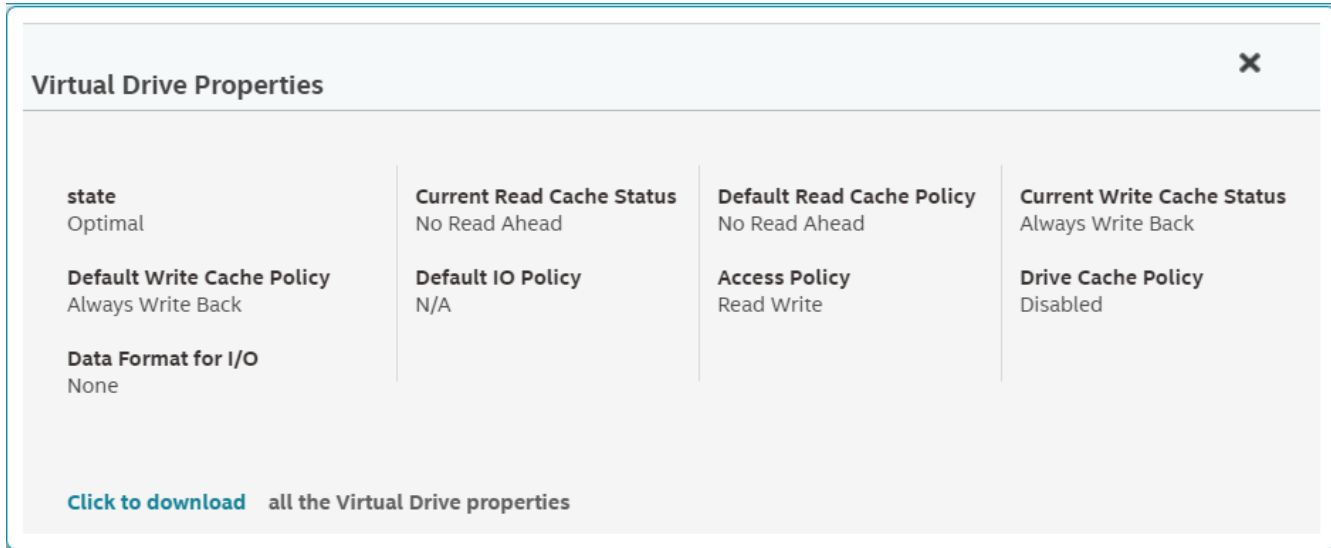


Table 10: Virtual Drive Properties


Property	Description
state	The current status of the virtual drive. These options are available: <ul style="list-style-type: none"> • Optimal • Partially Degraded • Degraded • Offline
Current Read Cache Status	The read cache policy for the virtual drive. These options are available: <ul style="list-style-type: none"> • No Read Ahead
Default Read Cache Policy	<ul style="list-style-type: none"> • Read Ahead • No Read Ahead
Current Write Cache Status	The write cache policy for the virtual drive. These are available: <ul style="list-style-type: none"> • Write Back • Write Through • Always Write Back
Default Write Cache Policy	<ul style="list-style-type: none"> • Default • Enabled • Disabled

Property	Description
Default IO Policy	<ul style="list-style-type: none"> • Direct IO • Cached IO
Access Policy	<p>The access policy for the virtual drive. These options are available:</p> <ul style="list-style-type: none"> • Read Write • Read Only • Hidden <p>The Hidden policy is applicable for only hidden virtual drives. No other access policies are applicable after you select Hidden as the access policy.</p>
Drive Cache Policy	<p>The virtual drive cache setting. These options are available:</p> <ul style="list-style-type: none"> • Unchanged • Enable • Disable
Data Format I/O	<p>The current physical drive format for the virtual drive I/O performed by the hardware.</p>

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, and 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 Properties** dialog appears.

Figure 53: Modify Virtual Drive Properties Dialog

Modify Virtual Drive:
VDName_00 Properties

Virtual Drive Name
VDName_00

Read Policy
No Read Ahead

A controller attribute indicating the current Read Policy mode

No Read Ahead
In No Read Ahead mode, read ahead capability is disabled.

Write Policy
Write Back

Drive Write Cache Policy
Disabled

Save settings

4. Change the virtual drive properties as needed.
For information about these properties, see [Adding Virtual Drives](#).
5. Click **Save settings**.

Start and Stop Locating a Virtual Drive

If the drives in the virtual drives reside 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, and 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. Select **Actions > Stop Locate** to stop the LEDs from blinking.

Erasing a Virtual Drive

The virtual drive erase function operates on a specified virtual drive and overwrites all user-accessible locations. It supports nonzero patterns and multiple passes. The virtual drive erase function optionally deletes the virtual drive and

erases the data within the virtual drive's LBA range. The virtual drive erase function is a background operation, and it posts events to notify users of the progress.

Perform the following steps to erase a virtual drive.


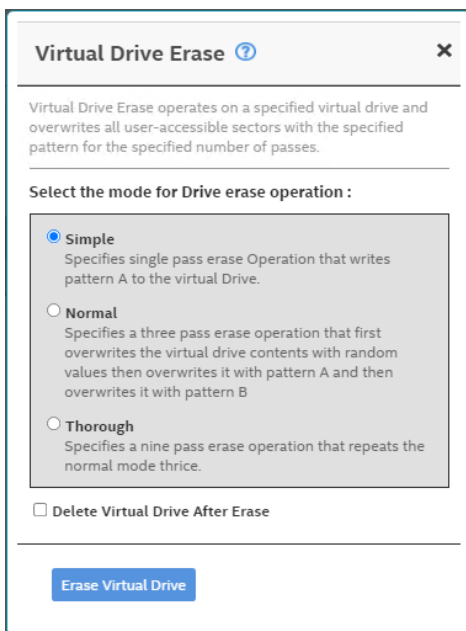
1. Navigate to the Controller dashboard, click a drive group name (for example, **DG_1**).
Click the  icon that corresponds to a drive group to display its contents.
The virtual drives and physical drives associated with the selected drive group appear.
2. Click the virtual drive whose content you want to erase.
3. Select **Element(s) Actions > Erase**.
The **Virtual Drive Erase** dialog appears.

Figure 54: Virtual Drive Erase Dialog



The dialog shows the following modes:

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


4. Select the **Delete Virtual Drive After Erase** check box to delete the virtual drive after the erase operation has been completed.
5. Select a mode and click **Erase Virtual Drive**.
A warning message appears asking for your confirmation.
6. Click **Yes, Erase Drive**.
After the virtual drive erase operation has started, the **Stop Erase** option is enabled in the **Element(s) Actions** menu. You can monitor the progress of the erase operation. See [Background Operations Support](#).

Initializing a Virtual Drive

When you create a new virtual drive using 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, and click a drive group name (for example, **DG_1**).

Click the  icon that corresponds 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.

NOTE

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](#).

Starting Consistency Check on a Virtual Drive

Perform the following steps to start consistency check on a virtual drive. For more information about consistency check, see [Running the Consistency Check](#).

1. Navigate to the Controller dashboard, click a drive group name (for example, **DG_1**).

Click the  icon that corresponds 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.

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 that corresponds 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.

NOTE

You cannot delete operating system drives. If you try to do so, an error message appears.

Behavior of Virtual Drive Operations on VMware

On the LSA client, when you select multiple virtual drives on which to perform operations such as Drive Initialization, Consistency Check, Drive Erase, and so on, the VMware ESXi Server might time out, resulting in a delay of the operation that is being performed.

Because the requested operations are process-intensive and take a long time to complete, the LSA GUI can transition into an idle state until either the requested operation completes or the VMware ESXi Server times out. During this state (idle state of LSA GUI), you cannot perform any other action on the LSA GUI.

While the LSA client is processing the requested operation in the background and is taking more than the usual time to complete, the LSA GUI transitions in to an idle state without letting you perform any other operation; you might assume that the LSA GUI has hung up because you cannot see what is happening in the background.

You can overcome this problem by enabling the **Schedule Panel** feature. The following are some of the advantages of enabling the Schedule Panel feature:

- When performing any process-intensive operations on multiple virtual drives such as Drive Initialization, Consistency Check, Drive Erase, and so on, a Schedule Panel screen is displayed, which shows the Target ID of the virtual drive and the status of the operation that is being performed.
- This Schedule Panel screen is updated on completion of each process, indicating whether the request is being processed or not, in a real time.
- The **Close** button is enabled on the Schedule Panel screen only when all the operations that are being performed are complete, allowing you to close the Schedule Panel screen only when the button is enabled.
- Until the operations are complete, you cannot perform any other action on the LSA GUI.

Enabling the Schedule Panel Feature

By default, the operating systems that run on the VMware ESXi OS are enabled with the **Schedule Panel** feature. However, if you are running Windows, Linux, or any other operating system, you can modify the parameters to enable this feature.

Perform the following steps to enable the **Schedule Panel** feature:

1. Browse to `C:\Program Files (x86)\LSI\LSIStorageAuthority\server\html\files`.
2. Open the `configfile.json` file.
3. Search for the `server` field in the `configfile.json` file.
4. Depending on the operating system you are using, change the value present in the `server` field to:
 - VMware ESXi Operating System – 0 .
 - Windows and Linux Operating Systems – 1 .
 - All other operating systems – 2 .

LWA is a stand-alone installation. To enable the **Schedule Panel** feature for the LWA server, change the value in the present server field to 1 or 2 .

Disabling the Schedule Panel Feature

By default, systems using the VMware ESXi OS are enabled with the **Schedule Panel** feature. However, you can also modify the parameters to disable this feature.

Perform the following steps to disable the **Schedule Panel** feature:

1. Browse to `C:\Program Files (x86)\LSI\LSIStorageAuthority\server\html\files`.
2. Open the `configfile.json` file.
3. Search for the `schedule` field in the `configfile.json` file.
4. Modify the value present in the `schedule` field to:

- To disable the **Schedule Panel** feature – 0 .

Disabling this feature might result in the VMware ESXi server timing out if you have selected multiple virtual drives to perform process-intensive operations. See [Behavior of Virtual Drive Operations on VMware](#) for more information.

- To enable the **Schedule Panel** feature – 1 .

Modifying the Time Limit of the Schedule Panel

The **Close** button on the **Schedule Panel** screen is only enabled when all the operations that are being performed are complete; until the operations are complete, you cannot perform any other action on the LSA GUI. However, using this Time Limit feature, you can configure to close the **Schedule Panel** screen if operation did not complete within the expected time. The maximum wait time for each virtual drive operation to complete is set to four seconds by default.

Perform the following steps to modify the time limit of the Schedule Panel:

1. Browse to `C:\Program Files (x86)\LSI\LSIStorageAuthority\server\html\files`.
2. Open the `configfile.json` file.
3. Search for the `maxTime` field in the `configfile.json` file.
4. Modify the value present in the `maxTime` field according to your requirement.

By default, the `maxTime` value is set to 4000 milliseconds (four seconds) for each virtual drive operation.

5. Save your new settings before closing the `configfile.json` file.

Managing Physical Drives

The LSI Storage Authority software lets you manage all the physical drives connected to the controller.

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 55: Physical Drive Properties Window

Physical Drive Properties			
Device ID 329	Enclosure : Slot EN_362 : 10	Status Unconfigured Good	Exposed As PHYSICAL-DEVICE
Model HUC156030CS4204	Vendor HGST	Serial Number 0TGR945V	Shield Counter 0
Coerced Size 278.937GiB	Raw Capacity 279.461GiB		
General Properties			
SAS Address 0 0x5000cca05b2889d2	Capable Speed UNKNOWN	Temperature 34C (93.2 F)	Firmware Revision Level C703
Power State Powersave	Unmap Capable No	Physical Sector Size 4KiB	Media Error Count 0
Multipath No	Multiport No	Current Write Cache Disabled	Default Write Cache Yes

Table 11: Physical Drive Properties

Property	Description
Device ID	The device ID of the physical drive that is assigned by the manufacturer.
Enclosure : Slot	The slot number of the enclosure in which the drive is connected.
Status	The current status of the physical drives.
Exposed As	To differentiate the physical drives, the drives are exposed as one of the following drive: <ul style="list-style-type: none"> • JBOD • PHYSICAL-DEVICE
Model	The model number of the physical drive.
Vendor	The vendor of the physical.
Serial Number	The serial number of the physical drive.
Shield Counter	The shield counter value

Property	Description
Coerced Size	A drive property that indicates the capacity to which a drive has been coerced (forced) to make it compatible with other drives that are nominally the same capacity.
Raw Capacity	The actual full capacity of the drive before any coercion mode is applied to reduce the capacity.
General Properties	
SAS Address 0	The World Wide Name (WWN) for the physical drive.
Capable Speed	The maximum speed supported by the device.
Temperature	The temperature of the physical drive.
Firmware Revision Level	The revision level of the physical drive's firmware.
Power State	The Power Status displays the following status: <ul style="list-style-type: none"> • On – A physical drive is spun up. • Powersave – A physical drive is spun down.
Unmap Capable	The device maximum supported speed.
Physical Sector Size	The size of the physical sector of the drive. The possible options are 4 KiB or 512 B.
Media Error Count	The number of media errors detected while communicating with LUNs.
Multipath	EID and slot number of the path if the drive is in multipath.
Multiport	EID and slot number of the path if the drive is in multiport.
Current Write Cache	The current write cache setting of the LUNs.
Default Write Cache	The default write cache setting of the LUNs.
Commissioned Spare	Indicates whether the drive is a commissioned spare.
Emergency Spare	Indicates whether the drive is an emergency spare.
Shield Diagnostics Completion Time	Indicates the shield diagnostics completion time. This is displayed only if the controller supports shield state.

Locating Tape Drives



If your system is connected to tape drives, LSA lists those connected tape drives. The tape drive is represented with a special (tape ) icon in the **Type** column of the **Physical Drives** tab.

Figure 56: Tape Drive

1 Unconfigured Drives		1 Unconfigured good						
	Enclosure : Slot	Device ID	Type	Interface	Capacity	Sector Size	Status	Model
	EN_0 : 8	1	TAPE	SAS	0B	512B	Unconfigured good	ULTRIUM5

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. Navigate to the physical drive on the Controller dashboard, and select the drive you want to identify, such as Unconfigured Good drive, online physical drive, configured drive, and so on.
2. Select **Element(s) Actions > Start Locate**.
The corresponding LED on the physical drive starts blinking.
3. Select **Element(s) Actions > Stop Locate** to stop the LED from blinking.

Making a Drive Offline

Perform the following steps to place 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, and click a drive group name (for example, **DG_1**).
Click the  icon that corresponds 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 place offline.
3. Select **Element(s) Actions > Make Offline**.
The drive status changes to Offline.

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, and click a drive group name (for example, **DG_1**).
Click the  icon that corresponds 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 **Element(s) Actions > Make Drive Online**.
The drive status changes to Online.

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.

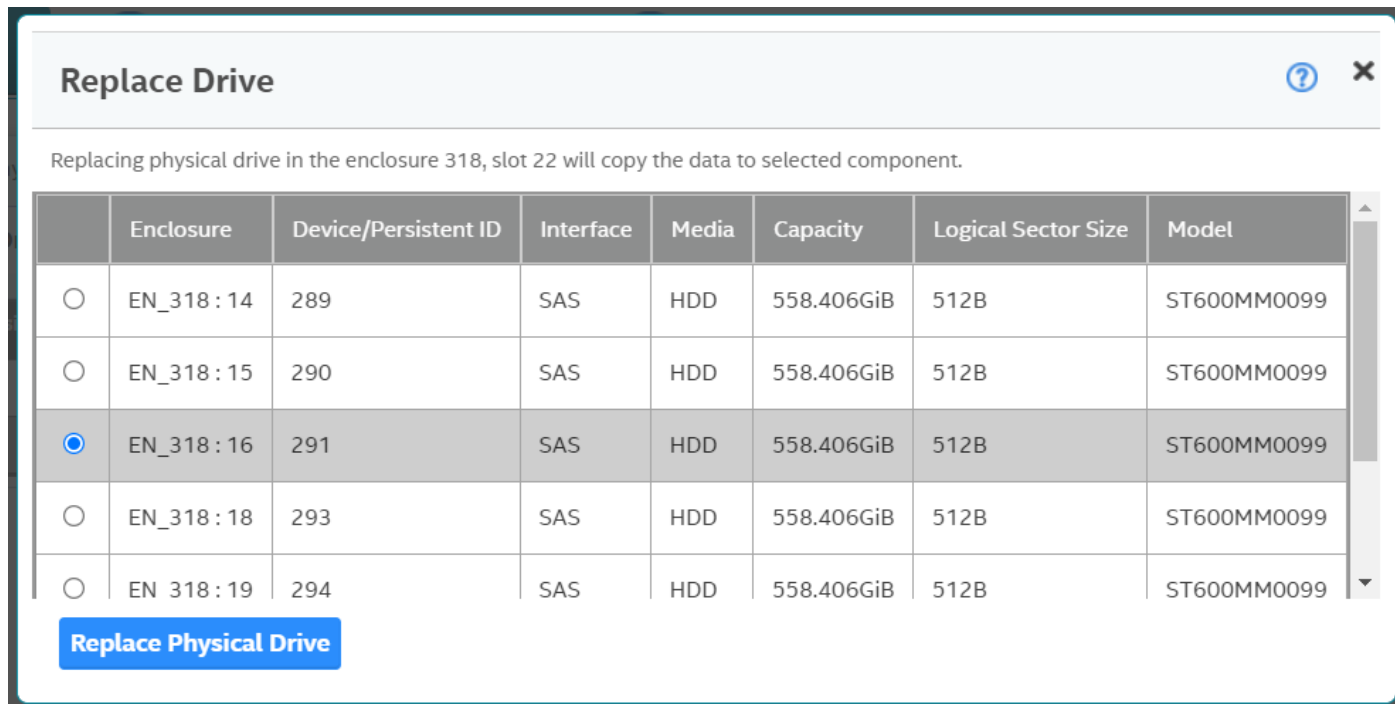
ATTENTION

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

1. Navigate to the Controller dashboard, and click a drive group name (for example, **DG_1**).
Click the  icon that corresponds 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 **Element(s) Actions > Start Replace Drive**.
The **Replace Drive** dialog appears.

Figure 57: Replace Drive Dialog



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.

Marking a Drive as a Missing Drive

If a drive is currently part of a redundant configuration and if the drive is displaying signs of failure, you can mark the drive as missing and start rebuilding data on that drive.

1. Navigate to the Controller dashboard and select **Drive Groups**.
2. Click a drive group name (for example, **DG_1**).
3. Click the **+** icon that corresponds to a drive group to display its contents.
The virtual drives and physical drives associated with the selected drive group appear.
4. Click the **Physical Drive** tab, and select a drive which you want to mark as missing.
5. Select **Element(s) Actions > Mark Drive Offline**.
A confirmation dialog appears.
6. Select **Confirm** and click **Yes, Mark Drive Offline** to proceed towards marking the drive offline.
The drive is marked as offline as shown in the following figure.

Figure 58: Mark Drive Offline Dialog

The screenshot displays the Storage Authority software interface. At the top, there are three summary cards: '2 Drive Groups' (6 Virtual Drives), '19 Drives' (8 Unconfigured Drives), and '2 Other Hardware' (Includes Energy Pack). Below these, the 'DG_0' RAID 1 configuration is expanded, showing '1 Virtual Drives' and '2 Physical Drives'. A table lists the physical drives:

Enclosure:Slot	Device/Persistent ID	Media	Interface	Capacity	Logical Sector Size	Model	NS/LU Count
EN_318:23	298	HDD	SAS	278.875GiB	512B	AL145EB030N	1
EN_318:15	missing 1						

Below the table, the 'DG_1' RAID 6 configuration is partially visible, showing '5 Virtual & 4 Physical Drives' and 'Used 557.749 GiB of 557.75 GiB Available'.

7. Navigate to the **Drives** tab and expand the **Configured Drives** section to see the drives that are offline.
8. Select a drive whose status is offline and go to **Element(s) Actions > Mark Missing**.
9. Navigate to the **Drives** tab.
10. Select an Unconfigured Good drive from the list of Unconfigured Good drives, and go to **Element(s) Actions > Replace Missing Drive**.

The **Replace Missing Drive** dialog appears.

Figure 59: Replace Missing Drive Dialog

The 'Replace Missing Drive' dialog box is shown. It contains the following text: 'Replaces a drive of a degraded array that is marked missing with an Unconfigured Good drive'. Below this is a table with the following data:

	Drive Groups	Raid Level	Missing Position	Enclosure:Slot
<input type="radio"/>	Drive Group 0	Raid 1	1	EN_318:15

At the bottom of the dialog, there is a button labeled 'Replace Missing Drive'.

11. Select the drive and click **Replace Missing Drive**.
12. Navigate to the **Drive Groups** tab and select a new drive.
13. Click **Element(s) Actions > Start Rebuild**.

Replacing a Missing Drive

1. Navigate to the Controller dashboard and select **Drive Groups**.
2. Create a new drive group for any RAID level with two drives.
3. Navigate to the drive group, and mark one of the physical drives from disk group 0 as offline.
4. Select the physical drive that is marked as offline, and click **Mark the drive as missing**.
5. Select **Unconfigured drives**, then select the physical drive, and then **Replace Missing Drive**.
6. Select the drive group where you want to replace the missing physical drive.
7. Click **Ok**.

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 **Element(s) Actions > Assign Global Hot Spare**.
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.

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 **Element(s) Actions > Remove Global Hot Spare**.
The hot spare drive is removed and is listed in the **Unconfigured Drives** section as an unconfigured good drive.

Assigning Dedicated Hot Spares

Dedicated hot spare drives provide protection to one or more specified drive groups. A dedicated hot spare can only be assigned to one drive group at a time.

If you select an Unconfigured Good drive, you have the option of assigning it as a dedicated hot spare drive.

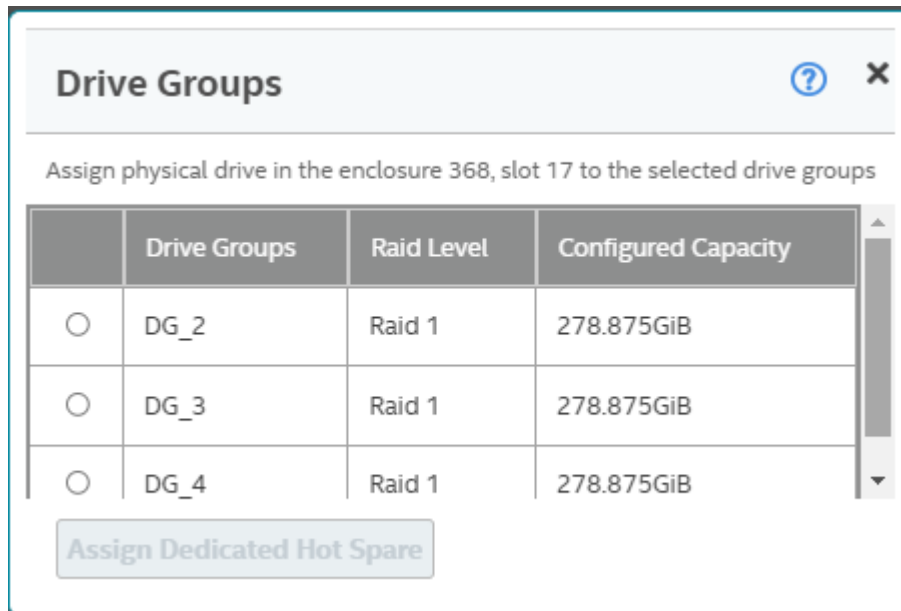
Perform these steps to assign a dedicated hot spare.

1. Navigate to the Controller dashboard, and click the **Drives** tab.
All of the associated drives appear.

- Expand **Unconfigured Drives**, and select an unconfigured good drive.
- Select **Element(s) Actions > Assign Dedicated Hot Spare**.

The **Drive Groups** dialog appears.

Figure 60: Drive Groups Dialog



- Select the corresponding radio button of a drive group and click **Assign Dedicated Hot Spare**.

A confirmation message appears. The unconfigured good drives are now changed to dedicated hot spares. The status of the unconfigured good drives appear as dedicated hot spares in the **Hot Spares** section.

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](#).

Converting an Unconfigured Bad Drive to an Unconfigured Good Drive

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

- Navigate to the Controller dashboard, and click the **Drives** tab.
All of the associated drives appear.
- Expand **Unconfigured Drives**, and select an unconfigured bad drive.
- Select **Element(s) Actions > Make Unconfigured**.
A confirmation message appears.
- Select **Confirm** and click **Yes, Make Unconfigured** to proceed with the operation.
A confirmation message appears.
- Select **Confirm** and click **Yes, Make 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.

Removing a Drive

You might 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 **Element(s) 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 **Element(s) Actions > Undo Prepare for Removal**.

Making Different Types of Drives

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.

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

You can make the following types of drives:

- Unconfigured
- Good
- JBOD

Making an Unconfigured Drive

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

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 **Element(s) Actions > Make Unconfigured**.
A confirmation message appears.
4. Select **Confirm** and click **Yes, Make Unconfigured** to proceed with the operation.
The JBOD drive is changed to an unconfigured drive.

Making a Good Drive

Perform the following steps to change the status of JBOD drives to 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 **Element(s) Actions > Make Good**.
A confirmation message appears.
4. Select **Confirm** and click **Yes, Make Good** to proceed with the operation.
The JBOD drive is changed to a good drive.

Making a JBOD Drive

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 **Element(s) Actions > Make JBOD**.
The unconfigured good drive is changed to a JBOD drive.

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 **Element(s) Actions > More Actions > Drive Erase**.
The **Physical Drive Erase** dialog appears.

Figure 61: 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 **Element(s) Actions** menu. You can monitor the progress of the erase operation. See [Background Operations Support](#).

Erasing a Drive Securely

The Instant Secure Erase feature 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 **Element(s) 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 **Element(s) Actions** menu. You can monitor the progress of the erase operation. See [Background Operations Support](#).

Managing Hardware Components

When you select the **Other Hardware** tab from the Controller dashboard, the hardware components associated with the controller are displayed.

Figure 62: Other Hardware

Enclosure/Slot	Device/Persistent ID	Media	Interface	Capacity	Logical Sector Size	Status	Model	NS/LU Count
EN_318 : 0	275	HDD	SAS	278.937GiB	4KiB	Online	HUC156030CS4204	1
EN_318 : 6	281	HDD	SAS	136.219GiB	512B	Online	ST91467035S	1
EN_318 : 8	283	HDD	SAS	465.25GiB	512B	Failed	ST95004315S	1
EN_318 : 9	284	HDD	SAS	278.937GiB	4KiB	Unconfigured Good	HUC156030CS4204	1
EN_318 : 10	285	HDD	SAS	278.937GiB	4KiB	Unconfigured Good	HUC156030CS4204	1
EN_318 : 11	286	HDD	SAS	278.937GiB	4KiB	Unconfigured Good	HUC156030CS4204	1
EN_318 : 12	287	HDD	SAS	558.406GiB	512B	Unconfigured Good	ST600MM0099	1
EN_318 : 13	288	HDD	SAS	558.406GiB	512B	Online	ST600MM0099	1
EN_318 : 14	289	HDD	SAS	558.406GiB	512B	Unconfigured Good	ST600MM0099	1
EN_318 : 15	290	HDD	SAS	558.406GiB	512B	Unconfigured Good	ST600MM0099	1
EN_318 : 16	291	HDD	SAS	558.406GiB	512B	Unconfigured Good	ST600MM0099	1
EN_318 : 17	292	HDD	SAS	558.406GiB	512B	Unconfigured Good	ST600MM0099	1
EN_318 : 19	294	HDD	SAS	558.406GiB	512B	Unconfigured Good	ST600MM0099	1
EN_318 : 20	295	HDD	SAS	558.406GiB	512B	Unconfigured Good	ST600MM0099	1
EN_318 : 21	296	HDD	SAS	278.875GiB	512B	Unconfigured Good	AL145EB030N	1
EN_318 : 22	297	HDD	SAS	278.875GiB	512B	Unconfigured Good	AL145EB030N	1
EN_318 : 23	298	HDD	SAS	278.875GiB	512B	Unconfigured Good	AL145EB030N	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 the controller performs periodically to determine the condition of the energy pack. You can start 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).

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 63: Energy Pack Properties Window

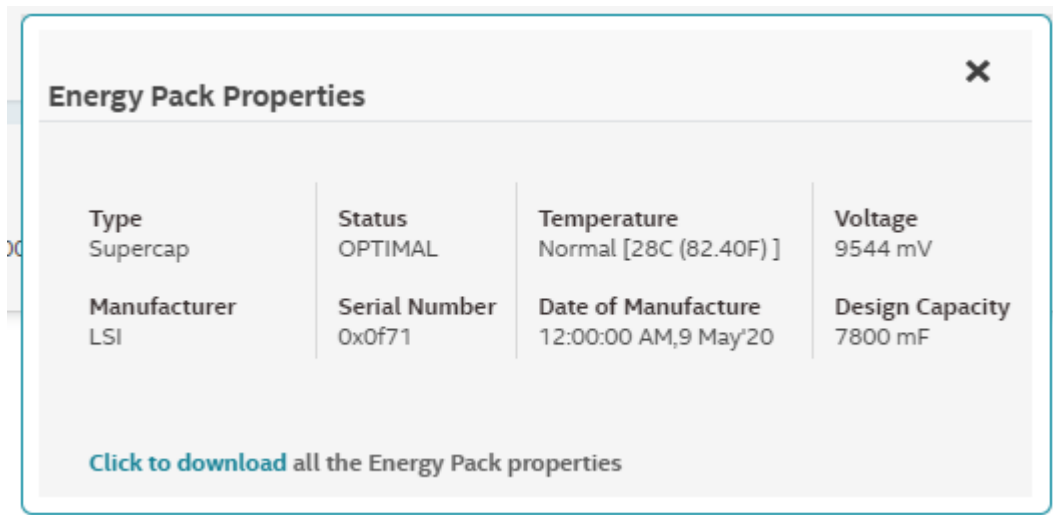


Table 12: Energy Pack Properties

Property	Description
Type	Type of the Energy Pack; for example, TTMC.
Status	Current status of the Energy Pack. The Energy Pack status field has the following states: <ul style="list-style-type: none"> • Optimal • Missing • Failed • Degraded • Degraded [Needs Attention] • Unknown
Temperature	Indicates the current temperature of the Energy Pack. Also indicates whether the current temperature of the Energy Pack is normal or high.
Voltage	Voltage level of the Energy Pack, in mV. Also indicates if the current Energy Pack voltage is normal or low.
Manufacturer	Manufacturer of the Energy Pack.
Serial Number	Serial number of the Energy Pack.
Date of Manufacture	Manufacturing date of the Energy Pack.
Design Capacity	Theoretical capacity of the Energy Pack.

Refreshing Properties

Some of the properties, such as temperature and voltage in the **Properties** section, do not refresh automatically. You must 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 appears.

2. Expand **Energy Pack**, and select a energy pack.
3. Select **Actions > Refresh Properties**.
The properties are updated.

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.

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 64: Enclosure Properties

Enclosure or Logical Enclosure Properties			
Vendor ID LSI	ID 318	Type SAS Expander	Serial Number
Product ID SAS3x40	Location EXTERNAL	Connector [C1.1 & C1.0]	Position 0
Product Revision Level 0601	Number of Slots 24	SAS Address 0x500304801f38e0fd	
Fans Count 0	Temperature Sensors Count 2	Power Supplies Count 0	Voltage Sensors Count 2

[Click to download](#) all the Enclosure or Logical Enclosure properties

If you have selected multiple virtual drives or multiple physical drives, you must click the (Expand button) to perform actions, such as starting 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 are performing certain actions through the **Actions** dialog.

Table 13: Enclosure Properties

Property	Description
Vendor ID	The vendor-assigned ID number of the enclosure.
ID	The ID of the enclosure in which the drive is located.
Type	The type of the enclosure.
Serial Number	The serial number of the enclosure.

Property	Description
Product ID	
Location	Indicates whether the drive is attached to an internal connector or an external connector of the enclosure.
Connector	Indicates the connector name and size of wide port. Single Path: [Port 0-3 x4] – Single x4 Multipath: [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]
Position	
Product Revision Level	The revision level of the enclosure's firmware.
Number of Slots	Total number of available slots.
SAS Address	
Fan count	Total number of fans that are connected.
Temperature Sensors Count	Total number of temperature sensors that are connected.
Power Supplies Count	Total number of power supplies that are connected.
Voltage Sensors Count	Total number of voltage sensors that are connected.

Show Events

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. Select **Actions > Show Events** in the Server dashboard or the Controller dashboard.

The **Show Events** 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 chronological order.

Figure 65: Show Events Window

The screenshot shows the 'Show Events' window with a table of log entries and an 'Actions' menu. The table displays the following data:

Severity Level	Event ID	Locate	Description	Time, Date
Information	357	Energy Pack	Energy pack is discharging	9:47:31 AM,25 Oct'21
Information	500	Controller	Host driver is loaded and operational	9:21:57 AM,25 Oct'21
Information	451	Physical Device SAS	Redundant path restored for PD 0x127(e0x13e/s20) Path (1) 0x5000c500af3fd8e1	9:20:24 AM,25 Oct'21
Information	548	Physical Device	Inserted: PD 0x127(e0x13e/s20)	9:20:24 AM,25 Oct'21
Information	451	Physical Device SAS	Redundant path restored for PD 0x11d(e0x13e/s10) Path (1) 0x5000c500af3fd345	9:20:24 AM,25 Oct'21
Information	548	Physical Device	Inserted: PD 0x11d(e0x13e/s10)	9:20:24 AM,25 Oct'21
Information	451	Physical Device SAS	Redundant path restored for PD 0x125(e0x13e/s18) Path (1) 0x5000c500af3fd6c9	9:20:24 AM,25 Oct'21
Information	548	Physical Device	Inserted: PD 0x125(e0x13e/s18)	9:20:24 AM,25 Oct'21
Information	451	Physical Device SAS	Redundant path restored for PD 0x123(e0x13e/s16) Path (1) 0x5000c500af4d1b1	9:20:24 AM,25 Oct'21
Information	548	Physical Device	Inserted: PD 0x123(e0x13e/s16)	9:20:24 AM,25 Oct'21
Information	451	Physical Device SAS	Redundant path restored for PD 0x12a(e0x13e/s23) Path (1) 0x5000c500af3e6dad	9:20:24 AM,25 Oct'21
Information	548	Physical Device	Inserted: PD 0x12a(e0x13e/s23)	9:20:24 AM,25 Oct'21
Information	451	Physical Device SAS	Redundant path restored for PD 0x129(e0x13e/s22) Path (1) 0x5000c500af3fd4e9	9:20:24 AM,25 Oct'21
Information	548	Physical Device	Inserted: PD 0x129(e0x13e/s22)	9:20:24 AM,25 Oct'21
Information	451	Physical Device SAS	Redundant path restored for PD 0x113(e0x13e/s0) Path (1) 0x5000c500af3e6791	9:20:24 AM,25 Oct'21
Information	548	Physical Device	Inserted: PD 0x113(e0x13e/s0)	9:20:24 AM,25 Oct'21
Information	451	Physical Device SAS	Redundant path restored for PD 0x117(e0x13e/s4) Path (1) 0x5000c500af4beb9	9:20:24 AM,25 Oct'21

The 'Actions' menu on the right includes the following options:

- Download Events
- Clear Events

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

NOTE

The View Event Log is empty when there are only process related events.

Downloading Events

To download the event logs, navigate to the **Show Events** window, and click **Download Events** to download the event log file.

Clearing Events

Perform the following steps to clear the event logs.

1. Click **Clear Events** in the **Show Events** window.
A confirmation dialog appears.
2. Select **Confirm**, and click **Yes, Clear Events**.
The event logs are cleared.

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

Viewing 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 14: Default Theme Settings

Theme	Default	Default File Name/Property Name
Logo		mainlogo.png <root>\LSI\LSIStorageAuthority \server\html\ui\images Dimensions <ul style="list-style-type: none"> • Width – 263 pixels • Height – 43 pixels • Bit depth – 32 LSI Storage Authority is present in <root>\LSI\LSIStorageAuthority \server\html\js\message_en.js in the form of <Key>:<Value> format. This value string can be customized.
Header		headbackground.png <root>\LSI\LSIStorageAuthority \server\html\ui\images Dimensions <ul style="list-style-type: none"> • Width – 1172 pixels • Height – 125 pixels • Bit depth – 32

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 pixels × 120 pixels)

- Width – 372 pixels
- Height – 120 pixels
- Bit depth – 32

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.

Customizing the Header Background Image

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 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 372 pixels × 120 pixels

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

Introduction to the Light Weight Monitor System

Overview

The Broadcom Light Weight Monitor (LWM) software provides server monitoring capabilities. The software monitors the status of the controller cards, virtual drives, drives, and other devices on the server. It alerts you of any issues that require immediate attention with system logs and real-time email notifications (based on the alert settings). For example, when the devices automatically go from an optimal state to a different state, such as a created virtual drive goes to a degraded state, the software gets those events from the controller. The software then sends a notification to you, using system logs and real-time email notifications based on different alert delivery methods.

Supported Operating Systems

The following table provides the OS support matrix for the Light Weight Monitor system.

Table 15: Supported Operating Systems

Operating System	Version
Microsoft	<ul style="list-style-type: none"> Windows Server 2022 Windows Server 2019 Windows 10 (RS5)
Linux	<ul style="list-style-type: none"> Red Hat Enterprise Linux 8.4 Red Hat Enterprise Linux 8.3 Red Hat Enterprise Linux 8.2 SUSE Linux Enterprise Server 15 SP3 SUSE Linux Enterprise Server 15 SP2
Ubuntu	<ul style="list-style-type: none"> Ubuntu 20.04 LTS

Alert Delivery Methods Based on Severity Levels

Based on the severity level (Fatal, Critical, Warning, and Information), the default alert delivery methods change. By default, each severity level has one or more alert delivery methods configured for it, as shown in the following table. To modify these alert delivery methods, see [Changing the Default Alert Delivery Method for Each Severity Level](#). The software supports the following alert delivery methods:

- System log
- Email notification

Table 16: Severity Level and Default Alert Delivery Methods

Severity Level	Default Alert Delivery Method	Description
Fatal	System log and email notification	A component has failed, and data loss has occurred or will occur.
Critical	System log	A component has failed, but no data loss has occurred.
Warning	Warning	A component might be close to a failure point.
Information	System log	Informational message. No user action is necessary.

System Log

By default, the system log is maintained in the local system log (syslog). Depending on the operating system you are using, the system log is maintained in the following syslog locations:

- Windows – **Event Viewer > Windows Logs > Application**
- Linux – `/var/log/messages`

Email Notification

By default, fatal events are notified through system log and email notifications. You can configure these notifications using this feature.

Based on your configuration, the email notifications are delivered to your inbox in a format as shown in the following figure. The email notification contains information, such as event's description, 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.

Figure 66: Email Notification Format

```

From monitor@server.com Thu Apr 25 05:38:58 2002
Return-Path: <monitor@server.com>
Received: from 135.24.227.207 (linux-7kul.lsi.com [135.24.227.207])
    by dhcp-135-24-227-243.lsi.com (8.13.8/8.13.8) with ESMTTP id g3P08v6T031
181;
    Thu, 25 Apr 2002 05:38:57 +0530
Date: Mon, 5 Oct 2015 17:18:09 +0530 (IST)
From: monitor@server.com
To: root@dhcp-135-24-227-243.lsi.com, hema@dhcp-135-24-227-243.lsi.com
Message-ID: <1290884027.3519.1444045689147.JavaMail.root@linux-7kul.lsi.com>
Subject: Warning | Event occured on: linux-7kul.lsi.com
MIME-Version: 1.0
Content-Type: text/plain;charset="UTF-8"
Content-Transfer-Encoding: 7bit

Controller ID: 0    PD Predictive failure:    Port 4 - 7:1:19
Event ID:96
Generated On: Mon Oct 05 17:18:19 IST 2015

System Details---
Server IP:135.24.227.207
OS name:Linux
OS Version:3.0
Driver Name:megaraid_sas
Driver Version:06.810.07.00

Image Details---
BIOS Version : 6.30.03.0_4.17.08.00_0x06130200
Firmware Package Version: 24.12.0-0011
Firmware Version : 4.620.00-4804

```

The Light Weight Monitor supports one of the SMTP authorization protocols called Auth Login. If your SMTP server is configured with Auth Login support, you must enable the Light Weight Monitor with login-specific information. For more information, see [Setting Up the Email Server](#).

By default, the Light Weight Monitor sends the email notification to the configured recipients. The recipients' information is based on your SMTP configuration. You can configure the SMTP configuration to add one or more recipient addresses. For more information, see [Adding the Email Addresses of Alert Notification Recipients](#).

Time Synchronization

A mismatch between the controller time settings and the OS time setting is possible during daylight savings time (DST) or known lower layer limitations. In this case, the timestamp will be different on the I/O. The Light Weight Monitor resolves this limitation by synchronizing the OS time and the controller time on an hourly basis.

NOTE

This is a default feature of the Light Weight Monitor software. This feature is not configurable.

Installing the Light Weight Monitor System

Two ways to install the Light Weight Monitor system exist:

- Through the LSI Storage Authority Master Setup
- Through the Light Weight Monitor stand-alone installation.

Installing the LSI Storage Authority Software on the Windows Operating System through the LSA Master Setup

Perform the following steps to install the LWM using the LSI Storage Authority Master Setup.

1. Run the LSI Storage Authority `setup.exe` file.
The **InstallShield Wizard** dialog appears.
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. (optional) Enter your user name and the organization name, and click **Next**.
The **Port Configuration Settings** dialog appears.

By default, LSA communicates on web server port 2463 and LSA server port 9000. Make sure 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 also can edit these port details after installation.

5. Click **Next** if you want to proceed with the default port configuration settings.
The **Destination Folder** dialog appears with the default file path.
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 or problem occurrences before they affect you.

Figure 67: Configure Range of Events to Generate Alert Notifications Dialog

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 change these configuration options as per your requirement, at any point in time.

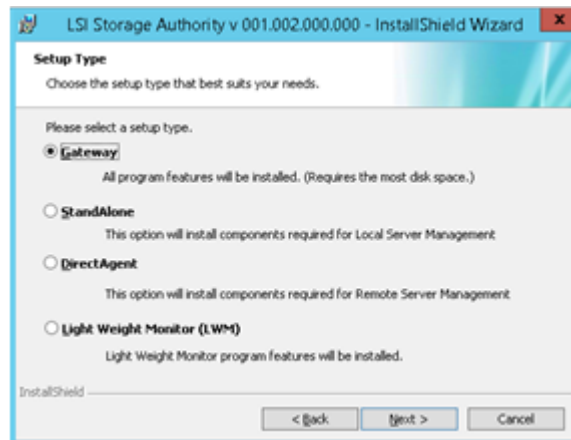
You can also change these configuration options per your requirement at any point in time by editing the `lsa.conf` file in the `LSI Storage Authority /conf` directory and choosing the required parameter. For example, if you have selected **Since Last Shutdown** as a configuration option to retrieve 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 68: Setup Type Dialog

9. Select Light Weight Monitor (LWM).

10. Click **Next**. The **Ready to Install the Program** window 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** check box to view the windows installer log file.

The log file (`LSA_install.log`) is created in the same folder from where the `setup.exe` command is executed.

12. Click **Finish**.

The Light Weight Monitor is successfully installed.

Uninstalling Light Weight Monitor Software on the Windows Operating System

You can uninstall the Light Weight Monitor software either through the **Control Panel** or through the application shortcut in the **Start** menu.

Uninstalling the Software through the Application Shortcut from the Start Menu

1. Select **Start > All Programs > LSI > LSISStorageAuthority > Uninstall LSA**.

The Light Weight Monitor package and all of the installed files are removed.

Uninstalling the Software through the Control Panel

1. To uninstall the software, select **Programs and Features** from the **Control Panel**.

2. Select the Light Weight Monitor software from the listed programs and click **Uninstall**.

The Light Weight Monitor package and all of the installed files are removed.

Installing the Light Weight Monitor on the Linux Operating System

Perform the following steps to install the LWM on the Linux OS.

1. Log on to the system with root privileges.

2. Run the following command:

```
rpm -ivh LightWeightMonitor-<x.xx.xx-xx>
```

Where: `x.xx.xx-xx` is the Light Weight Monitor version number.

The LWM agent is installed in the following directory path:

```
opt/lsi/LSIStorageAuthority/
```

Uninstalling the Light Weight Monitor on the Linux Operating System

1. Run the following command to search for the Light Weight Monitor package.

```
rpm -qa | grep Light*
```

2. Select the Light Weight Monitor package from the listed programs, and run the following command to uninstall the package:

```
rpm -e LightweightMonitor-<x.xx.xx-xx>
```

Where `x.xx.xx-xx` is the Light Weight Monitor version number.

Configuring the Light Weight Monitor System

This section provides information on how to configure the LWM agent.

NOTE

You must edit the `syslog.conf` and `config-current.JSON` files to configure the email server, alert settings, and system logs. These files have the write permissions. *Do not* edit any other files in the package.

NOTE

If you need to revert back to the original settings, you must copy the contents of the `config-default.JSON` file to the `config-current.JSON` file. This action restores the configuration settings to its original state.

Setting Up the Email Server

Perform the following steps to configure the email server.

1. Open the `config-current.JSON` file in one of the following directories:

- Windows OS: `<ProgramFilesFolder>\LSI\LSIStorageAuthority\conf\monitor`
- Linux OS: `cd /opt/lsi/LSIStorageAuthority/conf/monitor`
- Solaris x86 OS: `cd /opt/lsi/LSIStorageAuthority/conf/monitor`

2. Make the following changes in the "email" sections of the files.

- Windows OS, Linux OS, and Solaris x86 OS:

Default Configuration:

```
"email":
{
  "isActive": true,
  "type": "EMAIL",
  "sender": "lsa-monitor@server.com",
  "server": "127.0.0.1",
  "to":
  [
    "root@localhost"
  ],
  "authentication":
  {
    "type": "NONE"
  }
}
```

Changes to Be Made

Edit the `server` field with the IP address of your SMTP server. For example,

```
"server": "135.24.227.243"
```

3. If on your SMTP server, the Auth Login feature is enabled and if you want to enable this feature in the LWM agent, enter the following authentication details.

- Windows OS, Linux OS, and Solaris x86 OS:

Default Configuration:

```

"email":
{
  "isActive": true,
  "type": "EMAIL",
  "sender": "lsa-monitor@server.com",
  "server": "127.0.0.1",
  "to":
    [
      "root@localhost"
    ],
  "authentication":
    {
      "type": "NONE"
    }
}

```

Changes to Be Made

Enter the authentication details in the user name and password fields, and make the following changes in the "authentication" field.

```

"username": "lsi",
"password": "xxxx",
"authentication":
{
  "type": "AUTH-LOGIN"
}

```

Where: `lsi` represents SMTP server's user name and `xxxx` represents the Base 64 converted SMTP configuration password.

4. Save the `config-current.JSON` file.
5. Perform the following steps to restart the Light Weight Monitor services for the changes to take effect.

- Windows OS:

- a. In the command prompt, type the following command to stop the Light Weight Monitor service:

```
sc stop LSAService
```

- b. Type the following command to start the Light Weight Monitor service:

```
sc start LSAService
```

- Linux OS:

Run the following command:

```
/etc/init.d/LsiSASH restart
```

- Solaris x86 OS:

Run the following command:

```
/etc/init.d/LsiSASH restart
```

It is also recommended to restart the system log. To restart the system log, run the following command:

```
svcadm restart system-log
```

Adding the Email Addresses of Alert Notification Recipients

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

1. Open the `config-current.JSON` files in one of the following directories:

- Windows OS: `<ProgramFilesFolder>\LSI\LSIStorageAuthority\conf\monitor`
- Linux OS: `cd /opt/lsi/LSIStorageAuthority/conf/monitor`
- Solaris x86 OS: `cd /opt/lsi/LSIStorageAuthority/conf/monitor`

2. Make the following changes in the following sections of the files.

- Windows OS, Linux OS, and Solaris x86 OS:

In the "email" section and "to" field, add your email ID. You can add multiple email addresses by separating them with commas. For example,

```
"to": [ xxx@xx.com, abc@zyz.com, ... ]
```

Where: `xxx@xx.com` or `abc@zyz.com` represents your email ID.

3. Save the `config-current.JSON` file.

4. Perform the following steps to restart the Lightweight Monitor services for the changes to take effect.

- Windows OS:

a. Start the command prompt.

b. Type the following command to stop the Lightweight Monitor service:

```
sc stop LSAService
```

c. Type the following command to start the Lightweight Monitor service:

```
sc start LSAService
```

- Linux OS:

Run the following command:

```
/etc/init.d/LsisASH restart
```

Configuring Alert Settings

Changing the Default Alert Delivery Method for Each Severity Level

Perform the following steps to change the default alert delivery methods for each severity level.

1. Open the `config-current.JSON` file in the following directories:
 - Windows OS: `<ProgramFilesFolder>\LSI\LSIStorageAuthority\conf\monitor`
 - Linux OS: `cd /opt/lsi/LSIStorageAuthority/conf/monitor`
 - Solaris x86 OS: `cd /opt/lsi/LSIStorageAuthority/conf/monitor`
2. For example, if you want to edit the default alert delivery method for the warning severity level from the system log to both the system log and email notification, make the following changes.

- Windows OS, Linux OS, and Solaris x86 OS:

```
{
  "warning": [
    "systemlog",
    "email"
  ]
}
```

3. Save the `config-current.JSON` file.
4. Perform the following steps to restart the Light Weight Monitor services for the changes to take effect.

- Windows OS:
 - a. Start the command prompt.
 - b. Type the following command to stop the Light Weight Monitor service:

```
sc stop LSAService
```

- c. Type the following command to start the Light Weight Monitor service:

```
sc start LSAService
```

- Linux OS:
Run the following command:

```
/etc/init.d/LsiSASH restart
```

Changing the Alert Delivery Method for a Specific Event

Perform the following steps to change the alert delivery method for a specific event.

1. Open the `config-current.JSON` file in the following directories:
 - Windows OS: `<ProgramFilesFolder>\LSI\LSIStorageAuthority\conf\monitor`
 - Linux OS: `cd /opt/lsi/LSIStorageAuthority/conf/monitor`
 - Solaris x86 OS: `cd /opt/lsi/LSIStorageAuthority/conf/monitor`
2. Make the following changes to the `events []` array.
 - Windows OS, Linux OS, and Solaris x86 OS:


```
"events": [
  {
    "typeId": 4,
    "severity": "INFO",
    "actions": [
      "email"
    ]
  }
]
```

3. Save the `config-current.JSON` file.
4. Perform the following steps to restart the Light Weight Monitor services for the changes to take effect.

- Windows OS:
 - a. Start the command prompt.
 - b. Type the following command to stop the Light Weight Monitor service:

```
sc stop LSAService
```

- c. Type the following command to start the Light Weight Monitor service:

```
sc start LSAService
```

- Linux OS:
Run the following command:

```
/etc/init.d/LsiSASH restart
```

Changing the Severity Level for a Specific Event

Perform the following steps to change the severity level for a specific event.

1. Open the `config-current.JSON` file in the following directories:
 - Windows OS: `<ProgramFilesFolder>\LSI\LSIStorageAuthority\conf\monitor`
 - Linux OS: `cd /opt/lsi/LSIStorageAuthority/conf/monitor`
 - Solaris x86 OS: `cd /opt/lsi/LSIStorageAuthority/conf/monitor`
2. Make the following changes to the `events []` array:
 - Windows OS, Linux OS, and Solaris x86 OS:

```
"events": [
  {
    "typeId": 4,
    "severity": "CRITICAL",
    "actions": [
      "global"
    ]
  }
]
```

3. Save the `config-current.JSON` file.
4. Perform the following steps to restart the Light Weight Monitor services for the changes to take effect.
 - Windows OS:

- a. Start the command prompt.
- b. Type the following command to stop the Light Weight Monitor service:

```
sc stop LSAService
```

- c. Type the following command to start the Light Weight Monitor service:

```
sc start LSAService
```

- Linux OS:

Run the following command:

```
/etc/init.d/LsiSASH restart
```

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 to improve 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
- Performing patrol read
- 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-11 is 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

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.

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.

Physical Drive States

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

Table 17: 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.

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

Virtual Drive States

The virtual drive states are described in the following table.

Table 18: 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.

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 two 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.

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.

Wide Port

The term *port* identifies 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 targets.

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.

Replace

Replace lets you copy data from a source drive to a destination drive that is not a part of the virtual drive. Replace 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 Replace 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. Replace runs as a background activity, and the virtual drive is still available online to the host.

Replace 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 Replace. This situation avoids putting the drive group in Degraded status.

NOTE

During Replace, if the drive group involved in Replace 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 Replace:

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

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 5 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 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.

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.


Disk Striping

Disk striping lets you write data across multiple drives instead of just one drive. Disk striping involves partitioning each drive storage space into stripes. These stripes are interleaved in a repeated sequential manner. The combined storage space is composed of stripes from each drive. Use 64k for drive groups consisting of SSDs and 256k for drive groups consisting of HDDs.

For example, in a four-disk system 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. Supported stripe sizes are 64k for drive groups consisting of SSDs and 256k for drive groups consisting of HDDs.

Strip Size

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


Disk Mirroring

With disk mirroring (used in RAID 1, 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)



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 19: 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)



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

3_01081-00

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 of 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, RAID 60, and spanned PRL-11 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 20: Spanning for RAID 10, RAID 50, and RAID 60

Level	Description
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 10, RAID 50, RAID 60, and Spanned PRL-11), the span numbering starts from Span 0, Span 1, Span 2, and so on.

Hot Spares

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

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

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

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

NOTE

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

The hot 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-11 drive 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.

Disk Rebuilds

When a drive in a RAID drive group fails, you can rebuild the drive by re-creating the data that was stored on the drive before it failed. The RAID controller re-creates the data using the data stored on the other drives in the drive group. Rebuilding can be 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-11 drive groups.

The RAID controller uses hot spares to rebuild failed drives automatically and transparently, at user-defined rebuild rates. If a hot spare is available, the rebuild 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.)

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.

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.

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.

RAID Levels

The subsequent sections describe the RAID levels in detail.

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.

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

RAID 0 Drive Groups

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 21: 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




RAID 1 Drive Groups

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 22: 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



RAID 5 Drive Groups

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 23: 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.
Weak points	Not well suited to tasks requiring lots of small writes or small block write operations. Suffers more impact if no cache is used. Drive performance is reduced if a drive is being rebuilt. Environments with few processes do not perform as well because the RAID drive group overhead is not offset by the performance gains in handling simultaneous processes.
Drives	3 through 32.

Figure 75: RAID 5 Drive Group with Six Drives



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

3_01085-00

RAID 6 Drive Groups

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 24: 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 requiring a lot of small and/or random write operations. A RAID 6 virtual drive must generate two sets of parity data for each write operation, which results in a significant decrease in performance during write operations. Drive performance is reduced during a drive Rebuild operation. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes. A RAID 6 drive group costs more because of the extra capacity required by using two parity blocks per stripe.
Drives	4 through 32.

The following figure shows a RAID 6 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)



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

3 0108600

RAID 10 Drive Groups

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.

NOTE

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

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

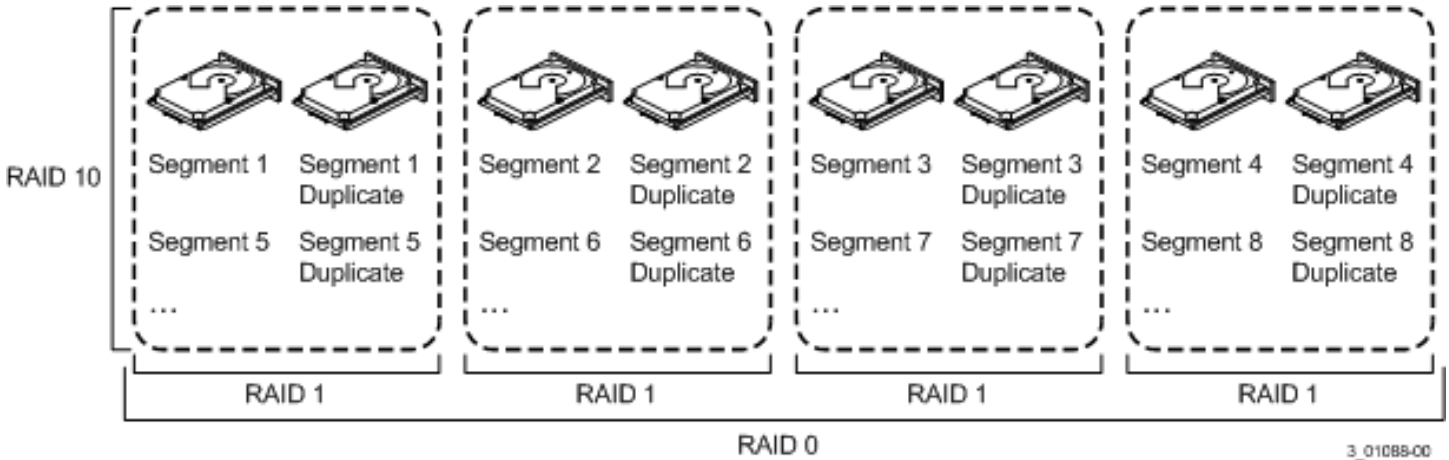
The following table provides an overview of RAID 10.

Table 25: 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.
Weak points	Requires twice as many drives as all other RAID levels except in RAID 1 drive groups.
Drives	36 to 240 – The number of drives in each span must be an even number, and the same number of drives in each span. The maximum number of drives supported by the controller (using an even number of drives in each RAID 10 virtual drive in the span). The MegaRAID 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 77: RAID 10 Level Virtual Drive



RAID 50 Drive Groups

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 26: 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	A minimum of two spans of three drives per span, up to eight spans of 3 to 32 drives per span (limited by the maximum number of drives supported by the controller). Each span must contain the same number of drives. Eight spans of RAID 5 drive groups that contain 3 to 32 drives each (limited by the maximum number of devices supported by the controller).

Figure 78: RAID 50 Level Virtual Drive



RAID 60 Drive Groups

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 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 27: RAID 60 Overview

Uses	<p>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.</p> <p>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.</p> <p>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.</p>
Strong points	<p>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.</p>
Weak points	<p>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.</p> <p>Drive performance is reduced during a drive rebuild. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.</p> <p>A RAID 60 drive group costs more because of the extra capacity required by using two parity blocks per stripe.</p>
Drives	<p>Eight spans of RAID 6 drive groups that contain 4 to 32 drives each (limited by the maximum number of devices supported by the controller).</p>

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

Figure 79: RAID 60 Level Virtual Drive



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

3 01090-00

RAID Configuration Strategies

The following factors in 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.

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 28: 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.
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.

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 29: 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 64 KiB to 256 KiB. 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.
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.

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

RAID Availability

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-11 virtual 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.

Configuration Planning

The factors to consider when planning a configuration are the number of drives 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.

Possible Raid Levels

While creating a new storage configuration, depending on the personality mode, LSI Storage Authority software displays the following possible RAID levels for that personality mode.

JBOD Mode:

The possible RAID levels in JBOD mode are:

- RAID 0
- RAID 1
- RAID 10
- PRL-11
- Spanned PRL-11

RAID Mode:

The possible RAID levels in RAID mode are the default RAID levels supported by the firmware.

None Mode:

- RAID 0
- RAID 1
- RAID 5
- RAID 6
- RAID 10
- RAID 50
- RAID 60
- PRL-11
- Spanned PRL-11

Events and Messages

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).

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 31: 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.

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 32: 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, for example, 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.	
302 Found	The requested resource resides temporarily under a different URI.	
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 a consistency check on RAID 0 virtual drive. The status code is MFI_STAT if coming from firmware.
401 Unauthorized	The authentication credentials are missing or invalid.	GET /servers/{id}/controllers# Run before authentication.
403 Forbidden	The server recognized the credentials in the request, but those credentials do not possess authorization to perform this request. For example, 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 no controller with ID 99.
405 Method Not Allowed	The URI is valid but the HTTP verb specified in the request (for example, 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 (for example, the TTY log) but the file is being deleted and no longer available (and will never be available).
422 Unprocessable Entity	Semantical error (see status code 400 for details).	Try to run consistency check on RAID 0 virtual drive.

Code	Description	Example
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 or an error in backend itself, for example, memory allocation error, or initialization error, and so on.
501 Not Implemented	Implies future availability. The current server lacks the ability to fulfill the request.	None so far.

SAS Address Assignment Rule

The PHY SAS address is calculated by incrementing the controller SAS address by one, based on the number of PHYs.

Suppose you are using 16 or 8 PHY cards and four connectors exist: C3, C2, C1, and C0. Each connector has four PHYs, and the autoport configuration is always enabled. Connector C3 has PHYs 0 to 3, Connector C2 has PHYs 4 to 7, Connector C1 has PHYs 8 to 11, and Connector C0 has PHYs 12 to 15.

- If you are connecting four different target devices and want to plug a cable into Connector 1, the SAS address for this port is 0x5000_0000_8000_0008 because the connector's first PHY is 8.
- Furthermore, when you plug a cable into Connector 0, the SAS address for this port is 0x5000_0000_8000_0009.
- Assuming nothing is connected to the HBA and you plug a cable into Connector 0, the SAS address assigned to this port is 0x5000_0000_8000_0008.
- Again, assuming nothing is connected to the HBA and you plug a cable into Connector 3, the SAS address assigned to this port is 0x5000_0000_8000_0000.
- Next, when a cable is plugged into Connector 2, the SAS address assigned to this port is 0x5000_0000_8000_0001.

This logic is also applicable for cards with eight PHYs.

Controllers have two SAS cores; each core can have a wide port, with at the most x8 connections. While connectors C0 and C1 can belong to one core, connectors C2 and C3 can belong to another core.

Multi-Selection Threshold for Virtual and Physical Drives

While selecting virtual drives or physical drives, you can only select up to 32 VDs or 32 PDs.

If you want to select more than 32 VDs or PDs, perform the following steps:

1. Navigate to the directory `\LSIStorageAuthority\server\html\files`
2. Open the `Configfile.json` file.
3. In the `Configfile.json` file, search for the `maxVDs` field.
The `maxVDs` field is set by default to a value of 32 to accept 32 VDs/PDs.
4. Modify this value to 64 .
5. Clear your browser's history.

Now, you can select more than 32 VDs or PDs, up to a maximum of 64 VDs or PDs.

Known Issues and Workarounds

The following is a list of known issues and workarounds.

- **Issue:** Light Weight Agent only allows two snapdump files to be downloaded at a time.
Workaround: None.
- **Issue:** An IR/IT firmware downgrade is not supported from one phase to another phase due to limitations in underlying layers.
Workaround: None.
- **Issue:** LSA does not detect all the controllers in a HyperV Environment when the controller passthrough is enabled or disabled.
Workaround: Restart the LSA services to reload and update the library.
- **Issue:** LSA may be inaccessible after a successful firmware update while IO's are occurring.
Workaround: Restart the LSA services.
- **Issue:** LSA may hang when downloading support logs on multiple clients.
Workaround: Restart the LSA services. Collect logs from one client at a time during non-heavy IO or drive/blackplane operations.

NOTE

Downloading the support log is available only for admin users.

- **Issue:** Allows the *Guest* user to log in when the *Guest* user is disabled through the **User Accounts**.
Workaround:
 1. Open the Command Prompt.
 2. Enter `lusrmgr.msc`.
 3. Select **Users**, then **Guest**.
 4. Right-click on the **Guest User**, and select the Properties option.
 5. Select the check box, **Account is Disabled**, if not already selected.
- **Issue:** The server response of IPv4 and IPv6 addresses groups are intermixed in the presence of multi NIC cards.
Workaround: None.
- **Issue:** When auto rebuild is enabled, multiclick PD actions are not updated properly.
Workaround: Manually refresh the page.
- **Issue:** Downgrading from 004.003.000.00 and higher is not recommend as the uninstall can be of the previous version can be erratic.
Workaround: Uninstall the latest version and reinstall an earlier version.
Version: 004.003.000.00 and higher
- **Issue:** Google Chrome may not position popup windows correctly.
Workaround: None.
Version: 61.0.3163.100 and later
- **Issue:** When using Mozilla FireFox, do not save the user name and password, or click the user name text box to enable saving.
Workaround: None.
- **Issue:** Updating or Erasing the UEFI or BIOS from any utility other than LSA, will result in LSA displaying old UEFI and BIOS information.
Workaround: Reboot, so LSA will display the correct BIOS details
Version: IT controllers only.
- **Issue:** The user cannot disable or modify the security, cannot delete the virtual drive, or clear the configuration.

- Workaround:** Clear the TR from DG and perform the respective operations.
- **Issue:** No matter the state of the controller or virtual drive, the state will be optimal and some operations may fail on TR DG/VD, which is beyond the scope of LSA.
Workaround: Clear the TR from DG and perform the respective operations.
 - **Issue:** Operations performed during an online controller reset fail.
Workaround: Do not perform any operation in LSA during an online controller reset.
 - **Issue:** Clicking between servers nodes in remote and manage discovery pages results in intermittent disappearance of servers or server icons.
Workaround: After clicking between server nodes wait for a few seconds for the screen to refresh and see all of the server nodes and icons.
Version: Mozilla and Windows Server 2012
 - **Issue:** Zoom operations.
Workaround: Do not zoom operations on a browser until the monitor resolution is low.
 - **Issue:** Performing any action (for example, Configuration) from the Server summary page, then manually refreshing the page will cause the user to be redirected to the initially selected Action page.
Workaround: Do not perform a manual refresh.
 - **Issue:** Converting a JBOD PD from JBOD to UG will cause applications to display different action menu names. LSA displays it as **Make unconfigured good**.
 - **Issue:** If the same dedicated hot spare is assigned to multiple drive groups, you may see inconsistency in the Element Count and DHSP Element selection check boxes on the Controller page.
 - **Issue:** Discovery of the Server through Manual Discovery option has the following limitations.
 - The Discovery server port should be 2463.
 - The Discovery server protocol should be HTTP.
 - **Issue:** In MegaRAID, when the patrol read is running at the physical drive level, it is a controller level operation. Each individual physical drive patrol read progress bar will not disappear after completing 100%.
Workaround: Wait for all of the physical drive progress bars to complete. Once all of the physical drive progress bars have reached 100%, they will disappear.
 - **Issue:** During installation or uninstallation, the publisher can show as unknown on the **User Account Control** message box.
 - **Issue:** If you select **Add the Virtual Drive(s)** from existing free space on a drive group or **Delete virtual drive(s)** from an existing drive group, LSA will refresh the complete controller page to update configuration information. Due to page refresh, the mouse reference is removed, so the page does not scroll up or down if you scroll the page using the mouse wheel.
Workaround: Click once anywhere on the page so the page scroll works. Alternatively, you can use the scrollbar.
 - **Issue:** LSA does not allow the physical drive to be selected from non-spanned virtual drives or spanned virtual drives.
 - **Issue:** The **Modify** option for the existing `setup.exe` will not work.
Workaround: Uninstall and reinstall the build instead of using the **Modify** option.

Glossary

This glossary defines the terms and acronyms used in this document.

access policy	A virtual drive property that indicates what type of access is allowed for a particular virtual drive. The possible values are <i>Read/Write</i> , <i>Read Only</i> , or <i>Blocked</i> .
APIPA	automatic Private IP addressing. A Windows-based operating system feature that enables a computer to automatically assign itself an IP address when no DHCP server is available to perform that function.
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.
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 read operations to see whether 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 operation 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 that indicates 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 4196 MB, and a 4-GB from another manufacturer might be 4128 MB. These drives could be coerced to a usable capacity of 4088 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.
current	Measure of the current flowing to (+) or from (–) the energy pack, 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. <ul style="list-style-type: none"> • In Write Back mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. • In Write Through mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction.
device ID	A controller or drive property indicating the manufacturer-assigned device ID.
DHCP	dynamic host configuration protocol
drive group	A group of drives attached to a RAID controller on which one or more virtual drives can be created. All virtual drives in the drive group use all of the drives in the drive group.

drive state	<p>A physical drive or a virtual drive property indicating the status of the appropriate drive.</p> <p>Physical Drive State</p> <p>A physical drive can be in any one of the following states:</p> <ul style="list-style-type: none"> • Unconfigured Good – A drive accessible to the RAID controller but not configured as a part of a virtual drive or as a hot spare. • 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. <p>An Unconfigured Good or Offline drive that has completed the prepare for removal operation.</p> <p>Virtual Drive State</p> <p>A virtual drive can be in any one of the following states:</p> <ul style="list-style-type: none"> • Optimal – A virtual drive whose members are all online. • Partially Degraded – A virtual drive with a redundant RAID level that is capable of sustaining more than one member drive failure. <p>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.</p> <ul style="list-style-type: none"> • 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.
energy pack	Refers to a battery backup unit or a CacheVault.
fast initialization	The firmware quickly writes zeros to the first and last 8-MB regions of the new virtual drive, and then completes the initialization in the background or with next scheduled Consistency Check. This allows you to start writing data to the virtual drive immediately.
fault tolerance	The capability of the drive subsystem to undergo a single drive failure per drive group without compromising data integrity and processing capability. SAS RAID controllers provide fault tolerance through redundant drive groups in RAID levels 1, 5, 6, 10, 50, and 60. They also support hot spare drives and the auto-rebuild feature.
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 lets you 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 performed only if a drive generates many media errors.
GUI	graphical user interface.
GT/s	giga transfers per second.
HDD	hard disk drive.

hot spare	<p>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.</p> <p>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.</p>
HPC	high performance computing.
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 read operations are buffered in cache memory. In Direct I/O mode, read operations 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 read operations on a specific virtual drive. It does not affect the read-ahead cache.)
learning cycle	An energy pack calibration operation performed by a RAID controller periodically to determine the condition of the energy pack. You can start energy pack 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 maximizes resource use, throughput, or response time.
manufacturing date	Date on which the energy pack assembly was manufactured.
manufacturing name	Device code that indicates the manufacturer of the components used to make the energy pack 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.
NIC	network interface card.
NAT	network address translator.
NVMe	nonvolatile memory. NVMe is a logical device interface specification for accessing nonvolatile storage media attached using the PCI Express (PCIe) bus.
offline	A drive is offline when it is part of a virtual drive but its data is not accessible to the virtual drive.
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.
PCI	Peripheral Component Interface.
PCIe	PCI Express.

RAID	<p>redundant array of independent disks.</p> <p>A group of multiple, independent drives that provide high performance by increasing the number of drives used for saving and accessing data.</p> <p>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.</p>
RAID 0	Uses data striping on two or more drives to provide high data throughput, especially for large files in an environment that requires no data redundancy.
RAID 1	Uses data mirroring on pairs of drives so that data written to one drive is simultaneously written to the other drive. RAID 1 works well for small databases or other small applications that require complete data redundancy.
RAID 5	Uses data striping and parity data across three or more drives (distributed parity) to provide high data throughput and data redundancy, especially for applications that require random access.
RAID 6	Uses data striping and parity data across three or more drives (distributed parity) to provide high data throughput and data redundancy, especially for applications that require random access. RAID 6 can survive the failure of two drives.
RAID 10	A combination of RAID 0 and RAID 1 that uses data striping across two mirrored drive groups. 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	<p>A virtual drive property indicating the RAID level of the virtual drive.</p> <p>Broadcom SAS RAID controllers support RAID levels 0, 1, 5, 6, 10, 50, and 60.</p>
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 that indicates 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 the requested data and stores the additional data in cache memory, anticipating that the data will be needed soon. This speeds up read operations for sequential data, but you will see 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	<p>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.</p> <p>A redundant configuration protects the data in case a drive fails in the configuration.</p>

replace	<p>The procedure used to replace data from a source drive of a virtual drive to a destination drive that is not a part of the virtual drive. The replace 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 replace operation can be run automatically or manually.</p> <p>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 replace operation runs as a background activity, and the virtual drive is still available online to the host.</p>
SAS	Serial-Attached SCSI. SAS is a serial, point-to-point, enterprise-level device interface that leverages the 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.
SASL	simple authentication and security layer.
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	Small Computer System Interface.
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.
SSD	solid state drive.
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	<p>A technique used to write data across all drives in a virtual drive.</p> <p>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.</p>
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.
temperature	Temperature of the energy pack, measured in Celsius.
URI	uniform resource identifier.
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.
VIPA	virtual IP address. An IP address assigned to multiple domain names or servers that share an IP address based on a single network interface card.
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 <i>Optimal</i> and <i>Degraded</i> .
write-back	<p>In Write-Back Caching mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a drive write transaction. Data is written to the drive subsystem in accordance with policies set up by the controller.</p> <p>These policies include the amount of dirty/clean cache lines, the number of cache lines available, and elapsed time from the last cache flush.</p>
write policy	See <i>Default Write Policy</i> .
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.

Revision History

Version 2.2, November 11, 2021

The following changes were made:

- Updated [Support Matrix](#).
- Updated [LSI Storage Authority Features](#).
- Updated [Types of Installation](#).
- Removed Managing LSA on the VMware ESXi Operation System chapter.
- Updated [Server Dashboard](#).
- Updated [Controller Dashboard](#).
- Removed UNMAP Capability Feature section and subsections.
- Updated [Personality Management](#).
- Updated [Changing Behavior Modes](#) section.
- Removed Changing Profiles section.
- Updated [Background Operations Support](#).
- Updated [Managing Controllers](#).
- Removed Downloading the TTY Log section.
- Updated [Advanced Software](#).
- Removed Using the CacheCade Pro 2.0 Feature section and subsections.
- Updated [Managing Drive Groups](#).
- Removed RAID Level Migration section and subsections.
- Updated [Managing Virtual Drives](#).
- Removed Hiding and Unhiding a Virtual Drive or a Drive Group section and subsections.
- Updated [Managing Physical Drives](#).
- Removed Sanitizing a Drive section.
- Updated [Managing Hardware Components](#).
- Removed Setting Learn Cycle Properties section.
- Removed Starting a Learn Cycle Manually section.
- Removed Viewing Tape Drives section.
- Updated [Show Events](#).
- Updated [Supported Operating Systems](#) for Light Weight Monitor System.
- Removed Installing the Installing the Light Weight Monitor on the Windows Operating System through the LSA Stand-Alone Installation section.
- Removed Installing Light Weight Monitor on the Oracle Solaris x86 Operating System section.
- Removed Uninstalling the Light Weight Monitor on the Oracle Solaris x86 Operating System section.
- Updated [Configuring the Light Weight Monitor System](#).
- Removed Setting Up the Custom Facility Level in the System Log File for the Solaris x86 Operating System section.
- Updated [Introduction to RAID](#).
- Updated [Events and Messages](#).
- Removed Event Messages section.
- Updated the screenshots throughout the guide.
- Minor rewrites for clarity and consistency.

Version 2.1, September 22, 2021

The following changes were made:

- Updated [Changing Profiles](#).
- Updated [Enabling the Schedule Panel Feature](#).
- Updated [Setting Up the Email Server](#).
- Updated [Adding Email Addresses of Recipients of Alert Notifications](#).
- Minor rewrites for clarity and consistency.

Version 2.0, June 17, 2021

The following changes were made:

- Updated [Support Matrix](#).
- Updated [Limitations of Installation and Configuration](#).
- Updated [Changing Behavior Modes](#).
- Updated [Viewing Controller Properties](#).
- Updated [Updating the Controller Firmware](#).
- Updated [Changing Drive Security Settings](#).
- Changed **Copyback** operation name to **Replace**.
- Updated [Event Messages](#).
- Updated [Known Issues and Workarounds](#).
- Minor rewrites for clarity and consistency.

Version 1.16, March 5, 2021

The following changes were made:

- Updated [Managing LSA on the VMware ESXi Operating Systems](#).
- Added [Replacing a Missing Drive](#).
- Minor rewrites for clarity and consistency.

Version 1.15, October 30, 2020

The following changes were made:

- Updated [Creating a New Storage Configuration Using the Advanced Configuration Option](#).
- Updated [Server Dashboard](#).
- Updated [Managing LSA on the VMware ESXi Operating Systems](#).
- Updated [Known Issues and Workarounds](#).
- Minor rewrites for clarity and consistency.

Version 1.14, July 24, 2020

The following changes were made:

- Updated [Support Matrix](#).
- Updated [Configuring HTTPS](#).
- Updated [Resetting Encryption Keys](#).
- Updated [Installing in the Interactive Mode](#).
- Updated [Upgrading and Downgrading the Firmware on IT Controllers \(VMware\)](#).
- Updated [Updating the Controller Firmware](#).
- Updated [Changing Drive Security Settings](#)
- Added [Known Issues and Workarounds](#).
- Minor rewrites for clarity and consistency.

Version 1.13, December 23, 2019

The following changes were made:

- Updated [Support Matrix](#).
- Updated [Scheduling a Patrol Read](#).
- Updated [Preinstallation Requirements](#)
- Minor rewrites for clarity and consistency.

Version 1.12, September 17, 2019

The following changes were made:

- Added [Snapdump Feature](#).
- Added [UNMAP Capability Feature](#).
- Replaced the screen captures that follow.
 - [Figure 12, Alert Settings Window](#)
 - [Figure 15, Server Dashboard](#)
 - [Figure 16, View Server Profile Window](#)
 - [Figure 17, Controller Dashboard](#)
 - [Figure 28, Background Processes in Progress Window](#)
 - [Figure 52, Virtual Drive Properties Window](#)
 - [Figure 63, Energy Pack Properties Window](#)
 - [Figure 64, Enclosure Properties](#)
- Minor rewrites for clarity and consistency.

Version 1.11, May 30, 2018

The following changes were made:

- Removed Private IP Range Restrictions.
- Updated [Managing Servers from the Remote Server Discovery Page](#).
- Added [Manually Discovering Servers](#).
- Added [Marking a Drive as a Missing Drive](#).

Version 1.10, February 16, 2018

The following changes were made:

- Added a procedure of *Optimizing Event Notifications* to Section 8, Managing LSA on VMware ESXi Operating Systems.
- Updated [Creating a New Storage Configuration Using the Advanced Configuration Option](#).
- Updated [Adding Physical Drives](#).
- Updated [Adding Hot Spares to the Existing Drive Group](#).
- Updated [Viewing Controller Properties](#).
- Updated [Scheduling a Consistency Check](#).
- Added Section 18.9, Behavior of Virtual Drive Operations on VMware.
- Updated Section 19.9, Assigning Dedicated Hot Spares.

Version 1.9, November 30, 2017

The following changes were made:

- Updated support for Tape Drive in Section 19.2, Tape Drives.
- Updated images.

Version 1.8, September 11, 2017

The following changes were made:

- Merged the *Light Weight Monitor (LWM)* User Guide into this document from Chapter 16 through Chapter 18.
- Updated Section 2.1, Support Matrix.
- Updated Section 6, Installing the LSI Storage Authority Software on the Microsoft Windows Operating System.
- Added Section 9.4, Hiding an Empty Backplane.
- Updated [Updating the Controller Firmware](#).

Version 1.7, June 21, 2017

The following changes were made:

- Added Section 8.1.1, Increasing the Memory Limit of Host Hardware RAID Controller.
- Updated Section 8.3, Configuring the Firewall on Various LSA Installers
- Updated Section 8.4, Collecting LSA Logs on Windows and Linux Operating Systems.
- Updated Section 8.5, Collecting LSA Logs on VMware Operating Systems.
- Updated Section 8.6, Logout and Reboot Requirements on VMware.
- Updated Section 8.7, Behavior of Event History.
- Updated Section 8.8, Behavior of Event Monitoring on Non-ESXi versus ESXi Server.
- Updated Section 10.2, Displaying or Blocking a Private IP Address.

Version 1.6, March 24, 2017

The following changes are made:

- Added NVMe support.
- Updated Section 6, Installing the LSI Storage Authority Software on the Microsoft Windows Operating System.
- Added [Managing PCIe Storage Lane Speed](#) .
- Updated [Glossary](#).

Version 1.5, November 18, 2016

The following change was made:

- Updated Section 8, Managing LSA on VMware ESXi Operating Systems.

Version 1.4, October 6, 2016

The following changes are made:

- Updated [LSI Storage Authority Overview](#).
- Updated Section 2.1, Support Matrix.
- Updated [LSI Storage Authority Features](#).
- Updated the descriptions in [Server Dashboard](#).
- Added [Personality Management](#).
- Updated [Table 8, Basic and Advanced Controller Properties](#).
- Updated [Enabling Drive Security](#).
- Updated [Changing Drive Security Settings](#).
- Updated [Viewing Physical Drive Properties](#).
- Updated [Viewing Enclosure Properties](#).
- Added [Appendix A](#) and [HTTP Status Codes and Description](#).

Version 1.3, November 2015

Added the following sections:

- LSI Storage Authority Feature Comparison Matrix
- Accessing LSA Over Network Address Translation (NAT)
- Installing the LSI Storage Authority Software on the Linux Operating System
- Viewing Controller Properties
- Viewing Drive Group Properties
- Viewing Virtual Drive Properties
- Viewing Physical Drive Properties
- Viewing Energy Pack Properties
- Monitoring Enclosures
- Viewing Enclosure Properties
- Customizing the Theme of the LSI Storage Authority Software

Version 1.2, July 2015

Added the following sections:

- Using LDAP Authentication
- Managing Servers from the Server Discovery Page
- Adding Managed Servers
- Removing Managed Hosts
- Managing Power Save Settings
- Enabling and Disabling SSD Guard
- Discarding Pinned Cache
- Creating Simple Configuration
- Using the MegaRAID CacheCade Pro 2.0 Feature
- Creating a CacheCade Virtual Drive
- Modifying CacheCade Virtual Drive Properties
- Enabling SSD Caching on a Virtual Drive
- Disabling SSD Caching on a Virtual Drive
- Enabling or Disabling SSD Caching on Multiple Virtual Drives
- Clearing Configuration on CacheCade Virtual Drives
- Removing Blocked Access
- Deleting a Virtual Drive with SSD Caching Enabled
- Fast Path Advanced Software
- MegaRAID SafeStore Encryption Services
- Enabling Drive Security
- Changing Security Settings
- Import or Clear a Foreign Configuration

Version 1.1, April 2015

Added the following chapters:

- Background Operations Support
- Managing Controllers
- MegaRAID Advanced Software Options
- MegaRAID Advanced Software
- Managing Drive Groups
- Managing Virtual Drives
- Managing Physical Drives
- Monitoring Energy Packs

Version 1.0, November 2014

Initial document release.

