



Hewlett Packard
Enterprise

HPE MR Storage Administrator User Guide

Abstract

This document includes feature, installation, and configuration information for HPE MR Storage Administrator users. It is intended for users with a good working knowledge of storage hardware and configuration of logical drives and arrays. Hewlett Packard Enterprise assumes you are qualified in the servicing of computer equipment and trained in recognizing hazards in products with hazardous energy levels.

For a comprehensive list of changes to this document, see the [Revision History](#).

© Copyright 2017-2026 Hewlett Packard Enterprise Development LP.

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein. Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website. Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

MegaRAID®, CacheCade™, FASTPATH®, and SafeStore™ are among the trademarks of Broadcom in the United States, the EU, and/or other countries. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.

Table of Contents

HPE MR Storage Administrator Application Overview.....	9
Support Matrix.....	10
HPE MR Storage Administrator Feature Support Matrix.....	12
Upgrade Requirements.....	16
Browser Cache.....	16
StandAlone Installer.....	17
Installing the HPE MR Storage Administrator Software on the Microsoft Windows Operating System.....	18
Installing in Noninteractive Mode.....	22
Uninstalling in Interactive Mode.....	22
Uninstalling in Noninteractive Mode.....	22
Installing the HPE MR Storage Administrator Software on the Linux Operating System.....	24
Installing in the Interactive Mode.....	24
Installing in the Noninteractive Mode.....	25
Uninstalling the LSI Storage Authority Software on the Linux Operating System.....	26
Performing Initial Configuration.....	27
Changing the HPE MR Storage Administrator Software Port Number.....	27
Changing the Nginx Web Server Port Numbers.....	27
Changing the Nginx Read Timeout.....	27
Performing the Initial Setup.....	29
Displaying or Blocking a Private IP Address.....	29
Alert Settings.....	30
Setting Up the Email Server.....	31
Adding the Email Addresses of Alert Notification Recipients.....	32
Configuring Different Types of Access.....	33
Server Dashboard.....	34
Controller Dashboard.....	36
Controller Configurations.....	37
Creating a New Storage Configuration Using the Simple Configuration Option.....	37
Creating a Storage Configuration Using the Advanced Configuration Option.....	38
Selecting Available Unconfigured Drives.....	40
Selecting Volume Settings.....	41
Clearing the Configuration.....	42
Importing or Clearing Foreign Configurations.....	42
UNMAP Capability Feature.....	42

UNMAP Capability Feature Behavior.....	43
UNMAP Feature Support.....	43
Personality Management.....	44
Function Profile Management.....	45
Auto Assign Policy.....	46
Background Operations Support.....	48
Managing Controllers.....	49
Viewing Controller Properties.....	49
Running Consistency Checks.....	49
Set Consistency Check Properties.....	50
Schedule Consistency Check Operation.....	50
Manage SID Ownership.....	51
Device Reporting Order.....	52
Running a Patrol Read Operation.....	52
Setting the Patrol Read Properties.....	52
Starting a Patrol Read Operation.....	53
Stopping a Patrol Read Operation.....	53
Managing SAS Storage Link Speed.....	53
Managing PCIe Storage Interface.....	56
Set Adjustable Task Rate.....	57
Setting the Task Information.....	59
Managing Power-Save Settings.....	60
Discarding Pinned Cache.....	61
Spin Down Drives at Shutdown.....	62
NVMe Thermal Poll Interval.....	62
Download Serial Output Log.....	63
Updating the Controller Firmware.....	63
Firmware Activation Status.....	64
Factory Repurpose.....	64
Factory Defaults.....	64
MegaRAID Advanced Software Features.....	66
Fast Path Advanced Software.....	66
SafeStore Encryption Services.....	66
Enable Security.....	67
Changing Drive Security Settings.....	69
Disabling Drive Security.....	71
Importing or Clearing a Foreign Configuration – Security-Enabled Drives.....	71
Managing Arrays.....	72
Viewing Array Properties.....	72

Adding a Volume to an Array.....	72
RAID Level Transformation.....	73
Migrating the RAID Level of an Array.....	73
Adding Drives to a Configuration.....	74
Removing Drives from a Configuration.....	75
Migrating the RAID Level Without Adding or Removing Drives.....	75
Managing Volumes.....	76
Viewing Volume Properties.....	76
Modifying Volume Properties.....	77
Start and Stop Locating a Volume.....	78
Erasing a Volume.....	79
Initializing a Volume.....	80
Starting Consistency Check on a Volume.....	80
Expanding the Online Capacity of a Volume.....	81
Expanding the Online Capacity of a Volume for MR200/MR4000 Controllers.....	81
Expanding the Online Capacity of a Volume for MR932 Controllers.....	81
Deleting a Volume.....	83
Managing Drives.....	84
Viewing Drive Properties.....	84
Locating Tape Drives.....	86
Start and Stop Locating a Drive.....	86
Making a Drive Offline.....	86
Making a Drive Online.....	86
Replacing a Drive.....	87
Marking a Drive as a Missing Drive.....	88
Replacing a Missing Drive.....	90
Viewing Protected Arrays.....	90
Assigning Global Spare Drives.....	91
Removing a Global Spare Drive.....	91
Assigning Dedicated Spare Drives.....	91
Rebuilding a Drive.....	92
Converting an Unconfigured Bad Drive to an Unconfigured Good Drive.....	92
Removing a Drive.....	92
Make Unconfigured Good Drives and Make JBOD Drives.....	92
Making Unconfigured Good Drives.....	93
Making a JBOD Drive.....	93
Erasing a Drive.....	93
Erasing a Drive Securely.....	94
Sanitizing a Drive.....	95

Managing Hardware Components.....	98
Monitoring the HPE Smart Storage Energy Pack.....	98
Monitoring Enclosures.....	98
Viewing Enclosure Properties.....	99
Physical Function Information.....	99
Operations on Physical Functions.....	100
NVMe Drive Format.....	102
Viewing Event Logs.....	104
Downloading Logs.....	104
Clearing the Event Logs.....	104
Known Issues and Workarounds.....	105
Appendix A: Support and Other Resources.....	107
Accessing Hewlett Packard Enterprise Support.....	107
Accessing Updates.....	107
Customer Self Repair.....	108
Remote Support.....	108
Warranty Information.....	108
Regulatory Information.....	109
Documentation Feedback.....	109
Appendix B: Glossary.....	110
Appendix C: Revision History.....	115

HPE MR Storage Administrator Application Overview

The HPE MR Storage Administrator application is a web-based application that lets you monitor, maintain, troubleshoot, and configure MegaRAID products. The HPE MR Storage Administrator graphical user interface (GUI) helps you to view, create, and manage storage configurations.

- **Monitoring and Configuring:** The HPE MR Storage Administrator application lets you monitor the controllers and configure the drives on the controller.
The application displays the status of the controller cards, volumes, and drives on the controller. The device status icons are displayed on their respective pages to notify you when there is 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 the HPE MR Storage Administrator application, you can perform system maintenance tasks, such as updating the controller firmware.
- **Troubleshooting:** The HPE MR Storage Administrator application displays information that is related to drive failures, device failures, and so on.

The application also provides recommendations and displays contextual links, helping you to easily locate drives and devices that have issues and troubleshoot them. In addition, you can download a complete report of all the devices and their configurations, properties, and settings and send it to the support teams for further analysis and troubleshooting.

ATTENTION

The HPE MR Storage Administrator application is a unified application supporting Gen10, Gen10 Plus, Gen11, and Gen12 controllers. This document covers the operations and properties of the listed controllers and provides information for the operation and properties that are supported by each controller.

Support Matrix

The following table provides the support requirements for the HPE MR Storage Administrator application.

Table 1: Hardware and Software Support Matrix

Support	Version/Flavors
Supported Controllers	• HPE MR932i-p Gen12 • HPE MR416i-p Gen11 • HPE MR416i-o Gen11 • HPE MR408i-o Gen11 • HPE MR408i-9 Gen11 • HPE MR216i-p Gen11 • HPE MR216i-o Gen11 • HPE MR416i-p Gen 10+ • HPE MR416i-a Gen 10+ • HPE MR216i-p Gen 10+ • HPE MR216i-a Gen 10+ • HPE Smart Array P824i-p MegaRAID Gen10 Controller
Supported operating systems	
Microsoft	• Microsoft Hyper-V Server 2022 • Microsoft Hyper-V Server 2019 • Microsoft Hyper-V Server 2016 • Microsoft Windows Server 2025 • Microsoft Windows Server 2019 • Microsoft Windows Server 2016 (Datacenter) • Microsoft Windows Server 2016 (Standard) • Microsoft Windows Server 2016 (Essentials) • Microsoft Hyper-V Server 2012 R2 • Microsoft Windows Server 2016 RS3 • Microsoft Windows Server 2012 R2 (Foundation) • Microsoft Windows Server 2012 R2 (Essentials) • Microsoft Windows Server 2012 R2 (Standard) • Microsoft Windows Server 2012 R2 (Datacenter)

Support	Version/Flavors
Linux	• Red Hat Enterprise Linux 10 • Red Hat Enterprise Linux 9.6 • Red Hat Enterprise Linux 9.5 • Red Hat Enterprise Linux 9.2 • Red Hat Enterprise Linux 9.1 • Red Hat Enterprise Linux 9 • Red Hat Enterprise Linux 8.7 • Red Hat Enterprise Linux 8.6 • Red Hat Enterprise Linux 8.5 • Red Hat Enterprise Linux 8.4 • Red Hat Enterprise Linux 8.3 • Red Hat Enterprise Linux 8.2 • Red Hat Enterprise Linux 8.1 • Red Hat Enterprise Linux 7.8 • Red Hat Enterprise Linux 7.8 • Red Hat Enterprise Linux 7.7 • SUSE Linux Enterprise Server 15 SP7 • SUSE Linux Enterprise Server 15 SP6 • SUSE Linux Enterprise Server 15 SP5 • SUSE Linux Enterprise Server 15 SP4 • SUSE Linux Enterprise Server 15 SP3 • SUSE Linux Enterprise Server 15 SP2 • SUSE Linux Enterprise Server 15 SP1 • SUSE Linux Enterprise Server 12 SP5 • SUSE Linux Enterprise Server 12 SP4 • SUSE Linux Enterprise Server 12 SP3
Ubuntu	• Ubuntu 22.04 LTS • Ubuntu 20.04 LTS • Ubuntu 18.04 LTS
openEuler	• openEuler 20.03 LTS
Supported web browsers	• Windows Internet Explorer 9.0 and later • Mozilla Firefox version 9.0 and later • Google Chrome version 16.0 and later
Supported networks	• Internet Protocol versions 4 and 6 • Network Address Translation • Domain • HTTP, HTTPS

HPE MR Storage Administrator Feature Support Matrix

The following tables outline the feature support for HPE Smart Array MegaRAID controllers with respect to software features and firmware features.

Table 2: MegaRAID Firmware Feature Support Matrix

Feature Name	MegaRAID Firmware
RAID level	RAID 0, RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, and RAID 60
Maximum drives	240
Maximum spans	8
Maximum volumes	240
Maximum media errors	256
Drive-mixing support	No.
Strip size support	64 KB, 128 KB, 256 KB, 512 KB, and 1024 KB
Maximum volumes per array	64
Multipath	No
Controller reset support	Yes

Table 3: Software Feature Support

Feature Name	Description
Server Dashboard	The Server Dashboard is the default landing page in the application. 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. See Server Dashboard for more information.
Controller Dashboard	The Controller Dashboard lets you perform controller related actions and view all the information pertaining to a controller. See Controller Dashboard for more information
Simple Configuration	The Simple Configuration option 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. See Creating a New Storage Configuration Using the Simple Configuration Option for more information.
Advanced Configuration	The Advanced Configuration option provides an easy way to create a new storage configuration. The Advanced Configuration option gives you greater flexibility than Simple Configuration because you can select the physical drive and volume parameters when you create a volume. In addition, you can use the Advanced Configuration option to create spanned arrays (parity groups). See Creating a Storage Configuration Using the Advanced Configuration Option for more information.
Foreign Configuration (Import/Clear)	A <i>foreign configuration</i> is a RAID configuration that already exists on a replacement set of drives that you install in a storage system. You can use the application to import the foreign configuration to the controller or clear the foreign configuration so that you can create a new configuration using these drives. See Importing or Clearing Foreign Configurations for more information.
Clear Configuration	The Clear Configuration feature lets you clear all existing configurations on a selected controller. See Clearing the Configuration for more information.

Feature Name	Description
Update Firmware	The Update Firmware feature lets you update the controller firmware. See Updating the Controller Firmware for more information.
Online Firmware Update	The Online Firmware Update feature lets you update the controller firmware. See Updating the Controller Firmware for more information.
Controller Operations	
Setting Consistency Check Properties	The Consistency Check operation verifies the correctness of the data in volumes that use RAID levels 1, 5, 6, 10, 50, and 60, configurations. For example, in a system with parity, checking the consistency means calculating the data on one drive and comparing the results to the contents of the parity drive. See Running Consistency Checks for more information.
Scheduling Consistency Check	The Scheduling Consistency Check feature lets you periodically run a consistency check on fault-tolerant volumes. See Schedule Consistency Check Operation for more information.
Setting Patrol Read Properties	A Patrol Read operation 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 operation for all RAID levels and for all spare drives. A Patrol Read operation is initiated only when the controller is idle for a defined period and has no other background activities. See Setting the Patrol Read Properties for more information.
Starting Patrol Read	A Starting Patrol Read operation lets you start a patrol read operation. See Starting a Patrol Read Operation for more information.
Stopping Patrol Read	A Stopping Patrol Read operation lets you stop an already started patrol read operation. See Stopping a Patrol Read Operation for more information.
Managing Link Speed	A Managing Link Speed operation 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. See Managing SAS Storage Link Speed for more information.
Setting Adjustable Task Rates	A Setting Adjustable Task Rates operation lets you change the Rebuild Rate, Transformation Rate, Patrol Read Rate, BGI Rate, and Consistency Check Rate for a controller. See Set Adjustable Task Rate for more information.
Discarding Preserved Cache	If the controller loses access to one or more volumes, the controller preserves the data from the volume. This preserved cache is called <i>Pinned Cache</i> . This cache is preserved until you import the volume or discard the cache. As long as pinned cache exists, you cannot perform certain operations on the volume. See Discarding Pinned Cache for more information.
Downloading Serial Output Log	The Serial Output Log file contains the firmware terminal log entries for the controller. The log information is shown as total number of entries available on the firmware side. See Download Serial Output Log for more information.
Background Operations	Provides information on Background Operations Support, such as Pause, Resume, Abort, and so on. See Background Operations Support for more information.
Advanced Software Features	
Fast Path	The MegaRAID FastPath software is a high-performance I/O accelerator for solid state drive (SSD) arrays connected to a MegaRAID controller card. This advanced software is an optimized version of MegaRAID technology that can dramatically boost storage subsystem and overall application performance; particularly those that demonstrate high random read/write operation workloads – when deployed with a MegaRAID SATA+SAS controllers connected to SSDs. See Fast Path Advanced Software for more information.
RAID 5 and RAID 6	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.

Feature Name	Description
Volume Operations	
Volume Settings/Modifying Volume Properties	A Volume Settings/Modifying Volume Properties operation lets you configure the volumes. See Selecting Volume Settings for more information.
Start and Stop Locating a Volume	If the drives reside in a disk enclosure, you can identify them by making their LEDs blink. See Start and Stop Locating a Drive for more information.
Erasing a Volume	An Erasing a Volume operation lets you erase data on Non SEDs (normal HDDs) using the Drive Erase option. The Erase operation is performed as a background task. See Erasing a Drive for more information.
Initializing a Volume	An Initializing a Volume operation lets you select the Fast Initialization or Full Initialization option to initialize a drive immediately under the Advanced Configuration wizard. See Initializing a Volume for more information.
Starting Consistency Check on a Volume	A Consistency Check operation verifies whether all stripes in a volume with a redundant RAID level have correct parity or mirror values. The Consistency Check operation involves mirroring data when an inconsistent stripe is detected for a RAID 1 configuration, and re-creating the parity from the peer disks in the case of a RAID 5 and RAID 6 configuration. This mechanism applies to variants and secondary RAID levels based on RAID 1 and RAID 5 configurations. See Starting Consistency Check on a Volume for more information.
Expanding the Online Capacity of a Volume	The transformation feature lets you expand the capacity of a volume by adding new drives or making use of unused space on existing disks, without requiring a reboot. See Expanding the Online Capacity of a Volume for MR200/MR4000 Controllers for more information.
Deleting a Volume	The Deleting a Volume feature lets you delete a volume. See Deleting a Volume for more information.
Drive Operations	
Assign Global Spare Drives	A global spare drive replaces a failed drive in any redundant array, as long as the capacity of the global spare drive is equal to or greater than the coerced capacity of the failed drive. See Assigning Global Spare Drives for more information.
Remove Global Spare Drives	A Remove Global Spare Drives operation lets you remove global spare drives. See Removing a Global Spare Drive for more information.
Assign Dedicated Spare Drives	A dedicated spare drive provides protection to one or more specified arrays on the controller. See Assigning Dedicated Spare Drives for more information.
Start and Stop Locating Drive	If the drives are in a disk enclosure, you can identify them by making their LEDs blink. See Start and Stop Locating a Drive for more information.
Making a Drive Online and Offline	The Making a Drive Online and Offline feature lets you change the state of a drive. See Making a Drive Offline and Making a Drive Online for more information.
Replacing a Drive	The Replacing a Drive feature lets you replace a drive if the drive shows signs of failing. See Replacing a Drive for more information.
Rebuilding a Drive	If a drive that is configured as RAID 1, 5, 6, 10, 50, or 60 fails, the firmware automatically rebuilds the data on a spare drive to prevent data loss. The Rebuild operation is a fully automatic process. You can monitor the progress of drive rebuilds in the Background Processes in Progress window. See Rebuilding a Drive for more information.
Erasing a Drive	The Erasing a Drive feature lets you erase data on Non SEDs (normal HDDs). The Erase operation is performed as a background task. See Erasing a Drive for more information.
Sanitizing a Drive	The Sanitizing a Drive feature lets you erase the data that resides on a drive using the Sanitize feature. The Sanitize feature is similar to the Drive Erase feature that is already supported by your controller, except that the Sanitize function is performed by the drive firmware, whereas the Drive Erase function is performed by the controller firmware. See Sanitizing a Drive for more information.

Feature Name	Description
Converting Unconfigured Bad Drive to Unconfigured Good Drive	When you force a drive offline, it enters the Unconfigured Bad state. If a drive contains valid disk data format (DDF) metadata, its drive state is Unconfigured Good. See Converting an Unconfigured Bad Drive to an Unconfigured Good Drive for more information.
Make Unconfigured Good Drive	When you power down a controller and insert a new drive, and if the inserted drive does not contain valid DDF metadata, the drive status is listed as <i>JBOD</i> (Just a Bunch of Drives) when you power up the system again. When you power down a controller and insert a new drive, and if the drive contains valid DDF metadata, its drive state is listed as Unconfigured Good. A new drive in the JBOD drive state is exposed to the host operating system as a stand-alone drive. See Make Unconfigured Good Drives and Make JBOD Drives for more information.
Make JBOD	The Make JBOD feature lets you create JBODs. See Making a JBOD Drive for more information.
Event Logs	
Viewing Event Logs	The application monitors the activity and performance of the server and all of the controllers attached to it. See Viewing Event Logs for more information.

Upgrade Requirements

Complete the tasks that follow while upgrading the HPE MR Storage Administrator application.

Browser Cache

Clear the browser cache on the client on which you are using MRSA for the following reasons:

- If you are upgrading from a previous version of MRSA.
- To clear the saved passwords, so passwords are not saved in the browser.

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.

Installing the HPE MR Storage Administrator Software on the Microsoft Windows Operating System

This section provides the procedure for installing HPE MR Storage Administrator Software on the Microsoft Windows operating system.

NOTE


The `LSA_HOME` directory is only accessible by administrators.

Perform the following steps to install the HPE MR Storage Administrator .

1. Run the HPE MR Storage Administrator `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

HPE MR Storage Administrator v 008.006.018.000 - InstallShield Wizard

Port Configuration Settings

Provide Port configuration settings.

Web Server Port

HPE MR Storage Administrator Server Port

InstallShield


< Back Next > Cancel

By default, MRSA communicates on **Web Server Port 2463** and **LSA Server Port 9000**. Ensure that these ports are available to be used by MRSA. Depending on your environment, if these ports are not available, specify the port details here. You can also edit the details of this port after installation. See [Changing the HPE MR Storage Administrator Software 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 are part of the event history. If the sequence numbers are less than the last log that was cleared (**Since Log Clear** option), MRSA 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 are not 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 are not part of the event history. If the sequence numbers are less than the last log that was cleared (**Since Log Clear** option), MRSA 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 **Ready to Install the Program** windows appears. Click **Next**.

Depending on the setup type you have selected, the **InstallShield Wizard Completed** dialog appears.

9. (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.
10. Ensure that port 2463 is not blocked by a firewall.
The Windows Firewall settings are located under **Control Panel > Windows Firewall**.
11. 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 MRSA 2.2 and later. While installing MRSA, if OpenSLP is not installed ensure that you select the option to install OpenSLP, and MRSA seamlessly installs the required version of OpenSLP.

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 MRSA.

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

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

Before you begin to uninstall MRSA, 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 HPE MR Storage Administrator Software on the Linux Operating System

The HPE MR Storage Administrator software supports both the interactive and the noninteractive modes of Linux installation.

1. Run the `rpm -ivh MRStorageAdministrator-xx.rpm` command from the installation disk.
2. Extract the contents of the zip file and install the appropriate package on the operating system.

The following is the corresponding OS support information:

- `gcc_4.8.x` — RHEL 7.x and supported operating systems of `gcc_8.3.x` and `gcc_11.2.x`
- `gcc_8.3.x` — RHEL 8.x, SLES 15 SPx, and supported operating systems of `gcc_11.2.x`
- `gcc_11.2.x` — RHEL 9.x, RHEL 10, Ubuntu 20.04 LTS, Ubuntu 22.04 LTS
- `gcc_4.8.x` — MRSA Installer in this folder supports `gcc_4.8.x` and higher
- `gcc_8.3.x` — MRSA Installer in this folder supports `gcc_8.3.x` and higher
- `gcc_11.2.x` — MRSA Installer in this folder supports `gcc_11.2.x` and higher

64-bit Linux operating systems

This package contains `gcc_4.8.x` , `gcc_8.3.x` and `gcc_11.2.x`.

Ubuntu

Use `install_deb.sh` to install the HPE MR Storage Administrator .

NOTE

Ensure the `Connect automatically` checkbox is selected. The checkbox is available under **Network Connections**.

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 = 2` (`retrieve_range_of_events_since = 2`). You must restart the LSI Storage Authority service for the configuration changes to take effect.

4. Enter the nginx server port number. The port range is from 1 to 65535. The default port number is 2463.
5. 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, MRSA communicates on Web Server Port 2463 and MRSA Server Port 9000. Ensure that these ports are available to be used by MRSA. Depending on your environment, if these ports are not available, specify the port details here. You can also edit this port details after installation.

6. Ensure that port 2463 is not blocked by a firewall.
7. 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

MRSA only supports 64-bit operating systems.

NOTE

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

8. To launch the application, 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

To install the LSI Storage Authority Software in a noninteractive or silent mode, use the following commands.

1. (Optional) If VC Redist 2010 and VC Redist 2015 are not installed on the server, complete the following steps.

- a) Install the VC Redist 2010 package from the command line `vc_redist_x64.exe /Q`.

VC Redist 2010 (`vc_redist_x64.exe`) is available in the `<Package_Dir>\ISSetupPrerequisites\{7f66a156-bc3b-479d-9703-65db354235cc}` directory.

- b) Install the VC Redist 2015 package from the command line `vc_redist.x64.exe /Q`.

VC Redist 2015 (`vc_redist.x64.exe`) is available in the `<Package_Dir>\ISSetupPrerequisites\{D093EE4D-527D-4CC7-AB3C-DCE3219FA508}` directory.

2. If OpenSLP is not installed, install the OpenSLP from the command line `openslp_3.0.0_0_x64`.

OpenSLP is available in the `<Package_Dir>\ISSetupPrerequisites\{23401E90-6962-476F-9D92-F9027E91A490}` directory.

3. Use one of the following modes to install the software.

- **Gateway** – `setup.exe /s /v"/qn ALLUSERS=1
ADDLOCAL=Gateway INSTALLATIONCHOICES=0 EVENTNOTIFICATIONCHOICES=3`
- **StandAlone** – `setup.exe /s /v"/qn ALLUSERS=1
ADDLOCAL=StandAlone INSTALLATIONCHOICES=1 EVENTNOTIFICATIONCHOICES=3`
- **DirectAgent** – `setup.exe /s /v"/qn ALLUSERS=1
ADDLOCAL=DirectAgent INSTALLATIONCHOICES=2 EVENTNOTIFICATIONCHOICES=3`
- **Light Weight Monitor** – `setup.exe /s /v"/qn ALLUSERS=1
ADDLOCAL=LightWeightMonitor INSTALLATIONCHOICES=129 EVENTNOTIFICATIONCHOICES=3`

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

4. (Optional) To change the default directory structure, provide the following input with each mode of installation.

```
Setup.exe /s /v"/qn ALLUSERS=1 ADDLOCAL=LightWeightMonitor INSTALLATIONCHOICES=129  
INSTALLDIR=CustomDirecotryLocation
```

5. (Optional) To redirect the MRSA installer logs, use the following command.

```
setup.exe /s /v"/l*v \"<PATH_TO_LOG_FILE>\<FILE_NAME>.log\" /v"/qn ALLUSERS=1 ADDLOCAL=Gateway  
INSTALLATIONCHOICES=0 EVENTNOTIFICATIONCHOICES=3 INSTALLDIR=CustomDirecotryLocation
```

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/LSIStorageAuhority/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 LSIStorageAuhority-1.00xx.xxxx-xxxx`

Performing Initial Configuration

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

Changing the HPE MR Storage Administrator Software Port Number

Perform the following steps to change the HPE MR Storage Administrator 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 HPE MR Storage Administrator Service.

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 HPE MR Storage Administrator.

Changing the Nginx Read Timeout

On VMware, when you request process-intensive operations such as creating **240 Volumes**, **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 and the HPE MR Storage Administrator.

Performing the Initial Setup

After you successfully log into the HPE MR Storage Administrator application, it is recommended that you perform the initial setup tasks before proceeding.

Displaying or Blocking a Private IP Address

This section outlines the strategy that the application follows to display or block a private IP address in a corresponding sub-net.

- **Private IP address** – A private IP address is a non-Internet facing IP address on an internal network. Private IP addresses are provided by network devices, such as routers, using network address translation (NAT).
- **Virtual IP address** – A virtual IP address (VIPA) is an IP address that is assigned to multiple domain names or servers that share an IP address based on a single network interface card (NIC). VIPAs are allocated to virtual private servers, websites, or any other application that resides on a single server. The host server for these applications has a network IP address that is assigned by a network administrator, whereas the different server applications have VIPAs. VIPAs enhance network load balancing and redundancy.
- **Automatic Private IP Addressing** – Automatic Private IP Addressing (APIPA) is a feature of Windows-based operating systems that enable a computer to automatically assign itself an IP address when no Dynamic Host Configuration Protocol (DHCP) server is available to perform that function. APIPA serves as a DHCP server failover mechanism and makes it easier to configure and support small local area networks.
- **Private IP Address Range** – The following is the IP address range which falls under either the private, (or) Virtual, (or) APIPA category:
 - **NAT** – 10.0.0.0 - 10.255.255.255
 - **Private (or) Virtual** – 172.16.0.0 - 172.31.255.255 or 192.168.0.0 - 192.168.255.255
 - **APIPA** – 169.254.0.0 to 169.254.255.255

The use cases that follow provide details on how the application behaves in various situations:

Table 4: Use Case #1: Without Blocking the Private IP

Use Case	Standalone / Client	Remarks
No NIC CARD (Windows)	Loopback (or) 127.0.0.1	As the server is not in network, the gateway cannot access the standalone server.
No NIC CARD (Linux)	Loopback (or) 127.0.0.1	Because the server is not in network, the gateway cannot access the standalone server.
Static IP	Using Static IP	—
DHCP IP	Using the DHCP IP	—
Private IP	Using the Private IP	In a more secured environment, private IP address cannot be accessed outside the server.

Table 5: Use Case #2: After Blocking the Private IP


Use Case	Standalone / Client	Remarks
No NIC CARD (Windows)	Loopback (or) 127.0.0.1	As the server is not in network, the gateway cannot access the standalone server.
No NIC CARD (Linux)	Loopback (or) 127.0.0.1	Because the server is not in network, the gateway cannot access the standalone server.
Static IP	Using Static IP	—
DHCP IP	Using the DHCP IP	—
Private IP	If a valid IP exists, it is displayed. If no valid IP exists, Loopback (or) 127.0.0.1 is displayed.	In a more secured environment, because a private IP address cannot be accessed outside the server, the application does not populate a private IP address.

Why does the application blocks certain IP addresses: In an enterprise world, when a computer is assigned a private IP address, the local devices see this computer through its private IP address. However, devices residing outside of your local network cannot directly communicate through the private IP address, but uses your router's public IP address to communicate. You must use a NAT router to directly access a local device assigned a private IP address.

In a more secure environment, although the application is able to discover and display the private IP address through the gateway server, when a request is made through the gateway server, the private IP is not accessible. Because the application cannot access the private IP, the application is unable to service the requests which are meant for the private IP.

Because of the previously mentioned reason, when the installation is a gateway, the corresponding gateway server is not able to communicate with the private IP address which in turn becomes an issue. The application works if the private IP addresses are behind the NAT router, which is the most preferable option in an enterprise world.

The diagram that follows shows how a private IP address should be accessed in enterprise networks and the problems with the private IP address:

Figure 5: Private IP Address Access

Alert Settings

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

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

1. Click **Settings** in the Server Dashboard.
The **Alert Settings** page appears, with the default alert delivery methods for each severity level.

Figure 6: Alert Settings Page

The screenshot shows the 'Alert Settings' page with tabs for 'Alert Settings', 'Mail Server', and 'Email'. The 'Alert Settings' tab is active. Below the tabs, it says 'Choose the alert delivery method for each severity level' and 'Displaying default alert settings'. The main content area is a table with four sections: 'Fatal', 'Critical', 'Needs Attention', and 'Information'. Each section has a description and two rows of settings: 'Within Application' and 'Outside Application'. Each row has three checkboxes for 'System Log', 'Event Log', and 'System Messages'. The 'Fatal' section has all three checked. The 'Critical' section has 'System Log', 'Event Log', and 'System Messages' checked, and 'Email' unchecked. The 'Needs Attention' section has 'System Log', 'Event Log', and 'Email' checked, and 'System Messages' unchecked. The 'Information' section has 'System Log', 'Event Log', and 'Email' checked, and 'System Messages' unchecked. On the right side, there is an 'Actions' panel with two buttons: 'Save Alert Settings' and 'Restore Default Alert Settings'.

Severity Level	Description	Within Application	Outside Application	System Log	Event Log	System Messages	Email
Fatal	when a component fails and data loss occurs			☒	☒	☒	
Critical	when a component fails			☒	☒	☒	☐
Needs Attention	when a component is close to failure point			☒	☒	☐	☒
Information	informational message where no user action is necessary			☒	☒	☐	☒

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

Setting Up the Email Server

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

1. Click the **Mail Server** tab on the **Settings** page.
The **Mail Server** page displays the current mail server settings.

Figure 7: Mail Server Window

The screenshot shows the 'Mail Server' tab in a settings application. At the top, there are three tabs: 'Alert Settings', 'Mail Server', and 'Email'. Below the tabs, a message reads: 'Provide mail and server settings from which the application will send alert notifications. Displaying current mail server settings'. The form contains the following fields: 'Sender Email Address' with the value 'isa-monitor@server.com', 'SMTP Server' with the value '127.0.0.1', 'Port' with the value '25' and a checked 'Use Default' checkbox, a section for authentication with a checked 'This server requires authentication' checkbox, 'User Name' and 'Password' fields, and 'Save' and 'Cancel' buttons at the bottom.

2. Specify the details in the respective fields as per your requirement.
3. Select the **This server requires authentication** check box on your SMTP server if the authentication login feature is enabled. To enable this feature, specify the authentication details in the **User Name** and **Password** fields respectively.
4. Click **Save**.

Adding the Email Addresses of Alert Notification Recipients

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

1. Click the **Email** tab in the **Setting** page.

Figure 8: Email Window

The screenshot shows the 'Email' tab in a settings application. At the top, there are three tabs: 'Alert Settings', 'Mail Server', and 'Email'. Below the tabs, a message reads: 'Provide email addresses to which the email alert notifications will be sent. Displaying current email settings'. The form contains the following elements: an 'Add Email Address' section with a text input field and an 'Add' button, a section titled 'Email alerts will be sent to the following email ids' with a list containing 'root@localhost' and a 'Remove' button, a 'Send Test Mail' button, and 'Save' and 'Cancel' buttons at the bottom.

2. Specify the details in the respective fields as per your requirement.
3. Click **Save**.

Configuring Different Types of Access

To enable logging in with a user name and password, complete the following steps.

NOTE

By default, the application is installed without a login so the user can access the application without entering a username and password.

1. Stop the MRSA services and nginx services by running the `/etc/init.d/LsiSASH stop` command.
2. Go to the file `<MRSA_Home>/LSIStorageAuthority\conf`.
3. Open the file `LSA.conf`.
4. Update the following field value to 0.

```
# bypass authentication (use with caution)
bypass_authentication = 1
```
5. Restart the MRSA services and nginx services by running the `/etc/init.d/LsiSASH restart` command.
6. Verify the status of the services by running the `/etc/init.d/LsiSASH status` command.

NOTE

For a more secure environment, use the `bypass_authentication = 0` command. Setting `bypass_authentication = 1` allows non-administrator users to log in and access MRSA, which may not be desired for all setups.

Server Dashboard

The Server Dashboard is the default landing page for the software. The **Server Dashboard** displays the overall summary of the server and the devices that are attached to it. You can troubleshoot, configure, maintain, and monitor the controllers from the **Server Dashboard**. The following figure and table describes this page.

Figure 9: Server Dashboard Window




Table 6: Server Dashboard Description

Callout	Description
1	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: • : 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. – Click @Settings to perform initial settings. – 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 software that is installed. You can also view the controller information such as controller hardware, enclosure of the controller, and information about the drives and volumes associated with the controller. • : Lets you enable or disable system messages. • : Displays the application context-sensitive help.
2	Controller Status – Description as follows: • Displays the status of all of the controllers that are connected to the server. It displays the total number of controllers and status icons based on their criticality: – Critical: Indicates that a critical error exists on the controller and the controller needs immediate attention. – Needs Attention: Indicates that an error exists on the controller that needs attention, however, not immediately. – Optimal: Indicates that the controller is operating in an optimal state. • Displays critical issues of failed devices and provides recommendations for troubleshooting. You can also see contextual links, which help you to easily locate the device and initiate troubleshooting. Based on the criticality of the controller, the application displays information about that particular controller in the controller information pane. For example, if a controller is in the critical state, that controller is opened by default. If you want to view information about other controllers, click the respective Controller Status icon. Click View All Controllers to view information about all of the controllers. OS Information – Displays the server's operating system information. Download Support Log – Lets you download the support log, which contains consolidated information about the server and all the devices to which it is connected.
3	• Controller Information: Displays information about 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 • Lets you perform these tasks: – Configure the controllers. See Controller Configurations. • 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 the appropriate controller.

Controller Dashboard

You can perform controller related actions and view all the information pertaining to a controller from the Controller Dashboard. The figure and table that follows describes this page.

Figure 10: Controller Dashboard Window




Table 7: Controller Dashboard Description

Callout	Description
1	Controller Summary – Displays the name of the MegaRAID controller card. The color-coded icons indicate the status of the controller card. Displays the basic controller properties, such as the controller serial number, vendor ID, SAS address, driver version, device ID, host interface, and so on. Click the  icon to view the advanced properties of the controller, such as the NVRAM details, BIOS version, firmware properties, emergency spare properties, and so on.
2	Controller Views – Displays all of the configured arrays, volumes, and drives associated with the selected controller card. It also displays the hardware, such as enclosures and backplanes associated with the controller. All these views are displayed as tabs. Click the  icon to view to view detailed information about the device. For example, click an array to view the associated volumes and drives. Select any device from the expanded view to perform relevant actions and view device properties.
3	Controller Actions – Lets you perform the following actions: • Create a configuration • Clear a configuration • Update the controller firmware • Import or clear foreign configurations • View premium features • View the event log • Add personality management • Function profile management • Auto assign

Controller Configurations

You can use the application to create and modify storage configurations on systems with Hewlett Packard Enterprise controllers.

You can create RAID 0, RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, and RAID 60 storage configurations.

The supported RAID levels differ or might not be supported for some controllers. For more information, see [HPE MR Storage Administrator Feature Support Matrix](#).

You can create these 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 volume. See [Creating a New Storage Configuration Using the Simple Configuration Option](#) for details.
- **Advanced Configuration**
Lets you choose additional settings and customize volume creation. This option provides greater flexibility when creating volumes for your specific requirements. See [Creating a Storage Configuration Using the Advanced Configuration Option](#) for details.

Creating a New Storage Configuration Using the Simple Configuration Option

The Simple Configuration option is the quickest and easiest way to create a new storage configuration. When you select Simple Configuration option, 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 volume using that physical drive. To create a volume when the physical drive is in the Prepare for Removal state, you must manually undo the operation by using the Undo Remove option.

Perform these steps to create a simple storage configuration:

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

The **Simple Configuration** page opens.

Figure 11: Simple Configuration Page

 Simple Configuration ⓘ

Step 1/1 : Choose your configuration settings

1. RAID Level Setting [\(Compare and select\)](#)

RAID 0

This RAID level is suitable for high performance with zero data redundancy. Choose this option only for non-critical data.

2. How many Volumes do you wish to create?

1

each with capacity of

1.818 TB

3. Miscellaneous Drive Attributes

☐ Assign Spare

Spare will be assigned depending upon the availability of eligible spare candidate drives. A spare drive will take over for a drive if a failure happens, ensuring the data remain intact

Finish

2. Select a RAID level for the array from the drop-down box.
3. (Optional) – Click **Compare and Select** to view detailed information about each RAID level.
When you use the Simple Configuration option, the RAID controller supports RAID levels 0, 1, 5, and 6. The window text provides 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 volumes that you want to create.
5. Select the capacity of the volumes.
Each volume has the same capacity.
6. Select the **Assign Spare** check box if you want to assign a dedicated spare drive to the new volume.
If an Unconfigured Good drive is available, that drive is assigned as a spare drive. Spare drives are drives that are available to replace failed drives automatically in a fault tolerant volume (RAID 1, RAID 5, or RAID 6).
7. Click **Finish**.
A message appears stating that the configuration is successfully created.

Creating a Storage Configuration Using the Advanced Configuration Option

The Advanced Configuration option provides an easy way to create a storage configuration. The Advanced Configuration option gives you greater flexibility than simple configuration because you can select the drives and the volume parameters when you create a volume. In addition, you can use the Advanced Configuration procedure to create spanned arrays (parity groups). Perform these steps to create an advanced storage configuration.

1. Select **Configure > Advanced Configuration** from the Server Dashboard or the Controller Dashboard.
The **Advanced Configuration** page is displayed.

Figure 12: Advance Configuration Page

The screenshot shows the 'Advanced Configuration' page, Step 2/2: Configure Drives and Create Volumes. At the top, there's a header with a wrench icon, the title 'Advanced Configuration', and a help icon. Below the header, there are 'Back' and 'Finish' buttons. The main content area is divided into two sections. The left section, titled '2 Drives', shows a list of two drives: 'Port 2I,Box=1,Bay=4, Model-LK0480GFJSK SATA, 447.131 GB' and 'Port 2I,Box=1,Bay=3, Model-MK0120GFDKP SATA, 111.79 GB'. Each drive has a blue 'X' icon to its right. Above the list is an 'Add Drives' button. The right section, titled '0 Volumes', shows a list with the text 'No Volume added yet'. Above this list is an 'Add Logical Drives' button. At the bottom of the page, there are 'Back' and 'Finish' buttons.

2. Select a RAID level for the array from the drop-down box.
3. (Optional) – Click **Compare and Select** to view the detailed information on each RAID level.
When you use the Advanced Configuration option, the RAID controller supports RAID levels 10, 50, and 60. The **Compare and Select** option provides you 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. Click **Next**.
 5. Click **Add Drives** to add drives to the array.
 6. (Optional) – Select the span depth using the slider bar.
 7. Click **Add Drives** to add drives to the array.
- The **Available Unconfigured Drive** window appears.

Figure 13: Available Unconfigured Drive Window

0 Foreign Drives

5 Unconfigured Drives

0 Configured Drives

0 Spares

0 JBOD

5 Unconfigured good

Enclosure Bay	Device ID	Type	Interface	Capacity	Sector Size	Status	Model
Port 3L,Box=1,Bay=1	56	HDD	SAS	300GB	512B	Unconfigured good	EG000300.JWBHR
Port 3L,Box=1,Bay=2	55	HDD	SAS	300GB	512B	Unconfigured good	EG000300.JWBHR
Port 3L,Box=1,Bay=3	59	HDD	SAS	300GB	512B	Unconfigured good	EG000300.JWBHR
Port 3L,Box=1,Bay=4	57	HDD	SAS	300GB	512B	Unconfigured good	EG000300.JWBHR
Port 4L,Box=1,Bay=5	58	HDD	SAS	300GB	512B	Unconfigured good	EG000300.JWBHR

Actions

Properties

Select any Drive to view its actions & properties

For information on adding unconfigured drives to the array, see [Selecting Available Unconfigured Drives](#).

8. Select the drives from the list of available unconfigured drives and click **Add Drives**.
 9. Click **Add Volumes** to add volumes to the array.
- The **Volume Settings** window appears.

Figure 14: Volume Settings Window

The screenshot shows the 'Volume Settings' window. At the top, it states '55.375 GB available across 1 selected drive' and '16 more Volumes can be added'. Below this, it asks 'How many Volumes do you wish to create?' with a dropdown set to '1'. Next to it, 'each with capacity of' is followed by a dropdown set to '55.375' and a unit dropdown set to 'GB'. There are two input fields: 'Volume Name' with the text 'volumeName' and 'Strip Size' with a dropdown set to '256 KB'. On the left, there is a sidebar with five sections: 'Initialization State' (No Initialization), 'Read Policy' (Read Ahead), 'Write Policy' (Write Back), 'IO Policy' (Direct IO), and 'Drive Write Cache Policy' (Disabled). The 'IO Policy' section is currently selected. To the right of this sidebar, there is a description: 'The IO policy applies to reads on a specific Volume. It does not affect the read ahead cache.' Below this, the 'Direct IO' option is selected with a radio button, and its description is: 'In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. This option is the default setting.' At the bottom left of the window is a blue 'Add Volume' button.

For information on configuring volumes, see [Selecting Volume Settings](#).

10. Specify all the required details and click **Add Volumes**.

11. Click **Finish**.

A message appears confirming that your configuration is complete.

Selecting Available Unconfigured Drives

The **Available Unconfigured Drive** window lets you add drives and spare drives to the array.

Perform these steps to add drives and spare drives to the array.

1. Select the drives to add from the **Available Unconfigured Drives** window, and click **Add Drives**.

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

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

2. Click **Add Spares** to add dedicated spare drives to the array.

The **Available Unconfigured Drives** window appears.

3. Select the drives you want to add as spare drives and click **Add Spare Drives**.

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

Selecting Volume Settings

The **Volume Settings** window enables you to configure the volumes. Detailed descriptions for all of the parameters are present in the **Volume Settings** window.

The volume settings differ or might not be supported for some controllers. For more information, see [HPE MR Storage Administrator Feature Support Matrix](#).

Perform these steps to configure a volume:

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

Each volume has the same capacity. If you specify the capacity first and then the number of volumes, the volume capacity is adjusted with the available capacity.

3. Specify a name for the volume in the **volume Name** field.

The volume name can have a maximum of 15 characters.

4. Select a strip size from the **Strip Size** drop-down list.

Strip sizes of 64 KB , 128 KB , 256 KB , 512 KB , and 1 MB are supported.

5. Specify the Initialization State.

The options follow:

- **Fast Initialization**
- **Full Initialization**
- **No Initialization**

6. Specify the read policy for the volume.

The options follow:

- **No Read Ahead**
- **Read Ahead**

7. Specify the Write Policy for the volume.

The options follow:

- **Write Through**
- **Write Back**
- **Always Write Back**

8. Specify the IO policy for the volume.

The options follow:

- **Cached IO**
- **Direct IO**

9. Specify a Drive Write Cache Policy for the volume.

The options follow:

- **Unchanged**
- **Disabled**
- **Enabled**

10. Click **Add volumes**.

The newly created volume appears in the **Advanced Configuration** window just below the **volumes** section.

NOTE

You will lose some drive capacity if you choose different size drives volume while creating a volume.

If you want to modify the volume settings before completing the configuration, click the  icon.

The **volume Settings** window opens.

Modify the settings as desired and click **Modify volume**.

Clearing the Configuration

You can clear all existing configurations on a selected controller.

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

1. Navigate to the Controller Dashboard.
2. Click **Configure**, then click **Clear Configuration**.
A confirmation message appears.
3. Select **Confirm** and click **Yes, Clear configuration** to clear existing configurations on the controller.

NOTE

The operating system or file system drives cannot be cleared. The following error code appears:

```
The Operation is not allowed because one or more Volumes or Drive(s) has an OS / FS / Unknown
Boot partition(cannot be read)." User can go to individual Drives or Volumes and delete the
Volumes or change the JBOD to UG.
```

Importing or Clearing Foreign Configurations

A foreign configuration is a RAID configuration that already exists on an added drive. You can use the application 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 these steps to import or clear foreign configurations.

1. Navigate to the Controller Dashboard.
2. Click **Configure**, then click **Foreign Configuration**.
The **Foreign Configuration** window appears, which lists all of the foreign configurations.
3. Click one of these 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.

UNMAP Capability Feature

The UNMAP capability feature is a SCSI command that is used to reclaim unused LBAs from the SSD. The UNMAP feature allows an application or OS to tell the storage array that the disk blocks contain deleted data so the array can deallocate the blocks. Deallocation reclaims storage space and helps in wear leveling management. Using the UNMAP capability feature extends the SSD lifespan.

NOTE

The UNMAP Capability feature can be enabled on the volume using SAS and NVMe SSD only for MR200/400.

The UNMAP Capability feature can be enabled on the volume using SAS SSD only for MR932.

UNMAP Capability Feature Behavior

MRSA behavior for PR5 designs and later include the behaviors that follow.

- Display the PD Property, whether the PD (physical drive) is UNMAP capable or not.
- Display the PD Capability, whether the PDs can be used for volumes for the UNMAP feature.
- Lets users create an UNMAP supported volume.

MRSA behavior for PR5 designs include the following limitations.


- The UNMAP feature is not supported for JBOD designs.
- Host software applications cannot support firmware in designs earlier than PR5 because of the change in the firmware API.

UNMAP Feature Support

When using the UNMAP feature, you can perform the actions that follow.

- Enable the UNMAP capability during volume creation.

Figure 15: Enable the UNMAP Feature During Volume Creation



- Show the PD level UNMAP properties.

Figure 16: Drive Level UNMAP Properties Window

Physical Drive Properties

General Properties

SAS Address 0 0x0	SAS Address 1 0x5000cca04a71a0fa	Negotiated Link Speed 12G	Drive Speed 12G
Temperature 28C	Revision Level A2C0	Power Status ON	Native Command Queuing Capable Enabled
Unmap Capable Yes	Unmap Capable for Lds Yes	Physical Sector Size 4KB	

Enclosure Properties

Enclosure ID 69	Enclosure Model VirtualSES	Enclosure Location Internal	Enclosure Connector [C1 x1]
---------------------------	--------------------------------------	---------------------------------------	---------------------------------------

Drive Security Properties

Full Disk Encryption	Full Disk Encryption Type
-----------------------------	----------------------------------

- Let the user modify the UNMAP capability of the Volume using the VD Modify (Enable/Disable) Properties option.

Figure 17: Enable or Disable the SCSI UNMAP Feature

Configured Capacity 1.272 TB of 1.453 TB

Serial No: FW-00000000 | SAS Address: 0x5000000012345678 | Alarm: | Driver Version: 7.710.0

2 Drive Groups (2 Virtual Drives) | **4 Drives** (0 Unconfigured Drives) | **1 Other Hardware**

DG_0 RAID 1 | 1 Virtual & 2 Physical Drives | Used 558.41 GB of 558.41 GB Available

DG_1 RAID 0 | 1 Virtual & 1 Physical Drives | Used 185.78 GB of 185.78 GB Available

1 Virtual Drives | **1 Physical Drives**

Enclosure:Slot	Device ID	Type	Interface	Capacity	Sector Size	Model
EN_69:6	0	SSD	SAS	185.78GB	512B	HUSMH8020BSS200

Context Menu:

- Schedule Consistency Check (Starts on Saturday : 3 AM)
- Set Consistency Check Properties
- Set Patrol Read Properties (Starts on Friday : 2 PM)
- Start Patrol Read
- Manage SAS Storage Link Speed
- Premium Features
- Set Adjustable Task Rate
- Disable SSD Guard
- Disable SCSI Unmap** (highlighted with a red circle)
- Download TTY Log
- Manage Power Save Setting

Profile Management

- Configure
- Update Firmware (Firmware Version 5.110.01-2585)
- View Event Log
- More Actions

Element(s) Actions

Selected Elements: 0

Select any Drive Group, Virtual or Physical Drive to view its actions

Personality Management

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

- In the **Change Personality** page, select the **Change Auto-configure behavior** check box.
- From the **Select Behavior Mode** drop-down list, select an appropriate behavior mode.
The available behavior modes are based on the current firmware support.

- **NONE (Unconfigured Good)** – If a user selects this option when a new disk or an old disk with no DDF (Data Disk Format) metadata is inserted in the system, it becomes Unconfigured Good Drive. The Unconfigured Good Drive keeps in Unconfigured Good after a reboot.
- **JBOD** – If all the drives are in unconfigured good state, once you select the JBOD mode, all the drives are automatically converted to JBOD drives. When a new disk or old disk with new DDF metadata is inserted in the system, it becomes a JBOD. The Unconfigured Good Drive is automatically converted to JBOD after a reboot.
- **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.

NOTE

For firmware versions earlier than 52.26.3-5250, Auto-Configure behavior is set to JBOD and cannot be configured for all the controller models.

For firmware versions 52.26.3-5250 or later, Auto-Configure behavior can be configured.

For HPEMR200 controllers, by default it is set to JBOD.

For HPEMR400 controllers, by default it is set to None.

Figure 18: Change Personality Dialog

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

Function Profile Management

1. Select **Actions > Function Profile Management** on the Controller dashboard.

The **Function Profile Management** page appears.

Figure 19: Function Profile Management

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

Close

Function Profile Management

	Profile ID	Is Current	Is Requested	Selectable	Selectable With Force	Function Type	Start Function ID	Min Controller Functions	Max Controller Functions	Queues/Vectors
☒	1	Yes	No	No	No	MPI		1	4	256

Save

Server Reboot will be required after saving the changes

Table 8: Function Profile Management Properties

Property	Description
Profile ID	Displays the unique identifier for the profile.
Is Current	Indicates whether the profile is the current profile.
Is Requested	Indicates whether the profile is the requested profile.
Selectable	Indicates whether the user can select this profile.
Selectable With Force	Indicates whether the user can select this profile only with a force option.
Function Type	Displays the function type for the profile.
Start Function ID	Displays the starting function ID of the profile.
Min Controller Functions	Displays the minimum supported controller functions.
Max Controller Functions	Displays the maximum supported controller functions.
Queues/Vectors	Displays the supported queues and vectors.

2. Select the radio button in the first column, and then click **Save** to change the current profile.
3. Reboot the system for the changes to take effect.

Auto Assign Policy

Select **Actions > Auto Assign** on the Controller dashboard.
The **Change Auto assign policy** page appears.

Figure 20: Change Auto assign policy Dialog

Change Auto assign policy

Current Auto-assign policy : Supervisor

Default Auto-assign policy : Supervisor

☒ Change Auto-assign policy ☐ Auto-assign policy (Execute Once)

Auto-assign policy

Supervisor

UnAssigned

Supervisor

Weighted Round Robin

Change

Background Operations Support

The application provides background operations to Pause, Resume, Abort, Pause All, Resume All, and Abort All. The features enhance the functionality where operations running in the background on a drive or a volume can be paused for some time, and resumed later.

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

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

Figure 21: Background Processes in Progress Window



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


Managing Controllers

The HPE MR Storage Administrator lets you monitor the activity of all the controllers and the attached devices.

Viewing Controller Properties

The Controller Dashboard displays basic controller properties. Click the  icon to view the advanced properties of the controller.

Figure 22: Basic and Advanced Controller Properties Window



Running Consistency Checks

The Consistency Check operation verifies the correctness of the data in volumes 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 volumes.

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

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

Set Consistency Check Properties

Perform these 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 these two options:
 - **Continue Consistency Check and Fix Error** – The RAID controller continues the consistency check, and if it finds any errors, fixes them.
 - **Stop Consistency Check On Error** – The RAID controller stops the consistency check operation if it finds any errors.
3. Click **Save**.

Schedule Consistency Check Operation

Perform these steps to schedule a Consistency Check operation:

1. In the Controller Dashboard, select **More Actions > Schedule Consistency Check**.
The **Schedule Consistency Check** page appears.

Figure 23: Schedule Consistency Check Dialog

← Go back to Array, Drives and Other Hardware list Close

Schedule Consistency Check ⓘ Next

1. Set Consistency Check Mode

Concurrent ▾

2. Schedule Consistency Check every

Daily ▾ From December 19, 2017 ⓘ at 07:00 PM ▾

Next

2. Set the **Consistency Check Mode**.
The available options are:
 - **Concurrent** – Run a Consistency Check operation concurrently on all volumes.
 - **Sequential** – Run a Consistency Check operation on one volume at a time.
 - **Disable** – Disables the Consistency Check feature.
3. Set the desired interval at which you want to check the consistency of a drive.
The available options are:
 - **Hourly, Daily, Weekly, Monthly, and Continuously**.
 - Select an appropriate date and time range.
4. Click **Next**.
The **Schedule Consistency Check** page appears, letting you add the volumes on which you want to perform a Consistency Check operation.

5. Click **Add Volumes**.

The **Available Volume** dialog appears which lists all the volumes present in the selected array.

6. Select the volume on which you want to run the a Consistency Check operation.

7. Click **Save**.

The consistency check runs based on the frequency/interval chosen by you. You can also monitor the progress of the consistency check operation. See [Background Operations Support](#)

8. (Optional) – Select the volume, from the Controller View section, on which you want to immediately perform a Consistency Check operation, then go to **More Actions > Start Consistency Check**.

NOTE

If you try to run a Consistency Check operation on a volume that has not been initialized, a confirmation dialog appears, asking for your confirmation.

Manage SID Ownership

The Manage SID Ownership feature allows you to set the SID credentials of SED drives to be managed by MSID or by firmware. For managing the SID ownership, navigate to **More Actions > Manage SID Ownership**.

Figure 24: Manage SID Ownership

The screenshot shows the 'Manage SID Ownership' dialog box in two states. The top part shows the initial state with 'Current SID Credential : MSID', 'Requested SID Credential: NONE', 'Default SID Credential : MSID', and 'Supported SID Credential: MSID, FW'. Below this, 'Drive Security Firmware Managed: 0' is shown with a warning message. The bottom part shows 'Drive Security Default: 0' with another warning message. The bottom part of the screenshot shows the state after a successful change, with 'Requested SID Credential: FW' and 'Supported SID Credential: MSID, FW'. A green checkmark and a message 'Change Successful, but reboot is required to take effect' are displayed. The 'Drive Security Firmware Managed' value is now 2, and the 'Drive Security Default' value is also 2. Two tables are shown, both displaying the same physical disk information.

Enclosure	Interface	Media	Capacity	Logical Sector Size	Model
EN_406 : 5	SATA	SSD	238.474GiB	512B	KINGSTON SKC600256G
EN_406 : 8	SATA	SSD	238.474GiB	512B	KINGSTON SKC600256G

Enclosure	Interface	Media	Capacity	Logical Sector Size	Model
EN_406 : 5	SATA	SSD	238.474GiB	512B	KINGSTON SKC600256G
EN_406 : 8	SATA	SSD	238.474GiB	512B	KINGSTON SKC600256G

The **Manage SID Ownership** dialog box has two tables. The **Drive Security Default** table displays the MSID/default owned physical disk information. The **Drive Security Firmware Managed** table displays the firmware that is owned. When all of the physical disks are owned by MSID, the physical disk count is displayed next to the table name.

NOTE

The MSID/default physical disks are secured JBOD disks only. The firmware-owned physical disks are secured JBOD, but owned by firmware.

Device Reporting Order

Perform these steps to set the device reporting order.


NOTE

Device reporting order for non supervisor physical functions can be performed from the **Physical Functions** tab (**Select Physical Function > Go to Actions > Device Reporting Order**).

1. In the Controller dashboard, select **More Actions > Device Reporting Order**.

The **Device Reporting Order** page appears.

Figure 25: Device Reporting Order Dialog



Device Reporting Order

Device Reporting Order

Logical drives are reported after JBOD devices

Logical drives are reported after JBOD devices

Logical drives are reported prior to JBOD devices

First DevicePersistentID

1

View List

2. In the Device Reporting Order dialog, select an order for device reporting. The following options are available.

- Logical drives are reported after JBOD devices
- Logical drives are reported before JBOD devices

3. Set the **First DevicePersistentID**.

Running a Patrol Read Operation

The Patrol Read option lets you periodically verify 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 operation for all RAID levels and for all spare drives. A Patrol Read operation 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 operation without changing the properties.

Setting the Patrol Read Properties

Perform these steps to set the Patrol Read properties.

1. Select **More Actions > Set Patrol Read Properties** in the Controller Dashboard.
The **Available Volumes** dialog appears.
2. Select the volumes for which you want to set the Patrol Read properties and click **Add Volumes**.
The **Set Patrol Read Properties** dialog appears.

3. Click **Select Volumes**.

Click the  icon to remove the volumes you have already added.

4. Click **Next**.

5. Perform these steps to set the properties:

- a) Select an operation mode for patrol read from the **Set Patrol Read Mode** drop-down list.

The options follow:

- **Automatic** – The Patrol Read operation runs automatically at the time interval you specify.
- **Manual** – The Patrol Read operation runs only when you manually start it, by selecting **Start Patrol Read** from the Controller Dashboard.
- **Disabled** – The Patrol Read operation does not run.

- b) (Optional) – Specify a maximum number of drives to include in the Patrol Read operation concurrently.

The count must be a number from 1 to 240.

- c) Select the frequency at which the Patrol Read operation runs from the drop-down list.

The default frequency is **Weekly** (168 hours), which is suitable for most configurations. The other options are **Hourly**, **Daily**, and **Monthly**.

- d) Select the month, day, and year on which to start the Patrol Read operation.

- e) Select the time of day to start the Patrol Read operation.

- f) (Optional) – Select the **Start Patrol Read Now** check box.

6. Click **Finish**.

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

Starting a Patrol Read Operation

Perform these steps to start a Patrol Read operation.

1. Select **More Actions** > **Start Patrol Read** on the Controller Dashboard.

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 this step to stop a Patrol Read operation.

Select **More Actions** > **Stop Patrol Read** on the Controller Dashboard.

Managing SAS Storage Link Speed

- **For direct-attached drives:**

- For drives directly connected to the controller, 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.

- **For UBM backplanes:**

- Lane speeds for controllers connected to UBM backplanes cannot be modified.

Perform these steps to change the link speed.

1. Select **More Actions > Manage SAS Storage Link Speed** on the Controller dashboard.
The **Manage SAS Storage Link Speed** dialog appears.

Figure 26: Manage SAS Storage Link Speed Window

Manage SAS Storage Link Speed 

Phy	Status	Port Number	Select Link Speed	Current Speed	Connected Device
0	OPTIMAL		MAX 	MAX	
1	OPTIMAL		MAX 	MAX	
2	OPTIMAL		MAX 	MAX	
3	OPTIMAL		MAX 	MAX	
4	OPTIMAL	0	12.0Gb/s 	12.0Gb/s	***
5	OPTIMAL	0	12.0Gb/s 	12.0Gb/s	***
6	OPTIMAL	0	12.0Gb/s 	12.0Gb/s	***
7	OPTIMAL	0	12.0Gb/s 	12.0Gb/s	***
8	OPTIMAL		MAX 	MAX	
9	OPTIMAL		MAX 	MAX	
10	OPTIMAL		MAX 	MAX	
11	OPTIMAL		MAX 	MAX	
12	OPTIMAL		MAX 	MAX	
13	OPTIMAL		MAX 	MAX	
14	OPTIMAL		MAX 	MAX	
15	OPTIMAL		MAX 	MAX	

Save

Server Reboot will be required after saving the changes

- The **Phy** column displays the system-supported phy link values.
The phy link values range from 0 through 7.
 - The **Status** column displays the status of the link speed.
 - The **Port Number** column displays the port numbers.
 - The **Select Link Speed** column displays the phy link speeds.
 - The **Current Speed** column displays the current negotiated logical link rate of the phy.
 - The **Connected Device** column displays the device that is 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**, **3.0Gb/s**, **6.0Gb/s**, **12.0Gb/s**, or **22.5Gb/s**.
By default, the link speed in the controller is set to **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 Interface

A lane represents a set of differential signal pairs, one pair for transmission and one pair for reception, similar to SAS phys.

The Managing PCIe Storage Interface feature allows you to change the lane speed between a controller and an expander or between the controller and a drive that is directly connected to the controller. PR5 and later versions support both SAS/SATA topologies and PCIe topologies using the same device phys to manage the lane speed.

NOTE

Lane speeds for UBM backplanes cannot be modified.

Perform the following steps to change the lane speed.

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

Figure 27: Manage PCIe Storage Interface Dialog

Manage PCIe Storage Interface ⓘ

lane	Status	Link Number	Lane Speed	Current Speed	Connected Device
0	OPTIMAL	0	2.5GT/s	N/A	
1	OPTIMAL	0	2.5GT/s	N/A	
2	OPTIMAL	0	2.5GT/s	N/A	
3	OPTIMAL	0	2.5GT/s	N/A	
4	OPTIMAL	1	2.5GT/s	N/A	
5	OPTIMAL	1	2.5GT/s	N/A	
6	OPTIMAL	1	2.5GT/s	N/A	
7	OPTIMAL	1	2.5GT/s	N/A	
8	OPTIMAL	2	2.5GT/s	N/A	
9	OPTIMAL	2	2.5GT/s	N/A	
10	OPTIMAL	2	2.5GT/s	N/A	
11	OPTIMAL	2	2.5GT/s	N/A	
12	OPTIMAL	3	2.5GT/s	N/A	
13	OPTIMAL	3	2.5GT/s	N/A	
14	OPTIMAL	3	2.5GT/s	N/A	
15	OPTIMAL	3	2.5GT/s	N/A	

Save

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 **Link Number** column displays the link numbers.
- The **Lane Speed** column displays the lane speed.
- The **Current Speed** column displays the current negotiated logical link rate of the phy.

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** for MR Gen10 plus and Gen11 controllers.

The lane speed values are **Unknown**, **2.5GT/s**, **5GT/s**, **8GT/s**, **16GT/s**, and **32GT/s** for MR Gen10 plus and Gen12 controllers.

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.

Set Adjustable Task Rate

Perform these steps to set the adjustable task rate.

1. Select **More Actions > Set Adjustable Task Rate** on the Controller Dashboard.

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

Figure 28: Set Adjustable Task Rate Dialog

Set Adjustable Task Rate [?](#)

Task	Priority
Rebuild Rate (%)	33 
Patrol Read Rate (%)	30 
BGI Rate (%)	30 
Consistency Check Rate (%)	30 
Transformation Rate (%)	30 

Save

2. Enter changes, as needed, in the following task rates:

NOTE

Setting any of these rates to perform faster can result in the system I/O rate being slower.

- **Rebuild Rate (%)** – Enter a number from 1 to 100 to control the rate at which a rebuild is performed on a drive when it is necessary.
The higher the number, the faster the rebuild occurs.
- **Patrol Rate (%)** – Enter a number from 1 to 100 to control the rate at which Patrol Read operations are performed.
The Patrol Read function monitors drives to find and resolve potential problems such as media problems. The higher the number, the faster the Patrol Read operation occurs.
- **Background Initialization (BGI) Rate (%)** – Enter a number from 1 to 100 to control the rate at which volumes are initialized in the background.
Background initialization establishes mirroring or parity for a RAID volume while allowing full host access to the volume. The higher the number, the faster the initialization occurs.
- **Consistency Rate Check (%)** – Enter a number from 1 to 100 to control the rate at which a consistency check is performed.
A Consistency Check operation scans the consistency data on a fault tolerant volume to determine whether the data has become inconsistent. The higher the number, the faster the Consistency Check operation is performed.
- **Transformation Rate (%)** – Enter a number from 1 to 100 to control the rate at which transformation of a volume occurs.
The higher the number, the faster the transformation occurs.

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

Setting the Task Information

Perform the following steps to set the adjustable task rates.

1. In the Controller dashboard, select **More Actions** > **Set Adjustable Task Rate** or **Set Task Priority** depending on the firmware version.

Figure 29: Set Adjustable Task Rate Dialog

Set Adjustable Task Rate 

Task	Priority
Rebuild Rate (%)	33 
Patrol Read Rate (%)	30 
BGI Rate (%)	30 
Consistency Check Rate (%)	30 
Transformation Rate (%)	30 

Save

Figure 30: Set Task Priority Dialog

Set Task Priority 

Task	Priority
Rebuild Operating Mode Priority	Rebuild 
Patrol Read Rate (%)	30 
BGI Rate (%)	30 
Consistency Check Rate (%)	30 
Transformation Rate (%)	30 

Save

2. Enter changes, as needed, in the following fields:
 - Depending on the firmware version, complete one of the following:
 - Rebuild Rate** – Enter a number from 1 to 100 to control the rate at which the rebuild is performed on a drive when it is necessary.
 - Rebuild Operating Mode Priority** – Select either **I/O** or **Rebuild** to prioritize the operating mode.
 - **Patrol Read Rate (%)** – Enter a number from 1 to 99 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 (%)** – Enter a number from 1 to 99 to control the rate at which are initialized in the background.

Background initialization establishes mirroring or parity for a RAID while allowing full host access to the . 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 1 to 99 to control the rate at which a consistency check is performed.

A consistency check scans the consistency data on a fault-tolerant 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 1 to 99 to control the rate at which online capacity expansion occurs.

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

Managing Power-Save Settings

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 that are connected to a MegaRAID controller. Reducing the power consumption helps to share resources more efficiently and lowers the cost.

Power consumption helps to share resources more efficiently and lowers the cost.

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

The **Manage Power Save Settings** dialog appears.

Figure 31: Manage Power Save Settings Dialog

Manage Power Save Settings

Power Save(Dimmer Switch) technology that conserves energy by spinning down idle drives. The controller will automatically spin up those drives from power save mode whenever necessary.

Specify the Power Save Settings below:

☒ Unconfigured Drives

☐ Spare Drives

Spin Down Time:

30 mins

Ensure that if the drives are idle for the specified time, then the drives will go to power save mode.

Set Drives Smart Poll Internal Interval: 300 Seconds

Set Drives Smart Poll External Interval: 300 Seconds

Enable Drives Smart Polling: ☒ Advanced Host Drives

Enable Drives Temperature Polling: ☒ Advanced Host Drives

Finish

2. Select the **Unconfigured Drives** check box to let the controller enable the unconfigured drives to enter the Power Save mode.
3. Select the **Spare Drives** check box to let the controller enable the spare drives to enter the Power Save mode.
4. Select the spin down time using the drop-down list from the **Spin Down Time:** field.
The drive standby time can be 30 minutes, 1 hour, and 2 hours through 24 hours.
5. Set the **Smart Poll Internal Interval** in seconds.
6. Set the **Smart Poll External Interval** in seconds.
7. To enable smart polling for drives, check the **Enable Drives Smart Polling** checkbox.
8. To enable drive temperature polling, check the **Enable Drives Temperature Polling** checkbox.
9. Click **Finish** to save the settings.
A confirmation message appears.

Discarding Pinned Cache

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

ATTENTION

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

Perform these steps to discard the pinned cache.

1. Select **More Actions > Discard Preserved Cache** on the Controller Dashboard.

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**.

Spin Down Drives at Shutdown

The application allows the controller operations to set spin down drives during a shutdown operation.

1. In the Controller dashboard, select **More Actions** > **Spin Down Drives At Shutdown**.
2. Select the **Drive Type** to spin down at shutdown.



The dialog titled "Spin Down Drives At Shutdown" features a teal header bar with a back arrow and the text "Go back to Drive Group, Drives and Other Hardware list" on the left, and a red "Close" button on the right. Below the header, the title "Spin Down Drives At Shutdown" is displayed. A table with two columns, "Drive Type" and "Set value", lists five drive types: SATA HDD, SAS HDD, SATA SSD, SAS SSD, and NVME SSD. The "Set value" column contains checkboxes, with SATA HDD, SAS SSD, and NVME SSD checked. A "Save" button is located at the bottom right of the table.

Drive Type	Set value
SATA HDD	☒
SAS HDD	☐
SATA SSD	☒
SAS SSD	☒
NVME SSD	☒

3. Click **Save**.


NVMe Thermal Poll Interval

Perform these steps to set the increase the frequency of temperature polling for NVMe drives. Adjusting the frequency of temperature polling helps achieve better cooling.

1. Select **More Actions** > **Set NVMe Thermal Poll Interval** on the Controller Dashboard.

The **Set NVMe Thermal Poll Interval** dialog appears.

Figure 32: Set NVMe Thermal Poll Interval Dialog



The dialog titled "Set NVMe Thermal Poll Interval" has a light gray header bar with the title and a close button (X). The main content area shows the label "NVMe Thermal Poll Interval" followed by a text input field containing the value "83", a spinner control, and the unit "Seconds". A blue "Save" button is positioned at the bottom left.

2. Enter the thermal poll interval, in seconds.
3. Click **Save**.

Download Serial Output Log

You can download the **Serial Output Log** file, which contains the firmware terminal log entries for the controller. The log information is shown as total number of entries available. Perform this step to download the Serial Output Log file.

Select **More Actions > Download Serial Output Log** on the Controller Dashboard.

The `Serial_Output_Log` file is downloaded.

Updating the Controller Firmware

The application lets you update the controller firmware.

NOTE


HPE does not recommend this method of updating the controller firmware. Instead use SPP, ILO, or PLDM firmware upgrade methods.

Perform these steps to update the controller firmware.

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

The **Update Firmware** dialog appears and displays the components, current version, selected version, current security version, and the selected security version, when applicable.

Figure 33: Update Firmware Window



3. Click **Browse** to locate and open the `.rom` file.

During a PSOC firmware upgrade (or downgrade), a message is displayed if the same firmware version exists on the controller or if the selected version is lower than the current version. Reconfirm the process by clicking the confirm check box (force option).

4. Click **Update**.
5. Select either **Online Activation** or **Offline Activation**.

By default, **Online Activation** is selected.

NOTE

The **Online Activation** and **Offline Activation** radio buttons are applicable only to MR9XX controllers. MR2XX and MR4XX controllers depend on the Online Firmware Update bit, which is usually set to **Online Activation**.

An **Offline Activation** requires a system reboot.

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

After the update is complete, a message appears to confirm the success of the update. The message also displays the new version of the controller firmware if the firmware is supported.

Firmware Activation Status

The LSI Storage Authority software enables you to monitor the component version information and firmware activation status.

Perform the following steps to access the component version information and firmware activation status.

1. Navigate to the Controller dashboard.
2. Click **More Actions**.

The **Firmware Activation Status** window appears. The **Firmware Activation Status** window displays the component version information and activation status.

Factory Repurpose

The application lets you clear the NVRAM.

NOTE


The Factory Repurpose function cannot be performed when the controller has drives that are configured in either an array or JBOD configuration.

1. Navigate to the Controller Dashboard.
2. Select **More Actions > Factory Repurpose** on the Controller Dashboard.

A message appears, prompting you to confirm your choice.

3. Select **Confirm > Yes**.

Figure 34: Perform Factory Repurpose Dialog



Factory Defaults

The HPE MR Storage Administrator software enables you to view the list of modified factory values and provides an option to restore the modified values to the default factory settings.

Perform the following steps to restore factory defaults.

1. Navigate to the Controller dashboard.
2. Select **Element(s) Actions > Factory Defaults**.

The **Factory Defaults** window appears. The **Factory Defaults** window displays the modified factory default properties.

Figure 35: Factory Defaults

[← Go back to Drive Group, Drives and Other Hardware list](#)

Factory Defaults [?](#)

Below are the list of modified factory default properties

Element	Properties	Current	DEFAULT
Controller	SES Association Type In MultiPath Config	LUN	Target Port
Controller	Patrol Read Percentage	70	30
Controller	BGI Percentage	80	30
Controller	Consistency Check Percentage	99	30

[Restore Factory Defaults](#)[View All](#)

- 3. (Optional) Click **View All** to view the default values.
- 4. Click **Restore Factory Defaults** to restore any modified settings back to the default factory setting.

MegaRAID Advanced Software Features

The MegaRAID Advanced Software (Premium) are features that the HPE MR Storage Administrator application supports on certain HPE Smart Array MR controllers.

The MegaRAID advanced software includes these features:

- MegaRAID FastPath
- RAID 5 and RAID 6


The MegaRAID software licensing authorizes you to enable the MegaRAID advanced software features. By default, the MegaRAID Advanced Software (Premium Features) is enabled.

The **Premium Features** option on the Controller Dashboard lets you use the MegaRAID Advanced Software features.

Perform these steps to use the advanced controller features:

1. Select **Actions > Premium Features** on the Controller Dashboard.
The **Premium Features** window opens.

Figure 36: Premium Features Window



Fast Path Advanced Software

The FastPath software is a high-performance I/O accelerator for solid state drive (SSD) arrays connected to a MegaRAID controller. This advanced software is an optimized version of MegaRAID technology that can dramatically boost storage subsystem and overall application performance. Particularly those that demonstrate high random read/write operation workloads – when deployed with a 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. Data on the drive is encrypted, so the drive becomes inaccessible to anyone attempting to access the drive 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 that 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.

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 SafeStore.

SED drives support reprovision, which transforms SED drives into unsecured and unlocked. The drive is cryptographically erased as part of this operation irrespective of the previous state of the drive.

Reprovision with Physical Presence SID (PSID) Revert allows drives to be cryptographically erased using a PSID key. The erase returns the drive to the factory default state.

Enable Security

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

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

Figure 37: Enable Security

Enable Drive Security ⓘ

Controller ID: 0 AVAGO MegaRAID SAS 9380-8e
Enabling drive security on this controller will have the option to create securevirtual drives using a security key.

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

Security Key Identifier:
AVAGO_SDS_SVS2876301_1ee1d712

Specify a security key identifier. The controller has provided a default identifier for you. You may use this string or enter your own identifier. If you have multiple security keys, the identifier will help you determine which security key to enter.

Security Key:
[Text Field]

Confirm:
[Text Field]

☐ Show Key

☐ Pause for password at boot time

☐ Enforce strong password security

Password:
[Text Field]

Confirm:
[Text Field]

☐ Show Password

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 at least 1 number, 1 lowercase letter, 1 uppercase letter and 1 non-alphanumeric character (e.g. >?@).

Optionally, You may enter a password to provide additional security. If you choose "Pause for password at boot time", you must enter it whenever you boot the server.
Note:
The password is case sensitive and must be between 8 and 32 characters.
If enforce strong password security is selected, then password field should contain at least 1 number, 1 lowercase letter, 1 uppercase letter and 1 non-alphanumeric character (e.g. >?@).

Are you sure you want to enable drive security?
☐ Confirm

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

- **Security Key Identifier** – The controller, by default, assigns a security key identifier.

However, you can change this security key identifier as per your requirement. If you have more than one security key identifier, the controller helps you to determine which security key identifier to enter.

- **Security Key** – Provides you with an option to create secure volumes by specifying the security key. The security key provided by you locks each SED drive attached to the controller.
- **Suggest Security Key** – Alternatively, you can click this option to have the system create a security key for you.
- **Password** – You can also specify a password to provide additional drive security.
- **Pause for password at boot time** and **Enforce strong password security** – If you select the **Pause for password at boot time**, you need to provide the password each time during the reboot to unlock the secure volume. To enter the password during the reboot, press **F9** to enter **System Utilities**, go to **System Configuration**. Select the controller and provide the password in the **Enter Boot Time Password** tab. 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:

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

NOTE

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

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

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

NOTE

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

The security key is case-sensitive. It must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one 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 select the **Pause for password at boot time**, you need to provide the password each time during the reboot to unlock the secure volume. To enter the password during the reboot, press **F9** to enter **System Utilities**, go to **System Configuration**. Select the controller and provide the password in the **Enter Boot Time Password** tab.
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.



CAUTION

Make sure to write down this password somewhere for future reference. If you are unable to provide the password when it is required by the system, you will lose access to your data.

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

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 MRSA.

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

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

The **Change Drive Security** dialog appears.

Figure 38: Change Drive Security Dialog



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

When LKMS is enabled, MRSA will show the current drive security mode as **FW/USER** instead of LKM.

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

Figure 39: Change Drive Security Dialog Options

Change Drive Security ⓘ

Controller ID: 0 AVAGO MegaRAID SAS 9300-8e
Drive security is currently enabled. Change the drive security settings on this controller or switch between key management.

Current drive security mode is NONE
☒ Change current security settings ☐ Switch to External Key Management(EKM) mode (EKM is currently not supported by LSA)

☒ **Use the existing security key identifier**
Current Security Key Identifier:

☐ **Enter a new security key identifier**
New 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 enter your own identifier. If you have multiple security keys, the identifier will help you determine which security key to enter.

☒ **Use the existing drive security key**
☐ **Enter a new drive security key**

Security Key:

Confirm:

☐ Show Key

☐ Pause for password at boot time
☐ Enforce strong password security

Password:

Confirm:

☐ Show Password

—Security Key—

The security key will be used to lock each self encrypted drive attached to the controller.
For maximum security, use 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 at least 1 number, 1 lowercase letter, 1 uppercase letter and 1 non-alphanumeric character (e.g. >?@).

—Password—

Optionally, you may enter a password to provide additional security. If you choose "Pause for password at boot time", you must enter it whenever you boot the server.
Note:
The password is case sensitive and must be between 8 and 32 characters.
If enforce strong password security is selected, then password field should contain at least 1 number, 1 lowercase letter, 1 uppercase letter and 1 non-alphanumeric character (e.g. >?@).

Are you sure you want to change the current security settings?

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 drive security key** option or select the **Enter a new drive security key** to specify a new security key and confirm once again.
5. Either click **Suggest Security Key** to have the system create a security key, or you can enter a new security key in the **Security Key** text field.
6. (Optional) – Select the **Show Key** 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 select this option, you need to provide the password each time during the reboot to unlock the secure volume. To enter the password during the reboot, press **F9** to enter **System Utilities**, go to **System Configuration**. Select the controller and provide the password in the **Enter Boot Time Password** tab.
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 volumes. 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 arrays 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 volumes on all of the secure arrays.

Perform the following steps to disable drive security:

1. In the Controller dashboard, select **More Actions > Disable Drive Security**.
A warning message appears asking for your confirmation.
2. Select **Confirm** and click **Yes, Disable Drive Security**.
The software disables drive security.

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

Managing Arrays


The HPE MR Storage Administrator application lets you monitor the status of arrays and spanned arrays (parity groups). The following lists the array states:

- Optimal: An array whose members are all online.
- Partially Degraded: An array with a redundant RAID level that can sustain one or more member disk failure.
- Degraded: An array with a redundant RAID level with one or more member failures and can no longer sustain a subsequent drive failure, with the exception in PRL-11 (RAID 10).
- Offline: An array with one or more member disk failures and data is no longer available

Viewing Array Properties

Select an array in the Controller Dashboard to view its properties.

Figure 40: Array Properties Window



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

Table 9: Array Properties Description

Property	Description
Free Capacity	Indicates the free space available in the array.
Secured	Indicates whether the array is secured.
Free Capacity Count	The number of holes present on the array. Selecting the count lists the holes and the capacity of the array.

Adding a Volume to an Array

You can add volumes to an existing array if sufficient storage space is present in the existing volumes of the array.

Perform these steps to add a volume to an existing array:

1. Navigate to the Controller Dashboard and click an array name (for example, **Array_1**).
In the right pane, under **Actions**, the **Add Volumes** option appears.
2. Click **Add Volumes**.
The **Volumes Settings** window appears.
3. Specify the settings for the volumes you want to create.
See [Selecting Volume Settings](#) for details on creating volumes.
4. Click **Volumes Settings**.
The newly created volume gets added to the selected array.

RAID Level Transformation

RAID level transformation is the process of converting one RAID configuration to another. You can perform RAID level transformation at the array level. The table that follows describes the valid RAID level transformation matrix.

Table 10: Array – RAID Level Transformation Description

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

Migrating the RAID Level of an Array

Perform these steps to migrate the RAID level of an array.

1. Navigate to the Controller Dashboard and click an array name (for example, **Array_1**).
In the right pane, under **Actions**, the **Modify Array** option appears.
2. Click **Modify Array**.
The **Modify Array** window appears.

Figure 41: Modify Array Window

Modify Array ⓘ
Step 1/2: Choose your array settings

Next

Array 0

1. RAID Level Setting (Compare and select)

RAID 0

This RAID level is suitable for high performance with zero data redundancy. Choose this option only for non-critical data.

☒ It is advisable to backup data before you proceed. Are you sure you want to continue?

Next

3. Select the RAID level to which you want to migrate the array from the **RAID Level Setting** drop-down menu.

It is recommended you back up the data *before* you change the RAID levels.

ATTENTION

Checking the “**It is advisable to back up data before you proceed. Are you sure you want to continue?**” checkbox does *NOT* launch a backup. You must follow the prescribed process to perform an array backup.

4. Click **Next**.

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

Figure 42: Modify Array Settings Dialog

Adding Drives to a Configuration

For example, if you migrate the RAID level of a array from RAID 0 to RAID 5, the **Modify Array** wizard lets you add unconfigured drives to the existing configuration to enable the RAID level transformation.


1. Click **Add Drives** in the **Modify Array** window.

NOTE

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

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

Figure 43: Available Unconfigured Drive Window



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

The RAID level is migrated. A confirmation message appears. You can monitor the progress of the transformation. See [Background Operations Support](#).

Removing Drives from a Configuration

For example, if you migrate the RAID level of a array from RAID 5 to RAID 0, the **Modify Array** wizard lets you remove drives from the existing configuration to enable the RAID level transformation.

1. Select **Remove drives** in the **Modify Array** window, and click **Next**.

The **Modify Array** window appears and it states the number of drives that you must remove to change the RAID level from the current level to a new RAID level and the maximum number of drives that can be removed.

2. Click the  icon to remove the drives.
3. Click **Finish**.

The RAID Level is migrated. A confirmation message appears. You can monitor the progress of the transformation. See [Background Operations Support](#).

Migrating the RAID Level Without Adding or Removing Drives

For example, if you migrate the RAID level of your array from RAID 5 to RAID 0, the **Modify Array** wizard lets you migrate the RAID level without adding or removing the drives.

Select **Migrate RAID level** in the **Modify Array**, and click **Next**.

The RAID level is migrated. A confirmation message appears. You can monitor the progress of the transformation. See [Background Operations Support](#).

Managing Volumes

The HPE MR Storage Administrator application lets you perform various operations on the volumes.

The firmware supports the following volume states on the controller:

- **Optimal** – A volume whose members are all online.
- **Partially Degraded** – A volume with a redundant RAID level that can sustain one or more member disk failure. This state also applies to the volume’s member drives. Currently, a RAID 6 or RAID 60 volume is the only volume that can be Partially Degraded.
- **Degraded** – A volume with a redundant RAID level with one or more member failures and can no longer sustain a subsequent drive failure, with the exception in PRL-11 (RAID 10).
- **Offline** – A volume with one or more member disk failures and data is no longer available (corresponds to DDF Failed state).

Viewing Volume Properties

Select a volume from an array in the Controller Dashboard to view its properties.

Figure 44: Volume Properties

Volume Properties			
State Optimal	Current Read Cache Status No Read Ahead	Default Read Cache Policy No Read Ahead	Current Write Cache Status Write Back
Default Write Cache Policy Write Back	Current IO Status Direct IO	Default IO Policy Direct IO	Access Policy Read Write
Write Cache Status Enabled	Drive Write Cache Policy Disabled		

Table 11: Volume Properties

Property	Description
State	The current status of the volume. These options are available: • Optimal • Partially Degraded • Degraded • Offline
Current Read Cache Status	The current read cache status for the volume. These options are available: • Read Ahead • No Read Ahead
Default Read Cache Status	The default read cache status for the volume. These options are available: • Read Ahead • No Read Ahead
Current Write Cache Status	The current write cache status for the volume. These options are available: • Write Through • Write Back • Always Write Back
Default Write Cache Status	The default write cache status for the volume. These options are available: • Write Through • Write Back • Always Write Back
Default IO Policy	The input/output policy for the volume. These options are available: • Direct IO
Access Policy	The access policy for the volume. These options are available: • Read Write • Read Only • Hidden – The Hidden policy is applicable for only hidden volumes. No other access policies are applicable after you select Hidden as the access policy.
Write Cache Status	The current state of the write-back cache. The status depends on the battery status and the controller status for write-back operations. • Enabled – When the current cache policy is <code>Write Back</code>. • Disabled – The current cache policy is <code>Write Through</code>, <code>JBOD</code>, or <code>Always Write Back</code>. • Temporarily Disabled – The firmware is moving to <code>Write Back</code> or <code>Write Through</code>. For example, pinned cache, the firmware is downloading, or reconstruction and charging is occurring.
Drive Write Cache Policy	The volume cache setting. These options are available: • Unchanged • Enable • Disable

Modifying Volume Properties

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

1. Navigate to the Controller Dashboard, click an array name (for example, **Array_1**).

Click the  icon that corresponds to the array to display its contents.

The volumes and drives associated with the selected array appear.

2. Click the volume whose settings you want to change.
3. Select **More Actions > Modify Properties**.

The **Modify <Volume Name>** dialog appears.

Figure 45: Modify Volume Dialog

Modify Volume:
volumeName_00 Properties

Volume Name
volumeName_00

Read Policy
No Read Ahead

Write Policy
Write Back

IO Policy
Direct IO

Drive Write Cache Policy
Disabled

The IO policy applies to reads on a specific Volume. It does not affect the read ahead cache.

☒ **Direct IO**
In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. This option is the default setting.

Save settings

4. Change the volume properties as needed.
For information about these properties, see [Selecting Volume Settings](#).
5. Click **save settings**.

Start and Stop Locating a Volume

If the drives that contain the volumes are located in a disk enclosure, you can identify them by making their LEDs blink. Perform these steps to identify the volumes:

1. Navigate to the Controller Dashboard, click an array name (for example, **Array_1**).

Click the  icon that corresponds to the array to display its contents.

The volumes and drives associated with the selected array appear.

2. Click the volume that you want to locate in the disk enclosure.

3. Select **Actions > Start Locate**.

The LEDs on the drives in the volume start blinking.

4. To stop the LEDs from blinking, select **Actions > Stop Locate**.

Erasing a Volume

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

NOTE

Use disk management tools within the operating system to first unmount the volume before performing an erase.

Perform these steps to erase a volume.

1. Navigate to the Controller Dashboard, click a array name (for example, **Array**).

Click the  icon corresponding to a array to display its contents.


The volumes and drives associated with the selected array appear.

2. Click the volume whose content you want to erase.

3. Select **Actions > Erase**.

The **Volume Erase** dialog appears.

Figure 46: Volume Erase Dialog



The dialog box is titled "Volume Erase" with a question mark icon and a close button (X). Below the title bar, there is a descriptive text: "Volume Erase operates on a specified Volume and overwrites all user-accessible sectors with the specified pattern for the specified number of passes." Below this, a section titled "Select the mode for Drive erase operation :" contains three radio button options: "Simple" (selected), "Normal", and "Thorough". Each option has a description of the erase operation. Below these options is a checkbox labeled "Delete Volume After Erase". At the bottom of the dialog is a blue button labeled "Erase Logical Drive".

Volume Erase ⓘ ✕

Volume Erase operates on a specified Volume and overwrites all user-accessible sectors with the specified pattern for the specified number of passes.

Select the mode for Drive erase operation :

☒ **Simple**
Specifies single pass erase Operation that writes pattern A to the Volume.

☐ **Normal**
Specifies a three pass erase operation that first overwrites the Volume contents with random values then overwrites it with pattern A and then overwrites it with pattern B

☐ **Thorough**
Specifies a nine pass erase operation that repeats the normal mode thrice.

☐ **Delete Volume After Erase**

Erase Logical Drive

The dialog shows these modes:

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

Initializing a Volume

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

Perform these steps to initialize a volume after completing the configuration process.

1. Navigate to the Controller Dashboard, click an array name (for example, **Array_1**).
Click the  icon that corresponds to the array to display its contents.
The volumes and drives associated with the selected array appear.
2. Click the volume that you want to initialize.
3. Select **Actions > Start Initialize**.
A warning message appears.
ATTENTION
Initialization erases all data on the volume. Make sure to back up any data you want to keep before you initialize a volume. Make sure the operating system is not installed on the volume 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 volume.
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 Volume

Perform the following steps to start consistency check on a volume. For more information of consistency check, see [Running Consistency Checks](#).

1. Navigate to the Controller Dashboard, click an array name (for example, **Array_1**).
Click the  icon that corresponds to that array to display its contents.
The volumes and drives associated with the selected array appear.
2. Click the volume 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.

Expanding the Online Capacity of a Volume

The Online Capacity Expansion (OCE) function lets you expand the capacity of a virtual disk by adding new physical disks or using unused space on existing disks, without requiring a reboot.

ATTENTION

Make sure to back up the data on the volume before you proceed with the OCE.

Expanding the Online Capacity of a Volume for MR200/MR4000 Controllers


Perform the following steps to expand the capacity of a volume.

ATTENTION

This feature is available only for MR200/MR4000 controllers. For previous generations, the online capacity of a volume cannot be expanded when the existing array size is used. See [RAID Level Transformation](#) for information on expanding the capacity by doing a RAID level migration.

1. Navigate to the Controller Dashboard, click an array name (for example, **Array_1**) then click the  icon that corresponds to an array to display its contents. .
The volumes and drives associated with the selected array appear.
2. Click the volume whose capacity you want to expand.
3. Select **More Actions > Expand**.
The **Expand Volume** dialog appears.

Figure 47: Expand Volume Dialog



4. Select the percentage of the available capacity that you want the volume to use.
5. Click **Expand**.
The volume expands by the selected percentage of the available capacity.

Expanding the Online Capacity of a Volume for MR932 Controllers

Perform the following steps to expand the capacity of a volume.

ATTENTION


This feature is available only for MR932 controllers.

1. Navigate to the Controller dashboard, select an array name (for example, **DG_1**) then click the  icon that corresponds to an array to display its contents.
The volumes and physical drives that are associated with the selected array appear.

- 2. Select the volume whose capacity you want to expand.
- 3. Select **More Actions > Expand**.


The **Expand Volume** dialog appears.

Figure 48: Expand Volume Dialog




- 4. Select the **It is advisable to backup data before you proceed. Are you sure you want to continue?** check box and click **Next**.
- 5. Click **Add Physical Drives** and select an unconfigured drive.

Figure 49: Available Unconfigured Drive(s) Dialog



- 6. Click **Finish**.

Figure 50: Expand Volume - Add Drives Dialog



Deleting a Volume

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



CAUTION

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

Perform the following steps to delete a volume.

1. Navigate to the Controller Dashboard, click an array name (for example, **Array_1**).

Click the  icon that corresponds to the array to display its contents.

The volumes and drives associated with the selected array appear.

2. Click the volume 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.

NOTE

You can delete an operating system or file system volume. However, if you try to do so, the following message appears.

Selected Volume has an OS/FS, are you sure you want to delete it?

Managing Drives

The HPE MR Storage Administrator application lets you manage all the drives connected to the controller.

The firmware defines the following states for the physical disks connected to the controller:

- **Unconfigured Good** – A drive accessible to the RAID controller but not configured as a part of a virtual drive.
- **Online** – A disk accessible to the RAID controller and configured as part of a virtual drive.
- **Failed** – A disk drive that is part of a virtual drive, but has failed and is no longer usable.
- **Rebuild** – A disk drive to which data is being rebuilt to restore full redundancy to a virtual drive.
- **Unconfigured Bad** – A disk drive that is not a part of an array and is known to be bad. This state is typically assigned to a drive that has `Failed`, but is no longer part of a configured virtual drive, because it has been replaced by a spare drive.

On hot-plug, or discovery during startup, if the firmware cannot communicate with a drive, or drive initialization fails, the drive is marked `Unconfigured Bad`. A hot-plugged drive might continue to be in the `Unconfigured Bad` state, if it had `Failed` previously. The firmware can be configured to remember the `Failed` drive information.

- **Shield state** – The shield state is an interim state of a physical drive for diagnostic operations. The results of the diagnostic tests determine if the physical drive is good or bad. If any of the diagnostic tests fail, the physical drive transitions to a bad state (`FAILED` or `UNCONF BAD`).
- **Foreign** – While importing disks from a different RAID controller (foreign metadata), the physical disk is marked as Foreign until the configuration on the disks is added to the existing configuration on the controller.

Foreign is not actually a drive state, it indicates that a drive is derived from another configuration. The foreign drives typically remain in an `Unconfigured Good` state until they are imported into the current configuration. A Foreign drive is any disk that has a disk data format (DDF) configuration record and is not a part of the current set of configured disks. Even if the disk is removed from the current configuration, it is still considered Foreign until it is imported. A Foreign drive cannot be configured as a part of a virtual drive or a spare unless the Foreign configuration on the drive has been explicitly cleared by the user.

Users migrating the volumes from an old controller to a new controller must ensure that all the controller settings and NVDATA settings that they had earlier on the old controller are present in the new controller because if incompatible settings are found, the foreign import of the drive might fail.

- **Spare** – A disk drive that is configured as a spare. If the spare is not activated, the status light emitting diode (LED) state corresponds to `Online`.
- **Copyback** – A disk drive serving as a Copyback destination drive. The drive's state transitions to `Online` when the Copyback operation completes and the source drive transitions to an unconfigured good or bad drive.
- **Offline** – A disk drive that is still part of a configured volume, but is not active now. This state is used to represent a configured drive for which the data is not valid. This state can occur as a transition state, or because of any action performed by the user.
- **Just a Bunch of Disks (JBOD)** – Drives that are marked JBOD cannot be part of any configuration because they are stand-alone drives that are exposed to the operating system. Because the operating system and/or applications manage these drives through the controller, it is not appropriate for the firmware to perform RAID operations on JBOD drives. To include a JBOD drive in a RAID configuration, its state must be transitioned to `Unconfigured Good`. If users enable the JBOD support, the firmware marks new drives as JBOD unless the drive contains a valid DDF record.

NOTE

The **Make Failed** drive operation is not available for MR200/MR400 controllers.

Viewing Drive Properties

Select a drive from an array in the Controller dashboard to view its properties.

Figure 51: Drive Properties

0 Foreign Drives							
1 Unconfigured Drives							
1 Unconfigured good							
4 Configured Drives							
4 Online							
Enclosure : Bay	Device ID	Type	Interface	Capacity	Sector Size	Status	Model
Port 3i,Box1, Bay=1	56	HDD	SAS	300GB	512B	Online	EG000300JWBHR
Port 3i,Box1, Bay=2	55	HDD	SAS	300GB	512B	Online	EG000300JWBHR
Port 3i,Box1, Bay=4	57	HDD	SAS	300GB	512B	Online	EG000300JWBHR
Port 4i,Box1, Bay=5	58	HDD	SAS	300GB	512B	Online	EG000300JWBHR
0 Spares							
0 JBOD							

Actions

- Make Drive Offline
- Start Locating
- Stop Locating
- Replace Drive

Properties

Status
Online

Exposed As
PHYSICAL-DEVICE

Product ID
EG000300JWBHR

Vendor ID
HP

Serial Number
3770A05JFXFD1710

Shield Counter
0

Device ID
56

Usable Capacity
279.87GB

Raw Capacity
300GB

[more properties](#)

Table 12: Drive Properties

Property	Description
Status	The current status of the drive.
Exposed As	To differentiate the drives, the drives are exposed as one of the following: - JBOD - PHYSICAL-DEVICE
Product ID	The product ID of the drive.
Vendor ID	The ID assigned to the drive by the vendor.
Serial Number	The serial number of the drive.
Shield Counter	The shield counter value.
Device ID	The device ID of the drive that is assigned by the manufacturer.
Usable Capacity	The usable storage capacity, based on the RAID level used.
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 drive.
SAS Address 1	The WWN for the drive.
Negotiated Link Speed	The negotiated link speed for data transfer to and from the drive.
Drive Speed	The speed of the drive.
Temperature	The temperature of the drive.
Revision Level	The revision level of the drive's firmware.
Power Status	The power status displays the following status: On – when a drive is spun up.
Native Command Queueing	Indicates if the Native Command Queueing (NCQ) function is enabled. NCQ enables the drive to queue the I/O requests and reorder them for efficiency.
Sector Size	The size of the sector of the drive. The possible options are 4 KB or 512 B.
Enclosure Properties	

Property	Description
Enclosure ID	The ID of the enclosure in which the drive is located.
Enclosure Location	The port number of the enclosure to which the drive is connected.

Locating Tape Drives

If your system is connected to tape drives, the application 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 52: 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 drives are in a disk enclosure, you can identify them by making their LEDs blink. Perform the following steps to identify the drives:

1. Navigate to the drive on the Controller dashboard, and select the drive you want to identify such as, Unconfigured Good drive, Online drive, Configured drive, and so on.
2. Select **Actions > Start Locating**.
The corresponding LED on the drive starts blinking.
3. To stop the LED from blinking, select **Actions > Stop Locating**.

Making a Drive Offline

Perform the following steps to make a drive offline.

ATTENTION

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

1. Navigate to the Controller Dashboard, click an array name (for example, **Array_1**).
Click the  icon corresponding to an array to display its contents.
The volumes and drives associated with the selected array appear.
2. Click the **Drive** tab, and select the drive that you want to make offline.
3. Select **Actions > Make Drive Offline**.
A confirmation message appears.
4. Select **Confirm** and click **Yes, Make Drive Offline** to make the selected drive *Offline*.

Making a Drive Online

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

ATTENTION

When transitioning a drive to an online state manually or forcefully, you lose data on the drive. When adding a drive to an existing volume or replacing a drive in an existing volume, you must manually start the rebuild process, if the rebuild does not automatically start. When the rebuild process is completed, the drive automatically transitions to an online state.

1. Navigate to the Controller Dashboard, click an array name (for example, **Array_1**).

Click the  icon corresponding to an array to display its contents.

The volumes and drives that are associated with the selected array appear.

2. Click the **Drive** tab, and select the offline drive that you want to make online.

3. Select **Actions > Make Drive Online**.

The drive status changes to *Online*.

Replacing a Drive

You might want to replace a drive if the drive shows signs of failing. Before you start this operation, be sure that an available unconfigured good replacement drive is available. The replacement drive must have at least as much capacity as the drive you are replacing. Perform the following steps to replace a drive.

ATTENTION

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

1. Navigate to the Controller dashboard, click a array name (for example, **Array_1**). Click the  icon corresponding to a array to display its contents.

The volumes and drives associated with the selected array appear.

2. Click the **Drive** tab, and select a drive which you want to replace.

3. Select **Actions > Replace Drive**.

The **Replace Drive** dialog appears.

Figure 53: Replace Drive

Replace Drive

Port 3I,Box=1,Bay=2 will copy the data to selected component.

	Enclosure	Device ID	Interface	Type	Capacity	Sector Size	Model
☐	Port 3I,Box=1,Bay=4	57	SAS	HDD	300GB	512B	EG000300JWBHR
☐	Port 4I,Box=1,Bay=5	58	SAS	HDD	300GB	512B	EG000300JWBHR
☐	Port 3I,Box=1,Bay=3	59	SAS	HDD	300GB	512B	EG000300JWBHR

Replace Drive


4. Select a replacement drive and click **Replace 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 **Arrays**.
2. Click an array name (for example, **Array_1**).
3. Click the  icon that corresponds to an array to display its contents.
The volumes and physical drives associated with the selected array appear.
4. Click the **Physical Drive** tab, and select a drive which you want to mark as missing.
5. Select **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 54: Mark Drive Offline Dialog



- 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 **Actions > Mark Drive as Missing**.
- 9. Navigate to the **Drives** tab.
- 10. Select an Unconfigured Good drive from the list of Unconfigured Good drives, and go to **Actions > More Actions > Replace Missing Drive**.

The **Replace Missing Drive** dialog appears.

Figure 55: Replace Missing Drive Dialog



11. Select the drive and click **Replace Missing Drive**.
12. Navigate to the **Arrays** tab and select a new drive.
13. Click **Actions > More Actions > Start Rebuild**.

Replacing a Missing Drive

1. Navigate to the Controller dashboard and select **Arrays**.
2. Create a new array for any RAID level with two drives.
3. Navigate to the array, 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 array where you want to replace the missing physical drive.
7. Click **Ok**.

Viewing Protected Arrays

Perform the following steps to view a list of protected arrays.

1. Navigate to the Controller dashboard, and click the **Drives** tab.
All of the associated drives appear.
2. Expand **Hot Spares**, and select a PD.

Figure 56: Show Protected Arrays

The screenshot shows the Dell iDRAC Controller dashboard with the **Drives** tab selected. The top navigation bar includes links for **1 Arrays**, **7 Drives** (3 Unconfigured Drives), **2 Other Hardware** (Includes Energy Pack), and **Physical Function Information**.

The main content area displays a summary of drives:

- 4 Unconfigured Drives**: 2 Unconfigured Good, 1 Unusable (Unknown), 1 Unconfigured Bad (Degraded). A **Filter** button is present.
- 2 Configured Drives**: 2 Online.
- 1 Spares**: 1 Dedicated spare. A **Filter** button is present.
- 0 JBOD**.

The table below lists the drives:

Enclosure : Bay	Device/Persistent ID	Media	Interface	Capacity	Sector Size	Status	Model	NS/LU Count
Port1,Box=1,Bay=2	292	SSD	SAS	1.92TB	512B	Unconfigured Good	KPM7XRUG1T92	1
Port1,Box=1,Bay=3	293	Unknown	NVMe	0KB	0KB	Unusable (Unknown)		0
Port1,Box=1,Bay=4	294	SSD	SAS	1.92TB	512B	Unconfigured Good	KPM7XRUG1T92	1
Port1,Box=1,Bay=8	298	HDD	SATA	1TB	512B	Unconfigured Bad (Degraded)	MM1000GBKAL	1

Below the table, the **1 Spares** section shows a single drive:


Enclosure : Bay	Device/Persistent ID	Media	Interface	Capacity	Sector Size	Status	Model	NS/LU Count
Port1,Box=1,Bay=1	291	SSD	SAS	1.92TB	512B	Dedicated spare	KPM7XRUG1T92	1

On the right side, the **Element(s) Actions** menu is open, showing the following options:

- Show Protected Array
- Unassign Dedicated Spare
- Start Locate
- Stop Locate

3. Select **Element(s) Actions > Show Protected Arrays**.
The **Arrays** lists appears.

Figure 57: Protected Arrays



The screenshot shows a dialog box titled "Arrays" with a close button (X) in the top right corner. Below the title bar, it says "Port1,Box=1,Bay=1 protected drive group list". Inside the dialog is a table with three columns: "Arrays", "Raid Level", and "Configured Capacity". There is one row in the table with the values "Array_0", "Raid 1", and "399.532GB".

Arrays	Raid Level	Configured Capacity
Array_0	Raid 1	399.532GB

Assigning Global Spare Drives

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

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

Removing a Global Spare Drive

Perform the following steps to remove a global spare drive.

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

Assigning Dedicated Spare Drives

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

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

4. Select an array and click **Add Dedicated Spare Drive**.

A confirmation message appears.

5. Click **Done**.

The unconfigured good drive is changed to a dedicated spare drive. The status of the unconfigured good drive appears as a dedicated spare drive in the **Spare Drives** section.

Rebuilding a Drive

If a drive configured as RAID 1, 5, 6, 10, 50, or 60 fails, the firmware automatically rebuilds the data on a spare drive to prevent data loss. The rebuild operation 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.

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

All of the associated drives appear.

2. Expand **Unconfigured Drives** and select an unconfigured bad drive.

3. Select **Actions > Make Unconfigured Good**.

A confirmation message appears.

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

The unconfigured bad drive is changed to an 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 drive for removal spins the drive into a power save mode.

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

All of the associated drives appear.

2. Expand **Unconfigured Drives**, and select a drive that you want to remove.

3. Select **Actions > Prepare for Removal**.

The drive is in the power save mode and is ready for removal.

4. Wait until the drive spins down and then remove it.

If you do not want to remove the drive, select **Actions > Undo Prepare for Removal**.

Make Unconfigured Good Drives and Make JBOD Drives

Unconfigured Good: The drive is unconfigured and hidden from the OS.

JBOD: The drive is exposed to the host OS as a physical drive. The user cannot use a JBOD drive to create a RAID configuration, because it is exposed to the host OS. A JBOD drive has to be converted to Unconfigured Good before creating a RAID configuration on the drive.

NOTE

The default JBOD behavior for the MR400 controller was changed in firmware 52.26.3-5250 and later versions. See [Personality Management](#) for details.

Making Unconfigured Good Drives


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

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

Making a JBOD Drive

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

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**.



- **Secure** (Default) – Make the physical disk secure.

NOTE

Secure is checked by default when one UG physical disk is already secured and when the user wants to convert the drive to JBOD.


- **Cryptographic Erase (FIPS compliant, Warning Loss of Data)**. – Erase the PD and convert to JBOD.

Erasing a Drive

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

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

Figure 58: Drive Erase Dialog



The dialog shows the following modes:

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

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

A warning message appears asking for your confirmation.

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

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

Erasing a Drive Securely

ATTENTION

The following procedure is applicable only to MR416i-p, MR416i-o, MR416i-a, MR216i-p, MR216i-o, MR216i-a, MR408i-o, and MR932i-p.

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 **Actions > Instant Secure Erase**.
A confirmation message appears.
4. Select **Confirm** and click **Yes, Securely Erase Drive** to proceed with the operation.
After the secure erase operation has started, the **Stop Erase** option is enabled in the **Actions** menu. You can monitor the progress of the erase operation. See [Background Operations Support](#).

Sanitizing a Drive

You can erase the data residing on a drive using the **Sanitize** feature. The **Sanitize** option is similar to the *Drive Erase* feature that is already supported by your controller, except that the **Sanitize** option is performed by the drive firmware, whereas the *Drive Erase* feature is performed by the controller firmware.

The Sanitize option is an industry standard SCSI feature. It uses industry standard Sanitize SCSI Block command. The Sanitize operation is constantly monitored the by controller firmware and the drive sanitization progress events are notified to you through Background Operations Support.

ATTENTION

The following procedure is applicable only to P824i-p and newer controllers.

To Sanitize a drive, you must make sure that:

- The selected drive is in an Unconfigured Good state.
- The selected drive is not a JBOD drive.
- The selected drive is not part of any array, dedicated spare drive, or global spare drive.

Sanitize operation is enabled only when no other operation is in progress on the selected drive.

When the Sanitize operation is in progress, you cannot perform any other operation on the drive that is being sanitized.


Perform the following steps to sanitize a drive:

1. Navigate to the Controller dashboard and click the **Drives** tab.
All of the associated drives appear.
2. Expand **Unconfigured Drives** and choose an unconfigured good drive.
 - You can run drive sanitization on multiple Unconfigured Good drives at the same time.
However, the Sanitize option is only enabled when the same type of sanitize operation is supported on all the selected drives. For example, on solid state drives (SSDs), **Block Erase** is allowed, and on hard disk drives (HDDs), **Overwrite** is allowed.
 - You cannot run the Sanitize operation on mixed drive types.
For example, you have selected two drives to run the Sanitize operation; one of them is an SSD and the other one is an HDD. In this scenario, you will not be able to run the Sanitize operation because they are not the same drive type, nor are they of the same sanitize operation type.
3. Select **Actions > More Actions > Start Sanitize**.
The **Sanitize Drive** dialog appears.

NOTE

After you start the drive sanitize operation, you cannot stop or pause the operation until it is complete.

Figure 59: Sanitize Dialog

The image shows a 'Sanitize Drive' dialog box. At the top, the title 'Sanitize Drive' is in bold, followed by a question mark icon and a close 'X' icon. Below the title bar, a descriptive text states: 'Sanitize operation causes the device server to alter information on the logical unit'. A horizontal line separates this from the next section, which is titled 'Select the mode for sanitize operation'. Under this title, there are two options. The first is 'Overwrite', which is selected with a radio button. Its description is: 'Overwrite operation causes the device server to alter information by writing a data pattern to the medium one or more times;'. The second option is 'Allow Unrestricted Sanitize Exit', which is not selected (checkbox is empty). At the bottom of the dialog, there is a blue button labeled 'Sanitize Drive'.

Depending on the drives you have selected for sanitization (SSDs or HDDs), the following options are available:

- **Overwrite** – If you have selected HDD, you can sanitize the physical using the Overwrite option.
This option writes a particular data pattern on the drive one or more times.
- **Block Erase** – If you have selected SSDs, you can sanitize the drives using the Block Erase option.
This option sets the physical blocks on the drive to a vendor-specific value.
- **Allow Unrestricted Sanitize Exit** – If, for some reason, the Sanitize operation fails, the system tries to bring the drive out of the failure mode irrespective of whether you select this check box not.
However, if this check box is selected, and if the system succeeds in bringing the drive out of the failure mode, the drive is then returned as an Unconfigured Good drive. If you do not select this check box, and if the Sanitize operation fails, the system places the drive in an Unconfigured Bad state.

4. Click **Sanitize Drive**.


A confirmation message appears.

5. Click **Yes, Sanitize Drive(s)** to start sanitizing the selected drives.

You can monitor the progress of the Sanitize operation in the Background Operations section. The status of the drive is also displayed as **Sanitize** until the sanitization operation completes.

The following figure displays the Background Operations section where the sanitize operation is in progress. The figure also displays the status of the drive that is being sanitized.


Figure 60: Background Operations and Drive Sanitize Dialog



Managing Hardware Components

When you select the **Other Hardware** tab from the Controller dashboard, the hardware components window appears.

Figure 61: Other Hardware Window




Monitoring the HPE Smart Storage Energy Pack

When the application is running, you can monitor the status of the HPE Smart Storage Energy Pack.

Also, if the HPE Smart Storage Energy Pack is in an Optimal state, **WriteCache Policy** is enabled. If the HPE Smart Storage Energy Pack is in not in an optimal state, **WriteCache Policy** is disabled.

To view the **WriteCache Policy** status, go to the **Array** tab, select an **Array**, then select a **Volume**. The **WriteCache Policy** status is displayed under the **Properties** section, as shown in the following figure.

Figure 62: WriteCache Policy Window




Monitoring Enclosures

When the HPE MR Storage Administrator application 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 **Box** to view its properties.

Figure 63: Enclosure Properties Window



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

Table 13: Enclosure Properties

Property	Description
Name	Indicates the name of the enclosure.
Bay count	Indicates the number of bays.
Location	Indicates the location of the enclosure.

Physical Function Information


Depending on the allowed operations, the application displays the physical function information.

NOTE

MRSA is used exclusively for managing supervisor controllers and managing physical functions can only be accessed through the supervisor controller.

1. Navigate to the Controller dashboard, and click the **Physical Function Information** tab
2. Expand the **Physical Function Information**.

Figure 64: Physical Function Information Dialog



3. Click to open the **Physical Function Properties**.
The **Physical Function Properties** window appears.

Figure 65: Physical Function Properties

Physical Function Properties			
Function ID	Name	AU Assign Capable	Max MSIX Vector
0	Supervisor	No	65

Operations on Physical Functions

Perform these steps to modify physical functions.

- 1. Navigate to the Controller Dashboard.
- 2. Select the **Physical Functions Information > Element(s) Actions**.

Figure 66: Physical Function Information

1 Critical Issue(s), 2 Need Attention

Controller Information

Configured Capacity 3.84 TB of 9.078 TB

Serial Number
PYPTHX08KCY00N

SAS Address
0x500062b222801600

Driver Version
8.16.00.00

2 Arrays
2 Volumes

7 Drives
2 Unconfigured Drives

2 Other Hardware
Includes Energy Pack

Physical Function Information

4 Physical Functions

Function ID	PCI Info	Supervisor	State	Is Current	Function Type	IO Capable	Max Queues	
0	0x0000:11:00:0	Yes	Enabled	Yes	MPI	Yes	256	***
1	0x0000:12:00:0	No	Enabled	No	MPI	Yes	256	***
2	0x0000:13:00:0	No	Enabled	No	MPI	Yes	256	***
3	0x0000:14:00:0	No	Enabled	No	MPI	Yes	256	***

Controller Actions

Personality Management

Function Profile Management

Configure

Auto Assign

Update Firmware
Package Version 8.15.0.150-01003-00005

Show Events

More Actions

Selected Element(s): 1

Element(s) Actions

Modify Physical Function

View/Remove Mapped Volumes

Map Volumes

View/Remove Mapped Drives

Map Drives

- 3. Click **Physical Function Row > Select Actions**.
- 4. Choose one of the following actions to modify.
 - Modify Physical Function

×

Modify Physical Function

Physical Function Name

Supervisor

Physical Function Enabled

Specify Physical Function Status

☒ **Enabled**
 Enable Physical Function

☐ **Disabled**
 Disable Physical Functon

Save settings

- Map Volumes

×

Map Volumes

Map's selected Volumes to the Physical Function

		ID ↕	Name ↕	Capacity ↕	Raid Level ↕
☒		3	volumeName_00	1.919TB	0

Map Volumes

- View or Remove Mapped Volumes

Remove Mapped Volumes

Remove selected Volumes from the Physical Function

		ID	Name	Capacity	Raid Level
☒		3	volumeName_00	1.919TB	0

Remove Mapped Volumes

- Map Drives

Map Drives

Map's selected Drives to the Physical Function

☐	Port,Box,Bay	Device/Persistent ID	Media	Interface	Capacity	Sector Size	Model
☐	[Port1],Box=1,Bay=7	297	SSD	SAS	400.088GB	512B	EO0400JEFPE

Map Drives

- View or Remove Mapped Drives

Remove Mapped Drives

Remove selected Drives from the Physical Function

☐	Port,Box,Bay	Device/Persistent ID	Media	Interface	Capacity	Sector Size	Model
☐	[Port1],Box=1,Bay=7	297	SSD	SAS	400.088GB	512B	EO0400JEFPE

Remove Mapped Drives

NVMe Drive Format

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

2. Expand **Unconfigured Drives**, and select a drive to reconfigure.
3. Select **Element(s) Actions > Start Reconfigure**.
4. Update the fields in **Reconfigure** dialog.
5. Click **Reconfigure**.

NOTE

The **Reconfigure** button is only enabled if the user changes any values.

Viewing Event Logs

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

1. Select **More Actions > View Event Log** on the Server or Controller dashboard.

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

Figure 67: View Event Log Window

Go back to Array, Drives and Other Hardware list

Close

Show Events ?

Displaying Latest Event Entries

Severity Level	Event Id	Locale	Description	Time, Date
Critical	251	Volume	PCI Slot: 12 Volume is now DEGRADED Volume 236	10:02:25 PM,14 Nov'2024
Information	81	Volume	PCI Slot: 12 State change on Volume: 236 Previous: Optimal; Current: Degraded;	10:02:25 PM,14 Nov'2024
Information	114	Physical Device	PCI Slot: 12 Port 21,Box=1,Bay=3 (DeviceId: 23) - Drive: State change - Previous: Online, Current: Offline	10:02:25 PM,14 Nov'2024
Information	370	Volume	PCI Slot: 12 Volume is available. Volume: 236	10:01:22 PM,14 Nov'2024
Information	138	Volume Configuration	PCI Slot: 12 Created Volume: 236	10:01:22 PM,14 Nov'2024
Information	249	Volume	PCI Slot: 12 Volume is now OPTIMAL Volume 236	10:01:22 PM,14 Nov'2024
Information	114	Physical Device	PCI Slot: 12 Port 21,Box=1,Bay=4 (DeviceId: 20) - Drive: State change - Previous: UnConfigured Good; Current: Online	10:01:22 PM,14 Nov'2024
Information	114	Physical Device	PCI Slot: 12 Port 21,Box=1,Bay=3 (DeviceId: 23) - Drive: State change - Previous: UnConfigured Good; Current: Online	10:01:22 PM,14 Nov'2024
Information	114	Physical Device	PCI Slot: 12 Port 21,Box=1,Bay=3 (DeviceId: 23) - Drive: State change - Previous: Online, Current: UnConfigured Good	10:01:03 PM,14 Nov'2024
Information	139	Volume Configuration	PCI Slot: 12 Deleted Volume: 237	10:01:03 PM,14 Nov'2024
Information	114	Physical Device	PCI Slot: 12 Port 21,Box=1,Bay=2 (DeviceId: 22) - Drive: State change - Previous: Offline, Current: UnConfigured Good	10:00:27 PM,14 Nov'2024
Information	231	Physical Device Configuration	PCI Slot: 12 Port 21,Box=1,Bay=2 (DeviceId: 22) - Drive: Marked Missing on array: 2 Row 0	10:00:27 PM,14 Nov'2024
Critical	251	Volume	PCI Slot: 12 Volume is now DEGRADED Volume 237	10:00:17 PM,14 Nov'2024
Information	81	Volume	PCI Slot: 12 State change on Volume: 237 Previous: Optimal; Current: Degraded;	10:00:17 PM,14 Nov'2024
Information	114	Physical Device	PCI Slot: 12 Port 21,Box=1,Bay=2 (DeviceId: 22) - Drive: State change - Previous: Online, Current: Offline	10:00:17 PM,14 Nov'2024
Information	370	Volume	PCI Slot: 12 Volume is available. Volume: 237	10:00:04 PM,14 Nov'2024
Information	138	Volume Configuration	PCI Slot: 12 Created Volume: 237	10:00:04 PM,14 Nov'2024
Information	249	Volume	PCI Slot: 12 Volume is now OPTIMAL Volume 237	10:00:04 PM,14 Nov'2024
Information	114	Physical Device	PCI Slot: 12 Port 21,Box=1,Bay=3 (DeviceId: 23) - Drive: State change - Previous: UnConfigured Good; Current: Online	10:00:04 PM,14 Nov'2024
Information	114	Physical Device	PCI Slot: 12 Port 21,Box=1,Bay=2 (DeviceId: 22) - Drive: State change - Previous: UnConfigured Good; Current: Online	10:00:04 PM,14 Nov'2024
Information	114	Physical Device	PCI Slot: 12 Port 21,Box=1,Bay=3 (DeviceId: 23) - Drive: State change - Previous: Online, Current: UnConfigured Good	9:59:41 PM,14 Nov'2024

Actions

Download Events

Clear Events

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

Downloading Logs

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

Clearing the Event Logs

Perform the following steps to clear the event logs.

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

Known Issues and Workarounds

The following is a list of known issues and workarounds.

- **Issue:** The desktop shortcut is linked with a batch file. If the Windows User Account Control (UAC) is enabled, the privilege escalation prompt is not presented for users who have administrative privileges, but not administrator resulting in access denied message.
Workaround: Enter the MRSA URL, `localhost:port` directly into the browser
- **Issue:** The event API length for the serial number field is 10 characters. Only 10 characters are displayed in the event data for the PD serial number.
Workaround: None.
- **Issue:** When heavy I/Os are active, MRSA may be slow and the user may see stale data.
Workaround: Refresh the browser or restart the MRSA service to retrieve the updated data.
- **Issue:** When the user enables or disables a controller, the controller is not loaded or unloaded in MRSA by default.
Workaround: Restart the MRSA service to retrieve the updated controller list.
- **Issue:** A repurposed event occurs.
Workaround: Restart the MRSA service.
- **Issue:** Only the GUI is branded per HPE terminologies. Downloaded files are not branded.
Workaround: None.
- **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:** MRSA does not detect all the controllers in a HyperV Environment when the controller passthrough is enabled or disabled.
Workaround: Restart the MRSA service to reload and update the library.
- **Issue:** MRSA may be inaccessible after a successful firmware update while I/O's are occurring.
Workaround: Restart the MRSA services.
- **Issue:** MRSA may hang when downloading support logs on multiple clients.
Workaround: Restart the MRSA 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 an auto rebuild is enabled, multiclick PD actions are not updated properly.
Workaround: Manually refresh the page.
- **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:** Operations performed during an online controller reset fail.

Workaround: Do not perform any operation in MRSA during an online controller reset.

- **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 causes 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 causes the application to display different action menu names. MRSA displays it as **Make unconfigured good**.
- **Issue:** If the same dedicated spare is assigned to multiple arrays, you may see inconsistency in the Element Count and DHSP Element selection check boxes on the Controller page.
- **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 disappear.

- **Issue:** During installation or uninstallation, the publisher can show as unknown on the **User Account Control** message box.

- **Issue:** MRSA does not allow the physical drive to be selected from non-spanned volumes or spanned volumes.

- **Issue:** The **Modify** option for the existing `setup.exe` does not work.

Workaround: Uninstall and reinstall the build instead of using the **Modify** option.

Support and Other Resources

Accessing Hewlett Packard Enterprise Support

For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:

<http://www.hpe.com/assistance>

To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:

<http://www.hpe.com/support/hpesc>

Information to collect:

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

Accessing Updates

Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.

To download product updates:

Hewlett Packard Enterprise Support Center

www.hpe.com/support/hpesc

Hewlett Packard Enterprise Support Center: Software downloads

www.hpe.com/support/downloads

My HPE Software Center

www.hpe.com/support/softwaredepot

To subscribe to eNewsletters and alerts:

www.hpe.com/support/e-updates

To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center More Information on Access to Support Materials page:

www.hpe.com/support/AccessToSupportMaterials

NOTE

Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HPE Passport set up with relevant entitlements.

Customer Self Repair

Hewlett Packard Enterprise customer self repair (CSR) programs allow you to repair your product. If a CSR part needs to be replaced, it will be shipped directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your Hewlett Packard Enterprise authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider or go to the CSR website: www.hpe.com/support/selfrepair

Remote Support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which will initiate a fast and accurate resolution based on your product's service level. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

If your product includes additional remote support details, use search to locate that information.

Remote support and Proactive Care information

- **HPE Get Connected**
www.hpe.com/services/getconnected
- **HPE Proactive Care Services**
www.hpe.com/services/proactivecare
- **HPE Proactive Care Service: Supported Products List**
www.hpe.com/services/proactivecaresupportedproducts
HPE Proactive Care Advanced Service: Supported Products List
www.hpe.com/services/proactivecareadvancedsupportedproducts

Proactive Care Customer Information

- **Proactive Care Central**
- **Proactive Care Service Activation**
www.hpe.com/services/proactivecarecentralgetstarted

Warranty Information

To view the warranty for your product or to view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products* reference document, go to the Enterprise Safety and Compliance website:

www.hpe.com/support/Safety-Compliance-EnterpriseProducts

Additional warranty information

- **HPE ProLiant and x86 Servers and Options**
www.hpe.com/support/ProLiantServers-Warranties
- **HPE Enterprise Servers**
www.hpe.com/support/EnterpriseServers-Warranties
- **HPE Storage Products**
www.hpe.com/support/Storage-Warranties
HPE Networking Products
www.hpe.com/support/Networking-Warranties

Regulatory Information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at the Hewlett Packard Enterprise Support Center:

www.hpe.com/support/Safety-Compliance-EnterpriseProducts

Additional regulatory information

Hewlett Packard Enterprise is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements such as REACH (Regulation EC No 1907/2006 of the European Parliament and the Council). A chemical information report for this product can be found at:

www.hpe.com/info/reach

For Hewlett Packard Enterprise product environmental and safety information and compliance data, including RoHS and REACH, see:

www.hpe.com/info/ecodata

For Hewlett Packard Enterprise environmental information, including company programs, product recycling, and energy efficiency, see:

www.hpe.com/info/environment

Documentation Feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (docsfeedback@hpe.com). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.

Glossary

This glossary defines the terms used in this document.

access policy	A volume property indicating what kind of access is allowed for a particular volume. The possible values are <i>Read/Write</i> , <i>Read Only</i> , or <i>Blocked</i> .
array	A group of drives attached to a RAID controller on which one or more volumes can be created. All volumes in the array use all of the drives in the array.
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 reads to see if the required data is already stored in cache memory. If it is already in cache memory (a cache hit), it is read from cache memory immediately and the main memory read is aborted (or not started). If the data is not cached (a cache miss), it is fetched from main memory and saved in cache memory.
caching	The process of using a high speed memory buffer to speed up a computer system's overall read/write performance. The cache can be accessed at a higher speed than a drive subsystem. To improve read performance, the cache usually contains the most recently accessed data, as well as data from adjacent drive sectors. To improve write performance, the cache can temporarily store data in accordance with its write back policies.
capacity	A property that indicates the amount of storage space on a drive or volume.
coerced capacity	A drive property indicating the capacity to which a drive has been coerced (forced) to make it compatible with other drives that are nominally the same capacity. For example, a 4-GB drive from one manufacturer might be 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 array 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 volume with a redundant RAID level are consistent and that automatically fixes any errors. For RAID 1 arrays, 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. HPE Smart Array P824i-p MR Gen10 Controllers perform RAID functions such as striping and mirroring to provide data protection.
copyback	The procedure used to copy data from a source drive of a volume to a destination drive that is not a part of the volume. The copyback operation is often used to create or restore a specific physical configuration for a array (for example, a specific arrangement of array members on the device I/O buses). The copyback operation can be run automatically or manually. Typically, a drive fails or is expected to fail, and the data is rebuilt on a spare drive. The failed drive is replaced with a new drive. Then the data is copied from the spare drive to the new drive, and the spare drive reverts from a rebuild drive to its original spare drive status. The copyback operation runs as a background activity, and the volume is still available online to the host.

current write policy	A volume property that indicates whether the volume currently supports Write Back mode or Write Through mode. • In Write Back mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. • In Write Through mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction.
device ID	A controller or drive property indicating the manufacturer-assigned device ID.
DDF	Data disk format.
drive type	A drive property indicating the characteristics of the drive.
fast initialization	A mode of initialization that quickly writes zeros to the first and last sectors of the volume. This allows you to immediately start writing data to the volume while the initialization is running in the background.
fault tolerance	The capability of the drive subsystem to undergo a single drive failure per array without compromising data integrity and processing capability. HPE Smart Array MR Controllers provide fault tolerance through redundant arrays in RAID levels 1, 5, 6, 10, 50, and 60. They also support spare drive 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 a drive or from a network, then passes control to the operating system.
formatting	The process of writing a specific value to all data fields on a drive, to map out unreadable or bad sectors. Because most drives are formatted when manufactured, formatting is usually done only if a drive generates many media errors.
FS	File system.
GUI	Graphical user interface.
HPE Smart Storage Energy Pack	Refers to a energy pack backup unit.
JBOD	Just a bunch of disks. JBOD generally refers to a collection of hard disks that are directly managed by the host. JBOD is an alternative to using a RAID configuration. Rather than configuring a storage array to use a RAID level, the disks within the array are treated as independent disks.
initialization	The process of writing zeros to the data fields of a volume and, in fault-tolerant RAID levels, generating the corresponding parity to put the volume in a Ready state. Initialization erases all previous data on the drives. Arrays will work without initializing, but they can fail a consistency check because the parity fields have not been generated.
IO policy	A volume property indicating whether Cached I/O or Direct I/O is being used. In Cached I/O mode, all reads are buffered in cache memory. In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to cache and the host concurrently. If the same data block is read again, it comes from cache memory. (The IO Policy applies to reads on a specific volume. It does not affect the read ahead cache.)
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.
LDF	Logical disk format.
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 HPE Smart Array P824i-p MR Gen10 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.

offline	A drive is offline when it is part of a volume but its data is not accessible to the volume.
OS	Operating system.
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.
physical drive or disk (PD)	A disk used to emphasize a contract with virtual disks.
RAID	A group of multiple, independent drives that provide high performance by increasing the number of drives used for saving and accessing data. A RAID array 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 volumes. Data throughput improves because several drives can be accessed simultaneously. RAID configurations also improve data storage availability and fault tolerance. Redundant RAID levels (RAID levels 1, 5, 6, 10, 50, and 60) provide data protection.
RAID 0	Uses data striping on two or more drives to provide high data throughput, especially for large files in an environment that requires no data redundancy.
RAID 1	Uses data mirroring on pairs of drives so that data written to one drive is simultaneously written to the other drive. RAID 1 works well for small databases or other small applications that require complete data redundancy.
RAID 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 arrays. 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 arrays 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 arrays 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 array (parity group).
RAID level	A volume property indicating the RAID level of the volume. HPE Smart Array MR Controllers support RAID levels 0, 1, 5, 6, 10, 50, and 60.
RAID transformation	A feature in RAID subsystems that allows changing a RAID level to another level without powering down the system.
raw capacity	A drive property indicating the actual full capacity of the drive before any coercion mode is applied to reduce the capacity.
read policy	A controller attribute indicating the current Read Policy mode. In Always Read Ahead mode, the controller reads sequentially ahead of requested data and stores the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data. In No Read Ahead mode (known as Normal mode in WebBIOS), read ahead capability is disabled.
rebuild	The regeneration of all data to a replacement drive in a redundant volume after a drive failure. A drive rebuild normally occurs without interrupting normal operations on the affected volume, 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 volume	A method of undoing the configuration of a new volume. If you highlight the volume in the Configuration wizard and click Reclaim , the individual drives are removed from the volume configuration.

redundancy	A property of a storage configuration that prevents data from being lost when one drive fails in the configuration.
redundant configuration	A volume that has redundant data on drives in the array that can be used to rebuild a failed drive. The redundant data can be parity data striped across multiple drives in a array, or it can be a complete mirrored copy of the data stored on a second drive. A redundant configuration protects the data in case a drive fails in the configuration.
SAS	Acronym for Serial-Attached SCSI. SAS is a serial, point-to-point, enterprise-level device interface that leverages the Small Computer System Interface (SCSI) protocol set. The SAS interface provides improved performance, simplified cabling, smaller connectors, lower pin count, and lower power requirements when compared to parallel SCSI.
SATA	Acronym for Serial Advanced Technology Attachment. A physical storage interface standard. SATA is a serial link that provides point-to-point connections between devices. The thinner serial cables allow for better airflow within the system and permit smaller chassis designs.
SCSI device type	A drive property indicating the type of the device, such as drive.
serial no.	A controller property indicating the manufacturer-assigned serial number.
spare drive	A standby drive that can automatically replace a failed drive in a volume and prevent data from being lost. A spare drive can be dedicated to a single redundant array or it can be part of the global spare drive pool for all arrays controlled by the controller. When a drive fails, the application automatically uses a spare drive to replace it and then rebuilds the data from the failed drive to the spare drive. Spare Drives can be used in RAID 1, 5, 6, 10, 50, and 60 storage configurations.
stripe size	A volume property indicating the length of the interleaved data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB, and the strip size is 16 KB. The user can select the stripe size.
striping	A technique used to write data across all drives in a volume. Each stripe consists of consecutive volume data addresses that are mapped in fixed-size units to each drive in the volume using a sequential pattern. For example, if the volume includes five drives, the stripe writes data to drives one through five without repeating any of the drives. The amount of space consumed by a stripe is the same on each drive. Striping by itself does not provide data redundancy. Striping in combination with parity does provide data redundancy.
strip size	The portion of a stripe that resides on a single drive in the array.
subvendor ID	A controller property that lists additional vendor ID information about the controller.
transformation	The process of moving volumes and spare drive 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 volume information on the drives.
transformation rate	The user-defined rate at which an array modification operation is carried out.
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.
volume	An entity within a SCSI target that executes I/O commands. A storage unit created by a RAID controller from one or more drives. Although a volume can be created from several drives, it is seen by the operating system as a single drive. Depending on the RAID level used, the volume can retain redundant data in case of a drive failure.
volume state	A volume property indicating the condition of the volume. Examples include Optimal and Degraded.
write-back	In Write-Back Caching mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a drive write transaction. Data is written to the drive subsystem in accordance with policies set up by the controller. These policies include the amount of dirty/clean cache lines, the number of cache lines available, and elapsed time from the last cache flush.

write policy	See <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 1.12, November 2025

Updated UNMAP Capability Feature, Set Adjustable Task Rate, Personality Management, and Removing a Drive sections.

Added Configuring Different Types of Access, Setting the Task Information, Managing Power-Save Settings, and Auto Assign Policy.

Version 1.11, September 2025

Updated the terminology for drive group. Updated the Support Matrix, Configuring Different Types of Access, Creating a New Storage Configuration Using the Simple Configuration Option, Managing SAS Storage Link Speed, Expanding the Online Capacity of a Volume for MR200/MR400 Controllers, Updating the Controller Firmware, Marking a Drive as a Missing Drive, Making a JBOD Drive, and SafeStore Encryption Services sections.

Added Function Profile Management, Manage SID Ownership, Firmware Activation Status, Device Reporting Order, Managing Factory Defaults, Expanding the Online Capacity of a Volume, Expanding the Online Capacity of a Volume for MR932 Controllers, Locating Tape Drives, Viewing Protected Arrays, NVMe Drive Format, Physical Function Information, and Operations on Physical Functions.

Removed Configuring Different Types of Access and Multi-Selection Threshold Physical Drives sections.

Version 1.10, May 2025

Updated the terminology for virtual drive. Updated Changing Behavior Modes and Make Unconfigured Good Drives and Make JBOD Drives section.

Added Clearing NVRAM.

Version 1.9, February 2025

Updated the terminology for logical drive, hot spare, and OCE. Updated the UNMAP Capability Feature, Managing Arrays, and Make Unconfigured Good Drives and Make JBOD Drives sections.

Added Spin Down, Replacing a Missing Drive, and NVMe Thermal Poll Interval.

Removed CacheCade content.

Version 1.8, January 2024

Updated the Abstract and Support Matrix.

Added Upgrade Requirements, Standalone Installer, Installing the LSI Storage Authority Software on the Microsoft Windows Operating System, Installing the LSI Storage Authority Software on the Linux Operating System, Changing the LSI Storage Authority Application Port Number, Changing the Nginx Web Server Port Numbers, and Changing the Nginx Read Timeout.

Version 1.7, September 2023

Updated the Making a Drive Online and Known Issues and Workarounds sections.

Added Changing Behavior Modes.

Version 1.6, May 2023

Updated the Using the MegaRAID CacheCade Pro 2.0 Feature and Known Issues and Workarounds sections.

Added Configuring Different Types of Access.

Version 1.5, October 2022

Updated the Support Matrix, Erasing a Drive Securely, Sanitizing a Drive, Known Issues and Workarounds, and Glossary sections.

Version 1.4, April 2022

Added Known Issues and Workarounds, Marking a Drive as a Missing Drive, Removing a Drive, Erasing a Drive Securely, MegaRAID SafeStore Encryption Services, UNMAP Capability Feature, and Multi-Selection Threshold for Virtual and Physical Drives sections.

Updated Clearing the Configuration, Deleting a Logical Drive, Performing Initial Configuration, Glossary, and Support Matrix sections.

Version 1.3, February 2021

Added new Known Issues and Workarounds, Marking a Drive as a Missing Drive, Removing a Drive, Erasing a Drive Securely, MegaRAID SafeStore Encryption Services, UNMAP Capability Feature, and Multi-Selection Threshold for Virtual and Physical Drives sections.

Updated Clearing the Configuration, Deleting a Logical Drive, Performing Initial Configuration, Glossary, and Support Matrix sections.

Preliminary, Version 1.1, January 2020

Updated the Overview, Server Dashboard, and Configuration sections.

Preliminary, Version 1.0, December 25, 2017

Initial document release.