



HPE MR932i-p PCIe SPDM PI Storage Controller User Guide

Part Number: 30-47C5C88C-003
Published: April 2026
Edition: 3

HPE MR932i-p PCIe SPDM PI Storage Controller User Guide

Abstract

This document includes feature, installation, and configuration information about HPE MR controller and is for the person who installs, administers, and troubleshoots servers and storage systems. Hewlett Packard Enterprise assumes you are qualified in the servicing of computer equipment and trained in recognizing hazards in products with hazardous energy levels.

Part Number: 30-47C5C88C-003

Published: April 2026

Edition: 3

© Copyright 2026 Hewlett Packard Enterprise Development LP

Notices

The information provided here is subject to change without notice. Hewlett Packard Enterprise's products and services are covered only by the express warranty statements that come with them. This document does not constitute an additional warranty. Hewlett Packard Enterprise is not responsible for any technical or editorial errors or omissions in this document.

Confidential computer software. You must have a valid license from Hewlett Packard Enterprise to possess, use, or copy the software. In accordance 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 the vendor's standard commercial license.

Links to third-party websites will take you outside of the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for the information outside the Hewlett Packard Enterprise website.

Acknowledgments

Intel®, Itanium®, Optane™, Pentium®, Xeon®, Intel Inside®, Intel® VMD, Intel® Virtual RAID on CPU (Intel® VROC), and the Intel Inside logo are trademarks of Intel Corporation in the U.S. and other country/regions.

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

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

VMware®, VMware NSX®, VMware vCenter®, and VMware vSphere® are registered trademarks or trademarks of VMware, Inc. and its subsidiaries in the United States and other jurisdictions.

MegaRAID™ and CacheVault™ are the registered trademark of Broadcom, Inc.

All third-party marks are the property of their respective owners.

Table of contents

HPE MR932i-p PCIe SPDM PI Storage Controller.....	9
HPE MR932i-p PCIe SPDM PI Storage Controller naming framework.....	10
Comparison of HPE MR932i-p PCIe SPDM PI Storage Controller and HPE MR416i-p Gen11 SPDM Storage Controller.....	11
Component identification for HPE MR932i-p PCIe SPDM PI Storage Controller.....	14
Supported features and management tools.....	17
Management and configuration tools.....	17
Operating environments.....	18
RAID technologies.....	18
Transformation features.....	19
Drive technologies.....	20
Security features.....	20
Reliability features.....	21
Performance features.....	21
RAID concepts and terminologies.....	22
RAID levels.....	23
RAID 0 - Striping.....	23
RAID 1 and RAID 1+0 (RAID 10) - Mirroring.....	24
RAID 5 - Parity.....	26
RAID 50 - Parity group.....	28
RAID 6 - Double parity.....	29
RAID 60 - Double parity group.....	30
RAID levels comparison.....	31
RAID selection guide.....	32
Selecting RAID for fault tolerance.....	32
Selecting RAID for write performance.....	33
Selecting RAID for usable capacity.....	35
Selecting RAID for the storage solution.....	36
Striping.....	37
Mirroring.....	37
Parity.....	38
Parity groups.....	38
Spare drives.....	39
Dedicated spare.....	39

Global spare.....	40
Initialization operations.....	40
Fast initialization.....	41
Background initialization.....	41
Full initialization.....	41
No initialization.....	42
Write operations.....	42
Regenerative writes.....	42
Backed-out writes.....	43
Full-stripe writes.....	43
Mixed mode - RAID and JBOD simultaneously.....	43
Read load balancing.....	45
Drive rebuild.....	45
Foreign configuration import.....	46
Array transformations.....	46
Expand array.....	46
Volume transformations.....	47
Transportable controller.....	47
Expand volume.....	48
Drive health and maintenance.....	48
Hot-plug drive LED definition.....	48
EDSFF Carrier LED definition.....	50
Consistency check.....	51
Replace drive.....	52
Dynamic sector repair.....	52
Online drive firmware update.....	53
Predictive drive failure.....	53
Patrol read.....	54
Security.....	54
Drive erase.....	55
Sanitize.....	55
Self-encrypting drive.....	56
Cache offload encryption.....	57
Hardware secure boot.....	57
Security Protocol and Data Model.....	58
Reliability.....	58
Recovery ROM.....	58
Cache Error Checking and Correction.....	59

Thermal monitoring.....	59
Performance.....	59
Manage SAS storage link speed.....	60
Manage PCIe storage interface.....	60
I/O performance mode.....	60
Cache.....	61
Read cache policy	61
Write cache policy.....	61
Flash-backed write cache.....	62
I/O policy.....	62
Drive caching.....	63
Strip size selection.....	63
Multiple physical functions.....	63
Installation.....	64
Installing in a configured server.....	65
Installing in an unconfigured server.....	66
Installing the OS with the controller driver.....	67
Installing an MR controller on the server.....	68
Connecting storage devices.....	68
Connecting internal storage.....	68
Cable part numbers.....	69
Configuration and maintenance tools for HPE MR932i-p PCIe SPDM PI Storage Controller.....	69
Redfish.....	71
UEFI System Utilities.....	72
Controller boot process in UEFI.....	72
Launching the system utilities	73
Resolving controller boot blocking events in UEFI.....	74
HPE MR Storage Administrator.....	75
HPE StorCLI2.....	76
Configuration and management using UEFI System Utilities.....	76
UEFI System Utilities GUI.....	77
Navigating the UEFI dashboard view sections.....	81
Main Menu.....	81
Actions.....	81
Configuration management.....	82
Creating volume.....	82
Selecting drives for volume creation.....	83
Creating profile-based volume.....	84

Viewing array properties.....	85
Converting multiple drives to JBOD.....	85
Converting multiple drives to unconfigured.....	86
Clearing a configuration.....	87
Viewing and importing foreign configuration.....	87
Clearing Escrow key.....	88
Clearing foreign configuration.....	88
Controller management.....	88
Viewing controller properties.....	89
Advanced controller management.....	89
Setting factory defaults.....	90
Setting factory repurpose.....	90
Saving persistent events.....	91
Clearing persistent events.....	91
Enabling security.....	92
Disabling drive security.....	93
Changing security settings.....	93
Spinning up multiple drives.....	94
Managing Snapdump.....	94
Managing SAS storage link speed.....	95
Managing PCIe storage interface.....	95
Managing controller personality.....	96
Managing auto-assign policy.....	97
Managing function profile.....	98
Managing controller function.....	98
Managing advanced software options.....	99
Advanced controller properties.....	99
Configuring cache and memory settings.....	100
Configuring patrol read settings.....	100
Configuring spare settings.....	101
Configuring task rates.....	102
Volume management.....	103
Viewing and configuring volume properties.....	103
Deleting a volume.....	104
Expanding a volume.....	104
Initializing a volume.....	105
Locating a volume.....	105
Erasing a volume.....	106

Mapping a volume to a function.....	106
Securing a volume.....	107
Removing a volume from function.....	107
Locating a drive associated with a volume.....	108
Making a drive associated with a volume offline.....	109
Checking the consistency of a volume.....	110
Device management.....	110
Viewing logical enclosure properties.....	111
Viewing drive properties.....	111
Locating a physical drive associated with a logical enclosure.....	111
Initializing or clearing a drive.....	112
Making a drive offline.....	113
Replacing a drive.....	113
Converting a drive to JBOD.....	114
Converting a drive to unconfigured.....	115
Assigning a global spare.....	115
Unassigning a global spare.....	116
Preparing a drive for removal.....	116
Erasing a drive cryptographically.....	117
Sanitizing an unconfigured good drive.....	117
Reconfiguring a drive.....	118
Viewing logical unit information of a drive.....	119
Viewing logical unit information of a drive.....	119
Energy pack management.....	120
Viewing the energy pack.....	120
Maintenance.....	120
Updating software and firmware.....	121
Error reporting.....	121
Diagnostic tools.....	122
Storage reference.....	124
Memory and storage capacity conventions.....	124
RAID conventions.....	125
HPE Storage Controller documentation resources.....	125
HPE Storage Controller documentation quick links.....	127

HPE MR932i-p PCIe SPDM PI Storage Controller

Provides an overview of the controller architecture, capabilities, supported interfaces, and core hardware characteristics.

The HPE MR932i-p x32 Lanes PCIe Gen5 SPDM Plug-in Storage Controller operates in multi-mode. It supports simultaneous operation in both SAS and PCIe (NVMe) modes, enabling concurrent connectivity to different types of storage devices. The controller negotiates speeds and protocols to ensure compatibility and performance across the supported interfaces.

The following table is a quick reference of the controller hardware specifications, supported features, and capabilities:

Feature	Support
Ports	32 internal
I/O Processor	SAS5132W
Host Interface	x16 PCIe 5.0
 IMPORTANT Only x16 link width is supported.	
Storage Interfaces	• SAS: 22.5Gb/s, 12Gb/s, 6Gb/s, 3Gb/s • PCIe (NVMe): 32 GT/s, 16 GT/s, 8 GT/s, 5 GT/s, 2.5 GT/s
Drive Support	Only supports SAS and NVMe SSD drives. • NVMe: 32 SSDs • SAS: 32 drives without expander  IMPORTANT Does not support SATA drives (SSD or HDD) and SAS HDDs.
Form Factor	PCIe FH/HL
RAID Levels	0, 1, 5, 6, 10, 50, 60
Cache Memory	Integrated RAID Cache (IRC)
Storage Interface Connectors	Four MCIO (SFF-TA-1016) x8

Feature	Support
Cache Protection	Yes
Energy Pack	embedded

The HPE MR932i-p PCIe SPDM PI Storage Controller uses an `embedded Energy Pack` for write-cache protection. It requires no battery power cable and is not a user-replaceable part.



NOTE

The previous MR-generation controllers used an HPE Smart Storage Energy Pack for write-cache protection.

Subtopics

[HPE MR932i-p PCIe SPDM PI Storage Controller naming framework](#)

[Comparison of HPE MR932i-p PCIe SPDM PI Storage Controller and HPE MR416i-p Gen11 SPDM Storage Controller](#)


More information

[Supported features and management tools](#)

[Component identification for HPE MR932i-p PCIe SPDM PI Storage Controller](#)

HPE MR932i-p PCIe SPDM PI Storage Controller naming framework

The controller naming framework provides a structured methodology for assigning clear, consistent, and uniform patterns across different controllers and generations.



More information

[HPE MR932i-p PCIe SPDM PI Storage Controller](#)

[Component identification for HPE MR932i-p PCIe SPDM PI Storage Controller](#)

[Comparison of HPE MR932i-p PCIe SPDM PI Storage Controller and HPE MR416i-p Gen11 SPDM Storage Controller](#)

Comparison of HPE MR932i-p PCIe SPDM PI Storage Controller and HPE MR416i-p Gen11 SPDM Storage Controller

The following table compares and highlights the key differences between the HPE MR932i-p PCIe SPDM PI Storage Controller and HPE MR416i-p Gen11 SPDM Storage Controller. It provides a consolidated comparison of their hardware, RAID capabilities, security features, performance functions, and management utilities.

Features	HPE MR932i-p PCIe SPDM PI Storage Controller	HPE MR416i-p Gen11 SPDM Storage Controller
Controller Hardware		

Features	HPE MR932i-p PCIe SPDM PI Storage Controller	HPE MR416i-p Gen11 SPDM Storage Controller
Host Interface	x16 PCIe 5.0	X8 PCIe 4.0
	 IMPORTANT Only x16 link width is supported.	
Storage Interfaces	• PCIe (NVMe): 32 GT/s, 16 GT/s, 8 GT/s, 5 GT/s, 2.5 GT/s • SAS: 22.5Gb/s, 12Gb/s, 6Gb/s, 3Gb/s	• PCIe (NVMe): 16 GT/s, 8 GT/s, 5 GT/s, 2.5 GT/s • SAS: 12Gb/s, 6Gb/s, 3Gb/s • SATA: 6Gb/s, 3Gbs
Drive Support	Only supports SAS and NVMe SSD drives. • NVMe: 32 SSDs • SAS: 32 drives without expander	• NVMe: 16 SSDs • SAS, SATA: 16 HDD/SSDs without expander
	 IMPORTANT Does not support SATA drives (SSD or HDD) and SAS HDDs.	
Form Factor	FH/HL	HH/HL
Cache Memory	Integrated RAID Cache (IRC)	8G x 72 flash-backed write cache
Storage Interface Connectors	Four SFF-TA-1016 x8	Two SFF-8654 x8
Energy Pack	embedded	HPE Energy Pack Options
RAID Features		
Maximum logical drives	16 non-single-drive RAID 0, and 64 single-drive RAID 0	240
Maximum logical drives per array	8	64
Expand array	Yes for no-span RAID levels. Only one first virtual disk consumes all available raw space.	✓

Features	HPE MR932i-p PCIe SPDM PI Storage Controller	HPE MR416i-p Gen11 SPDM Storage Controller
Migrate RAID level	Not supported	✓
<u>Drive rebuild</u>	Rebuild rate: Priority method. Rebuild vs I/O. Two drives of RAID 6 rebuild at the same time.	Rebuild rate: Configurable zero to 100; default 30. Two drives of RAID 6 rebuild; one at a time.
<u>Foreign configuration import</u>	Supported only for the drive configured by HPE MR932i-p PCIe SPDM PI Storage Controller.	Supported only for the drive configured by HPE MR932i-p PCIe SPDM PI Storage Controller, MR200, and MR400 controllers.
Volume block size	4KiB $\frac{1}{2}$	512B, 4KiB
Autoconfiguration	Not applied to the drive that is already known to the controller.	Applied to all unconfigured drives.
Background Initialization	Enabled by default on all RAID levels, except RAID 0, with no minimum drive count requirement.	Enabled on RAID 5, RAID 6, RAID 50, and RAID 60 when the minimum drive count requirement is met.
Security Features		
<u>Cache offload encryption</u>	✓	Not supported
<u>Self-encrypting drive</u>	Cannot disable drive security when any MR-secured physical drives are present.	Can disable drive security even when MR-secured physical drives are present.
Reliability Features		
<u>Recovery ROM</u>	✓	Not supported
Performance Features		
<u>I/O performance mode</u>	Always enabled for all RAID levels	Not for all RAID levels
<u>Read cache policy</u>	Always set to No read ahead	Configurable: Read ahead and No read ahead
<u>Drive caching</u>	Set to Disable by default.	Set to Enable by default.
<u>Strip size selection</u>	64KiB	64KB, 128KB, 256KB, 512KB, and 1MB
<u>Multiple physical functions</u>	4	1
Utilities and Drivers		
CLI	HPE StorCLI2	HPE StorCLI
Drivers	Linux: MPI3MR	Linux: megaraid_sas

Features	HPE MR932i-p PCIe SPDM PI Storage Controller	HPE MR416i-p Gen11 SPDM Storage Controller
	• Windows: mpi3drv • VMware: bcm_mpi3	• Windows: Megasas35 • VMware: lsi_mr3

¹ The controller supports only 4096-byte block size. The drives you use may support 4096-byte, 512-byte native, or 512-byte emulated block sizes. However, the mixing of 512-byte native or emulated drives with 4096-byte drives in the same array is not allowed.

More information

[HPE MR932i-p PCIe SPDM PI Storage Controller](#)

Component identification for HPE MR932i-p PCIe SPDM PI Storage Controller

Provides visual and descriptive identification of controller components, connectors, and LEDs for hardware recognition.

This section provides a detailed overview of the physical components of the HPE MR932i-p PCIe SPDM PI Storage Controller. It includes a labeled diagram and descriptions to help you identify connectors, ports, and LEDs.




Table 1. Connectors

ID	Type	Description
1	Onboard serial UAR T connector	4-pin connector Reserved for Broadcom use.

ID	Type	Description
2	embedded Energy Pack	Supercapacitor Energy Pack that supplies backup power to preserve cached data if there is host power loss or server failure.
7, 8, 9, and 10	Storage interface connector	Four MCIO (SFF-TA-1016) internal connectors. Connect the adapter by cable to the storage devices. Port number mapping is as follows: • 7: Port 3 • 8: Port 4 • 9: Port 1 • 10: Port 2

Table 2. LEDs

ID	Color	Type	Description
3	Green	ONFI activity	Indicates when the ONFI is active for cache offload or recovery.
4	Yellow	Controller over temperature	Stays on solid to indicate that the SAS5132W device temperature sensor is over the temperature threshold. When the device is in the proper temperature range, this LED is off.
5	Yellow	Supercapacitor fault	Indicates that the CacheVault power module is in a fault state or is over the temperature threshold.
6	Green	System heartbeat	Indicates that the SAS5132W RoC ASIC is operating normally. This LED flashes at 1 Hz.

More information

[Installation](#)

[Hot-plug drive LED definition](#)

[EDSFF Carrier LED definition](#)

Supported features and management tools

Provides details about supported features, operating environments, and management utilities available for the controller.

This section outlines the key capabilities of the HPE MR932i-p PCIe SPDM PI Storage Controller and the tools available for configuring, monitoring, and maintaining storage environments.

It provides a quick reference to help match controller functionality with workload requirements.

Subtopics

Management and configuration tools

Operating environments

RAID technologies

Transformation features

Drive technologies

Security features

Reliability features

Performance features

More information

RAID concepts and terminologies

Management and configuration tools

You can use the following tools and utilities to manage and configure the controller.

- [Redfish](#)
- [UEFI System Utilities](#)
- [HPE MR Storage Administrator](#)



IMPORTANT

Manage HPE MR932i-p PCIe SPDM PI Storage Controller using the HPE MR Storage Administrator. Smart Storage Administrator does not manage HPE MR storage controllers.

- [HPE StorCLI2](#)



IMPORTANT

To manage the HPE MR932i-p PCIe SPDM PI Storage Controller, use HPE StorCLI2. For HPE MR Gen10 Plus and Gen11 controllers, use HPE StorCLI.

- Intelligent Provisioning



IMPORTANT

Intelligent Provisioning does not currently support HPE MR932i-p PCIe SPDM PI Storage Controller. Support will be added in a future release.

Operating environments

The following operating systems are supported.

Operating environments	HPE MR932i-p PCIe SPDM PI Storage Controller
Microsoft Windows	✓
Linux	✓
VMware vSphere	✓
UEFI Boot mode	✓

More information

[Installing the OS with the controller driver](#)

RAID technologies

The following RAID technologies are supported.

Features	HPE MR932i-p PCIe SPDM PI Storage Controller
RAID levels	0, 1, 5, 6, 10, 50, 60
Maximum volumes	16 non-single-drive RAID 0 64 single-drive RAID 0
Maximum volumes per array	8
Maximum physical drives without expander	32
Mixed mode - RAID and JBOD simultaneously	✓
Read load balancing	✓
Parity groups	✓
Initialization operations	✓
Regenerative writes	✓
Backed-out writes	✓
Full-stripe writes	✓
Dedicated spare	✓
Global spare	✓
Drive rebuild	✓
Foreign configuration import	✓

More information

[RAID levels comparison](#)

[RAID selection guide](#)

Transformation features

This topic identifies the transformation features that the controller supports and does not support.

Features	HPE MR932i-p PCIe SPDM PI Storage Controller
<u>Expand array</u>	✓
<u>Transportable controller</u>	✓
<u>Expand volume</u>	✓
Migrate RAID level	Not supported
Migrate Stripe	Not supported

Drive technologies

The following drive technology features are supported.

Features	HPE MR932i-p PCIe SPDM PI Storage Controller
<u>Hot-plug drive LED</u>	✓
<u>Consistency check</u>	✓
<u>Replace drive</u>	✓
<u>Dynamic sector repair</u>	✓
<u>Online drive firmware update</u>	✓
<u>Predictive drive failure</u>	✓
<u>Patrol read</u>	✓

Security features

The following security features are supported.

Features	HPE MR932i-p PCIe SPDM PI Storage Controller
<u>Drive erase</u>	✓
<u>Sanitize</u>	✓
<u>Self-encrypting drive</u>	✓
<u>Cache offload encryption</u>	✓
<u>Hardware secure boot</u>	✓
<u>Security Protocol and Data Model</u>	✓

Reliability features

The following reliability features are supported.

Features	HPE MR932i-p PCIe SPDM PI Storage Controller
<u>Recovery ROM</u>	✓
<u>Cache Error Checking and Correction</u>	✓
<u>Thermal monitoring</u>	✓

Performance features

This topic identifies the performance features that the controller supports and does not support.

Features	HPE MR932i-p PCIe SPDM PI Storage Controller
<u>Manage SAS storage link speed</u>	✓
<u>Manage PCIe storage interface</u>	✓
<u>I/O performance mode</u>	✓

Features	HPE MR932i-p PCIe SPDM PI Storage Controller
Read cache policy	Not supported
Write cache policy	✓
Flash-backed write cache	✓
I/O policy	✓
Drive caching	✓
Strip size selection	Not supported
Multiple physical functions	✓

RAID concepts and terminologies

Provides foundational RAID concepts, terminology, and selection guidance to support storage design decisions.

This section introduces the core concepts of RAID, explaining its fundamental principles, key terminologies, and RAID levels. It covers concepts such as striping, mirroring, and parity, along with their role in enhancing storage performance, capacity, and fault tolerance.

It also introduces advanced RAID configurations and explains how different RAID levels balance redundancy and efficiency.

Subtopics

[**RAID levels**](#)

[**RAID levels comparison**](#)

[**RAID selection guide**](#)

[**Striping**](#)

[**Mirroring**](#)

[**Parity**](#)

[**Parity groups**](#)

[**Spare drives**](#)

[**Initialization operations**](#)

[**Write operations**](#)

[**Mixed mode - RAID and JBOD simultaneously**](#)

[**Read load balancing**](#)

[**Drive rebuild**](#)

[**Foreign configuration import**](#)

[**Array transformations**](#)

[**Volume transformations**](#)

Drive health and maintenance

Security

Reliability

Performance

RAID levels

HPE MR932i-p PCIe SPDM PI Storage Controller supports different RAID levels. Each level combines techniques like striping, mirroring, and parity to balance performance, capacity, and fault tolerance.

Common levels are RAID 0, 1, 5, 6, 10, 50, and 60.

Subtopics

RAID 0 - Striping

RAID 1 and RAID 1+0 (RAID 10) - Mirroring

RAID 5 - Parity

RAID 50 - Parity group

RAID 6 - Double parity

RAID 60 - Double parity group

More information


RAID levels comparison

RAID selection guide

RAID 0 - Striping

RAID 0 configuration provides data striping. There is no protection against data loss when a drive fails. However, it is useful for rapid storage of large amounts of noncritical data. For example, printing or image editing, or when cost is the most important consideration.

- Requires a minimum of one drive.
- Supports up to 32 drives.



Benefits

- Prioritizes performance and low cost over data protection.
- Provides the highest write performance of all RAID methods.
- Offers the lowest cost per unit of stored data among RAID methods.
- Uses the entire drive capacity for data. No capacity is reserved for fault tolerance.

More information


Striping

RAID 1 and RAID 1+0 (RAID 10) - Mirroring

RAID 1 and RAID 1+0 (RAID 10) configurations use disk mirroring to provide data redundancy. All data is duplicated to a second drive, so usable capacity is $C \times (n / 2)$, where C is the drive capacity with n drives in the array.


RAID 1

- Applies when the array has exactly two physical drives.



RAID 10

- Requires a minimum of four drives.
- Supports up to 32 drives.
- Uses the 1E_RLQ algorithm. Does not use spans. Forms mirror pairs in the order that the drives appear in the array.
- Mirrors drives in pairs. If one drive in a pair fails, its partner provides the data.
- Tolerates multiple drive failures as long as no two failed drives belong to the same mirror pair. Drive counts increase by increments of two.



Benefits

- Prioritizes performance and data protection over usable capacity.
- Offers the highest write performance among fault-tolerant configurations.
- Prevents data loss when a drive fails as the failed drive is not mirrored to another failed drive.
- Tolerates failure of up to half the drives that is one per mirror pair without data loss.

More information

Mirroring


Read load balancing

RAID 5 - Parity

RAID 5 protects data with distributed parity, represented as $P_{x,y}$ in the figure. The controller computes parity summing the XOR data corresponding blocks from each drive in a stripe and evenly writes parity blocks across all drives within the volume.

If one drive fails, the controller rebuilds its data from the remaining data and parity. Usable capacity is $C \times (n - 1)$, where C is one drive's capacity and n is the drive count.

- Requires a minimum of three drives.
- Supports up to 32 drives.



Benefits

- Balances usable capacity, write performance, and data protection.
- Provides the highest usable capacity of fault-tolerant RAID levels.
- Tolerates a single drive failure without data loss.

More information

Parity

Parity groups


RAID 50 - Parity group

RAID 50 combines multiple identical RAID 5 parity groups and stripes data across them. The smallest RAID 50 configuration uses six drives organized into two parity groups of three drives each.

For a given number of drives, arranging drives into the largest possible number of parity groups minimizes the risk of data loss. For example, four parity groups of three drives are more secure than three groups of four drives. However, this arrangement reduces usable capacity.

RAID 50 loses all data if a second drive in the same parity group fails before the first failed drive finishes rebuilding. Compared with single RAID 5 sets, RAID 50 uses a larger portion of the capacity for parity.

- Requires a minimum of six drives.
- Supports up to 32 drives without expander.



Benefits

- Improves RAID 5 performance, especially for writes.
- Provides better fault tolerance than RAID 0 or RAID 5.
- Tolerates up to n physical drives failure without loss of data, as long as the failed drives are in different parity groups. Here, n is the number of parity groups.

More information

Parity groups


Parity

RAID 6 - Double parity

RAID 6 protects data with double parity. The controller stores two different parity sets, represented as $P_{x,y}$ and $Q_{x,y}$ in the figure, letting the array survive any two drive failures. Each parity set uses capacity equal to one drive, so usable capacity is $C \times (n - 2)$, where C is one drive's capacity and n is the drive count.

Use RAID 6 when preventing data loss is critical, but cost still matters. RAID 6 with Advanced Data Guarding (ADG) has a lower probability of data loss than RAID 5.

- Requires a minimum of four drives.
- Supports up to 32 drives.



Benefits

- Prioritizes data protection and usable capacity over write performance.
- Tolerates any two drive failures without data loss.

More information

Parity

Parity groups

RAID 60 - Double parity group

RAID 60 combines multiple identical RAID 6 parity groups and stripes data across them. The smallest possible RAID 60 configuration has eight drives organized into two parity groups of four drives each.

For a given number of drives, you can reduce the risk of data loss by using the largest possible number of parity groups. For example, configuring five parity groups of four drives offers better protection than four parity groups of five drives. However, increasing the number of parity groups reduces usable capacity because more drives are reserved for parity.

All data is lost if a third drive in a parity group fails before one of the other failed drives in the parity group has finished rebuilding. A greater percentage of array capacity is used to store redundant or parity data than with non-nested RAID methods.

Configuration guidelines

- Divide the drives evenly across the parity groups. The total number of drives must be exactly divisible by the number of parity groups.
- Determine the number of parity groups that you can configure based on the total number of drives.
- Calculate the maximum number of parity groups possible with `Total drives ÷ minimum drives per group`. Use four drives per group for RAID 60.
- Increase the number of parity groups to improve fault tolerance, but this method reduces usable capacity as more drives are used for parity.
- Requires a minimum of eight drives.
- Supports up to 32 drives without expander.

Benefits

- Provides higher performance than RAID 6, especially during writes.
- Offers better fault tolerance than RAID 0, 5, 50, or 6.

- Tolerates up to $2n$ drive failures, where n is the number of parity groups, as long as no more than two failed drives are in the same parity group.

More information

[Parity groups](#)

[Parity](#)

RAID levels comparison

This table helps you quickly evaluate how each RAID level balances speed, redundancy, and usable space, making it easier to choose the right configuration for your workload.

RAID level	Technique	Minimum drives	Fault tolerance	Storage capacity ¹
RAID 0	Striping	2	None	100%
RAID 1	Mirroring	Exactly 2	1 drive failure	C
RAID 5	Distributed parity	3	1 drive failure	$C \times (n - 1)$
RAID 50	Parity group: Striping data across RAID 5 volume sets.	6	1 drive failure per RAID 5 group	$C \times (n-1) \times S$
RAID 6	Double distributed parity	4	2 drive failures	$C \times (n - 2)$
RAID 60	Parity group: Striping data across RAID 6 volume sets.	8	2 drives failure per RAID 6 group	$C \times (n - 2) \times S$
RAID 10 (1+0)	Mirroring+Striping	4	1 drive failure per mirrored pair	$C \times (n / 2)$

¹ C : Drive Capacity, n : Drive number per Parity Group (Span), S : The number of Parity Groups (Span)

More information

[RAID selection guide](#)

[RAID 0 - Striping](#)

[RAID 1 and RAID 1+0 \(RAID 10\) - Mirroring](#)

[RAID 5 - Parity](#)

[RAID 50 - Parity group](#)

[RAID 6 - Double parity](#)

[RAID 60 - Double parity group](#)

[RAID levels](#)

RAID selection guide

Choosing the right RAID is essential for designing a reliable and efficient storage solution. The following three factors primarily drive RAID selection:

- **Fault tolerance:** Determines how well a RAID configuration can withstand disk failures and keep systems operational.
- **Write performance:** Defines data writing speed, crucial for workloads like databases and virtualization.
- **Usable capacity:** Reflects available storage after redundancy overhead such as mirroring or parity.

Subtopics

[Selecting RAID for fault tolerance](#)

[Selecting RAID for write performance](#)

[Selecting RAID for usable capacity](#)

[Selecting RAID for the storage solution](#)

More information

[RAID levels comparison](#)

[RAID 0 - Striping](#)

[RAID 1 and RAID 1+0 \(RAID 10\) - Mirroring](#)

[RAID 5 - Parity](#)

[RAID 50 - Parity group](#)

[RAID 6 - Double parity](#)

[RAID 60 - Double parity group](#)


[RAID levels](#)

Selecting RAID for fault tolerance

If your environment requires a high level of fault tolerance, select a RAID level optimized for fault tolerance.

The chart illustrates how RAID-level fault tolerance changes as storage array size increases. It includes RAID 0, 5, 50, 10, 6, and 60, and shows reliability on a scale from one to one billion and array sizes from zero to 96 drives. The chart assumes two parity groups for RAID 50 and RAID 60 and shows:

- RAID 10 is 30,000 times more reliable than RAID 0.
- RAID 5, RAID 50, RAID 6, and RAID 60 become less reliable as fault tolerance decreases when array size increases.



More information

[RAID levels comparison](#)


[RAID selection guide](#)

[RAID levels](#)

Selecting RAID for write performance

If your environment requires high write performance, select a RAID level optimized for write performance.

The following chart compares the write performance of RAID 10, 5, 50, 6, and 60 relative to RAID 0. The data in the chart assumes that drive performance is the limiting factor and that drive-write speeds match drive-read speeds.



Considerations for write performance

- RAID 5, 50, 6, and 60 performance assumes that parity initialization is complete.
- Write performance decreases as fault tolerance improves due to extra I/O.
- Read performance remains consistent across all RAID levels, except for small RAID 5 and 6 arrays.

Disk I/O per host write by RAID level

The following table shows the disk I/O required for each host write operation:



NOTE

Supported RAID levels can vary based on the controller model.

RAID level	Disk I/O for each host write
RAID 0	1
RAID 10	2
RAID 5	4
RAID 6	6

More information


[RAID levels comparison](#)

[RAID selection guide](#)

[RAID levels](#)

Selecting RAID for usable capacity

If your environment requires a high usable capacity, select a RAID level that is optimized for usable capacity. The chart shows the relationship between the number of drives in the array and the usable capacity over the capacity for RAID 0. The chart assumes two parity groups for RAID 50 and RAID 60.



Considerations for usable capacity

- Usable capacity decreases as fault tolerance improves due to an increase in parity data.
- Usable capacity for RAID 10 remains flat with larger arrays.
- Usable capacity for RAID 5, 50, 6, and 60 increases with larger arrays.

Minimum drive requirements by RAID level

The following table shows the minimum drive requirements for each RAID level:



NOTE

Supported RAID levels can vary based on the controller model.

RAID level	Minimum number of drives
RAID 0	1
RAID 10	4
RAID 5	3
RAID 6	4
RAID 50	6
RAID 60	8

More information

[RAID levels comparison](#)


[RAID selection guide](#)

[RAID levels](#)

Selecting RAID for the storage solution

The chart shows how each RAID level aligns with the requirements and how its relevance changes based on the needs of your environment. Depending on your requirements, you can optimize RAID selection as follows:

- **RAID 6 and 60:** Optimize for fault tolerance and usable capacity.
- **RAID 1 and 10:** Optimize for write performance.
- **RAID 5 and 50:** Optimize for usable capacity.



More information

[RAID levels comparison](#)

[RAID selection guide](#)

[RAID levels](#)

Striping

Striping is a data-distribution technique used to split data into fixed-size units called strips and write the data across all drives in a volume. Each stripe consists of consecutive volume data addresses, which are mapped in fixed-size units to each drive in the array following a sequential pattern. For example, if a volume includes five drives, a stripe distributes data blocks to drives one through five in order, without repeating any drive within that stripe. The amount of space allocated by each stripe is equal on every drive.

Striping improves performance by enabling parallel read and write operations across multiple drives. However, striping by itself does not provide data redundancy. When striping is combined with parity, it enables both data redundancy and fault tolerance. This combination is implemented in RAID levels such as RAID 5 and 6.

More information

[RAID 0 - Striping](#)

Mirroring

Mirroring is a data protection technique that creates an exact copy of all the data on one drive to another drive in the array. This method ensures high redundancy and fault tolerance, as the system can continue operating if one drive fails.

Mirroring provides fast read performance because data can be read from either copy. However, it requires double the storage capacity and does not improve write speed. It is commonly used in RAID 1 and combined with striping in RAID 10 for both performance and reliability.

More information

[RAID 1 and RAID 1+0 \(RAID 10\) - Mirroring](#)

[Read load balancing](#)

Parity

Parity is a technique that generates redundancy data from two or more parent data sets. This redundancy data allows the system to reconstruct one of the parent data sets if a drive failure occurs. Unlike mirroring, parity does not fully duplicate the parent data sets. Instead, it uses mathematical algorithms to create a smaller set of recovery information.

While parity provides fault tolerance, generating and writing parity data can slow down the write process compared to non-redundant methods. In RAID arrays, parity is typically applied to entire drives or to stripes distributed across all drives in the array.

More information

[RAID 5 - Parity](#)

[RAID 6 - Double parity](#)

[RAID 50 - Parity group](#)

[RAID 60 - Double parity group](#)

Parity groups

When you create a RAID 50 or 60 configuration, you must also set the number of parity groups, commonly referred to as spans.

The number of parity groups, also known as spans, must satisfy the following requirements:

- The number of parity groups must be greater than one.
- The total number of physical drives must be evenly divisible by the number of parity groups.
- Each parity group must contain at least the minimum number of drives required for the RAID level.
 - For RAID 50, the minimum is three drives per group.
 - For RAID 60, the minimum is four drives per group.

These requirements also determine the maximum number of parity groups, calculated as $\frac{\text{Total number of drives}}{\text{minimum number of drives necessary for that RAID level}}$.

Benefits

- Supports RAID 50 and 60.
- Increases fault tolerance due to higher number of parity groups.

More information

[RAID 50 - Parity group](#)

[RAID 5 - Parity](#)

[RAID 60 - Double parity group](#)

[RAID technologies](#)

Spare drives

Spare drives are unused disks in a RAID array that remain on standby to replace a failed drive automatically or manually. When a drive fails, a hot spare can immediately start rebuilding the data from parity or mirrored copies without administrator intervention, reducing downtime and risk of data loss.

Spares can be global that is available to any array or dedicated that is assigned to a specific array.

Subtopics

[Dedicated spare](#)

[Global spare](#)

More information

[Configuring spare settings](#)

Dedicated spare

A dedicated spare is a spare drive that is dedicated to one array.

It supports any fault tolerant volume such as RAID 1, 10, 5, 6, 50, 60.

The dedicated spare drive activates anytime a drive within the array fails.

More information

[RAID technologies](#)

[Configuring spare settings](#)

Global spare

A global spare drive replaces a failed drive in any array, as long as:

- The drive type is the same.
- The capacity of the global spare drives is equal to or larger than the capacity of the failed drive.

A global spare drive activates anytime a drive fails within a fault tolerant volume. For RAID 0 volumes, the global spare is active when a member drive reports a predictive failure.

More information

[RAID technologies](#)

[Configuring spare settings](#)

[Assigning a global spare](#)

[Unassigning a global spare](#)

Initialization operations

Initialize the logical drive after you create it. Initialization prepares the drive for use and ensures that the storage space is ready for data.



CAUTION

All data on the volume is erased during initialization. Before starting this operation, back up any data you want to keep.

Subtopics

[Fast initialization](#)

[Background initialization](#)

[Full initialization](#)

[No initialization](#)

More information

[RAID technologies](#)

[Configuring task rates](#)

[Initializing a volume](#)

Fast initialization

During fast initialization, the firmware quickly overwrites the first and last 8 MB regions of the new logical drive and clears any boot records or partition information. It then completes the initialization in the background. Monitor the progress of the initialization process using the progress indicator.

RAID levels that use parity (RAID 5, RAID 6, RAID 50, and RAID 60) require that the parity blocks be initialized to valid values. Valid parity data is required to enable enhanced data protection through background controller surface scan analysis and higher write performance (backed out write).

After parity initialization is complete, the controller does not read the entire stripe (regenerative write) to update the parity data. Due to this process, writes to a RAID 5 or RAID 6 volume are typically faster. This feature initializes parity blocks in the background while the volume is available for access by the OS.

Parity initialization takes several hours to complete. The time it takes depends on the size of the volume and the load on the controller. While the controller initializes the parity data in the background, the volume has full fault tolerance. You can start writing data to the volume immediately using this method.

More information

Initialization operations

Background initialization

Background initialization checks the drives for media errors when you create a volume. It starts automatically five minutes after volume creation and ensures that striped data segments are synchronized across all drives in the array.

Background initialization is enabled by default on all RAID levels except RAID 0, and it does not require a minimum drive count.

More information

Configuring task rates

Full initialization

During full initialization, a complete initialization is done on the new configuration. You cannot write data to the new volume until the initialization is complete. This process can take a long time if the drives are large. This initialization overwrites all blocks and destroys all the data on the volume.

Monitor the progress of the initialization process using the progress indicator.

More information

No initialization

If you select this option, the new configuration is not initialized, and the existing data on the drives is not overwritten. You can initialize the volume at a later time.

Write operations

Write operations is the process of storing data on a volume. In RAID-based systems, write operations involve writing user data along with parity or mirrored information to maintain redundancy and fault tolerance.

Subtopics

Regenerative writes

Backed-out writes

Full-stripe writes

Regenerative writes

Volumes can be created with background parity initialization so that they are available almost instantly. During this temporary parity initialization process, writes to the volume are performed using regenerative writes or full stripe writes. Anytime a member drive within an array fails, all writes that map to the failed drive are regenerative. A regenerative write is slower because it must read from nearly all the drives in the array to calculate new parity data. The write penalty for a regenerative write is $n + 1$ drive operations, where n is the total number of drives in the array.

As the array size grows, the write penalty increases, and write performance slows.

This method has the following benefits:

- Allows the volume to remain accessible before parity initialization completes.
- Allows the volume to be accessible when the array is degraded.

More information

RAID technologies

Backed-out writes

After parity initialization completes, random writes to a RAID 5, 50, 6, or 60 volume can use a faster backed-out write operation. A backed-out write uses the existing parity to calculate the new parity data.

As a result, the write penalty for:

- RAID 5 and 50 are always four drive operations.
- RAID 6 and 60 are always six drive operations.

The number of drives in the array does not influence the write penalty. This method improves random-write performance for RAID 5, 50, 6, and 60 volumes.

Backed-out writes are also known as — read-modify-write.

More information

[RAID technologies](#)

Full-stripe writes

A full-stripe write occurs when writes to the volume are sequential or when multiple random writes in the flash-backed write cache align sequentially. In a full-stripe write, the controller calculates new parity using only the new data being written to the drives. Because it does not read old data to generate parity, the write penalty is minimal. As the array size increases, the write penalty decreases by the ratio p/n , where p is the number of parity drives and n is the total number of drives.

This method delivers faster sequential-write performance for RAID 5, 6, and 60 volumes.

More information

[RAID technologies](#)

Mixed mode - RAID and JBOD simultaneously

Mixed mode lets you configure any drive as RAID or Just a Bunch of Drives (JBOD). This capability reduces the number of required controllers and maximizes drive-bay use within a backplane. For example, you can present all drives as JBOD while reserving a two-drive mirror for boot, using a single controller on a single backplane.

In the controller, a drive can be configured as:

- **A member of a volume:** The drive is created as a volume. It contains valid Disk Data Format (DDF) metadata.

- **Unconfigured Good:** The drive is unconfigured and hidden from the OS.
- **Spare:** The drive is on standby to provide extra protection for the array. It can automatically replace the failed drive in an array.
- **JBOD:** The drive is exposed to the host OS as a physical drive. User cannot use a JBOD drive to create a RAID configuration, because it is exposed to the host OS. It has to be converted to unconfigured good drives before creating a RAID configuration on it.

Autoconfiguration

The autoconfiguration feature simplifies the configuration workflow. When you insert a new drive, the controller automatically configures it based on the autoconfiguration option you select. This feature supports the following configuration options.

- **UGOOD:** Configures the newly inserted drive as an unconfigured good drive. You can use this drive later when creating a configuration.
- **JBOD:** Configures the newly inserted drive as a JBOD drive.
- **SecureJBOD:** Configures the newly inserted drive as secured JBOD if it is SED-capable. Otherwise, configures it as JBOD.
- **R0:** Configures the newly inserted drive as a single-drive RAID 0 write-through volume.
- **SecureR0:** Configures the newly inserted drive as a single-drive RAID 0 write-through volume. If the drive is SED capable, enables security on the volume.
- **ROWB:** Configures the newly inserted drive as a single-drive RAID 0, write-back volume.
- **SecureROWB:** Configures the newly inserted drive as a single-drive RAID 0 write-back volume. If the drive is SED capable, enables security on the volume.

The autoconfiguration option does not apply to the drives that the adapter already recognizes. For example, if autoconfiguration is set to JBOD and you remove and reinsert an unconfigured good drive, the drive remains unconfigured good. The adapter already knows the drive and does not reapply the autoconfiguration setting.

More information

RAID technologies

Converting multiple drives to JBOD

Converting multiple drives to unconfigured

Converting a drive to JBOD

Converting a drive to unconfigured

Reconfiguring a drive

Read load balancing

Read load balancing is a performance optimization feature in RAID controllers that distributes read requests across all drives in a mirrored set. In each mirrored pair, the controller monitors the workload on each drive and directs incoming read operations to the drive that currently has the least load.

This method has the benefit of enabling higher read performance and lower read latency.

More information

[Mirroring](#)

[RAID 1 and RAID 1+0 \(RAID 10\) - Mirroring](#)

[RAID technologies](#)

Drive rebuild

If a drive configured with RAID 1, 5, 6, 10, 50, or 60 fails, the firmware automatically rebuilds the data on a spare or replacement drive to prevent data loss. This rebuild process is fully automated. You can monitor the progress of drive rebuilds.

A manual rebuild is necessary if hot spares with enough capacity to rebuild the failed drives are not available. You must insert a drive with enough storage into the subsystem before rebuilding the failed drive. HPE MR932i-p PCIe SPDM PI Storage Controller supports rebuilding 2 drives in RAID 6 volume at the same time.

The rebuild operating mode specifies whether to prioritize rebuild or host I/O operations. You can choose from the following options:

- **Rebuild:** Gives priority to the rebuild operation over host I/O operations, resulting in faster rebuild completion but with some degradation in the host I/O performance.
- **Host I/Os:** Gives priority to host I/O operations over the rebuild process. Host and flush I/Os are not confined or diverted in this mode.



NOTE

If the failed drive is replaced while the server is powered off, and the newly inserted drive contains prior configuration metadata from this controller, the rebuild process will not start automatically.

More information

[RAID technologies](#)

[Configuring spare settings](#)

[Configuring task rates](#)

Foreign configuration import

Foreign configuration refers to the RAID configuration that exists on a replacement set of drives installed in a computer system. Using the HPE MR Storage Administrator, you can import the foreign configuration to the controller or clear it to create a new configuration on these drives.



NOTE

If the drives are configured using an HPE MR932i-p PCIe SPDM PI Storage Controller, you can import a foreign configuration. Importing a foreign configuration created on other controllers is not supported.

More information

[RAID technologies](#)

[Viewing and importing foreign configuration](#)

[Clearing foreign configuration](#)

Array transformations

Array transformation is the process of changing an existing RAID array to a different RAID level or configuration without losing data. It enables you to optimize performance, expand storage, or improve fault tolerance without recreating the array.

Subtopics

[Expand array](#)

More information

[Configuring task rates](#)

Expand array

Increase the capacity of an existing array by adding existing unassigned drives to it. Any drive you add must meet the following criteria:

- It must be an unassigned drive.
- It must be the same type as the existing drives in the array. For example, NVMe SSD or SAS SSD.

- It must have a capacity no less than the capacity of the smallest drive in the array.

The expand array feature is merged into the expand volume feature because expansion is supported only on non-span RAID levels. In these configurations, there is one first volume that consumes all the available raw capacity in the array. This method allows the controller to extend that volume as new drives are added.

More information

Transformation features

Volume transformations

Volume transformation is the process of modifying an existing logical volume within a RAID array to meet new requirements without deleting or recreating it. This transformation can include operations such as expanding capacity, changing RAID level, or migrating to different drive types. With volume transformation you can optimize storage resources, while preserving data.

Subtopics

Transportable controller

Expand volume

More information

Configuring task rates

Transportable controller

The controller firmware supports a transportable battery-backed cache memory to recover the data from a faulty server. This feature enables recovery by physically moving the entire controller to a replacement server.

The controller firmware assumes that the new server has the same configuration as the faulty server, including the same server generation and family. When the data is restored, the volumes are migrated to the new target server to facilitate cache flush.

More information

Transformation features

Expand volume

The expand volume feature allows you to increase the capacity of a volume by using unused space on existing disks, without requiring a system reboot. Extending a volume is not supported on striped RAID levels such as RAID 10, 50, and 60.

Increase the capacity of an existing volume by specifying a new size. Once the task is performed, use operating system partitioning tools to utilize the additional space.

Use the **Expand** option within the volume **Actions** menu in HPE MR Storage Administrator to enable this feature. For more information, see [**HPE MR Storage Administrator User Guide**](#).

More information

[Transformation features](#)

Drive health and maintenance

This section introduces the technologies and processes that maintain drive health and ensure data integrity in RAID environments. It includes features for monitoring, error correction, and updates to minimize downtime and prevent data loss.

Subtopics

[Hot-plug drive LED definition](#)

[EDSFF Carrier LED definition](#)

[Consistency check](#)

[Replace drive](#)

[Dynamic sector repair](#)

[Online drive firmware update](#)

[Predictive drive failure](#)

[Patrol read](#)

Hot-plug drive LED definition

Figure 1. LFF Low Profile (LP)





Figure 2. SFF Basic Carrier (BC)



ID	LED	Status	Definition
1	Fault\Locate ¹	Solid amber	The drive has failed, unsupported, or invalid.
		Solid blue	The drive is operating normally and being identified by a management application.
		Flashing amber/blue (one flash per second)	The drive has failed, or a predictive failure alert has been received for this drive; it also has been identified by a management application.
		Flashing amber (one flash per second)	A predictive failure alert has been received for this drive. Replace the drive as soon as possible.
2	Online\Activity	Solid green	The drive is online and has no activity. ²

ID	LED	Status	Definition
		Flashing green (4 flashes per second)	The drive is operating normally and has activity.
		Flashing green (one flash per second)	The drive is doing one of the following: • Rebuilding • Performing a capacity expansion • Performing a volume extension • Erasing • Spare part activation
		Off	The drive is not configured by a RAID controller or a spare drive.

- ¹ If for a failed drive the link is not working and the controller cannot detect the link, the Fault\Locate LED is Off.
- ² The solid green online LED indicates that the drive is a member of one or more logical drives.

More information


Component identification for HPE MR932i-p PCIe SPDM PI Storage Controller
Drive technologies

EDSFF Carrier LED definition

**Enterprise and Data Center Standard Form Factor (EDSFF) Carrier LED definition (Devon) - E3S
(76x112.75x7.5 mm)**



**Enterprise and Data Center Standard Form Factor (EDSFF) Carrier LED definition (Hereford) - E3S2T
(76x112.75x16.8 mm)**



ID	LED	Status	Definition
1	Fault\Locate	Solid amber	The drive has failed, unsu
		Solid blue	The drive is operating no management application.
		Flashing amber/blue (1Hz)	The drive has failed, or a received for this drive; it agement application.
		Flashing amber (1Hz)	A predictive failure alert l Replace the drive as soon
		Off	The drive is operating no by any management appl
2	Activity	Solid green	The drive is powered on a
		Flashing green	The drive is powered on a
		Off	The drive is powered off.

More information

[Component identification for HPE MR932i-p PCIe SPDM PI Storage Controller](#)

Consistency check

A consistency check operation verifies the correctness of the data in volumes that use RAID levels 1, 5, 6, 10, 50, and 60. For example, in parity-based systems, the check calculates data on one drive and compares the results with the contents of the parity drive.

Run consistency checks periodically on fault-tolerant volumes. RAID 0 does not provide data redundancy, so you cannot run the consistency check on RAID 0 volumes.

More information

Drive technologies

Configuring task rates

Checking the consistency of a volume

Replace drive

The replace drive operation allows you to replace failed physical drives in the array with healthy drives. The original array and volume numbering remain after the replacement.

Conditions and restrictions

- The replacement drives must use the same interface type as the original drives.
- The operation is available only if enough unassigned drives of the correct size are available.
- The array must not be undergoing a transformation process. For example, rebuilding to a spare.
- The array must have a working cache, making it capable of transformation.

More information

Drive technologies

Replacing a drive

Dynamic sector repair

Disk drive media can develop defects caused by variances in the drive mechanisms under normal operating conditions. The dynamic sector repair feature helps protect the controller from media defects.

Dynamic sector repair proactively addresses this issue by detecting and repairing defective sectors on the drive. When a media defect is identified, the controller automatically remaps the affected sector to a healthy location, ensuring continued access to data and preventing potential failures. This process helps maintain optimal performance and protects the storage system from media-related errors without requiring manual intervention.

Benefits

- Maintains data integrity by proactively addressing media defects.
- Reduces the risk of data loss and unplanned downtime.
- Extends the usable life of disk drives by isolating and remapping bad sectors.

More information

Online drive firmware update

Online drive flashing saves time by reducing downtime when updating drive firmware. Instead of taking drives offline to load a new firmware image, you can download an updated drive firmware to the controller and update all the drives while the server remains online.

Benefits

- Minimizes or eliminates system downtime during drive firmware updates.
- Simplifies maintenance and improves operational efficiency.
- Ensures that drives are running the latest firmware, which can enhance performance, reliability, and compatibility.

More information

Drive technologies

Predictive drive failure

Predictive drive failure is a monitoring feature that uses Self-Monitoring, Analysis, and Reporting Technology (SMART) to detect abnormal drive behavior that is likely to lead to failure.

SMART technology is embedded within the drive firmware. It continuously monitors a range of internal parameters, such as performance metrics, calibration data, and error rates. These checks are specific to each drive type. When these monitored values exceed predefined thresholds, the drive reports a predictive failure condition to the controller. Upon detection, the controller informs the host system and can trigger automated actions.

Benefits

- Predictive failure alerts allow administrators to replace drives before failure occurs⁴.
- SMART polling intervals and corrective actions can be configured through management utilities.
- Predictive failure events are integrated with system event logs and can be monitored.

More information

Drive technologies

Patrol read

Patrol read periodically verifies all sectors of the drives connected to a controller, including the system-reserved area in the RAID configured drives. You can run a patrol read on all RAID levels and spare drives.

Benefits

- **Comprehensive scanning:** Patrol read checks every sector on all drives, including spare drives, regardless of RAID level.
- **Idle-time operation:** The controller initiates the process only when the controller is idle for a defined period and has no background activities.
- **Configurable rate:** Adjust the patrol read rate from 1% to 100% using management tools. Higher rate results in faster scanning but may reduce system I/O performance.
- **On-demand or scheduled:** Start patrol read manually or schedule it to run at regular intervals.
- **Automatic repair:** Detects and repairs latent media defects before they cause data loss.

More information

[Drive technologies](#)

[Configuring patrol read settings](#)

Security

This section explains the technologies and protocols that ensure data security and compliance in storage environments. These features help protect sensitive information, prevent unauthorized access, and maintain regulatory compliance.

Subtopics

[Drive erase](#)

[Sanitize](#)

[Self-encrypting drive](#)

[Cache offload encryption](#)

[Hardware secure boot](#)

[Security Protocol and Data Model](#)

Drive erase

The drive erase option securely removes data from drives by overwriting every user-accessible sector with specified patterns. This process ensures that previously stored data cannot be recovered. The erase operation runs as a background task and can be repeated in multiple passes for enhanced security.

You can use the drive erase option on physical drives and volumes. For the physical drive, change the drive state to **Unconfigured Good**.

After the erase operation has started, you can stop the erase option using the **Actions** menu. To erase a drive securely, see “Erasing a Drive Securely” in the [HPE MR Storage Administrator User Guide](#).

Erase types

Simple

Writes a single pattern to the volume in one pass.

Normal

A three-pass operation that first overwrites the drive contents with random values, then overwrites it twice with patterns.

Thorough

Repeats the Normal erase operation three times for maximum security.

More information

Security features

Erasing a drive cryptographically

Sanitize

The sanitize crypto erase feature uses cryptographic technology to permanently remove all user data from a drive, ensuring that sensitive information cannot be recovered. This method has the following benefits:

- Removes all sensitive information from the drive ensuring compliance with data security standards.
- Completes within seconds offering fast and reliable data removal options.
- Provides multiple methods for different drive types and security needs.

Sanitize methods

Sanitize methods vary based on drive type and security requirements.

Sanitize Crypto Erase (Instant Secure Erase)

This method uses cryptographic technology to perform an instant secure erase of all user data.

Sanitize Overwrite (Hard Drives)

This method fills every physical sector of the drive with a predefined pattern. Once started, the drive continues sanitizing regardless of resets or power cycles.

Sanitize Block Erase (SSDs)

This method sets all blocks on the SSD to a vendor-specific value, effectively removing all user data. Once started, the drive continues sanitizing regardless of resets or power cycles.

More information

Security features

Sanitizing an unconfigured good drive

Self-encrypting drive

Self-Encrypting Drive (SED) secures data from unauthorized access or modification by encrypting all data stored on the drive. Even if the SED is removed from its storage system, the data remains inaccessible without the appropriate security authorization.

The controller does not allow you to disable drive security function when any secured JBOD or drive is present. To disable drive security, you must delete or unsecure the drives or remove them from the controller. If autoconfiguration is set to secure JBOD, the controller also prevents drive security from being disabled.

Key management options

Host key management

- Enable the SED drive as JBOD and expose it to the operating system.
- Manage the SED using third-party key management tools such as `SEDutil`.
- Monitoring is available through HPE MR Storage Administrator, HPE StorCLI2, and UEFI System Utilities.

Local key management

- Enable SED security using HPE MR Storage Administrator, HPE StorCLI2, or UEFI System Utilities.
- Provide a controller-wide security key identity and security key.
- During boot, the controller uses the stored security key to unlock the drive. When powered down, the encryption key is locked, protecting against theft.

Remote Key Management (External Key Management)

- Configure HPE iLO Key Manager before enabling remote key management.

- Enable remote key management using UEFI System Utilities. For more information, see [Enabling security](#).
- Configuration utility in UEFI System Utilities works with iLO key manager to create and store the security key identity and key on a remote key manager server.
- When drive is powered down, the encryption key is locked. During boot, the security key is retrieved from the remote key manager server to unlock the drive.
- The controller does not allow you to disable drive security if there are any MR-secured physical drives present in the controller. To disable drive security, delete or unsecure all JBODs and physical drives or remove them from the controller.
- If autoconfiguration is set to secure JBOD, disabling drive security is not allowed.

More information

[Enabling security](#)

[Disabling drive security](#)

[Changing security settings](#)

[Security features](#)

Cache offload encryption

Cache offload encryption protects data stored in the controller's write cache during an unexpected power loss. When the controller offloads cached data to non-volatile memory, it encrypts the data to prevent unauthorized access.

After power is restored, the controller decrypts and reloads the data to complete pending write operations, ensuring both data integrity and security.

More information

[Security features](#)

Hardware secure boot

Hardware secure boot permits only authenticated firmware to run on the controller. The controller boots from an internal boot ROM, which establishes the initial root of trust (RoT). Using this RoT, the controller verifies each subsequent firmware image and builds a chain of trust, ensuring that only authorized firmware is used.

The hardware secure boot process is transparent unless the adapter encounters a counterfeit firmware image. If the controller downloads a counterfeit image, the image authentication fails and the download utility displays failure messages.

More information

Security features

Security Protocol and Data Model

Security Protocol and Data Model (SPDM) is an industry-standard security framework developed by the Distributed Management Task Force (DMTF). It enables system hardware components such as PCIe cards, NVMe drives to authenticate their identities and verify their integrity.

SPDM-capable components contain strong cryptographic identities and can provide cryptographically signed attestations of their security state. During server startup, the system cryptographically authenticates the components and collects measurements of their security-relevant properties to confirm that each device is operating in its intended state. After successful verification, the system passes control to the operating system.

More information

Security features

Reliability

Reliability includes cache error checking and correction, which detects and fixes memory errors to protect data integrity. It also covers thermal monitoring, which tracks temperature to prevent overheating and maintain performance. Together, these features help reduce hardware failures and keep the system running efficiently.

Subtopics

Recovery ROM

Cache Error Checking and Correction

Thermal monitoring

Recovery ROM

The controllers store a redundant copy of the controller firmware image to protect against data corruption. If the active firmware image becomes corrupt, the controllers use the redundant firmware image and continue

operating. Once the redundant firmware is activated, the corrupted firmware image is then updated with valid firmware by the active firmware. The recovery ROM provides protection against power outages during firmware flashing.

More information

[Reliability features](#)

Cache Error Checking and Correction

Cache Error Checking and Correction (ECC) DRAM technology protects the integrity of data stored in the controller's cache memory. The ECC mechanism creates extra check bits for each block of data and uses them to detect and correct single-bit errors automatically.

It also detects multi-bit errors that cannot be corrected, preventing corrupted data from being written to disk. This protection helps ensure reliable cache operation and maintains overall system data integrity.

More information

[Reliability features](#)

Thermal monitoring

The controller monitors the temperature of each drive in the server. HPE iLO periodically collects these drive temperatures from the controller to control the fan speed. The fan speed is optimized so that each drive is maintained below its maximum continuous operating temperature regardless of the workload.

This method has the benefit of saving cost by allowing the fans to run at an optimal setting while ensuring that drives do not overheat.

More information

[Reliability features](#)

Performance

The performance features of the controller help optimize how data is transferred, cached, and processed across connected drives. These features enable you to administer different workloads, reduce latency, and ensure performance under varying I/O conditions.

Subtopics

[Manage SAS storage link speed](#)

[Manage PCIe storage interface](#)

[I/O performance mode](#)

[Cache](#)

[Multiple physical functions](#)

Manage SAS storage link speed

The Manage SAS Storage Link Speed feature displays 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 (physical links) in a SAS port can have different link speeds or can have the same link speed.

More information

[Performance features](#)

[Managing SAS storage link speed](#)

Manage PCIe storage interface

The Manage PCIe storage interface feature displays the lane speed between the controller and an expander or between the controller and a drive that is directly connected to the controller. A lane represents a set of differential signal pairs similar to SAS phys, one pair is for transmission, and one pair is for reception.

More information

[Performance features](#)

[Managing PCIe storage interface](#)

I/O performance mode

I/O performance mode, also known as FastPath, is an intelligent I/O passthrough mechanism to improve performance of SSD arrays. FastPath boosts storage subsystem and overall application performance, especially for applications that demonstrate high-random read/write operation workloads.

I/O performance mode is enabled for a volume when the volume is created with properties **Direct I/O**, **Write-Through**, and **No Read Ahead**.

When the volume is in optimal state, I/O performance mode supports all RAID levels.

More information

Cache

The cache features improve performance by temporarily storing frequently accessed or write-pending data in high-speed memory. These settings allow the controller to manage how data is read, written, and protected in the cache, helping balance speed, reliability, and data integrity.

Subtopics

[Read cache policy](#)

[Write cache policy](#)

[Flash-backed write cache](#)

[I/O policy](#)

[Drive caching](#)

[Strip size selection](#)

More information

[Configuring cache and memory settings](#)

Read cache policy

The read cache policy option is always set to **No Read Ahead**, which provides the best performance for SSD.

More information

[Performance features](#)

Write cache policy

The write cache policy options for the volume are:

- **Write-Through:** Sends a data transfer completion signal to the host when the drive subsystem has received all the data in a transaction. This mode may result in slower performance as it does not use the controller cache.
- **Write-Back:** Sends a data transfer completion signal to the host when the controller cache has received all the data in a transaction. This option provides a good balance between data protection

and performance as the controller switches between write-back and write-through depending on the controller status.

- **Always Write-Back:** Sends a data transfer completion signal to the host when the controller cache receives all the data in a transaction. If you select the always write-back policy and the energy pack is absent, the firmware is forced to use the write-back policy.

More information

Performance features

Flash-backed write cache

Write-back caching scheme allows the host applications to continue running without waiting for data to be written to disk. With write-back caching, the controller stores incoming data in a high-speed cache and immediately reports completion to the OS. This reduces write latency from milliseconds to microseconds. The controller then writes the cached data to the drives later, at the most optimal time.

Flash-Backed Write Cache (FBWC) uses flash devices to retain cache data and the embedded Energy Pack to supply power during a power loss. Unlike Battery-Backed Write Cache (BBWC), which must provide power for the entire outage, FBWC needs power only long enough to copy data from DRAM to flash. The controller writes the saved data to the drives on the next power-up.

HPE MR932i-p PCIe SPDM PI Storage Controller embedded Energy Pack is not a user-replaceable part. Do not hot-plug the embedded Energy Pack. Removing or inserting an embedded Energy Pack while the controller is powered on can damage the board or the supercapacitor.

More information

Performance features

I/O policy

The I/O policy controls how a volume handles read operations. It does not affect the read-ahead cache. The direct I/O option is available for volumes.

In direct I/O mode, the controller does not buffer reads in cache. Instead, it sends data to the host and cache at the same time. If the same data is requested again, it is served from the cache.

Direct I/O is the default setting.

More information

Performance features

Drive caching

Drive cache is disabled by default. You can choose one of the following drive cache settings for a volume:

- **Default:** Uses the factory-recommended write cache setting for the drive.
- **Enable:** Turns on the drive write cache to improve write performance.
- **Disable:** Turn on the drive write cache to prioritize data integrity over performance.

More information

[Performance features](#)

Strip size selection

When you create a volume, the controller stores data on each drive in units called strips. The strip size is fixed at `64 KiB`, which is supported only for SSDs.

More information

[Performance features](#)

Multiple physical functions

Multiple physical functions (MPF) improve performance by allowing the controller to expose several true PCIe physical functions to the system. This capability increases I/O parallelism and enables high performance in both standard and virtualized server environments. In virtualized environments, MPF allows virtual machines to achieve near bare-metal performance. MPF also supports better server consolidation and storage isolation.

Each physical function (PF) uses ID-based routing as defined in the PCI Express Base Specification 4.0. The bus number in the configuration cycle TLP header identifies the PCIe function. As the physical functions appear through a virtual switch, the function number for each device is always zero.

PF type	Total number in the controller	Device ID	Description
Supervisor	One	0x00B3	Manages configuration and exposes assigned JBODs and volumes. As the supervisor PF controls configuration, allow only trusted applications or VMs to access it. Use UEFI, HPE MR Storage Administrator, or HPE StorCLI2 to identify the supervisor PF.
Nonsupervisor	Three	0x00B5	Exposes the JBODs or volumes assigned to them, and cannot perform configuration tasks. You cannot enable or disable nonsupervisor PF.

More information

- [Mapping a volume to a function](#)
- [Removing a volume from function](#)
- [Performance features](#)
- [Managing function profile](#)
- [Managing controller function](#)

Installation

Provides instructions and requirements for installing the controller, connecting storage, and preparing the system for operation.

Use the topics in this section to install the controller into a server that is already configured or a server that is not yet configured.

Subtopics

- [Installing in a configured server](#)
- [Installing in an unconfigured server](#)
- [Installing an MR controller on the server](#)

More information

Component identification for HPE MR932i-p PCIe SPDM PI Storage Controller

Installing in a configured server

Procedure

1. Back up data on the system.
2. Close all applications.
3. Update the server firmware, if it is not the latest revision.
4. If the new controller is the new boot device, install the device drivers.
5. Ensure that users are logged off and all tasks are completed on the server.
6. Power down the server.
7. Power down all peripheral devices that are attached to the server.
8. Disconnect the power cord from the power source.
9. Disconnect the power cord from the server.
10. Remove or open the access panel.

For server-specific procedures, see the server user guide.



WARNING

To reduce the risk of personal injury from hot surfaces, allow the drives and the internal system components to cool before touching them.

1. Remove the riser.
2. Select an available x16 CEM slot.



TIP

Only x16 link width is supported.

- .3. Remove the slot cover.

Save the retaining screw, if one is present.

- .4. Slide the controller along the slot alignment guide, if one is present, and then press the board firmly into the expansion slot so that the contacts on the board edge are seated properly in the slot.
- .5. Secure the controller in place with the retaining screw. If the slot alignment guide has a latch near the rear of the board, close the latch.
- .6. Connect storage devices to the controller.

For cabling information, see the server user guide.
- .7. Reinstall the riser.
- .8. Connect the power cord to the server.
- .9. Connect the power cord to the power source.
- .10. Power up all peripheral devices.
- .11. Power up the server.

Installing in an unconfigured server

Procedure

1. [Update the server firmware.](#)
2. Select an available x16 CEM slot.



TIP

Only x16 link width is supported.

3. Remove the slot cover.

Save the retaining screw, if one is present. For server-specific procedures, see the server user guide.

4. Slide the controller along the slot alignment guide, if one is present. Press the board firmly into the expansion slot so that the contacts on the board edge are seated properly in the slot.
5. Secure the controller in place with the retaining screw. If the slot alignment guide has a latch near the rear of the board, close the latch.

6. Install physical drives, as required.
7. Connect the drive backplane to the controller.

For cabling information, see the server user guide.

8. Power up the server.

If this controller is the only one installed in the server, and it is the boot device configure a volume, or select the physical drives for JBOD mode using the configuration utility in UEFI System Utilities.

9. Proceed to install the OD on the controller driver.

Subtopics

Installing the OS with the controller driver

Installing the OS with the controller driver

Prerequisites

Ensure that you have the controller driver available. Obtain it by extracting it from the Service Pack for ProLiant (SPP) (<https://www.hpe.com/servers/spp>) or by downloading it from the Hewlett Packard Enterprise website.

Procedure

1. Power on the server.
2. Configure a volume or select physical drives for JBOD mode using the configuration utility in UEFI System Utilities.
3. Launch the OS installation and point to the controller driver when prompted.
4. If the OS utility does not see the drive, repeat the step to configure the volume or select physical drives for JBOD and retry the installation.
5. After the installation is complete, use the iLO GUI or the SPP to update the controller firmware if it is not the latest revision. For more information on firmware update through iLO, see the HPE iLO User Guide.

More information

Operating environments

Installing an MR controller on the server

For product-specific installation information:

1. Go to Hewlett Packard Enterprise Support Center <https://www.hpe.com/support/hpesc>
2. Search for the user guide of your specific server and open it.
3. Go to **Hardware options installation** and select **Storage controller options**. You will find the required installation information here.

Connecting storage devices

For more information about supported drive models on specific HPE ProLiant servers, see [QuickSpecs](#) for the specific server.

Subtopics

[Connecting internal storage](#)

Connecting internal storage

Procedure

1. Power down the server.
2. Install drives, if necessary.

Hewlett Packard Enterprise recommends drives of similar type. All drives grouped in a volume must meet the following criteria:

- They must be either NVMe or SAS.
- They must be either all hard drives or all solid state drives.
- For the most efficient use of drive space, the drives must have comparable capacity.

For more information about drive installation, see Server documentation.

3. Use the internal SAS cable identified in the server that is compatible with the controller. For more information, see [HPE Compute MR Controllers QuickSpecs](#).
 - If the drives are hot-plug capable, connect the internal connector of the controller to the connector on the hot-plug drive cage.

- If the drives are not hot-plug capable, connect the internal connector of the controller to the non-hot-plug drives.

4. Close or install the access panel, and secure it with thumbscrews, if any are present.



CAUTION

Do not operate the server for long periods with the access panel open or removed. Operating the server in this manner results in improper airflow and improper cooling that can lead to thermal damage.

5. Power up the server.

Cable part numbers

For more information on cables, see the server [QuickSpecs](#).

Configuration and maintenance tools for HPE MR932i-p PCIe SPDM PI Storage Controller

Provides an overview of supported configuration and maintenance tools used to manage the controller.

After initial provisioning of the server or compute module, you can use any of the following tools to set up, configure, and maintain the controllers.

Tools	Category	User Guides
	Pre-OS & Post-OS	<u>HPE Storage Controllers – Management overview</u>



[Redfish](#)

Tools	Category	User Guides
	Pre-OS & Post-OS	Configuration and management using UEFI System Utilities
UEFI System Utilities		
	Post-OS	HPE MR Storage Administrator User Guide  IMPORTANT Manage HPE MR932i-p PCIe SPDM PI Storage Controller using the HPE MR Storage Administrator. Smart Storage Administrator does not manage HPE MR storage controllers.
HPE MR Storage Administrator		
HPE StorCLI2	Post-OS	HPE StorCLI2 User Guide  IMPORTANT To manage the HPE MR932i-p PCIe SPDM PI Storage Controller, use HPE StorCLI2. For HPE MR Gen10 Plus and Gen11 controllers, use HPE StorCLI.
	Pre-OS & Post-OS	 IMPORTANT Intelligent Provisioning does not currently support HPE MR932i-p PCIe SPDM PI Storage Controller. Support will be added in a future release.
Intelligent Provisioning		

Server management tools

For initial provisioning of the server or compute module, use the following server management tools.



HPE iLO

[HPE iLO Documentation Quick Links](#)

HPE OneView

[HPE OneView Documentation Quick Links](#)

Subtopics

[Redfish](#)[UEFI System Utilities](#)[HPE MR Storage Administrator](#)[HPE StorCLI2](#)

Redfish

These controllers support the DMTF standard known as PLDM for Redfish device enablement in HPE ProLiant servers. This open standard API allows HPE option cards, storage controllers to host their own set of Redfish resources and capabilities, which are rooted under the `iLO /redfish/v1` service root. As a result, the option card firmware owns the feature and capabilities.

For more information on Redfish APIs, see [HPE Storage Controllers: Management overview](#).

More information

[Error reporting](#)

UEFI System Utilities

The UEFI System Utilities is embedded in the system ROM. Its features enable you to perform a wide range of configuration activities, including the following:

- Configuring system devices and installed options
- Enabling and disabling system features
- Displaying system information
- Selecting the primary boot controller
- Configuring memory options
- Launching other preboot environments

For onscreen help, press **F1**.

For more information on the UEFI System Utilities, see [Configuration and management using UEFI System Utilities](#).

Subtopics

Controller boot process in UEFI

Launching the system utilities

Resolving controller boot blocking events in UEFI

Controller boot process in UEFI

This topic explains how the controller behaves during the boot process and describes conditions that can prevent the controller from completing initialization.

Boot process

During system startup, the controller firmware attempts to complete initialization automatically. By default, the firmware handles non-critical errors without your involvement by continuing, ignoring, or bypassing the issue.

However, when the firmware detects a critical condition that cannot be bypassed, the system displays a message that configuration is required and either:

- Boots in safe mode, or
- Halts the boot process entirely

These behaviors protect data integrity and prevent the controller from operating in an unsupported or unstable state.

Also, if the system requires the passphrase the system displays a message that configuration is required and pauses the boot process.

Blocking boot events

Certain critical conditions trigger blocking boot events. When these events occur, the controller stops the boot process and requires your intervention before normal operation can resume.

Some hardware-level issues, particularly the ones related to controller configuration or NV cache, may be unrecoverable and prevent the controller from initializing successfully.

Blocking boot events can occur due to any of the following reasons:

- Configuration errors, such as premium features not being programmed.
- Detection of preserved cache.
- Non-volatile cache issues, such as capacity being full, back up failure due to high bad block count, or initialization failure.
- Controller reaching fatal temperature.
- Communication errors with the option ROM.

Passphrase at boot time

While enabling security, if **Passphrase at Boot Time** is set, the system displays a message during boot indicating that configuration is required. The system does not prompt for passphrase entry during boot. To continue, follow the standard boot event resolution procedure.

More information

[Resolving controller boot blocking events in UEFI](#)

[Enabling security](#)

[Changing security settings](#)

Launching the system utilities

Procedure

1. (Optional) If you access the server remotely, start an iLO remote console session.
 - a. Open a browser and enter `https://<iLO host name or IP address>` to log on to the iLO web interface.
 - b. On the login page, enter a directory or local user account name and password, and then click **Log In**.
 - c. Launch **HTML5 Remote Console**.

- i. Select **Host** in the iLO navigation tree, and then select the **Remote Console**.
 - ii. Select the **Console** thumbnail at the lower left corner of the iLO web interface.
2. Restart or power on the server.

The server restarts and the POST screen appears.

3. (Optional) If the system detects any blocking boot events or requires the passphrase, follow this [procedure](#) to resolve the condition and continue the boot process.
4. Press **F9**.

The **System Utilities** screen appears.



NOTE

Using System Utilities requires BIOS administrator authorization. If the BIOS administrator requires a password, the server prompts for the password to be entered prior to launching the System Utilities. For information on setting the administrator password, see “Server Security options” in the [HPE UEFI System Utilities User Guide](#).

More information

[Controller boot process in UEFI](#)

[Resolving controller boot blocking events in UEFI](#)

[Enabling security](#)

[Changing security settings](#)

Resolving controller boot blocking events in UEFI

About this task

During system boot, the system displays a configuration required message in the following scenarios:

- The controller detects a critical condition, cannot complete initialization, and starts in safe mode with limited functionality.
- The system requires the passphrase to continue.

Procedure

1. From the **System Utilities** screen, select **System Health**.

2. Select **View System Health**.
3. In the **System Health Summary** screen, select **HPE MRXXX**.
4. Follow the instructions on the screen to enter the passphrase or resolve the blocking event.

More information

[Controller boot process in UEFI](#)

[Changing security settings](#)

[Enabling security](#)

HPE MR Storage Administrator

HPE MR Storage Administrator is a web-based application that enables you to monitor, configure, maintain, and troubleshoot MR controllers. HPE MR Storage Administrator enables you to view, create, and manage storage configurations.



IMPORTANT

The HPE MR Storage Administrator manages **only** the MR controllers. It does not manage SR controllers.

- **Monitoring and configuring:** HPE MR Storage Administrator enables you to monitor the controllers and configure the drives on the controller. It displays the status of the controller cards, logical drives, and drives on the controller. The device status icons notify you if there are drive failures and other events that require your immediate attention. Email notifications about 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. You can also import or clear foreign configurations.



IMPORTANT

Limited alerting and monitoring is available through iLO when HPE Agentless Management Service is installed. This controller supports limited alerts through iLO using the AMS agent.

- **Maintaining:** Using HPE MR Storage Administrator, you can perform system maintenance tasks, such as updating the controller firmware.
- **Troubleshooting:** HPE MR Storage Administrator displays information related to drive failures, device failures, and other issues. It also provides recommendations and displays contextual links, helping you to locate the drives/devices that have issues and troubleshoot them. You can also download a report of

the devices and their configurations, properties, and settings and send it to Hewlett Packard Enterprise Support for further troubleshooting.

Obtain HPE MR Storage Administrator installation files through the Service Pack for ProLiant (SPP), which you can download from the Hewlett Packard Enterprise website (<https://www.hpe.com/servers/spp/download>). Be sure to use the latest SPP version for the server.

For more information about the HPE MR Storage Administrator, see [**HPE MR Storage Administrator User Guide**](#).

More information

[**Diagnostic tools**](#)

HPE StorCLI2

The Storage CLI (HPE StorCLI2) tool is the command-line management software designed for the HPE MR Controller. HPE StorCLI2 is designed to be easy to use, consistent, and easy to script.

Obtain HPE StorCLI2 through the HPE Service Pack for ProLiant (SPP), which you can download from <https://www.hpe.com/servers/spp/download>. Be sure to use the latest SPP version for the server.

For more information about HPE StorCLI2, see [**HPE StorCLI2 User Guide**](#).

More information

[**Diagnostic tools**](#)

Configuration and management using UEFI System Utilities

Provides guidance for configuring, managing, and monitoring the controller using UEFI System Utilities.

This section tells you how to configure controllers, physical drives, virtual drives, and perform other configuration tasks using the **UEFI System Utilities** interface.

Subtopics

[**UEFI System Utilities GUI**](#)

[**Navigating the UEFI dashboard view sections**](#)

[**Configuration management**](#)

[**Controller management**](#)

[**Advanced controller management**](#)


[**Advanced controller properties**](#)

[**Volume management**](#)

UEFI System Utilities GUI

Both mouse and keyboard devices are supported on the UEFI System Utilities GUI.

Regions



Item	Region	Description
1	Caption Bar	UEFI form title and the system buttons. The Form title shows the name of the form that you are currently navigating.
2	Navigation History	Forms to which you navigated previously. A Navigation History node is added to the navigation h

Item	Region	Description
		istory each time you visit a new system utility form.
3	Server Information	Server and function key information.
4	System Utilities Form	Menu options available on the current form.
5	Activity Bar	System-wide functions, such as function keys and the system status indicator.


Keyboard support in the GUI

The GUI has support for basic keys to navigate the system utilities form.

Keyboard key	Behavior
Tab	Changes focus to the different regions.
Up and Down arrows	Selects a menu option. When selected, the color of a menu option changes to green.
Enter	Selects an entry. The selected option changes color to green. When a submenu is chosen, the submenu appears.
Esc	Exits full screen mode.
F1	Displays online help in text mode.  NOTE To access online help on the GUI mode, click ? icon.
F7	Loads default UEFI BIOS configuration settings.  NOTE Pressing F7 only resets the BIOS configuration. It does not reset other entities, such as option cards or HPE iLO.
F10	Prompts you to save changed settings.

Keyboard key	Behavior
F12	Prompts you to save changed settings, and then exits the System Utilities.
Reboot Required	Radio button turns red when changes require a reboot the server.
Changes Pending	Radio button turns red when pending changes are to be saved to take effect.


Navigation history region and keyboard support



Navigation History shows the forms that you have previously opened. The utility adds a new Navigation History node each time you visit a form. You can click any node to return to that form.

If the bar contains more nodes than it can display, the home node collapses. Selecting the home node opens a pop-up list of all visited forms, and you can click any item in the list to return to that form.

Features



Item	Feature	Description
1	Mouse speed	Control the speed of the mouse.
2	Help	Context-sensitive help for the page.
3	Keyword search	Search for an RBSU option.
4	Select Language	Choose the language for the current page.
5	Forced Write Information	Displays options that are forced to change due to changing an option.
6	Pending Changes	Lists changes that have not been saved.

Navigating the UEFI dashboard view sections

Procedure

1. From the **System Utilities** screen, select **System Configuration**.
2. Select the **HPE MRXXX** controller that you want to view or configure.

The following sections of the dashboard view are displayed.

- **Main Menu:** Click to access the [main menu](#) options.
- **Help:** Click to view the help content.
- **PROPERTIES:** Displays the controller details.
- **ACTIONS:** Lists the common [actions](#) for the controller.
- **BACKGROUND OPERATIONS:** Displays volume and drive operations that are in progress.

Subtopics

[Main Menu](#)

[Actions](#)

Main Menu

Main Menu is one of the sections in the dashboard view, and it offers the following management options:

- [Configuration Management](#)
- [Controller Management](#)
- [Volume Management](#)
- [Device Management](#)
- [Energy Pack Management](#)

Actions

ACTIONS is one of the sections in the dashboard view, and it gives you quick access to key tasks:

- [View Foreign Configuration](#) - This option is available when a foreign drive is available.
- [Configure](#)
- [Manage Controller Personality](#)
- [Manage Advanced Software Options](#)
- [Set Factory Defaults](#)

Configuration management

The **Configuration Management** screen is an interface for creating volumes and selecting the drives to use. It allows managing foreign configurations, clearing or converting configurations, and performing setup tasks for RAID and JBOD environments.

Subtopics

[Creating volume](#)

[Creating profile-based volume](#)

[Viewing array properties](#)

[Converting multiple drives to JBOD](#)

[Converting multiple drives to unconfigured](#)

[Clearing a configuration](#)

[Viewing and importing foreign configuration](#)

[Clearing Escrow key](#)

[Clearing foreign configuration](#)

Creating volume

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Configuration Management > Create Volume**.
3. In the **Create Volume** screen, set the required values.
4. To select drives, complete the steps [here](#), and return to the **Create Volume** screen to configure the next fields.
5. (Optional) Select the checkbox to **Enable Unmap/Write Same Unmap** for a non-NVMe drive.

6. (Optional) Select the checkbox to **Secure Virtual Drive** with SED.
7. (Optional) Select the checkbox to **Ignore Auto-assign Policy**.

**TIP**

If **Ignore Auto-assign Policy** is set, the system does not assign the volume to any physical function. The host cannot discover the volume until you manually map the volume to a physical function.

8. Click **Save Configuration**.
9. Review and address any warning messages, as required.
10. Select the **Confirm** checkbox, and click **Yes** to proceed.

Subtopics**Selecting drives for volume creation****More information****RAID levels****Managing auto-assign policy**

Selecting drives for volume creation

About this task

Use the steps in this topic to select the unconfigured drives that you want to include in the volume.

Procedure

1. In the **Select Drives** screen, filter the type of drive by selecting an option from the **Select Media Type** drop-down list.

Option is **SSD**.
2. Select the interface type from the **Select Interface Type** drop-down list to filter based on interface.

Options include **SAS**, **NVMe**, or **All**.
3. Specify the sector size from the **Logical Sector Size** drop-down list to further filter based on size.

Options include **512B**, **4KiB**, or **Both**.

4. From the **CHOOSE UNCONFIGURED DRIVES** list, select the checkbox for the drives you want or click **Select All** to add all the drives.

5. Click **OK**.

The **Success** screen appears to confirm that the selection was performed successfully.

6. Click **OK**.

The **Create Volume** screen displays.

7. In the **Create Volume** screen, continue to configure the rest of the fields. Proceed to the next field [here](#).

Creating profile-based volume

About this task



WARNING

Creating a volume will permanently delete any data on an associated drive.

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Configuration Management > Create Profile Based Volume**.
3. In the **Create Profile Based volume** screen, select a RAID level.
4. In the **Generic RAID <level>** screen, select the criteria from the **Drive Selection Criteria** drop-down list.
5. Confirm the **PROFILE PARAMETERS**.
6. (Optional) For RAID5, configure **File Server**, or **Web/Generic Server** build an array that provides high data throughput and data redundancy for small, random access workloads.
7. (Optional) For RAID 10, configure **Database** to build an array that provides excellent fault tolerance and low latency.
8. Click **Save Configuration**.
9. Review and address any warning messages, as required.
10. Select the **Confirm** checkbox, and click **Yes** to proceed.

More information

RAID levels

Read cache policy

Write cache policy

Flash-backed write cache

Strip size selection

Drive caching

Initialization operations

Viewing array properties

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Configuration Management**.
3. Select **View Array Properties**.

Results

All the properties for the array are displayed.

Converting multiple drives to JBOD

Prerequisites

Ensure that you have backed up data before changing the state of the drives.

About this task

You can convert the state of physical drives from an unconfigured state to JBOD mode. In JBOD mode, each disk is presented individually to the host operating system without RAID features such as striping, mirroring, or parity.



WARNING

Converting drives to JBOD will permanently delete any data on the drives.

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.

2. Select **Main Menu > Configuration Management > Convert to JBOD**.
3. (Optional) In the **Convert to JBOD** screen, select the checkbox to **Ignore Auto-Assign Policy**.
4. From the **SELECT DRIVES TO CONVERT TO JBOD** list, select the checkbox for the drives you want or select the **Check All** checkbox to select all the drives.
5. Click **OK**.
6. Review and address any warning messages, as required.
7. Click **OK**.

More information

[Mixed mode - RAID and JBOD simultaneously](#)
[Managing auto-assign policy](#)

Converting multiple drives to unconfigured

Prerequisites

Ensure that you have backed up data before changing the state of the drives.

About this task

You can convert the state of physical drives from their current state to an unconfigured state.



WARNING

Converting drives to unconfigured will permanently delete any data on the drives.

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Configuration Management > Convert to Unconfigured**.
3. From the **SELECT DRIVES TO CONVERT TO UNCONFIGURED** list, select the checkbox for the drives you want to choose or select the **Check All** checkbox to select all the drives.
4. Click **OK**.
5. Review and address any warning messages, as needed.
6. Click **OK**.

More information

Mixed mode - RAID and JBOD simultaneously

Clearing a configuration

About this task



WARNING

Clearing a configuration deletes all the volumes and spare drives attached to the controller.

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Configuration Management > Clear Configuration**.
3. In the warning message screen, click **Confirm** and **Yes** to proceed.

The **Success** screen appears to notify that the operation is complete.

Viewing and importing foreign configuration

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Configuration Management > Manage Foreign Configuration**.
3. Select **Import Foreign Configuration**.
4. Review and address any warning messages, as required.
5. Select the **Confirm** checkbox, and click **Yes** to proceed.

More information

Foreign configuration import

Clearing Escrow key

About this task

When you clear the escrow key, the controller deletes all escrow keys and then locks any foreign secured drives.

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Configuration Management > Manage Foreign Configuration**.
3. Select **Clear Escrow Key**.
4. Review and address any warning messages, as required.
5. Select the **Confirm** checkbox, and click **Yes** to proceed.

Clearing foreign configuration

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Configuration Management > Manage Foreign Configuration**.
3. Select **Clear Foreign Configuration**.
4. Review and address any warning messages, as required.
5. Select the **Confirm** checkbox, and click **Yes** to proceed.

More information

[Foreign configuration import](#)

Controller management

The **Controller Management** screen provides an interface to view controller status and properties, perform advanced management tasks, adjust controller properties and behavior. The advanced options are nested in this screen.

Subtopics

Viewing controller properties

Viewing controller properties

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Controller Management**.
3. In the **Controller Management** screen, view the basic properties of the controller:
4. (Optional) Click **Advanced Controller Management** to perform advanced controller tasks. For more information, see Advanced controller management.
5. (Optional) Click **Advanced Controller Properties** to configure advanced controller settings.

More information

Advanced controller management

Advanced controller properties

Advanced controller management

The **Advanced Controller Management** screen allows you to reset the controller to factory defaults or repurpose mode, manage persistent events and snapdumps, adjust controller personality and function settings. It also provides the option to configure link-speed interfaces, spin up drives, and control other maintenance-oriented behaviors.

Subtopics

Setting factory defaults

Setting factory repurpose

Saving persistent events

Clearing persistent events

Enabling security

Disabling drive security

Changing security settings

Spinning up multiple drives

Managing Snapdump

Managing SAS storage link speed

Managing PCIe storage interface
Managing controller personality
Managing auto-assign policy
Managing function profile
Managing controller function
Managing advanced software options

Setting factory defaults

About this task

Setting factory defaults restores all controller properties to their factory values. This action also erases log information and trace buffers.

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Controller Management > Advanced Controller Management**.
3. In the **Advanced Controller Management** screen, select **Set Factory Defaults**.

The warning message appears to request confirmation.

4. Select the checkbox for the **Reset Profile** option.
5. Select the checkbox for the **Confirm** option, and click **Yes**.

Setting factory repurpose

About this task

Setting factory repurpose resets the controller to its factory state by removing all data and settings.

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Controller Management > Advanced Controller Management**.
3. In the **Advanced Controller Management** screen, select **Set Factory Repurpose**.

The warning message appears to request confirmation.

4. Select the checkbox for the **Confirm** option, and click **Yes**.

Saving persistent events

About this task

Saves a copy of the persistent events to a file using the specified file name and path.

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Controller Management > Advanced Controller Management**.
3. In the **Advanced Controller Management** screen, select **Save Persistent Events**.
4. In the **Save Persistent Events** screen, select the required settings.
5. Click **Save Event**.
6. Select the checkbox for the **Confirm** option, and click **Yes**.
7. Click **OK**.

Clearing persistent events

About this task

Deletes all the existing persistent events.

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Controller Management > Advanced Controller Management**.
3. In the **Advanced Controller Management** screen, select **Clear Persistent Events**.

The Success screen appears to notify that the operation is complete.

4. Click **OK**.

Enabling security

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Controller Management > Advanced Controller Management**.
3. In the **Advanced Controller Management** screen, select **Enable Security**.
4. For the **Key Management Mode**, select either **External Key Management (EKM)** or **Local Key Management (LKM)**.
5. Click **OK**.
6. (Optional) If you have selected **Local Key Management (LKM)**, do the following:
 - a. Enter the required data on the screen. You are allowed to set the **Security Key Identifier**, **Security Key**, and **Passphrase at Boot Time**.



TIP

If you set **Passphrase at Boot Time**, the system displays a message during boot indicating that configuration is required. For more information, see, [Controller boot process in UEFI](#), and [Resolving controller boot blocking events in UEFI](#).

- b. Select the **I Recorded the Security Settings for Future Reference** check box.
 - c. Click **Enable Drive Security**.
7. (Optional) If you have selected **External Key Management (EKM)**, do the following:
 - a. Make sure that the iLO key manager is properly configured only then you can select **External Key Management (EKM)**. For more information, see [HPE iLO User Guide](#).
 - b. Click **Enable Drive Security**.
 - c. Reboot the system.

The controller retrieves the SED drive keys from the remote HPE Enterprise Secure Key Manager server only during power on. It communicates to HPE iLO through the BIOS.

More information

Self-encrypting drive

Disabling drive security

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Controller Management > Advanced Controller Management**.
3. In the **Advanced Controller Management** screen, select **Disable Security**.

The warning message appears requesting confirmation.

4. Review the warning message, click **Confirm** and **Yes** to proceed.

More information

Self-encrypting drive

Changing security settings

About this task

When the Local Key Management (LKM) mode is set, changes to the security settings are applied immediately.

When External Key Management (EKM) mode is set, changes to the security settings are applied when you reboot the system.

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Controller Management > Advanced Controller Management**.
3. In the **Advanced Controller Management** screen, select **Change Security Settings**.
4. Change the required fields. You are allowed to change the **Security Key Identifier**, **Security Key**, and **Passphrase at Boot Time**.



TIP

If you set **Passphrase at Boot Time**, the system displays a message during boot indicating that configuration is required. For more information, see, [Controller boot process in UEFI](#), and [Resolving controller boot blocking events in UEFI](#).

5. Click **Save Security Settings**.

More information

[Self-encrypting drive](#)

Spinning up multiple drives

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Controller Management > Advanced Controller Management**.
3. In the **Advanced Controller Management** screen, select **Start/Spin-Up Drives**.
4. From the drives list, select the checkbox for the drives you want or select the **Check All** checkbox to select all the drives.
5. Click **Go**.
6. Review the warning message, select the **Confirm** checkbox, and click **Yes** to proceed.
7. Click **OK**.

Managing Snapdump

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Controller Management > Advanced Controller Management**.
3. In the **Advanced Controller Management** screen, select **Manage Snapdump**.
4. In the **Manage Snapdump** screen, set the required values.



TIP

If you select the checkbox for **Retrieve Extended TTY Logs**, the **Select Controller Function** option is not shown.

5. (Optional) If you click **Save new On-Demand Snapdump**, click **OK** on the success screen to return to the **Manage Snapdump** screen.
6. (Optional) If you click **Clear All Snapdumps**:
 - a. Select the checkbox for the **Confirm** option, and click **Yes**.
 - b. Click **OK**.

More information

Diagnostic tools

Managing SAS storage link speed

Procedure

1. From the **System Utilities** screen, select **System Configuration** > **HPE MRXXX**.
2. Select **Main Menu** > **Controller Management** > **Advanced Controller Management**.
3. In the **Advanced Controller Management** screen, select **Manage SAS Storage Link Speed**.
4. For the ports, select the required speed for each port from the drop-down list.
5. Click **Apply Changes**.
6. Select the **Confirm** checkbox, and click **Yes** to proceed.

More information

Manage SAS storage link speed

Managing PCIe storage interface

Procedure

1. From the **System Utilities** screen, select **System Configuration** > **HPE MRXXX**.
2. Select **Main Menu** > **Controller Management** > **Advanced Controller Management**.
3. In the **Advanced Controller Management** screen, select **Manage PCIe Storage Interface**.

4. (Optional) In the **VIEW THE PHY INFORMATION** section, view the current negotiation speed for the ports.
 - a. Click **View/Change Link Speed**.
 - b. In the **View/Change Link Speed** screen, select the required speed for each port from the drop-down list.
 - c. Click **Apply Changes**.
 - d. Click **OK**.
5. (Optional) In the **MANAGE PCIe WIDTH** section, view the width of the port.
 - a. For the ports, select the link width for each port from the drop-down list.
 - b. Click **Apply Changes**.
 - c. Select the **Confirm** checkbox, and click **Yes** to proceed.
 - d. Click **OK**.

More information

Manage PCIe storage interface

Managing controller personality

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Controller Management > Advanced Controller Management**.
3. In the **Advanced Controller Management** screen, select **Manage Controller Personality**.
4. In the **Manage Controller Personality** screen, select the required settings.
5. Click **Apply Changes**.
6. Select the **Confirm** checkbox, and click **Yes** to proceed.

The **Success** screen appears to notify that the operation is complete.

7. (Optional) In the **CONFIGURE NOW** section, select an option for **Auto-Configure Behavior (Execute Once)** to immediately configure eligible drives as per the selected configuration.
 - a. Select the **Ignore Auto-Assign Policy** checkbox to override the existing auto-assign policy.

- b. Click **Apply Now**.
- c. Select the **Confirm** checkbox, and click **Yes** to proceed.

More information

Managing auto-assign policy

Managing auto-assign policy

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Controller Management > Advanced Controller Management**.
3. In the **Advanced Controller Management** screen, select **Manage Auto-Assign Policy**.

The current and default policies are displayed in **Current Auto-Assign Policy** and **Default Auto-Assign Policy**.

4. (Optional) Select the **Change Auto-Assign Policy** checkbox to make change the existing policy.
 - a. From the drop-down list, select an option for **Auto-Assign Policy**.
 - b. Click **Apply Changes**.
 - c. Click **OK**.
5. (Optional) Select the **Auto-Assign Policy (Execute Once)** checkbox to immediately update the policy.
 - a. From the drop-down list, select an option for **Auto-Assign Policy**.
 - b. Click **Apply Changes**.
 - c. Click **OK**.

More information

Creating volume

Converting multiple drives to JBOD

Managing controller personality

Converting a drive to JBOD

Managing function profile

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Controller Management > Advanced Controller Management**.
3. In the **Advanced Controller Management** screen, select **Manage Function Profile**.
4. Select an option for **Choose Profile ID** from the drop-down list.
5. Assess all the details of the profile.

More information

Multiple physical functions

Managing controller function

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Controller Management > Advanced Controller Management**.
3. In the **Advanced Controller Management** screen, select **Manage Controller Function**.
4. In the **Manage Controller Function** screen, select a function from the list.
5. Assess all the details that are displayed.
6. (Optional) Click **Select First Device** to choose a device that has to be reported as the first device by the controller function.
 - a. Assess the warning, and click **Yes** to proceed.

The **Success** screen appears to notify that the operation is complete.
7. (Optional) Click **Map JBODs** to associate JBODs to the controller function.
 - a. Select the checkbox for the JBOD you want to disassociate in the list.
 - b. Click **OK**.
 - c. Click **OK**.
8. (Optional) Click **Remove Mapped JBODs** to disassociate JBODs from the controller function.

- a. Select the checkbox for the JBOD you want to disassociate in the list.
- b. Click **OK**.
- c. Click **OK**.

More information

Multiple physical functions

Managing advanced software options

About this task

This page displays all activated Advanced Software Options available on the controller. These options allow you to enable and configure advanced MegaRAID features.

To use any advanced feature, you must first activate the corresponding software activation key. The MegaRAID Advanced Software Options section is shown only when the controller supports MegaRAID software licensing. If the controller does not support licensing, this section will not appear.

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Controller Management > Advanced Controller Management**.
3. In the **Advanced Controller Management** screen, select **Manage Advanced Software Options**.

Results

All the activated advanced software options are displayed.

Advanced controller properties

The **Advanced Controller Properties** screen provides the settings to manage cache and memory usage, schedule patrol read operations, configure spare-drive behavior, tune background task rates, and enable external key management for secure SED operations.

Subtopics

Configuring cache and memory settings

Configuring patrol read settings

Configuring spare settings

Configuring task rates

More information

Enabling security

Configuring cache and memory settings

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Controller Management > Advanced Controller Properties**.
3. In the **Advanced Controller Properties** screen, select **Cache and Memory**.
4. In the **Cache and Memory** screen, update the cache flush interval as required.

The cache flush interval is the interval (in seconds) at which the contents of the onboard data cache are flushed.

The remaining settings on this screen are not configurable in this release.

More information

Cache

Configuring patrol read settings

About this task

The patrol read operation scans and resolves potential problems on configured drives.

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Controller Management > Advanced Controller Properties**.
3. In the **Advanced Controller Properties** screen, select **Patrol Read**.
4. In the **Patrol Read** screen, start, suspend, resume, or stop the patrol read operation.

View state and iteration information in the **State** and **Iterations Completed** fields.

5. Configure the **Rate**.

Rate is the percentage of system resources dedicated to perform a patrol read operation on configured drives.

6. Disable or enable the **Corrects Unconfigured Areas**.
7. Click **Apply Changes**.

More information

Patrol read

Configuring spare settings

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Controller Management > Advanced Controller Properties**.
3. In the **Advanced Controller Properties** screen, select **Spare**.
4. In the **Spare** screen, do the following:
 - **Emergency Spare:** Select to commission global spare drives, unconfigured good drives, or both as emergency hot spare drives.
 - **Emergency Spare for SMARTer:** Enable or disable commissioning and emergency hot spare drives for PFA events.
 - **Replace Drive** - Enable or disable to copy data from a hot spare drive to a physical drive.
 - **Replace Drive on SSD SMART Error** - Enable or disable to start a Replace operation if a Self-Monitoring Analysis and Report Technology (SMART) error is detected on a drive.
5. Click **Apply Changes**.

More information

Drive rebuild

Spare drives

Dedicated spare

Global spare

Configuring task rates

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Controller Management > Advanced Controller Properties**.
3. In the **Advanced Controller Properties** screen, select **Task Rates**.
4. In the **Task Rates** screen, configure the following task rates:
 - **Background initialization (BGI) Rate:** The percentage of system resources dedicated to perform a background initialization on a redundant volume.
 - **Patrol Read Rate:** The percentage of system resources dedicated to perform a Patrol Read operation on configured drives.
 - **Consistency Check Rate:** The percentage of system resources dedicated to perform a consistency check operation on a redundant volume.
 - **Transformation:** The percentage of system resources dedicated to perform a RAID level migration or online capacity expansion on a volume.
 - **Rebuild Operating Mode Priority:** Select **Rebuild** or **Host I/O**.

Rebuild mode gives priority to rebuilding RAID data, completing the rebuild faster, and slowing host I/O.

Host I/O mode gives priority to host traffic, so host and flush I/Os continue without being limited, and the rebuild runs at a lower priority.
5. Click **Apply Changes**.

More information

[Background initialization](#)

[Initialization operations](#)

[Consistency check](#)

[Array transformations](#)

[Volume transformations](#)

[Drive rebuild](#)

Volume management

Volume Management refers to the process of creating, configuring, and managing logical volumes on physical storage devices.

Subtopics

Viewing and configuring volume properties

Deleting a volume

Expanding a volume

Initializing a volume

Locating a volume

Erasing a volume

Mapping a volume to a function

Securing a volume

Removing a volume from function

Locating a drive associated with a volume

Making a drive associated with a volume offline

Checking the consistency of a volume

Viewing and configuring volume properties

Procedure

1. Select **Main Menu > Volume Management**.
2. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
3. In the **Volume Management** screen, select the volume.
4. (Optional) Click **View Associated Drives** to see the properties of the drives associated with the volume.
 - a. Select the drives and click **View Drive Properties**.
5. Click **Advanced** to configure volume and volume drive properties.
6. Set the required values.
7. Click **Apply Changes**.
8. Click **OK**.

More information

Viewing drive properties

Deleting a volume

Procedure

1. Select **Main Menu > Volume Management**.
2. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
3. In the **Volume Management** screen, select the volume.
4. In the **Volume** screen, select **Delete Volume** from the **Operation** drop-down list.
5. Click **Go**.
6. Review the warning message, and click **Confirm** and **Yes** to proceed.

Expanding a volume

Procedure

1. Select **Main Menu > Volume Management**.
2. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
3. In the **Volume Management** screen, select the volume.
4. In the **Volume** screen, select **Expand Volume** from the **Operation** drop-down list.
5. Click **Go**.
6. In the **Expand Volume** screen, click **Add Drives**.
7. Select the checkbox for the drives you want to add or click **Select All** to add all the drives.
8. Click **OK**.
9. On the success screen, click **OK**.

Initializing a volume

Procedure

1. Select **Main Menu > Volume Management**.
2. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
3. In the **Volume Management** screen, select the volume.
4. In the **Volume** screen, select **Fast Initialization** or **Slow Initialization** from the **Operation** drop-down list.
5. Click **Go**.
6. Review the warning message, and select the checkbox for **Confirm**.
7. Click **Yes** to proceed.
8. Click **OK**.

More information

Initialization operations

Full initialization

Locating a volume

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Volume Management**.
3. In the **Volume Management** screen, select the volume.
4. In the **Volume** screen, select **Start Locate** from the **Operation** drop-down list.
5. Click **Go**.

The LEDs for the drive associated with the volume begin flashing.

- a. (Optional) To stop locating the drive, select **Stop Locate** from the **Operation** drop-down list.

The LEDs stop flashing.

6. Click **OK**.

Erasing a volume

Procedure

1. Select **Main Menu > Volume Management**.
2. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
3. In the **Volume Management** screen, select the volume.
4. In the **Volume** screen, select **Volume Erase** from the **Operation** drop-down list.
5. Select one of the following options from the **Erase Mode** drop-down list:
 - **Simple** - Writes a pattern to the volume in a single pass.
 - **Normal** - A three-pass operation that first overwrites the drive contents with random values then overwrites it twice with patterns.
 - **Thorough** - Repeats the Normal drive erase operation three times.
6. If you want to delete the volume after the erase operation is completed, select the checkbox for **Delete After Erase**.
7. Click **Go**.
8. Review the warning message, and select the checkbox for **Confirm**.
9. Click **Yes** to proceed.
10. Click **OK**.

Mapping a volume to a function

Procedure

1. Select **Main Menu > Volume Management**.
2. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
3. In the **Volume Management** screen, select the volume.
4. In the **Volume** screen, select **Map to Function** from the **Operation** drop-down list.

5. Set a value for **Select Controller Function** drop-down list.
6. Click **Go**.
7. Click **OK**.

More information

Multiple physical functions

Securing a volume

About this task

The secure volume operation enables security on the volume. Deleting the volume is the only way to disable this security.

Procedure

1. Select **Main Menu > Volume Management**.
2. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
3. In the **Volume Management** screen, select the volume.
4. In the **Volume** screen, select **Secure Volume** from the **Operation** drop-down list.
5. Click **Go**.
6. Review the warning message, and select the checkbox for **Confirm**.
7. Click **Yes** to proceed.
8. Click **OK**.

Removing a volume from function

Procedure

1. Select **Main Menu > Volume Management**.
2. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
3. In the **Volume Management** screen, select the volume.

4. In the **Volume** screen, select **Remove from function** from the **Operation** drop-down list.
5. Click **Go**.
6. Click **OK**.

More information

Multiple physical functions

Locating a drive associated with a volume

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Volume Management**.
3. In the **Volume Management** screen, select the volume.
4. In the **Volume** screen, click **View Associated Drives**.
5. Select the checkbox for the drive you want.
6. Click **View Drive Properties**.
7. From the **Operation** drop-down list, select **Start Locate**.
8. Click **Go**.



TIP

The LEDs for the drive associated with the volume begin flashing.

- a. Optional. To stop locating the drive, select **Stop Locate** from the **Operation** drop-down list.



TIP

The LEDs stop flashing.

9. Click **OK**.

Making a drive associated with a volume offline

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Volume Management**.
3. In the **Volume Management** screen, select the volume.
4. In the **Volume** screen, click **View Associated Drives**.
5. Select the checkbox for the drive you want.
6. Click **View Drive Properties**.
7. From the **Operation** drop-down list, select **Make Offline**.
8. Click **Go**.



TIP

The associated drive status shows offline.

- a. Optional. To move the drive back to online status, select **Force Online** from the **Operation** drop-down list.



TIP

The associated drive status shows online.

- b. Optional. To move the drive to missing status, select **Make Missing** from the **Operation** drop-down list.



TIP

The associated drive status shows unconfigured good.

9. Click **OK**.

Checking the consistency of a volume

Procedure

1. Select **Main Menu > Volume Management**.
2. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
3. In the **Volume Management** screen, select the volume.
4. In the **Volume** screen, select **Check Consistency** from the **Operation** drop-down list.
5. Click **Go**.
6. Click **OK**.

The progress indicator displays the percentage of completion for the consistency check.

7. (Optional) To stop or suspend the consistency check, click **Stop**, or **Suspend**.
 - a. To resume the suspended consistency check, click **Resume**.

More information

Consistency check

Device management

This screen allows you to view enclosures and the drives attached to it and perform operations like locating the drives. You can also view more properties of the drive using the **Advanced** option.

Subtopics

Viewing logical enclosure properties

Viewing drive properties

Locating a physical drive associated with a logical enclosure

Initializing or clearing a drive

Making a drive offline

Replacing a drive

Converting a drive to JBOD

Converting a drive to unconfigured

Assigning a global spare

Unassigning a global spare

Preparing a drive for removal

Erasing a drive cryptographically

Sanitizing an unconfigured good drive

Reconfiguring a drive

Viewing logical unit information of a drive

Viewing logical unit information of a drive

Viewing logical enclosure properties

Procedure

1. From the **System Utilities** screen, select **System Configuration** > **HPE MRXXX**.
2. Select **Main Menu** > **Device Management**.
3. Select the required logical enclosure from the list.

All the properties of the logical enclosure are displayed.

4. (Optional) Click **Next Logical Enclosure** to see the properties of the next enclosure.

Viewing drive properties

Procedure

1. From the **System Utilities** screen, select **System Configuration** > **HPE MRXXX**.
2. Select **Main Menu** > **Device Management**.
3. Select the required logical enclosure from the list.
4. Scroll to the **AVAILABLE BAYS** section, and select the required drive.

All the properties of the drive are displayed.

Locating a physical drive associated with a logical enclosure

Procedure

1. From the **System Utilities** screen, select **System Configuration** > **HPE MRXXX**.
2. Select **Main Menu** > **Device Management**.

3. Select the required logical enclosure from the list.
4. Scroll to the **AVAILABLE BAYS** section, and select the required drive.
5. From the **Operation** drop-down list, select **Start Locate**.
6. Click **Go**.
 - a. (Optional) To stop locating the drive, select **Stop Locate** from the **Operation** drop-down list.
7. Click **OK**.

Initializing or clearing a drive

Prerequisites

Ensure that you have backed up data before changing the state of the drive.

About this task



NOTE

Initializing a drive will delete any data on the drive.

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Device Management**.
3. Select the required logical enclosure from the list.
4. Scroll to the **AVAILABLE BAYS** section, and select the required drive.

All the properties are displayed.
5. From the **Operation** drop-down list, select **Clear**.
6. Click **Go**.
7. Review the message, select the **Confirm** checkbox, and click **Yes** to proceed.
8. Click **OK**.

Making a drive offline

Prerequisites

Ensure you have backed up data before changing the state of the drive.

Procedure

1. Select **Main Menu > Device Management**.
2. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
3. Select the required logical enclosure from the list.
4. Scroll to the **AVAILABLE BAYS** section, and select the required drive.
5. From the **Operation** drop-down list, select **Making Offline**.
6. Click **Go**.
7. Review the message, select the **Confirm** checkbox, and click **Yes** to proceed.
8. Click **OK**.

Replacing a drive

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Device Management**.
3. Select the required logical enclosure from the list.
4. Scroll to the **AVAILABLE BAYS** section, and select the required drive.

All the properties are displayed.
5. From the **Operation** drop-down list, select **Replace Drive**.
6. Click **Go**.
7. Select a drive from the **Select Replacement Drive** drop-down list.
8. Click **Replace Drive**.
9. Click **OK**.

More information

Replace drive

Converting a drive to JBOD

Prerequisites

Ensure that you have backed up data before changing the state of the drive.

About this task

You can convert the state of a physical drive from an unconfigured state to JBOD mode. In JBOD kmode, each disk is presented individually to the host operating system without RAID features such as striping, mirroring, or parity.



WARNING

Converting a drive to JBOD will permanently delete any data on the drive.

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Device Management**.
3. Select the required logical enclosure from the list.
4. Scroll to the **AVAILABLE BAYS** section, and select the required drive.
5. From the **Operation** drop-down list, select **Convert to JBOD**.
6. Select **Main Menu > Configuration Management > Convert to JBOD**.
7. (Optional) In the **Convert to JBOD** screen, select the checkbox to **Ignore Auto-Assign Policy**.
8. Click **OK**.
9. Review and address any warning messages, as required.
10. Click **OK**.

More information

Mixed mode - RAID and JBOD simultaneously
Managing auto-assign policy

Converting a drive to unconfigured

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Device Management**.
3. Select the required logical enclosure from the list.
4. Scroll to the **AVAILABLE BAYS** section, and select the required bay.
5. From the **Operation** drop-down list, select **Convert to Unconfigured**.
6. Click **Go**.
7. Review the message, select the **Confirm** checkbox, and click **Yes** to proceed.
8. Click **OK**.

More information

Mixed mode - RAID and JBOD simultaneously

Assigning a global spare

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Device Management**.
3. Select the required logical enclosure from the list.
4. Scroll to the **AVAILABLE BAYS** section, and select the required drive.
5. From the **Operation** drop-down list, select **Assign Global Spare**.
6. Click **Go**.
7. Click **OK**.

More information

Unassigning a global spare

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Device Management**.
3. Select the required logical enclosure from the list.
4. Scroll to the **AVAILABLE BAYS** section, and select the required drive.
5. From the **Operation** drop-down list, select **Unassign Spare**.
6. Click **Go**.
7. Review the message, select the **Confirm** checkbox, and click **Yes** to proceed.
8. Click **OK**.
9. On the success screen, click **OK**.

More information

Global spare

Preparing a drive for removal

About this task

Prepare for removal feature safely shuts down the unconfigured good drive so you can unplug it without causing hardware damage or system errors. If you do not remove the drive, you can undo the operation to reconnect the drive and return it to a ready state.

Procedure

1. Select **Main Menu > Device Management**.
2. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
3. Select the required logical enclosure from the list.
4. Scroll to the **AVAILABLE BAYS** section, and select the required drive.

5. From the **Operation** drop-down list, select **Prepare for removal**.
6. Click **Go**.
7. Click **OK**.
8. (Optional) To reverse the action, from the **Operation** drop-down list, select **Undo Prepare for removal**, and repeat steps 6 and 7.

Erasing a drive cryptographically

Prerequisites

Ensure that you have backed up data before changing the state of the drive.

Procedure

1. Select **Main Menu** > **Device Management**.
2. From the **System Utilities** screen, select **System Configuration** > **HPE MRXXX**.
3. Scroll to the **AVAILABLE BAYS** section, and select the required drive.
4. From the **Operation** drop-down list, select **Cryptographic Erase**.
5. Select the **Uninterrupted Erase** checkbox.
6. Click **Go**.
7. Review the warning message, select the **Confirm** checkbox, and **Yes** to proceed.
8. Click **OK**.

More information

Drive erase

Sanitizing an unconfigured good drive

Prerequisites

- Back up of the data before changing the state of the drive.
- Convert the drive to an unconfigured good drive. If not done, you are not allowed to sanitize the drive.

About this task



WARNING

During a sanitize operation, any data on the drive is permanently lost.

Procedure

1. Select **Main Menu > Device Management**.
2. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
3. Select the required logical enclosure from the list.
4. Scroll to the **AVAILABLE BAYS** section, and select the required drive.
5. From the **Operation** drop-down list, select **Sanitize**.
6. Select a value from the **Erase Mode** drop-down list.
7. Select the **Uninterrupted Erase** checkbox.
8. Click **Go**.
9. Review the message, select the **Confirm** checkbox, and click **Yes** to proceed.
10. Click **OK**.

More information

Sanitize

Reconfiguring a drive

Prerequisites

Ensure that you have backed up data before changing the state of the drive.

Procedure

1. Select **Main Menu > Device Management**.
2. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
3. Scroll to the **AVAILABLE BAYS** section, and select the required drive.
4. From the **Operation** drop-down list, select **Reconfigure**.

5. From the **Usage** drop-down list, select **For RAID** or **For JBOD**.
6. From the **Options** drop-down list, select a value.
7. Click **Go**.
8. Review the warning message, select the **Confirm** checkbox, and **Yes** to proceed.
9. Click **OK**.

More information

Mixed mode - RAID and JBOD simultaneously

Viewing logical unit information of a drive

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Device Management**.
3. Select the required logical enclosure from the list.
4. Scroll to the **AVAILABLE BAYS** section, and select the required bay.
5. Scroll down and click **Logical Unit Information**.

All the properties of the logical unit are displayed.

Viewing logical unit information of a drive

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Device Management**.
3. Select the required logical enclosure from the list.
4. Scroll to the **AVAILABLE BAYS** section, and select the required bay.
5. Scroll down and click **Advanced**.

All the advanced properties of the drive are displayed.

Energy pack management

Energy Pack Management focuses on monitoring and maintaining backup power units, such as batteries or super capacitors, that safeguard cached data during power outages. It includes health checks, charge status monitoring, and replacement alerts to ensure data integrity and system reliability.

Subtopics

Viewing the energy pack

Viewing the energy pack

Procedure

1. From the **System Utilities** screen, select **System Configuration > HPE MRXXX**.
2. Select **Main Menu > Energy Pack Management**.

The properties of the energy pack are displayed.

Maintenance

Provides information for maintaining the controller, including updates, diagnostics, and error reporting.

This section provides the insights on maintaining the controller.

Subtopics

Updating software and firmware

Error reporting

Diagnostic tools

Updating software and firmware

Server and controller firmware must be updated before using the controller for the first time. To update system software and firmware, download the Service pack for HPE ProLiant (SPP) from the Hewlett Packard Enterprise website: <https://www.hpe.com/servers/spp/download>.

For information about the SPP, see the SPP product documentation here: <https://hpe.com/support/servicepackforhpeproliant-quicklinks>.

Hewlett Packard Enterprise distributes drivers and other support software for servers and server blades through the SPP, which you can download from here: <https://www.hpe.com/servers/spp/download>.

Always use the latest SPP version that supports your server or server blade.

After updating HPE MR Storage Administrator, clear the browser cache.

Error reporting

Integrated Management Log

The controller reports diagnostic error messages (POST messages) during boot. It logs these messages to the UEFI Health Log and also the Integrated Management Log (IML) within HPE iLO. Many POST messages suggest corrective actions. For more information about POST messages, see **Integrated Management Log Messages and Troubleshooting Guide for HPE ProLiant GenXXX servers and HPE Synergy**.

SNMP traps

The controller supports SNMP traps documented in the `cpqida.mib` and `cpqstsys.mib` MIBs. SNMP traps are sent as part of the HPE iLO SNMP management function. For information on configuring HPE iLO SNMP traps and a full description of supported SNMP traps, see the **HPE Storage Controllers: Management overview**.

Redfish alerts

The controller supports sending alerts through the HPE iLO Redfish API. For information on configuring Redfish alerts, see the **HPE Storage Controllers: Management overview**.

Application event log

HPE MR Storage Administrator reports array events with detailed diagnostic information to the system event logs on each platform:

- **Windows:** Records events in the application area of the Microsoft Windows system event Log.
- **Linux:** Records events to the system event log that is located at `/var/log/messages`.
- **VMware:** Records events to the system event log that is located at `/var/log/vmkernel.log`.

System event log

The controller uses PLDM over Redfish device enablement to send Redfish events, which generate error messages in the HPE IML. The device firmware generates these events that may occur during power-up or run-time. When the HPE Agentless Management Service (AMS) is installed, the HPE IML events are also stored in the Windows System Event Log and Linux `/var/log/messages`, and ESXi `/var/log/syslog.log`.

More information

Redfish

Diagnostic tools

To troubleshoot array problems and generate feedback about arrays, use the following diagnostic tools:

- HPE MR Storage Administrator
- HPE StorCLI2

HPE MR Storage Administrator

HPE MR Storage Administrator displays event log messages and system messages. To view the event log, in the controller dashboard click **View Event Log** under the **Actions** menu. In the **Actions** menu for the event log, you can download the log file and clear the log file. To view the system messages, click the bell icon, and the messages appear on the top of the window.

HPE MR Storage Administrator records consolidated information about the server and all the devices to which it is connected. To download the support log, in the **Server Dashboard** click **Download Support Log**.

HPE StorCLI2

If you cannot use HPE MR Storage Administrator, use HPE StorCLI2 to collect the required logs to analyze issues.

Run `storcli2 show` to get controller index first, then run the following commands, where `x` represents the controller index.

Command	Description
<code>storcli2 /cx show events file=events.log</code>	Collects all events
<code>storcli2 /cx show events type=sincereboot file=SBevents.log</code>	Collects events since the last reboot
<code>storcli2 /cx show all > show_all.log</code>	Collects general device information

Command	Description
<code>storcli2 /cx show termlog file=show_termlog.log</code>	Collects firmware logs
<code>storcli2 /cx get snapdump ondemand force</code>	Generates an on-demand snapdump
<code>storcli2 /cx get snapdump all</code>	Captures all available snapdump data

Snapdump trigger types

The Snapdump debugging feature automatically collects all controller debug information when the controller faults. Snapdump is an automated way for the controller to capture diagnostic data when certain types of failures occur.

The system supports several Snapdump trigger types:

ONDEMAND

The controller generates Snapdump data at runtime, collecting basic debug information. Snapdumps can only be generated once every 10 minutes because the process takes time and may cause I/O timeouts. Enhanced Snapdump lifts this restriction.

- Generates Snapdump data manually during runtime.
- Collects only basic Snapdump information.
- Allows Snapdump generation only after a 10-minute interval, because the operation is time-consuming and may cause I/O timeouts.
- Removes the 10-minute interval requirement when using enhanced Snapdump.

RUNTIME

- Generates Snapdump data for failures not caused by a firmware fault.
- Covers events such as foreign import failures, pinned-cache generation failures, or failures triggered through the debug interface.
- Allows Snapdump generation only after a 10-minute interval, similar to ONDEMAND.

MPI3.0 Diagnostic Fault Reset

- Triggered automatically by the driver when the controller becomes unresponsive.
- Captures diagnostic data around an MPI3.0 Diagnostic Fault Reset event.

Controller Fault

- Triggered when the controller firmware has an invalid condition and faults.
- Provides detailed data to help diagnose firmware errors.

More information

[HPE StorCLI2](#)

[HPE MR Storage Administrator](#)

[Managing Snapdump](#)

Storage reference

Provides reference information for storage terminology, capacity conventions, and RAID conventions.

Subtopics

[Memory and storage capacity conventions](#)

[RAID conventions](#)

Memory and storage capacity conventions

Memory capacities are specified using binary prefixes:

- KiB = 2^{10} bytes
- MiB = 2^{20} bytes
- GiB = 2^{30} bytes
- TiB = 2^{40} bytes

Storage capacities are specified using SI prefixes:

- KB = 10^3 bytes
- MB = 10^6 bytes
- GB = 10^9 bytes
- TB = 10^{12} bytes

Older and other documentation might use SI prefixes for binary values.

Actual available memory capacity and actual formatted storage capacity for devices are less than specified values.

RAID conventions

Hewlett Packard Enterprise uses the following naming convention for RAID levels:

- RAID 0
- RAID 1
- RAID 10
- RAID 5
- RAID 50
- RAID 6
- RAID 60

RAID 50 and RAID 60 are also known in the industry as RAID 5+0 and RAID 6+0, respectively.

HPE Storage Controller documentation resources

Provides references to official documentation resources for HPE storage controllers.

General websites

Websites	Links
Hewlett Packard Enterprise Support Center	https://support.hpe.com
Hewlett Packard Enterprise Worldwide	https://www.hpe.com/assistance
Subscription Service/Support Alerts	https://www.hpe.com/support/e-updates
Other websites	Support and other resources

User guides for hardware RAID controllers, boot devices, and virtual RAID

For the list of guides and links, see [HPE Storage Controller documentation quick links](#).

User guides for configuration tools

For the list of guides and links, see [HPE Compute Software and Firmware Product Documentation Quick Links](#).

QuickSpecs for hardware RAID controllers

Content	Microchip® SmartR OC 3100	Microchip® SmartR OC 3200	Broadcom® SAS3900	Broadcom® SAS5132W
QuickSpecs - Gen10 and Gen10 Plus	https://www.hpe.com/psnow/doc/a00047736enw	https://www.hpe.com/psnow/doc/a50002562enw	https://www.hpe.com/psnow/doc/a50002563enw	Not applicable
QuickSpecs - Gen11	Not applicable	https://www.hpe.com/psnow/doc/a50004312enw	https://www.hpe.com/psnow/doc/a50004311enw	Not applicable
QuickSpecs - MR93 2i-p	Not applicable	Not applicable	Not applicable	https://www.hpe.com/psnow/doc/a50004311enw

QuickSpecs for Boot devices and Virtual RAID

Content	HPE NS204 boot devices	Intel VROC
QuickSpecs - Gen10 Plus and Gen11	https://www.hpe.com/psnow/doc/a00094638enw	https://www.hpe.com/psnow/doc/a50002570enw

Technical papers

Content	Links
Encryption Overview	https://www.hpe.com/info/SCEO
Management Overview	https://www.hpe.com/info/SCMO

Training videos

Content	Links
Automate the deployment of the HPE MRXXX Storage Controllers at scale	https://hpedemoportal.ext.hpe.com/
Automate the deployment of the HPE MRXXX Storage Controllers using the DMTF RDE Redfish APIs	https://hpedemoportal.ext.hpe.com/
Automate the deployment of the HPE MRXXX Storage Controllers using the HPE Deployment Automation solution with Ansible	https://hpedemoportal.ext.hpe.com/

Content	Links
Boot Windows OS from Intel VROC RAID volume	https://hpedemoportal.ext.hpe.com/
How to Configure HPE MRXXX Storage Controllers using the MR Storage Administrator (MRSA) GUI	https://hpedemoportal.ext.hpe.com/
How to Configure HPE MRXXX Storage Controllers using StorCLI	https://hpedemoportal.ext.hpe.com/
How to Configure HPE MRXXX Storage Controllers using HPE UEFI/BIOS	https://hpedemoportal.ext.hpe.com/
Install and boot VMWare ESXi from Intel VROC RAID 1 Volume	https://hpedemoportal.ext.hpe.com/
The top key features of the 3rd Gen Intel Xeon Scalable Processors on HPE ProLiant DL380 Gen10 Plus Server Demo 1: Intel Virtual RAID on CPU (Intel VROC) usage and benefits	https://hpedemoportal.ext.hpe.com/

HPE Storage Controller documentation quick links

Provides quick access links to commonly used HPE storage controller documentation.

This table is a comprehensive list of all the documents for HPE storage controllers and boot devices. It provides quick links to all the versions of each document.

Table 1. HPE SR Controllers

Documents	Versions		
User Guide	Gen10	Gen10 Plus	Gen11
GUI User Guide	Multi-gen	Multi-gen	Multi-gen
CLI User Guide	Multi-gen	Multi-gen	Multi-gen
Alerts - Customer Advisories	—	Gen10 Plus	Gen11

Table 2. HPE MR Controllers

Documents	Versions			
User Guide	Gen10	Gen10 Plus	Gen11	MR932i-p

Documents	Versions			
GUI User Guide	<u>Multi-gen</u>	<u>Multi-gen</u>	<u>Multi-gen</u>	<u>Multi-gen</u>
CLI User Guide	<u>Multi-gen</u>	<u>Multi-gen</u>	<u>Multi-gen</u>	<u>HPE StorCLI2</u>
Alerts - Customer Advisories	<u>Gen10</u>	<u>Gen10 Plus</u>	<u>Gen11</u>	<u>MR932i-p</u>

Table 3. Intel Virtual Raid on CPU for HPE

Documents	Versions	
User Guide	<u>Gen10 Plus - Linux</u>	<u>Multi-gen</u>
	<u>Gen10 Plus - VMware</u>	
	<u>Gen10 Plus - Windows</u>	
Alerts - Customer Advisories	<u>Multi-gen</u>	<u>Multi-gen</u>

Table 4. HPE NS204i Boot Devices

Documents	Versions
User Guide	<u>Multi-gen</u>
Alerts - Customer Advisories	<u>Multi-gen</u>