



Hewlett Packard
Enterprise

HPE MR Gen11 Controller User Guide

Part Number: 30-07388CA7-003
Published: January 2024
Edition: 3

HPE MR Gen11 Controller User Guide

Abstract

This document includes feature, installation, and configuration information about HPE MR Gen11 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: 30-07388CA7-003

Published: January 2024

Edition: 3

© Copyright 2022–2024 Hewlett Packard Enterprise Development LP

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

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

Acknowledgments

MegaRAID™ is the registered trademark of Broadcom, Inc.

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.

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

Table of contents

- HPE MR Gen11 Controller
 - 200 series
 - 400 series
- Features
 - Controller supported features
 - Operating environments
 - RAID technologies
 - Transformation
 - Drive technology
 - Security
 - Reliability
 - Performance
 - RAID technologies
 - Selecting the right RAID type for your IT infrastructure
 - Selecting RAID for fault tolerance
 - Selecting RAID for write performance
 - Selecting RAID for usable capacity
 - Selecting RAID for the storage solution
 - Mixed mode - RAID and JBOD simultaneously
 - Striping
 - RAID 0
 - Mirroring
 - RAID 1 and RAID 1+0 (RAID 10)
 - Read load balancing
 - Parity
 - RAID 5
 - RAID 50
 - RAID 6
 - RAID 60
 - Parity groups
 - Initialize
 - Fast initialization
 - Full initialization
 - No initialization
 - Regenerative writes
 - Backed-out writes
 - Full-stripe writes
 - Spare drives
 - Dedicated spare

- Global spare
- Drive rebuild
- Foreign configuration import
- Transformation
 - Array transformations
 - Expand array
 - Replace drive
 - Volume transformations
 - Transportable controller
 - Expand volume
 - Migrate RAID level
- Drive technology
 - Hot-plug drive LED
 - Consistency check
 - Dynamic sector repair
 - Online drive firmware update
 - Predictive drive failure
 - Patrol read
- Security
 - Drive erase
 - Simple
 - Normal
 - Thorough
 - Sanitize
 - Sanitize overwrite
 - Sanitize block erase
 - Sanitize crypto erase
 - Self-encrypting drive
 - Security Protocol and Data Model
- Reliability
 - Cache Error Checking and Correction
 - Thermal monitoring
- Performance
 - Manage SAS storage link speed
 - Manage PCIe storage interface
 - I/O performance mode
 - Cache
 - Read cache policy
 - Write cache policy
 - Flash-backed write cache
 - I/O policy

- Drive caching
- Strip size selection
- Installation
 - Installing in a configured server
 - Installing in an unconfigured server
 - Installing the OS with the controller driver
 - Installing MR controller on the server
 - Connecting storage devices
 - Connecting internal storage
 - Cable part numbers
- Configuration
 - Configuration tools for HPE MR Controllers
 - Array and controller configuration
 - HPE iLO GUI
 - Intelligent Provisioning
 - HPE MR Storage Administrator
 - HPE OneView
 - Redfish
 - StorCLI
 - UEFI System Utilities
 - Using UEFI System Utilities
 - Configuration in UEFI System Utilities
 - Viewing controller information and performing common actions
 - Configuration management
 - Creating a logical drive
 - Selecting drives to include in a logical drive
 - Creating a profile-based logical drive
 - Importing secured foreign drive
 - Viewing array properties
 - Viewing global spare drives
 - Making a JBOD
 - Making an unconfigured good drive
 - Clearing a configuration
 - Controller management
 - Managing the controller
 - Advanced controller management
 - Clearing controller events
 - Saving controller events
 - Saving a serial log
 - Enabling drive security
 - Disabling drive security

- Changing drive security settings
 - Changing drive security key management mode
 - Managing link speed
 - Managing advanced software options
 - Scheduling a consistency check
 - Setting factory defaults
 - Managing auto-config
 - Configuring advanced controller properties
 - Configuring cache and memory settings
 - Configuring patrol read settings
 - Configuring spare settings
 - Configuring Task Rates
 - Logical drive management
 - Viewing and configuring properties of a logical drive
 - Deleting a logical drive
 - Initializing a logical drive
 - Locating a physical drive associated with a logical drive
 - Erasing a logical drive
 - Drive management
 - Viewing drive properties
 - Locating a drive
 - Initializing a drive
 - Placing a drive offline
 - Erasing a drive
 - Making a JBOD
 - Making an unconfigured good drive
 - Making a bootable drive
 - Assigning a global spare drive
 - Unassigning a global spare drive
 - Sanitizing an unconfigured good drive
- Maintenance
 - Updating software and firmware
 - Error reporting
 - Diagnostic tools
 - Troubleshooting resources
- Models
 - Modular controller (-o)
 - HPE MR216i-o Gen11 SPDM Storage Controller ports and connectors
 - HPE MR408i-o Gen11 SPDM Storage Controller ports and connectors
 - HPE MR416i-o Gen11 SPDM Storage Controller ports and connectors
 - Standup PCIe Plug-In Controller (-p)

- HPE MR216i-p Gen11 SPDM Storage Controller ports and connectors
- HPE MR416i-p Gen11 SPDM Storage Controller ports and connectors
- Additional hardware and options
 - Energy pack options
 - HPE Smart Storage Battery
 - HPE Smart Storage Hybrid Capacitor
- Storage reference
 - Memory and storage capacity conventions
 - RAID conventions
- Websites
- Support and other resources
 - Accessing Hewlett Packard Enterprise Support
 - Accessing updates
 - Remote support
 - Warranty information
 - Regulatory information
 - Documentation feedback

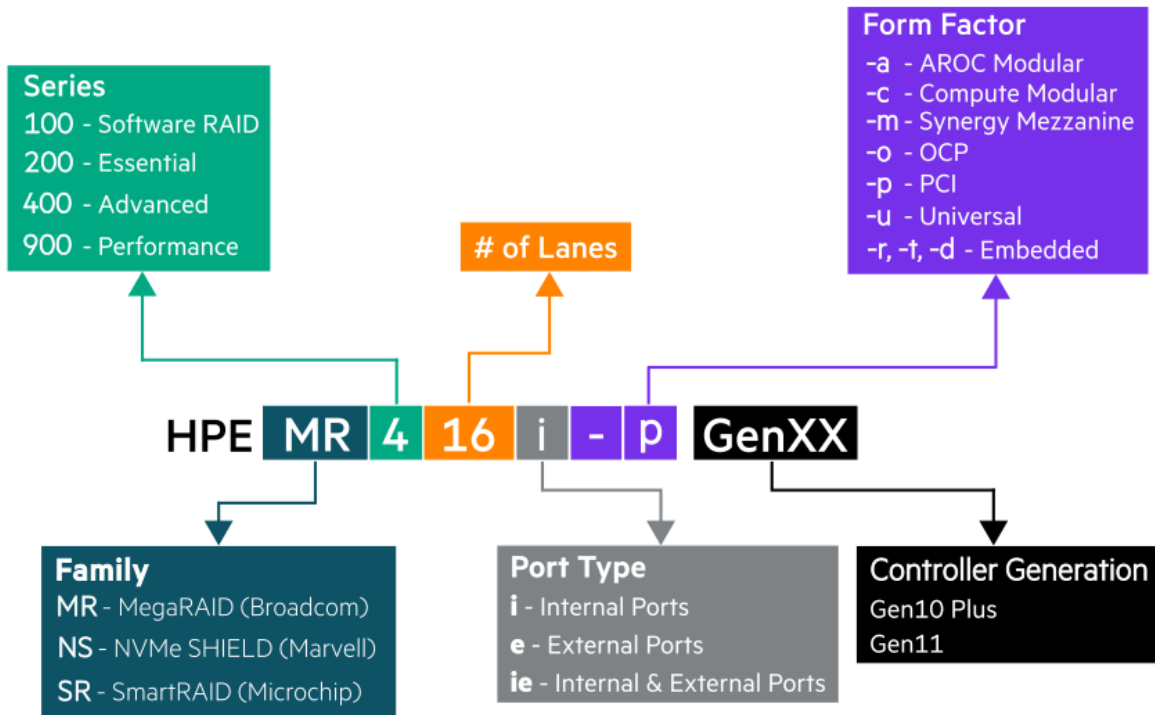
HPE MR Gen11 Controller

The HPE MR Gen11 Controller is a family of tri-mode controllers ideal for maximizing performance while supporting advanced RAID levels. This controller operates in mixed mode, which combines RAID and JBOD operations simultaneously.

This controller provides:

- Up to 16 lanes across 2 x8 internal SlimSAS (SFF-8654) ports
- 16 Gbs NVMe, 12 Gbs SAS, 6 Gbs SATA
- Mixed mode RAID and JBOD functionality simultaneously
- UEFI boot mode
- Support for Self-Encrypting Drive (SED)
- 8 lane PCIe Gen4 host interface
- Security Protocol and Data Model (SPDM) authentication

Figure 1. HPE MR Controller naming framework



Management and configuration tools for HPE MR Gen11 Controller

- [HPE iLO GUI](#)
- [HPE MR Storage Administrator](#)
- [HPE OneView 8.4 or later](#)
- [Intelligent Provisioning 4.31 or later](#)
- [Redfish](#)
- [StorCLI](#)
- [UEFI System Utilities](#)



HPE MR Storage Administrator



i IMPORTANT:

Manage HPE MR Storage Controllers using HPE MR Storage Administrator. You cannot manage it using Smart Storage Administrator.

Subtopics

[200 series](#)

[400 series](#)

200 series

The following controllers are supported.

- HPE MR216i-o Gen11 16 Internal Lanes/No Cache SPDM OCP Storage Controller
- HPE MR216i-p Gen11 16 Internal Lanes/No Cache SPDM PCI Plug-in Storage Controller

These controllers are ideal for all virtualized environments where HBA / Pass-through mode is applicable offering high-bandwidth and up to 3M IOPs (I/O per Second) of 4KiB random read performance.

- RAID levels — 0, 1, 10

Supported servers	HPE MR216i-o Gen11 SPDM Storage Controller	HPE MR216i-p Gen11 SPDM Storage Controller
HPE ProLiant DL320 Gen11	✓	✓
HPE ProLiant DL325 Gen11	✓	✓
HPE ProLiant DL345 Gen11	✓	✓
HPE ProLiant DL360 Gen11	✓	✓
HPE ProLiant DL365 Gen11	✓	✓
HPE ProLiant DL380 Gen11	✓	✓
HPE ProLiant DL380a Gen11	✓	✓
HPE ProLiant DL385 Gen11	✓	✓
HPE ProLiant DL560 Gen11	✓	✓
HPE ProLiant DL20 Gen11	✓	✓
HPE ProLiant ML30 Gen11	✓	✓
HPE ProLiant ML110 Gen11	✓	✓
HPE ProLiant ML350 Gen11	✓	✓
HPE Alletra 4110 Server	✓	✓
HPE Alletra Storage Server 4120	✓	✓

400 series

The following controllers are supported.

- HPE MR408i-o Gen11 8 Internal Lanes/4GB Cache SPDM OCP Storage Controller
- HPE MR416i-o Gen11 16 Internal Lanes/8GB Cache SPDM OCP Storage Controller
- HPE MR416i-p Gen11 16 Internal Lanes/8GB Cache SPDM PCI Plug-in Storage Controller

These controllers are ideal for all data center environments where RAID is applicable offering reduced rebuild times, high-bandwidth, and up to 4.86 GBs of RAID 5 sequential write performance.

- RAID levels—0, 1, 10, 5, 6, 50, 60
- 8G x72 flash-backed write cache for MR416 and 4G x72 flash-backed write cache for MR408 controllers



Supported servers	HPE MR408i-o Gen11 SPDM Storage Controller	HPE MR416i-o Gen11 SPDM Storage Controller	HPE MR416i-p Gen11 SPDM Storage Controller
HPE ProLiant DL320 Gen11	✓	✓	✓
HPE ProLiant DL325 Gen11	✓	✓	✓
HPE ProLiant DL345 Gen11	✓	✓	✓
HPE ProLiant DL360 Gen11	✓	✓	✓
HPE ProLiant DL365 Gen11	✓	✓	✓
HPE ProLiant DL380 Gen11	✓	✓	✓
HPE ProLiant DL380a Gen11	✓	✓	✓
HPE ProLiant DL385 Gen11	✓	✓	✓
HPE ProLiant DL560 Gen11	✓	✓	✓
HPE ProLiant DL20 Gen11	✓	✓	✓
HPE ProLiant ML30 Gen11	✓	✓	✓
HPE ProLiant ML110 Gen11	✓	✓	✓
HPE ProLiant ML350 Gen11	✓	✓	✓
HPE Alletra 4110 Server	✓	✓	✓
HPE Alletra Storage Server 4120	✓	✓	✓

Features

Subtopics

[Controller supported features](#)

[RAID technologies](#)

[Transformation](#)

[Drive technology](#)

[Security](#)

[Reliability](#)

[Performance](#)

Controller supported features

This section lists the features supported by this controller. For additional information about the features, see [HPE MR Storage Administrator User Guide](#).

Subtopics

[Operating environments](#)

[RAID technologies](#)

[Transformation](#)

[Drive technology](#)

[Security](#)

[Reliability](#)

[Performance](#)



Operating environments

Operating System	200 series	400 series
Windows	✓	✓
Linux	✓	✓
VMware	✓	✓
UEFI Boot mode	✓	✓

RAID technologies

The following RAID technologies are supported:

Feature	200 series	400 series
RAID levels	0, 1, 10	0, 1, 5, 6, 10, 50, 60
Max logical drives ¹	240	240
Max logical drives per array	64	64
Max physical drives without expander	16	16
Mixed mode - RAID and JBOD simultaneously	✓	✓
Read load balancing	✓	✓
Parity groups	✓	✓
Initialize	✓	✓
Regenerative writes	✓	✓
Backed out writes	✓	✓
Full-stripe writes	✓	✓
Dedicated spare	✓	✓
Global spare	✓	✓
Drive rebuild	✓	✓
Foreign configuration import	✓	✓

¹ The count 240 contains RAID Virtual Drive and JBOD.

Transformation

The following transformation features are supported:

Feature	200 series	400 series
Expand Array	✓	✓
Replace drive	✓	✓
Transportable controller	✓	✓
Expand volume	✓	✓
Migrate RAID level	✓	✓
Migrate Stripe	—	—

Drive technology

The following drive technology features are supported:

Feature	200 series	400 series
<u>Hot-plug drive LED</u>	✓	✓
<u>Consistency check</u>	✓	✓
<u>Dynamic sector repair</u>	✓	✓
<u>Online drive firmware update</u>	✓	✓
<u>Predictive drive failure</u>	✓	✓
<u>Patrol read</u>	✓	✓

Security

The following security feature is supported:

Feature	200 series	400 series
<u>Drive erase</u>	✓	✓
<u>Sanitize</u>	✓	✓
<u>Self-encrypting drive</u>	✓	✓
<u>Security Protocol and Data Model</u>	✓	✓

Reliability

The following reliability features are supported:

Feature	200 series	400 series
<u>Recovery ROM</u>	—	—
<u>Cache Error Checking and Correction</u>	✓	✓
<u>Thermal monitoring</u>	✓	✓

Performance

The following performance features are supported:

Feature	200 series	400 series
<u>Manage SAS storage link speed</u>	✓	✓
<u>Manage PCIe storage interface</u>	✓	✓
<u>I/O performance mode</u>	✓	✓
<u>Read cache policy</u>	✓	✓
<u>Write cache policy</u>	—	✓
<u>Flash-backed write cache</u>	—	✓
<u>I/O policy</u>	✓	✓
<u>Drive caching</u>	✓	✓
<u>Strip size selection</u>	✓	✓



RAID technologies

Subtopics

[Selecting the right RAID type for your IT infrastructure](#)

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

[Striping](#)

[Mirroring](#)

[Parity](#)

[Spare drives](#)

[Drive rebuild](#)

[Foreign configuration import](#)

Selecting the right RAID type for your IT infrastructure

The RAID setting that you select is based upon the following:

- The fault tolerance required
- The write performance required
- The amount of usable capacity that you need

Subtopics

[Selecting RAID for fault tolerance](#)

[Selecting RAID for write performance](#)

[Selecting RAID for usable capacity](#)

[Selecting RAID for the storage solution](#)

Selecting RAID for fault tolerance

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


This chart shows the relationship between the RAID level fault tolerance and the size of the storage array. The chart includes RAID 0, 5, 50, 10, 6, and 60. It also shows the percent reliability in increments between 1 and one billion and the storage array drive increments between 0 and 96.

This chart assumes that two parity groups are used for RAID 50 and RAID 60.

This chart shows that:

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





Selecting RAID for write performance


If your environment requires high write performance, select a RAID type that is optimized for write performance

The chart below shows how RAID 10, 5, 50, 6, and 60 compare to the percent write performance of RAID 0.

The data in the chart assumes that the performance is drive limited and that drive write performance is the same as drive read performance.

Consider the following points:

- RAID 5, 50, 6, and 60 performance assumes parity initialization has completed.
- Write performance decreases as fault tolerance improves due to extra I/O.
- Read performance is generally the same for all RAID levels except for smaller RAID 5\6 arrays.



The table below shows the Disk I/O for every host write:

Supported RAID levels may vary based on the controller model.

RAID type	Disk I/O for every host write
RAID 0	1
RAID 10	2
RAID 5	4
RAID 6	6

Selecting RAID for usable capacity

If your environment requires a high usable capacity, select a RAID type that is optimized for usable capacity. The chart in this section demonstrates the relationship between the number of drives in the array and the percent usable capacity over the capacity for RAID 0.


Consider the following points when selecting the RAID type:

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

Note the minimum drive requirements for the RAID types, as shown in the table below.

RAID type	Minimum number of drives
RAID 0	1
RAID 10	2
RAID 5	3
RAID 6	4
RAID 50	6
RAID 60	8

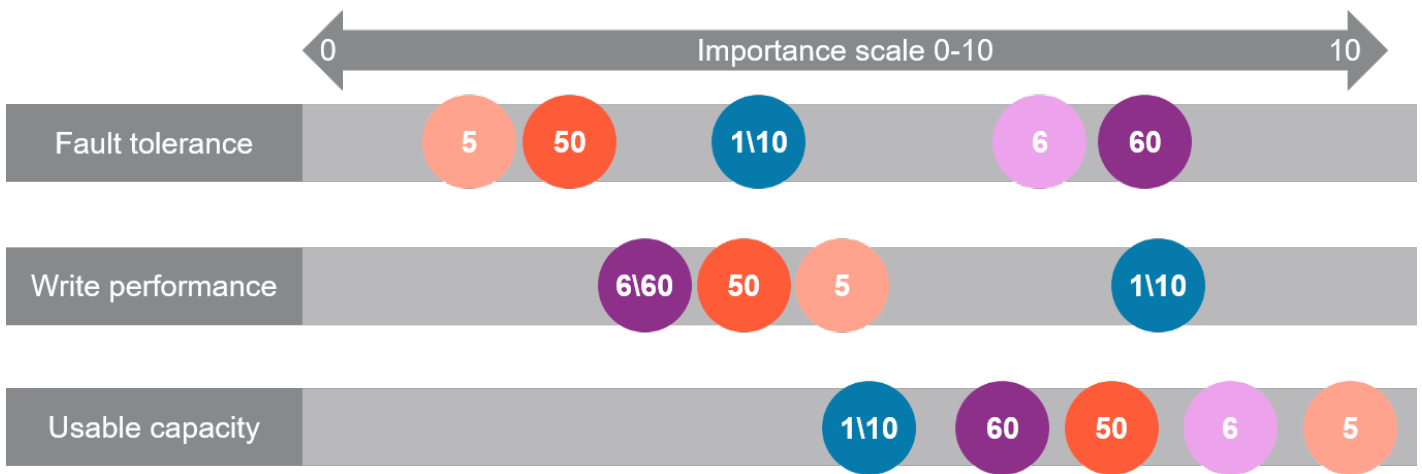
Supported RAID levels may vary based on the controller model.



Selecting RAID for the storage solution

The chart in this section shows the relevance of the RAID type to the requirements of your environment. Depending on your requirements, you should optimize the RAID types as follows:

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



Mixed mode - RAID and JBOD simultaneously

Mixed mode allows any drive to be configured as RAID or Just a Bunch of Drives (JBOD). This option can reduce the number of controllers in the system and efficiently use drive bays within a backplane. For example, a solution that needs all the drives to be presented as JBOD (except a two-drive mirror for boot support) can be achieved with a single controller attached to a single backplane.

In HPE MR controller, a drive can be configured as:

- A member of a logical drive: The drive is created as a logical drive. 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.

When a new drive is added, the Auto-Config option is used by the controller firmware to decide the drive state. If Auto-Config is set to JBOD, the new drive status is automatically configured as JBOD. If Auto-Config is set to None, the new drive status stays in Unconfigured Good.

NOTE:

- For firmware versions earlier than 52.26.3-5250, Auto-Config is set to JBOD and cannot be configured for all the controller models.
- For firmware versions 52.26.3-5250 or later, Auto-Config can be configured.
 - For HPE MR200 controllers, by default it is set to JBOD.
 - For HPE MR400 controllers, by default it is set to None.

More information

[Managing auto-config](#)

Striping


Subtopics

[RAID 0](#)

RAID 0

A RAID 0 configuration provides data striping, but 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. The minimum number of drives required is one.

The maximum number of drives supported for RAID 0 is 32.



This method has the following benefits:

- It is useful when performance and low cost are more important than data protection.
- It has the highest write performance of all RAID methods.
- It has the lowest cost per unit of stored data of all RAID methods.
- It uses the entire drive capacity to store data (none allocated for fault tolerance).

Mirroring

Subtopics

RAID 1 and RAID 1+0 (RAID 10)


Read load balancing

RAID 1 and RAID 1+0 (RAID 10)

In RAID 1 and RAID 1+0 (RAID 10) configurations, data is duplicated to a second drive. The usable capacity is $C \times (n / 2)$ where C is the drive capacity with n drives in the array. A minimum of two drives is required.


When the array contains only two physical drives, the fault-tolerance method is known as RAID 1.

The maximum number of drives supported for RAID 1 is 32.



When the array has more than two physical drives, drives are mirrored in pairs, and the fault-tolerance method is known as RAID 1+0 or RAID 10. If a physical drive fails, the remaining drive in the mirrored pair can still provide all the necessary data. Several drives in the array can fail without incurring data loss, as long as no two failed drives belong to the same mirrored pair. The total drive count must increment by 2 drives. A minimum of four drives is required.

The maximum number of drives supported for RAID 10 is 32.



This method has the following benefits:

- It is useful when high performance and data protection are more important than usable capacity.
- This method has the highest write performance of any fault-tolerant configuration.
- No data is lost when a drive fails, as long as no failed drive is mirrored to another failed drive.
- Up to half of the physical drives in the array can fail.

Read load balancing

In each mirrored pair or trio, the controller balances read requests between drives based upon individual drive load.

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



Parity

Subtopics

[RAID 5](#)

[RAID 50](#)

[RAID 6](#)

[RAID 60](#)

[Parity groups](#)

[Initialize](#)

[Regenerative writes](#)


[Backed-out writes](#)

[Full-stripe writes](#)

RAID 5

RAID 5 protects data using parity (denoted by $P_{x,y}$ in the figure). Parity data is calculated by summing (XOR) the data from each drive within the stripe. The strips of parity data are distributed evenly over every physical drive within the logical drive. When a physical drive fails, data that was on the failed drive can be recovered from the remaining parity data and user data on the other drives in the array. The usable capacity is $C \times (n - 1)$ where C is the drive capacity with n drives in the array. A minimum of three drives is required.

The maximum number of drives supported for RAID 5 is 32.




This method has the following benefits:

- It is useful when usable capacity, write performance, and data protection are equally important.
- It has the highest usable capacity of any fault-tolerant configuration.
- Data is not lost if one physical drive fails.

RAID 50

RAID 50 is a nested RAID method in which the constituent drives are organized into several identical RAID 5 logical drive sets (parity groups). The smallest possible RAID 50 configuration has six drives organized into two parity groups of three drives each.





For any given number of drives, data loss is least likely to occur when the drives are arranged into the configuration that has the largest possible number of parity groups. For example, four parity groups of three drives are more secure than three parity groups of four drives. However, less data can be stored on the array with the larger number of parity groups.

All data is lost if a second drive fails in the same parity group before data from the first failed drive has finished rebuilding. A greater percentage of array capacity is used to store redundant or parity data than with non-nested RAID methods (RAID 5, for example). A minimum of six drives is required.

The maximum number of drives supported for RAID 50 is 256.

This method has the following benefits:


- Higher performance for RAID 5, especially during writes.
- Better fault tolerance than either RAID 0 or RAID 5.
- Up to n physical drives can fail (where n is the number of parity groups) without loss of data, as long as the failed drives are in different parity groups.

RAID 6

RAID 6 protects data using double parity. With RAID 6, two different sets of parity data are used (denoted by $P_{x,y}$ and $Q_{x,y}$ in the figure), allowing data to still be preserved if two drives fail. Each set of parity data uses a capacity equivalent to that of one of the constituent drives. The usable capacity is $C \times (n - 2)$ where C is the drive capacity with n drives in the array.

A minimum of 4 drives is required.

The maximum number of drives supported for RAID 6 is 32.



This method is most useful when data loss is unacceptable but cost is also an important factor. The probability that data loss will occur when an array is configured with RAID 6 (Advanced Data Guarding (ADG)) is less than it would be if it were configured with RAID 5.

This method has the following benefits:

- It is useful when data protection and usable capacity are more important than write performance.
- It allows any two drives to fail without loss of data.

RAID 60

RAID 60 is a nested RAID method in which the constituent drives are organized into several identical RAID 6 logical drive sets (parity groups). The smallest possible RAID 60 configuration has eight drives organized into two parity groups of four drives each.

For any given number of hard drives, data loss is least likely to occur when the drives are arranged into the configuration that has the largest possible number of parity groups. For example, five parity groups of four drives are more secure than four parity groups of five drives. However, less data can be stored on the array with the larger number of parity groups.

The number of physical drives must be exactly divisible by the number of parity groups. Therefore, the number of parity groups that you can specify is restricted by the number of physical drives. The maximum number of parity groups possible for a particular number of physical drives is the total number of drives divided by the minimum number of drives necessary for that RAID level (three for RAID 50, 4 for RAID 60).

A minimum of 8 drives is required.

The maximum number of drives supported for RAID 60 is 256.

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.

This method has the following benefits:

- Higher performance for RAID 6, especially during writes.
- Better fault tolerance than RAID 0, 5, 50, or 6.
- Up to $2n$ physical drives can fail (where n is the number of parity groups) without loss of data, as long as no more than two failed drives are in the same parity group.

Parity groups

When you create a RAID 50 or RAID 60 configuration, you must also set the number of parity groups.

You can use any integer value greater than 1 for this setting, with the restriction that the total number of physical drives in the array must be exactly divisible by the number of parity groups.

The maximum number of parity groups possible for a particular number of physical drives is the total number of drives divided by the minimum number of drives necessary for that RAID level (three for RAID 50, four for RAID 60).

This feature has the following benefits:

- It supports RAID 50 and RAID 60.
- A higher number of parity groups increases fault tolerance.

Initialize

Initialize a logical drive after you configure it. When you initialize the logical drive, you prepare the storage medium for use.

CAUTION:

All data on the logical drive is lost when you initialize it. Before you start this operation, back up any data that you want to keep.

Subtopics

[Fast initialization](#)

[Full initialization](#)

[No initialization](#)

Fast initialization

During fast initialization, firmware quickly overwrites the first and last 8 MB regions of the new logical drive, clears any boot records or partition information, and 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, writes to a RAID 5 or RAID 6 logical drive are typically faster because the controller does not read the entire stripe (regenerative write) to update the parity data. This feature initializes parity blocks in the background while the logical drive is available for access by the OS.

Parity initialization takes several hours to complete. The time it takes depends on the size of the logical drive and the load on the controller. While the controller initializes the parity data in the background, the logical drive has full fault tolerance.

This method has the benefit of allowing you to start writing data to the logical drive immediately.

Background initialization

To access the background initialization (BGI) rate, select **More Actions > Set Adjustable Task Rate**, and then locate it under the **Priority Percentage** column.

Enter a number from **1–100**. The higher the number, the faster the initialization occurs and the system I/O rate might be slower as a result.

If you use RAID 5, you must have a minimum of five drives for a background initialization to start.

If you use RAID 6, you must have at least seven drives for a background initialization to start.

RAID	Minimum drives	Minimum drives for Auto BGI to set to ON status by default
RAID 5	3 ¹	5
RAID 50	6 ¹	10 (two parity groups)
RAID 6	4 ¹	7
RAID 60	8 ¹	14 (two parity groups)

¹ When creating smaller volumes, BGI can be started manually. For more information, see `storcli /cx/vx set autobgi=<on|off>` in the [HPE MR StorCLI User Guide](#).

NOTE:

- Minimum disk values were decided as a trade-off where smaller volumes do not see a performance issue when the volume is not consistent and does not require background initialization. This method allows the user to use the volume more quickly and after consistency check is complete, performance is fully restored.
- User can manually initialize the parity blocks, but it causes performance issue. Otherwise, system runs consistency check automatically and generates parity.

Full initialization

During full initialization, a complete initialization is done on the new configuration. You cannot write data to the new logical drive 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 data on the logical drive.

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

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 logical drive at a later time.

Regenerative writes

Logical drives can be created with background parity initialization so that they are available almost instantly. During this temporary parity initialization process, writes to the logical drive 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 much 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 you can see, the write penalty is greater (slower write performance) with larger arrays.

This method has the following benefits:

- It allows the logical drive to be accessible before parity initialization completes.
- It allows the logical drive to be accessible when degraded.

Backed-out writes

After parity initialization is complete, random writes to a RAID 5, 50, 6, or 60 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 RAID 50 is always four drive operations, and the write penalty for a RAID 6 and RAID 60 is always six drive operations. As you can see, the write penalty is not influenced by the number of drives in the array.

Backed-out writes is also known as "read-modify-write."

This method has the benefit of faster RAID, 5, 50, 6, or 60 random writes.

Full-stripe writes

When writes to the logical drive are sequential or when multiple random writes that accumulate in the flash-backed write cache are found to be sequential, a full-stripe write operation can be performed. A full-stripe write allows the controller to calculate new parity using new data being written to the drives. There is almost no write penalty because the controller does not need to read old data from the drives to calculate the new parity. As the size of the array grows larger, the write penalty is reduced by the ratio of p / n where p is the number of parity drives and n is the total number of drives in the array.

This method has the benefit of faster RAID 5, 6, or 60 sequential writes.

Spare drives

Subtopics

[Dedicated spare](#)

[Global spare](#)



Dedicated spare

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

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

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

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 logical drive. For RAID 0 logical drives, the global spare is active when a member drive reports a predictive failure.

Drive rebuild

If a drive, which is configured as 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. The rebuild is a fully automatic process. Monitor the progress of drive rebuilds in the Background Processes in Progress window.

Access the drive rebuild rate by selecting **Set Adjustable Task Rate** under the **More Actions** menu then locating it under the **Priority Percentage** column. Enter a number from 1 to 100. The higher the number, the faster the rebuild will occur (and the system I/O rate might be slower as a result).

Foreign configuration import

A foreign configuration import is a RAID configuration that exists on a replacement set of drives that you install in a computer system. You can use the **HPE MR Storage Administrator** to import the foreign configuration to the controller or clear the foreign configuration so that you can create a configuration using these drives.

Transformation

Subtopics

[Array transformations](#)

[Volume transformations](#)

Array transformations

Subtopics

[Expand array](#)

[Replace drive](#)

Expand array

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



- It must be an unassigned drive.
- It must be of the same type as existing drives in the array example, NVMe SSD, SAS HDD, SAS SSD, SATA HDD, or SATA SSD.
- It must have a capacity no less than the capacity of the smallest drive in the array.

This operation uses the **Modify Array** option in the HPE MR Storage Administrator user interface. This feature is supported when there is a single logical drive configured in the array.

Replace drive

The replace drive operation allows you to replace failed physical drives in the array with healthy physical drives. The original array and logical drive numbering is unaffected after the replacement. Note the following conditions and restrictions for the **Heal Array** operation:

- The replacement physical drives and original drives must be the same interface type such as SAS or SATA.
- The operation is available only if enough unassigned physical drives of the correct size are available.
- The array is not transforming for example, rebuilding to a spare.
- The array has a working cache, making it capable of transformation.

Volume transformations

Subtopics

[Transportable controller](#)

[Expand volume](#)

[Migrate RAID level](#)

Transportable controller

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

In this design, the controller firmware assumes that the new server has the same configuration. That is, the configuration includes the same server generation and family, and logical drives are migrated to the new target server to facilitate cache flush when the data is restored.

Expand volume

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

Increase the capacity of an existing logical drive by specifying a new size. Once the task is performed, use operating system partitioning software to take advantage of the extended space available.

Enable this feature in the HPE MR Storage Administrator user interface using the **Expand** option within the logical drive **Actions** menu.

Migrate RAID level

RAID level transformation is the process of converting one RAID configuration to another. You can perform RAID level transformation at the array level.

The RAID level transformation feature allows you to change the current level of fault tolerance (RAID type) for your logical drive. When the fault tolerance changes, you may have more or less unused space, depending on the fault tolerance with which you started.

This operation uses the **Modify Array** option in the HPE MR Storage Administrator user interface. This feature is supported when there is a single logical drive configured in the array.



The following table describes the valid RAID level transformation matrix.

Initial RAID level	Migrated RAID level
RAID 0	RAID 1
RAID 0	RAID 5
RAID 0	RAID 6
RAID 1	RAID 0
RAID 1	RAID 5
RAID 1	RAID 6
RAID 5	RAID 0
RAID 5	RAID 6
RAID 6	RAID 0
RAID 6	RAID 5

Drive technology

Subtopics

[Hot-plug drive LED](#)

[Consistency check](#)

[Dynamic sector repair](#)

[Online drive firmware update](#)

[Predictive drive failure](#)

[Patrol read](#)

Hot-plug drive LED

Figure 1. LFF Low Profile (LP)





Figure 2. SFF Basic Carrier (BC)





Item	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 (1 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 (1 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.
		Flashing green (4 flashes per second)	The drive is operating normally and has activity.
		Flashing green (1 flash per second)	The drive is doing one of the following: <ul style="list-style-type: none"> Rebuilding Performing a RAID migration Performing a strip size migration Performing a capacity expansion Performing a logical drive 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.

Consistency check

A consistency check operation verifies the correctness of the data in logical drives that use RAID levels 1, 5, 6, 10, 50, and 60. For example, in a system with parity, checking consistency means to calculate the data on one drive and comparing the results to the contents of the parity drive.

Periodically run a consistency check on fault-tolerant logical drives. Because RAID 0 does not provide data redundancy, you cannot run a consistency check on RAID 0 logical drives.

To run a consistency check, you must first set the consistency check properties, and then you can either:

- Schedule a consistency check to be run at an interval that you define.
- Start the consistency check operation immediately.

The consistency check priority value ranges from 1 to 100.

You can use the following modes for the consistency check:

- Concurrent - Run consistency check concurrently on all logical drives.
- Sequential - Run consistency check on one logical drive at a time.
- Disable - Disables consistency check.

Access the consistency check rate by selecting Set Adjustable Task Rate under the More Actions menu and then locating it under the Priority Percentage column. Enter a number from 1 to 100. The higher the number, the faster the consistency check is performed (and the system I/O rate might be slower as a result).

Dynamic sector repair

Disk drive media can develop defects caused by variances in the drive mechanisms under normal operating conditions. To protect data from media defects, Hewlett Packard Enterprise has built a dynamic sector repair feature into these controllers.

Online drive firmware update

These controllers support online drive flashing, which saves time when updating drive firmware. Instead of taking the drive offline before loading a new firmware image, you can download an updated drive firmware image to the controller and update all the drives while the server is online.

Predictive drive failure

These controllers use Self-Monitoring and Reporting Technology (SMART) to inform the host when a drive is experiencing abnormal operation likely to lead to drive failure.

SMART places the monitoring capabilities within the drive. These monitoring routines have direct access to internal performance, calibration, and error measurements for a specific drive type.

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 for all RAID levels and for all spare drives. A patrol read is initiated only when the controller is idle for a defined period and has no other background activities. You can set the patrol read properties and start the patrol read operation, or you can start the patrol read without changing the properties.

Access the patrol rate by selecting **Set Adjustable Task Rate** under the **More Actions** menu then locating it under the **Priority Percentage** column. Enter a number from 1 to 100. Setting an higher rate means faster patrol read process and as a result the system I/O rate might be slower.

Patrol read:

- scans drives and repairs media defects
- detects and repairs media defects when accessing a bad sector during busy periods

Security

Subtopics

[Drive erase](#)

[Sanitize](#)

[Self-encrypting drive](#)

[Security Protocol and Data Model](#)

Drive erase

Erase data on drives by using the drive erase option. The erase operation consists of a series of write operations to a drive that overwrite every user-accessible sector of the drive with specified patterns. The erase operation can be repeated in multiple passes using different data patterns for enhanced security. The erase operation is performed as a background task.

This operation can be performed on the physical drive or logical drive. For the physical drive, you must change the drive state to Unconfigured good.

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



Subtopics

Simple

Normal

Thorough

Simple

A Simple erase writes a pattern to the logical drive in a single pass.

Normal

A normal erase operation is a three-pass operation that first overwrites the drive contents with random values then overwrites it twice with patterns.

Thorough

A thorough drive erase operation repeats the Normal drive erase operation three times.

Sanitize

Sanitize crypto erase also known as Cryptographic erase uses cryptographic technology to perform an instance secure erase of all user data. This method has the following benefits:

- Removes all sensitive information from the drive.
- Completes within seconds.

Subtopics

Sanitize overwrite

Sanitize block erase

Sanitize crypto erase

Sanitize overwrite

Sanitize overwrite (hard drive) fills every physical sector of the drive with a pattern.

This method has the following benefits:

- Removes all sensitive information from the drive.
- Once started, the drive will continue to sanitize regardless of resets and power cycles.

Sanitize block erase

Sanitize block erase (SSD) sets the blocks on the drive to a vendor-specific value, removing all user data.

This method has the following benefits:

- Removes all sensitive information from the drive.
- Once started, the drive continues to sanitize regardless of resets and power cycles.

Sanitize crypto erase

Sanitize crypto erase also known as instant secure erase uses cryptographic technology to perform an instant secure erase of all user data.

This method has the following benefits:

- Removes all sensitive information from the drive.
- Completes within seconds.

Self-encrypting drive

The HPE MR Gen11 Controller supports Self-Encrypting Drive (SED) that secures the drive data from unauthorized access or modification of data. As the data on the drive is encrypted even if the SED drive is removed from its storage system, it cannot be accessed without appropriate security authorization.

Host key management

To use host key management, enable the SED drive as JBOD and expose the drive to OS. This method allows you to manage SED using third-party key management like `SEDutil`. SED monitoring is also available in HPE MR Storage Administrator, Storage Command Line Interface (StorCLI) tool, and configuration utility in UEFI System Utilities..

Local key management

You can enable SED drive security for local key management using the HPE MR Storage Administrator, StorCLI tool, and configuration utility in UEFI System Utilities. You must provide a controller-wide security key identity and security key. While boot up, the security key stored in the controller is used to unlock the drive. Whenever the drive is powered down, the security-enabled drive data encryption key is locked. This action protects the drives or systems against any theft.

Remote key management

Remote key management is also known as external key management.



NOTE:

You can enable SED drive security for remote key management using the configuration utility in UEFI System Utilities. For more information, see [Enabling drive security](#).

The configuration utility in UEFI System Utilities works with iLO key manager to create the security key identify and security key in the remote key manager server. iLO key manager needs to be configured before enabling remote key management in the configuration utility. Whenever the drive is powered down, the security-enabled drive data encryption key is locked. While boot up, the security key is retrieved from the remote key manager server to unlock the drive.

Security Protocol and Data Model

Security Protocol and Data Model (SPDM) is a security standard developed by Distributed Management Task Force (DMTF). It enables system hardware components such as PCIe cards, NVMe drives to have their identity authenticated and their integrity verified.

SPDM-capable components have strong cryptographic identities and can provide cryptographically signed attestations of their security state. When the server starts, SPDM-capable components are authenticated cryptographically. Measurements of their security-relevant properties are obtained to determine whether they operate at their intended state and then control is passed to the OS.

Reliability

Subtopics

[Cache Error Checking and Correction](#)

[Thermal monitoring](#)

Cache Error Checking and Correction

Error checking and correction (ECC) DRAM technology protects the data while it is in cache. The ECC scheme generates 8 bits of check data for every 64 bits of regular data transferred. The memory controller uses this information to detect and correct data errors originating inside the DRAM chip or across the memory bus.

Thermal monitoring

The controller monitors the temperature of each drive in the server. 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.

Performance

Subtopics

[Manage SAS storage link speed](#)

[Manage PCIe storage interface](#)

[I/O performance mode](#)

[Cache](#)

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.

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.

I/O performance mode

The I/O performance mode feature also known as FastPath is an intelligent I/O passthrough mechanism for Solid State Drive (SSD) arrays. This advanced software is an optimized version of controller technology that can dramatically boost 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 logical drive when the logical drive is created with properties Direct I/O, Write Through, and No Read Ahead.

The logical drive must be in an optimal state for I/O performance mode to be enabled.

I/O performance mode allows the following I/O scenarios:

- Read I/Os to all RAID levels when I/O size is less than or equal to the strip size.
- Write I/Os to RAID 0 when I/O size is less than or equal to the strip size.
- All I/Os targeted to a single drive RAID 0.

Write I/Os to RAID 1, 5, and 6 do not use I/O performance mode.

Cache

Subtopics

[Read cache policy](#)

[Write cache policy](#)

[Flash-backed write cache](#)

[I/O policy](#)

[Drive caching](#)

[Strip size selection](#)

Read cache policy

The read cache policy options for the logical drive are:

- No Read Ahead - In no read ahead mode, read ahead capability is disabled. This setting is the default option.
- Read ahead - Read ahead capability allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This process speeds up reads for sequential data, but there is little improvement when accessing random data.

Write cache policy

The write cache policy options for the logical drive are:

- Write Through - In this mode, the controller 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 - In this mode, the controller 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 - In this mode, the controller sends a data transfer completion signal to the host when the controller cache receives all the data in a transaction. If you select the Always Write Back policy and the energy pack is absent, the firmware is forced to use the Write Back policy.

This controller now supports two power options to enable the controller cache features. Collectively, these two options are known as the energy pack and include the following,

- HPE Smart Storage Battery
- HPE Smart Storage Hybrid Capacitor

Since the interface does not display the cache policy used during the creation of the logical drive, you can identify the policy by:

- Viewing the energy pack charge status in iLO and the write cache status.
- Viewing cache policy information in the event log.
- Temporarily disabling the cache, which displays the original cache setting.

Flash-backed write cache

The HPE MR controllers use a write-back caching scheme that allows host applications to continue without waiting for write operations to be completed on the disk. A controller without a write-back cache returns completion status to the OS after it writes the data to the drives. A controller with write-back caching can post write data to high-speed cache memory, and return completion status to the OS immediately. The write operation can be completed in microseconds rather than milliseconds. The controller writes the data from its write cache to the disk later, at an optimal time for the controller.

HPE MR 400 controllers support Flash-Backed Write Cache (FBWC), which uses flash devices to retain cache data and energy pack to provide power during a power loss. The FBWC feature offers significant advantages over the earlier Battery-Backed Write Cache (BBWC) systems. While a BBWC requires backup power during the entire power loss, FBWC only needs power during the time it takes to backup from DRAM to flash. As the FBWC writes the contents of memory to flash devices, there is no longer a 48-hour energy pack life limitation. The data is posted to the disk drive on the next power-up.

I/O policy

The I/O policy applies to reads on a specific logical drive. It does not affect the read ahead cache. Direct IO option is available for logical drive.

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

Drive caching

The drive caching options for the logical drive are:

- Unchanged - Leave the current drive cache policy as is.
- Enabled - Enable the drive cache.
- Disabled - Disable the drive cache.

Strip size selection

When a logical drive is created, the unit of data that is stored on each drive is defined as a “strip” (ranging in size from 64 KiB to 1 MiB). These strips are distributed across the physical drives in the array.

Best performance is obtained by aligning and sizing the strip size to the application I/O request size. The smaller (≤ 64 KiB) the strip size, the longer the background parity scans and rebuilds take and the more impact to the host I/O during these operations. However, multiple stripes can be checked or rebuilt concurrently.

Installation

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 MR controller on the server](#)

[Connecting storage devices](#)

[Cable part numbers](#)

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.

See the [server user guide](#), for server-specific procedures.

**WARNING:**

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

11. Remove the riser.
12. Select an available x8 or larger PCIe expansion slot.
13. Remove the slot cover.

Save the retaining screw, if one is present.
14. 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.
15. 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.
16. Connect storage devices to the controller.

For cabling information, see the [server user guide](#).
17. Reinstall the riser.
18. Connect the power cord to the server.
19. Connect the power cord to the power source.
20. Power up all peripheral devices.
21. Power up the server.

More information

[Updating software and firmware](#)

[Array and controller configuration](#)

[Connecting internal storage](#)

Installing in an unconfigured server

Procedure

1. [Update the server firmware](#).
2. Select an available x8 or larger PCIe expansion slot.
3. Remove the slot cover.

Save the retaining screw, if one is present. See the [server user guide](#), for server-specific procedures.
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. Install physical drives, as needed.
7. Connect 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, you must [configure a logical drive](#) 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](#)

More information

[Updating software and firmware](#)

[Array and controller configuration](#)

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 logical drive 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 logical drive or select physical drives for JBOD and retry the installation.
5. After the installation is complete, use 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 "Updating iLO or server firmware by using the flash firmware feature " in [HPE iLO 6 User Guide](#).

More information

[Creating a logical drive](#)

[Making a JBOD](#)

Installing MR controller on the server

For product-specific installation information, refer to the respective product guide.

Server/ Platform	Path to user guide
HPE ProLiant DL320 Gen11	https://www.hpe.com/info/dl320gen11-docs
HPE ProLiant DL325 Gen11	https://www.hpe.com/info/dl325gen11-docs
HPE ProLiant DL345 Gen11	https://www.hpe.com/info/dl345gen11-docs
HPE ProLiant DL360 Gen11	https://www.hpe.com/info/dl360gen11-docs
HPE ProLiant DL365 Gen11	https://www.hpe.com/info/dl365gen11-docs
HPE ProLiant DL380 Gen11	https://www.hpe.com/info/dl380gen11-docs
HPE ProLiant DL380a Gen11	https://www.hpe.com/info/dl380agen11-docs
HPE ProLiant DL385 Gen11	https://www.hpe.com/info/dl385gen11-docs
HPE ProLiant DL560 Gen11	https://www.hpe.com/info/dl560gen11-docs
HPE ProLiant DL20 Gen11	https://www.hpe.com/info/dl20gen11-docs
HPE ProLiant ML30 Gen11	https://www.hpe.com/info/ml30gen11-docs
HPE ProLiant ML110 Gen11	https://www.hpe.com/info/ml110gen11-docs
HPE ProLiant ML350 Gen11	https://www.hpe.com/info/ml350gen11-docs
HPE Alletra Storage 4120 Server	https://www.hpe.com/info/alletra4120-docs
HPE Alletra 4110 Server	https://www.hpe.com/info/alletra4110-docs

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 logical drive must meet the following criteria:

- They must be either NVMe, SAS, or SATA.
- 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 [HPE MR Gen11 Storage Controllers QuickSpecs](#) that is compatible with the controller:
 - 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.

More information

[Array and controller configuration](#)

Cable part numbers

For more information on cables, see the server [HPE MR Gen11 Storage Controllers QuickSpecs](#).

Configuration

Subtopics

[Configuration tools for HPE MR Controllers](#)

[Array and controller configuration](#)

[Configuration in UEFI System Utilities](#)

Configuration tools for HPE MR Controllers

You can use these tools to configure the HPE MR controllers.





[HPE iLO GUI](#)

[HPE iLO 6 User Guide](#)



[HPE MR Storage Administrator](#)

[HPE MR Storage Administrator User Guide](#)



[HPE OneView](#)

[HPE OneView User Guide for VM](#)

[HPE OneView User Guide for Synergy](#)



[Intelligent Provisioning](#)

[Intelligent Provisioning 4.31 User Guide for HPE ProLiant and Synergy Gen11 Servers](#)



[Redfish](#)

[HPE Storage Controllers – Management overview](#)

[HPE StorCLI](#)

[HPE StorCLI User Guide](#)



[UEFI System Utilities](#)

[UEFI System Utilities User Guide for HPE ProLiant Gen11 Servers, and HPE Synergy](#)

Array and controller configuration

During the initial provisioning of the server or compute module, you must configure the controller using the configuration utility in UEFI System Utilities.

After the initial provisioning of the server or compute module, you can use any of the following options to configure the arrays and controllers:

- UEFI System Utilities
- HPE MR Storage Administrator
- StorCLI

HPE MR Storage Administrator and StorCLI are available in the HPE Service Pack for ProLiant.

For more information about using each configuration utility, see the documentation for the configuration utility.

 **NOTE:**

- Any RAID configuration created for the MR controller cannot be moved to SR controllers.
- The message

Data Protection disabled

in the logical drive properties can be ignored as it refers to a feature not currently supported by the HPE MR Storage Administrator product.

Subtopics

[HPE iLO GUI](#)

[Intelligent Provisioning](#)

[HPE MR Storage Administrator](#)

[HPE OneView](#)

[Redfish](#)

[StorCLI](#)

[UEFI System Utilities](#)

HPE iLO GUI

HPE iLO 6 is a remote server management processor embedded on the system boards of supported HPE servers and compute modules. HPE iLO enables the monitoring and controlling of servers from remote locations. It is a powerful management tool that provides multiple ways to configure, update, monitor, and repair servers remotely.

For more information, see [HPE iLO 6 User Guide](#)

Intelligent Provisioning

Intelligent Provisioning is a single-server deployment tool embedded in HPE ProLiant servers. Intelligent Provisioning simplifies server setup, providing a reliable and consistent way to deploy servers.

Intelligent Provisioning prepares the system for installing original, licensed vendor media and Hewlett Packard Enterprise-branded versions of OS software. Intelligent Provisioning also prepares the system to integrate optimized server support software from the HPE Service Pack for ProLiant (SPP). SPP is a comprehensive systems software and firmware solution for HPE ProLiant servers, server blades, their enclosures, and HPE Synergy compute modules. These components are preloaded with a basic set of firmware and OS components that are installed along with Intelligent Provisioning.

After the server is running, you can update the firmware to install additional components. You can also update any components that have been outdated since the server was manufactured.

To access Intelligent Provisioning :

- Press F10 from the POST screen and enter either Intelligent Provisioning or HPE Server Message Block (SMB) Setup.
- From the iLO web interface access Lifecycle Management. It allows you to access Intelligent Provisioning without rebooting your server.

HPE MR Storage Administrator

HPE MR Storage Administrator is a web-based application that enables you to monitor, configure, maintain, and troubleshoot the HPE MR Gen11 Controller. 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.

i 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](#).

HPE OneView

HPE OneView is a secure virtual appliance that runs on a virtual machine host. It enables you to comprehensively manage a converged data center throughout the hardware life cycle. HPE OneView makes it possible to easily monitor, configure, and manage physical and logical compute resources using either the GUI or REST API.

For more information, see [HPE OneView User Guide](#).

Redfish

These controllers support the DMTF standard known as PLDM for Redfish device enablement in HPE ProLiant Gen10 server and later version 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 feature and capabilities are owned by the option card firmware.

GET requests

The following table lists the Redfish resources for GET requests:

Redfish Resource	Method	URL
Storage	GET	/redfish/v1/Systems/{ID}/Storage/{ID}
Storage Controller Collection	GET	/redfish/v1/Systems/{ID}/Storage/{ID}/Controllers
Storage Controller	GET	/redfish/v1/Systems/{ID}/Storage/{ID}/Controllers/{ID}
Port Collection	GET	/redfish/v1/Systems/{ID}/Storage/{ID}/Controllers/{ID}/Ports
Port	GET	/redfish/v1/Systems/{ID}/Storage/{ID}/Controllers/{ID}/Ports/{ID}
Volume Collection	GET	/redfish/v1/Systems/{ID}/Storage/{ID}/Volumes
Volume Capabilities	GET	/redfish/v1/Systems/{ID}/Storage/{ID}/Volumes/Capabilities
Volume	GET	/redfish/v1/Systems/{ID}/Storage/{ID}/Volumes/{ID}
Drive	GET	/redfish/v1/Systems/{ID}/Storage/{ID}/Drives/{ID}

Write requests

Redfish write is supported in firmware versions 52.24.3-4948 and later.

The following table lists the Redfish resources for the write requests:



Redfish Resource	Method	URL
Volume Create	POST	/redfish/v1/Systems/{ID}/Storage/{ID}/Volumes
Volume Delete	DEL	/redfish/v1/Systems/{ID}/Storage/{ID}/Volumes/{ID}
Drive Secure Erase	POST	/redfish/v1/Systems/{ID}/Storage/{ID}/Drives/{ID}/Actions/Drive.SecureErase
Volume Initialize	POST	/redfish/v1/Systems/{ID}/Storage/{ID}/Volumes/{ID}/Actions/Volume.Initialize
Storage Reset to Defaults	POST	/redfish/v1/Systems/{ID}/Storage/{ID}/Actions/Storage.ResetToDefaults
Drive Hotspare Type	PATCH	/redfish/v1/Systems/{ID}/Storage/{ID}/Drives/{ID}
Drive Location Indicator Active	PATCH	/redfish/v1/Systems/{ID}/Storage/{ID}/Drives/{ID}
Drive Write Cache Enabled	PATCH	/redfish/v1/Systems/{ID}/Storage/{ID}/Drives/{ID}
Volume Write Cache Policy	PATCH	/redfish/v1/Systems/{ID}/Storage/{ID}/Volumes/{ID}
Volume Read Cache Policy	PATCH	/redfish/v1/Systems/{ID}/Storage/{ID}/Volumes/{ID}
Volume Display Name	PATCH	/redfish/v1/Systems/{ID}/Storage/{ID}/Volumes/{ID}
Volume Dedicated Spare Drives	PATCH	/redfish/v1/Systems/{ID}/Storage/{ID}/Volumes/{ID}

For more details on Redfish management, see [HPE Storage Controllers Management Overview](#).

StorCLI

The Storage Command Line Interface (StorCLI) tool is the command line management software designed for the HPE MR Gen11 Controller. StorCLI is a command line interface that is designed to be easy to use, consistent, and easy to script.

Obtain StorCLI through the 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 StorCLI, see [HPE MR StorCLI User Guide](#).

UEFI System Utilities

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

- Configuring system devices and installed options
- Enabling and disabling system features
- Displaying system information
- Selecting the primary boot controller
- Configuring memory options
- Selecting a language
- Launching other pre-boot environments such as the Embedded UEFI Shell and Intelligent Provisioning

For more information on the UEFI System Utilities, see [UEFI System Utilities User Guide for HPE ProLiant Gen11 Servers and HPE Synergy](#).

For onscreen help, press **F1**.

Subtopics

[Using UEFI System Utilities](#)

Using UEFI System Utilities

To use the System Utilities, use the following keys.

Action	Key
Access System Utilities	F9 during server POST
Navigate menus	Up and Down arrows
Select items	Enter
Save selections	F10
Access Help for a highlighted configuration option ¹	F1

¹ Scan the QR code on the screen to access online help for the UEFI System Utilities and UEFI Shell.

Default configuration settings are applied to the server at one of the following times:

- Upon the first system power-up
- After defaults have been restored

Default configuration settings are sufficient for typical server operations. However, you can modify configuration settings as needed. The system prompts you for access to the UEFI System Utilities each time the system is powered up.

Configuration in UEFI System Utilities

This section contains information about using the Smart Array Configuration utility within the UEFI System Utilities to manage the controller.

For more information about the options mentioned in this section, see other sections in this guide and [HPE MR Storage Administrator User Guide](#).

Subtopics

[Viewing controller information and performing common actions](#)

[Configuration management](#)

[Controller management](#)

[Advanced controller management](#)

[Configuring advanced controller properties](#)

[Logical drive management](#)

[Drive management](#)

Viewing controller information and performing common actions

About this task

Use the Dashboard View screen to view the properties of the controller, view the server profile, and perform common actions.

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11.
2. In the Dashboard View panel, click Main Menu to access the main menu.
3. Click Help to view the help content.
4. View the properties.



Setting	Definition
Status	Status of the controller.
Backplane	Number of backplanes connected to the controller.
Enclosure	Number of enclosures connected to the controller.
Drives	Number of drives connected to the controller.
Arrays	Number of arrays on the controller.
Logical Drives	Number of logical drives on the controller.
ROC Temperature (C)	The temperature of the ROC.
View Server Profile	Shows the UEFI spec version that this system supports and menu options such as Controller Management, Hardware Components, Drive Management, and Logical Drive Management.

5. Perform the common actions, as needed:
 - [Configuration management](#)
 - [Setting factory defaults](#)
6. View any background operations in progress.
7. View the status of the advanced software options.



NOTE:

In this version of the controller, all advanced software options are enabled.

8. Optional. Click [Manage MegaRAID Advanced Software Options](#) to view the list of activated advanced software options.

Configuration management

Subtopics

[Creating a logical drive](#)

[Creating a profile-based logical drive](#)

[Importing secured foreign drive](#)

[Viewing array properties](#)

[Viewing global spare drives](#)

[Making a JBOD](#)

[Making an unconfigured good drive](#)

[Clearing a configuration](#)

Creating a logical drive

About this task



WARNING: Creating a logical drive will permanently delete any data on an associated drive.

Only Unconfigured Good drives can be used in a logical drive. If you have JBOD drives, you can choose to convert them to Unconfigured Good drives for use in your logical drive.

Procedure

1. From the System Utilities screen, select [System Configuration](#) > [HPE MRXXX Gen11](#) > [Main Menu](#) > [Configuration Management](#) > [Create Logical Drive](#).

2. If you have JBOD drives, they appear in the Create Logical Drive screen. Decide whether you want to convert them to JBOD before proceeding. To proceed without converting JBOD drives to Unconfigured Good drives, click Skip.
3. In the Create Logical Drive screen, select from the following options.

Option	Description
RAID level	Depending on the number of drives you have available, select RAID0, RAID1, RAID5, RAID6, RAID10, RAID50, or RAID60.
Data Integrity	Not supported
Select Drives From	Enables the physical drive selection option: <ul style="list-style-type: none"> • Free Capacity uses unused (free) drive capacity that is already part of a logical drive. • Unconfigured Capacity creates a logical drive on unconfigured drives.
Select Drive	View steps here .
Logical Drive Name	Name of the logical drive
Logical Drive Size	Shows and/or assigns the amount of logical drive storage space. By default, it assigns the maximum capacity available for the logical drive. <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p>NOTE:</p> <p>Only three decimal places are accepted.</p> </div>
Logical Drive Size Unit	Shows the amount of logical drive storage space in megabytes, gigabytes, and terabytes.
Strip size	Selects the strip element size for the logical drive. Drive striping involves partitioning each drive storage space into strips of different sizes that the controller firmware supports.
Read cache policy	Selects the read cache policy for the logical drive. The read cache policy of a logical drive determines how the controller handles reads to that logical drive. Possible settings are: <ul style="list-style-type: none"> • No read ahead - Specifies that the controller does not use read ahead for the current logical drive. • Read ahead - Allows the controller to read sequentially ahead of requested data and store the additional data in cache memory, anticipating that the data is required soon.
I/O Policy	Shows and/or selects the I/O policy for the logical drive. The possible settings are Direct and Cached.
Access policy	Shows and/or selects the access policy for the logical drive. The possible settings are Read/Write, Read Only, or Blocked.
Drive cache	Drive cache setting - Can be set to Unchanged, Enable, or Disable.
Disable background initialization	Background initialization status: <ul style="list-style-type: none"> • No - Default leaves the background initialization enabled. This option means that a new configuration can be initialized in the background while you use the app to do other configuration tasks. • Yes - Disables background initialization for configurations on this controller.
Default initialization	Logical drive initialization setting. Possible options are: <ul style="list-style-type: none"> • No - Do not initialize the logical drive. • Fast - Initializes the first 100 MB on the logical drive. • Full - Initializes the entire logical drive.
Emulation type	Specifies the emulation type policy for the logical drive. Possible options are Default, Disable, and Force.

4. Click Save Configuration.



5. Review and address any warning messages, as needed.
6. Click Confirm and Yes to proceed.

Subtopics

Selecting drives to include in a logical drive

Selecting drives to include in a logical drive

About this task

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

Procedure

1. In the Select Drives screen, specify the type of drive by selecting an option from the Select Media Type drop-down menu.
Options include SSD, HDD, or Both.
2. Select the interface type from the Select Interface Type drop-down menu.
Options include SAS, SATA, NVMe or All.
3. Specify the sector size from the Logical Sector Size drop-down menu.
Options include 512 B, 4 KB, or Both.
The unconfigured drives that match the criteria you specified are listed.
4. Select one or more unconfigured drives.
5. Click Apply Changes.
The Success screen appears to confirm that the selection was performed successfully.
6. Click OK.
The Create Logical Drive screen appears.

Creating a profile-based logical drive

About this task

 **WARNING:** Creating a logical drive will permanently delete any data on an associated drive.

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Configuration Management > Create Profile Based Logical Drive.
2. In the Create Profile Based Logical Drive screen, select a RAID level.
3. In the Generic R<level> screen, select the drive criteria.
4. View the profile parameters.

Parameter	Description
Logical Drive Name	Name of the logical drive
RAID Level	RAID level-based on the profile selected.
Logical Drive Size	Amount of logical drive storage space. By default, the maximum capacity available for the logical drive is displayed.
Strip Size	Stripe element size for the logical drive. Drive striping involves partitioning each drive storage space into strips of different sizes the controller firmware supports.
Read Cache Policy	Read cache policy for the logical drive. For any profile, if the drive is an SSD, the No Read Ahead option is displayed. If the drive is not SSD, the default option is displayed. Possible options are Read Ahead or No Read Ahead.
Write Cache Policy	Write cache policy for the logical drive. For any profile, if the drive is an SSD, the Write Through option is displayed. If the drive is not SSD, the default option is displayed. Possible options are Write Through or Write Back.
I/O Policy	I/O policy for the logical drive. For any profile, if the drive is an SSD, the Direct I/O option is displayed. If the drive is not SSD, the default option is displayed. Possible options are Direct I/O or Cached I/O.
Access Policy	Access policy for the logical drive. Possible settings are Read/Write, Read Only, or Blocked.
Drive Cache	Drive cache setting for the logical drive. Possible values are Unchanged, Enable, or Disable.
Default Initialization	Logical drive initialization setting. Possible options are No, Fast, and Full.

5. Click Save Configuration.
6. Review and address any warning messages, as needed.
7. Click Confirm and Yes to proceed.

Importing secured foreign drive

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Configuration Management.
2. In the Configuration Management screen, select Manage Foreign Configuration.
3. Select Preview Configuration.
All foreign logical drives are shown.
4. Select Import Foreign Configuration and do one of the following:
 - a. For Local Key Management (LKM), enter the security key for locked drives.
 - b. For External Key Management, restart to unlock the drive.

Viewing array properties

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Configuration Management > View Array Properties .
2. In the View Array Properties screen, for each array, view the following properties:
 - Capacity Allocation - Associated logical drives and available free capacity
 - Protected - Not currently supported

Viewing global spare drives

Prerequisites

- You must have previously created spare drives to view this option.

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Configuration Management > View Global Spare Drives.
2. In the View Global Spare Drives screen, view the list of spare drives.

Making a JBOD

About this task

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

Procedure

1. Select System Configuration > HPE MRXXX Gen11 > Main Menu > Configuration Management > Make JBOD.
2. In the Make JBOD screen, select the Unconfigured Good drive that you want to convert to JBOD.
3. Click OK.
4. Review and address any warning messages, as needed.
5. Click Confirm and Yes to proceed.

The Success screen appears to confirm that the operation is complete. For the drive status change to reflect on the Drive Management screen, refresh the view.

Making an unconfigured good drive

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Configuration Management .
2. In the Configuration Management screen, select Make Unconfigured Good from the Operation drop-down menu.
3. In the Make Unconfigured Good screen, select the JBOD drive that you want to convert to unconfigured good.
4. Click OK.
5. The Warning screen appears, click Confirm and then click Yes.

The Success appears to confirm that the operation is complete.

Clearing a configuration

About this task

 **WARNING:** Clearing a configuration deletes all the logical drives and spare drives attached to the controller.

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Configuration Management > Clear Configuration.
2. In the warning message screen, click Confirm and Yes to proceed.

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

Controller management

Subtopics

[Managing the controller](#)

Managing the controller

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Controller Management.
2. In the Controller Management screen, view the basic properties of the controller:

Property	Description
Product name	The name of the controller.
Serial number	Controller serial number.
Controller status	Possible options are Optimal, Needs Attention, Failed, or Safe Mode.
Select Boot Device	The selected primary bootable device. This setting is applicable for legacy boot mode only.
PCI ID	Manufacturer-assigned ID.
PCI Slot number	ID of the PCI Slot that contains the controller.
Active package version	Active version of the controller package.
Backup version	Backup version of the controller package.
PSOC firmware version	PSOC version.
Connector	Number of host data ports or connectors on the controller.
Drive count	Number of drives currently attached to the controller.
Logical drive count	Number of logical drives on the controller.

3. Click Advanced Controller Management to perform advanced controller tasks.
For more information, see [Advanced Controller Management](#).
4. Click Advanced Controller Properties to configure advanced controller settings.

Advanced controller management

Subtopics

[Clearing controller events](#)

[Saving controller events](#)

[Saving a serial log](#)

[Enabling drive security](#)

[Disabling drive security](#)

[Changing drive security settings](#)

[Changing drive security key management mode](#)

[Managing link speed](#)

[Managing advanced software options](#)

[Scheduling a consistency check](#)

Clearing controller events

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Controller Management.
2. In the Controller Management screen, select Advanced Controller Management.
3. In the Advanced Controller Management screen, select Clear Controller Events.

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

Saving controller events

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Controller Management.
2. In the Controller Management screen, select Advanced Controller Management.
3. In the Advanced Controller Management screen, select Save Controller Events.

4. In the Save Controller Events screen, select the file system for saving the logs.

5. Select the directory to save the logs.

The root directory is selected by default.

6. Specify a filename, with a .txt file extension.

7. Click Save Events.

The Success screen appears to notify that the operation has completed successfully.

Saving a serial log

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Controller Management.

2. In the Controller Management screen, select Advanced Controller Management.

3. In the Advanced Controller Management screen, select Save Serial Log.

4. In the Save Serial Log screen, select the file system for saving the serial log.

5. Select the directory for saving the log.

The current directory is the default directory.

6. Specify the filename for the log, using a `.txt` file extension.


7. Select the number of entries (in KB) to save in the log.

8. Click Save Log.

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

Enabling drive security

Procedure



1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Controller Management.
2. In the Controller Management screen, select Advanced Controller Management.
3. In the Advanced Controller Management screen, select Enable Drive Security.
4. For the Key Management Mode, select either Local Key Management (LKM) or External Key Management (EKM).
5. 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 Password at Boot Time.
 - b. Select the I Recorded the Security Settings for Future Reference check box.
 - c. Click Enable Drive Security.
6. 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 "Using Enterprise Secure Key Manager with iLO" in the [HPE iLO 6 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-up. It communicates to iLO through the BIOS.

Disabling drive security

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Controller Management.
2. In the Controller Management screen, select Advanced Controller Management.
3. In the Advanced Controller Management screen, select Disable Drive Security.

The warning message appears requesting confirmation.
4. Review the warning message, click Confirm and Yes to proceed.

Changing drive security settings

About this task

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

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

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Controller Management.
2. In the Controller Management screen, select Advanced Controller Management.
3. In the Advanced Controller Management screen, select Change Current Security Settings.
4. Change the required fields. You are allowed to change the Security Key Identifier, Security Key, and Password at Boot Time.
5. Click Save Security Settings.

Changing drive security key management mode

About this task

If the iLO key manager is properly configured, you are allowed to change the drive security mode from Local Key Management (LKM) to External Key Management (EKM).

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Controller Management.
2. In the Controller Management screen, select Advanced Controller Management.
3. In the Advanced Controller Management screen, select Switch to External Key Management (EKM) Mode , and click OK.
4. Enter the Current Security Key Identifier and Current Security Key, and click OK.
5. Restart the system for changes to take effect.

Managing link speed

About this task

You can view the link speed of the controller.

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Controller Management.
2. In the Controller Management screen, select Advanced Controller Management.
3. In the Advanced Controller Management screen, select Manage Link Speed.
4. In the Manage Link Speed screen, view the PHY settings for the controller.

Managing advanced software options

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Controller Management.
2. In the Controller Management screen, select Advanced Controller Management.
3. In the Advanced Controller Management screen, select Manage MegaRAID Advanced Software Options .
4. In the Manage MegaRAID Advanced Software Options screen, view the list of advanced options that are currently enabled.



NOTE: In this version of the controller, all of the advanced software options are activated.

Scheduling a consistency check

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Controller Management.
2. In the Controller Management screen, select Advanced Controller Management.
3. In the Advanced Controller Management screen, select Schedule Consistency Check .
4. In the Schedule Consistency Check screen, select the frequency.
5. Assign the start date for the consistency check.
6. Assign the start time for the consistency check.
7. Select the consistency check mode. The options are:
 - Sequential - Checks the logical drives one at a time.
 - Concurrent - Checks all the logical drives simultaneously.
8. Select the Start Immediately option to start the consistency check on all logical drives except any excluded drives.
9. Click Exclude Logical Drives to specify the logical drives to exclude from the consistency check.
10. Click Apply Changes.

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

Setting factory defaults

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Controller Management.
2. In the Controller Management screen, select Advanced Controller Management.
3. In the Advanced Controller Management screen, select Set Factory Defaults.
The warning message appears to request confirmation.
4. Click Confirm and Yes.

Managing auto-config

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Controller Management.
2. In the Controller Management screen, select Advanced Controller Management.
3. In the Advanced Controller Management screen, select Manage Personality Mode.
4. In the **Manage Personality Mode** screen, select the Auto-Configure Behavior. The options are:
 - a. **None:** If Auto-Config is set to None, the new drive status stays in Unconfigured Good.
 - b. **JBOD:** If Auto-Config is set to JBOD, the new drive status is automatically configured as JBOD.
5. Click Apply Changes.

More information

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

Configuring advanced controller properties

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Controller Management.
2. In the Controller Management screen, select Advanced Controller Properties .
3. In the Advanced Controller Properties screen, click Cache and Memory to configure the cache and memory settings.
For more information, see [Configuring cache and memory settings](#).
4. Click Patrol Read to configure the patrol read settings.
For more information, see [Configuring patrol read settings](#).
5. Click Spare to configure the spare settings.
For more information, see [Configuring spare settings](#).
6. Click Task Rates to configure the task rates settings.
For more information, see [Configuring Task Rates](#).
7. Modify any of the following controller properties:



Property	Description
Auto Import Foreign Configuration	Enables or disables the automatic import of foreign configurations without any user intervention
Coercion Mode	Drive coercion forces the drives of varying capacities to the same size so they can be used in an array. The coercion mode options are None, 128 MB, and 1 GB.
Boot Mode	Specifies options to handle errors (if they occur) during boot time. Possible options are: <ul style="list-style-type: none"> • Stop on error - Shows error and waits for your input • Pause on error - Noncritical errors show up briefly, giving you a choice to take action before the firmware proceeds with the boot. Critical errors wait for your response and do not proceed with the boot. • Ignore errors - Errors are ignored and the firmware proceeds with the boot. • Safe mode on errors - The application loads but with a limited set of operations. You cannot use this adapter as a boot adapter while in safe mode.
Controller BIOS	Enables or disables the controller BIOS. If the boot device is on the selected RAID controller, the BIOS must be enabled.
ROC Temperature (C)	Temperature of the ROC. (read only)
Shield State Supported	Whether shield state is supported in the controller or not supported.
Drive Security	Status of the drive security (encryption) feature on the controller.
T10-PI	Status of the data protection feature on the controller.
Extended Logical Drive Support	Allows you to enable or disable the controller's max logical drive limit. Any change takes effect after the system reboots.
Maintain Drive Fail History	Enables you to track the bad drives.
SMART Polling	Determines how frequently the controller polls for drives reporting a Predictive Drive Failure. The default is 300 seconds.
Stop Consistency Check on Error	Enables or disables the option to stop the consistency check operation on a redundant logical drive if there is an inconsistency found in the data.
JBOD Mode	Enables or disables the JBOD mode.
Write Verify	Enables or disables the write verification during the cache flush.
Large I/O Support	Enables or disables large I/O support. Any changes take effect after a system reboot.

8. Click Apply Changes.

Subtopics

[Configuring cache and memory settings](#)

[Configuring patrol read settings](#)

[Configuring spare settings](#)

[Configuring Task Rates](#)

Configuring cache and memory settings

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Controller Management.
2. In the Controller Management screen, select Advanced Controller Properties .
3. In the Advanced Controller Properties screen, click Cache and Memory to configure the cache and memory settings.
4. In the Cache and Memory screen, update the cache flush interval as needed.

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.

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 Gen11 > Main Menu > Controller Management.
2. In the Controller Management screen, select Advanced Controller Properties .
3. In the Advanced Controller Properties screen, click Patrol Read to configure the patrol read settings.
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 fields.

5. Configure the mode using the Mode drop-down menu:
 - Auto - Patrol read runs continuously on the system based on a schedule.
 - Manual - Enables patrol read to be started or stopped manually.
 - Disabled - Patrol read operation is disabled.
6. Configure the rate.

Rate is the percentage of system resources dedicated to perform a patrol read operation on configured drives.
7. Configure the setting for unconfigured space.
8. Click Apply Changes.

Configuring spare settings

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Controller Management.
2. In the Controller Management screen, select Advanced Controller Properties .
3. In the Advanced Controller Properties screen, click Spare to configure the spare settings.
4. In the Spare screen, enable or disable the following settings:
 - Persistent Spare - The option to have the system backplane or the storage enclosure drive slots dedicated as spare slots. If enabled, the replacement of a spare drive in the same slot will automatically configure the drive as a spare.
 - Replace Drive - The option to copy data from a spare drive to a drive.
 - Replace Drive on SMART Error - The option to start a replace operation if a SMART error is detected on a drive.
5. Click Apply Changes.

Configuring Task Rates

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Controller Management.
2. In the Controller Management screen, select Advanced Controller Properties .
3. In the Advanced Controller Properties screen, click Task Rates to configure the Task Rates.
4. In the Spare screen, configure the following task rates:



- Background initialization (BGI) Rate - The percentage of system resources dedicated to perform a background initialization on a redundant logical drive.
- Consistency Check Rate - The percentage of system resources dedicated to perform a consistency check operation on a redundant logical drive.
- Patrol Read Rate - The percentage of system resources dedicated to perform a Patrol Read operation on configured drives.
- Rebuild rate - The percentage of system resources dedicated to rebuild data on a new drive after a storage configuration drive has failed.
- Transformation Rate - The percentage of system resources dedicated to perform a RAID level migration or online capacity expansion on a logical drive.

5. Click Apply Changes.

Logical drive management

Subtopics

[Viewing and configuring properties of a logical drive](#)

[Deleting a logical drive](#)

[Initializing a logical drive](#)


[Locating a physical drive associated with a logical drive](#)

[Erasing a logical drive](#)

Viewing and configuring properties of a logical drive

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Logical Drive Management.
2. In the Logical Drive Management screen, select the logical drive.
3. View the basic properties of the logical drive:
 - Name - View or modify the name of the logical drive.
 - Raid Level - RAID level of the logical drive.
 - Status - Current status of the logical drive.
 - Size - Size of the logical drive in MB/GB/TB
4. To view the drives associated with the logical drive, click View Associated Drives. See [Viewing drive properties](#).
5. To view the advanced logical drive properties, click Advanced....



Property or policy	Description
Logical sector size	Logical sector size of the logical drive. Possible options are 4 KB or 512 B.
Strip size	Strip element size for the logical drive.
Protected	Whether or not the logical drive uses protection.
Bad blocks	Whether the logical drive has bad blocks.
SSD Caching	Whether or not SSD caching is enabled on this logical drive
Access	Specifies the access policy for the logical drive. Possible settings are Read/Write, Read Only, or Blocked.
Write Cache Status	Current write cache status for the logical drive. Possible settings are Enabled, Disabled, and Temporarily Disabled.
Current Write Cache Policy	Current write cache policy for the logical drive. Possible settings are Write Through, Write Back, and Always Write Back. Any change takes effect after the system reboots.
Default Write Cache Policy	Specifies the default write cache policy for the logical drive. Possible settings are Write Through, Write Back, and Always Write Back. Any change takes effect after the system reboots.
Disable Background Initialization (BGI)	Disables or enables background initialization. Selecting No enables background initialization and allows a new configuration to be initialized in the background while you use the application to perform other configuration tasks. This is the default option. Select Yes if you do not want to allow background initialization for configurations on this controller.
Read Cache Policy	Specifies the read cache policy for the logical drive. Possible settings are No Read Ahead and Read Ahead.
Drive Cache	Specifies the drive cache policy for the logical drive. Possible settings are Unchanged, Enable, and Disable.
Input/Output	Specifies the I/O policy for the logical drive. Possible settings are Direct and Cached.
Emulation Type	Specifies the emulation type policy for the logical drive. Possible settings are default, Disable, and Force.

Deleting a logical drive

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Logical Drive Management.
2. In the Logical Drive Management screen, select the logical drive.
3. In the Logical Drive screen, select Delete Logical Drive from the Operation drop-down menu.
4. Click Go.
5. Review warning message, and click Confirm and Yes to proceed.

Initializing a logical drive

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Logical Drive Management.
2. In the Logical Drive Management screen, select the logical drive.
3. In the Logical Drive screen, select Fast Initialization or Slow Initialization from the Operation drop-down menu.
4. Click Go.
5. Review warning message, and click Confirm and Yes to proceed.



Locating a physical drive associated with a logical drive

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Logical Drive Management.
2. In the Logical Drive Management screen, select the logical drive.
3. In the Logical Drive screen, select Start Locate from the Operation drop-down menu.
4. Click Go.

The LEDs for the drive associated with the logical drive begin blinking.

To stop the LEDs from blinking, select Stop Locate from the Operation drop-down menu.

Erasing a logical drive

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Logical Drive Management.
2. In the Logical Drive Management screen, select the logical drive.
3. In the Logical Drive screen, select Logical Drive Erase from the Operation drop-down menu.
4. Select one of the following options from the Erase Mode drop-down menu:
 - Simple - Writes a pattern to the logical drive 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.
5. If you want to delete the logical drive after the erase is completed, click Delete After Erase.
6. Click Go.
7. Review warning messages and click Confirm and Yes to proceed.

Drive management

Subtopics

[Viewing drive properties](#)

[Locating a drive](#)

[Initializing a drive](#)

[Erasing a drive](#)

[Making a JBOD](#)

[Making an unconfigured good drive](#)

[Making a bootable drive](#)

[Assigning a global spare drive](#)

[Unassigning a global spare drive](#)

[Sanitizing an unconfigured good drive](#)

Viewing drive properties

Procedure



1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Drive Management.
2. In the Drive Management screen, select the drive.
3. In the drive screen, view the basic properties:

Property	Description
Drive ID	ID of the drive
Status	Current drive status
Size	Size of the drive in MB/GB/TB
Type	Device type
Model	Model number of the drive
Hardware Vendor	Hardware manufacturer
Associated Logical Drive	List of logical drives associated with the drive

4. To view advanced properties, click Advanced....

Property	Description
Certified	Whether the selected drive is vendor-certified.
Logical Sector Size	Logical sector size of this drive. Possible options are 4 KB or 512 B.
Physical Sector Size	Physical sector size of this drive. Possible options are 4 KB or 512 B.
SMART Status	Self-Monitoring Analysis and Reporting Technology (SMART) status of a drive. This feature monitors the internal performance of all motors, heads, and drive electronics to detect predictable drive failures.
Revision	Firmware revision of the drive.
Media Errors	Physical errors that are detected on the disk media.
SAS Address	World Wide Name (WWN) for a drive.
Drive Power State	Indicates the power condition (On or Power Save) of the drive.
Cache Setting	Disk cache setting of the drive.
Available Size	Available size of the drive.
Used Space	Configured space of the drive.
Disk Protocol	Type of hard disk drive used.
Device Speed	Speed of the physical disk.
Negotiated Drive Transfer Speed	Negotiated link speed for the data transfer to/from the drive.
Number of Connections	Number of connections of the drive.
FDE Capable	Whether the drive is encryption capable.
Data Integrity Capable	Whether the drive is capable of protection.
Temperature (C)	Temperature of the drive.

Locating a drive

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Drive Management.
2. In the Drive Management screen, select the drive that you want to locate.
3. In the Drive screen, select Start Locate or Stop Locate from the operation drop-down menu.

4. Click Go.

The Success screen appears to confirm that the operation is successful.

The LEDs on the drive starts blinking.

Initializing a 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 Gen11 > Main Menu > Drive Management.
2. In the Drive Management screen, select the drive that you want to initialize.
3. In the drive screen, select Initialize Drive from the Operation drop-down menu.
4. Click Go.

The Success screen appears to confirm that the operation is successful.

5. Review the warning message and click Confirm and Yes to proceed.

Subtopics

Placing a drive offline

Placing a drive offline

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Drive Management.
2. In the Drive Management screen, select the online drive that you want to place offline.
3. In the Drive screen, select Place Drive Offline from the Operation drop-down menu.
4. Click Go.
5. Review the warning message and click Confirm and Yes to proceed.

Erasing a drive

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Drive Management.
2. In the Drive Management screen, select the drive that you want to erase.
3. In the Drive screen, select Drive Erase from the Operation drop-down menu.
4. Select one of the following options from the Erase Mode drop-down menu:
 - Simple - Writes a pattern to the logical drive 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.
5. Click Go.
6. Review the warning message and click Confirm and Yes to proceed.

Making a JBOD

About this task

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

Procedure

1. Select System Configuration > HPE MRXXX Gen11 > Main Menu > Drive Management.
2. In the Drive Management screen, select the Unconfigured Good drive that you want to convert to JBOD.
3. In the drive screen, select the Make JBOD from the Operation drop-down menu.
4. Click GO.

The Success screen appears to confirm that the operation is complete. For the drive status change to reflect on the Drive Management screen, refresh the view.

Making an unconfigured good drive

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Drive Management .
2. In the Drive Management screen, select the JBOD drive that you want to convert to Unconfigured Good.
3. In the drive screen, select Make Unconfigured Good from the Operation drop-down menu.
4. Click Go.

The Success screen appears to confirm that the operation is complete. For the drive status change to reflect on the Drive Management screen, refresh the view.

Making a bootable drive

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Drive Management.
2. In the Drive Management screen, select the drive that you want to use as a bootable drive.
3. In the Drive screen, select Make Bootable Drive from the Operation drop-down menu.
4. Click Go.

The Success screen appears to confirm that the operation completed successfully. For the status change to appear in the Drive Management screen, refresh the view.

5. Review the warning message and click Confirm and Yes to proceed.

Assigning a global spare drive

Procedure

1. From the System Utilities screen, select System Configuration > HPE MRXXX Gen11 > Main Menu > Drive Management.
2. In the Drive Management screen, select the drive that you want to use as a global spare drive.
3. In the drive screen, select Assign Global Spare Drive from the Operation drop-down menu.
4. Click Go.

The Success screen appears to confirm that the operation completed successfully.

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

Unassigning a global spare drive

Procedure

1. From the System Utilities screen, select **System Configuration > HPE MRXXX Gen11 > Main Menu > Drive Management**.
2. In the Drive Management screen, select the spare drive that you want to unassign.
3. In the drive screen, select **Unassign Global Spare Drive** from the Operation drop-down menu.
4. Click **Go**.

The Success screen appears to confirm that the operation completed successfully.

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

Sanitizing an unconfigured good drive

About this task

Cryptographic erase is supported on SSDs, and some HDDs. To understand if your drive supports crypto erase, see [HPE MR Gen11 Storage Controllers QuickSpecs](#).



WARNING:

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

Procedure

1. From the System Utilities screen, select **System Configuration > HPE MRXXX Gen11 > Main Menu > Drive Management**.
2. In the Drive Management screen, select the JBOD drive that you want to convert to **Unconfigured Good**.
3. In the drive screen, select **Make Unconfigured Good** from the Operation drop-down menu. Click **Go**, and follow the instructions on the screen to proceed.
4. Click **Advanced** and check the **Cryptographic Erase Capable** field shows **Yes** to make sure you can perform cryptographic erase on the drive.
5. In the drive screen, select **Cryptographic Erase** from the Operation drop-down menu. Click **Go**, and follow the instructions on the screen to proceed.

Maintenance

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. For system software and firmware updates, download the Service Pack for ProLiant (SPP) from the Hewlett Packard Enterprise website <https://www.hpe.com/servers/spp/download>. For information about the SPP, see the product documentation at the [Hewlett Packard Enterprise website](#).

Hewlett Packard Enterprise now distributes drivers and other support software for servers and server blades through SPP, which you can download from <https://www.hpe.com/servers/spp/download>. Be sure to use the latest SPP version for the server or server blade.

If you installed an OS by using the Intelligent Provisioning software, its configure and install feature may have provided the latest driver support.

After you update HPE MR Storage Administrator, clear the browser cache.

Error reporting

- **Integrated Management Log (IML)**

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 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 iLO SNMP management function. The most common SNMP traps include:

<code>cpqDa6CntlRStatusChange</code>	Controller status change
<code>cpqDa6LogDrvStatusChange</code>	Logical drive status change
<code>cpqDa7PhyDrvStatusChange</code>	Drive status change
<code>cpqDa6AccelStatusChange</code>	Accelerator status change
<code>cpqDa6AccelBadDataTrap</code>	Accelerator bad data

For information on configuring iLO SNMP traps and a full description of supported SNMP traps, see the [HPE iLO 6 User Guide](#).

- **Redfish alerts**

The controller supports sending alerts through the iLO Redfish API. These alerts are defined in the <https://redfish.dmtf.org/registries/StorageDevice.1.1.0.json> message registry. The Redfish alerts include:

- ControllerFailure
- DriveOK
- DrivePredictiveFailure
- DriveFailure
- DriveFailureCleared
- DriveInserted
- DriveRemoved
- VolumeOK
- VolumeDegraded
- VolumeFailure
- WriteCacheProtected
- WriteCacheTemporarilyDegraded
- WriteCacheDegraded
- WriteCacheDataLoss

- **Application event log**

HPE MR Storage Administrator reports array events to the application area inside of the Microsoft Windows system event log, which includes detailed diagnostic information of the most recent events encountered by the controller.

On Linux, the system event log is located at `/var/log/messages`.

On VMware, the system event log is located at `/var/log/vmkernel.log`.

- **System event log**

The controller uses PLDM over Redfish device enablement to send Redfish events, which result in error messages in the HPE IML. These events are generated by the device firmware and 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`.

For more information, see <https://www.hpe.com/info/SCMO>.

Diagnostic tools

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

- **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**.

- **StorCLI**

If you cannot use MRSA, use StorCLI to collect the required logs to analyze issues.

Run `storcli64 show` to get controller index first, then run the following commands, where x means the controller index.

Command	Description
<code>storcli64 /cx show events file=events.log</code>	All events
<code>storcli64 /cx show events type=sincereboot file=SBevents.log</code>	Events since last reboot
<code>storcli64 /cx show all > show_all.log</code>	General device information
<code>storcli64 /cx show termlog > show_termlog.log</code>	Firmware log
<code>storcli64.exe /cx get snapdump</code>	Capture snapdump

Subtopics

[Troubleshooting resources](#)

Troubleshooting resources

Troubleshooting resources are available for HPE Gen11 server products in the following documents:

- **Troubleshooting Guide for HPE ProLiant Gen11 servers** provides procedures for resolving common problems and comprehensive courses of action for fault isolation and identification, issue resolution, and software maintenance.

<https://www.hpe.com/info/gen11-troubleshooting>

- **Integrated Management Log Messages for HPE ProLiant Gen10, Gen10 Plus, and Gen11 servers and HPE Synergy** provides IML messages and associated troubleshooting information to resolve critical and cautionary IML events.

<https://www.hpe.com/info/Troubleshooting-IML-en>

Models

Subtopics

[Modular controller \(-o\)](#)

[Standup PCIe Plug-In Controller \(-p\)](#)


Modular controller (-o)

Subtopics

[HPE MR216i-o Gen11 SPDM Storage Controller ports and connectors](#)


[HPE MR408i-o Gen11 SPDM Storage Controller ports and connectors](#)

HPE MR216i-o Gen11 SPDM Storage Controller ports and connectors



Item	Description
1	Internal x8 LP SlimSAS Connector, Port 1i
2	Internal x8 LP SlimSAS Connector, Port 2i


HPE MR408i-o Gen11 SPDM Storage Controller ports and connectors



Item	Description
1	Internal x8 LP SlimSAS Connector, Port 1i
2	Backup power cable connector

HPE MR416i-o Gen11 SPDM Storage Controller ports and connectors





Item	Description
1	Internal x8 LP SlimSAS Connector, Port 1i
2	Internal x8 LP SlimSAS Connector, Port 2i
3	Backup power cable connector

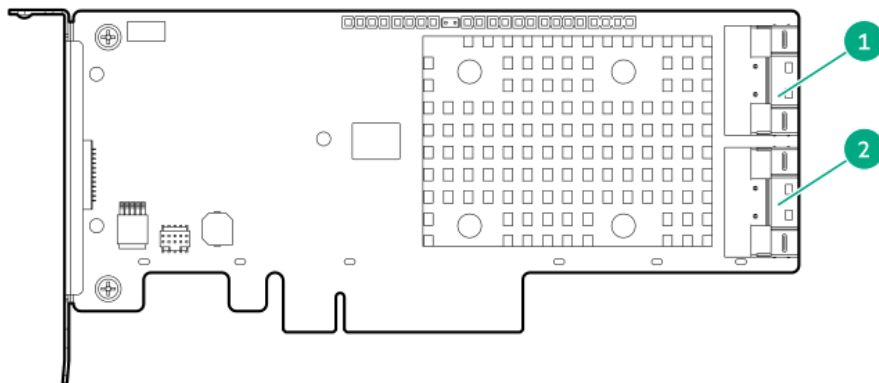
Standup PCIe Plug-In Controller (-p)

Subtopics

[HPE MR216i-p Gen11 SPDM Storage Controller ports and connectors](#)

[HPE MR416i-p Gen11 SPDM Storage Controller ports and connectors](#)


HPE MR216i-p Gen11 SPDM Storage Controller ports and connectors



Item	Description
1	Internal x8 SlimSAS Connector, Port 1i
2	Internal x8 SlimSAS Connector, Port 2i



HPE MR416i-p Gen11 SPDM Storage Controller ports and connectors



Item	Description
1	Internal x8 SlimSAS Connector, Port 1i
2	Internal x8 SlimSAS Connector, Port 2i
3	Backup power cable connector

Additional hardware and options

Subtopics

[Energy pack options](#)

Energy pack options

Hewlett Packard Enterprise offers a variety of intelligent energy packs ranging in cell chemistry, power output, and cable lengths to fit within the broad range of servers. The centralized energy pack supports flash-backed write cache (FBWC) or SmartCache across storage controllers. It is installed at the front of the server and plugs into a 14-pin (2x7) connector on the server motherboard. Power is routed through the motherboard and PCI risers and delivered to the storage controller using a three-pin controller backup power cable. The power cable is included with the purchase of supported storage controllers.

The health of energy pack can be viewed using any of the following options:

- In iLO GUI, select **Power & Thermal >Power** and check statistics in **Smart Storage Energy Pack** panel.
- Along with the controller GUI & CLI tools, you can check iLO Redfish `CacheSummary` property under the `StorageController` resource.
- iLO Redfish `SmartStorageBattery` property under `Chassis` resource.

Upon starting, the storage controllers exchange information with HPE BIOS during POST and monitor the voltage received from the three-pin controller backup power cable.



NOTE:

The energy pack must only be installed, removed, or replaced while the server is powered off and AC power cords are removed.

Subtopics

[HPE Smart Storage Battery](#)

[HPE Smart Storage Hybrid Capacitor](#)

HPE Smart Storage Battery

HPE Smart Storage Battery is an optional lithium-ion, low-halogen centralized backup source. It supports unlimited number of devices of 96W battery or two devices of 12W battery. The time required to recharge is two hours for 96W battery and one hour for 12W battery. For more information, see [HPE Smart Storage Batteries and Hybrid Capacitors QuickSpecs](#).

HPE Smart Storage Hybrid Capacitor

HPE Smart Storage Hybrid Capacitor is a battery-free technology for power storage while eliminating the environmental impact of lithium-ion batteries. It supports up to three MR416 storage controllers. The time required to recharge takes less than one minute. For more information, see [HPE Smart Storage Batteries and Hybrid Capacitors QuickSpecs](#).

Storage reference

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.



Websites

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

Hardware RAID controllers

Content	Microchip® SmartROC 3100	Microchip® SmartROC 3200	Broadcom® Aero-16
User guides - Gen10 and Gen10 Plus	https://www.hpe.com/support/SR-Gen10-UG	https://www.hpe.com/support/SR-Gen10Plus-UG	https://www.hpe.com/support/MR-Gen10Plus-UG
User guides - Gen11	—	https://www.hpe.com/support/SR-Gen11-UG	https://www.hpe.com/support/MR-Gen11-UG
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
QuickSpecs - Gen11	—	https://www.hpe.com/psnow/doc/a50004312enw	https://www.hpe.com/psnow/doc/a50004311enw
GUI User Guide	https://www.hpe.com/support/SSA-UG	https://www.hpe.com/support/SSA-UG	https://www.hpe.com/support/MRSA
CLI User Guide	https://www.hpe.com/support/SSACLI-UG	https://www.hpe.com/support/SSACLI-UG	https://www.hpe.com/support/StorCLI

Boot devices and Virtual RAID

Content	HPE NS204 boot devices	Intel VROC
User guides - Gen10 Plus and Gen11	https://www.hpe.com/support/NS204-UG	https://www.hpe.com/support/IntelVROC-Gen10Plus-docs https://hpe.com/support/VROC-Gen11-UG https://www.intel.com/vroc
QuickSpecs - Gen10 Plus and Gen11	https://www.hpe.com/psnow/doc/a00094638enw	https://www.hpe.com/psnow/doc/a50002570enw

Technical paper and other references

Content	Links
Encryption Overview	https://www.hpe.com/info/SCEO
Management Overview	https://www.hpe.com/info/SCMO
iLO 5 User Guide	https://www.hpe.com/support/ilo5-ug-en
iLO6 User Guide	https://www.hpe.com/support/ilo6-ug-en
UEFI Gen10 and Gen10 Plus User Guide	https://www.hpe.com/support/UEFIGen10-UG-en
UEFI Gen11 User Guide	https://www.hpe.com/support/UEFIGen11-UG-en

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/
Boot Windows OS from Intel VROC RAID volume	https://hpedemoportal.ext.hpe.com/
How to Manage HPE SRxxx Gen10+ Storage Controllers	https://youtu.be/NsoDI9-FheU
How to Manage HPE MRxxx Gen10+ Storage Controllers	https://youtu.be/Xh5FA8YjgRk
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 RAID1 Volume	https://hpedemoportal.ext.hpe.com/
Management of Redfish Device Enabled Storage Controllers	https://youtu.be/Ju-r-xhfzKU
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/

For additional websites, see [Support and other resources](#).

Support and other resources

Subtopics

[Accessing Hewlett Packard Enterprise Support](#)

[Accessing updates](#)

[Remote support](#)

[Warranty information](#)

[Regulatory information](#)

[Documentation feedback](#)

Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:
<https://www.hpe.com/info/assistance>
- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:
<https://www.hpe.com/support/hpesc>

Information to collect

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

- Add-on products or components
- Third-party products or components

Accessing updates

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

- To download product updates:

Hewlett Packard Enterprise Support Center

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

My HPE Software Center

<https://www.hpe.com/software/hpesoftwarecenter>

- To subscribe to eNewsletters and alerts:

<https://www.hpe.com/support/e-updates>

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

<https://www.hpe.com/support/AccessToSupportMaterials>

IMPORTANT:

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

Remote support

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

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

HPE Get Connected

<https://www.hpe.com/services/getconnected>

HPE Tech Care Service

<https://www.hpe.com/services/techcare>

HPE Complete Care

<https://www.hpe.com/services/completecure>

Warranty information

To view the warranty information for your product, see the [warranty check tool](#).

Regulatory information

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

<https://www.hpe.com/support/Safety-Compliance-EnterpriseProducts>

Additional regulatory information

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

<https://www.hpe.com/info/reach>

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

<https://www.hpe.com/info/ecodata>

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

<https://www.hpe.com/info/environment>

Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, use the Feedback button and icons (at the bottom of an opened document) on the Hewlett Packard Enterprise Support Center portal (<https://www.hpe.com/support/hpesc>) to send any errors, suggestions, or comments. This process captures all document information.

