



Hewlett Packard
Enterprise

HPE SR Gen11 Controller User Guide

Part Number: 30-32ECA3FF-005

Published: October 2024

Edition: 5

HPE SR Gen11 Controller User Guide

Abstract

This document includes feature, installation, and configuration information about HPE SR Gen11 controllers 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: 30-32ECA3FF-005

Published: October 2024

Edition: 5

© 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

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 the property of their respective owners.

Table of contents

- HPE SR Gen11 Controller
 - 400 series
 - 900 series
- Supported features for HPE SR Gen11 Controller
 - Features support
 - OS 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 HBA simultaneously)
 - Striping
 - RAID 0
 - Mirroring
 - RAID 1 and RAID 1+0 (RAID 10)
 - RAID 1 (Triple) and RAID 10 (Triple)
 - Read load balancing
 - Mirror splitting and recombining
 - Parity
 - RAID 5
 - RAID 50
 - RAID 6
 - RAID 60
 - Parity groups
 - Background parity initialization
 - Rapid parity initialization
 - Regenerative writes
 - Backed-out writes
 - Full-stripe writes
 - Spare drives
 - Dedicated spare

- Predictive spare activation
- Failure spare activation
- Auto-replace spare
- Drive rebuild
 - Rapid rebuild
 - Puncture
 - Rebuild priority
 - Before replacing drives
- Transformation
 - Array transformations
 - Expand array
 - Move array
 - Replace array
 - Shrink array
 - Mirror array
 - Heal array
 - Volume transformations
 - Expand volume
 - Migrate RAID level
 - Migrate strip size
 - Transformation priority
- Drive technology
 - Predictive drive failure
 - Online drive firmware update
 - Dynamic sector repair
 - Controller surface scan
 - Shingled magnetic recording
 - Hot-plug drive LED
 - SSD over-provisioning optimization
 - SSD wear gauge reports
- Security
 - Controller-based encryption
 - Local Key management mode
 - Remote Key management mode
 - Self-Encrypting Drive
 - Host Key management
 - Local Key management
 - Remote Key management
 - Sanitize erase
 - Sanitize overwrite
 - Sanitize block erase

- Sanitize crypto erase
 - Sanitize freeze lock
 - Signed firmware
 - Hardware-based Root of Trust
 - Secure Boot
 - Security Protocol and Data Model
- Reliability
 - Link error monitoring
 - Recovery ROM
 - Cache error checking and correction
 - Thermal monitoring
- Performance
 - HPE SR SmartCache
 - Elements of HPE SR SmartCache
 - HPE SR SmartCache write policy and RAID type
 - HPE SR SmartCache line size
 - HPE SR SmartCache volume capacity
 - Features and benefits of HPE SR SmartCache
 - Features not supported with HPE SR SmartCache
 - Cache
 - Read cache
 - Flash-backed write cache
 - Cache ratio selection
 - Write cache bypass threshold
 - No-battery write cache
 - Drive write cache control
 - I/O performance mode
 - Video on demand
 - Strip size selection
 - Power modes
- Installation instructions for HPE SR Gen11 Controller
 - Supported servers
 - Installing in an unconfigured server
 - Installing in a previously configured server
 - Installing an HPE Synergy mezzanine controller (-m)
 - Installing a standup PCIe plug-in controller (-p)
 - Connecting storage devices
 - Connecting internal storage
 - Cable part numbers
- Configuration of HPE SR Gen11 Controller
 - Configuration tools for HPE SR controllers

- Array and controller configuration
 - Comparison of SSA and UEFI System Utilities
- HPE iLO GUI
- Intelligent Provisioning
- HPE OneView
- Redfish
 - DMTF Redfish Storage Model
- Smart Storage Administrator
- UEFI System Utilities
 - Using UEFI System Utilities
- Powering on and selecting boot options
- Maintenance
 - Updating software and firmware
 - Error reporting
 - Diagnostic tools
 - Troubleshooting resources
- Models
 - Standup PCIe Plug-In Controller (-p)
 - HPE SR932i-p Gen11 SPDM Storage Controller ports and connectors
 - HPE Synergy mezzanine controller (-m)
 - HPE SR416ie-m 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
- HPE Storage Controller documentation resources
 - HPE Storage Controller documentation quick links
- Support and other resources
 - Accessing Hewlett Packard Enterprise Support
 - HPE product registration
 - Accessing updates
 - Remote support
 - Warranty information
 - Regulatory information
 - Documentation feedback

HPE SR Gen11 Controller

This controller offers a reliable family of RAID controllers that attach to:

- Internal hot-plug drives
- RAID levels 0, 1, 5, 6, 10, 50, 60, 1 Triple, and 10 Triple
- Mixed mode (RAID and HBA pass-through functionality simultaneously)
- Controller Based Encryption (CBE) for RAID volumes (HBA drives not supported)
- UEFI boot mode
- Smart Storage Administrator (SSA)

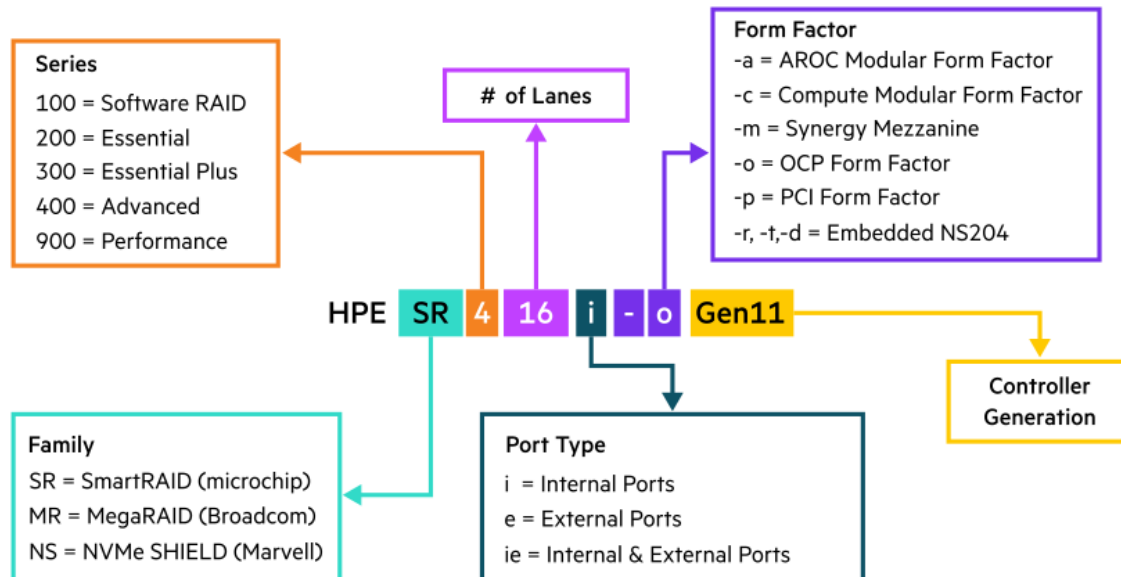


NOTE:

RAID1/10 Triple is previously known as RAID 1/10 ADM.

The HPE SR Gen11 Controller family supports 400 series and 900 series.

Figure 1. HPE Gen11 storage controller naming framework



Management and configuration tools for HPE SR Gen11 controller

- [HPE iLO GUI](#)
- [HPE OneView](#)
- [Intelligent Provisioning](#)
- [Redfish](#)
- [Smart Storage Administrator GUI](#)
- [Smart Storage Administrator CLI](#) & [Smart Storage Administrator Scripting](#)
- [UEFI System Utilities](#)

Subtopics

[400 series](#)

[900 series](#)

400 series

400 series controllers provide an enterprise-level solution for supported RAID levels and software-defined storage solutions. These controllers operate in Mixed Mode, which combines RAID and HBA operations simultaneously. They offer encryption for data-at-rest on any drive with HPE SR Secure Encryption and provide enterprise-class reliability, security, and efficiency.

This controller provides:

- Good RAID performance capabilities with 4 GB x72 flash-backed write cache
- Up to 8 SAS or SATA or NVMe lanes for internal drives
- x8 PCIe Gen4 host interface
- 16G NVMe or 24G SAS or 6G SATA support
- U.3 NVMe drives Support

Name	Supported HPE Gen11 compute modules
HPE SR416ie-m Gen11 8 Int 8 Ext Lanes/4GB Cache SPDM Mezzanine Storage Controller	HPE Synergy

900 series

900 series controllers are ideal for maximizing performance while supporting advanced RAID levels. These controllers operate in Mixed Mode which combines RAID and HBA operations simultaneously. They offer encryption for data-at-rest on any drive with HPE SR Secure Encryption. They offer flash-backed write cache, read-ahead cache, and provide enterprise-grade storage performance, reliability, security, and efficiency.

These controllers provide:

- Best RAID performance capabilities with 8 GB x144 flash-backed write cache
- Up to 32 SAS or SATA or NVMe lanes for internal drives
- x16 PCIe Gen4 host interface
- 16G NVMe or 24G SAS or 6G SATA support
- U.3 NVMe drives Support

Name	Supported HPE Gen11 servers
HPE SR932i-p Gen11 32 Internal Lanes/8GB Wide Cache SPDM PCI Plug-in Storage Controller	HPE ProLiant and HPE Apollo



NOTE: For any restrictions on where the card can be installed, consult the server [QuickSpecs](#) and user guide available here: <https://www.hpe.com/support/hpesc>.

Supported features for HPE SR Gen11 Controller

This chapter tells you the features you can use with the controller and explains how each feature works and why it is useful.

Subtopics

- [Features support](#)
- [RAID technologies](#)
- [Transformation](#)
- [Drive technology](#)
- [Security](#)
- [Reliability](#)
- [Performance](#)

Features support

This section lists the features supported for each controller class. For the latest information about the features supported by each individual controller, see the [HPE Compute SR Gen11 Storage Controllers QuickSpecs](#).



Subtopics

[OS environments](#)
[RAID technologies](#)
[Transformation](#)
[Drive technology](#)
[Security](#)
[Reliability](#)
[Performance](#)

OS environments

OS	400 series	900 series
Windows	✓	✓
Linux	✓	✓
VMware	✓	✓
UEFI boot mode	✓	✓

RAID technologies

Feature	400 series	900 series
RAID levels	0, 1, 5, 6, 10, 50, 60, 1T, 10T	0, 1, 5, 6, 10, 50, 60, 1T, 10T
Max volumes	64	64
Max physical drives	238	238
Max physical drives per volume	64	64
Drive protocol	SATA, SAS, NVMe	SATA, SAS, NVMe
Mixed mode (RAID and HBA simultaneously)	✓	✓
Read load balancing	✓	✓
Mirror splitting and recombining	✓	✓
Rapid Parity Initialization	✓	✓
Regenerative Writes	✓	✓
Backed out Writes	✓	✓
Full Stripe Writes	✓	✓
Dedicated spare	✓	✓
Predictive Spare Activation	✓	✓
Failure Spare Activation	✓	✓
Auto-replace Spare	✓	✓
Rapid Rebuild	✓	✓
Rebuild Priority	✓	✓

Transformation



Feature	400 series	900 series
<u>Expand Array</u>	✓	✓
<u>Move Array</u>	✓	✓
<u>Replace Array</u>	✓	✓
<u>Shrink Array</u>	✓	✓
<u>Mirror Array</u>	✓	✓
<u>Heal Array</u>	✓	✓
<u>Expand volume</u>	✓	✓
<u>Migrate RAID Level</u>	✓	✓
<u>Migrate Strip Size</u>	✓	✓
<u>Transformation Priority</u>	✓	✓

Drive technology

Feature	400 series	900 series
<u>Predictive drive failure</u>	✓	✓
<u>Online drive firmware update</u>	✓	✓
<u>Dynamic sector repair</u>	✓	✓
<u>Controller surface scan</u>	✓	✓
<u>Shingled magnetic recording</u>	✓	✓
<u>Hot-plug drive LED</u>	✓	✓
<u>SSD over-provisioning optimization</u>	✓	✓
<u>SSD wear gauge reports</u>	✓	✓

Security

Feature	400 series	900 series
<u>Controller Based Encryption LKM</u>	✓	✓
<u>Controller Based Encryption RKM</u>	✓	✓
<u>Self-Encrypting Drive HKM</u>	✓	✓
<u>Self-Encrypting Drive LKM</u>	✓	✓
<u>Self-Encrypting Drive RKM</u>	✓	✓
<u>Sanitize Erase</u>	✓	✓
<u>Sanitize Freeze Lock</u>	✓	✓
<u>Signed Firmware</u>	✓	✓
<u>Hardware-based Root of Trust</u>	✓	✓
<u>Secure Boot</u>	✓	✓
<u>Security Protocol and Data Model</u>	✓	✓

Reliability



Feature	400 series	900 series
Link Error Monitoring	✓	✓
Recovery ROM	✓	✓
Cache Error Checking and Correction	✓	✓
Thermal Monitoring	✓	✓

Performance

Feature	400 series	900 series
HPE SR SmartCache	✓	✓
SSD Smart Path	✓	✓
Read Cache	✓	✓
Flash-Backed Write Cache	✓	✓
Cache Ratio Selection	✓	✓
Write Cache Bypass Threshold	✓	✓
Drive Write Cache Control	✓	✓
Video on Demand	✓	✓
Strip Size Selection	✓	✓
Power Modes	✓	✓

RAID technologies

Subtopics

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

[Mixed mode \(RAID and HBA simultaneously\)](#)

[Striping](#)

[Mirroring](#)

[Parity](#)

[Spare drives](#)

[Drive rebuild](#)

Selecting the right RAID type for your IT infrastructure

The RAID setting that you select is based on 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

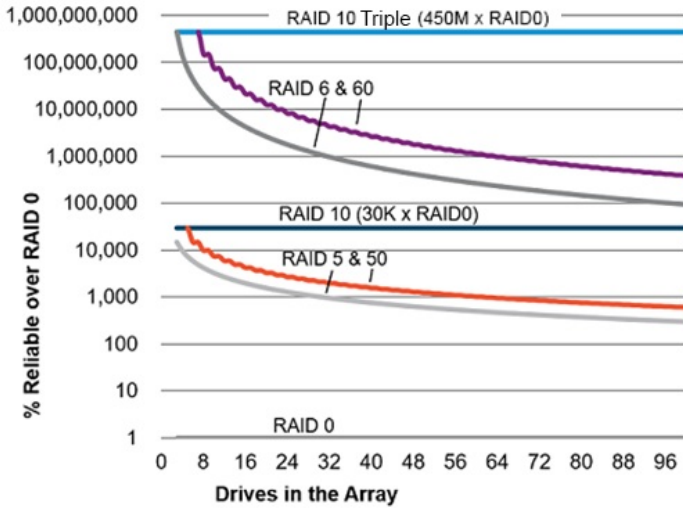
If your IT environment requires a high level of fault tolerance, select a RAID level that is optimized for fault tolerance.

This chart shows the relationship between the RAID level fault tolerance and the size of the storage array. The chart includes RAID 0, 1, 5, 50, 10, 6, 60, RAID 1 Triple, and RAID 10 Triple. 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.
- RAID 10 Triple is 450,000,000 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 following chart shows how RAID 10, 10 Triple, 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 that parity initialization has been 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 following table shows the Disk I/O for every host write:

RAID type	Disk I/O for every host write
RAID 0	1
RAID 1/10	2
RAID 1/10 Triple	3
RAID 5	4
RAID 6	6

Supported RAID levels can vary based on the controller model.

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 shows 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 and RAID 10 Triple remains flat with larger arrays.
- The usable capacity for RAID 5, 50, 6, and 60 increases with larger arrays.
- RAID 50 and RAID 60 assume two parity groups.

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

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

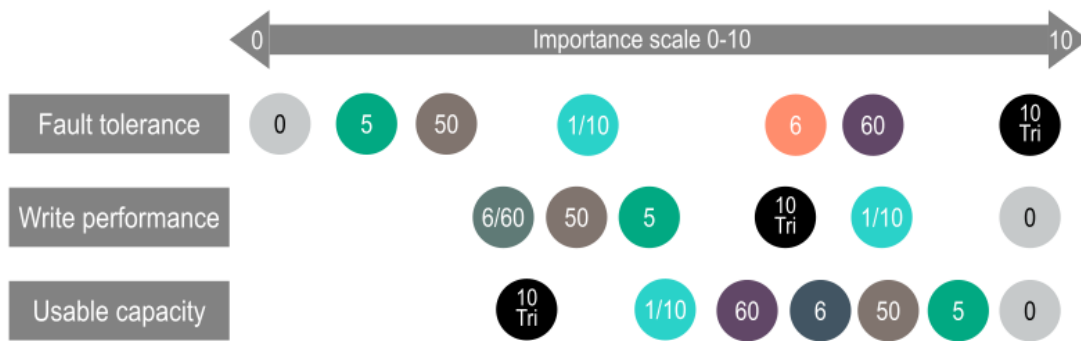
Supported RAID levels can vary based on the controller model.



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 can optimize the RAID types as follows:

- RAID 1/10 Triple: Optimize for fault tolerance and write performance.
- 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 HBA simultaneously)

Any drive that is not a member of a logical drive or assigned as a spare is presented to the operating system. This mode occurs by default without any user intervention. Logical drives are also presented to the operating system.

Controllers that support mixed mode 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 presented as HBA (except a two-drive mirror for boot support) can be accomplished with a single controller attached to a single backplane.

Drive LED	Method	HBA	RAID
Locate LED (Solid Blue)	SSACLI	Yes	Yes
	Virtual SCSI Enclosure Services (SES)	Yes	No
Drive Failure LED (Solid Amber)	Auto	Yes	Yes
	Virtual SES	Yes	No
Predictive Drive Failure LED (Flashing Amber)	Auto	No	Yes ¹
	Virtual SES	Yes	No
Reporting	See Diagnostic tools	Yes	Yes

¹ Applicable to any fault tolerant logical drive such as RAID 1, 10, 5, 6, 50 and 60. Not applicable to RAID 0.

Virtual SES is a computer protocol hosted by the controller driver. It is used with disk storage devices/enclosures to report and access drive bay locations, and control LEDs. The Virtual SES SCSI devices appear as a normal enclosure and support host tools such as the SG_UTIL Linux package, which contains the SG_SES tool.

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 printing or image editing, for example) or when cost is the most important consideration. The minimum number of drives required is one.





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 \(RAID 10\)](#)

[RAID 1 \(Triple\) and RAID 10 \(Triple\)](#)


[Read load balancing](#)

[Mirror splitting and recombining](#)


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.



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.



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.

RAID 1 (Triple) and RAID 10 (Triple)


In RAID 1 Triple and RAID 10 Triple configurations, data is duplicated to two additional drives. The usable capacity is $C \times (n / 3)$ where C is the drive capacity with n drives in the array. A minimum of 3 drives is required.

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





When the array has more than six physical drives, drives are mirrored in trios, and the fault-tolerance method is known as RAID 10 Triple. If a physical drive fails, the remaining two drives in the mirrored trio can still provide all the necessary data. Several drives in the array can fail without incurring data loss, as long as no three failed drives belong to the same mirrored trio. The total drive count must increment by 3 drives.



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 read performance of any configuration due to load balancing.
- This method has the highest data protection of any configuration.
- No data is lost when two drives fail, as long as no two failed drives are mirrored to another failed drive.
- Up to two-thirds 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 on individual drive load.

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



Mirror splitting and recombining

The split mirrored array feature splits any mirrored array (RAID 1, 10, 1 Triple, or 10 Triple) into multiple RAID 0 logical drives containing identical drive data.

The following options are available after creating a split mirror backup:

- Re-mirror the array and preserve the existing data. Discard the contents of the backup array.
- Re-mirror the array and roll back to the contents of the backup array. Discard existing data.
- Activate the backup array.

The re-mirrored array combines two arrays that consist of one or more RAID 0 logical drives into one array consisting of RAID 1 or RAID 1+0 logical drives.

For controllers that support RAID 1 Triple and RAID 10 Triple, this task can be used to combine:

- one array with RAID 1 logical drives and one array with RAID 0 logical drives into one array with RAID 1 Triple logical drives
- one array with RAID 1+0 logical drives and one array with RAID 0 logical drives into one array with RAID 10 Triple logical drives

This method allows you to clone drives and create temporary backups.

Parity

Subtopics

[RAID 5](#)

[RAID 50](#)

[RAID 6](#)

[RAID 60](#)

[Parity groups](#)

[Background parity initialization](#)

[Rapid parity initialization](#)


[Regenerative writes](#)

[Backed-out writes](#)

[Full-stripe writes](#)

RAID 5

RAID 5 protects data using parity (denoted by P_{x,y} in the figure). The system calculates parity data 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, its data 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.




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.


This method has the following benefits:

- Higher performance than 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.



This method is most useful when data loss is unacceptable but cost is also an important factor. The probability of data loss is less when an array is configured with a RAID 6 Advanced Data Guarding (ADG) setup than with a RAID 5 configuration.

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.

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 than 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.

For this setting, you can use any integer value greater than 1 with a restriction. The rule is 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 that is 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.

Background parity initialization

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, RAID 6, RAID 50, and RAID 60 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 operating system. Parity initialization takes several hours or days 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 feature has the benefit of allowing the logical drive to become usable sooner.

Rapid parity initialization

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. This is because the controller does not read the entire stripe (regenerative write) to update the parity data.

The rapid parity initialization method works by overwriting both the data and parity blocks in the foreground. The logical drive remains invisible and unavailable to the operating system until the parity initialization process completes. Keeping the logical volume offline eliminates the possibility of I/O activity. This speeds up the initialization process, and enables other high performance initialization techniques that wouldn't be possible if the volume was available for I/O. Once the parity is complete, the volume is brought online and becomes available to the operating system

This method has the following benefits:

- It speeds up the parity initialization process.
- It ensures that parity volumes use backed-out writes for optimized random write performance.

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 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. The write penalty for a RAID 6 and RAID 60 is always six drive operations. As you can see, the number of drives in the array does not influence the write penalty.

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

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



Full-stripe writes

A full-stripe write operation can be performed when writes to the logical drive are sequential. It can also be performed when multiple random writes that accumulate in the flash-backed write cache are found to be sequential. 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 . Here, 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](#)

[Predictive spare activation](#)

[Failure spare activation](#)

[Auto-replace spare](#)

Dedicated spare

A dedicated spare is a spare drive that is shared across multiple arrays within a single RAID controller.

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

The dedicated spare drive activates any time a drive within the array fails.

Predictive spare activation

Predictive spare activation mode will activate a spare drive anytime a member drive within an array reports a predictive failure. The data is copied to the spare drive while the RAID volume is still healthy.

Assigning one or more online spare drives to an array enables you to postpone replacement of faulty drives.

The predictive failure drive is marked as failed and ready for removal and replacement after the copy is complete. After you install a replacement drive, the controller will restore data automatically from the activated spare drive to the new drive.



This method has the following benefits:

- It is up to four times faster than a typical rebuild.
- It can recover bad blocks during spare activation.
- It supports all RAID levels including RAID 0.

Failure spare activation

Failure spare activation mode activates a spare drive when a member drive within an array fails to use fault tolerance methods to regenerate the data.

Assigning one or more online spare drives to an array enables you to postpone replacement of faulty drives.

Auto-replace spare

Auto-replace spare allows an activated spare drive to become a permanent member of the drive array. The original drive location becomes the location of the spare drive.

This method has the benefit of avoiding the copy-back operation after replacing the failed drive.

Drive rebuild

Subtopics

[Rapid rebuild](#)

[Puncture](#)

[Rebuild priority](#)

Rapid rebuild

These controllers include rapid rebuild technology for accelerating the rebuild process. Faster rebuild time helps restore logical drives to full fault tolerance before a subsequent drive failure can occur, reducing the risk of data loss.

Generally, a rebuild operation requires approximately 15 to 30 seconds per gigabyte for RAID 5 or RAID 6. Actual rebuild time depends on several factors, including the amount of I/O activity occurring during the rebuild operation, the number of disk drives in the logical drive, the rebuild priority setting, and the disk drive performance.

This feature is available for all RAID levels except RAID 0.

Puncture

Puncture is a controller feature which allows a drive rebuild to complete despite the loss of a data stripe caused by a fault condition that the RAID level cannot tolerate. When the RAID controller detects this type of fault, the controller creates a "puncture" in the affected stripe and allows the rebuild to continue. Puncturing keeps the RAID volume available and the remaining volume can be restored.

Future writes to the punctured stripe will restore the fault tolerance of the affected stripe. To eliminate the punctured stripe, the affected volume should be deleted and recreated using Rapid Parity Initialization (RPI) or Erasing the drive before creating the logical drive. The data affected by the punctured stripe must be restored from a previous backup.

Punctures may be minimized by performing the following:

- Update drivers and firmware.
- Increase surface scan priority to high.
- Review IML and OS system event logs for evidence of data loss or puncture.

Rebuild priority

The Rebuild Priority setting determines the urgency with which the controller treats an internal command to rebuild a failed logical drive.

- Low setting: Normal system operations take priority over a rebuild.
- Medium setting: Rebuilding occurs for half of the time, and normal system operations occur for the rest of the time.
- Medium high setting: Rebuilding is given a higher priority over normal system operations.
- High setting: The rebuild takes precedence over all other system operations.

If the logical drive is part of an array that has an online spare, rebuilding begins automatically when drive failure occurs. If the array does not have an online spare, rebuilding begins when the failed physical drive is replaced.

Subtopics

Before replacing drives

- Open Systems Insight Manager, and inspect the Error Counter window for each physical drive in the same array to confirm that no other drives have any errors. For more information about Systems Insight Manager, see the documentation on the Insight Management DVD or on <https://www.hpe.com/support/hpesc>.
- Be sure that the array has a current, valid backup.
- Confirm that the replacement drive is of the same type as the degraded drive (either SAS or SATA and either hard drive or solid-state drive).
- Use replacement drives that have a capacity equal to or larger than the capacity of the smallest drive in the array. The controller immediately fails drives that have insufficient capacity.

In systems that use external data storage, be sure that the server is the first unit to be powered down and the last unit to be powered up. Taking this precaution ensures that the system does not, erroneously, mark the drives as failed when the server is powered up.

In some situations, you can replace more than one drive at a time without data loss. For example:

- In RAID 1 configurations, drives are mirrored in pairs. You can replace a drive if it is not mirrored to other removed or failed drives.
- In RAID 10 configurations, drives are mirrored in pairs. You can replace several drives simultaneously if they are not mirrored to other removed or failed drives.
- In RAID 50 configurations, drives are arranged in parity groups. You can replace several drives simultaneously, if the drives belong to different parity groups. If two drives belong to the same parity group, replace those drives one at a time.
- In RAID 6 configurations, you can replace any two drives simultaneously.
- In RAID 60 configurations, drives are arranged in parity groups. You can replace several drives simultaneously, if no more than two of the drives being replaced belong to the same parity group.
- In RAID 1 Triple and RAID 10 Triple configurations, drives are mirrored in sets of three. You can replace up to two drives per set simultaneously.

To remove more drives from an array than the fault tolerance method can support, follow the previous guidelines for removing several drives simultaneously, and then wait until rebuild is complete (as indicated by the drive LEDs) before removing additional drives.

However, if fault tolerance has been compromised, and you must replace more drives than the fault tolerance method can support, delay drive replacement until after you attempt to recover the data.

Transformation

Subtopics

[Array transformations](#)

[Volume transformations](#)

Array transformations

Subtopics

[Expand array](#)

[Move array](#)

[Replace array](#)

[Shrink array](#)

[Mirror array](#)

[Heal array](#)

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 (for example, 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.

For more information, see [Smart Storage Administrator - How to Expand an Array and Extend the Logical Drive](#).

Move array

The Move Array operation allows you to transfer the contents of a disk array from one set of physical drives to a second set of physical drives. Note the following conditions and restrictions for the Move Array operation:

- The destination physical drive set must have the same number of drives as the source physical drives set.
- The array type (SAS or SATA) must remain the same.
- The destination drive must have enough capacity to hold all the logical drives present in the source array.

Replace array

The Replace Array operation enables you to transfer the contents of an array to an existing empty array or a new array. All logical drives from the source array are transferred. The original array is deleted and its data drives are freed as unassigned drives. The drive types at source and destination arrays can be different. Note the following conditions and restrictions for the Replace Array operation:

- The destination array must have the same number of physical drives as the source array to be replaced.
- Both the source and the destination arrays must be in the OK state. All the existing logical drives in the source arrays must be in the OK state.
- The destination array must have enough capacity to hold all the logical drives present in the source array.

Shrink array

The Shrink Array operation allows you to remove drives from an existing array. The following conditions apply:

- The array must have enough free space to accommodate all existing logical drives.
- You may not remove drives from the array if the resulting number of drives will not support the fault tolerance (RAID level) of any existing logical drive. For example, if you have an array with four physical drives and a RAID 5 logical drive, you can remove at most one drive since RAID 5 requires at least three physical drives.
- If the array contains a RAID 1+0 logical drive, you can only remove even number of drives.
- If the array contains a compound RAID (RAID 50 or RAID 60) logical drive, drives can only be removed in multiples of the number of parity groups. For example, an array with 10 physical drives and a RAID 50 logical drive can be shrunk by removing two or four disks only.

Mirror array

The Mirror Array operation allows you to double the number of data drives in the array and convert all logical drives in the array to RAID 1 or RAID 1+0.

Keep the following points in mind:

- This option is available only if the array contains only RAID 0 drives.
- When the total number of data drives in the resulting array is two, the resulting RAID level is RAID 1. When the total number of data drives is four or more, the resulting RAID level is RAID 1+0.

Heal array



The Heal Array 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 the original drives must be the same interface type (such as SAS or SATA) as the original drives.
- The operation is available only if enough unassigned physical drives of the correct size are available.
- The array has at least one failed drive.
- 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

[Expand volume](#)

[Migrate RAID level](#)

[Migrate strip size](#)

[Transformation priority](#)

Expand volume

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.

For more information, see [Smart Storage Administrator - How to Expand an Array and Extend the Logical Drive](#).

Migrate RAID level

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

Migrate strip size

The migrate strip size feature allows you to change the current strip size for your logical drive. When the strip size changes, you will have more or less unused space, depending on the strip size with which you started. For migration to a larger strip size to be possible, the array needs to contain unused drive space. This extra space is necessary because some of the larger data stripes in the migrated array are likely to be filled inefficiently.

Transformation priority

As the transformation priority level increases, the rate at which requests from the operating system are processed decreases. Transformation refers to array expansions, logical drive extensions, logical drive migrations, and array shrink and move operations.

- High: Transformation will complete as fast as possible at the expense of normal I/O.
- Medium: Transformation will perform with some impact on normal I/O.
- Low: Transformation will perform only when normal I/O is not occurring. This level will cause the transformation to take the most time to complete.

Drive technology

Subtopics



[Predictive drive failure](#)
[Online drive firmware update](#)
[Dynamic sector repair](#)
[Controller surface scan](#)
[Shingled magnetic recording](#)
[Hot-plug drive LED](#)
[SSD over-provisioning optimization](#)
[SSD wear gauge reports](#)

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.

Online drive firmware update

These controllers support online drive flashing, which saves time when updating drive firmware. Instead of taking the drive offline to load 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.

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 built a dynamic sector repair feature into these controllers:

- Perform a background surface analysis during inactive periods, continually scanning all drives for media defects
- Detect media defects when accessing a bad sector during busy periods
- Automatically remap the bad sector to a reserve area on the disk drive
- In a fault-tolerant configuration, automatically regenerate the data and write it to the remapped reserved area on the disk drive

Controller surface scan

Controller surface scan analysis is an automatic background process that ensures that you can recover data if a drive failure occurs. The controller scanning process:

- Verifies physical drives in fault-tolerant logical drives for bad sectors.
- Verifies the consistency of parity data in RAID 5 or RAID 6 Advanced Data Guarding (ADG) configurations.

You can disable the surface scan analysis, set it to high, or specify a time interval that the controller is inactive before a surface scan analysis is started on the physical drives that are connected to it.

- **Disabled:** Disabling the controller surface scan can decrease the potential latency impacts that might occur due to waiting for a scanning I/O to complete, but at the cost of not detecting the growth of bad blocks on the media before a data loss situation.
- **High:** Setting the controller surface scan to high increases the probability of detecting a bad block before it becomes a data loss situation.
- **Idle:** Setting the controller surface scan to idle and setting the corresponding surface scan delay can decrease the potential latency impacts, but still allow the scanning of bad blocks during the idle time.

Parallel surface scans count controls how many controller surface scans can operate in parallel per array. The feature is used when there is more than one logical drive on a controller on more than one array configured. This setting allows the controller to detect bad blocks on multiple logical drives on different arrays in parallel. It can significantly decrease the time it takes to detect back, especially for logical drives using large capacity drives on multiple arrays.



Shingled magnetic recording

Shingled Magnetic Recording (SMR) is a magnetic storage data recording technology for HDD that allows up to 30% higher capacity by overlapping the previous drive tracks. Thus, the tracks partially overlap, similar to roof shingles. The overlapping tracks slow down random write performance since the operating system must perform a read modify write of the entire zone. SAS SMR drives use the Zoned Block Command (ZBC) set. SATA SMR drives use the Zoned ATA Command (ZAC) set.

Drives	Host Managed (HM)	Host Aware (HA)	Device Managed (DM)
SAS SMR	HBA Only (ZBC)	HBA Only (ZBC)	Not supported
SATA SMR	HBA Only (ZAC)	HBA Only (ZAC)	SATA SMR + DM is not supported

This method has the following benefits:

- Support for HDD with higher storage density
- Supports for HDD with lower cost per GB
- Support for HDD with lower power per GB

Hot-plug drive LED

Figure 1. LFF Low Profile (LP)




Figure 2. SFF Basic Carrier (BC)





Figure 3. E3S EDSFF Carrier



Figure 4. E3S2T EDSFF Carrier



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 (one flash per second)	The drive has failed, or a predictive failure alert has been received for this drive; it also has been identified by a management application.
		Flashing amber (one flash per second)	A predictive failure alert has been received for this drive. Replace the drive as soon as possible.
2	Online\Activity	Solid green	The drive is online and has no activity.
		Flashing green (4 flashes per second)	The drive is operating normally and has activity.
		Flashing green (one flash per second)	The drive is doing one of the following: <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.
3	Power\Activity	Solid green	The drive is powered on and has no activity
		Flashing green (4 flashes per second)	The drive is operating normally and has activity

¹ If for a failed drive the link is not working and the controller cannot detect the link, the Fault\Locate LED is Off.

SSD over-provisioning optimization

Solid state drive manufacturers reserve an additional percentage of the total drive capacity for over-provisioning. The over-provisioned capacity is used to manage writes and wear leveling. SSD over-provisioning can increase the endurance of an SSD by distributing the total number of writes and erases across a larger population of NAND flash blocks and pages.

Over-provision optimization is an optional feature that initializes the drive to use the entire capacity to manage writes and wear leveling. As logical drives are created and data is written, this over-provisioned capacity is reduced. The optimization process is performed when the first logical drive in an array is created, and when a physical drive is used to replace a failed drive.

This feature provides the following benefits:

- Improved SSD write performance
- Improved SSD endurance

SSD wear gauge reports

These reports contain information about the current usage level and remaining expected lifetime of SSDs attached to the system.

When running a report, you can either view a graphic representation of the report with SSD usage and estimated lifetime information, or generate a report without

a graphical display, with the option of saving the report.

Security

Subtopics

- [Controller-based encryption](#)
- [Self-Encrypting Drive](#)
- [Sanitize erase](#)
- [Sanitize freeze lock](#)
- [Signed firmware](#)
- [Hardware-based Root of Trust](#)
- [Secure Boot](#)
- [Security Protocol and Data Model](#)

Controller-based encryption

Controller-Based Encryption (CBE) is an enterprise-grade data encryption solution that protects data at rest on any SAS/SATA/NVMe drive configured as a member of a RAID volume. Controller-based encryption is also known as HPE SR Secure Encryption. The solution is available for both local and remote deployments.

CBE is configured using the Smart Storage Administrator (SSA).

Prerequisites:

- Only supports drives in RAID mode.
- A valid secure encryption license for each server to be encrypted.

Subtopics

- [Local Key management mode](#)
- [Remote Key management mode](#)

Local Key management mode

Local Key management mode, or Local mode, is a solution designed for small to medium-size data centers. The solution uses a passphrase password, or Master Encryption Key name, to set the security on the controller and enable encryption. The Master Encryption Key must be tracked independently of the controllers in case the controller needs replacement or drive migration is required among controllers with different passwords.

This method has the following benefits:

- Encrypts data on both the attached bulk storage and the cache memory of the controllers.
- Supports any HDD or SSD in the HPE server portfolio.
- Does not require ESKM.

Remote Key management mode

In Remote Key management mode, keys are imported and exported between the controller and the Enterprise Secure Key Manager (ESKM), which provides a redundant, secure store with continuous access to the keys. To enable key exchanges between the controller and the ESKM, a network connection is required both during pre-OS boot time and during OS operations. Because the controller does not have direct network access capabilities, iLO provides the necessary network access to facilitate key exchanges between the controller and the ESKM. For more information see, "Using key managers with iLO" in [HPE iLO 6 User Guide](#).

Prerequisites:

- Integrated Lights Out (iLO) Advanced or Scale Out Edition license, per HPE ProLiant server
- Network availability
- Remote ESKM

This method has the following benefits:

- Encrypts data on both the attached bulk storage and the cache memory of the controllers
- Supports any HDD or SSD in the HPE server portfolio
- Keys are kept in separate storage from servers to protect against physical removal

Self-Encrypting Drive

Self-Encrypting Drive (SED) 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.

HPE SR controller supports the following Trusted Computing Group (TCG) storage security subsystem classes:

- Enterprise Standard version 1.01
- Opal Standard version 2.01

Guidelines

- Secured logical drives must be deleted before returning to Original Factory State (OFS), which also removes all the data on the logical drive.
- Import option is only available for a secured SED configured from other HPE SR controllers.
- Reset controller settings or clear controller configuration option deletes all secret keys, passwords, and identifications on the controller without modifying the drives.
- SED management is not available when Controller-Based Encryption (CBE) is enabled.
- Unsecured drives must be in OK state.
- If the controller password is enabled, unlock the controller before performing any drive removal or reinsertion operations, or boot from a secured SED.

Subtopics

[Host Key management](#)

[Local Key management](#)

[Remote Key management](#)

Host Key management

To use host key management, enable the Self-Encrypting Drive (SED) as an unassigned drive and expose the drive to the OS. This method allows you to manage SED using third-party key management tools like SEDutil. SED monitoring is also available in Smart Storage Administrator (SSA), Smart Storage Administrator CLI (SSACL) tool, and configuration utility in UEFI System Utilities.

Host key management requires controller firmware version 3.01.14.62 or later.



NOTE: SED in host key management mode is not recommended for RAID volumes.

Local Key management

To enable local key management, you must provide a controller wide security key identity and security key. While booting 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 theft.

The requirements for local key management are:

- Controller firmware version 3.01.17.56 or later
- Windows driver version v1010.42.0.1020 or later.
- Linux driver version v2.1.18-045 or later.
- VMware driver version v4330.0.116 or later.

Remote Key management

The configuration utility in UEFI System Utilities works with the iLO key manager to create the security key identify and security key in the remote key manager server. iLO key manager must 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.

The requirements for remote key management are:

- Controller firmware version 3.01.23.72 or later.
- Windows driver version v1010.74.0.1020 or later.
- Linux driver version v2.1.24-046 or later.
- VMWare driver version v4530.0.104 or later.

Sanitize erase

When you sanitize erase a drive, you remove all sensitive information from a physical drive. This includes non-volatile media, non-volatile cache, bad blocks, and overprovisioned areas. Sanitize erase operations cannot be stopped after starting, and the drive will continue to sanitize after a hot-plug or server reboot. During the sanitize erase operation, the drive is unusable until after the process is complete.

Sanitize erase methods:

- **Restricted:** Using the restricted sanitize method means that until a drive successfully completes the sanitize operation, it will be unusable. If a restricted sanitize operation fails, you are only allowed to start another sanitize operation, or, if the drive is under warranty, you can return it to HPE.
- **Unrestricted:** Using the unrestricted sanitize method means that the drive will be recoverable in the case that the sanitize erase operation fails. User data might still be present on the drive. Not all drives support the unrestricted sanitize method.



NOTE:

These sanitize erase methods satisfy the requirements for the purge action set by the National Institute of Standards and Technology. For more information about the purge action, see "Guidelines for Media Sanitization" at the U.S. Department of Commerce website (<https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-88r1.pdf>).

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 (SSD) also known as instant secure 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.

Sanitize freeze lock

Sanitize freeze lock and anti-freeze lock allows you to control whether the sanitize commands for SATA drives are allowable by the operating system and Hewlett Packard Enterprise tools after a system boot or drive hot-plug.

This feature has three settings:

- None-This state is the normal state of the physical disk. No freeze or anti-freeze commands are sent to any of the drives.
- Freeze-This setting prevents a drive sanitize operation.
- Anti-freeze-This setting prevents physical disks from being frozen. This setting enables drive sanitize operations.

This setting is applicable only to SATA drives connected to an HPE SR controller.

Signed firmware

Signature ensures that the firmware is authentic and is not altered in any way. A signature is added to the firmware file to protect you from receiving a corrupted or hacked version of the firmware. If a hacked version of firmware is placed on the controller, data is at risk. Signed firmware protects your data.

Hardware-based Root of Trust

Root of Trust (RoT) is a source that can always be trusted within a cryptographic system. It safeguards the security of data and helps in building trust in the overall ecosystem. RoT is the critical fundamental building block that contains the keys used for cryptographic functions during the secure boot process. These cryptographic functions provide strong protection for product life cycle management during all the operation phases, such as power-up, run-time operations.

Secure Boot

Secure boot is a security standard developed by members of the server industry to help ensure that a device boots using only software that is trusted by the OEM. When the server starts, the firmware checks the signature of each piece of boot software, including UEFI firmware drivers, also known as Option ROMs, EFI applications, and the OS. If the signatures are valid, the server boots, and the firmware gives control to the OS.

Security Protocol and Data Model

Security Protocol and Data Model (SPDM) is a security standard developed by the 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.

The 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



- [Link error monitoring](#)
- [Recovery ROM](#)
- [Cache error checking and correction](#)
- [Thermal monitoring](#)

Link error monitoring

This controller monitors and reports link errors within the SAS topology. A SAS link is a serial connection between devices such as the controller, expander, or drive. Each of these devices communicate using one or more transmitter and receiver pairs. Each receiver counts the number of link errors that it receives since power up. Normally, a link error is recoverable within the SCSI or ATA protocol. The controller collects these counters in the background and evaluates how many link errors accumulate within a 1 hour time period. If the number of link errors exceed a threshold, the controller reports the error to the System Event Log.

This method has the benefit of allowing you to identify faulty hardware such as controllers, SAS cables, or I/O modules.

Recovery ROM

The controllers store a redundant copy of the controller firmware image to protect against data corruption. If the active firmware image becomes corrupt, the controllers use the redundant firmware image and continue operating. Once the redundant firmware is activated, the corrupted firmware image is then updated with valid firmware by the active firmware. The recovery ROM provides protection against power outages during firmware flashing.

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

- [HPE SR SmartCache](#)
- [Cache](#)
- [I/O performance mode](#)
- [Video on demand](#)
- [Strip size selection](#)
- [Power modes](#)

HPE SR SmartCache

HPE SR SmartCache (also known as Microchip maxCache 4.0) creates an SSD cache volume that is used to cache frequently accessed data from a data volume. The SSD cache volume uses the same controller-based encryption methods and keys as the data volume. After creating the data volume, the Cache Manager tool within Smart Storage Administrator (SSA) can be used to enable HPE SR SmartCache.

In the first screen, select drives for the cache volume. In the second screen, the cache volume is created and data volume is assigned to it. The cache volume supports RAID 0, 1, or 5 using either Write-Through or Write-Back. HPE SR SmartCache requires an energy pack and HPE SR SmartCache License. The HPE SR SmartCache License is included and preinstalled on the HPE Smart Array P816i-a SR Gen10 and all Gen10 Plus or later controller models.



Subtopics

[Elements of HPE SR SmartCache](#)

[HPE SR SmartCache write policy and RAID type](#)

[HPE SR SmartCache line size](#)

[HPE SR SmartCache volume capacity](#)

[Features and benefits of HPE SR SmartCache](#)

[Features not supported with HPE SR SmartCache](#)

Elements of HPE SR SmartCache

- **Data Volume** - The first element is the data volumes, which are composed of any supported drives typically using HDD. SSD Smart Path that is `IOPerfModeEnabled` in Redfish on the data volumes must be disabled. Deleting the data volume automatically results in the deletion of the cache volume.
- **Cache Volume** - The second element is the cache volumes, which are composed of any supported SSD drives used to cache the data volume. The capacity of the cache volume must be less than or equal to the capacity of the data volume. If multiple cache volumes are present, they must be on the same array of SSD drives.
- **Cache Module** - The final element is the cache module, which contains metadata. User data for the cache volume may also be in the cache module.

HPE SR SmartCache write policy and RAID type

HPE SR SmartCache supports a cache write policy of either write-back or write-through.

- **Write-through cache policy** - Write-through caching accelerates frequently accessed read operations using a RAID 0 or RAID 1 cache volume. All write operations go to the data volume.
- **Write-back cache policy** - Write-back caching accelerates frequently accessed read and write operations using a RAID 0, RAID 1, or RAID 5 cache volume. Writes may be cached on the cache volume and written to the data volume later. The cache write policy must be set to write-through to allow deletion of the cache volume.

HPE SR SmartCache line size

Cache line size is the amount of data allocated for each line in the controller cache. Supported values are 64KiB or 256KiB, but some controllers may only support 64KiB. It can impact the performance and maximum size of both the data volume and cache volume.

Cache line size is equivalent to the strip size or stripe size depending upon the cache volume RAID level.

- **RAID 0 & RAID 1 cache volumes** - Cache line size is the strip size. Also known as `StripSizeBytes` in Redfish.
- **RAID 5 cache volumes** - Cache line size is the full stripe size. The controller automatically generates the strip size by dividing the cache line size by the number of data drives that is total drive count minus 1. This calculation must result in a valid strip size of 16KiB, 32KiB, 64KiB, or 128KiB. Else, the creation of the cache volume fails. As a result, RAID5 cache volumes are limited to the following combinations of total drive count and cache line size.

Total drives count	Cache line size KiB	Generated RAID 5 strip size KiB
3	64	32 (64 / 2)
3	256	128 (256 / 2)
5	64	16 (64 / 4)
5	256	64 (256 / 4)
9	256	32 (256 / 8)

HPE SR SmartCache volume capacity

The cache volume capacity must be less than or equal to the capacity of the data volume. The maximum capacity of the cache volume and data volume is based on the cache line sizes and cache module sizes as listed in the following table.

Cache Line Size	Cache Module Size	Min Cache Volume Size	Max Cache Volume Size		Max Data Volume Size		Required Cache Ratio
			GB	GiB	TB	TiB	
64	1	16	1074	1024	274	256	0/100
64	2, 4	16	2147	2048	274	256	0/100
64	8	16	2147	2048	274	256	50/50
256	1	16	4294	4096	1099	1024	0/100
256	2, 4	16	8589	8096	1099	1024	0/100
256	8	16	8589	8096	1099	1024	50/50



NOTE:

Servers that are configured with HPE SmartCache write-back must be shutdown gracefully to avoid reports of inconsistent parity-repaired messages. If the server is shutdown ungracefully when using HPE SmartCache write-back cache, it increases the possibility of inconsistent parity over flash-backed write cache (FBWC). If no-battery write caching is enabled, the 8 GiB cache modules can be configured for a 0/100 cache ratio.

Features and benefits of HPE SR SmartCache

- **Workload Acceleration** - Performance by caching hot data on SSDs.
- **Simplified Deployment** - Seamless integration into your data center with no application or OS changes.
- **Efficient Operations** - Performance and efficiency gains using your existing HPE ProLiant server technology.

Features not supported with HPE SR SmartCache

When using HPE SR SmartCache, the following features are not available. If you have to use a feature, disable HPE SR SmartCache, complete the operation, and then re-enable.

- Expand Array
- Move Array
- Replace Array
- Shrink Array
- Mirror Array
- Heal Array
- Extend Logical Drive
- Migrate RAID Level
- Migrate Strip Size
- Transformation Priority
- Mirror Splitting and Recombining
- Change Cache Ratio

Cache

Subtopics

[Read cache](#)

[Flash-backed write cache](#)

[Cache ratio selection](#)

[Write cache bypass threshold](#)

[No-battery write cache](#)

[Drive write cache control](#)

Read cache

The controllers use an adaptive read-ahead algorithm that

- Detects sequential read activity on single or multiple I/O threads
- Predicts when sequential read requests will follow
- Reads ahead from the disk drives

When the read request occurs, the controller retrieves the data from high-speed cache memory in microseconds rather than from the disk drive in milliseconds. This adaptive read-ahead scheme provides excellent performance for sequential small block read requests.

This algorithm anticipates data needs and reduces wait time.

The controller disables read-ahead when it detects nonsequential read activity. The controller adaptive read-ahead caching eliminates issues with fixed read-ahead schemes that increase sequential read performance but degrade random read performance.

Read cache can only increase performance if read data has previously been stored in the cache. Since the size of the disk array is many orders of magnitude larger than the size of the cache, the probability that a random read would already be in the cache is small. For this reason, the controllers do not store random read data in the cache.

Read cache is most effective in increasing the performance for sequential small-block read workloads and, in particular, read workloads at low queue depth. The controller differentiates between sequential and random workloads. It uses read cache in a predictive capacity to prefetch data when it detects sequential workloads. It identifies the pattern of the read commands, and then reads ahead on the drives. After reading the data, the controller puts that data into the cache, so it is available if the upcoming read commands call for it.

You can use the Smart Storage Administrator (SSA) utility to configure the percentage of the cache to use for read caching. The default configuration on these controllers assigns 10% of the available cache space for read cache.

Flash-backed write cache

These controllers use a write-back caching scheme that lets host applications continue without waiting for write operations to complete to 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 then immediately return completion status to the OS. The write operation completes in microseconds rather than milliseconds. The controller writes data from the controller’s write cache to disk later, at an optimal time for the controller.

Once the controller locates write data in the cache, subsequent reads to the same disk location come from the cache. Subsequent writes to the same disk location will replace the data held in cache. This is a “read cache hit.” It improves bandwidth and latency for applications that frequently write and read the same area of the disk.

The write cache will typically fill up and remain full usually in high-workload environments. The controller uses this opportunity to analyze the pending write commands to improve their efficiency. The controller can

- Use write coalescing that combines small writes to adjacent logical blocks into a single larger write for quicker execution
- Perform command reordering, rearranging the execution order of the writes in the cache to reduce the overall disk latency
- Store and analyze a larger number of pending write commands, increasing the opportunities for write coalescing and command reordering while delivering better overall performance

When the controller has a large cache memory size, it can coalesce and reorder commands efficiently, which improves overall array performance.

You can use Smart Storage Administrator (SSA) to configure the percentage of the cache to use for write caching.

Flash-backed write cache (FBWC) uses flash devices to retain cache data and the energy pack to provide power during a power loss. The FBWC offers significant advantages over earlier BBWC systems. While a battery-backed write cache (BBWC) requires backup power during the entire power loss, an FBWC only needs power during the time it takes to backup from DRAM to flash. Since the FBWC writes the contents of memory to flash devices, there is no longer a 48-hour energy pack life limitation, and the data posts to the disk drive on the next power-up.

Cache ratio selection

The controller cache ratio setting determines the amount of memory allocated to read and write operations. Different types of applications have different optimum settings. You can change the ratio if the following are true:

- The controller has a cache that uses backup power from HPE Smart Storage Battery or HPE Smart Storage Hybrid Capacitor.

- There are logical drives configured on the controller.

The HPE SR932i-p Gen11 controller has a battery-backed cache size of 4GB, so the maximum write cache ratio is restricted to 50% of the 8GB cache size. This ratio suits most workloads. However, workloads that are characterized by highly sequential reads or reads from recent writes might benefit from a higher read percentage.

Write cache bypass threshold

All writes larger than the specified value will bypass the write cache and be written directly to the disk for non-parity RAID volumes.

A smaller value allows the controller to reserve write caching to I/O smaller than the threshold.

No-battery write cache

The no-battery write cache option (NBWC) is supported by these controllers that do not require an energy pack.

Drive write cache control

Drive write cache is cache within the physical drive. On controllers and drives that support physical drive write cache, you can enable or disable the write cache for all physical drives that are:

- Configured as part of a logical drive.
- Unconfigured and exposed to the host on the controller.

I/O performance mode

I/O performance mode is also known as SmartPath. SSDs require special tactics to capture the full advantage of their low-latency capabilities. I/O performance mode enables high performance of SSD-based logical volumes by allowing certain types of I/O requests to take a more direct path to the physical disks, bypassing most of the firmware layers of the RAID controller. When you create an array, the I/O performance mode is enabled by default.

The device driver software coordinates with the controller firmware to:

- Maintain the necessary disk mapping information.
- Decide which I/O requests are eligible for HPE SSD Smart Path.

All other requests and any error handling are still routed through the normal I/O path on the controller. This method has the following benefits:

- Benefits repetitive, read-heavy I/O workloads using the accelerated path.
- Frees up more I/O handling capacity on the normal I/O path.

Video on demand

Video streaming services, like Video On Demand (VOD), or Video Surveillance, typically require significant amounts of disk storage with predictable latency, high bandwidth, and generally using large size I/Os. It differs from low latency optimizations that prioritize absolute lowest latency at the expense of variability and bandwidth. The controller offers several video streaming applicable optimizations. Other system level optimizations should be evaluated, like I/O prioritization in the BIOS, block layer, and aligning file system allocations to RAID stripes.

This method has the following benefits:

Disabling the Elevator Sorting

Reduces maximum latency by processing I/Os in order.

Enabling the Degraded Performance Optimization

If using a parity protected RAID level such as RAID 5/50/6/60 optimizes for large block writes while in a degraded mode.

Setting the controller Cache Ratio to 100% write

The high stream count creates a random read I/O profile that will have little benefit of a read ahead cache.

Disabling the Controller Monitor and Performance collection

Reduces the latency spikes under consistent heavy I/O load caused by collecting management data.
Increasing the Surface Scan Delay to 30

Minimizes latency impact by controller media surface scans.
Using a Rebuild Priority of Medium or Medium High

Interleaves rebuild I/Os in a consistent way to have a more predicable latency during RAID rebuilds.
Enabling the Flexible Latency Scheduler

Reduces maximum latency for an individual I/O by prioritizing the I/O the longer it takes.

Strip size selection

When a controller makes an array, the unit of data that it manipulates 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.

The best performance and drive longevity 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.

Power modes

There are three available power modes:

- Maximum performance
- Minimum power
- Balanced

Maximum performance (default)

This is the default setting. All settings are selected based on maximum performance. Power savings options that affect performance are disabled.

Balanced

You can use this setting to save power with minimal effects on performance. For large queue depths, this setting affects throughput by 10% or less.

At lower queue depths or infrequent I/O, impacts on performance might be greater. This command is typically useful in environments using only hard drives, and is not recommended when using SSDs.

Settings are based on the user configuration, such as the number or types of drives, the RAID level, and storage topology. Significant changes to the configuration might require a reboot for optimal setting selection. If a reboot is required to change settings, SSA generates a warning.

Minimum power

When settings are selected without regard to system performance, maximum power savings is achieved. Hewlett Packard Enterprise recommends this setting for specific applications, but it is not appropriate for most customers. Most applications will suffer significant performance reduction.



IMPORTANT: A reboot might be required after switching power modes to optimize savings and performance.



IMPORTANT: When the power mode is set to Balanced, future controller configuration changes might require a reboot for optimal performance.

Installation instructions for HPE SR Gen11 Controller

This chapter tells you how to install the controller onto the server and connect the internal storage devices.

Subtopics

[Supported servers](#)

[Installing in an unconfigured server](#)

[Installing in a previously configured server](#)

[Installing an HPE Synergy mezzanine controller \(-m\)](#)

[Installing a standup PCIe plug-in controller \(-p\)](#)

Supported servers

For more information about installing the controller in a supported server, see the server user guide.

For the list of controllers supported on a server, see the server [QuickSpecs](#).

Installing in an unconfigured server

Procedure

1. Install the hardware.

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

2. For all the controllers except HPE SR416ie-m Gen11 SPDM Storage Controller, connect one end of the controller backup power cable to the backup power connector on the controller and the other end to the controller backup power connector on the system board or PCI riser board.
3. For all the controllers, install the optional energy pack.
4. Install physical drives, as needed, and attach the physical drives to the controller.
5. Power up the server.
6. Use the Service Pack for ProLiant (SPP) to deploy updated firmware, software, and device drivers to the server.

For more information about SPP, see the SPP website (<https://www.hpe.com/servers/spp>).

You might need to extract the controller driver from the SPP if your operating system installation files do not include the driver and if you do not plan to use Intelligent Provisioning to install the operating system.

7. Create a storage array using the Smart Storage Administrator or the configuration utility in UEFI System Utilities.
8. Install the operating system and device drivers.

If you use Intelligent Provisioning, select the Firmware Update option to apply the updated firmware. For more information about Intelligent Provisioning, see the product documentation on the [Hewlett Packard Enterprise website](#).

If you do not use Intelligent Provisioning to install the operating system, and if you are prompted for the driver during the installation, point to the driver that you extracted in step 6.


Installing in a previously configured server

Prerequisites

Before beginning this procedure, download the SPP from the Hewlett Packard Enterprise website <https://www.hpe.com/servers/spp/download>.

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. Do one of the following:
 - If the new controller is the new boot device, install the device drivers.
 - If the new controller is not the new boot device, go to the next step.
5. Ensure that users are logged off and all tasks are completed on the server.
6. Power down the server.



**CAUTION:**

In systems that use external data storage, be sure that the server is the first unit to be powered down and the last to be powered back up. Taking this precaution ensures that the system does not erroneously mark the drives as failed when the server is powered up.

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. Disconnect all peripheral devices.
11. Install the hardware.
For server-specific procedures, see the server user guide.
12. For all the controllers except HPE SR416ie-m Gen11 SPDM Storage Controller, connect one end of the controller backup power cable to the backup power connector on the controller and the other end to the controller backup power connector on the system board or PCI riser board.
13. For all the controllers, install the optional energy pack.
14. Connect storage devices to the controller.
15. Connect peripheral devices to the server.
16. Connect the power cord to the server.
17. Connect the power cord to the power source.
18. Power up all peripheral devices.
19. Power up the server.
20. If you are running the server in UEFI Boot Mode, power on and select the boot options.
21. Update the controller firmware if it is not the latest revision.
22. Update the drive firmware if it not the latest revision.
23. If the new controller is not the new boot device, install the device drivers.
24. (Optional) Create additional logical drives.

Installing an HPE Synergy mezzanine controller (-m)

About this task

For information on how to install the HPE SR416ie-m Gen11 SPDM Storage Controller, see the [HPE Synergy 480 Gen11 User Guide](#).

Installing a standup PCIe plug-in controller (-p)

About this task

**WARNING:**

To reduce the risk of personal injury or damage to the equipment, consult the safety information and user documentation provided with the server before attempting the installation. Some servers contain high energy circuits, high current circuits, moving parts (such as fan blades), or any combination of these hazards, that may be exposed if covers and access panels are removed while the product is connected to a power source. These products are intended to be serviced only by qualified personnel who have been trained to deal with these hazards. Do not remove enclosures or attempt to bypass any interlocks designed to guard against these hazardous conditions.

Procedure

1. Perform a complete backup of all server data.
2. Remove or open the access panel.

**WARNING:**

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

3. Select an available x16 for 900 series, or a larger PCIe expansion slot.



NOTE: For more information on where the card can be installed and restrictions, refer server [QuickSpecs](#) and user guide available here: <https://www.hpe.com/support/hpesc>.

4. Remove the slot cover.

Save the retaining screw, if one is present.

5. 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.
6. 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.
7. Connect one end of the controller backup power cable to the backup power connector on the controller and the other end to the controller backup power connector on the system board or PCI riser board. To determine the location of the connector, see the server user guide.
8. Connect storage devices to the controller.
9. Close or install the access panel.

Before powering on the system, be sure the energy pack is installed. For more information, see the server user guide.

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 SAS or SATA or NVMe.
- 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 the following resources:

- Server documentation
 - Drive documentation
3. Use the internal cable identified in the server 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.

Cable part numbers

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

Configuration of HPE SR Gen11 Controller

This chapter tells you about the tools and procedures to set up the controller.

Subtopics

[Configuration tools for HPE SR controllers](#)

[Array and controller configuration](#)

[HPE iLO GUI](#)

[Intelligent Provisioning](#)

[HPE OneView](#)

[Redfish](#)

[Smart Storage Administrator](#)

[UEFI System Utilities](#)

[Powering on and selecting boot options](#)

Configuration tools for HPE SR controllers

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





[HPE iLO GUI](#)

[HPE iLO 6 User Guide](#)



[HPE OneView](#)

[HPE OneView User Guide for VM](#)

[HPE OneView User Guide for Synergy](#)



[Intelligent Provisioning](#)

[Intelligent Provisioning 4.0 User Guide for HPE ProLiant Gen11 Servers and HPE Synergy](#)



[Redfish](#)

[HPE Storage Controllers – Management overview](#)



[Smart Storage Administrator GUI](#)

[HPE Smart Storage Administrator GUI User Guide](#)



[Smart Storage Administrator CLI & Smart Storage Administrator Scripting](#)

[HPE Smart Storage Administrator CLI User Guide](#)



[UEFI System Utilities](#)

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

Array and controller configuration

You can configure arrays and controllers during the initial provisioning of the server or compute module and at any time after the initial configuration. Configuration tasks can be initiated using Smart Storage Administrator (SSA) (accessible through Intelligent Provisioning) or the configuration menus of the UEFI System Utilities.

During the initial provisioning of the server or compute module, an array is required to be configured before the operating system can be installed. You can configure the array using either of the options below:

- When you launch Intelligent Provisioning, you can specify options that enable Intelligent Provisioning to poll for any drives that are present and build an appropriate array for those drives. For example, if two drives are connected to the card, the setup defaults to RAID 1. Hewlett Packard Enterprise recommends selecting this option when initially provisioning a server. For more information, see the Intelligent Provisioning documentation available here: <https://www.hpe.com/support/hpesc>.
- You can use the UEFI System Utilities to create the primary array that is required.

After the initial provisioning of the server or compute module, you can use either SSA or the UEFI System Utilities to configure the arrays and controllers.

Comparison of SSA and UEFI System Utilities

This controller can be configured by using either SSA or the configuration utility within the UEFI System Utilities. Both SSA and UEFI System Utilities can be used to configure the controller.

SSA provides a full set of array configuration features while the UEFI System Utilities provides a limited set of features. However, users may prefer using the UEFI System Utilities during the initial configuration of the server or compute module because the UEFI System Utilities loads faster than SSA during that step.

To identify the standard configuration tasks that are supported within each interface, review the table.

Task	SSA	UEFI System Utilities
Create or delete arrays and logical drives	+	+
Assign a RAID level to a logical drive	+	+
Identify devices by causing the LEDs to illuminate	+	+
Assign or delete a spare drive	+	+
Share a spare drive among several arrays	+	+
Assign multiple spare drives to an array	+	+
Set the spare activation mode	+	+
Specify the size of the logical drive	+	+
Create multiple logical drives per array	+	+
Set the strip size	+	+
Migrate the RAID level or strip size	+	
Expand an array	+	
Set the expand priority and migrate priority	+	
Set the cache ratio (accelerator) priority	+	+
Extend a logical drive	+	
Set the boot controller	+	
Enable HPE SR SmartCache	+	
Configure HPE SR SmartCache	+	+
Enable/Configure Controller Based Encryption (CBE)	+	+
Erase Drives	+	+

For specific information about how to use either SSA or the UEFI System Utilities, see the online help.

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 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

Subtopics

[DMTF Redfish Storage Model](#)

DMTF Redfish Storage Model

These controllers support the DMTF standard known as PLDM for Redfish Device Enablement in a HPE ProLiant Gen10 server and beyond. 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.

These controllers have implemented the DMTF Redfish storage data model for inventory (GET). They currently support Redfish write operations (POST, DEL, and PATCH).

GET requests

The following table lists the Redfish resources for the 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 ¹	GET	/redfish/v1/Systems/{ID}/Storage/{ID}/Volumes/{ID}
Drive	GET	/redfish/v1/Systems/{ID}/Storage/{ID}/Drives/{ID}

¹ For controller firmware version 3.01.23.72, a Redfish volume with RAIDType set to None is automatically generated for each HBA drive. When the corresponding HBA drive is assigned to a RAID volume, the None RAIDType volume is deleted automatically.

Write requests

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



Redfish Resource	Method	URL
Volume Create	POST	POST /redfish/v1/Systems/{ID}/Storage/{ID}/Volumes
Volume Delete	DEL	DEL /redfish/v1/Systems/{ID}/Storage/{ID}/Volumes/{ID}
Drive Secure Erase	POST	/redfish/v1/Systems/{ID}/Storage/{ID}/Drives/{ID}/Actions/Drive.SecureErase
Storage Reset to Defaults	POST	/redfish/v1/Systems/{ID}/Storage/{ID}/Actions/Storage.ResetToDefaults
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 IO Perf Mode Enabled	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}
Storage Controller Consistency Check Rate Percent	PATCH	/redfish/v1/Systems/{ID}/Storage/{ID}/Controllers/{ID}
Storage Controller Rebuild Rate Percent	PATCH	/redfish/v1/Systems/{ID}/Storage/{ID}/Controllers/{ID}
Storage Controller Transformation Rate Percent	PATCH	/redfish/v1/Systems/{ID}/Storage/{ID}/Controllers/{ID}

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

Smart Storage Administrator

Smart Storage Administrator (SSA) is the main tool for configuring arrays on these controllers. It exists in three interface formats: the SSA GUI, the SSA CLI, and SSA Scripting. All these formats supports configuration tasks. Some of the advanced tasks are available in only one format.

Smart Storage Administrator GUI

SSA GUI is an advanced utility that enables you to perform many complex configuration tasks. The SSA GUI is accessible both offline and online:

- **Accessing SSA in the offline environment:** Using one of multiple methods, you can run SSA before launching the host OS. In offline mode, users can configure or maintain detected and supported HPE ProLiant devices, such as optional controllers and integrated controllers. Some SSA features are only available in the offline environment, such as setting the boot controller or performing split-mirror operations.
- **Accessing SSA in the online environment:** This method requires an administrator to download the SSA executables and install them. You can run SSA online after launching the host OS.

For more information, see [HPE Smart Storage Administrator User Guide](#)

Smart Storage Administrator CLI (SSA CLI)

SSA CLI is a command line interface tool used to manage HPE SR Storage Controller products. The storage CLI tool supports scripting for mass deployment of server storage.

For more information, see [HPE Smart Storage Administrator CLI User Guide](#)

Smart Storage Administrator Scripting

SSA Scripting is a standalone application that is distributed with the HPE SSA CLI application and is used for configuring arrays on Smart Array devices.

UEFI System Utilities

The UEFI System Utilities is embedded in the system ROM. It 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

For more information on the UEFI System Utilities, see the [HPE UEFI System Utilities Documentation Quick Links](#).

For on-screen 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 ¹	F9 during server POST
Navigate menus	Up and Down arrows
Select items	Enter
Save selections	F10
Access Help for a highlighted configuration option ²	F1

¹ When boot password is enabled for a controller, there is no POST stop and enter boot password prompt. Users must press **F9** to enter RBSU, and unlock the encrypted drives. Otherwise, the system is booted with locked encrypted drives.

² 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.

Powering on and selecting boot options

About this task

HPE Gen11 servers operate in UEFI Boot Mode. The boot controller and boot order are set automatically.

Procedure

1. Press the Power On/Standby button.
2. During the initial boot:
 - a. To modify the server configuration ROM default settings, press **F9** in the ProLiant POST screen to enter the UEFI System Utilities screen. By default, the System Utilities menus are in the English language.
 - b. If you do not need to modify the server configuration and are ready to install the system software, press **F10** to access Intelligent Provisioning.

For more information on automatic configuration, see UEFI documentation on the <https://www.hpe.com/support/hpesc>.

Maintenance

Subtopics



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.

For information on how to update firmware for the HPE SR416ie-m Gen11 SPDM Storage Controller, see "Drivers, firmware, and software updates_IP and SPP" in the [HPE Synergy 480 Gen11 User Guide](#).

Error reporting

- **Integrated Management Log**

The controller uses PLDM over Redfish device enablement to send Redfish events, which result in error messages in the HPE Integrated Management Log (IML). These events are generated by the device firmware and may occur during power up or run-time. The IML log can be viewed from the HPE iLO GUI.

- **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>cpqDa6AccelStatusChange</code>	Accelerator status change
<code>cpqDa7PhyDrvStatusChange</code>	Drive 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:

- BatteryCharging
- BatteryFailure
- BatteryMissing
- BatteryOK
- ControllerDegraded
- ControllerFailure
- ControllerPasswordAccepted
- ControllerPasswordRequired
- ControllerPreviousError
- DriveFailure
- DriveFailureCleared
- DriveInserted
- DriveMissing

- DriveMissingCleared
- DriveOffline
- DriveOfflineCleared
- DriveOK
- DrivePredictiveFailure
- DriveRemoved
- VolumeDegraded
- VolumeFailure
- VolumeOffline
- VolumeOfflineCleared
- VolumeOK
- WriteCacheDataLoss
- WriteCacheDegraded
- WriteCacheProtected
- WriteCacheTemporarilyDegraded

- **System Event Log**

The controller uses PLDM over Redfish device enablement to send Redfish events, which results 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 on storage controller products, see <https://www.hpe.com/info/SCMO>.

Diagnostic tools

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

- **Smart Storage Administrator (SSA)**

SSA can be accessed offline using Intelligent Provisioning or booting from the SPP ISO image. It can also be accessed online by downloading the SSA version 6.10.14.0 or later. For more information on using SSA, see [HPE Smart Storage Administrator User Guide](#).

- **HPE iLO**

The iLO firmware continuously monitors the controller independent of the operating system and logs any failure events to the IML, iLO RESTful API, and SNMP. In addition, the iLO web interface allows users to view the status of the controller and its attached devices.

- **UEFI System Utilities**

The UEFI System Utilities is embedded in the system ROM. The UEFI System Utilities enable you to view controller configuration and settings. For more information, see [UEFI System Utilities User Guide for HPE ProLiant Gen11 Servers and HPE Synergy](#).

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.

Models

Subtopics

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


[HPE Synergy mezzanine controller \(-m\)](#)

Standup PCIe Plug-In Controller (-p)

Subtopics

[HPE SR932i-p Gen11 SPDM Storage Controller ports and connectors](#)

HPE SR932i-p Gen11 SPDM Storage Controller ports and connectors



Item Description

- | | |
|---|------------------------------|
| 1 | Internal x8 SlimSAS port 1 |
| 2 | Internal x8 SlimSAS port 2 |
| 3 | Internal x8 SlimSAS port 3 |
| 4 | Internal x8 SlimSAS port 4 |
| 5 | Backup power cable connector |


HPE Synergy mezzanine controller (-m)

Subtopics

[HPE SR416ie-m Gen11 SPDM Storage Controller ports and connectors](#)



HPE SR416ie-m Gen11 SPDM Storage Controller ports and connectors



Item Description

- 1 Internal x8 MCIO port 1

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.



NOTE:

- The energy pack must only be installed, removed, or replaced while the server is powered off and AC power cords are removed.
- HPE SR416ie-m Gen11 SPDM Storage Controller does not support the three-pin controller backup power cable as power is delivered through the PCI connector.

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 an unlimited number of devices with the 96W battery. A 12W battery supports up to two controllers. The time required to recharge is two hours for a 96W battery and one hour for a 12W battery. For more information, see [HPE Smart Storage Batteries and Hybrid Capacitors](#).

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 two controllers. The time required to recharge takes less than one minute. For more information, see [HPE Smart Storage Batteries and Hybrid Capacitors](#).

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 1 (Triple)
- RAID 10 (Triple)

RAID 1T and 10T are also known as RAID 1 Triple and RAID 10 Triple, and was previously known as RAID 1/10 Advanced Data Mirror (ADM).

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

HPE Storage Controller documentation resources

General websites

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

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

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

User guides for configuration tools

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

QuickSpecs for hardware RAID controllers

Content	Microchip® SmartROC 3100	Microchip® SmartROC 3200	Broadcom® Aero-16
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

QuickSpecs for Boot devices and virtual RAID

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

Technical papers

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

Training videos



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

Subtopics

[HPE Storage Controller documentation quick links](#)

HPE Storage Controller documentation quick links

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



Documents	Versions		
HPE SR Controllers			
User Guide	Gen10	Gen10 Plus	Gen11
GUI User Guide	Multi-gen	—	—
CLI User Guide	Multi-gen	—	—
Alerts - Customer Advisories	—	Gen10 Plus	Gen11
HPE MR Controllers			
User Guide	Gen10	Gen10 Plus	Gen11
GUI User Guide	Multi-gen	—	—
CLI User Guide	Multi-gen	—	—
Alerts - Customer Advisories	Gen10	Gen10 Plus	Gen11
Intel Virtual Raid on CPU for HPE			
User Guide	—	Gen10 Plus - Linux	Gen11
		Gen10 Plus - VMware	
		Gen10 Plus - Windows	
Alerts - Customer Advisories	Multi-gen		
HPE NS204i Boot Devices			
User Guide	Multi-gen	—	—
Alerts - Customer Advisories	Multi-gen		

Support and other resources

Subtopics

- [Accessing Hewlett Packard Enterprise Support](#)
- [HPE product registration](#)
- [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

HPE product registration

To gain the full benefits of the Hewlett Packard Enterprise Support Center and your purchased support services, add your contracts and products to your account on the HPESC.

- When you add your contracts and products, you receive enhanced personalization, workspace alerts, insights through the dashboards, and easier management of your environment.
- You will also receive recommendations and tailored product knowledge to self-solve any issues, as well as streamlined case creation for faster time to resolution when you must create a case.

To learn how to add your contracts and products, see <https://www.hpe.com/info/add-products-contracts>.

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 Service

<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.

