Hewlett Packard Enterprise

HPE Smart Array SR Gen10 User Guide

Abstract

This document includes feature, installation, and configuration information about Hewlett Packard Enterprise Smart Array SR Gen10 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: 879006-004a
Published: May 2019
Edition: 4a

© Copyright 2017, 2019 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

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

Contents

HPE Smart Array SR Gen10 6
S-class 6
E-class 7
P-class 7

Features 9
Features support 9
Operating environments 9
RAID technologies 9
Transformation 10
Drive technology 10
Security 11
Reliability 11
Performance 11
Controller supported features 12
RAID technologies 12
Selecting the right RAID type for your IT infrastructure 12
Mixed mode (RAID and HBA simultaneously) 15
Striping 16
Mirroring 16
Parity 19
Spare drives 23
Drive Rebuild 24
Transformation 26
Array transformations 26
Logical drive transformations 27
Transformation priority 28
Drive technology 28
Predictive drive failure 28
Online drive firmware update 28
Dynamic sector repair 28
Controller surface scan 28
Shingled magnetic recording 29
HPE SmartDrive LED 29
SSD over-provisioning optimization 31
SSD Wear Gauge reports 31
Security 31
HPE Smart Array SR Secure Encryption 31
Sanitize erase 32
Sanitize freeze lock and anti-freeze lock 33
Reliability 33
Dual Domain Topologies 33
Link error monitoring 34
Recovery ROM 34
Cache Error Checking and Correction (ECC) 34
Thermal monitoring 34
Performance 34
HPE Smart Array SR SmartCache 34
SSD Smart Path 35

Dual domain path selection....35
Cache....35
Drive write cache control....37
Video on demand....37
Stripe size selection....38
Power modes....38

Installation, configuration, and maintenance....40
Installation....40
Supported servers....40
Installing a Smart Array in an unconfigured server....40
Installing a Smart Array in a previously configured server....41
Installing a Smart Array....42
Connecting storage devices....45
Cable part numbers....47
Enabling Smart Array SW RAID....47
Device drivers....47
Windows operating systems....47
Configuration....48
Array and controller configuration....48
HPE Smart Storage Administrator....49
UEFI System Utilities....50
Intelligent Provisioning....50
Configuring boot controller options....51
System maintenance tools....52
Updating software and firmware....52
Diagnostic tools....53

Models....55
Modular Smart Array (-a/-b/-c)....55
HPE Smart Array E208i-a SR Gen10....55
HPE Smart Array P408i-a SR Gen10....56
HPE Smart Array P816i-a SR Gen10....59
HPE Smart Array P204i-b SR Gen10....62
HPE Smart Array E208i-c SR Gen10....65
HPE Smart Array P204i-c SR Gen10....65
HPE Smart Array P408i-c SR Gen10....66
Standup PCIe Plug-In Smart Array (-p)....66
HPE Smart Array E208i-p SR Gen10....66
HPE Smart Array E208e-p SR Gen10....67
HPE Smart Array P408i-p SR Gen10....69
HPE Smart Array P408e-p SR Gen10....72
Mezzanine controllers (-m)....75
HPE Smart Array P408e-m SR Gen10....75
HPE Smart Array P416ie-m SR Gen10....77

Additional hardware and options....81
Energy pack options....81
HPE Smart Storage Battery....81
HPE Smart Storage Hybrid Capacitor....81
Energy pack specifications....82
HPE 12G SAS Expander Card....82

Specifications.....83
Memory and storage capacity conventions.....83
RAID conventions.....83
Controller specifications.....83
Energy pack specifications.....83

Support and other resources.....85
Accessing Hewlett Packard Enterprise Support.....85
Accessing updates.....85
Customer self repair.....86
Remote support.....86
Warranty information.....86
Regulatory information.....87
Documentation feedback.....87

Websites.....88

HPE Smart Array SR Gen10

HPE Smart Array SR Gen10 offers a reliable family of RAID controllers that attach to:

- Internal hot-plug drives
- Internal non hot-plug drives
- External JBODs to HPE Gen10 ProLiant, Synergy, and Apollo servers

The HPE Smart Array SR Gen10 family includes S-class, E-class, and P-class integrated within a common set of Smart Array SmartRAID (SR) management tools. Each class is characterized according to software features of RAID support, RAID levels, features, and performance, and hardware features of SAS/SATA lanes, port type, and form factor.

The type-a, type-b, and type-c designations indicate the server and compute module platforms that are supported. Specifically:

- Type-a modular controllers are compatible with ProLiant DL, ProLiant ML, and Apollo platforms.
- Type-b modular controllers are compatible with ProLiant BL platforms.
- Type-c modular controllers are compatible with HPE Synergy platforms.

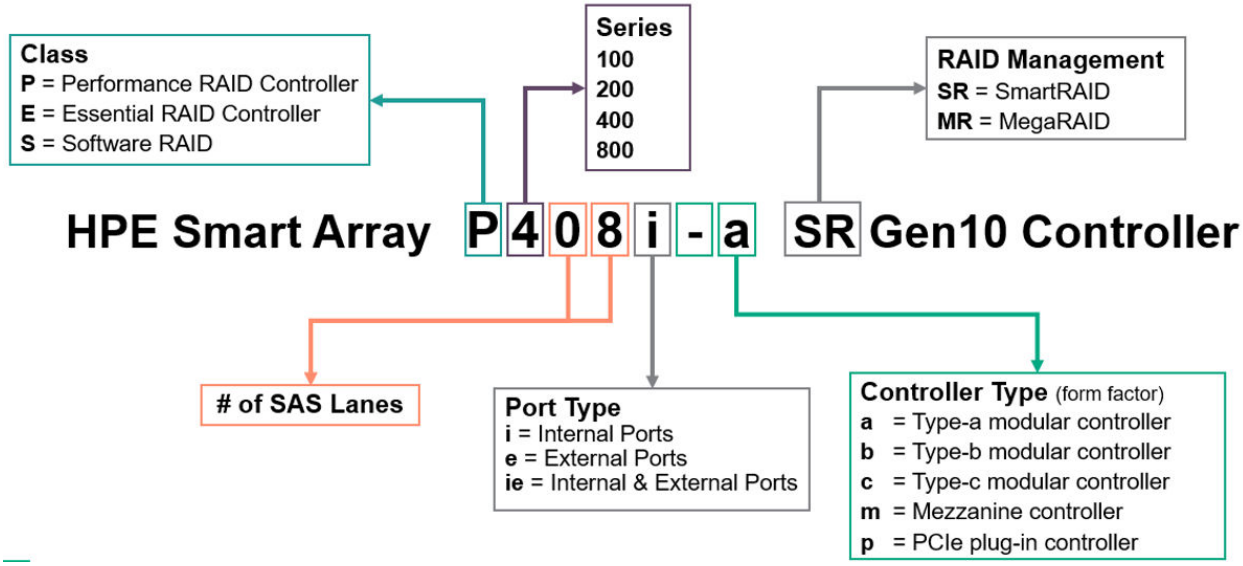

S-class

S-class provides software RAID capabilities for use with Microsoft Windows operating systems. HPE Smart Array S100i SR Gen10 SW RAID is an ideal entry-level solution that use SATA drives in basic RAID configurations.

S-class provides

- Up to 14 SATA lanes attached to internal drives
- RAID levels 0, 1, 5, and 10
- Hot-plug and non hot-plug SATA drive support
- 6G SATA support

- UEFI Boot Mode Only
- Windows Server 2012 R2, Windows Server 2016, and Windows Server 2019
- System memory used as read cache
- Smart Array management tools

NOTE: Currently only offline HPE Smart Storage Administrator is supported in a pre-boot environment. HPE Smart Storage Administrator available through Service Pack for ProLiant is not supported.

Name	Supported HPE Gen10 servers
HPE Smart Array S100i SR Gen10 SW RAID	ProLiant, Apollo, Synergy, BladeSystem

E-class

E-class Smart Array controllers provide an enterprise level, cost-effective solution for RAID 0, 1, 5, and 10 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 Smart Array SR Secure Encryption and provide enterprise-class reliability, security and efficiency.

E-class HPE Smart Array SR Gen10 provides:

- Up to 8 SAS/SATA lanes for internal or external drives
- RAID levels 0, 1, 5, and 10
- Mixed mode (RAID and HBA pass-through functionality simultaneously)
- Data-at-rest encryption
- 12G SAS support
- HPE 12G SAS Expander Card support
- UEFI and Legacy Boot modes
- No cache memory support
- Smart Array management tools

Name	Supported HPE Gen10 servers
HPE Smart Array E208i-a SR Gen10	ProLiant and Apollo
HPE Smart Array E208i-p SR Gen10	ProLiant and Apollo
HPE Smart Array E208e-p SR Gen10	ProLiant and Apollo
HPE Smart Array E208i-c SR Gen10	Synergy

Feature	S-class	E-class	P-class
Shingled Magnetic Recording (SMR)	--	✔	✔
HPE SmartDrive LED	✔	✔	✔
SSD Over-Provisioning Optimization	--	✔	✔
SSD Wear Gauge reports	✔	✔	✔

[1] Reboot is required

Security

Feature	S-class	E-class	P-class
HPE Smart Array SR Secure Encryption	--	✔	✔
Sanitize Erase	--	✔	✔
Sanitize Freeze Lock	--	✔	✔
Signed Firmware	not applicable	✔	✔
Drive Authentication	✔	✔	✔

Reliability

Feature	S-class	E-class	P-class
Dual Domain Topologies	--	✔	✔
Link Error Monitoring	--	✔	✔
Recovery ROM	not applicable	✔	✔
Cache Error Checking and Correction	not applicable	✔	✔
Thermal Monitoring	✔	✔	✔

Performance

Feature	S-class	E-class	P-class
HPE Smart Array SR SmartCache	--	--	✔
SSD Smart Path	✔	✔	✔

Table Continued

Feature	S-class	E-class	P-class
Dual Domain Path Selection	--	✔	✔
Read Cache	✔	--	✔
Flash-Backed Write Cache	--	--	✔
Cache Ratio Selection	--	--	✔
Write Cache Bypass Threshold	--	--	✔
Drive Write Cache Control	✔	✔	✔
Video on Demand	--	✔	✔
Stripe Size Selection	✔	✔	✔
Power Modes	--	✔	✔

Controller supported features

The features supported by each Smart Array controller are described in the **Controller Family Datasheet** and in the following QuickSpecs:

- **S-class controller S100i**
- **All E-class and P-class controllers**

RAID technologies

Selecting the right RAID type for your IT infrastructure

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

- The number of parity groups that you have
- The fault tolerance required
- The write performance required
- The amount of usable capacity that you have

Configuring the RAID fault tolerance

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

This chart shows the relationship between the RAID level fault tolerance and the size of the storage array. The chart includes RAID 0, 5, 50, 10, 6, 60, and RAID 10 ADM. 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 ADM 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.

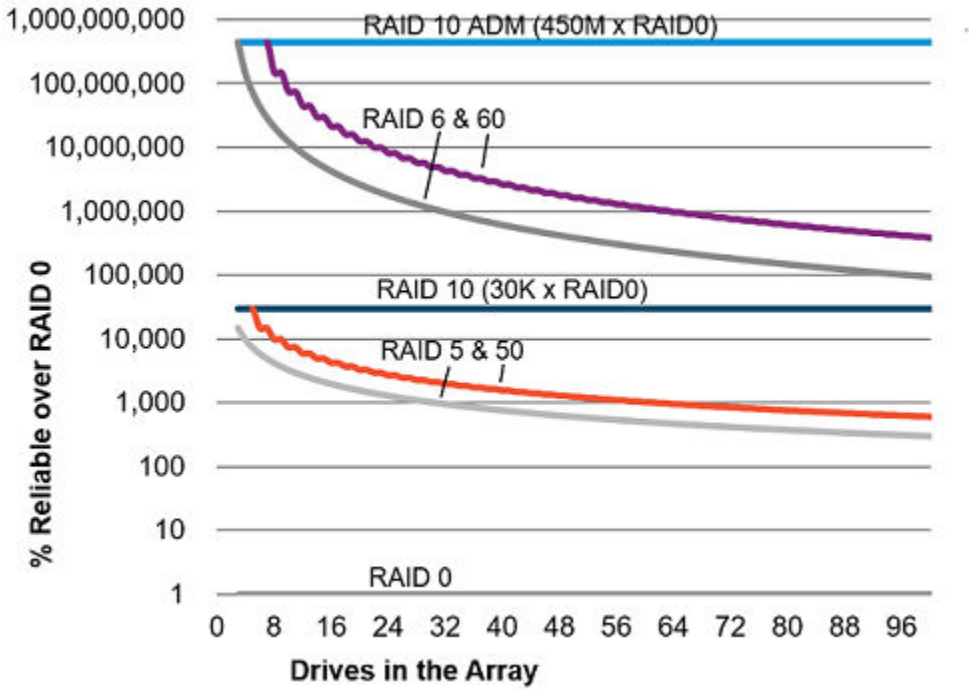

Configuring the RAID write performance

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

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

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

Consider the following points:

- Write performance decreases as fault tolerance improves due to extra I/O.
- Read performance is generally the same for all RAID levels except for smaller RAID 5\6 arrays.

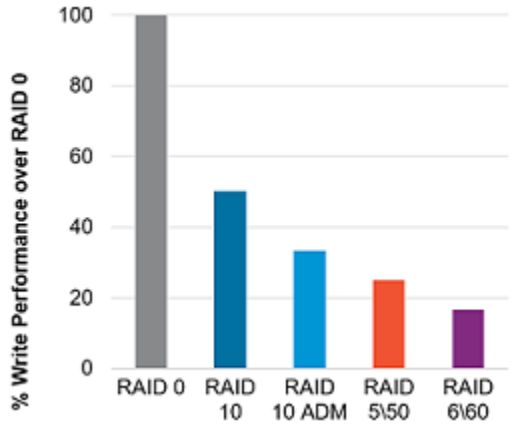

The table below shows the Disk I/O for every host write:

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

Configuring the RAID usable capacity

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

Consider the following points when selecting the RAID type:

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

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

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

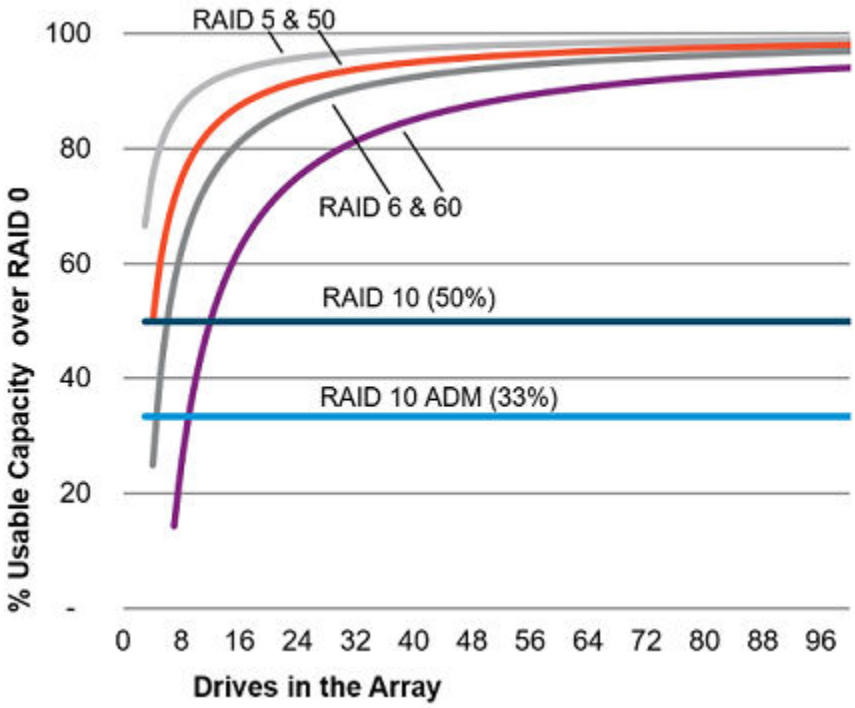

Configuring the storage solution

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

- RAID 10 ADM: 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.

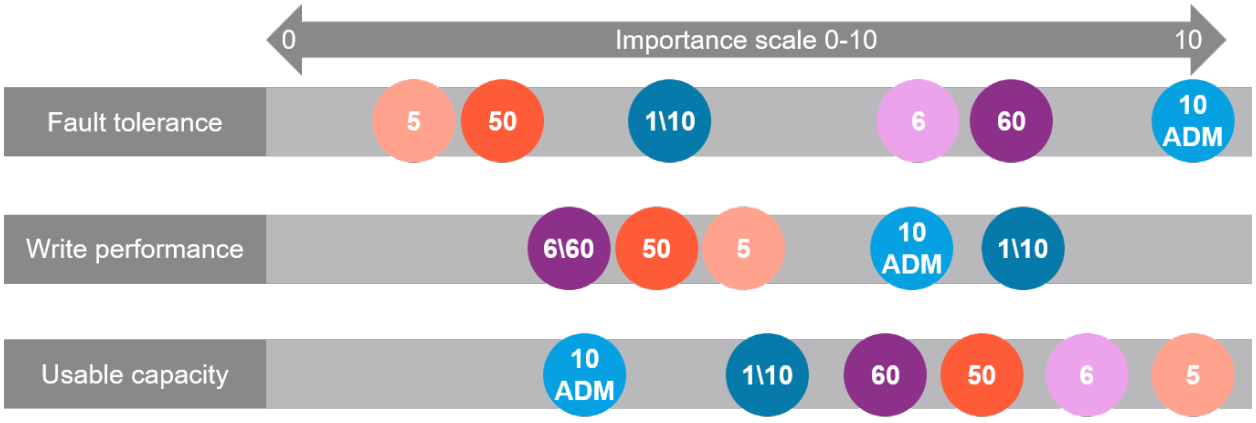

Mixed mode (RAID and HBA simultaneously)

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 and cannot be disabled. Logical drives are also presented to the operating system.

Controllers that support mixed mode (P-class and E-class) 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.

Striping

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.

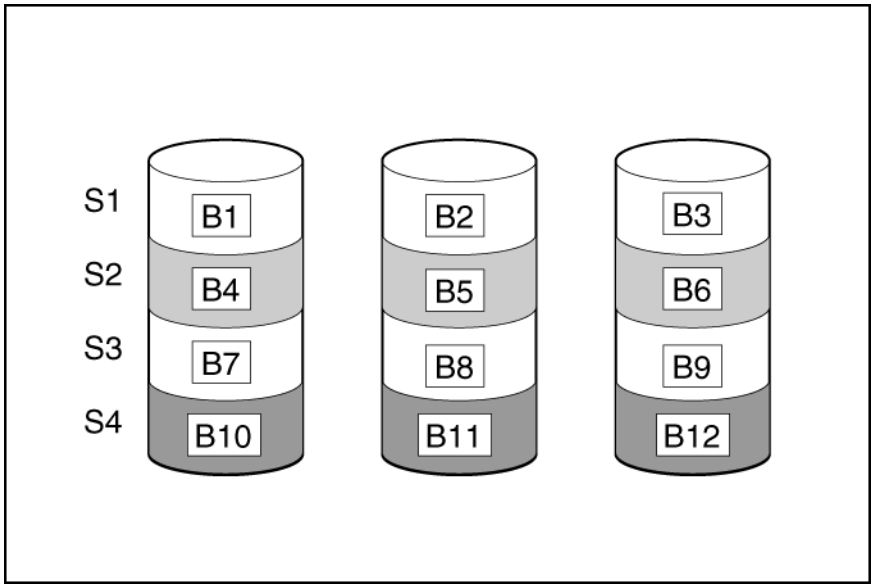

This method has the following benefits:

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

Mirroring

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 x (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.

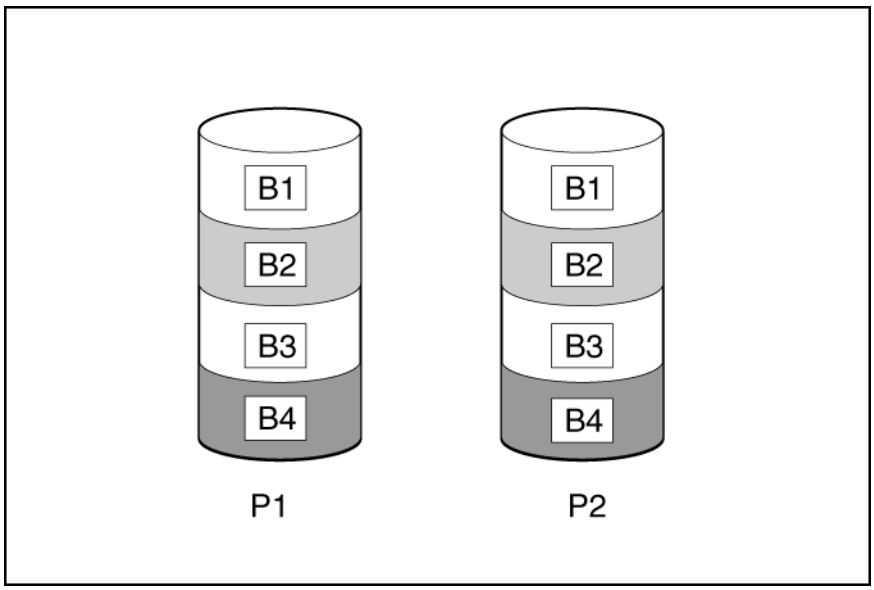

When the array has more than two physical drives, drives are mirrored in pairs, and the fault-tolerance method is known as RAID 1+0 or RAID 10. If a physical drive fails, the remaining drive in the mirrored pair can still provide all the necessary data. Several drives in the array can fail without incurring 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.

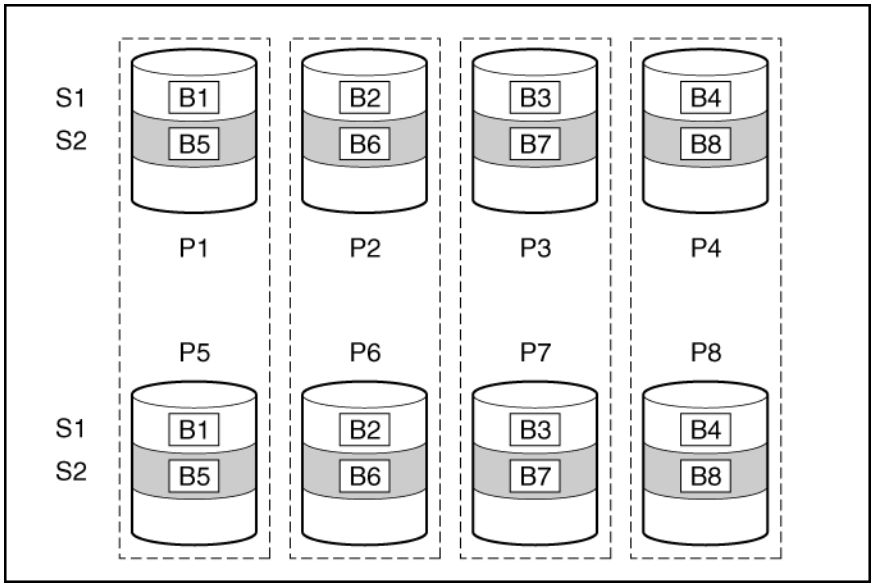

This method has the following benefits:

- 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 (ADM) and RAID 10 (ADM)

In RAID 1 (Advanced Data Mirroring (ADM)) and RAID 10 (ADM) configurations, data is duplicated to two additional drives. The usable capacity is C x (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 (ADM).

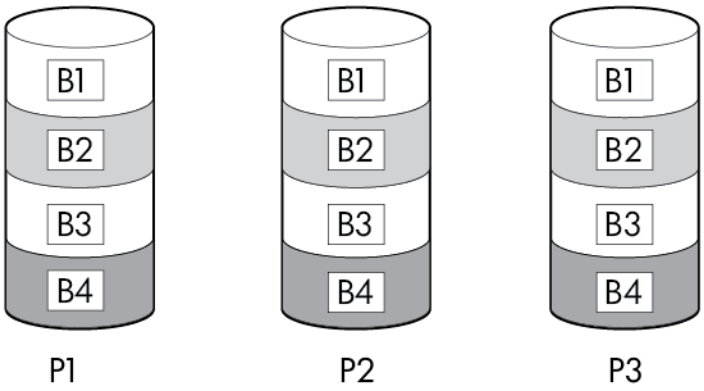

When the array has more than six physical drives, drives are mirrored in trios, and the fault-tolerance method is known as RAID 10 (ADM). If a physical drive fails, the remaining two drives in the mirrored trio can still provide all the necessary data. Several drives in the array can fail without incurring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

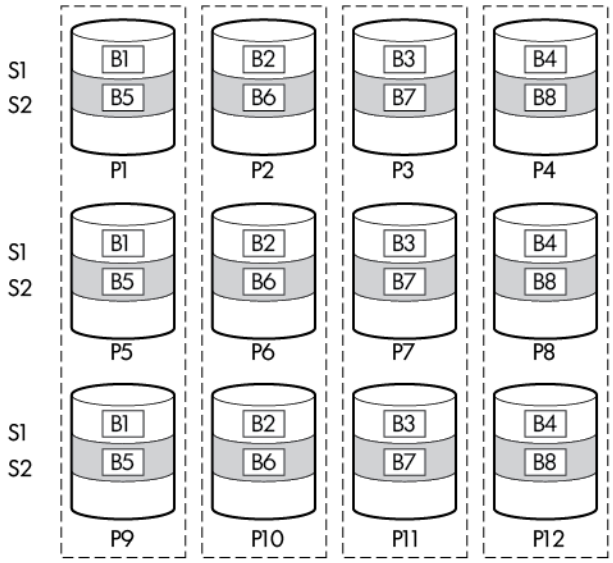

This method has the following benefits:

- 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, Smart Array balances read requests between drives based upon 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 ADM, or 10 ADM) 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 (ADM) and RAID 10 (ADM), 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 (ADM) logical drives
- one array with RAID 1+0 logical drives and one array with RAID 0 logical drives into one array with RAID 10 (ADM) logical drives

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

Parity

RAID 5

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

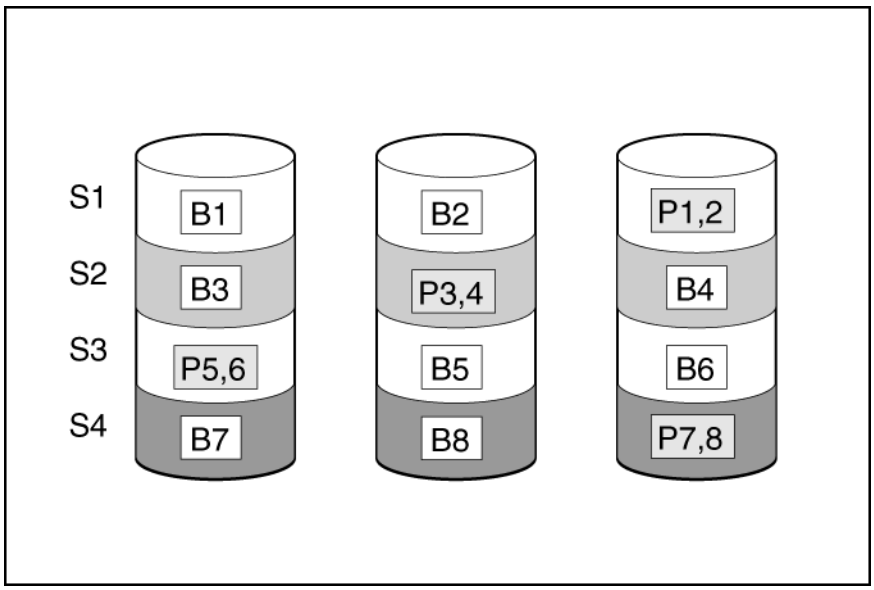

This method has the following benefits:

- 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.

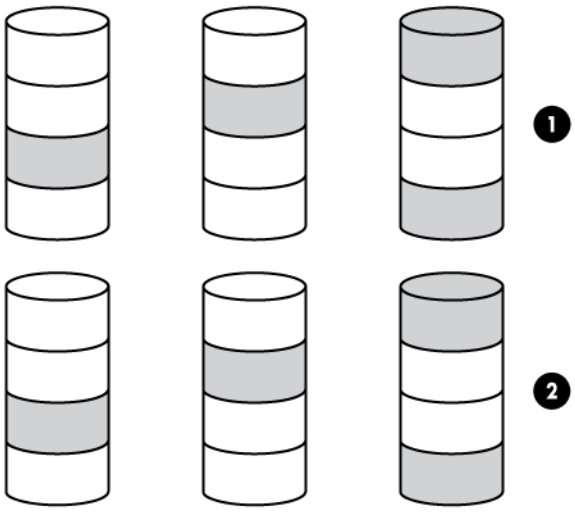

For any given number of drives, data loss is least likely to occur when the drives are arranged into the configuration that has the largest possible number of parity groups. For example, four parity groups of three drives are more secure than three parity groups of four drives. However, less data can 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 Px,y and Qx,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 x (n - 2) where C is the drive capacity with n drives in the array.

A minimum of 4 drives is required.

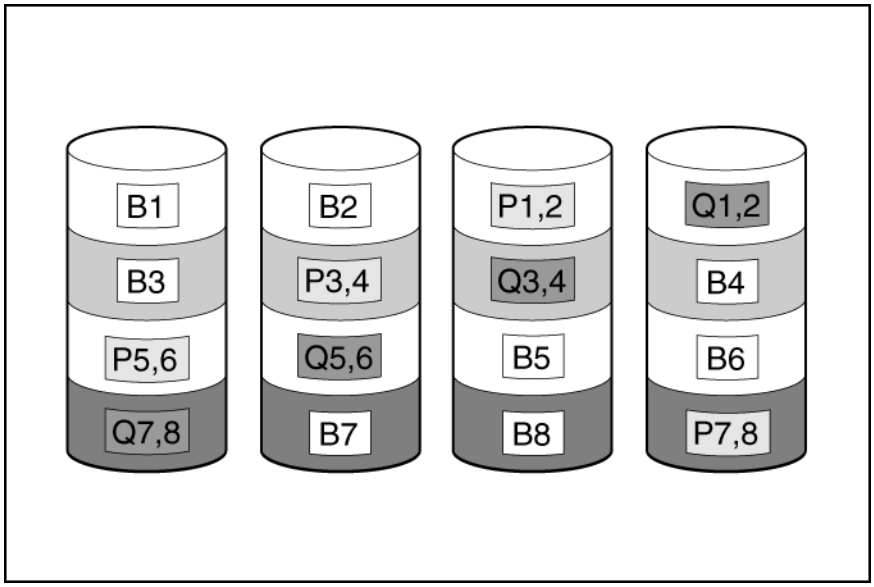

This method is most useful when data loss is unacceptable but cost is also an important factor. The probability that data loss will occur when an array is configured with RAID 6 (Advanced Data Guarding (ADG)) is less than it would be if it were configured with RAID 5.

This method has the following benefits:

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

RAID 60

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

For any given number of hard drives, data loss is least likely to occur when the drives are arranged into the configuration that has the largest possible number of parity groups. For example, five parity groups of

four drives are more secure than four parity groups of five drives. However, less data can be stored on the array with the larger number of parity groups.

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

A minimum of 8 drives is required.

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.

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

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

This feature has the following benefits:

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

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 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, thus speeding the initialization process, and enabling 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. Any time a member drive within an array is failed, all writes that map to the failed drive are regenerative. A regenerative write is much slower because it must read from nearly all of the drives in the array to calculate new parity data. The write penalty for a regenerative write is n + 1 drive operations where n is the total number of drives in the array. As you can see, the write penalty is greater (slower write performance) with larger arrays.

This method has the following benefits:

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

Backed-out writes

After parity initialization is complete, random writes to a RAID 5, 50, 6, or 60 can use a faster backed-out write operation. A backed-out write uses the existing parity to calculate the new parity data. As a result, the write penalty for RAID 5 and RAID 50 is always four drive operations, and the write penalty for a RAID 6 and RAID 60 is always six drive operations. As you can see, the write penalty is not influenced by the number of drives in the array.

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

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

Full-stripe writes

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

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

Spare drives

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 using 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

Rapid rebuild

Smart Array 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.

Rebuild priority

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

- At the low setting, normal system operations take priority over a rebuild.
- At the medium setting, rebuilding occurs for half of the time, and normal system operations occur for the rest of the time.
- At the medium high setting, rebuilding is given a higher priority over normal system operations.
- At the 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.

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 the **Hewlett Packard Enterprise website**.
- 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 (ADM) and RAID 10 (ADM) 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

Array transformations

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.

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(s) must have enough capacity to hold all the logical drives present in the source array.

Replace array

The Replace Array operation allows 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 OK state. All the existing logical drives in the source arrays must be in 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 may 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 may only remove an even number of drives.
- If the array contains a compound RAID (RAID 50 or RAID 60) logical drive, drives may 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 may 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 only available 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.

Logical drive transformations

Extend logical drive

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.

Migrate RAID level

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

Migrate stripe size

The migrate stripe size feature allows you to change the current strip size for your logical drive. When the strip size changes, you may have more or less unused space, depending on the strip size with which you started. For migration to a larger stripe size to be possible, the array might need 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

Predictive drive failure

HPE Smart Array controllers use Self-Monitoring and Reporting Technology (S.M.A.R.T to inform the host when a disk drive is experiencing abnormal operation likely to lead to drive failure.

S.M.A.R.T. places the monitoring capabilities within the disk drive itself. These monitoring routines have direct access to internal performance, calibration, and error measurements for a specific drive type.

Online drive firmware update

The latest generation HPE Smart Array controllers support online drive flashing, which saves time when updating disk drive firmware. Instead of taking the hard disk drive (HDD) offline before loading a new firmware image, you can download an updated HDD firmware image to the HPE Smart Array controller and update all of the HDDs the next time you reboot the server.

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, HPE built a dynamic sector repair feature into HPE Smart Array controllers.

HPE Smart Array 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

- checks physical drives in fault-tolerant logical drives for bad sectors
- in RAID 5 or RAID 6 (ADG) configurations, verifies the consistency of parity data

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 may 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 scan count allows the control of how many controller surface scans can operate in parallel. This is used when there is more than one logical drive on a controller. This setting allows the controller to detect bad blocks on multiple logical drives in parallel and can significantly decrease the time it takes to detect back, especially for logical drives using very large capacity drives.

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

HPE SmartDrive LED

HPE SmartDrives are the latest Hewlett Packard Enterprise drive technology. Identify a SmartDrive by its carrier, shown in the following illustration.

When a drive is configured as a part of an array and connected to a powered-up controller, the drive LEDs indicate the condition of the drive.

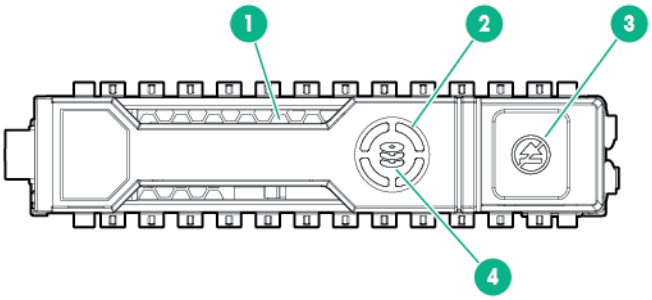

Item	LED	Status	Definition
1	Locate[1]	Solid blue	The drive is being identified by a host application.
		Flashing blue	The drive carrier firmware is being updated or requires an update.
2	Activity ring	Rotating green	Drive activity
		Off	No drive activity
3	Do not remove	Solid white	Do not remove the drive. Removing the drive causes one or more of the logical drives to fail.
		Off	Removing the drive does not cause a logical drive to fail.
4	Drive status	Solid green	The drive is a member of one or more logical drives.
		Flashing green	The drive is doing one of the following: • Rebuilding • Performing a RAID migration • Performing a strip size migration • Performing a capacity expansion • Performing a logical drive extension • Erasing • Spare drive activation
		Flashing amber/green	The drive is a member of one or more logical drives and predicts the drive will fail.
		Flashing amber	The drive is not configured and predicts the drive will fail.
		Solid amber	The drive has failed.
		Off	The drive is not configured by a RAID controller or is a spare drive.

[1] The blue Locate LED is behind the release lever and is visible when illuminated.

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 Smart Array 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 solid state drives 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

IMPORTANT: HPE Special Reminder: Before enabling encryption on the Smart Array controller module on this system, you must ensure that your intended use of the encryption complies with relevant local laws, regulations and policies, and approvals or licenses must be obtained if applicable.

For any compliance issues arising from your operation/usage of encryption within the Smart Array controller module which violates the above mentioned requirement, you shall bear all the liabilities wholly and solely. HPE will not be responsible for any related liabilities.

HPE Smart Array SR Secure Encryption

HPE Smart Array SR Secure Encryption is a controller-based, enterprise-class data encryption solution that protects data at rest on any SAS/SATA drive attached to the HPE Smart Array controller. The solution is available for both local and remote deployments.

HPE Smart Array SR Secure Encryption is configured using the HPE Smart Storage Administrator (SSA).

Prerequisites:

- An installed Smart Array Controller compatible with Secure Encryption
- A valid Secure Encryption license for each server to be encrypted.

Local Key Management Mode

Local Key Management Mode, or Local Mode, is a solution designed for small to medium-size data centers. The solution utilizes a paraphrase 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. For more information, see *HPE Smart Array SR Secure Encryption Installation and User Guide*.

This method has the following benefits:

- Encrypts data on both the attached bulk storage and the cache memory of Smart Array 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 Smart Array 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, *HPE Smart Array SR Secure Encryption Installation and User Guide*.

Prerequisites:

- Integrated Lights Out (iLO) Advanced or Scale Out Edition license, per 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 Smart Array Controllers
- Supports any HDD or SSD in the HPE server portfolio
- Keys are kept in separate storage from servers to protect against physical removal

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 may still be present on the drive. Not all drives support unrestricted.

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**.

Sanitize overwrite (HDD)

Sanitize overwrite 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 (SSD)

Sanitize block erase 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 freeze lock and anti-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.

For more information about sanitize, see "**Sanitize erase**".

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 a Smart Array SR controller.

Reliability

Dual Domain Topologies

Dual Domain support for SAS creates redundant pathways from servers to storage devices. These redundant paths reduce or eliminate single points of failure within the storage network and increase data availability. Dual Domain SAS implementations can tolerate HBA failure, external cable failure, expander failure, or failure in a spanned disk (JBOD).

NOTE: This feature is not supported for SATA drives.

Link error monitoring

Smart Array 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. Smart Array 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, Smart Array 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

HPE Smart Array controllers store a redundant copy of the controller firmware image to protect against data corruption. If the active firmware image becomes corrupt, HPE Smart Array controllers use the redundant firmware image and continue operating. The recovery ROM provides protection against power outages during firmware flashing.

Cache Error Checking and Correction (ECC)

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 Smart Array memory controller uses this information to detect and correct data errors originating inside the DRAM chip or across the memory bus.

Thermal monitoring

Smart Array monitors the temperature of each drive in the server. iLO periodically collects these drive temperatures from Smart Array 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

HPE Smart Array SR SmartCache

SmartCache enables solid-state drives to be used as caching devices for hard drive media. SmartCache provides the following features:

- Accelerates application performance.
- Provides lower latency for transactions in applications.
- Supports all operating systems without the need for changes to OS, driver, or applications.

SmartCache will be fully enabled after the first SmartCache is created on the controller.

The following features are not available unless SmartCache is disabled:

- Expand Array
- Advanced Capacity Expansion
- Move Logical Drive

- Change Array Drive Type
- Mirror Splitting and Recombining (offline only)
- Split Mirror Backup and Rollback (online and offline)
- Heal Array
- Extend Logical Drive
- Migrate RAID/Strip Size
- Change Cache Ratio
- Align Logical Drives

SmartCache requires an energy pack and a SmartCache license. If you purchase the HPE Smart Array P816i-a SR Gen10, a license is included.

For more information, see the **Hewlett Packard Enterprise website**.

SSD Smart Path

SSDs require special tactics to capture the full advantage of their low-latency capabilities. HPE SSD Smart Path enables the 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. SSD Smart Path is enabled by default when you create an array.

The HPE Smart Array device driver software coordinates with HPE Smart Array controller firmware to

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

All other requests, as well as any error handling, are still routed through the normal IO path on the HPE Smart Array controller. This method has the following benefits:

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

Dual domain path selection

When using a dual domain SAS topology, Smart Array has 2 paths that it could use to access each SAS drive. During power-up, Smart Array selects an optimal path based upon the path with the fewest number of SAS expanders between the controller and the drive. If both paths contain the same number of expander hops, then Smart Array distributes half of the drives to each path. For example, odd bays use path 1, while even bays use path 2.

This method has the benefit of optimizing performance by allowing the controller to use the additional bandwidth available in the secondary drive path.

Cache

Read cache

The HPE Smart Array 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. HPE Smart Array 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, HPE Smart Array 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 HPE Smart Array 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 HPE Smart Array controller puts that data into the cache, so it is available if the upcoming read commands call for it.

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

The read-ahead cache feature requires an energy pack.

Flash-backed write cache

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

You can use HPE Smart Storage Administrator to configure the percentage of the cache to use for write caching. The default configuration on Smart Array controllers assigns 90% of the available cache space for write cache.

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 controller has a cache that uses backup power (HPE Smart Storage Battery or HPE Smart Storage Hybrid Capacitor), and
- There are logical drives configured on the controller.

The default of 90% write to 10% read is the best ratio for most workloads. Workloads that are highly sequential reads or reads from most recent writes may 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/Os smaller than the threshold.

No-battery write cache

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

Drive write cache control

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

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

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. This differs from low latency optimizations that prioritize absolute lowest latency at the expense of variability and bandwidth. The HPE Smart Array offers several video streaming applicable optimizations. Additional 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 very 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 MediumHigh - 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.

Stripe size selection

When an HPE Smart Array 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. A "stripe" is one set of strips. HPE Smart Array controllers configure strips, not stripes. You can calculate the stripe size from:

- Strip size
- Number of physical drives in the logical drive
- RAID level

The best performance and drive longevity is obtained by aligning and sizing the strip and thus stripe size to the application I/O request size and alignment. 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 may 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 may require a reboot for optimal setting selection. If a reboot is required to change settings, HPE 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 may be required after switching power modes to optimize savings and performance.

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

Installation, configuration, and maintenance

Installation

Supported servers

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

For the latest support information, see "**Controller supported features**" for web links to the Quick Specs for each controller.

For a list of Gen10 supported servers, see the **Smart Array Controller Server Matrix**.

Installing a Smart Array in an unconfigured server

Procedure

1. Install the Smart Array hardware.

   For server-specific procedures, see the server user guide.
2. For P-series, perform the following:

   - Connect one end of the controller backup power cable to the backup power connector on the Smart Array and the other end to the controller backup power connector on the system board or PCI riser board.
   - Install the optional energy pack.
3. Install physical drives, as needed, and attach the physical drives to the Smart Array.
4. Power up the server.
5. 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 (**http://www.hpe.com/servers/spp**).

   You might need to extract the Smart Array 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.
6. Create a storage array using the HPE Smart Storage Administrator or the Smart Array configuration utility in UEFI System Utilities.
7. 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 *Intelligent Provisioning User Guide for HPE ProLiant Gen10 Servers and HPE Synergy* (**http://www.hpe.com/support/IPGen10-UG-en**).

   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 5.

More information

Updating software and firmware on page 52
Installing a Smart Array on page 42
Connecting storage devices on page 45

Array and controller configuration on page 48

Installing a Smart Array 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 Smart Array is the new boot device, install the device drivers.
   - If the new Smart Array 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 Smart Array hardware.

    For server-specific procedures, see the server user guide.
12. For P-series, perform the following:
    - Connect one end of the controller backup power cable to the backup power connector on the Smart Array and the other end to the controller backup power connector on the system board or PCI riser board.
    - Install the optional energy pack.
13. Connect storage devices to the controller.
14. Connect peripheral devices to the server.
15. Connect the power cord to the server.
16. Connect the power cord to the power source.
17. Power up all peripheral devices.
18. Power up the server.

19. If you are running the server in UEFI Boot Mode, power on and select the boot options.
20. Update the controller firmware if it is not the latest revision.
21. Update the drive firmware if it not the latest revision.
22. (Optional) If running the server in Legacy Boot Mode, set the controller as the boot controller.
23. (Optional) If running the server in Legacy Boot Mode, change the controller boot order.
24. If the new controller is not the new boot device, install the device drivers.
25. (Optional) Create additional logical drives.

More information

Powering on and selecting boot options in UEFI Boot Mode on page 51
Updating software and firmware on page 52
Array and controller configuration on page 48
Connecting internal storage on page 46

Installing a Smart Array

Installing a modular Smart Array (-a/-b)

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. If required for installation, remove the Smart Array air baffle.
4. Install the Smart Array by aligning the board with the alignment pins and pressing down.

   If the Smart Array has thumbscrews, tighten the thumbscrews. If the Smart Array has a curved handle, lift the handle before pressing down, then swing the handle back down to secure the connection. For more information, see the server user guide.
5. Connect storage devices to the controller.
6. 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.

Installing a modular Smart Array (-c)

CAUTION: Hewlett Packard Enterprise recommends performing a complete backup of all server or compute module data before performing a Smart Array installation or removal.

CAUTION: In systems that use external data storage, be sure that the server or compute module 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 or compute module is powered up.

For specific procedures, see the user guide at the **Hewlett Packard Enterprise website**.

Procedure

1. Power down the compute module.
2. Remove the compute module.
3. Place the compute module on a flat, level work surface.
4. Remove the access panel.
5. Remove all drives.
6. Remove all drive blanks.
7. Remove the front panel/drive cage assembly.

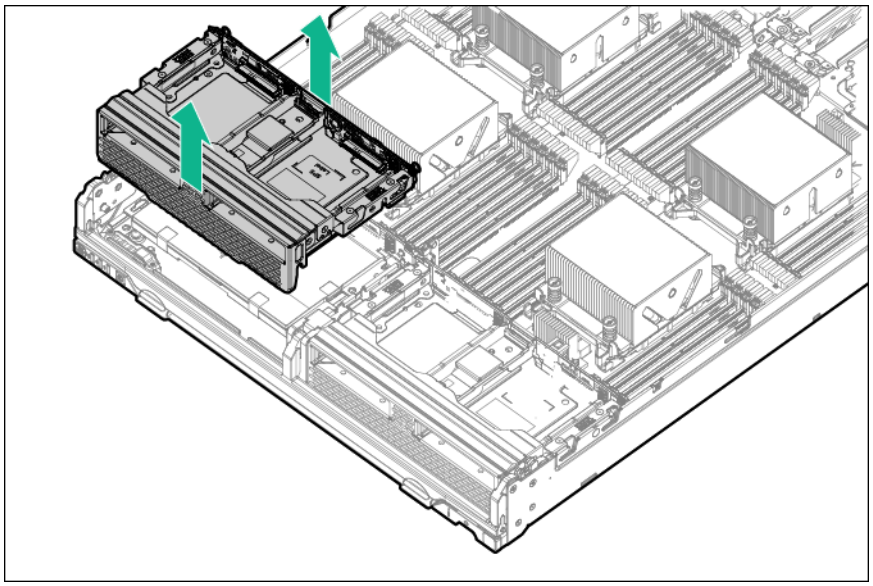

8. Install the Smart Array .

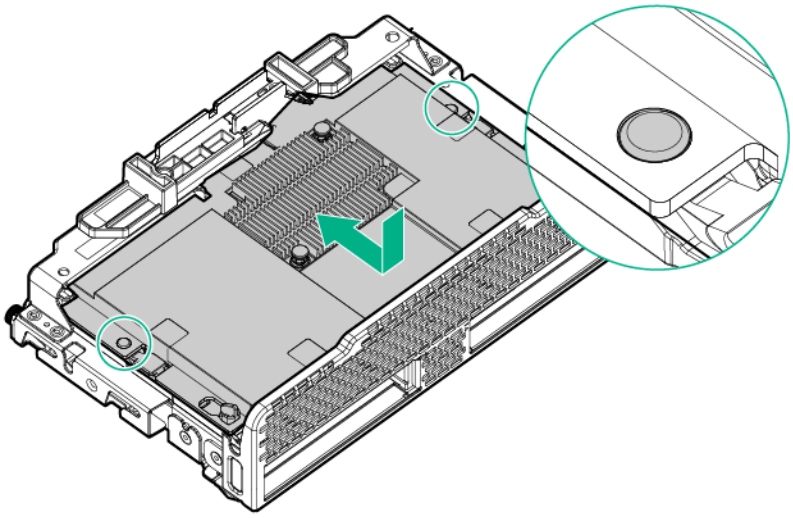

9. Install the front panel/drive cage assembly.
10. Install all drives.
11. Install all drive blanks.
12. Install the access panel.
13. Install the compute module.

Installing a standup PCIe Plug-In Smart Array (-p)

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 x8 or larger PCIe expansion slot.

   A x8 physical size slot is required, even though the slot width may be electrically x4 or x1. Hewlett Packard Enterprise recommends using a slot that is electrically x8.

4. Remove the slot cover.

   Save the retaining screw, if one is present.

5. Slide the Smart Array 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 Smart Array 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 Smart Array 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 Smart Array .
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.

Installing a mezzanine (-m) Smart Array

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 mezzanine connector on the system board.
4. Remove the connector cover, and then save it for future use.
5. Install the Smart Array by aligning the board with the alignment pins and pressing down.

   If the Smart Array has thumbscrews, tighten the thumbscrews. For more information, see the server user guide.
6. Connect storage devices to the Smart Array .
7. Close or install the access panel.

NOTE: 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.

Connecting storage devices

For more information about supported drive models on specific Gen10 HPE servers, see the QuickSpecs for the specific server.

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

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

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

5. Power up the server.

More information

Array and controller configuration on page 48

Connecting external storage

Procedure

1. Power down the server.
2. Connect an external SAS cable to the external port of the controller:

   a. Pull back the tab on the Mini-SAS x4 connector on the cable.

   b. Insert the cable connector into the external port of the controller.

   c. Release the tab.

3. Connect the other end of the cable to the SAS input connector of the external storage enclosure.

4. Power up the enclosure.
5. Power up the server.

Cable part numbers

For more information on cables, see the server QuickSpecs on the Hewlett Packard Enterprise website **http://www.hpe.com/info/qs**.

Enabling Smart Array SW RAID

Prerequisites

Server boot mode must be set to UEFI Mode.

Only SATA drives are supported.

Procedure

1. Reboot the server.
2. Start UEFI System Utilities by pressing **F9 (System Utilities)** during POST.
3. Select **System Configuration** > **BIOS/Platform Configuration (RBSU)** > **Storage Options** > **SATA Controller Options** > **Embedded SATA Configuration** > **Smart Array SW RAID Support** and press the **Enter** key.
4. In the SATA Controller Options screen, for the Embedded SATA configuration option, select **Smart Array SW RAID Support** from the drop-down menu and click **OK**.

   If you see the warning "Important: Smart Array SW RAID is not supported when the Boot Mode is configured in Legacy BIOS Mode", click **OK**.
5. Press **F12: Save and Exit**.
6. Save changes by clicking **Yes-Save Changes**.
7. Click **Reboot**.

Device drivers

A device driver is required for the operating system to communicate with the disk drive controller. Device drivers are provided with Service Pack for ProLiant. Use Intelligent Provisioning to install the drivers. For more information on supported operating systems, see the product page for the appropriate controller on the Hewlett Packard Enterprise website.

Driver updates are posted to the Hewlett Packard Enterprise website. When prompted for product information, enter the appropriate server model name or controller model.

Windows operating systems

The device driver for the HPE Smart Array S100i SR Gen10 SW RAID solution is not included with the off-the-shelf Microsoft Windows operating system media. If you are installing the operating system to a logical volume managed by the S100i SW RAID solution, then you must ensure that the device driver is installed so that the Windows installation will recognize the logical volume.

If you install the operating system using the assisted installation option in the Intelligent Provisioning software, the driver is automatically added during installation. No further action is required.

If you install the operating system without Intelligent Provisioning, you must manually inject the driver during the OS installation from a USB drive.

Manually injecting the device driver during OS installation

Procedure

1. Obtain the drivers and extract them to a USB drive.

   The drivers for the S100i SR SW RAID solution can be obtained from either Service Pack for ProLiant or the **Hewlett Packard Enterprise website**. The option to extract the files is presented when running the driver installation component.

2. Insert the USB drive containing the driver, and click **Browse**.

   For installations using a SATA optical drive to load the Windows installation media, Windows provides an early prompt for the optical device driver to continue installation. When the USB is inserted and the driver is loaded, the driver loads for the SATA optical device and for the SATA drives or SATA SSDs connected to the S100i SR SW RAID solution.

   For installations that do not use the SATA optical drive to load the Windows installation media, the **Where Do You Want to Install Windows** window displays.

3. Click **Load driver** and point to the extracted Smart Array SW RAID solution driver folder.

Configuration

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 HPE Smart Storage Administrator (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 Smart Array 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.
- 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 HPE SSA or the UEFI System Utilities to configure the arrays and controllers.

Comparison of HPE SSA and UEFI RAID configuration utility

HPE Smart Array SR Gen10 can be configured by using either HPE SSA or the Smart Array Configuration Utility within the UEFI System Utilities. Both HPE SSA and UEFI System Utilities can be used to configure the Smart Array S-class, E-class, and P-class.

HPE 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 HPE SSA during that step.

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

Task	HPE 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 stripe size	+	+
Migrate the RAID level or stripe 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	+	
Configure HPE Smart Array SR SmartCache	+	

For specific information about how to use either HPE SSA or the UEFI System Utilities, see *HPE Smart Array SR Gen10 Configuration Guide* on the Hewlett Packard Enterprise website **http://www.hpe.com/info/smartstorage-docs**.

HPE Smart Storage Administrator

HPE SSA is the main tool for configuring arrays on Smart Array controllers. It exists in three interface formats: the HPE SSA GUI, the HPE SSA CLI, and HPE SSA Scripting. All formats provide support for configuration tasks. Some of the advanced tasks are available in only one format.

The diagnostic features in HPE SSA are also available in the standalone software HPE Smart Storage Administrator Diagnostics Utility CLI.

HPE SSA is accessible both offline and online:

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

UEFI System Utilities

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

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

For more information on the UEFI System Utilities, see the *UEFI System Utilities User Guide for HPE ProLiant Gen10 Servers* on the **Hewlett Packard Enterprise website**.

Scan the QR code located at the bottom of the screen to access mobile-ready online help for the UEFI System Utilities and UEFI Shell. For on-screen help, press **F1**.

Using UEFI System Utilities

To use the System Utilities, use the following keys.

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

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

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

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

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

Intelligent Provisioning

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

Intelligent Provisioning 3.30 and later includes HPE SMB Setup . When you launch F10 mode from the POST screen, you are prompted to select whether you want to enter the Intelligent Provisioning or HPE SMB Setup mode.

NOTE: After you have selected a mode, you must reprovision the server to change the mode that launches when you boot to F10.

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 Service Pack for ProLiant (SPP). SPP is a comprehensive systems software and firmware solution for 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 .

IMPORTANT: HPE ProLiant XL servers do not support operating system installation with Intelligent Provisioning , but they do support the maintenance features. For more information, see "Performing Maintenance" in the Intelligent Provisioning user guide and online help.

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 SMB Setup .
- From the iLO web interface using **Always On**. **Always On** allows you to access Intelligent Provisioning without rebooting your server.

Configuring boot controller options

Configuration procedures vary if the server is running in UEFI Boot Mode or Legacy Boot Mode.

Selecting a boot mode

Procedure

1. From the System Utilities screen, select **System Configuration** > **Boot Options** > **Boot Mode**, and press the **Enter** key.
2. Select a setting and press the **Enter** key.
   - UEFI Mode (default) - Configures the system to boot to a UEFI-compatible operating system.

     NOTE: When booting to the UEFI Mode, configure the system to use native UEFI graphic drivers.
   - Legacy BIOS Mode - Configures the system to boot to a traditional operating system in Legacy BIOS compatibility mode.
3. Press the **F10** key to save your selection.
4. Reboot the server.

Powering on and selecting boot options in UEFI Boot Mode

On servers operating in UEFI Boot Mode, the boot controller and boot order are set automatically.

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

For more information on automatic configuration, see the UEFI documentation on the **Hewlett Packard Enterprise website**.

Changing the Legacy BIOS boot order

Prerequisite

Boot Mode is set to **Legacy BIOS Mode**.

Procedure

1. From the **System Utilities** screen, select **System Configuration** > **BIOS/Platform Configuration (RBSU)** > **Boot Options** > **Legacy BIOS Boot Order** and press **Enter**.
2. Use the arrow keys to navigate within the boot order list.
3. Press the **+** key to move an entry higher in the boot list.
4. Press the **-** key to move an entry lower in the list.
5. Press **F10**.

System maintenance tools

Updating software and firmware

Server and controller firmware should be updated before using the controller for the first time. For system software and firmware updates, download the SPP from the Hewlett Packard Enterprise website **http://www.hpe.com/servers/spp/download**. For information about the SPP, see the product documentation at **http://www.hpe.com/info/spp/docs**.

Hewlett Packard Enterprise now distributes drivers and other support software for servers and server blades through Service Pack for ProLiant, or SPP, which you can download from **http://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.

Online firmware activation readiness

Online firmware activation enables the update of controller firmware without a system reboot. You can now verify if your controller configuration is ready for future firmware updates that are enabled for online activation.

To verify if your controller configuration is ready for online firmware activation, use the Smart Storage Administrator user interface. Use the **Check Online Firmware Activation Readiness** button on the **Actions** menu. For more information, see *HPE Smart Array SR Gen10 Configuration Guide* at **http://www.hpe.com/info/SmartArrayGen10-docs**.

Online firmware activation requires that both the controller configuration *and* the update firmware are online firmware activation enabled. Proper enablement of a controller configuration is dependent upon several factors, including:

- A supported operating system.
- Number and state of the drives.
- Supported RAID configuration.
- Whether SmartCache or encryption is enabled.

One example of an enabled controller configuration is:

- installed on Linux
- has no more than eight drives attached
- RAID1 configured with two drives
- fully healthy (in optimal state with no drive firmware update pending or in progress)
- does not use SmartCache or encryption

When firmware updates enabled for online activation become available, review the firmware release notes for information about supported controller configurations.

Diagnostic tools

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

- **HPE Smart Storage Administrator (HPE SSA)**

  HPE SSA can be accessed offline using Intelligent Provisioning or booting from the SPP ISO image. It can also be accessed online by downloading the HPE SSA executables. For more information on accessing and using HPE SSA, see the *HPE Smart Array SR Controllers Configuration Guide*at the **Hewlett Packard Enterprise website**.

- **System Event Log**

  HPE Smart Array SR Event Notification Service for Windows reports array events to the Microsoft Windows system event log, and records the Smart Array serial log, which includes detailed diagnostic information of the most recent events encountered by the controller. The HPE ProLiant Agentless Management Service reports events to the Linux event log. You can obtain the utility from the **Hewlett Packard Enterprise website**. When prompted for product information, enter the server model name.

- **Integrated Management Log (IML)**

  Smart Array reports diagnostic error messages (POST messages) during boot and logs these messages to the UEFI Health Log and also the Integrated Management Log (IML) within iLO. Many POST messages suggest corrective actions. For more information about POST messages, see *Integrated Management Log Messages and Troubleshooting Guide for HPE ProLiant Gen 10 and HPE Synergy*.

- **HPE iLO**

  The iLO firmware continuously monitors Smart Array independent of the operating system and logs any failure events to the IML, iLO WebGUI, iLO RESTful API, and SNMP.

Troubleshooting resources

Troubleshooting resources are available for HPE Gen10 server products in the following documents:

- *Troubleshooting Guide for HPE ProLiant Gen10 servers* provides procedures for resolving common problems and comprehensive courses of action for fault isolation and identification, issue resolution, and software maintenance.
- *Error Message Guide for HPE ProLiant Gen10 servers and HPE Synergy* provides a list of error messages and information to assist with interpreting and resolving error messages.
- *Integrated Management Log Messages and Troubleshooting Guide for HPE ProLiant Gen10 and HPE Synergy* provides IML messages and associated troubleshooting information to resolve critical and cautionary IML events.

To access the troubleshooting resources, see the Hewlett Packard Enterprise Information Library (**http://www.hpe.com/info/gen10-troubleshooting**).

Models

Modular Smart Array (-a/-b/-c)

HPE Smart Array E208i-a SR Gen10

HPE Smart Array P408i-a SR Gen10

HPE Smart Array P816i-a SR Gen10

HPE Smart Array P204i-b SR Gen10

HPE Smart Array E208i-c SR Gen10

HPE Smart Array P204i-c SR Gen10

HPE Smart Array P408i-c SR Gen 10

HPE Smart Array E208i-a SR Gen10

E208i-a controller ports and connectors

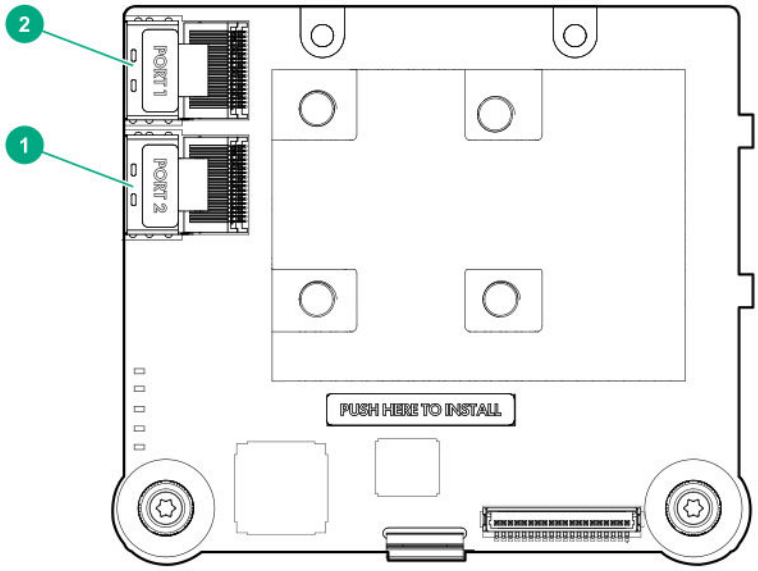

Item	Description
1	Internal x4 Mini-SAS port 2i
2	Internal x4 Mini-SAS port 1i

E208i-a controller status LEDs

Immediately after you power up the server, the controller runtime LEDs illuminate briefly in a predetermined pattern as part of the POST sequence. At all other times during server operation, the illumination pattern of the runtime LEDs indicates the status of the controller.

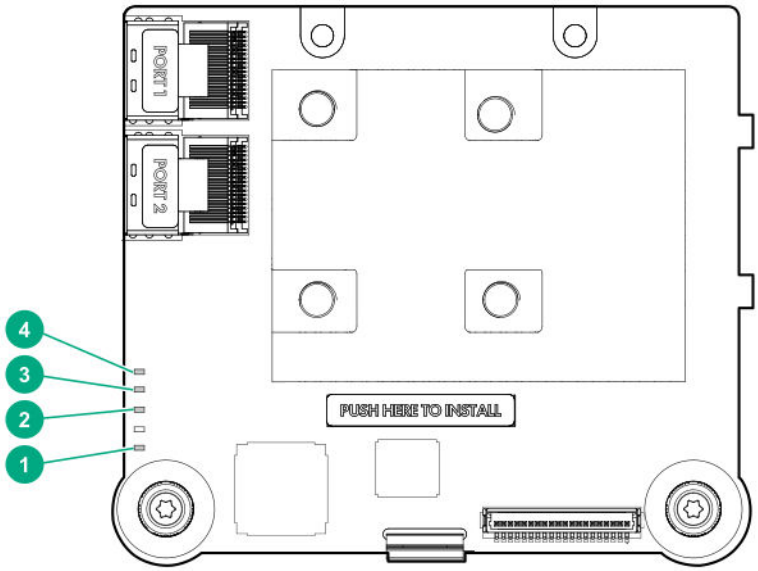

Item	Color	Name	Interpretation
1	Amber	Debug	On = Controller is in reset state. Off = Controller is in an idle or runtime state.
2	Green	Crypto	On = All attached volumes are encrypted Off = All attached volumes are plaintext Flashing = Both encrypted and plaintext volumes are present
3	Amber	Fault	When an error occurs, this LED is on. During power up, this LED is solid for up to 2 seconds.
4	Green	Heartbeat	When the controller is in good health, this LED flashes at 1 Hz. During power up, this LED is solid for up to 2 seconds.

HPE Smart Array P408i-a SR Gen10

P408i-a controller ports and connectors

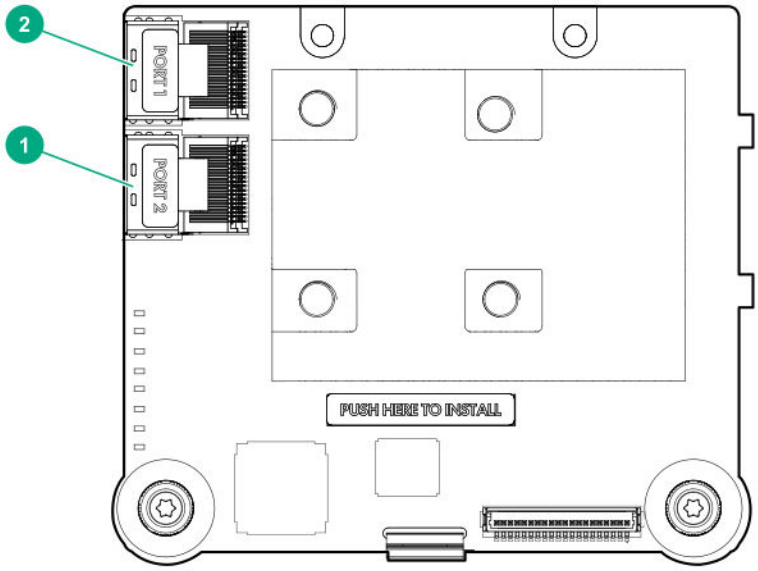

Item	Description
1	Internal x4 Mini-SAS port 2i
2	Internal x4 Mini-SAS port 1i

P408i-a controller status LEDs

Immediately after you power up the server, the controller runtime LEDs illuminate briefly in a predetermined pattern as part of the POST sequence. At all other times during server operation, the illumination pattern of the runtime LEDs indicates the status of the controller.

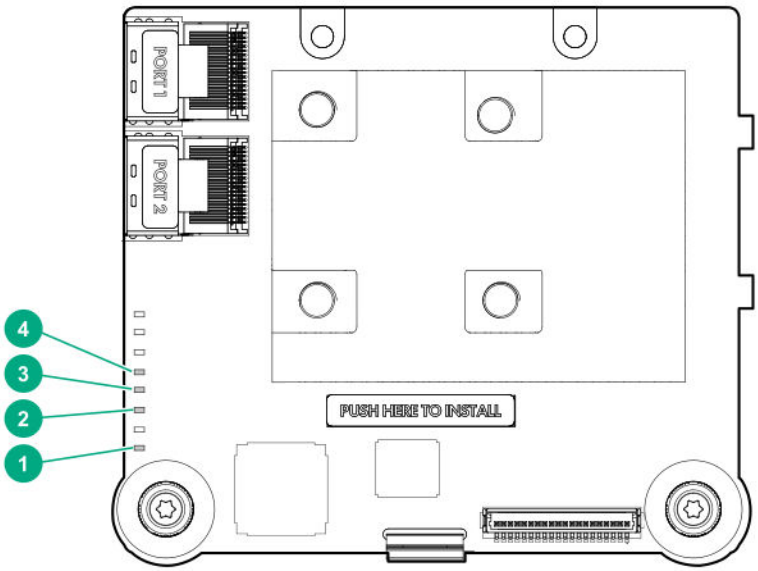

Item	Color	Name	Interpretation
1	Amber	Debug	On = Controller is in reset state. Off = Controller is in an idle or runtime state.
2	Green	Crypto	On = All attached volumes are encrypted Off = All attached volumes are plaintext Flashing = Both encrypted and plaintext volumes are present
3	Amber	Fault	When an error occurs, this LED is on. During power up, this LED is solid for up to 2 seconds.
4	Green	Heartbeat	When the controller is in good health, this LED flashes at 1 Hz. During power up, this LED is solid for up to 2 seconds.

P408i-a controller FBWC LEDs

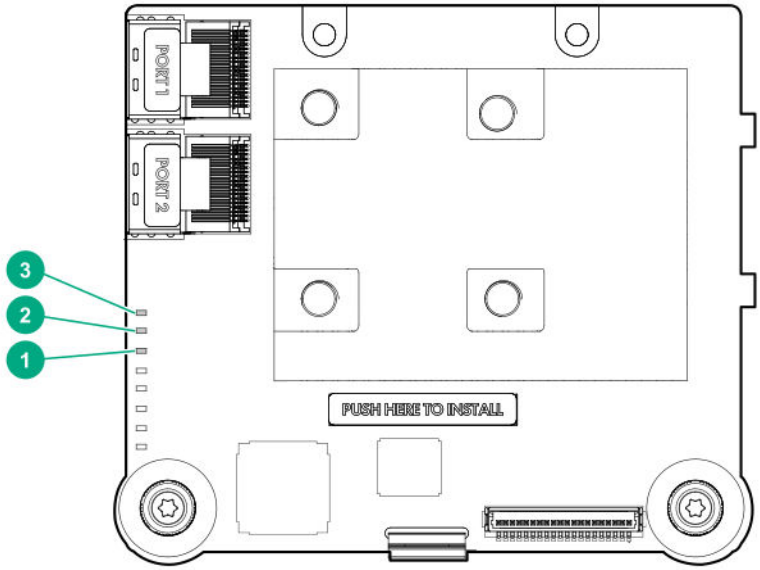

1 - DDR LED3 - Green	2 - DDR LED2 - Green	3 - DDR LED1 - Amber	Interpretation
Off	Off	Off	The FBWC module is not powered
Flashing once every 2 seconds	Flashing once every 2 seconds	Off	The cache microcontroller is executing from within its boot loader and receiving new flash code from the host controller.

Table Continued

1 - DDR LED3 - Green	2 - DDR LED2 - Green	3 - DDR LED1 - Amber	Interpretation
Flashing once per second	Flashing once per second	Off	The FBWC module is powering up, and waiting for backup power.
Flashing once per second	Off	Off	The FBWC module is idle, and waiting for backup power.
On	Off	Off	The FBWC module is idle, and backup power is ready.
On	On	Off	The FBWC module is idle, the backup power is ready, and the cache contains data that has not yet been written to the drives.
Off	Flashing once per second	Off	A backup of the DDR content on the FBWC module is in progress.
Off	On	Off	The current backup is complete with no errors.
Off	Flashing once per second	Flashing once per second	The current backup failed, and data has been lost.
On	Flashing once per second	Flashing once per second	A power error occurred during the previous or current boot. Data may be corrupt.
Off	On	Flashing once per second	An overtemperature condition exists.
Off	On	On	The current backup is complete, but power fluctuations occurred during backup.
On	On	On	The FBWC module microcontroller has failed.

HPE Smart Array P816i-a SR Gen10

P816i-a controller ports and connectors

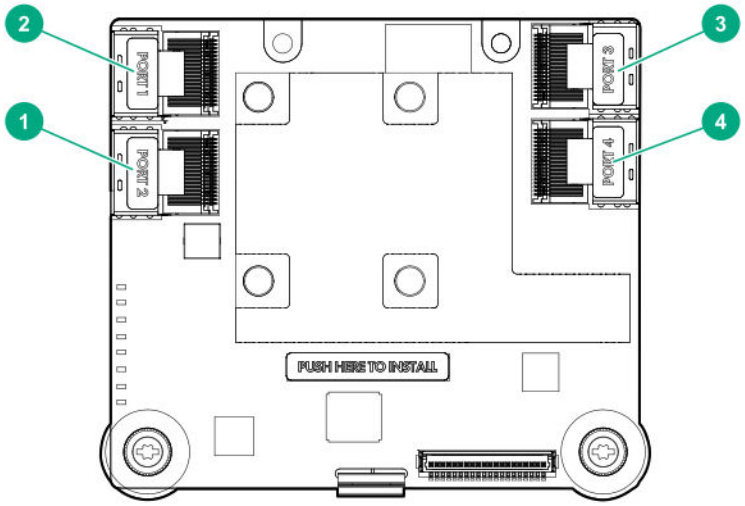

Item	Description
1	Internal x4 Mini-SAS port 2i
2	Internal x4 Mini-SAS port 1i
3	Internal x4 Mini-SAS port 3i
4	Internal x4 Mini-SAS port 4i

P816i-a controller status LEDs

Immediately after you power up the server, the controller runtime LEDs illuminate briefly in a predetermined pattern as part of the POST sequence. At all other times during server operation, the illumination pattern of the runtime LEDs indicates the status of the controller.

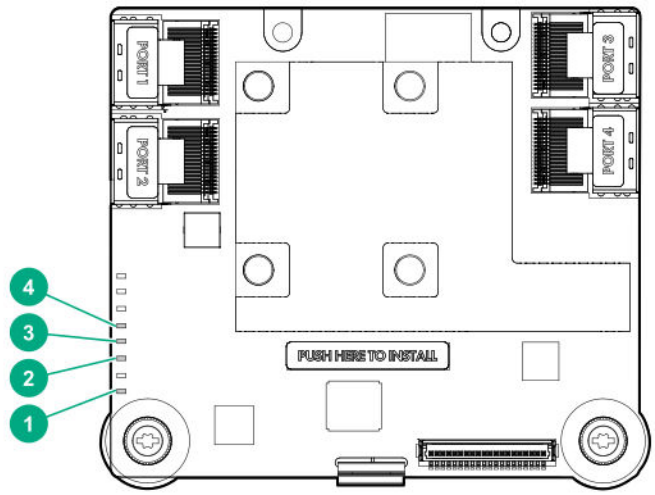

Item	Color	Name	Interpretation
1	Amber	Debug	On = Controller is in reset state. Off = Controller is in an idle or runtime state.
2	Green	Crypto	On = All attached volumes are encrypted. Off = All attached volumes are plaintext. Flashing = Both encrypted and plaintext volumes are present.
3	Amber	Fault	When an error occurs, this LED is on. During power up, this LED is solid for up to 2 seconds.
4	Green	Heartbeat	When the controller is in good health, this LED flashes at 1 Hz. During power up, this LED is solid for up to 2 seconds.

P816i-a controller FBWC LEDs

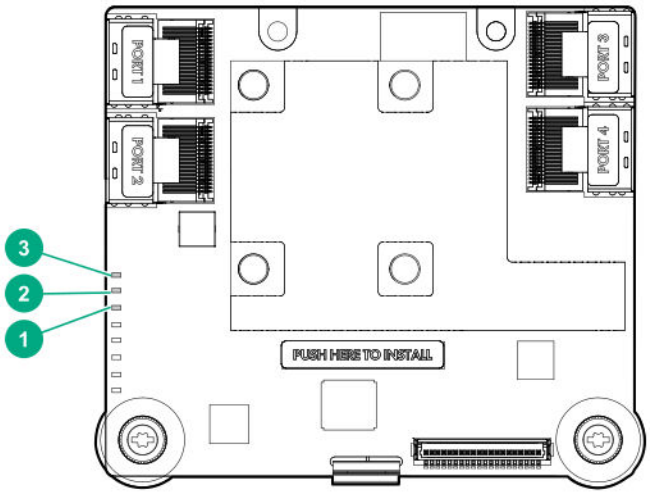

1 - DDR LED3 - Green	2 - DDR LED2 - Green	3 - DDR LED1 - Amber	Interpretation
Off	Off	Off	The FBWC module is not powered
Flashing once every 2 seconds	Flashing once every 2 seconds	Off	The cache microcontroller is executing from within its boot loader and receiving new flash code from the host controller.
Flashing once per second	Flashing once per second	Off	The FBWC module is powering up, and waiting for backup power.

Table Continued

1 - DDR LED3 - Green	2 - DDR LED2 - Green	3 - DDR LED1 - Amber	Interpretation
Flashing once per second	Off	Off	The FBWC module is idle, and waiting for backup power.
On	Off	Off	The FBWC module is idle, and backup power is ready.
On	On	Off	The FBWC module is idle, the backup power is ready, and the cache contains data that has not yet been written to the drives.
Off	Flashing once per second	Off	A backup of the DDR content on the FBWC module is in progress.
Off	On	Off	The current backup is complete with no errors.
Off	Flashing once per second	Flashing once per second	The current backup failed, and data has been lost.
On	Flashing once per second	Flashing once per second	A power error occurred during the previous or current boot. Data may be corrupt.
Off	On	Flashing once per second	An overtemperature condition exists.
Off	On	On	The current backup is complete, but power fluctuations occurred during backup.
On	On	On	The FBWC module microcontroller has failed.

HPE Smart Array P204i-b SR Gen10

P204i-b controller status LEDs

Immediately after you power up the server, the controller runtime LEDs illuminate briefly in a predetermined pattern as part of the POST sequence. At all other times during server operation, the illumination pattern of the runtime LEDs indicates the status of the controller.

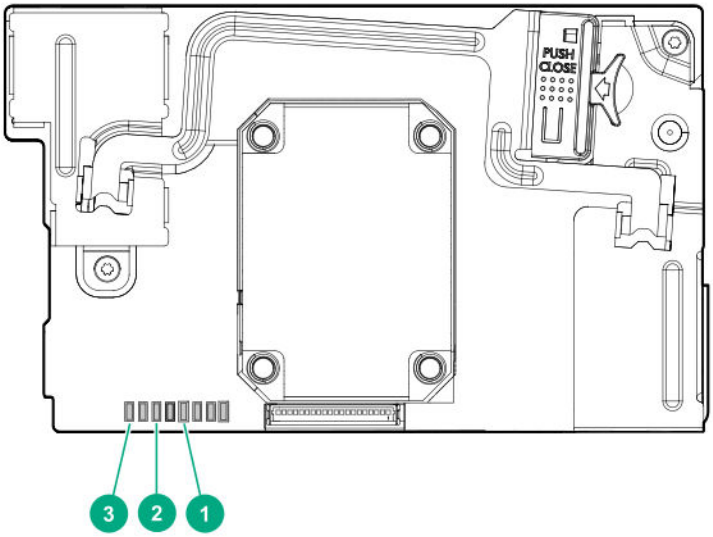

Item	Color	Name	Interpretation
1	Amber	Debug	On = Controller is in reset state. Off = Controller is in an idle or runtime state.
2	Green	Crypto	On = All attached volumes are encrypted. Off = All attached volumes are plaintext. Flashing = Both encrypted and plaintext volumes are present.
3	Green	Heartbeat	When the controller is in good health, this LED flashes at 1 Hz. During power up, this LED is solid for up to 2 seconds.

P204i-b controller FBWC LEDs

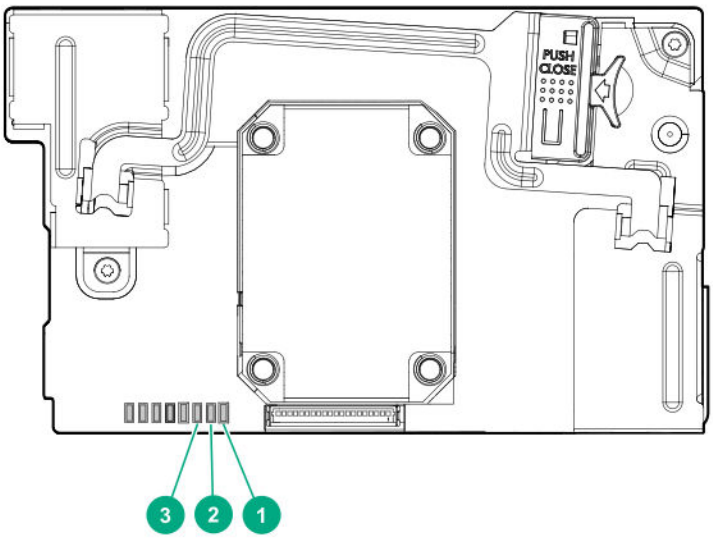

1 - DDR LED3 - Green	2 - DDR LED2 - Green	3 - DDR LED1 - Green	Interpretation
Off	Off	Off	The FBWC module is not powered
Flashing once every 2 seconds	Flashing once every 2 seconds	Off	The cache microcontroller is executing from within its boot loader and receiving new flash code from the host controller.
Flashing once per second	Flashing once per second	Off	The FBWC module is powering up, and waiting for backup power.
Flashing once per second	Off	Off	The FBWC module is idle, and waiting for backup power.
On	Off	Off	The FBWC module is idle, and backup power is ready.
On	On	Off	The FBWC module is idle, the backup power is ready, and the cache contains data that has not yet been written to the drives.
Off	Flashing once per second	Off	A backup of the DDR content on the FBWC module is in progress.
Off	On	Off	The current backup is complete with no errors.
Off	Flashing once per second	Flashing once per second	The current backup failed, and data has been lost.
On	Flashing once per second	Flashing once per second	A power error occurred during the previous or current boot. Data may be corrupt.
Off	On	Flashing once per second	An overtemperature condition exists.
Off	On	On	The current backup is complete, but power fluctuations occurred during backup.
On	On	On	The FBWC module microcontroller has failed.

HPE Smart Array E208i-c SR Gen10

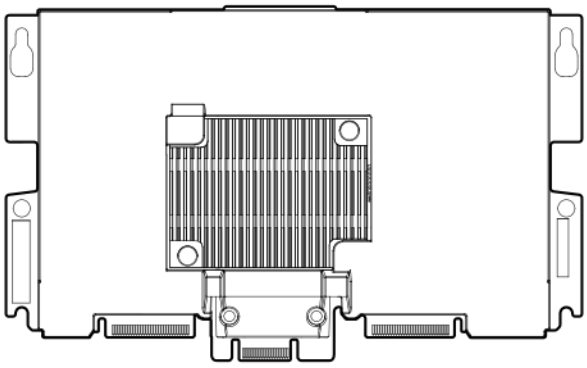

HPE Smart Array P204i-c SR Gen10

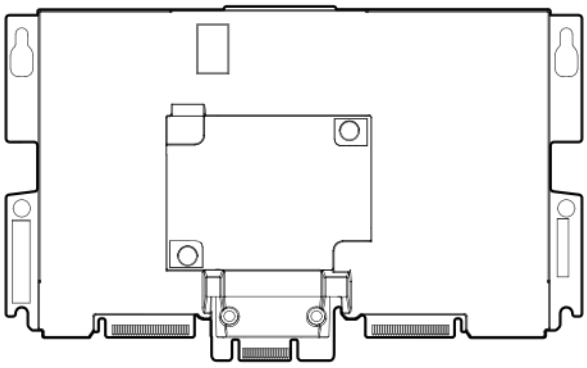

HPE Smart Array P408i-c SR Gen10

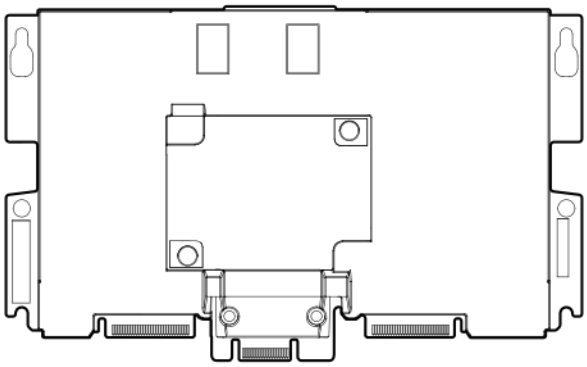

Standup PCIe Plug-In Smart Array (-p)

HPE Smart Array E208i-p SR Gen10

HPE Smart Array E208e-p SR Gen10

HPE Smart Array P408i-p SR Gen10

HPE Smart Array P408e-p SR Gen10

HPE Smart Array E208i-p SR Gen10

E208i-p controller ports and connectors

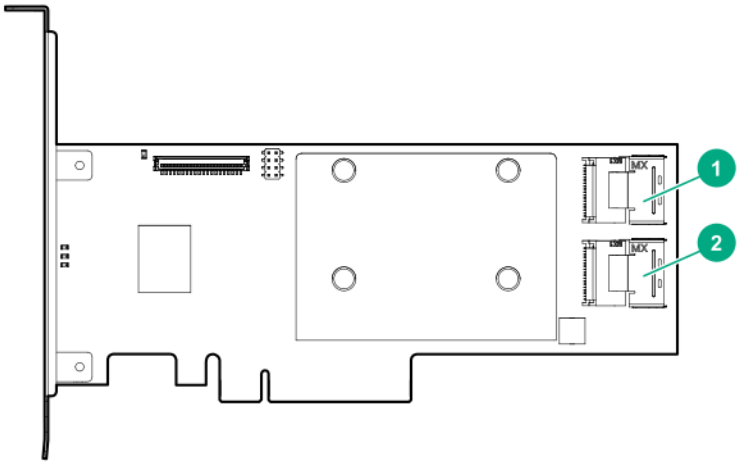

Item	Description
1	Internal x4 Mini-SAS port 1i
2	Internal x4 Mini-SAS port 2i

E208i-p controller status LEDs

Immediately after you power up the server, the controller runtime LEDs illuminate briefly in a predetermined pattern as part of the POST sequence. At all other times during server operation, the illumination pattern of the runtime LEDs indicates the status of the controller.

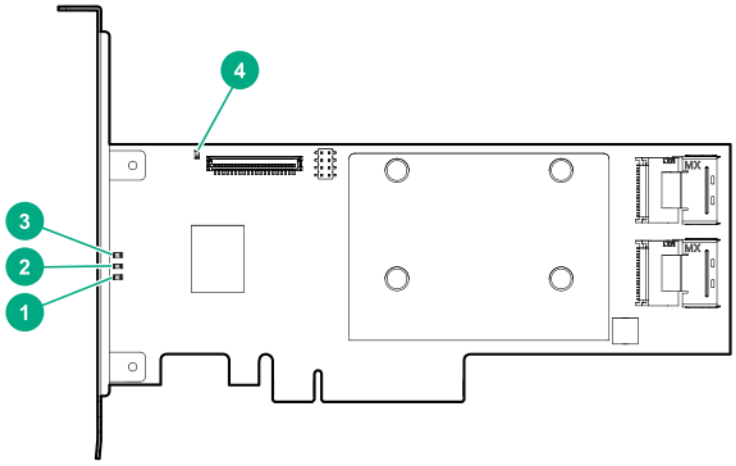

Item	Color	Name	Interpretation
1	Green	Crypto	On = All attached volumes are encrypted Off = All attached volumes are plaintext Flashing = Both encrypted and plain text volumes are present
2	Amber	Fault	When an error occurs, this LED is on. During power up, this LED is solid for up to 2 seconds.
3	Green	Heartbeat	When the controller is in good health, this LED flashes at 1 Hz. During power up, this LED is solid for up to 2 seconds.
4	Amber	Debug	On = Controller is in reset state. Off = Controller is in an idle or runtime state.

HPE Smart Array E208e-p SR Gen10

E208e-p controller ports and connectors

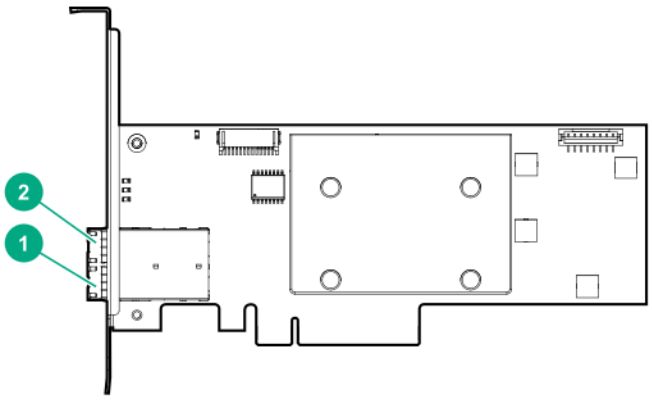

Item	Description
1	External x4 Mini-SAS HD port 2e
2	External x4 Mini-SAS HD port 1e

E208e-p controller status LEDs

Immediately after you power up the server, the controller runtime LEDs illuminate briefly in a predetermined pattern as part of the POST sequence. At all other times during server operation, the illumination pattern of the runtime LEDs indicates the status of the controller.

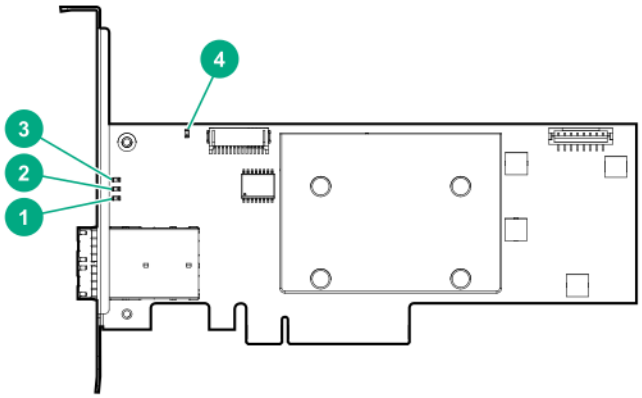

Item	Color	Name	Interpretation
1	Green	Crypto	On = All attached volumes are encrypted Off = All attached volumes are plaintext Flashing = Both encrypted and plaintext volumes are present
2	Amber	Fault	When an error occurs, this LED is on. During power up, this LED is solid for up to 2 seconds.
3	Green	Heartbeat	When the controller is in good health, this LED flashes at 1 Hz. During power up, this LED is solid for up to 2 seconds.
4	Amber	Debug	On = Controller is in reset state. Off = Controller is in an idle or runtime state.

HPE Smart Array P408i-p SR Gen10

P408i-p controller ports and connectors

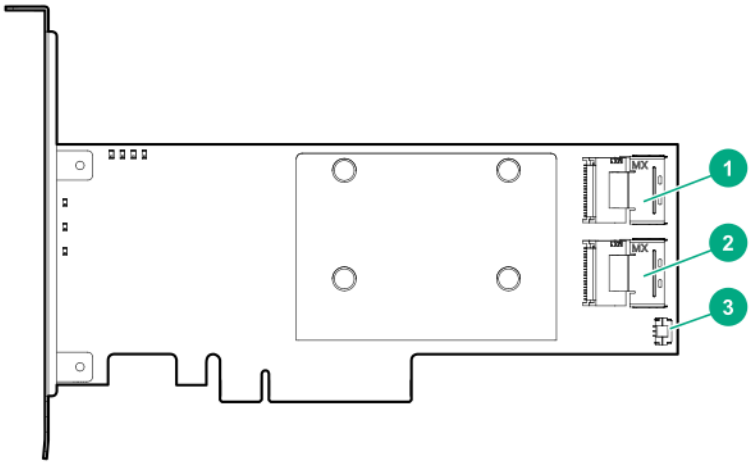

Item	Description
1	Internal x4 Mini-SAS port 1i
2	Internal x4 Mini-SAS port 2i
3	Controller backup power connector

P408i-p controller status LEDs

Immediately after you power up the server, the controller runtime LEDs illuminate briefly in a predetermined pattern as part of the POST sequence. At all other times during server operation, the illumination pattern of the runtime LEDs indicates the status of the controller.

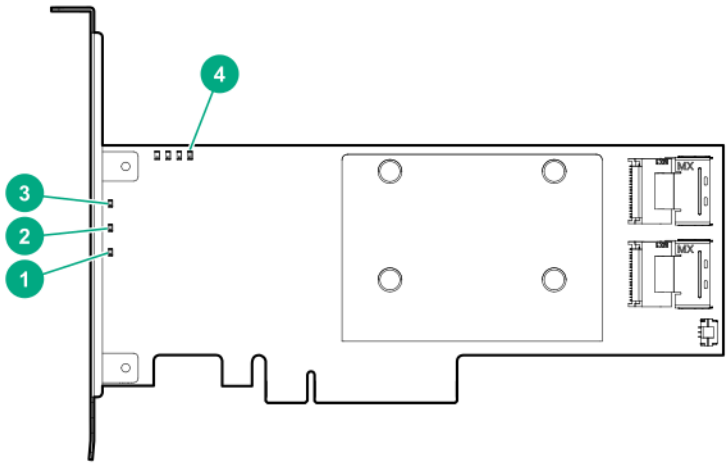

Item	Color	Name	Interpretation
1	Green	Crypto	On = All attached volumes are encrypted Off = All attached volumes are plaintext Flashing = Both encrypted and plaintext volumes are present
2	Amber	Fault	When an error occurs, this LED is on. During power up, this LED is solid for up to 2 seconds.
3	Green	Heartbeat	When the controller is in good health, this LED flashes at 1 Hz. During power up, this LED is solid for up to 2 seconds.
4	Amber	Debug	On = Controller is in reset state. Off = controller is in an idle or runtime state. Flashing 5 Hz = Controller and cache are performing a backup.

P408i-p controller flash-backed write cache LEDs

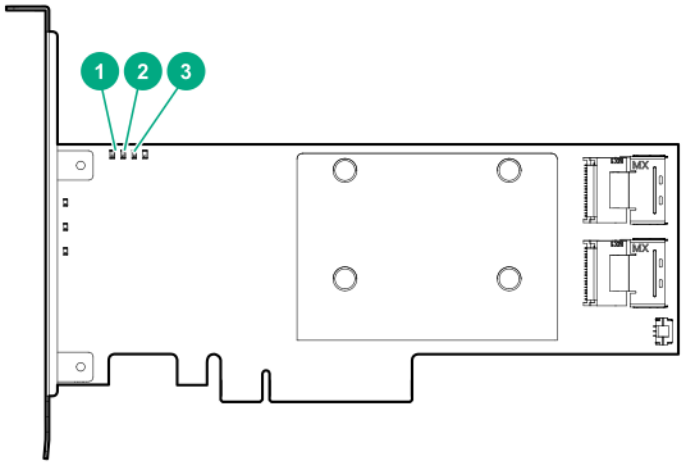

1 - DDR LED1 - Amber	2 - DDR LED2 - Green	3 - DDR LED3 - Green	Interpretation
Off	Off	Off	The FBWC module is not powered
Off	Flashing once every 2 seconds	Flashing once every 2 seconds	The cache microcontroller is executing from within its boot loader and receiving new flash code from the host controller.
Off	Flashing once per second	Flashing once per second	The FBWC module is powering up, and waiting for backup power.
Off	Off	Flashing once per second	The FBWC module is idle, and waiting for backup power.
Off	Off	On	The FBWC module is idle, and backup power is ready.
Off	On	On	The FBWC module is idle, the backup power is ready, and the cache contains data that has not yet been written to the drives.
Off	Flashing once per second	Off	A backup of the DDR content on the FBWC module is in progress.
Off	On	Off	The current backup is complete with no errors.
Flashing once per second	Flashing once per second	Off	The current backup failed, and data has been lost.

Table Continued

1 - DDR LED1 - Amber	2 - DDR LED2 - Green	3 - DDR LED3 - Green	Interpretation
Flashing once per second	Flashing once per second	On	A power error occurred during the previous or current boot. Data may be corrupt.
Flashing once per second	On	Off	An overtemperature condition exists.
On	On	Off	The current backup is complete, but power fluctuations occurred during backup.
On	On	On	The FBWC module microcontroller has failed.

HPE Smart Array P408e-p SR Gen10

P408e-p controller ports and connectors

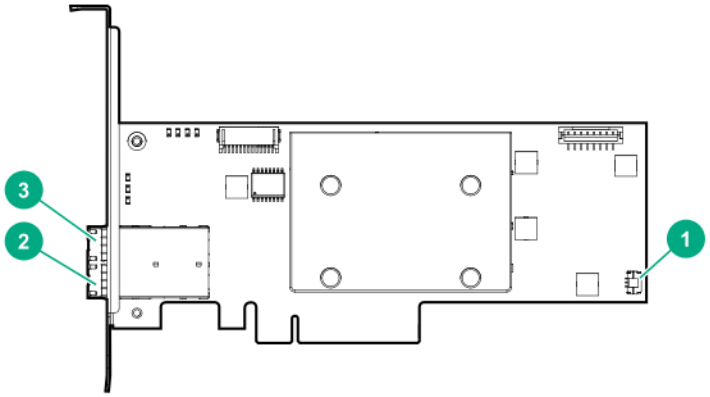

Item	Description
1	Controller backup power connector
2	External x4 Mini-SAS HD port 2e[1]
3	External x4 Mini-SAS HD port 1e[1]

[1] External SAS ports 1e and 2e comprise the single Mini-SAS HD receptacle connector, which can accept either two Mini-SAS HD x4 plug connectors or a single Mini-SAS HD x8 plug connector.

P408e-p controller status LEDs

Immediately after you power up the server, the controller runtime LEDs illuminate briefly in a predetermined pattern as part of the POST sequence. At all other times during server operation, the illumination pattern of the runtime LEDs indicates the status of the controller.

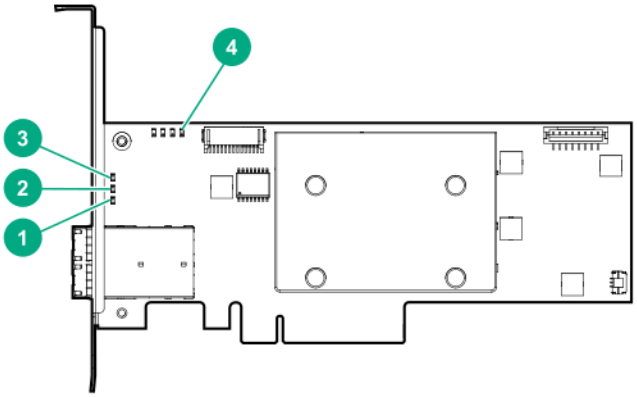

Item	Color	Name	Interpretation
1	Green	Crypto	On = All attached volumes are encrypted. Off = All attached volumes are plaintext. Flashing = Both encrypted and plaintext volumes are present.
2	Amber	Fault	When an error occurs, this LED is on. During power up, this LED is solid for up to 2 seconds.
3	Green	Heartbeat	When the controller is in good health, this LED flashes at 1 Hz. During power up, this LED is solid for up to 2 seconds.
4	Amber	Debug	On = Controller is in reset state. Off = controller is in an idle or runtime state. Flashing 5 Hz = Controller and cache are performing a backup.

P408e-p controller flash-backed write cache LEDs

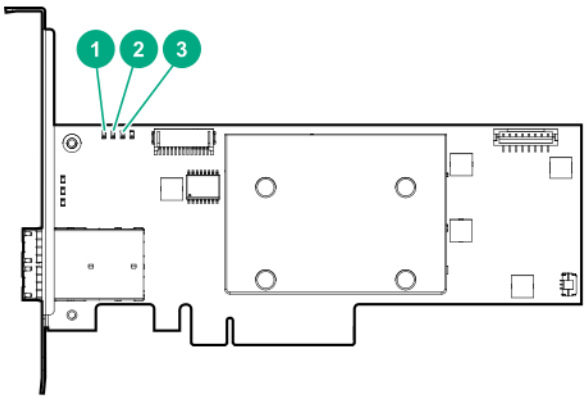

1 - DDR LED1 - Amber	2 - DDR LED2 - Green	3 - DDR LED3 - Green	Interpretation
Off	Off	Off	The FBWC module is not powered
Off	Flashing once every 2 seconds	Flashing once every 2 seconds	The cache microcontroller is executing from within its boot loader and receiving new flash code from the host controller.
Off	Flashing once per second	Flashing once per second	The FBWC module is powering up, and waiting for backup power.
Off	Off	Flashing once per second	The FBWC module is idle, and waiting for backup power.
Off	Off	On	The FBWC module is idle, and backup power is ready.
Off	On	On	The FBWC module is idle, the backup power is ready, and the cache contains data that has not yet been written to the drives.
Off	Flashing once per second	Off	A backup of the DDR content on the FBWC module is in progress.
Off	On	Off	The current backup is complete with no errors.
Flashing once per second	Flashing once per second	Off	The current backup failed, and data has been lost.

Table Continued

1 - DDR LED1 - Amber	2 - DDR LED2 - Green	3 - DDR LED3 - Green	Interpretation
Flashing once per second	Flashing once per second	On	A power error occurred during the previous or current boot. Data may be corrupt.
Flashing once per second	On	Off	An overtemperature condition exists.
On	On	Off	The current backup is complete, but power fluctuations occurred during backup.
On	On	On	The FBWC module microcontroller has failed.

Mezzanine controllers (-m)

HPE Smart Array P408e-m SR Gen10

HPE Smart Array P416ie-m SR Gen 10

HPE Smart Array P408e-m SR Gen10

P408e-m controller status LEDs

Immediately after you power up the server, the controller runtime LEDs illuminate briefly in a predetermined pattern as part of the POST sequence. At all other times during server operation, the illumination pattern of the runtime LEDs indicates the status of the controller.

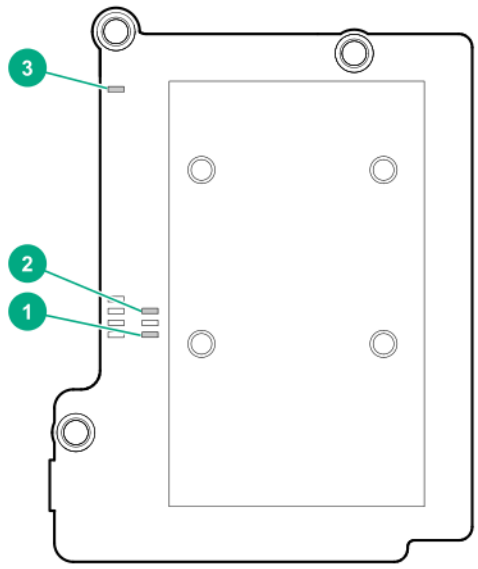

Item	Color	Name	Interpretation
1	Green	Crypto	On = All attached volumes are encrypted. Off = All attached volumes are plaintext. Flashing = Both encrypted and plaintext volumes are present.
2	Green	Heartbeat	When the controller is in good health, this LED flashes at 1 Hz. During power up, this LED is solid for up to 2 seconds.
3	Amber	Debug	On = Controller is in reset state. Off = Controller is in an idle or runtime state. Flashing 5 Hz = Controller and cache are performing a backup.

P408e-m controller flash-backed write cache LEDs

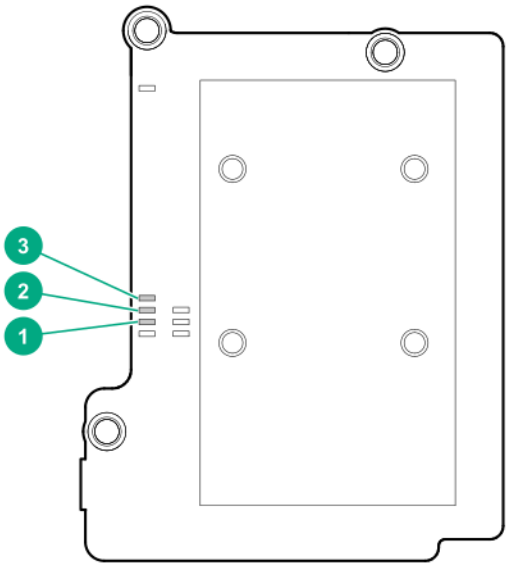

1 - DDR LED1 - Green	2 - DDR LED2 - Green	3 - DDR LED3 - Green	Interpretation
Off	Off	Off	The FBWC module is not powered
Off	Flashing once every 2 seconds	Flashing once every 2 seconds	The cache microcontroller is executing from within its boot loader and receiving new flash code from the host controller.
Off	Flashing once per second	Flashing once per second	The FBWC module is powering up, and waiting for backup power.
Off	Off	Flashing once per second	The FBWC module is idle, and waiting for backup power.

Table Continued

1 - DDR LED1 - Green	2 - DDR LED2 - Green	3 - DDR LED3 - Green	Interpretation
Off	Off	On	The FBWC module is idle, and backup power is ready.
Off	On	On	The FBWC module is idle, the backup power is ready, and the cache contains data that has not yet been written to the drives.
Off	Flashing once per second	Off	A backup of the DDR content on the FBWC module is in progress.
Off	On	Off	The current backup is complete with no errors.
Flashing once per second	Flashing once per second	Off	The current backup failed, and data has been lost.
Flashing once per second	Flashing once per second	On	A power error occurred during the previous or current boot. Data may be corrupt.
Flashing once per second	On	Off	An overtemperature condition exists.
On	On	Off	The current backup is complete, but power fluctuations occurred during backup.
On	On	On	The FBWC module microcontroller has failed.

HPE Smart Array P416ie-m SR Gen10

P416ie-m controller ports and connectors

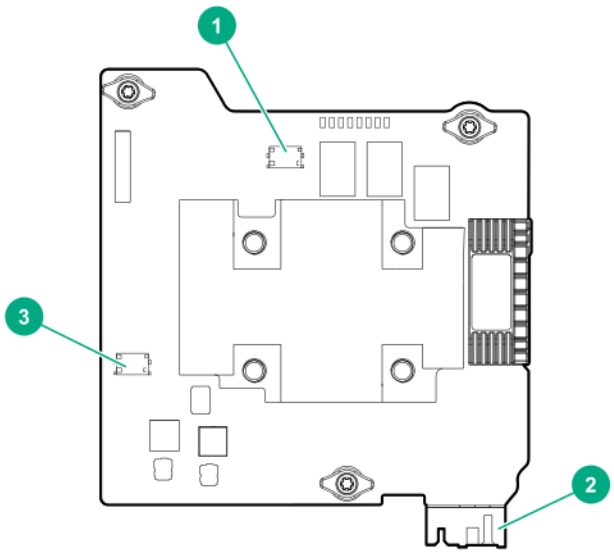

Item	Description
1	Internal x4 Slim SAS port 4i
2	Fabric connector
3	Internal x4 Slim SAS port 3i

P416ie-m controller status LEDs

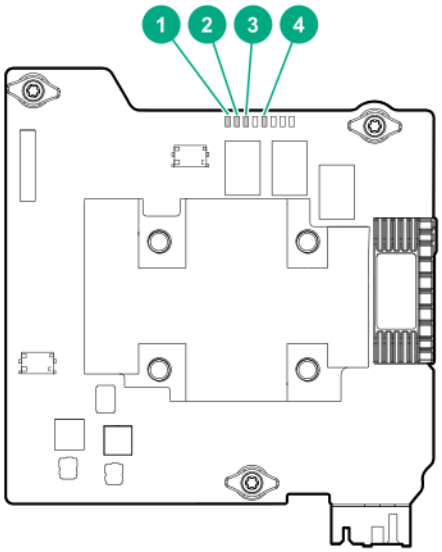

Immediately after you power up the server, the controller runtime LEDs illuminate briefly in a predetermined pattern as part of the POST sequence. At all other times during server operation, the illumination pattern of the runtime LEDs indicates the status of the controller.

Item	Color	Name	Interpretation
1	Green	Heartbeat	When the controller is in good health, this LED flashes at 1 Hz. During power up, this LED is solid for up to 2 seconds.
2	Amber	Fault	When an error occurs, this LED is on. During power up, this LED is solid for up to 2 seconds.
3	Green	Crypto	On = All attached volumes are encrypted. Off = All attached volumes are plaintext. Flashing = Both encrypted and plaintext volumes are present.
4	Amber	Debug	On = Controller is in reset state. Off = Controller is in an idle or runtime state. Flashing 5 Hz = Controller and cache are performing a backup.

P416ie-m controller flash-backed write cache LEDs

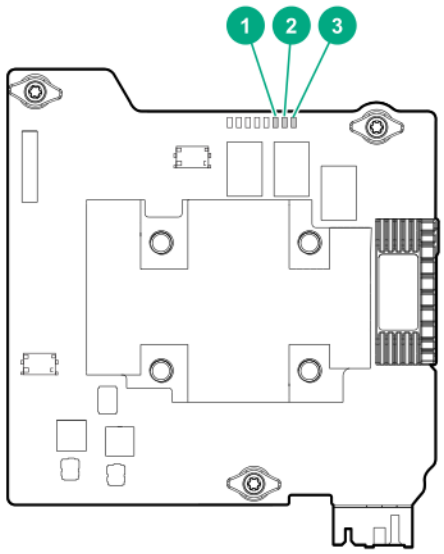

1 - DDR LED3 - Green	2 - DDR LED2 - Green	3 - DDR LED1 - Amber	Interpretation
Off	Off	Off	The FBWC module is not powered
Flashing once every 2 seconds	Flashing once every 2 seconds	Off	The cache microcontroller is executing from within its boot loader and receiving new flash code from the host controller.
Flashing once per second	Flashing once per second	Off	The FBWC module is powering up, and waiting for backup power.
Flashing once per second	Off	Off	The FBWC module is idle, and waiting for backup power.
On	Off	Off	The FBWC module is idle, and backup power is ready.
On	On	Off	The FBWC module is idle, the backup power is ready, and the cache contains data that has not yet been written to the drives.
Off	Flashing once per second	Off	A backup of the DDR content on the FBWC module is in progress.
Off	On	Off	The current backup is complete with no errors.
Off	Flashing once per second	Flashing once per second	The current backup failed, and data has been lost.

Table Continued

1 - DDR LED3 - Green	2 - DDR LED2 - Green	3 - DDR LED1 - Amber	Interpretation
On	Flashing once per second	Flashing once per second	A power error occurred during the previous or current boot. Data may be corrupt.
Off	On	Flashing once per second	An overtemperature condition exists.
Off	On	On	The current backup is complete, but power fluctuations occurred during backup.
On	On	On	The FBWC module microcontroller has failed.

Additional hardware and options

Energy pack options

Hewlett Packard Enterprise offers two centralized backup power source options to back up write cache content on P-class Smart Array controllers in case of an unplanned server power outage.

- **HPE Smart Storage Battery**
- **HPE Smart Storage Hybrid Capacitor**

> **IMPORTANT:** The HPE Smart Storage Hybrid Capacitor is only supported on Gen10 and later servers that support the 96W HPE Smart Storage Battery.

One energy pack option can support multiple devices. An energy pack option is required for P-class Smart Array controllers. Once installed, the status of the energy pack displays in HPE iLO. For more information, see the HPE iLO user guide on the Hewlett Packard Enterprise website (**http://www.hpe.com/support/ilo-docs**).

HPE Smart Storage Battery

The HPE Smart Storage Battery supports the following devices:

- HPE Smart Storage SR controllers
- NVDIMMs

> **IMPORTANT:** To support NVDIMMs, the HPE Smart Storage Battery must be installed.

A single 96W battery can support up to 24 devices.

After the battery is installed, it might take up to two hours to charge. Controller features requiring backup power are not re-enabled until the battery is capable of supporting the backup power.

HPE Smart Storage Hybrid Capacitor

The HPE Smart Storage Hybrid Capacitor supports the following devices:

HPE Smart Storage SR controllers

> **IMPORTANT:** NVDIMMs are only supported by the HPE Smart Storage Battery.

The capacitor pack can support up to three devices.

Before installing the HPE Smart Storage Hybrid Capacitor, verify that the system BIOS meets the minimum firmware requirements to support the capacitor pack.

> **IMPORTANT:** If the system BIOS or controller firmware is older than the minimum recommended firmware versions, the capacitor pack will only support one device.

The capacitor pack is fully charged after the system boots.

Energy pack specifications

Table 1: HPE Smart Storage Battery (96W)

Feature	Description
Time required to recharge Smart Storage Battery	96 W: 2 hours (For maximum load of 24 devices)
Duration of Smart Storage Battery backup	150 seconds (maximum support) The Smart Storage Battery provides a sufficient duration to transfer the cached data from DDR memory to flash memory, where the data remains indefinitely or until a controller retrieves the data.

For more information, see the **QuickSpecs document** for the Smart Storage Battery.

Table 2: HPE Smart Storage Battery (12W)

Feature	Description
Time required to recharge Smart Storage Battery	12 W: 1 hour (For maximum load of 3 devices)
Duration of Smart Storage Battery backup	150 seconds (maximum support) The Smart Storage Battery provides a sufficient duration to transfer the cached data from DDR memory to flash memory, where the data remains indefinitely or until a controller retrieves the data.

For more information, see the **QuickSpecs document** for the Smart Storage Battery.

Table 3: HPE Smart Storage Hybrid Capacitor

Feature	Description
Time required to recharge capacitor pack	Not applicable; charge is immediate
Duration of capacitor pack backup	60 seconds The capacitor pack provides a sufficient duration to transfer the cached data from DDR memory to flash memory, where the data remains indefinitely or until a controller retrieves the data.

The capacitor pack can be used in place of the 96 W Smart Storage Battery, but it is not supported on the Apollo servers.

For more information, see the **QuickSpecs document** for the capacitor pack.

HPE 12G SAS Expander Card

The HPE 12G SAS Expander connects to Smart Array E-class and P-class Gen10 controllers. Depending on the server configuration, the expander may support up to a maximum of 28 drives.

For more information about the 12G SAS Expander Card, see the **HPE 12G SAS Expander Card QuickSpecs**.

To install the 12G SAS Expander Card, see the server-specific documentation that ships with the card.

Specifications

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

Controller specifications

The specifications of the Smart Array controllers are in the QuickSpecs for each controller.

- **S-class controller S100i QuickSpecs**
- **All E-class and P-class controllers QuickSpecs**

Energy pack specifications

Table 4: HPE Smart Storage Battery (96W)

Feature	Description
Time required to recharge Smart Storage Battery	96 W: 2 hours (For maximum load of 24 devices)
Duration of Smart Storage Battery backup	150 seconds (maximum support) The Smart Storage Battery provides a sufficient duration to transfer the cached data from DDR memory to flash memory, where the data remains indefinitely or until a controller retrieves the data.

For more information, see the **QuickSpecs document** for the Smart Storage Battery.

Table 5: HPE Smart Storage Battery (12W)

Feature	Description
Time required to recharge Smart Storage Battery	12 W: 1 hour (For maximum load of 3 devices)
Duration of Smart Storage Battery backup	150 seconds (maximum support) The Smart Storage Battery provides a sufficient duration to transfer the cached data from DDR memory to flash memory, where the data remains indefinitely or until a controller retrieves the data.

For more information, see the **QuickSpecs document** for the Smart Storage Battery.

Table 6: HPE Smart Storage Hybrid Capacitor

Feature	Description
Time required to recharge capacitor pack	Not applicable; charge is immediate
Duration of capacitor pack backup	60 seconds The capacitor pack provides a sufficient duration to transfer the cached data from DDR memory to flash memory, where the data remains indefinitely or until a controller retrieves the data.

The capacitor pack can be used in place of the 96 W Smart Storage Battery, but it is not supported on the Apollo servers.

For more information, see the **QuickSpecs document** for the capacitor pack.

Support and other resources

Accessing Hewlett Packard Enterprise Support

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

  http://www.hpe.com/info/assistance

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

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

Information to collect

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

Accessing updates

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

  Hewlett Packard Enterprise Support Center

  www.hpe.com/support/hpesc

  Hewlett Packard Enterprise Support Center: Software downloads

  www.hpe.com/support/downloads

  Software Depot

  www.hpe.com/support/softwaredepot

- To subscribe to eNewsletters and alerts:

  www.hpe.com/support/e-updates

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

  www.hpe.com/support/AccessToSupportMaterials

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

Customer self repair

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

For more information about CSR, contact your local service provider or go to the CSR website:

http://www.hpe.com/support/selfrepair

Remote support

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

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

Remote support and Proactive Care information

HPE Get Connected

www.hpe.com/services/getconnected

HPE Proactive Care services

www.hpe.com/services/proactivecare

HPE Proactive Care service: Supported products list

www.hpe.com/services/proactivecaresupportedproducts

HPE Proactive Care advanced service: Supported products list

www.hpe.com/services/proactivecareadvancedsupportedproducts

Proactive Care customer information

Proactive Care central

www.hpe.com/services/proactivecarecentral

Proactive Care service activation

www.hpe.com/services/proactivecarecentralgetstarted

Warranty information

To view the warranty information for your product, see the links provided below:

HPE ProLiant and IA-32 Servers and Options

www.hpe.com/support/ProLiantServers-Warranties

HPE Enterprise and Cloudline Servers

www.hpe.com/support/EnterpriseServers-Warranties

HPE Storage Products

www.hpe.com/support/Storage-Warranties

HPE Networking Products

www.hpe.com/support/Networking-Warranties

Regulatory information

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

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

Additional regulatory information

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

www.hpe.com/info/reach

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

www.hpe.com/info/ecodata

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

www.hpe.com/info/environment

Documentation feedback

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

Websites

General websites

Hewlett Packard Enterprise Information Library

www.hpe.com/info/EIL

Subscription Service/Support Alerts

www.hpe.com/support/e-updates

Insight Remote Support

www.hpe.com/info/insightremotesupport/docs

For additional general support websites, see **Support and other resources**.

Product websites

HPE Smart Array

http://www.hpe.com/info/SmartArrayGen10-docs

HPE Integrated Lights-Out

www.hpe.com/info/ilo