Hewlett Packard Enterprise

HPE MSA 1060/2060/2062 Storage Troubleshooting Guide

Abstract

This document provides information about troubleshooting HPE MSA 1060/2060/2062 Storage Systems.

Part Number: R0Q73-90020
Published: September 2020

Notices

Contents

Maintenance best practices

To reduce the likelihood of needing to restore from a backup:

- Enable the scrub process for both disk and disk groups.. Enabling scrubs allows bad disk drive blocks to be detected and repaired proactively.
- Keep all disk groups fault tolerant.
- Keep firmware updated.
- Enable event notifications through SNMP, Email, or Syslog for better array monitoring.
- Enable alert notifications.
- Use the MSA Health Check utility (**https://www.hpe.com/storage/MSAHealthCheck**)
- Isolate any faults.
- Have spare drives installed in the array. Having an available spare drive allows disk group reconstruction to begin when a disk drive fails. If no spare drives are available, replace failed drives immediately as long as the disk group is not OFFL

MSA Health Check

Maintaining the health of an MSA ensures that maximal performance is continuously available. Any component of the array that is signaling poor health status can degrade the health of the array and affect array performance. The most common components that can affect the throughput of an MSA are:

- Degraded or nonfunctioning power supplies
- Unhealthy or low-performance drives
- Drives in a leftover state
- QTOF disk groups
- Controllers that are not operational or have degraded health status
- Initiators that have not been discovered or have lost discovery
- Down-rev controller firmware and/or drive firmware.

To obtain a complete diagnostic profile of an MSA storage configuration, obtain array logs from the MSA and submit them at **https://www.hpe.com/storage/MSAHealthCheck**. Hewlett Packard Enterprise recommends resolving any issues that are found prior to reassessing array performance.

Component and LED identification

Disk drive bay numbers

NOTE: In the following illustrations, the numbers shown on the disk drives are for reference only to indicate how the slots are numbered.

Figure 1: 24-drive enclosure or expansion enclosure—front panel with hubcaps removed

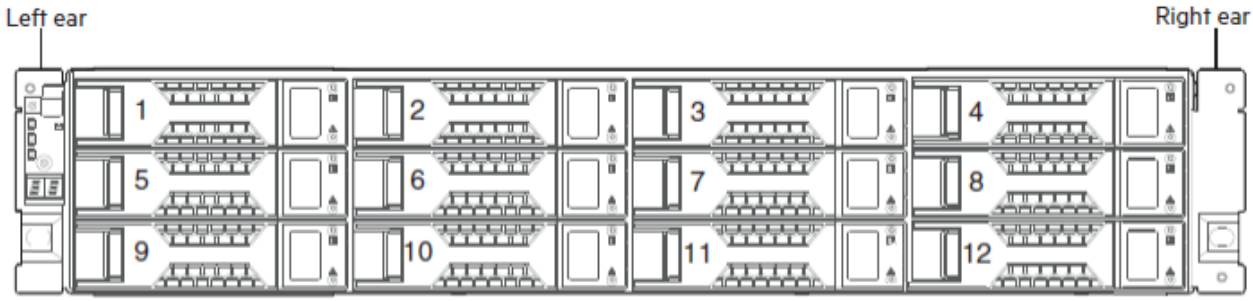

Figure 2: 12-drive controller or expansion enclosure—front panel

LED indicators

Enclosure LEDs

NOTE: In the following illustrations, the numbers shown on the disk drives are for reference only to indicate how the slots are numbered.

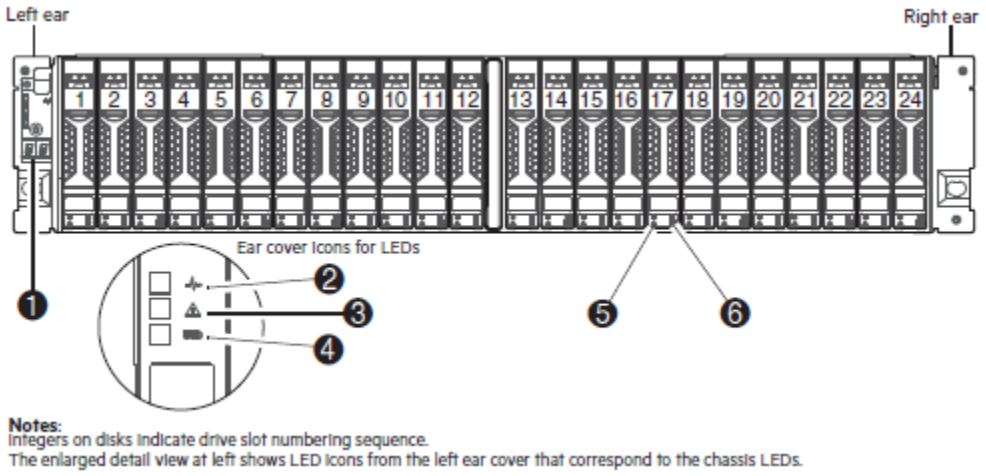

Figure 3: LEDs: 24-drive controller or expansion enclosure—front panel

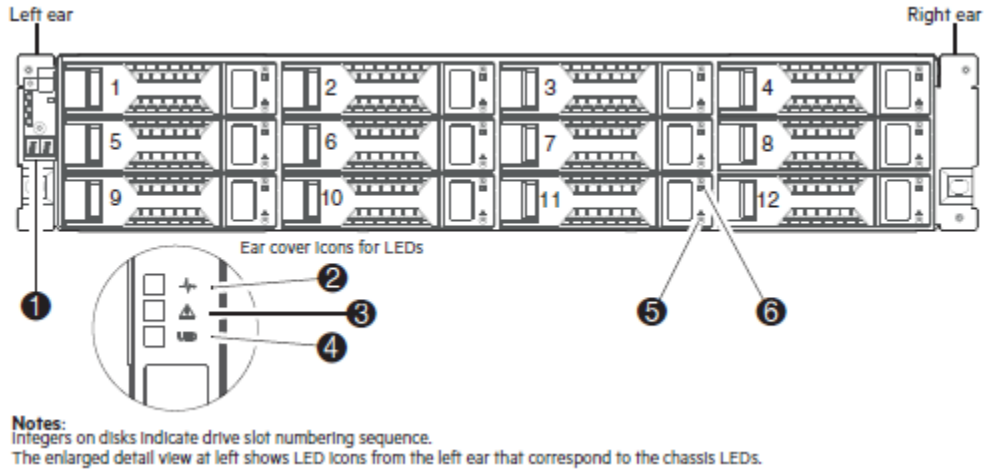

Figure 4: LEDs: 12-drive controller or expansion enclosure—front panel

Table 1: Enclosure LEDs

LED	Description	Definition
1	Enclosure ID	Green—On The enclosure ID value is shown using 7-segment display. Enables you to correlate the enclosure with logical views presented by management software. Sequential enclosure ID numbering of controller enclosures begins with the integer 1. The enclosure ID for an attached drive enclosure is nonzero.
2	System Power	Green—The enclosure is powered on with at least one power supply operating normally. Off—Both power supplies are off; the system is powered off.

Table Continued

LED	Description	Definition
3	Module Fault	Amber—Fault condition exists. The event has been identified, but the problem needs attention. Off—No fault condition exists.
4	Unit Identification (UID)	Blue—Blinking The enclosure is identified. Off—Identify LED is not illuminated.
5	Disk drive Fault/UID	See **Drive LEDs**.
6	Disk drive Online/Activity	See **Drive LEDs**.

Drive LEDs

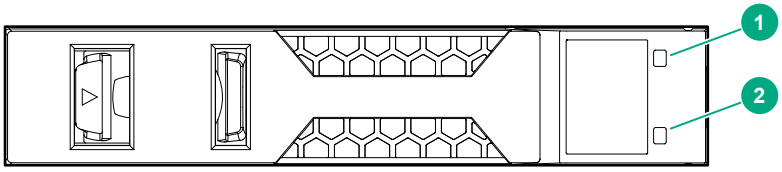

Figure 5: 2.5" SFF disk drive—front panel

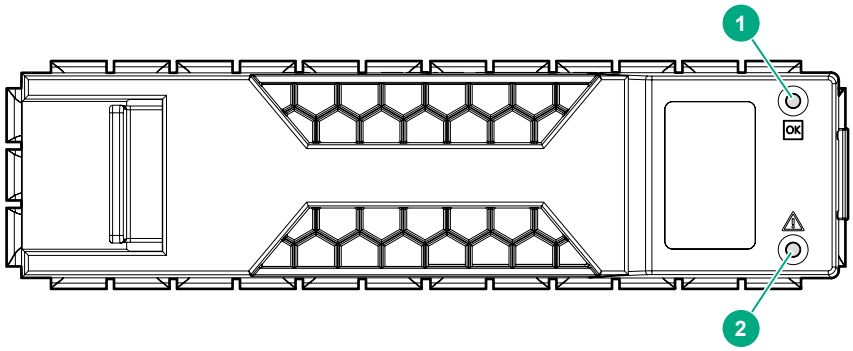

Figure 6: 3.5" LFF disk drive—front panel

Table 2: Disk drive LEDs

LED	Description
1	Online/Activity (green)
2	Fault/UID (amber/blue)

Activity LED (Green)	Fault LED (Amber)	Status/condition[1]
Off	Off	Drive module or enclosure is off, or the drive is not recognized.
Flicker with activity	Blink: (3 on, 1 off)	Identifying drive.
Blink with activity	On	One Drive Link (PHY lane) down
Off	On	Fault (leftover/faile/locked-out both Drive Link (PHY lanes) down
Off: Blink with activity	Off	Available
Blink with activity	Off	Storage system: Initializing
Blink with activity	Off	Storage system: Fault tolerant
Off: Blink with activity	Off	Storage system: Quarantined
Off: Blink with activity	Blink: (Is on/Is off) for the reconstructing drive only	Storage system: Reconstructing **NOTE:** Do not remove the disk drive.

1 If multiple conditions occur simultaneously, the LED state will behave as indicated by the condition listed earliest in the table, as rows are read from top to bottom.

CLI/SMU status indicators

Table 3: Drive status

Usage	Displayed in the CLI or SMU	Description
Available	AVAIL	The drive is available for use.
Failed	FAILED	The drive has failed due to errors. Replace the drive. Reasons for this status include excessive media errors, SMART errors, drive hardware failures, unsupported drives.
Spare	GLOBAL SP	The drive is assigned as a global spare.
Leftover	LEFTOVR	The drive is leftover. It is missing during a rescan of drives or it failed due to Unrecoverable Read Errors (UREs), SMART issues, or other errors.

Table Continued

Usage	Displayed in the CLI or SMU	Description
Pool	`VIRTUAL POOL` or `Pool <pool-id>, <tier-type>`	The drive is used in a disk group.
Unusable	`UNUSABLE`	The drive cannot be used in a disk group. Possible reasons include: the drive is locked with a different passphrase or the passphrase is unavailable, the system is secured but the drive is not FDE capable, or the drive is unsupported.

Table 4: Disk group status

Status	Displayed in the CLI or SMU	Description
Critical	`CRIT`	The disk group is online; however, one or more drives are down and the disk group is not fault tolerant.
Damaged	`DMGD`	The disk group is online and fault tolerant, but some of its drives are damaged.
Fault Tolerant with down drives	`FTDN`	The disk group is online and fault tolerant; however, some drives are down.
Fault Tolerant and online	`FTOL`	The disk group is online and fault tolerant.
Missing	`MSNG`	The disk group is online and fault tolerant, but some of its drives are missing.
Offline	`OFFL`	The disk group is offline because it is using offline initialization, or the drives are down and data loss is at risk.
Quarantined critical	`QTCR`	The disk group is in a critical state with at least one inaccessible drive.
Quarantined with a down disk drive	`QTDN`	The RAID 6 disk group has one inaccessible drive.

Table Continued

Status	Displayed in the CLI or SMU	Description
Quarantined offline	QTOF	The disk group is offline with multiple inaccessible drives causing user data to be incomplete, or is an NRAID or RAID 0 disk group.
Quarantined unsupported	QTUN	The disk group contains data in a format that is not supported by this system.
Unknown	UNKN	The status of the disk group is unknown.
Up	UP	The disk group is online and does not have fault-tolerant attributes.

Alerts and events

Alert notification

Hewlett Packard Enterprise strongly recommends that you review your current alert configuration to ensure that you are correctly monitoring your array.

Alert notification can be set to **all** or **none**:

- ALL—Sends notifications for all alerts (default).
- NONE—Disables email notification of alerts.

Alert notification levels:

INFORMATIONAL—A configuration or state change occurred, or a problem occurred that the system corrected. No action is required.

WARNING—A problem occurred that may affect system stability but not data integrity. Evaluate the problem and correct it if necessary.

ERROR—A failure occurred that may affect data integrity or system stability. Correct the problem as soon as possible.

CRITICAL—A failure occurred that may cause a controller to shut down. Correct the problem immediately

Event notification

Hewlett Packard Enterprise recommends that you review your current event configuration and levels to ensure that you are correctly monitoring your array.

Event notification can be set to four different levels:

- Informational
- Warning
- Error
- Critical

Best practice is to set the event notification level to **Warning**.

RESOLVED—A condition that caused an event to be logged has been resolved.

INFORMATIONAL—A configuration or state change occurred, or a problem occurred that the system corrected. No action is required.

WARNING—A problem occurred that may affect system stability but not data integrity. Evaluate the problem and correct if necessary.

ERROR—A failure occurred that may affect data integrity or system stability. Correct the problem as soon as possible.

CRITICAL—A failure occurred that may cause a controller to shut down. Correct the problem immediately.

Event codes

For a list of event codes and descriptions, see *HPE MSA 1060/2060/2062 Event Descriptions Reference Guide* on the Hewlett Packard Enterprise Support Center website:

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

Events sent as indications to SMI-S clients

If the storage systems SMI-S interface is enabled, the system will send events as indications to SMI-S clients. The SMI-S clients can then monitor the system performance. For information on enabling the SMI-S interface, see Configuring the system in the HPE MSA SMU guide for your product.

For information on the event categories pertaining to Field Replaceable Unit (FRU) assemblies and certain FRU components, see *Events sent as indications to SMI-S clients* in the *HPE MSA 1060/2060/2062 Event Descriptions Reference Guide*. A FRU is any HPE orderable replacement part.

Events requiring FRU replacement

Field Replaceable Unit (FRU) is any HPE orderable replacement part. Events requiring FRU replacements are explained in greater detail in the *HPE MSA Events Guide*. Documentation for HPE MSA Storage Systems is located on the Hewlett Packard Enterprise Support Center website:

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

Performing a power cycle

Enclosures only need to be powered off for routine maintenance. Under normal operations, because of the redundant nature of the MSA array, the system should not require a full system power cycle. Restarting each controller independently results in the same outcome as a full power cycle, while still maintaining host connectivity.

On the front of every MSA enclosure is an **Enclosure ID** LED. This LED will aid in identifying the enclosure. The Controller enclosure will report as "1" and any subsequent enclosures will have higher numbered IDs.

CAUTION: Never power cycle an array without knowing the status of the controllers. Removing power from a controller processing host IO could have negative consequences to data integrity.

If a power cycle is needed, follow these steps:

Procedure

1. Shut down or dismount the hosts (servers) with access to the MSA volumes. This allows any pending writes to flush to the MSA controllers write-back cache.
2. Shut down array controllers by using the CLI or SMU interface. Shutting down the array controllers allows the controllers to flush the controller write cache to the drives.

   NOTE: Never remove the power cords from the Power Supplies on the RAID enclosure without properly shutting down the array controllers. Shutting down the controllers allows for any write cache in the controller modules to be properly flushed down to the hard drives. Removing power from an enclosure may adversely affect disk groups that are shared between enclosures.
3. Power off the drive enclosures only after the controllers are properly shut down using the CLI or SMU. To power off the enclosures, shut off the switches.
4. Perform these steps when restoring power to the Storage infrastructure.
   a. Power on the drive enclosures from the bottom to the top.
   b. Allow the drives in each drive enclosure to power up before powering the next enclosure.
   c. If the enclosures are powered off, wait a minimum of one minute after powering on the last drive enclosure before powering up the MSA RAID enclosure. Waiting allows the hard drives enough time to spin up before the array controllers come online. (Larger capacity drives such as 6TB and 8TB drives may take additional time.)

Storage configurations

Disk groups

The MSA array groups Hard Disk Drives (HDDs) or Solid State Drives (SSDs) into disk groups. The array uses a small part of each drive in a disk group to store information (metadata) about the disk group. This information may include:

- The preferred controller for the disk group.
- The last controller that owned (managed and wrote to) the disk group.
- The drives that are members of the disk group.

Redundant disk groups

For disk groups that are redundant, such as RAID 1, RAID 5, RAID 6, RAID 10, and MSA-DP+, if there are insufficient drives to provide the data, the disk group is quarantined offline (QTOF). For virtual storage, this also causes the pool the disk group is contained in to fault and become read-only, and all volumes will become read-only.

Pools

Disk groups are further grouped into pools, and a volume may be spread across multiple disk groups in the pool. Volumes are mapped to hosts using one or more ports on one or both controllers. Best practice is to map to both controllers.

Performance troubleshooting

Substandard performance can stem from a variety of issues, including:

- **Disk groups**
  - Disk groups in the same tier have different performance and capacity.

    As much as possible, create disk groups with a set of drives that have the same capacity, performance, and media type.
  - Sequential write performance is degraded when disk groups do not adhere to the "power of 2" rule.
  - When creating a parity-based disk group, use the "power or 2" rule whenever possible. The "power of 2" rule states that the number of drives to configure in disk group, excluding parity and spare drives, should be equal to the relation 2^n where n = 1, 2, and so on. For example, in a RAID5 disk group the number of parity drives is 1, so RAID5 disk groups should be created with 3, 5, or 9 total drives. For RAID6 disk groups, the number of parity drives is 2, so RAID6 disk groups should be created with 4 or 10 total drives.
  - Disk group health is degraded.

    All disk groups should be healthy and in the FTOL state. Disk group health can degrade array performance. The following list describes common symptoms of a disk group that may be in an unhealthy state:
    - A physical drive in a disk group is in the Leftover state.
    - A physical drive has logged too many errors and has been put into a degraded status.
    - A physical drive in a disk group shows historically low throughput, even if the drive is not reporting errors.

- **Firmware**
  - Ensure that firmware is current on the MSA system, including controller firmware, enclosure firmware, and drive firmware. Utilize the HPE MSA Health Check website (**https://www.hpe.com/storage/MSAHealthCheck**) to verify all firmware is up to date.

- **Cache**
  - Unwritable cache is present.

    Run the following command to show the percentage of controller cache occupied by unwritten data. If you are unsure of what action to take, contact HPE Support.

    ```
    # show unwritable-cache
    ```

    Data that exists in controller cache which is associated with a volume that has become inaccessible due to failures in the pool or disk groups is deemed unwritable because it cannot be written to storage. To optimize array performance, it is necessary to clear unwritable cache. You must use the CLI `clear cache` command to remove unwriteable cache present in the controller.

    > **CAUTION:** Only use the `clear cache` command when all disk groups are online and accessible from the host. Clearing cache for a volume that is offline or quarantined could result in data loss.

- **Performance Monitoring Metrics**

  Balancing I/O may help in optimizing performance. A summary of usable metrics to assess array performance follows.
  - Host Port Utilization:
    - Are array ports used evenly to distribute host I/O? Is one array port servicing most host I/O while other configured array ports show little to none? Can additional array ports be utilized to increase performance?

CLI: `show host-port-statistics`

SMU: Select and graph **Host Port** in the expanded view of the **Performance Panel**.

- Controller Utilization:
  - Is the storage configuration balanced? Are controllers A and B both being used to service host I/O? Are the array controllers reporting forwarded commands?

    CLI: `show controller-statistics`

    SMU: Select and graph **Controller** in the expanded view of the **Performance Panel**.
- Disk Group Utilization:
  - Are the disk groups configured to have relatively comparable performance and capacity? Does one pool contain the bulk of created disk groups servicing host I/O?

    CLI: `show disk-group-statistics, show tier-statistics tier all, show pool-statistics`
- Volume Utilization
  - Is volume throughput evenly distributed across array controllers?

    CLI: `show volume-statistics`

    SMU: Select and graph **Volume** in the expanded view of the **Performance Panel**.
- Path Utilization
  - Is multipath configured on the host operating system to establish access redundancy for mapped volumes if there is an array controller failover?
  - Are network and/or fabric switches zoned correctly to balance host I/O distribution?

Viewing historical performance with Performance panel

The Performance Widget is used to diagnose and resolve MSA performance issues.

Procedure

1. In the SMU, click the slide-over arrow of the compact view of the Performance panel in the Dashboard.

2. Click **Add Graph**, and then select a system performance category.
3. Choose the associated metrics.

Firmware troubleshooting

Hewlett Packard Enterprise recommends keeping your array, enclosure, and drive firmware up to date with the latest release. This will ensure all the latest improvements and fixes are available for your system.

To check the overall health of your system and any available controller, enclosure, and drive firmware versions use the MSA Health Check Tool: **https://www.hpe.com/storage/MSAHealthCheck**.

Download your MSA Log File from your MSA array, and then upload the file to the MSA Health Check website.

Review the results by clicking through the tabs and saving the PDF report. Links to array, enclosure, and drive firmware are provided for easy access

Data availability troubleshooting

Disk drive troubleshooting

Disk drive is not detected after replacement

Symptom

After inserting a replacement for a failed drive, the replacement drive is not detected.

Solution 1

Cause

The replacement drive is not good.

Action

1. If there is another available disk drive slot, insert the replacement in the alternative slot.

   If the replacement drive is detected, use the CLI or SMU to assign it as a global spare.

2. If the replacement drive is not detected in the alternative slot, insert a known good drive in the alternative slot.

   If the known good drive is detected in the alternative slot, the replacement drive is not good. Use a new drive as a spare.

Solution 2

Cause

The drive slot is not good.

Action

1. Insert a known good drive in the slot where the replacement drive was not detected.

   If the known good drive is not detected, then the slot may have failed. If this is the case, replace the enclosure chassis.

Drive fails and there are issues with reconstruction

Symptom

A drive has failed and reconstruction does not complete.

Solution 1

Cause

A spare drive was used for reconstruction, but before the reconstruction completed, another drive from the same disk group (in the same subgroup for RAID 10) failed, or, in the case of a RAID 6 disk group, two or more drives from the same disk group failed.

Action

1. Attempt to de-quarantine the disk group by reseating the last failed drives.
2. If unable to de-quarantine the disk group by reseating the last failed drive, collect array logs and contact HPE Support.

Solution 2

Cause

Applies to RAID 6:

Disk group status is FTDN, a spare is used, and reconstruction begins. But the reconstruction stops before it completes.

Action

1. Consider replacing the drives if the drives failed due to hardware issues.
2. If the disk group goes QTOF after the failure of a third disk drive member and before the reconstruction completes, attempt to de-quarantine by reseating the last failed drives. If unable to de-quarantine the disk group, collect array logs and contact HPE Support.
3. If unable to de-quarantine the disk group, collect array logs and contact HPE Support.

A third disk drive member fails before reconstruction completes

Symptom

Disk group goes QTOF after the failure of a third disk drive member and before the reconstruction completes.

Action

1. Attempt to de-quarantine by reseating the last failed drives.
2. If unable to de-quarantine the disk group, collect array logs and contact HPE Support.

Disk drive fails and reconstruction does not start automatically

Symptom

Reconstruction does not start automatically after a drive failure.

Cause

No compatible spares are available.

Action

1. Replace each failed drive and then start reconstruction manually using one of the following methods:

- Add each new drive as a global spare.
- Enable the **Dynamic Spare Capability** option to use the new drives without designating them as spares.

NOTE: Depending on the disk group RAID level and size, disk speed, utility priority, and other processes running on the storage system, reconstruction can take hours or days to complete. You can stop reconstruction only by deleting the disk group. Deleting a disk group will cause permanent data loss.

Disk drive failure with dynamic spares enabled/disabled

Symptom

One or more drives in a disk group have failed. (Does not apply to MSA-DP+)

NOTE: Replacement drives must be compatible. Drives should be of the same type (SSD, Enterprise HDD for a standard tier disk group, or Midline HDD for an archive tier disk group) and of the same or larger capacity as the remaining drives in the disk group.

Solution 1

Cause

At the time a drive failed, the dynamic spares feature was enabled, and a properly sized disk drive was available to use as a spare.

Action

After the system completes reconstruction of the disk group, replace the failed drive.

Solution 2

Cause

At the time a drive failed, the dynamic spares feature was enabled but no compatible drive was available to use as a spare.

Action

1. Supply a compatible spare so the system can automatically use the new drive to reconstruct the disk group.
2. Replace the failed drive after reconstruction is complete.

Solution 3

Cause

At the time a drive failed, the dynamic spares feature was disabled, and no compatible global spare was available.

Action

1. Supply a compatible drive.
2. Use the CLI or SMU to assign it as a global spare.
3. Replace the failed drive after reconstruction is complete.

Solution 4

Cause

- RAID 1 or RAID 5: Two or more drives have failed in a disk group.
- RAID 6: Three or more drives have failed in a disk group.
- RAID 10: Two or more drives have failed in the same subgroup.

Any of the previous conditions makes the data in the disk group inaccessible and at risk.

Action

1. Attempt to recover by reseating the last failed drives.
2. If reseating the last failed drive does not resolve the issue, collect array logs and contact HPE Support.

Disk drive is marked as LEFTOVR

Symptom

The drive status of one or more drives in one enclosure is marked as LEFTOVR.

Solution 1

Cause

There are `MEDIUM` / `SMART` / `PROTOCOL` / `I/O TIMEOUT` errors for the drives.

Action

- If all disk groups are online, consider replacing the disk drive as an option to resolve the issue.
- If any disk group is offline, collect array logs and contact HPE support.

Solution 2

Cause

- Loose cabling
- Power loss or power supply issue
- A site issue

Action

- Resolve the cabling, power, or site issue.
- If all disk groups are online, consider clearing the metadata on the drive and use it as a spare to reconstruct disk groups as necessary.
- If you are unsure of the correct action to take, collect array logs and contact HPE Support.

Solution 3

Action

If this disk drive was a member of a disk group on another system, and that disk group does not exist on this system and, if all your disk groups are FTOL, clear the drive metadata if it is not required for a disk group from another array.

Drive is offline

Symptom

Drive is offline and the Fault/UID status (amber/blue) indicates a problem.

Solution 1

Cause

There is an error, failure, or critical fault with the drive.

Action

1. Review the event log for specific information regarding the fault.
2. Isolate the fault and replace the drive.
3. If you are unsure of what action to take, contact HPE Support.

Solution 2

Cause

There is no power or the drive is offline.

Action

1. Confirm that the disk drive is fully inserted and latched into place.
2. Verify that the enclosure is powered on.

Solution 3

Cause

An Event 8 error condition reports one of the following conditions for the drive:

- A hardware error
- An Illegal Request sense code for a command the disk drive supports
- A media error
- A SMART error

Action

- If all volumes and disk groups are online and available, then replace the drive.
- If the drive is marked as **Leftover** or **Failed**, and data recovery is needed, collect array logs and contact HPE Support.

  Related information: **Disk drive is marked as LEFTOVR**.

Solution 4

Cause

An Event 8 error condition reports that the RAID controller can no longer detect the drive.

Action

- Reseat the drive.
- If all the volumes and disk groups are online and available, then replace the drive.
- If data recovery is needed, collect array logs and contact HPE Support.

Solution 5

Cause

An Event 8 error condition reports that RAID 6 logic intentionally failed the drive.

Action

Replace the drive.

Failure reading data block

Symptom

An Event 542 or Event 543 was generated.

Cause

If the host did not get a read or write failure, the event occurred while reading or writing the metadata for the disk group. The event may have also appeared during a reconstruction.

Action

1. Do not reboot the array.
2. Gather logs, contact HPE Support, and provide the following information:
   - Event 542 information:
     - Drive name
     - Drive serial number
     - Logical Block Address (LBA) of the affected disk group

- LBA of the affected disk drive
- Enclosure slot number
- Enclosure number

- Event 543 information:
  - Name of affected volume
  - Serial number of the affected volume
  - LBA of the affected volume
  - Name of the affected disk group
  - Serial number of the affected disk group

3. If necessary, restore data from the last known good backup.

Multiple disk drive failures

Symptom

Two or more disk drives have failed.

Solution 1

Cause

- RAID 1: Two disk drive failures cause the disk group to enter a QTOF or OFFL state.
- RAID 5: Failure of two or more disk drives in a disk group causes the disk group to enter a QTOF or OFFL state.
- RAID 6: Failure of more than two disk drives in a disk group to enter a QTOF or OFFL state.
- RAID 10: Failure of both disk drives in the same disk group causes the disk group to enter a QTOF or OFFL state.
- DP+: When one drive fails, it is rebuilt on internal to the disk-group spare capacity. Likewise when a second drive fails, it will be rebuilt on internal disk-group capacity. When a third drive fails, the system will have a mixture of degraded and fault tolerant stripes of data. When a fourth drive fails, the system will have a mixture of critical, degraded, and fault tolerant stripes. In this case, the disk-group will go into a Rebalance Fault Tolerant (REFT) state. The system will degrade fault tolerant stripes to rebuild critical stripes, resulting in best system overall fault tolerance where it could survive another disk failure. As the number of disks increases the ability to withstand more drive faults increases.

Action

1. If the disk group is QTOF, it gets de-quarantined automatically after the drives are recognized. Review logs to determine if further action is required.

   If the virtual disk group is quarantined or offline, collect array logs and contact HPE Support.

Solution 2

Cause

RAID 6: Failure of two disk drives in a disk group causes the disk group to enter a CRIT state.

Action

If multiple spares are available, reconstruction begins automatically.

Solution 3

Cause

RAID 10: Two or more drives in different sub-disk groups have failed AND two or more disk drives in the same sub-disk group have failed.

Action

Collect array logs and contact HPE Support.

Issue with reported spare capacity in a DP+ configuration

Symptom

Configured spare capacity and actual spare capacity do not match.

Cause

One or more drives in the disk group has failed.

Action

Replace the failed drives.

Disk group member unavailable

Symptom

A disk group is in the offline (OFFL), critical (CRIT), or degraded (FTDN) state.

Cause

A failed disk causes a disk group to enter a critical state for RAID 1, RAID 5, RAID 10, or a degraded state for RAID 6, or two failed disks causes a RAID 6 disk group to enter a critical state.

Action

- If a spare is already available, reconstruction automatically begins.
- If a spare is not available, replace the failed drive, and add it as a spare.

Disk group quarantined during array boot

Symptom

During boot, one or more disk drives from the same disk group (or the same sub-disk group) are quarantined.

Solution 1

Cause

- RAID 5: During boot, multiple disk drives from the same disk group (or the same sub-disk group) go missing and are marked as QTOF.
- RAID 6: During boot, more than two disk drives go missing from the disk group. The disk group status is marked as QTOF.
- RAID 10: During boot, both the disk drives from the same sub-disk group go missing. The disk group status is marked as QTOF.

Action

1. Perform a manual rescan.
2. If the disk group does not automatically de-quarantine, or you are unsure of the correct action to take, collect array logs and contact HPE Support.

Solution 2

Cause

The wrong controller took ownership of the disk group during boot, and the last know cache and other disk group information is not available on the current controller.

Action

1. Shut down the system.
2. Perform one of the following actions:
   - If available, insert the controller that previously owned the disk group, remove the controller that took ownership of the disk group, and boot.
   - If the controller that was the previous owner is not available, then manually de-quarantine the disk group.
3. If you are unsure of the correct action to take, collect array logs and contact HPE Support.

Port troubleshooting

Expansion port not working

Symptom

A connected expansion port is not working and the Expansion Port Status LED is off.

Cause

The link is down.

Action

1. Inspect cable connections and reseat if necessary.
2. Inspect cables for damage.
   a. Swap cables to determine the fault.
   b. Replace cable if necessary.
3. In the SMU, review event logs for indicators of a specific fault in the expansion port. Follow any Recommended Actions.
4. If the issue is not resolved, contact collect array logs and contact HPE Support.

Host port not working

Symptom

A connected host port is not working, and the Host Link Status LED is off.

Cause

The link is down.

Action

1. Inspect cable connections and reseat if necessary.
2. Verify that the SFP is fully seated.
3. Inspect cables for damage.
   a. Swap cables to determine the fault.
   b. Replace cable if necessary.
4. Verify that the switch, if any, is operating properly. If possible, test with another port.
5. Verify that the HBA or NIC on the host is fully seated, and that the PCI slot is powered on and operational.
6. In the SMU, review event logs for indicators of a specific fault in a host datapath component. Follow any Recommended Actions.
7. If data hosts are having trouble accessing the storage system, and you cannot locate a specific fault or cannot access the event logs, see the *HPE MSA 1060/2060/2062 Installation Guide* for more information.
8. If the issue is not resolved, collect array logs and contact HPE Support.

iSCSI host port issues

Symptom

Cannot ping array iSCSI ports from the host.

Solution 1

Cause

Array host port issues.

Action

1. Check **Host port not working**.
2. Verify the array host port configuration: IP addresses, gateway, netmask, speed.

Solution 2

Cause

Network infrastructure configuration issues.

Action

Use standard networking troubleshooting procedures to isolate faults on the network.

Solution 3

Cause

Port configuration issues on the host.

Action

Check the configuration on the host for IP address, gateway, subnet mask, and speed.

Network port not working

Symptom

Cannot ping the management network ports.

Solution 1

Cause

Host network configuration or network switch configuration issues.

Action

1. Ping the gateway to verify host configuration.
2. Examine switch for VLAN and other configuration issues.

Solution 2

Cause

Network configuration on the array is not as expected.

Action

1. Log in to the other controller or connect a CLI cable to the unresponsive controller.
2. Verify that network port status is Up with expected IP address, gateway, and subnet mask settings.

Management controller troubleshooting

CLI is inaccessible using the CLI cable

Symptom

Cannot access the CLI when using the CLI cable.

Solution 1

Cause

There is an issue with the host terminal emulator.

Action

1. Connect a CLI cable to the array.
2. If the banner and login prompt do not appear in the terminal emulator, for Windows systems prior to Windows 10/Server 2016, ensure that the Windows USB device driver is installed, and for Linux systems, verify that the parameters are correct for the device driver.
3. If the banner and login prompt still do not appear in the terminal emulator for Windows systems, disconnect and quit the terminal emulator program, then disable and re-enable the USB port and restart the terminal emulator.

Solution 2

Cause

The management controller is unresponsive.

Action

1. Connect the CLI cable to the other controller.
2. Check the unresponsive controller status using the `show system` and `show redundancy-mode` commands.
3. If still unresponsive, reseat the unresponsive controller.
4. If still unresponsive, power cycle the array.

Configuration information is lost, and array management, alert and event messaging, and logging no longer function

Symptom

Host I/O is operating normally, but the following occurs:

- Configuration information is lost.
- Array management, alert and event messaging, and logging stop working.

Cause

When the Management Controller (MC) fails, the result may be a loss of the controller configuration information on the affected controller. The array cannot be managed from the impacted controller, but through the partner controller, if available.

Action

1. Restart the MC of the affected controller using either the CLI or SMU while logged in to the other controller.
2. If the configuration between the two controllers does not match, verify Partner Firmware Upgrade (PFU) is enabled.
3. After enabling PFU, if the configuration does not match, check the logs to determine if the array has had a memory card failure. If so, replace the controller during a scheduled maintenance window.

Management controller is not accepting user login

Symptom

The Storage Controller (SC) component is functioning properly, and all host I/O is being handled correctly, but the Management Controller (MC) component is not accepting a user login.

Action

1. Restart the MC from the other controller, if possible.
2. If the issue is not resolved, restart the corresponding SC from the other controller, and then restart the MC from the other controller.
3. If the issue is not resolved, halt all host I/O and then restart the array.

Management controller is unresponsive

Symptom

The MC is unresponsive. Array logs may contain a 139, 152, 153, 156, or 237 message.

Solution 1

Action

1. Log in to the MC on the other controller and restart the unresponsive MC.
2. Verify the network connectivity by issuing a `ping` command to the MC IP address of the controller. Ensure that an issue does not exist with the intranet on which the MSA is installed.

   Example setup:

   Controller A—IP 15.5.224.9

   Controller B—IP 15.5.224.10

   ```
   # ping 15.5.224.9 Reply from 15.5.224.9 bytes=32 time=60
   ```

3. If the `ping` command does not return a valid response, investigate the network or cabling and host-side issues. For example, verify that you can ping the gateway.
4. If Controller A responds to the ping but is not enabling logins through the CLI or SMU, log in to the other controller through the CLI.

5. Attempt to restart the MC on the unresponsive controller.

   Example:

   ```
   #restart mc A
   Continue? yes <enter return>
   Success: MC A restarted.
   ```

   NOTE: There may be up to a minute delay between the **Continue** prompt and the **Yes <enter return>** prompt to when the Success information message is displayed.

6. Wait a few moments after the Success information message is displayed for the restart process to complete, and then try to log in to the other controller.
7. If the MC remains unresponsive, shut down and restart the unresponsive controller using the other controller, and then recheck the MC.

   NOTE: During restart, you will briefly lose communication with the specified management controllers.

   Example: from Controller B, shut down A, wait for A to shut down, and then restart A.

   ```
   #shutdown A

   #restart sc a
   ```

   This action will restart the MC on controller A as well.

   NOTE: In a single controller environment, if the CLI command is not successful, quiesce all host I/O before restarting the controller.

Solution 2

Cause

Firmware Flash is blocked, which puts the system in a firmware upgrade loop.

Action

1. Verify that all host I/O to the array is quiesced or halted completely.
2. Restart both Storage Controllers using either the CLI or SMU.
3. Use the SmartComponent to update the firmware.

Solution 3

Cause

A SC is unable to communicate with the MC.

Action

1. Restart the MC using one of the following methods:

- Run the following CLI command:

  ```
  #restart mc A
  Continue? yes <enter return>
  Success: MC A restarted.
  ```

- Restart the MC using the SMU.

NOTE:

- During the restart process, you will briefly lose communication with the specified management controllers.
- If restarting the MC does not resolve the issue, follow the instructions in Solution 4 to reseat the offending controller during a maintenance window.

2. If the SC is still unable to communicate with the MC, then verify that all host I/O to the array is quiesced or halted.
3. Shut down the controller using the `shutdown` command.

   Syntax:

   ```
   shutdown a|b|both
   ```

   NOTE: Reseating a controller with active I/O in a dual controller array will cause I/O failover to the other controller.

4. Reseat the controller.

   IMPORTANT: Reseating a controller with active I/O in a dual controller array will cause I/O failover to the other controller.

User login troubleshooting

LDAP user cannot log in to storage system

Symptom

LDAP user is unable to log in to the storage system.

Action

1. Verify that user credentials are correct by logging to the Active Directory from another system.
2. Verify that other LDAP users are able to log in to the storage system.
3. Verify that the LDAP parameters are set correctly.
4. Verify that the LDAP User-Group name exactly matches the group name in the Active Directory.
5. Verify that the LDAP user is not a member of more than 100 Active Directory groups.

LDAP user gets inconsistent permissions

Symptom

LDAP user gets inconsistent permissions when logging in to the storage system.

Action

Verify that the user is not a member of more than one User-Group on the storage system.

User cannot change settings or load firmware

Symptom

- User is unable to make configuration changes to the storage system.
- User is unable to create/modify/delete users.
- User is unable to load firmware.

Cause

Incorrect role is assigned to user.

Action

Verify that the user has the `manage` role.

User cannot log in to one or more management interfaces

Symptom

User is unable to log in to any management interfaces, or user can access only one management interface.

Solution 1

Cause

Lost password.

Action

1. Bypass possible network issues by verifying the login credentials by using a USB CLI cable to log into the USB CLI console.

   To obtain a valid user name and password, contact your system administrator.

   For new systems, or systems that have been restored to defaults, you will be prompted to create a new user name and password combination using the **setup** user account.
2. If no user logins can be recovered, collect array logs and contact HPE Support.

Solution 2

Cause

The interface is not selected in User Management or a protocol is disabled/blocked.

Action

1. Verify that the user has the correct interfaces selected in User Management.
2. Verify that the management protocol is enabled (HTTPS, HTTP, SSH, TELNET).
3. Verify that the protocol is not blocked by a firewall by logging into the management interface from a system on the same subnet as the management port.
4. If the issue is still not resolved, collect array logs and contact HPE Support.

Solution 3

Cause

The interface login is not enabled for the user.

Action

Log in to another interface and verify that the interface for which the login failed is enabled for the user.

Power supply troubleshooting

For information on replacing an AC or DC Power and Cooling Module, see *HPE MSA 1060/2060/2062 Power and Cooling Module Replacement Instructions* on the Hewlett Packard Enterprise Support Center website:

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

Power supply warning or failure

Symptom

You receive a power supply warning with event code 551, or the power supply fails.

Action

1. Verify that the power supply modules are properly seated in their slots and that their latches are locked.
2. Verify that all power supply units are working.
3. Replace the power supply, if necessary.

IMPORTANT: Do not leave slots open for more than two minutes. If you must replace a module, leave the old module in place until you have the replacement ready. Leaving a slot open negatively affects the airflow and might cause the unit to overheat.

Power supply is off

Symptom

The power supply is off, and the power supply Input Power Source LED is off.

Solution 1

Cause

Power supply is not receiving adequate power.

Action

1. Verify that the power cable is properly connected to the power supply and the power source.
2. Verify that the power supply is firmly locked into position.
3. In the SMU, review event logs for specific information regarding the fault. Follow any Recommended Actions.
4. If the issue is not resolved, collect array logs and contact HPE Support.

Solution 2

Cause

Power supply module status is listed as failed or you receive a voltage event notification.

Action

1. Ensure that the switch is turned on.
2. Ensure that the power cables are firmly plugged into both the power supply and an appropriate electrical outlet.
3. Replace the power supply, if necessary.

IMPORTANT: Do not leave slots open for more than two minutes. If you must replace a module, leave the old module in place until you have the replacement ready. Leaving a slot open negatively affects the airflow and might cause the unit to overheat.

Power supply is not working properly

Symptom

The power supply Voltage / Fan Fault / Service Required LED is amber.

Cause

The power supply unit or fan is operating at an unacceptable voltage/RPM level, or has failed.

NOTE: When isolating faults in the power supply, remember that the fans in both modules receive power through a common bus on the midplane. If a power supply unit fails, the fans continue to operate normally.

Action

1. Verify that the power cable is properly connected to the power supply and the power source.
2. Verify that the power supply is firmly locked into position.
3. If the issue is not resolved, collect array logs and contact HPE Support. Replacement of the power supply may be necessary.

Controller module or I/O module troubleshooting

Controller module is off

Symptom

The controller OK LED is off.

Cause

The controller module is not powered on or the controller module has failed.

Action

1. Verify that the controller module is fully inserted and latched in place, and that the enclosure is powered on.
2. View the event log for specific information regarding the failure.
3. If the issue is not resolved, collect array logs and contact HPE Support.

Controller is not working

Symptom

The controller is not active, and the enclosure rear panel Fault LED is amber.

Cause

Fault detected or service action required.

Action

1. Restart the controller from the other controller using the CLI or SMU.
2. If the issue is not resolved, remove the controller and reinsert it.
3. If the issue is not resolved, collect array logs and contact HPE Support. Replacement of the controller may be required.

I/O module is offline

Symptom

The I/O module has failed and the enclosure front panel Module Fault LED is amber.

Cause

A fault condition exists. The I/O module has likely failed its self-test.

Action

1. Verify the LEDs on the back of the controller enclosure to narrow the fault to an I/O module, a power supply, or a SAS port.
2. View the event log for specific information regarding the fault, and follow any recommended actions.

3. Remove and reinstall the I/O module, then check the event log for errors.
4. If the issue is not resolved, isolate the fault, collect array logs, and contact HPEI Support. Replacement may be necessary.

Chassis troubleshooting

NOTE: Avoid unnecessary chassis replacement. Other than a mechanical failure, it is rare to have a chassis or midplane issue requiring replacement.

For information on replacing the chassis, see *HPE MSA Chassis Replacement Instructions* on the Hewlett Packard Enterprise Support Center website:

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

Issue installing a Field Replaceable Unit

Symptom

You have trouble inserting/installing a FRU.

Cause

- A FRU fails to seat properly, cannot be fully inserted, or once inserted will not slide all the way into the slot.
- The locking mechanism on a FRU is fully closed and the FRU does not lock in place or locked down properly.
- A chassis with a defect that prevents a FRU from being installed.

Action

1. Verify that a physical problem does not exist with the specific FRU before replacing the chassis. A mechanical issue of this type will not require log evaluation.
2. If there is a problem with the physical chassis itself, replacement may be necessary. Collect array logs and contact HPE Support.

Midplane issue diagnosis

Symptom

A 314 midplane FRU notification occurs, and any of the following events displays in the array events or logs: 274, 358, 495, 521, or 602.

Action

Collect array logs and contact HPE Support.

Obtaining replacement licenses

Prerequisites

IMPORTANT: Verify if any licenses must be obtained and installed on the new chassis. Existing license replacements are provided at no extra cost by HPE when executing a chassis replacement.

Procedure

1. Record the Licensing Serial Number of both the existing and replacement chassis to obtain new replacement licenses. The Licensing Serial Number is shown in the web-based Storage Management Utility (SMU) from the Install License panel (**Maintenance** > **Support** > **Licensing**), or by executing the `show license` command.

   Refer to the CLI Reference Guide or Storage Management Guide for details on how to locate the Licensing Serial Number, if necessary.

2. Obtain the new replacement licenses directly from the My License Portal (**https://myenterpriselicense.hpe.com/**) by clicking **Rehost Licenses** on the main page and choosing the appropriate locking id, which corresponds to the Licensing Serial Number of the existing array enclosure. The Rehost Licenses process will require the Licensing Serial Number of the existing and upgraded array enclosures to complete the process.

   NOTE: If multiple licenses need to be rehosted, make sure to rehost all licenses at the same time.

If there are problems or questions, HPE Support can be accessed through the License portal previously referenced.

Websites

GENERAL WEBSITES

Hewlett Packard Enterprise Information Library

https://www.hpe.com/info/EIL

Single Point of Connectivity Knowledge (SPOCK) Storage compatibility matrix

https://www.hpe.com/storage/spock

HPE Support

https://support.hpe.com/hpesc/public/home/

Storage white papers and analyst reports

https://www.hpe.com/storage/whitepapers

MSA WEBSITES

MSA QuickSpecs

https://www.hpe.com/info/qs

MSA Manuals

https://www.hpe.com/info/MSAdocs

MSA Firmware

https://www.hpe.com/storage/MSAFirmware

HPE MSA Support Material Access

https://www.hpe.com/storage/MSASupportMaterialAccess

MSA Health Check

https://www.hpe.com/storage/MSAHealthCheck

Support and other resources

Accessing Hewlett Packard Enterprise Support

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

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

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

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

Information to collect

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

Accessing updates

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

  Hewlett Packard Enterprise Support Center

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

  Hewlett Packard Enterprise Support Center: Software downloads

  https://www.hpe.com/support/downloads

  My HPE Software Center

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

- To subscribe to eNewsletters and alerts:

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

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

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

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

Remote support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which 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

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

HPE Proactive Care services

https://www.hpe.com/services/proactivecare

HPE Datacenter Care services

https://www.hpe.com/services/datacentercare

HPE Proactive Care service: Supported products list

https://www.hpe.com/services/proactivecaresupportedproducts

HPE Proactive Care advanced service: Supported products list

https://www.hpe.com/services/proactivecareadvancedsupportedproducts

Proactive Care customer information

Proactive Care central

https://www.hpe.com/services/proactivecarecentral

Proactive Care service activation

https://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

https://www.hpe.com/support/ProLiantServers-Warranties

HPE Enterprise and Cloudline Servers

https://www.hpe.com/support/EnterpriseServers-Warranties

HPE Storage Products

https://www.hpe.com/support/Storage-Warranties

HPE Networking Products

https://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:

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

Additional regulatory information

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

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

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

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

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

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

Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, 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.