Hewlett Packard Enterprise

HPE MSA Storage Troubleshooting Guide

Abstract

This document provides information about troubleshooting HPE MSA Storage Systems. Use this document to troubleshoot and maintain Hewlett Packard Enterprise MSA Storage Systems.

Part Number: Q1J79-62028
Published: August 2018
Edition: 3

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

Troubleshooting HPE MSA Storage

The following section details the events and their descriptions used to troubleshoot the following MSA Storage Systems:

- **Events and LEDs**
- **Disk drives**
- **Management controller issue**
- **User login issues**
- **Power supply faults and power cycle**
- **Chassis replacement**
- **Using the trust command**

Events and LEDs

Enable alerts

HPE strongly encourages you to enable alerts for better array monitoring. Set alerts to a Warning level to ensure that you are monitoring your system for possible failures.

Beginning with firmware versions VE270 and VL270 on the HPE MSA 1050/2050 systems, a **Welcome** panel appears on the **Home** topic until you complete a number of steps, including enabling alerts or acknowledging you do not want to receive alerts.

For older systems, HPE strongly recommends that you review your current alerts 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

INFORMATIONAL— used to inform all system events. This level will result in increased messages as all categories of messages will be reported. This level requires regular attention and review to ensure that **WARNING** (or **ERROR** or **CRITICAL**) events are not lost in the resulting expanded event report.

WARNING—recommended as it signals a problem exists that results in the loss of redundancy. If another similar failure occurs, a system outage (and resulting loss of access to data or possible loss of data integrity) occurs. HPE recommends that you investigate and resolve the event promptly.

ERROR—used to inform failure occurred that might affect data integrity or system stability. Correct the problem as soon as possible.

CRITICAL—used to inform failure occurred that might cause a controller to shut down. Correct the problem immediately.

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 Event Descriptions Reference Guide*. A FRU is any HPE orderable replacement part.

Documentation for HPE MSA Storage Systems is on the **Hewlett Packard Enterprise Support Center** website.

Events requiring FRU replacements

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.

LED indications and recommended actions

Table 1: LED indications and recommended actions

Indicator	Status	Cause	Action required
The enclosure front panel `Fault/Service Required` LED	Off	System is functioning properly.	No action required.
	Amber	A fault condition exists/ occurred. If installing an I/O module FRU, the module has not gone online and likely failed its self-test.	Verify the following: • The LEDs on the back of the controller enclosure to narrow the fault to a FRU, connection, or both • The event log for specific information regarding the fault; follow any Recommended Actions • If installing an IOM FRU, try removing and reinstalling the new IOM, and check the event log for errors If the previous actions do not resolve the fault, isolate the fault, and contact an authorized service provider or HPE technical support for assistance. Replacement of a FRU may be necessary.
The enclosure rear panel `FRU OK` LED	Green	Controller operating properly.	No action required.
	Blinking green	System is booting	Wait for system to boot.

Table Continued

Indicator	Status	Cause	Action required
	Off	The controller module is not powered on. The controller module has failed.	Verify the following: • The controller module is fully inserted and latched in place, and the enclosure is powered on • The event log for specific information regarding the failure
The enclosure rear panel `Fault/Service Required` LED	Off	System is functioning properly.	No action required.
	Amber	A fault has been detected or a service action is required.	• Restart this controller from the other controller using the SMU or the CLI • If the previous action does not resolve the fault, remove the controller and reinsert it If the previous actions do not resolve the fault, contact an authorized service provider or HPE technical support for assistance. Replacement of the controller may be necessary.
	Amber; blinking regularly	• Hardware-controlled power-up error • Cache flush error • Cache self-refresh error	
Connected host port `Host Link Status` LED	Blinking green	System is functioning properly.	No action required.

Table Continued

Indicator	Status	Cause	Action required
	Off	The link is down.	• Check cable connections and reseat if necessary. • Verify that the SFP is fully seated. • Inspect cables for damage. Replace cable if necessary. • Defective cable causes fault. Swap cables to determine the fault. Replace cable if necessary. • Verify that the switch, if any, is operating properly. If possible, test with another port. • Verify that the HBA is fully seated, and that the PCI slot is powered on and operational. • In the SMU, review event logs for indicators of a specific fault in a host datapath component; follow any Recommended Actions. • Contact an authorized service provider or HPE technical support for assistance.
A connected port `Expansion Port Status` LED	On	System is functioning properly.	No action required.

Table Continued

Indicator	Status	Cause	Action required
	Off	The link is down.	• Check cable connections and reseat if necessary • Inspect cables for damage. Replace cable if necessary • Defective cable causes fault. Swap cables to determine the fault. Replace cable if necessary • In the SMU, review event logs for indicators of a specific fault in a host datapath component; follow any Recommended Actions. • Contact an authorized service provider or HPE technical support for assistance
A connected port `Network Port Link Status` LED	Green	System is functioning properly.	No action required.
	Off	The link is down.	Use standard networking troubleshooting procedures to isolate faults on the network.
The power supply `Input Power Source` LED	Green	System is functioning properly.	No action required.

Table Continued

Indicator	Status	Cause	Action required
	Off	The power supply is not receiving adequate power.	• Verify that the power cable is properly connected and check the power source to which it connects • Check that the power supply FRU is firmly locked into position • In the SMU, check the event log for specific information regarding the fault; follow any Recommended Actions • If the previous action does not resolve or isolate the fault, contact an authorized service provider or HPE technical support for assistance

Table Continued

Indicator	Status	Cause	Action required
The power supply `Voltage/Fan Fault/ Service Required` LED	Off	System is functioning properly.	No action required.
	Amber	The power supply unit or a fan is operating at an unacceptable voltage/RPM level, or has failed.	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: Verify the following: • The power supply FRU is firmly locked into position. • The power cable is connected to a power source. • The power cable is connected to the power supply module. • If the previous action does not resolve the fault, isolate the fault and contact HPE technical support for assistance.

Isolating a host-side connection fault

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 following guides for more information.

- *HPE MSA P2000 G3 User Guide*
- *HPE MSA 1040 User Guide*
- *HPE MSA 2040 User Guide*
- *HPE MSA 1050 User Guide*
- *HPE MSA 2050 User Guide*

Disk drives

Disk drive, Vdisk, and disk group best practices

Enable background scrub and background disk scrub so that bad disk drive blocks can be detected and repaired proactively.

Keep all Vdisks and disk groups fault tolerant. Replace failed disk drives immediately, or have spare disk drives available in the array. Maintaining spares allows Vdisk and disk group reconstruction to begin as soon as a disk drive fails.

Disk drive LEDs

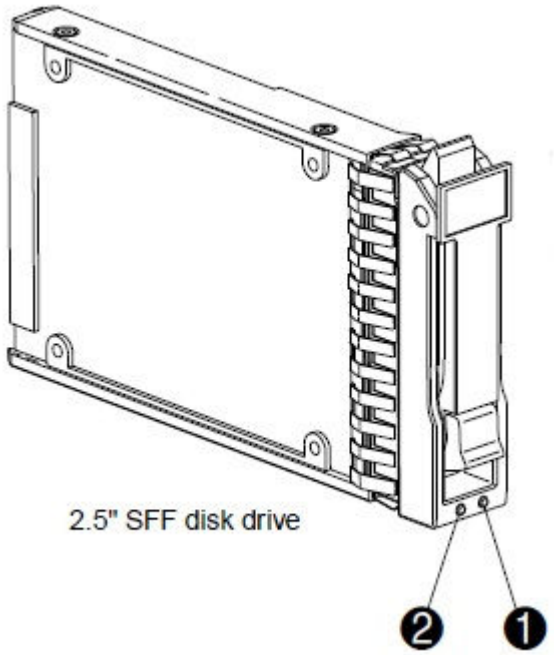

Figure 1: 2.5” SFF disk drive

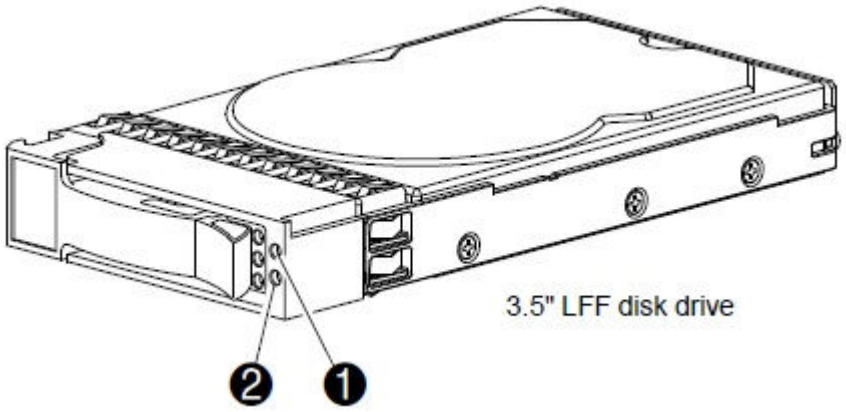

Figure 2: 3.5” LFF disk drive

Table 2: LEDs - disk drive LEDs

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

Table 3: LEDs - disk drive combinations

Online/Activity (green)	Fault/UID (amber/blue)	Description	Recommended action
On	Off	Normal operation. The drive is online, but it is not currently active.	No action required.
Blinking irregularly	Off	The drive is active and operating normally.	No action required.
Off	Amber; blinking regularly (1 Hz)	Offline; the drive is not being accessed. A drive error might be received for this device. Further investigation is required.	• Check the event log for specific information regarding the fault • Isolate the fault • Contact an authorized service provider or HPE technical support for assistance
On	Amber; blinking regularly (1 Hz)	Online; possible I/O activity. A drive error might be received for this device. Further investigation is required.	
Blinking irregularly	Amber; blinking regularly (1 Hz)	The drive is active, but a drive error might be received for this drive. Further investigation is required.	
Off	Amber; solid	Offline; no activity. A failure or critical fault condition is identified for this drive. This may indicate that the drive is leftover rather than a problem with the drive itself.	• Check the event log for specific information regarding the fault • Isolate the fault • Contact an authorized service provider or HPE technical support for assistance
Off	Blue; solid	Offline. A management application (for example, the SMU) selects the drive.	No action required.
On or blinking	Blue; solid	The controller is driving I/O to the drive. A management application (for example, the SMU) selects the drive.	No action required.

Table Continued

Online/Activity (green)	Fault/UID (amber/blue)	Description	Recommended action
Blinking regularly (1 Hz)	Off	**CAUTION:** Do not remove the drive. Removing a drive closes the current operation and causes data loss. The drive is reconstructing.	No action required.
Off	Off	Either there is no power, the drive is offline, or the drive is not configured.	Check that the disk drive is fully inserted, and latched in place. Verify that the enclosure is powered on.

Disk error conditions and recommended actions

Table 4: Disk error conditions and recommended actions

Symptom	Recommended Action
Event 8 reports that the RAID controller can no longer detect the disk.	Reseat the disk. If all the volumes and Vdisks or disk groups are online and available, then replace the disk.
Event 8 reports a media error for the disk.	If all the volumes and Vdisks or disk groups are online and available, then replace the disk. Otherwise, if the drive is marked as **Leftover** or **Failed**, and data recovery is needed, see **Using the trust command** on page 40.
Event 8 reports a hardware error for the disk.	If all the volumes and Vdisks or disk groups are online and available, then replace the disk. Otherwise, if the drive is marked as **Leftover** or **Failed**, and data recovery is needed, see **Using the trust command** on page 40.
Event 8 reports an Illegal Request sense code for a command the disk supports.	If all the volumes and Vdisks or disk groups are online and available, then replace the disk. Otherwise, if the drive is marked as **Leftover** or **Failed**, and data recovery is needed, see **Using the trust command** on page 40.
Event 8 reports that RAID 6 logic intentionally failed the disk.	Replace the disk.
Event 55 reports a SMART error for the disk.	If all the volumes and Vdisks or disk groups are online and available, then replace the disk. Otherwise, if the drive is marked as **Leftover** or **Failed**, and data recovery is needed, see **Using the trust command** on page 40.

Table Continued

Symptom	Recommended Action
At the time a disk failed, the dynamic spares feature was enabled, and a properly sized disk was available to use as a spare.	No action required; the system automatically uses that disk to reconstruct the Vdisk or disk group. Replace the failed disk when reconstruction has completed.
At the time a disk failed, the dynamic spares feature was enabled but no properly sized disk was available to use as a spare.	The best option is to supply an appropriately sized spare so the system can automatically use the new disk to reconstruct the Vdisk or disk group. Then replace the failed disk once reconstruction has completed.
At the time a disk failed, the dynamic spares feature was disabled and no dedicated spare or properly sized global spare was available.	The best option is to supply an appropriately sized disk, using the SMU or CLI to assign it as a dedicated Vdisk spare or global spare. The system will automatically reconstruct the Vdisk or disk group, you should then replace the failed disk once reconstruction is complete.
The status of the Vdisk or disk group that originally had the failed disk status is **FTOL** (Fault Tolerant Online). A dedicated Vdisk spare or global spare has been successfully integrated into the Vdisk or disk group. The replacement disk can be assigned as either a global spare or a dedicated Vdisk spare.	Use SMU or CLI to assign the new disk as either a global spare or a dedicated Vdisk spare.
The status of the disk installed is **LEFTOVR**	All the member disks in a Vdisk or disk group contain metadata in the first sectors. The storage system uses the metadata to identify Vdisk or disk group members after restarting or replacing enclosures. If this disk was a member of a Vdisk or disk group on another system, and that Vdisk or disk group does not exist on this system and, if all your Vdisks or disk groups are FTOL, clear the drive metadata if it is not required for a Vdisk or disk group from another array.
If the status of the Vdisk or disk group that originally had the failed disk status is **OFFL (Offline) or QTOF (Quarantined Offline)**, one or more disks have failed in a **RAID-0** Vdisk or disk group. Two or more disks have failed in a **RAID-1, 3, or 5** Vdisk or disk group. Three or more disks have failed in a **RAID 6** Vdisk or disk group.	The data in the Vdisk or disk group is inaccessible and at risk. Attempt to recover by reseating the last failed drives. If that does not resolve the issue, and data recovery is needed, see **Using the trust command** on page 40. If you are unsure of the correct action to take, collect array logs and contact HPE Support.

Table Continued

Symptom	Recommended Action
The status of the Vdisk or disk group that originally had the failed disk indicates that the Vdisk or disk group is being rebuilt.	Wait for the Vdisk or disk group to complete its operation. Do not remove any leftover or failed drives until the rebuild is complete. Once the rebuild has completed successfully, replace any failed drives.
The status of the Vdisk or disk group that originally had the failed disk is **CRIT (Critical) or QTDN (Quarantined with a down disk)**.	If after inserting a known good replacement disk it is not detected, the disk slot in the enclosure may be bad. If there is another available disk slot, insert the known good replacement disk in the alternative slot. If the disk is detected, using the SMU or CLI to assign it as a dedicated Vdisk spare or global spare. Once the reconstruction has completed, replace the enclosure. If the drive is not detected in the alternative slot, the drive is not a good disk and you will need to use a new drive as a spare.

For information on replacing the drive modules, see *HPE MSA Drive Module Replacement Instructions* on the **Hewlett Packard Enterprise Support Center** website.

Disk drive bay numbers

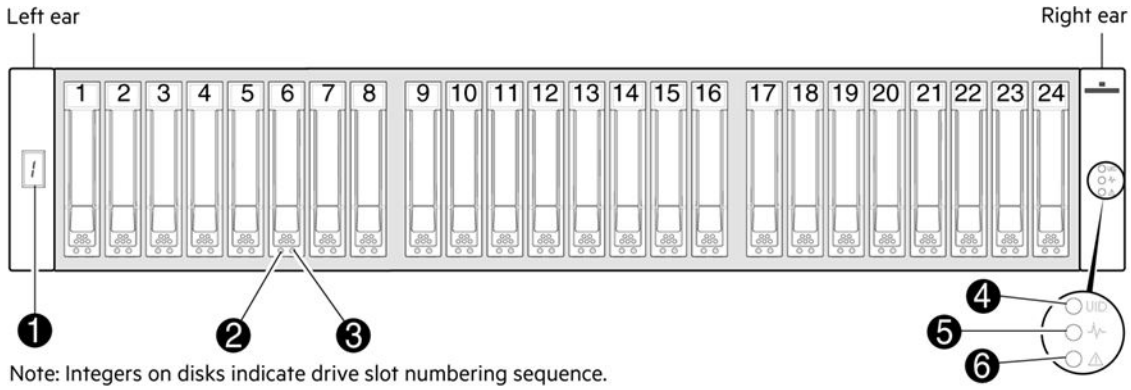

Figure 3: MSA 2040 SAN Array SFF enclosure

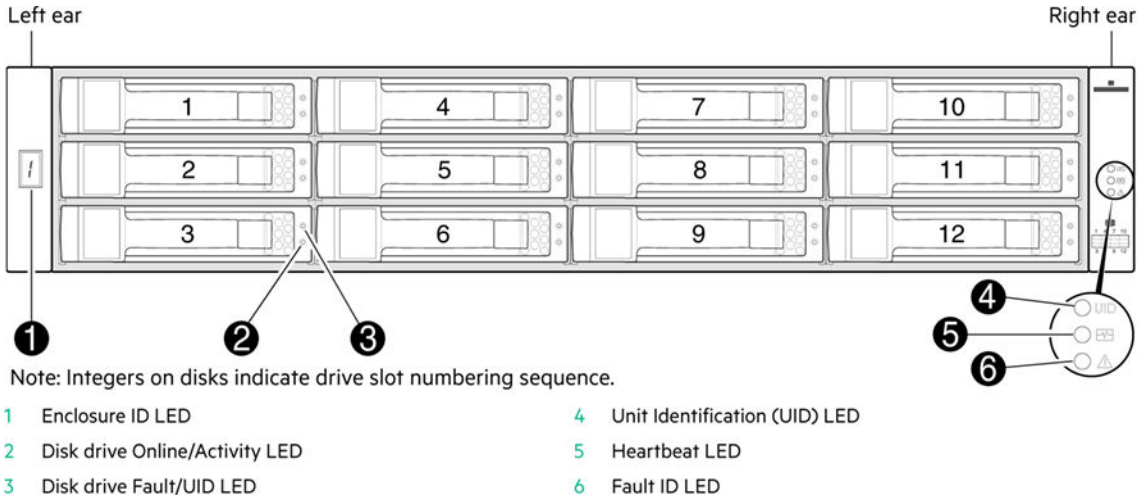

Figure 4: MSA 2040 SAN Array LFF or supported drive enclosure

Drive status

Table 5: Drive usage

Usage	Displayed in the CLI or SMU	Description
Available	`AVAIL`	The drive is available for use.
Failed	`FAILED`	The drive is unusable, replace the drive. Reasons for this status include excessive media errors, SMART errors, drive hardware failures, unsupported drives.
Global 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.
Vdisk or disk group	`VDISK or VIRTUAL POOL, <VDISK>:Linear, <pool-ID>:<tier-name>`	The drive is used in a Vdisk or disk group.
Vdisk spare	`VDISK SP`	The drive is a spare assigned to a Vdisk.

Disk drive leftover

Table 6: Hard disk drive leftover scenarios

Symptom	Cause	Action required
A disk drive is marked as **LEFTOVR**.	Due to `MEDIUM/SMART/ PROTOCOL/ I/O TIMEOUT` errors.	Check if any Vdisks or disk groups are **OFFL**; if disks are offline, see **Using the trust command** on page 40. If all Vdisk or disk groups are online, consider replacing the disk drive as an option to resolve the issue.
Multiple drives all in one enclosure are marked as **LEFTOVR**.	If a Vdisk or disk group is spanned across multiple enclosures, and all drives in an enclosure went missing due to a site issue, a power loss or power supply issue, or loose cabling, the Vdisk or disk group may go **OFFL**. Once the issue is corrected and the disk drives are detected, those disks are marked as **LEFTOVR**.	It is not necessary to replace all the disk drives that are marked **LEFTOVR,** as it is not a disk drive issue. Resolve the site power, or cabling issue. See **Using the trust command** on page 40. If you are unsure of the correct action to take, collect array logs and contact HPE Support.
A disk drive in a redundant Vdisk or disk group (in this scenario, other than RAID 6) fails. The Vdisk or disk group is marked as OFFL or **QTOF** (quarantined offline).	A disk has failed, and a spare disk was used and reconstruction began. Before the reconstruction completes, another disk drive from the same Vdisk or disk group (same subgroup in RAID 10 or RAID 50) fails, the status of the second failed disk drive along with the spare disk is marked as **LEFTOVR**.	Attempt to dequarantine the Vdisk or disk group by reseating the last failed drives. If unable to dequarantine the Vdisk or disk group, see **Using the trust command** on page 40. If you are unsure of the correct action to take, collect array logs and contact HPE Support.

Table Continued

Symptom	Cause	Action required
A disk drive in a RAID 6 Vdisk or disk group (with 1 spare) fails and is marked as **LEFTOVR**.	Vdisk or disk group status is **FTDN**, a spare is initiated, and reconstruction begins. Before the reconstruction completes, if another member of the same Vdisk or disk group fails, its status is marked as **LEFTOVR** and the Vdisk or disk group status is **CRIT**. The reconstruction for the first failed drive does not stop but continues to completion. If a second spare is available when the second drive fails, it will not be used for reconstruction until the previous reconstruction completes. Once the reconstruction completes, the status of the Vdisk or disk group is back to **FTDN**. If a spare is available, reconstruction of the Vdisk or disk group for the second failed HDD begins.	Consider replacing the drives if the drives failed due to hardware issues. If the Vdisk or disk group go **QTOF** after the failure of a third disk drive member and before the reconstruction completes, attempt to dequarantine by reseating the last failed drives. If unable to dequarantine the Vdisk or disk group, collect logs and contact HPE Support.
Reconstruction stops for a RAID 6 Vdisk or disk group when a second drive fails	If a RAID 6 Vdisk or disk group experiences one drive failure, it would use a spare to reconstruct the Vdisk or disk group. But if a second drive fails during the reconstruction and no more spares are available, then the reconstruction stops, the Vdisk or disk group status becomes **CRIT**, and the second failed drive is **LEFTOVR**. If one or more members of the same Vdisk or disk group fails and is marked **LEFTOVR**, the Vdisk or disk group goes **OFFL**.	Clear the disk metadata on the drive used for reconstruction. Replace the second failed hard drive if possible, and set both to be spares. Allow the reconstruction to restart. If the Vdisk or disk group goes **OFFL**, do not clear the metadata and see **Using the trust command** on page 40. If you are unsure of the correct action to take, collect array logs and contact HPE Support.
Reconstruction restarts for a RAID 6 Vdisk or disk group when a second drive fails.	If a RAID 6 Vdisk or disk group experiences one drive failure, it would use a spare to reconstruct the Vdisk or disk group. If a second drive fails during the reconstruction and a spare is available, then the reconstruction restarts from 0%, and the second failed drive is **LEFTOVR**. If one more member of the same Vdisk fails and is marked **LEFTOVR**, the Vdisk or disk group goes **OFFL**.	Wait for the Vdisk or disk group reconstruction to complete. If the Vdisk or disk group goes **OFFL**, see **Using the trust command** on page 40. If you are unsure of the correct action to take, collect array logs and contact HPE Support.

Disk drive data block fenced

Failure to read a data block from a disk drive will generate event 542. Event 542 provides the following:

- Name
- Serial number
- Logical Block Address (LBA) of the Vdisk or disk group affected

Event 542 may also provide the enclosure:

- Slot number
- Serial number
- LBA of the disk drive affected

If the block is in a volume, event 543 is generated. Event 543 provides the following:

- Name of the volume affected
- Serial number of the volume affected
- LBA of the volume affected
- Name of the Vdisk or disk group affected
- Serial number of the Vdisk or disk group affected
- Related event with the event ID 542

If you see event 542 or 543 in the array logs and the host did not get a read or write failure, the event occurred while reading or writing the metadata for the Vdisk or disk group. The event may have also appeared during a reconstruction. Do not reboot the array. Back up all data on the affected Vdisk or disk group, then contact technical support for assistance. You may need to restore data from a backup.

To reduce the likelihood of needing to restore from a backup, enable background scrub and background disk scrub. Enabling scrubs allows bad disk drive blocks to be detected and repaired proactively. Keep all Vdisks or disk groups fault tolerant. Replace failed disk drives immediately, or have spare disk drives available in the array. Maintaining spares allows Vdisk or disk group reconstruction to begin as soon as a disk drive fails.

Vdisk or disk group status

Table 7: Vdisk or disk group status

Status	Displayed in the CLI or SMU	Description
Critical	`CRIT`	The Vdisk or disk group is online; however, some drives are down and the Vdisk or disk group is not fault tolerant.
Fault Tolerant with down drives	`FTDN`	The Vdisk or disk group is online and fault tolerant; however, some drives are down.

Table Continued

Status	Displayed in the CLI or SMU	Description
Fault Tolerant and online	FTOL	The Vdisk or disk group is online and fault tolerant.
Offline	OFFL	The Vdisk or 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 Vdisk or disk group is in a critical state and quarantined because some drives are missing.
Quarantined offline	QTOF	The Vdisk or disk group is offline and quarantined because some drives are missing.
Quarantined unsupported	QTUN	The Vdisk or disk group contains data in a format that is not supported by this system.
Quarantined with down drives	QTDN	The Vdisk or disk group is offline and quarantined because at least one drive is missing. For example, one drive is missing from a RAID 6.
Up	UP	The Vdisk and disk group is online and does not have fault-tolerant attributes.

Vdisk or disk group quarantined

If a Vdisk or disk group is quarantined, see the following table. You cannot dequarantine virtual disk groups without the assistance of knowledgeable support personnel.

Table 8: Vdisk or disk group quarantine during array boot (not available for MSA 1050/2050)

Level	Symptom	Action required
RAID 5	During boot, multiple disk drives from the Vdisk or disk group go missing and the Vdisk or disk group status is marked as **QTOF**.	The system automatically de-quarantines the Vdisk or disk group once the missing disk drives are recognized after a rescan. If the Vdisk (linear only) is not automatically de-quarantined, perform a manual de-quarantine. The Vdisk status changes to OFFL (if the drives are not recognized during a rescan). If the status of the Vdisk changes to OFFL, follow the troubleshooting steps **Using the trust command** on page 40 for solution.
RAID 6	During boot, more than two disk drives go missing from the Vdisk or disk group. The Vdisk or disk group status is marked as **QTOF**.	

Table Continued

Level	Symptom	Action required
RAID 10	During boot, both the disk drives from the same sub-Vdisk or sub-disk group go missing. The Vdisk or disk group status is marked as **QTOF**.	If the virtual disk group does not automatically de-quarantine or you are unsure of the correct action to take, collect array logs and contact HPE Support.
RAID 50	During boot, multiple disk drives from the same sub-Vdisk or sub-disk group go missing and the Vdisk or disk group status is marked as **QTOF**.	
N/A	The wrong controller took ownership of the Vdisk or disk group on boot. The last known cache and other Vdisk or disk group information is not available on the current controller.	Shut down the system, remove the controller that took ownership of the Vdisk or disk group, insert the controller that previously owned the Vdisk or disk group, and boot. If the controller that was the previous owner is not available, then manually de-quarantine the Vdisk. If you are unsure of the correct action to take, collect array logs and contact HPE Support.

NOTE:

- In the previous scenarios, the drives that went down/missing are not spares or reconstruction targets.
- In the previous scenarios, if the disk drives go missing after a rescan, the Vdisk or disk group status is marked as **OFFL**. See **Using the trust command** on page 40.

Table 9: Quarantine upon rescan

Level	Symptom	Action required
RAID 5	A disk drive is marked as **LEFTOVR** or goes missing, making the Vdisk or disk group status **CRIT**. Another disk drive fails or is missing, the Vdisk or disk group is marked as **QTOF**.	If the disk drives are recognized during a rescan, the system will dequarantine the Vdisk or disk group after two minutes If the disk drives are not recognized during a rescan for a Vdisk, perform a manual de-quarantine to change the Vdisk status to **OFFL**. If the Vdisk status changes to **OFFL**, see **Using the trust command** on page 40. If the virtual disk group is quarantined or offline, collect array logs and contact HPE Support.
RAID 6	A disk drive is marked as **LEFTOVR** or goes missing making the Vdisk or disk group status as **FTDN**. Another disk drive from the same Vdisk or disk group is marked as **LEFTOVR** or goes missing. The status of the Vdisk or disk group is changed to **CRIT**. One more disk drive from the Vdisk or disk group is marked as **LEFTOVR** or goes missing, the Vdisk or disk group is marked as **QTOF**.	
RAID 10/50	A disk drive is marked as **LEFTOVR** or goes missing, making the Vdisk or disk group status **CRIT**. Another disk drive from the same sub-Vdisk or sub-disk group fails or is missing the Vdisk or disk group is marked as **QTOF**.	

NOTE:

- In the previous scenarios, if the disk drives fail, goes missing, or both during a boot and also upon a rescan, the status of the Vdisk or disk group is marked as **QTOF**.
- In the previous scenarios, the drives that went down or missing are not spares or reconstruction targets.

Disk drive failure during reconstruct

A compatible spare has a capacity equal to or greater than the smallest disk in the Vdisk or disk group and is the same disk type (SAS or SATA).

Table 10: RAID 6 reconstruction scenarios

Condition	Action required
No compatible spares are available, reconstruction does not start automatically	To start reconstruction manually, replace each failed disk, and then do one of the following: • Add each new disk as either a dedicated spare or a global spare. Remember that a different critical Vdisk or disk group might take the global spare than the one you intended. When a global spare replaces a disk in a Vdisk or disk group, the icon for the global spare in the enclosure view changes. The icon for the global spare will match the other disks in that Vdisk or disk group. • Enable the **Dynamic Spare Capability** option to use the new disks without designating them as spares.
One disk fails and a compatible spare is available	The system uses the spare to reconstruct the Vdisk or disk group. If a second disk fails during reconstruction, reconstruction continues until it is complete, regardless of whether a second spare is available. If the spare fails during reconstruction, reconstruction stops.
Two disks fail and only one compatible spare is available	The system may wait up to five minutes for a second spare to become available. After five minutes, the system uses the spare to reconstruct one disk in the Vdisk or disk group (referred to as **fail 2, fix 1 mode). If the spare fails during reconstruction, reconstruction stops.**
Two disks fail and two compatible spares are available	The system uses both spares to reconstruct the Vdisk or disk group. If one of the spares fails during reconstruction, reconstruction proceeds in **fail 2, fix 1** mode. If the second spare fails during reconstruction, reconstruction stops.

NOTE: Depending on the Vdisk or 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 Vdisk or disk group. Deleting a Vdisk or disk group will cause permanent data loss.

Vdisk or disk group member unavailable

RAID Level	Symptom	Action required
NRAID	Vdisk or disk group status is **OFFL** (Offline).	A single drive failure causes the Vdisk or disk group to go **OFFL** and `Trust` cannot recover the data. Data is restored from the backup.
RAID 0	Vdisk or disk group status is **OFFL** (Offline).	
RAID 1	Vdisk or disk group status is **CRIT** (Critical).	If a spare is already available (global or dedicated to the Vdisk or disk group), it is used for the Vdisk or disk group in **CRIT** or **FTDN** state and reconstruct begins. Replace the failed disk drive if necessary. If a spare is not available, consider the option to replace the failed disk drive or clear meta-data on `LEFTOVR` drive and add it as a spare.
RAID 5		
RAID 10		
RAID 50		
RAID 6	Vdisk or disk group status is **FTDN** (Degraded).	Do not use a leftover drive that failed due to drive errors such as media errors, smart trips, or hardware faults to use as a spare to reconstruct a Vdisk or disk group.

Two or multiple disk drive failures

RAID Level	Condition	Description/notes	Action required
RAID 1	Two disk drive failures in RAID 1 make it **QTOF** or **OFFL**.	RAID 1 supports a maximum of two disk drives.	If the Vdisk or disk group goes **QTOF**, the Vdisk or disk group gets de-quarantined automatically after the drives are recognized. Review logs and drives to determine if further action is required. • If the Vdisk status changes to **OFFL**, see **Using the trust command**. • If the virtual disk group is quarantined or offline, collect array logs and contact HPE Technical Support.
RAID 5	Failure of two or more disk drives in a RAID 5 Vdisk or disk group makes the Vdisk or disk group **QTOF** or **OFFL**.	NA	
RAID 6	Failure of two disk drives in a RAID 6 Vdisk or disk group makes the Vdisk or disk group **CRIT**.	If multiple spares are available for the Vdisk or disk group, two disk drive reconstruction begins.	
	Failure of multiple disk drives (greater than two) in a RAID 6 Vdisk or disk group makes the Vdisk or disk group **QTOF** or **OFFL**.	NA	

Table Continued

RAID Level	Condition	Description/notes	Action required
RAID 10	Failure of both the disk drives in the same sub-Vdisk or sub-disk group makes the Vdisk or disk group **QTOF**.	NA	
RAID 50	Failure of two or more drives, in the same sub-Vdisk or sub-disk group makes the Vdisk or disk group **QTOF**.	NA	
RAID 10, RAID 50	Failure of two or more drives, all in different sub-Vdisk or sub-disk group.	In a RAID sub-Vdisk or sub-disk group, if one drive fails, the sub-Vdisk or sub-disk group status, and therefore the Vdisk or disk group, is CRIT. If two or more sub-Vdisk or sub-disk groups have only one drive failed, the status of the Vdisk or disk group is still CRIT. Available spares will be used for reconstruction. If two or more disks fail in the same sub-Vdisk or sub-disk group, contact HPE Technical Support.	

Vdisk or disk group alerts

Symptom	Action required
Vdisk or disk group is critical	• For RAID types other than RAID 6, see: ◦ **Vdisk or disk group member unavailable** ◦ **Two or multiple disk drive failures** • For RAID 6, see **Two or multiple disk drive failures**.
Vdisk or disk group is quarantined	See **Vdisk or disk group quarantined**.
Vdisk or disk group is offline	See **Using the trust command**.
Vdisk or disk group reconstruction failed	See **Disk drive failure during reconstruct**.

Vdisk expansion frequently asked questions (FAQs)

Question	Answer
During a Vdisk expansion, multiple disk drives go missing while the Vdisk is quarantined. Will the expansion continue once the disk drives return and the Vdisk de-quarantines?	Yes, expansion resumes once the Vdisk is automatically de-quarantined. **NOTE:** Avoid using the manual de-quarantine option as it takes the Vdisk **OFFL**, which causes data loss.
During the expansion of a Vdisk (other than RAID 6), if a drive fails making the Vdisk CRIT and a spare is available can a reconstruction happen?	No, the spare is not used and the reconstruction does not begin immediately after the drive fails. A backup is recommended.
During RAID 6 Vdisk expansion if two drives fail, can it do two drive reconstruction?	No, the spares are not used and the reconstruction does not begin immediately after the drive fails. A backup is recommended.
Is it possible to create a backup of a Vdisk while an expansion is underway?	Yes, backup of a Vdisk can be performed during Vdisk expansion. **NOTE:** Best practice is to perform a backup of the data before initiating the expansion of a Vdisk.
During a Vdisk expansion, if the Vdisk goes **OFFL**, can we use TRUST to bring the Vdisk online? If so will the expansion then proceed?	Trust cannot be performed if a Vdisk goes **OFFL** during Vdisk expansion. Expansion is aborted and any data on the Vdisk is irrevocably lost. **NOTE:** Best practice is to perform a backup of the data before initiating the expansion of a Vdisk.
During a Vdisk expansion, if the controller owning the Vdisk crashes, will the expansion continue after the Vdisk is failed over to the partner controller?	Yes, Vdisk expansion continues after the Vdisk is failed over to the partner controller.
Can we create a secondary volume and create snapshots of a volume that is a member of a Vdisk under expansion? What about other activities like replication?	Yes, you can create a secondary volume and snapshots are created for a Vdisk under expansion. All the other activities work fine.
Can the volumes of an expanding Vdisk be remapped to different hosts?	Yes, the volumes can be remapped.

Management controller issue

Restarting the management controller

The Management controller (MC) is the array component that is responsible for functions related to the management of the array. The Storage controller (SC) is responsible for handling Host IO to the array. Both the MC and SC reside within the same physical controller module.

Some Management controller functions are:

- Managing array user accounts
- Managing array logins
- Managing array configuration information

Table 11: Management controller scenarios

Symptom	Action required
SC component is functioning properly. All Host I/O are being handled correctly while the MC component is not accepting a user login	• Restart MC from the other controller, if possible. If that does not resolve the issue, restart the corresponding SC from the other controller, then restart the MC from the other controller. • If that does not resolve the issue, restart the array.

The Management controller port is unresponsive.	Log in to the MC on the other SC and restart the unresponsive MC.
	For example, an array with the following setup:
	Controller A—IP 15.5.224.9
	Controller B—IP 15.5.224.10
	Verify the network connectivity by issuing a ping to the MC IP address of the controller. Ensure that an issue does not exist with the intranet on which the MSA is installed.
	`# ping 15.5.224.9 Reply from 15.5.224.9 bytes=32 time60`
	If the `ping` command does not return a valid response, investigate the network or cabling.
	If **controller A** responds to the ping but is not enabling logins through the CLI or SMU interface, log in to **controller B** through the CLI. Attempt to restart the MC on **controller A**.
	The following example shows a restart of a controller through the CLI interface:
	`#restart mc A`
	`Continue? yes <enter return>`
	`Success: MC A restarted.`
	There may be up to a minute delay between the **Continue** prompt and hitting the **Yes <enter return>** to when the Success information is displayed. Wait a few moments after the Success information is displayed for the restart process to complete and then try to log in to **controller A**.
	If the MC remains unresponsive, shutdown and restart **controller A** using the opposite **controller B**, then recheck the Management controller.
	NOTE: During restart, you will briefly lose communication with the specified management controllers:
	From **controller B**:
	`#shutdown A`
	Wait for **A** to shut down
	`#restart sc a`
	This will restart the MC on **controller A** as well.

Symptom	Action required
	NOTE: In a single controller environment, if the CLI command is not successful, quiesce all host I/Os before restarting the SC of the controller.

Instances when the management controller becomes unresponsive

Table 12: Instances when the management controller becomes unresponsive

Symptom	Cause	Action required

The following messages regularly appear in the array logs:
• 237 INFORMATIONAL Firmware update progress: The firmware was verified.
• 237 INFORMATIONAL Firmware update progress: The SC loader was updated.
Saved in primary location in flash. The firmware is the same so it was not flashed.
• 237 INFORMATIONAL Firmware update progress: The SC app was updated.
Saved in primary location in flash. The firmware is different so it was flashed; flashed successfully.
• 237 INFORMATIONAL Firmware update progress: The memory-controller FPGA was updated. Saved in primary location in flash. The firmware is the same so it was not flashed.
	Firmware Flash is blocked, upgrade retries are evident.	For systems that are in a firmware upgrade loop:
		1. Verify all host I/O to the array is quiesced or halted completely.
		Restart both Storage Controllers (SC) using either the SMU or CLI.
		2. Complete one of the following options:
		• Run the Online ROM Flash Component again to upgrade both the controllers.
		• Use the FTP option to upgrade both controllers:
		`ftp> put firmware_file flash`
		• Use the SFTP option to upgrade both controllers:
		`sftp> put firmware_file flash`
		3. Allow firmware upgrade to complete fully.
		4. Using either **SMU** or the **CLI**, check, and confirm all firmware versions for the Management Controller and Storage Controller matches between both controllers

The following sequence of messages appear in the array logs: • 152 INFORMATIONAL The Storage Controller is not	Management Controller and Storage Controller are unable to communicate.	For systems that have a communication error between the MC and SC, perform any one of the following steps:

Table Continued

Symptom	Cause	Action required

receiving data from the Management Controller. (This message is normal during firmware update.)
• 153 INFORMATIONAL The Storage Controller resumed communications with the Management Controller.
• 152 INFORMATIONAL The Storage Controller is not receiving data from the Management Controller.
• 156 WARNING The Management Controller was restarted by the Storage Controller.
• 139 INFORMATIONAL The Management Controller booted up. MC firmware version: XXXXXXXX (baselevel: XXXX)
• 153 INFORMATIONAL The Storage Controller resumed communications with the Management Controller.
• 152 INFORMATIONAL The Storage Controller is not receiving data from the Management Controller. (This message is normal during firmware update.)
• 153 INFORMATIONAL The Storage Controller resumed communications with the Management Controller
		• Run the following CLI command:
		`#restart mc A`
		`Continue? yes <enter return>`
		`Success: MC A restarted.`
		• Restart the MC using SMU.
		NOTE:
		• During the restart process, you will briefly lose communication with the specified management controllers
		• In some cases, if the restart function previously noted does not clear the MC hang issue, you might attempt to reseat the offending controller. Before reseating the offending controller, you must perform a couple of safety measures:
		1. Verify all host I/O to the array is quiesced or halted completely.
		2. Shut down the controller, for example, `shutdown a|b|both`
		Alternatively, shutdown the controller using the SMU. Shutting down both the controllers is an offline activity.
		3. Reseating a controller with active I/O in a dual controller array will cause I/O failover to the other controller. In a single controller array, this reseating should be

Table Continued

Symptom	Cause	Action required
		scheduled during a maintenance window.
The configuration information is lost, array management, event messaging, and logging cease to function, but the host I/O continues to operate normally	In both single and dual controller environments, when the 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.	In a dual controller scenario, restart the MC of the affected controller from the other controller using either SMU or CLI. In a single controller environment, you will need to halt all host IO and reseat the controller. In a dual controller environment if the configuration between the two controllers does not match, verify Partner Firmware Upgrade (PFU) is enabled. If after enabling PFU the configuration does not match, verify in logs if the array has had a Compact Flash (CF) card failure. If the CF Card has failed, it will need to be replaced. In a single controller array, if there is a CF Card failure the controller will enter write-through mode which may impact performance. Schedule a maintenance window to replace the CF Card.

User login issues

Table 13: User login issues

Symptom	Recommended action
Unable login to the management interfaces Lost password	Verify the login credentials by logging into the USB CLI console. Using the USB CLI cable bypasses possible network issues. The default user names and passwords are: • manage / !manage • monitor / !monitor If no user logins can be recovered, HPE Technical support may be required.
User login can access one management interface but not others	Check that the user has the correct Interfaces selected in User Management. Verify that the management protocol is enabled (HTTPS, HTTP, SSH, TELNET). 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.
User is unable to make configuration changes to the storage system.	Verify that the user has the `manage` role. Users with the monitor role are not able to make configuration changes.
User is unable to create/modify/delete users.	Verify that the user has the `manage` role.
User is unable to load firmware.	Verify that the user has the `manage` role.
LDAP user is unable to log in	Verify that user credentials are correct by logging to the Active Directory from another system. Verify that other LDAP users are able to log in to the storage system. Verify that the LDAP parameters are set correctly. Verify that the LDAP User-Group name exactly matches the group name in the Active Directory. Verify that the LDAP user is not a member of more than 100 Active Directory groups.
LDAP user gets inconsistent permissions when logging into storage system.	Verify that the user is not a member or more than one User-Group on the storage system.

Power supply faults and power cycle

Power cycle

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 is an **Enclosure ID** LED. This LED will aid in identifying the Controller shelf. The Controller Shelf will report as "1" and any subsequent enclosures will have higher numbered IDs.

If a power cycle is needed, follow these steps:

1. Shut down, dismount, or unmap the hosts (servers) with access to the MSA volumes. This allows any pending writes in Application cache or Server memory to flush to the MSA controllers write-back cache.
2. Shut down both array controllers in a dual controller system or a single controller system by using SMU or CLI interface. Shutting down the arrays allows the controllers to flush the controller write cache to the disks.

   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 single JBODs may adversely affect Vdisks or disk groups that are shared between enclosures.
3. JBOD Enclosures only need to be powered off for routine maintenance. The JBODs must be powered off only after the controllers are properly shut down using the CLI or SMU. To power off the JBOD enclosures either remove the power cables or shut off the switch, whichever step is applicable for your power supplies.

Perform the reverse steps when restoring power to the Storage infrastructure. Power on the JBOD enclosures from the bottom to the top. Allow the drives in each JBOD to power up before powering the next JBOD enclosure. If the enclosures are powered off, wait a minimum of one minute after powering on the last JBOD enclosure before powering up the MSA controllers. Waiting allows the Hard Drives enough time to spin up before the array controllers come online.

NOTE: Power on the JBOD enclosures from the bottom enclosure to the top, allowing enough time for the drives to spin up. Larger capacity drives such as 6TB and 8TB drives may take additional time.

⚠ **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.

Power supply faults and recommended actions

Table 14: Power supply faults and recommended actions

Symptom	Recommended action
Power supply warning or failure. Related Event code 551.	Verify the following: • All the power supply units are working. • No slots are left open for more than two minutes. If you must replace a module, leave the old module in place until you have the replacement. Leaving a slot open negatively affects the airflow and might cause the unit to overhead. • The controller modules are properly seated in their slots and that their latches are locked. • Replace the power supply module.
Power supply module status is listed as failed or you receive a voltage event notification. Related Event code 551.	• If the power supply module has a switch, ensure that the switch is turned on. • Ensure that the power cables are firmly plugged into both power supply and into an appropriate functional electrical outlet. • Replace the power supply module.
Power LED is off.	• If the power supply module has a switch, ensure that the switch is turned on. • Ensure that the power cables are firmly plugged into both power supply and into an appropriate electrical outlet. • Replace the power supply module.
Voltage/Fan Fault/Service Required LED is on.	Replace the power supply module.

For information on replacing the DC Power and Cooling Module, see *HPE MSA Power and Cooling Module Replacement Instructions* on the **Hewlett Packard Enterprise Support Center** website.

Chassis replacement

Table 15: Chassis replacement scenarios

Symptom	Action required
• 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 does not allow a FRU to be installed.	Where a problem exists with the physical chassis itself, replacement is necessary. 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.
An issue relating to the midplane	To determine if the midplane is the problem, collect and investigate the array logs. Also, look for a 314 midplane FRU notification in array events and a corresponding 247 related event. If event 274, 358, 495, 521, or 602 occurs, check the Recommended actions for further conditions to see if the midplane or chassis should be replaced.

IMPORTANT: Verify if any licenses will require to be obtained and installed on the new chassis. Replacements of existing licenses are provided at no extra cost by HPE when executing a chassis replacement.

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 action of the Home topic or the system's **Tools > Install License** page, or by executing the CLI `show license` command.

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

Obtain the new replacement licenses directly from the My License Portal (**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.

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.

NOTE:

- Avoid unnecessary chassis replacement. Other than a mechanical failure, it is rare to have a chassis or midplane issue requiring replacement.
- A FRU is any HPE orderable replacement part.

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

Using the `trust` command

The `trust` command resynchronizes time and date metadata on the drives, making **Leftover** or **Failed** drives that were once members of the Vdisk **active** members again. Use the `trust` command when a Vdisk is Offline and there is no data backup, or to recover the current data on a Vdisk. In these cases, the `trust` command works only if the drives continue to operate.

NOTE: The `trust` command should not be used on a virtual disk group. Contact HPE technical support if assistance is needed using the `trust` command.

After using the `trust` command:

1. After the trusted Vdisk is back online, back up the Vdisk and verify that the data is valid.
2. Delete the trusted Vdisk.
3. Replace any drive that was a member of the trusted Vdisk that **Failed** or went **Leftover** because of Unrecoverable Read Errors (UREs), SMART issues, or other errors.
4. Create a new Vdisk.
5. Restore data from a valid backup to the new Vdisk.

NOTE:

- The `trust` command is used as a last step in a disaster recovery situation. The `trust` command must be performed only by someone who knows how to use the `trust` command and has experience with Vdisk configurations and reviewing array logs.
- If you are not sure to take the corrective action, contact the Hewlett Packard Enterprise Support Center for further assistance.

CAUTION:

- The `trust` command can cause permanent data loss and unstable Vdisk operation.
- Use the `trust` command on a Vdisk as a disaster-recovery measure only; the Vdisk has no tolerance for any additional failures and must never be put back into a production environment.
- After the `trust` command is issued on a Vdisk, additional disaster recovery troubleshooting steps are limited.
- The `trust` command must be used only if the Vdisk has an Offline status; the `trust` command will not turn on a Vdisk with another status.
- Do not use the `trust` command when the storage system is:
  - Unstable
  - During power events
  - Events where enclosures or drives are added or removed unintentionally.
- Do not attempt to run the `trust` command on a Vdisk that is Quarantined critical, Quarantined offline, or Quarantined with down drives.
- When the Vdisk is Offline:
  - Never update the controller-module, expansion-module, or drive firmware.
  - Never clear unwritten cache data.
- You cannot use the `trust` command on a Vdisk that went Offline during Vdisk expansion.
- You cannot use the `trust` command on a Vdisk with a Critical status. Instead, add spares and let the system reconstruct the Vdisk.

Running the `trust` command

1. Disable the background scrub. For more information, see **CLI User Guide**.
2. Identify the cause for the Vdisk going **Quarantined offline** or **Offline**.
3. If an external issue, such as power failure or cable failures or disconnections, caused the Vdisk to go **Quarantined offline** or **Offline**, fix the external issue. Return the array to a stable state before continuing to the next step.
4. Determine the drives to be used with the `trust` command.
   - If the Vdisk went **Quarantined offline** or **Offline** during a reconstruction, determine if the drive that caused the Vdisk to go **Quarantined offline** or **Offline** went **Leftover** or **Failed**.
     - If the drive went **Leftover** and was not experiencing UREs, SMART issues, or other errors, unseat the spare drives that were brought into the Vdisk for reconstruction.
     - If the drive **Failed** or went **Leftover** due to UREs, SMART issues, or other errors (from now on known simply as the failed drive), determine whether to use the partially reconstructed target drive rather than the failed drive. For example, if the reconstruction is mostly done and the failed drive had numerous errors before failing, you may recover more data using the partially reconstructed drive.

Even though the data is not complete, because the failed drive may not stay up long enough to be useful. To determine how far the reconstruction has gone, consider the size of the Vdisk, the type of drives, and review the array logs to determine when the reconstruction started and the time a drive failure occurred.

Since using the `trust` command can irrevocably destroy data, or if you cannot determine the correct drives to use, or the status of the reconstruction, contact Hewlett Packard Enterprise Support Center for assistance.

- If the Vdisk went **Quarantined offline** or **Offline**, but not during a reconstruction, note the drives that caused the Vdisk to most recently go **Fault Tolerant with a down disk**, **Critical**, and **Quarantined offline** or **Offline**. Unseat the drives that caused the Vdisk to go **Fault Tolerant with a down disk** and **Critical**.

  - For a RAID-5 Vdisk, unseat the first **Failed** or **Leftover** drive that led to the Vdisk going **Quarantined offline** or **Offline**, according to the logs.
  - For a RAID 6 Vdisk, unseat the first two **Failed** or **Leftover** drives that led to the Vdisk going **Quarantined offline** or **Offline**, according to the logs.
  - For a RAID-50 Vdisk, if you know the sub-member set details, unseat the first **Failed** or **Leftover** drive that led to the RAID-5 sub-member set failing, according to the logs. If you do not have the sub-member set details, contact Hewlett Packard Enterprise Support Center for assistance. If you are uncertain about the order the drives **Failed** or went **Leftover**, or which drives were added into the Vdisk for reconstruction, contact Hewlett Packard Enterprise Support Center for further assistance.

IMPORTANT: If a drive **Failed** or went **Leftover** due to UREs, SMART issues or other errors, DO NOT attempt to clear the metadata on the drive and reuse the drive. **Failed** drives and drives that went **Leftover** due to UREs, Smart issues, or other errors are unstable and must be replaced.

5. Prevent all automatic reconstruction actions on a `trusted` Vdisk by following these steps:

   - Unseat the spare drives associated with the Vdisk. For more information, see CLI User Guide.
   - Unseat or delete global spares. For more information, see CLI User Guide.
   - Turn off the dynamic spares feature.

     Reconstruction causes heavy usage of drives that were already reporting errors. This usage can cause drives to fail during reconstruction, which can cause data to be unrecoverable.

6. Reseat the remaining affected drives.
7. Using the CLI, run the `enable trust` command.
8. Run the trust command on the Vdisk.

   If you get the following warning on the CLI while running the `Trust` command, contact the Hewlett Packard Enterprise Support Center for further assistance:

   ```
   Error: The trust operation failed because the disk group has an insufficient
   number of in-sync disks. - Please contact Support for further assistance.
   ```

   Or

   ```
   WARNING! Found out-of-sync disk(s). Using these disks for trust might cause
   data corruption.Do you want to include the out-of-sync disk(s)?• yes or y:
   ```

```
Allows the command to proceed, using these disks.• no or n: Allows the command
to proceed, without using these disks.• abort or a: Cancels the command.
```

Or

```
WARNING! Found partially reconstructed disk(s).Using these disks for trust
might cause data corruption.Do you want to include the partially reconstructed
disk(s)?• yes or y: Allows the command to proceed, using the out-of-sync
disks.• no or n: Allows the command to proceed, without using the out-of-sync
disks.• abort or a: Cancels the command.
```

After running the trust command

1. Perform a backup of the Vdisk and validate all backed-up data.
2. Delete the Vdisk.
3. Replace the **Failed** drives associated with this Vdisk, if necessary. If there are any leftover drives, verify that a full backup has been obtained before clearing the metadata on the drives. Do not clear the metadata if any Vdisk is in an offline state. If further help is needed contact HPE Support.
4. Recreate the Vdisk.
5. Restore the data from the most recent valid backup.
6. Re-enable automatic reconstruction by:
   - Reinserting or adding global spares that were unseated or deleted previously. For more information, see the *CLI User Guide*.
   - Re-enabling the spare drives associated with the Vdisk that were unseated previously.
   - Turning on the dynamic spares feature. For more information, see the *CLI User Guide*.
7. Re-enable background scrub. For more information, see the *CLI User Guide*.

Websites

MSA websites

MSA Manuals page:

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

MSA Firmware

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

MSA Drive Firmware

http://www.hpe.com/storage/MSADriveFirmware

MSA Forum

http://www.hpe.com/forum/MSA

HPE MSA Support Material Access

http://h20564.www2.hpe.com/hpsc/doc/public/display?docLocale=en_US&docId=emr_na-c05349541

General websites

Hewlett Packard Enterprise Information Library

www.hpe.com/info/EIL

Single Point of Connectivity Knowledge (SPOCK) Storage compatibility matrix

www.hpe.com/storage/spock

Storage white papers and analyst reports

www.hpe.com/storage/whitepapers

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/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 for your product or to view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products* reference document, go to the Enterprise Safety and Compliance website:

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

Additional warranty information

HPE ProLiant and x86 Servers and Options

www.hpe.com/support/ProLiantServers-Warranties

HPE Enterprise 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.