



Hewlett Packard
Enterprise

HPE MSA 2050 User Guide

Abstract

This document describes initial hardware setup for HPE MSA 2050 controller enclosures, and is intended for use by storage system administrators familiar with servers and computer networks, network administration, storage system installation and configuration, storage area network management, and relevant protocols.

Firmware Version: VL270

Part Number: Q1J79-62042

Published: February 2020

Edition: 1

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 U.S. trademarks of the Microsoft group of companies.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Java and Oracle are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

Revision History

723983-005	June 2017
Initial HPE release.	
Q1J79-62014	September 2017
Updated for MSA 2050 SAS.	
Q1J79-62026	September 2018
Updated release information.	
Q1J79-62042	February 2020
Updated power supply and user login/password information.	

Contents

1 Overview	9
MSA 2050 Storage models	9
MSA 2050 enclosure user interfaces	9
MSA 2050 SAN	9
Features and benefits	10
Product QuickSpecs	10
Related MSA documentation	10
2 Components	11
Front panel components	11
MSA 2050 Array SFF or supported 24-drive expansion enclosure	11
MSA 2050 Array LFF or supported 12-drive expansion enclosure	12
Disk drives used in MSA 2050 enclosures	12
Controller enclosure—rear panel layout	13
MSA 2050 SAN controller module—rear panel components	14
Drive enclosures	15
LFF and SFF drive enclosure — rear panel layout	15
Cache	15
Transportable CompactFlash	15
Supercapacitor pack	16
Upgrading to the MSA 2050	16
3 Installing the enclosures	17
Installation checklist	17
FDE considerations	17
Connecting controller and drive enclosures	18
Connecting the MSA 2050 controller to the LFF or SFF drive enclosure	19
Cable requirements for MSA 2050 enclosures	19
Testing enclosure connections	23
Powering on/powering off	23
AC power supply	24
DC power supply	25
4 Connecting hosts	27
Host system requirements	27
Connecting the enclosure to data hosts	27
MSA 2050 SAN	27
.....	29
Host connection configurations	29
Connecting direct attach configurations	30
Connecting switch attach configurations	31
Connecting remote management hosts	32
Connecting two storage systems to replicate volumes	32
Cabling for replication	33
Host ports and replication	33
Updating firmware	37

5	Connecting to the controller CLI port.....	38
	Device description	38
	Emulated serial port.....	38
	Preparing a Linux computer for cabling to the CLI port.....	38
	Preparing a Windows computer for cabling to the CLI port	39
	Obtaining IP values	39
	Setting network port IP addresses using DHCP	39
	Setting network port IP addresses using the CLI port and cable	39
	Using the CLI port and cable—known issues on Windows.....	43
	Problem	43
	Workaround.....	43
6	Basic operation.....	44
	Accessing the SMU	44
	Configuring and provisioning the storage system.....	44
7	Troubleshooting.....	45
	USB CLI port connection	45
	Fault isolation methodology.....	45
	Basic steps.....	45
	Options available for performing basic steps.....	45
	Performing basic steps	46
	If the enclosure does not initialize.....	47
	Correcting enclosure IDs.....	47
	Stopping I/O.....	47
	Diagnostic steps	48
	Is the enclosure front panel Fault/Service Required LED amber?.....	48
	Is the enclosure rear panel FRU OK LED off?.....	49
	Is the enclosure rear panel Fault/Service Required LED amber?.....	49
	Are both disk drive module LEDs off (Online/Activity and Fault/UID)?	49
	Is the disk drive module Fault/UID LED blinking amber?.....	49
	Is a connected host port Host Link Status LED off?	50
	Is a connected port Expansion Port Status LED off?.....	50
	Is a connected port Network Port Link Status LED off?.....	51
	Is the power supply Input Power Source LED off?.....	51
	Is the power supply Voltage/Current/Temperature/Fan Fault/Service Required LED amber?.....	51
	Controller failure.....	51
	If the controller has failed or does not start, is the Cache Status LED on/blinking?.....	52
	Transporting cache.....	52
	Isolating a host-side connection fault	52
	Host-side connection troubleshooting featuring host ports with SFPs	52
	Host-side connection troubleshooting featuring SAS host ports	54
	Isolating a controller module expansion port connection fault.....	55
	Isolating Remote Snap replication faults	55
	Replication setup and verification	56
	Diagnostic steps for replication setup	56
	Resolving voltage and temperature warnings	59
	Sensor locations.....	59
	Power supply sensors	59
	Cooling fan sensors.....	59
	Temperature sensors.....	60

Power supply module voltage sensors	60
8 Support and other resources	61
Accessing Hewlett Packard Enterprise Support	61
Information to collect	61
Accessing updates.....	61
Remote support	61
Remote support and Proactive Care information.....	62
Proactive Care customer information	62
Warranty information.....	62
Additional warranty information	62
Regulatory information	62
Additional regulatory information.....	63
Documentation feedback	63
A LED descriptions.....	64
Front panel LEDs.....	64
Enclosure bezel.....	64
MSA 2050 Array SFF or supported 24-drive expansion enclosure	65
MSA 2050 Array LFF or supported 12-drive expansion enclosure.....	66
Ear covers	66
Disk drive LEDs.....	67
Rear panel LEDs.....	68
Controller enclosure—rear panel layout	68
MSA 2050 LFF and SFF drive enclosures—rear panel layout	74
B Specifications and requirements.....	75
Safety requirements.....	75
Site requirements and guidelines.....	75
Site wiring and AC power requirements	75
Site wiring and DC power requirements	75
Weight and placement guidelines.....	76
Electrical guidelines.....	76
Ventilation requirements	76
Cabling requirements	76
Management host requirements.....	76
Physical requirements	77
Environmental requirements	78
Electrical requirements	78
Site wiring and power requirements.....	78
Power cord requirements.....	78
C Electrostatic discharge	79
Preventing electrostatic discharge	79
Grounding methods to prevent electrostatic discharge	79
D SFP option for host ports	80
Locate the SFP transceivers.....	80
Install an SFP transceiver	80
Verify component operation	81

Index 82

Figures

1	Bezel used with MSA 2050 enclosures: front panel	11
2	MSA 2050 Array SFF or supported 24-drive expansion enclosure: front panel	11
3	MSA 2050 Array LFF or supported 12-drive expansion enclosure: front panel	12
4	MSA 2050 Array: rear panel	13
5	MSA 2050 SAN controller module face plate (FC or 10GbE iSCSI)	14
6	MSA 2050 SAN controller module face plate (1 Gb RJ-45)	14
7	Supported drive enclosures: SFF/LFF rear panel	15
8	MSA 2050 CompactFlash memory card	16
9	Cabling connections between the MSA 2050 controller and a single drive enclosure	20
10	Cabling connections between MSA 2050 controllers and LFF and SFF drive enclosures	20
11	Fault-tolerant cabling connections showing maximum number of enclosures	22
12	AC power supply	24
13	DC power supply	25
14	DC power cable featuring sectioned D-shell and lug connectors	25
15	Connecting hosts: direct attach—one server/one HBA/dual path	30
16	Connecting hosts: direct attach—two servers/one HBA per server/dual path	30
17	Connecting hosts: direct attach—four servers/one HBA per server/dual path	31
18	Connecting hosts: switch attach—two servers/two switches	31
19	Connecting hosts: switch attach—four servers/multiple switches/SAN fabric	32
20	Connecting two storage systems for Remote Snap: multiple servers/one switch/one location	34
21	Connecting two storage systems for Remote Snap: multiple servers/switches/one location	35
22	Connecting two storage systems for Remote Snap: multiple servers/switches/two locations	35
23	Connecting two storage systems for Remote Snap: multiple servers/SAN fabric/two locations	36
24	Connecting a USB cable to the CLI port	40
25	Partial exploded view showing bezel alignment with 2U chassis	64
26	Detail views of enclosure ear cover mounting sleeves	64
27	LEDs: MSA 2050 Array SFF or supported 24-drive expansion enclosure: front panel	65
28	LEDs: MSA 2050 Array LFF or supported 12-drive expansion enclosure: front panel	66
29	Ear covers option to enclosure bezel	66
30	LEDs: Disk drive combinations — enclosure front panel	67
31	MSA 2050 SAN Array: rear panel	68
32	LEDs: MSA 2050 SAN controller module (FC and 10GbE SFPs)	69
33	LEDs: MSA 2050 SAN controller module (1 Gb RJ-45 SFPs)	70
34	LEDs: MSA 2050 SAS controller module (SFF-8644 12 Gb SAS connector)	71
35	LEDs: MSA 2050 Storage system enclosure power supply modules	73
36	LEDs: MSA 2050 3.5" 12-drive or 2.5" 24-drive enclosure rear panel	74
37	Install a qualified SFP option	80

Tables

1	Related MSA firmware documentation	10
2	Installation checklist	17
3	Supported terminal emulator applications	38
4	Terminal emulator display settings	38
5	Terminal emulator display settings	41
6	Terminal emulator connection settings	41
7	Diagnostics LED status: Front panel “Fault/Service Required”	48
8	Diagnostics LED status: Rear panel “FRU OK”	49
9	Diagnostics LED status: Rear panel “Fault/Service Required”	49
10	Diagnostics LED status: Front panel disks “Online/Activity” and “Fault/UID”	49
11	Diagnostics LED status: Front panel disks “Fault/UID”	49
12	Diagnostics LED status: Rear panel “Host Link Status”	50
13	Diagnostics LED status: Rear panel “Expansion Port Status”	50
14	Diagnostics LED status: Rear panel “Network Port Link Status”	51
15	Diagnostics LED status: Rear panel power supply “Input Power Source”	51
16	Diagnostics LED status: Rear panel power supply: “Voltage/Current/Temperature/Fan Fault/Service Required”	51
17	Diagnostics LED status: Rear panel “Cache Status”	52
18	Diagnostics for replication setup: Using Remote Snap feature	57
19	Diagnostics for replication setup: Creating a replication set	57
20	Diagnostics for replication setup: Replicating a volume	58
21	Diagnostics for replication setup: Checking for a successful replication	58
22	Power supply sensor descriptions	59
23	Cooling fan sensor descriptions	59
24	Controller platform temperature sensor descriptions	60
25	Power supply temperature sensor descriptions	60
26	Voltage sensor descriptions	60
27	Cache Status LED – power on behavior	72
28	Rackmount enclosure dimensions	77
29	Rackmount enclosure weights	77

1 Overview

HPE MSA Storage models are high-performance storage solutions combining outstanding performance with high reliability, availability, flexibility, and manageability. MSA 2050 enclosure models are designed to meet NEBS Level 3, MIL-STD-810G (storage requirements), and European Telco specifications.

MSA 2050 Storage models

The MSA 2050 enclosures support large form factor (LFF 12-disk) and small form factor (SFF 24-disk) 2U chassis, using either AC or DC power supplies. The MSA 2050 SAN controllers are introduced below.

NOTE: For additional information about MSA 2050 controller modules, see the following subsections:

- [“Controller enclosure—rear panel layout” \(page 68\)](#)
 - [“MSA 2050 SAN controller module—rear panel LEDs” \(page 69\)](#)
-

The MSA 2050 enclosures support virtual storage. For virtual storage, a group of disks with an assigned RAID level is called a *virtual disk group*. This guide uses the term *disk group* for brevity.

MSA 2050 enclosure user interfaces

The MSA 2050 enclosures support the Storage Management Utility (SMU), which is a web-based application for configuring, monitoring, and managing the storage system. Both the SMU and the command-line interface (CLI) are briefly described.

- The SMU is the primary web interface to manage virtual storage.
 - The CLI enables you to interact with the storage system using command syntax entered via the keyboard or scripting.
-


NOTE: For more information about the SMU, see the SMU Reference Guide or online help. For more information about the CLI, see the CLI Reference Guide. See also [“Related MSA documentation” \(page 10\)](#).

MSA 2050 SAN

MSA 2050 SAN controllers use Converged Network Controller technology, allowing you to select the desired host interface protocol from the available Fibre Channel (FC) or Internet SCSI (iSCSI) host interface protocols supported by the system. You can set all controller module host ports to use either FC or iSCSI protocol using the `set host-port-mode` CLI command. MSA 2050 SAN controllers support the following link speeds:

- 16 Gb FC
- 8 Gb FC
- 10 GbE iSCSI
- 1 GbE iSCSI

Alternatively, you can use the CLI to set Converged Network Controller ports to support a combination of host interface protocols. When configuring a combination of host interface protocols, host ports 1 and 2 are set to FC (either both 16 Gb/s or both 8 Gb/s), and host ports 3 and 4 must be set to iSCSI (either both 10 GbE or both 1 GbE), provided the Converged Network Controller ports use the qualified SFP connectors and cables required for supporting the selected host interface protocol. See [“MSA 2050 SAN controller module—rear panel LEDs” \(page 69\)](#) for more information.

 **TIP:** See the topic about configuring host ports within the SMU Reference Guide for information about configuring Converged Network Controller ports with host interface protocols of the same type or a combination of types.

Features and benefits

Product features and supported options are subject to change. Online documentation describes the latest product and product family characteristics, including currently supported features, options, technical specifications, configuration data, related optional software, and product warranty information.

Product QuickSpecs

Check the QuickSpecs for a complete list of supported servers, operating systems, disk drives, and options. See <https://www.hpe.com/support/MSA2050QuickSpecs>. If a website location has changed, an Internet search for “HPE MSA 2050 quickspecs” will provide a link.

Related MSA documentation

Related support information is provided in the “[Support and other resources](#)” chapter. Firmware-related MSA documentation titles directly pertaining to this guide are provided in the table below.

Table 1 Related MSA firmware documentation

For information about	See
Using the Storage Management Utility (SMU) web interface to configure and manage the product	HPE MSA 1050/2050 SMU Reference Guide
Using the command-line interface (CLI) to configure and manage the product	HPE MSA 1050/2050 CLI Reference Guide
Event codes and recommended actions	HPE MSA Event Descriptions Reference Guide

To access the above MSA documentation, see the Hewlett Packard Enterprise Information Library:

<https://www.hpe.com/support/msa2050>

NOTE: The table above provides complete titles of MSA firmware documents used with this guide. Within this guide, references to the documents listed are abbreviated as follows:

- SMU Reference Guide
 - CLI Reference Guide
 - Event Descriptions Reference Guide
-

2 Components

Front panel components

HPE MSA 2050 models support small form factor (SFF) and large form factor (LFF) enclosures. The SFF chassis, configured with 24 2.5" SFF disks, and the LFF chassis, configured with 12 3.5" LFF disks, are used as either controller enclosures or drive enclosures.

Supported drive enclosures, used for adding storage, are available in LFF or SFF chassis. The MSA 2050 LFF Disk Enclosure is the large form factor drive enclosure and the MSA 2050 SFF Disk Enclosure is the small form factor drive enclosure used for storage expansion.

HPE MSA 2050 models use either an enclosure bezel or traditional ear covers. The 2U bezel assembly is comprised of left and right ear covers connected to the bezel body subassembly. A sample bezel is shown below.





Figure 1 Bezel used with MSA 2050 enclosures: front panel

The front panel illustrations that follow show the enclosures with the bezel removed, revealing ear flanges and disk drive modules. Two sleeves protruding from the backside of each ear cover component of the bezel assembly push-fit onto the two ball studs shown on each ear flange to secure the bezel. Remove the bezel to access the front panel components.

TIP: See “Enclosure bezel” (page 64) for bezel attachment and removal instructions, and pictorial views.

MSA 2050 Array SFF or supported 24-drive expansion enclosure




Notes:

Integers on disks indicate drive slot numbering sequence.
 The enlarged detail view at right shows LED icons from the bezel that correspond to the chassis LEDs.
 The detail view locator circle (above right) identifies the ear kit that connects to LED light pipes in the bezel (or ear cover).

- | | |
|----------------------------------|---------------------------------|
| 1 Enclosure ID LED | 4 Unit Identification (UID) LED |
| 2 Disk drive Online/Activity LED | 5 Heartbeat LED |
| 3 Disk drive Fault/UID LED | 6 Fault ID LED |

Figure 2 MSA 2050 Array SFF or supported 24-drive expansion enclosure: front panel

MSA 2050 Array LFF or supported 12-drive expansion enclosure



Notes:

Integers on disks indicate drive slot numbering sequence.

The enlarged detail view at right shows LED icons from the bezel that correspond to the chassis LEDs.

The detail view locator circle (above right) identifies the ear kit that connects to LED light pipes in the bezel (or ear cover).

1 Enclosure ID LED

2 Disk drive Online/Activity LED

3 Disk drive Fault/UID LED

4 Unit Identification (UID) LED

5 Heartbeat LED

6 Fault ID LED

Bezel icons for LEDs

Figure 3 MSA 2050 Array LFF or supported 12-drive expansion enclosure: front panel

NOTE: Either the bezel or the ear covers should be attached to the enclosure front panel to protect ear circuitry.

You can attach either the enclosure bezel or traditional ear covers to the enclosure front panel to protect the ears, and provide label identification for the chassis LEDs. The bezel and the ear covers use the same attachment mechanism, consisting of mounting sleeves on the cover back face:

- The enclosure bezel is introduced in [Figure 1 \(page 11\)](#).
- The ear covers are introduced in [Figure 26 \(page 64\)](#).
- The ball studs to which the bezel or ear covers attach are labeled in [Figure 2 \(page 11\)](#) and [Figure 3 \(page 12\)](#).
- Enclosure bezel alignment for attachment to the enclosure front panel is shown in [Figure 25 \(page 64\)](#).
- The sleeves that push-fit onto the ball studs to secure the bezel or ear covers are shown in [Figure 26 \(page 64\)](#).

Disk drives used in MSA 2050 enclosures

MSA 2050 enclosures support LFF/SFF Midline SAS, LFF/SFF Enterprise SAS, and LFF/SFF SSD disks. They also support LFF/SFF Midline SAS and LFF/SFF Enterprise self-encrypting disks that work with the Full Disk Encryption (FDE) feature. For information about creating disk groups and adding spares using these different disk drive types, see the SMU Reference Guide.

NOTE: In addition to the front views of SFF and LFF disk modules shown in the figures above, see [Figure 30 \(page 67\)](#) for pictorial views.

Controller enclosure—rear panel layout

The diagram and table below display and identify important component items comprising the rear panel layout of the MSA 2050 controller enclosure (MSA 2050 SAN configured with SFPs is shown in the example).




Figure 4 MSA 2050 Array: rear panel

A controller enclosure accommodates two power supply FRUs of the same type—either both AC or both DC—within the two power supply slots (see two instances of callout 1 above). The controller enclosure accommodates two controller module FRUs of the same type within the I/O module slots (see callouts 2 and 3 above).

-
- ⓘ **IMPORTANT:** MSA 2050 controller enclosures support dual-controller only. If a partner controller fails, the array will fail over and run on a single controller until the redundancy is restored. A controller module must be installed in each IOM slot to ensure sufficient airflow through the enclosure during operation.
-

The diagrams with tables that immediately follow provide descriptions of the different controller modules and power supply modules that can be installed into the rear panel of an MSA 2050 controller enclosure. Showing controller modules and power supply modules separately from the enclosure provides improved clarity in identifying the component items called out in the diagrams and described in the tables. Descriptions are also provided for optional drive enclosures supported by MSA 2050 controller enclosures for expanding storage capacity.

NOTE: MSA 2050 controller enclosures support hot-plug replacement of redundant controller modules, fans, power supplies, disk drives, and I/O modules. Hot-add of drive enclosures is also supported.

MSA 2050 SAN controller module—rear panel components

Figure 5 shows host ports configured with either 8/16 Gb FC or 10GbE iSCSI SFPs. The SFPs look identical. Refer to the LEDs that apply to the specific configuration of your Converged Network Controller ports.




Figure 5 MSA 2050 SAN controller module face plate (FC or 10GbE iSCSI)

Figure 6 shows Converged Network Controller ports configured with 1 Gb RJ-45 SFPs.

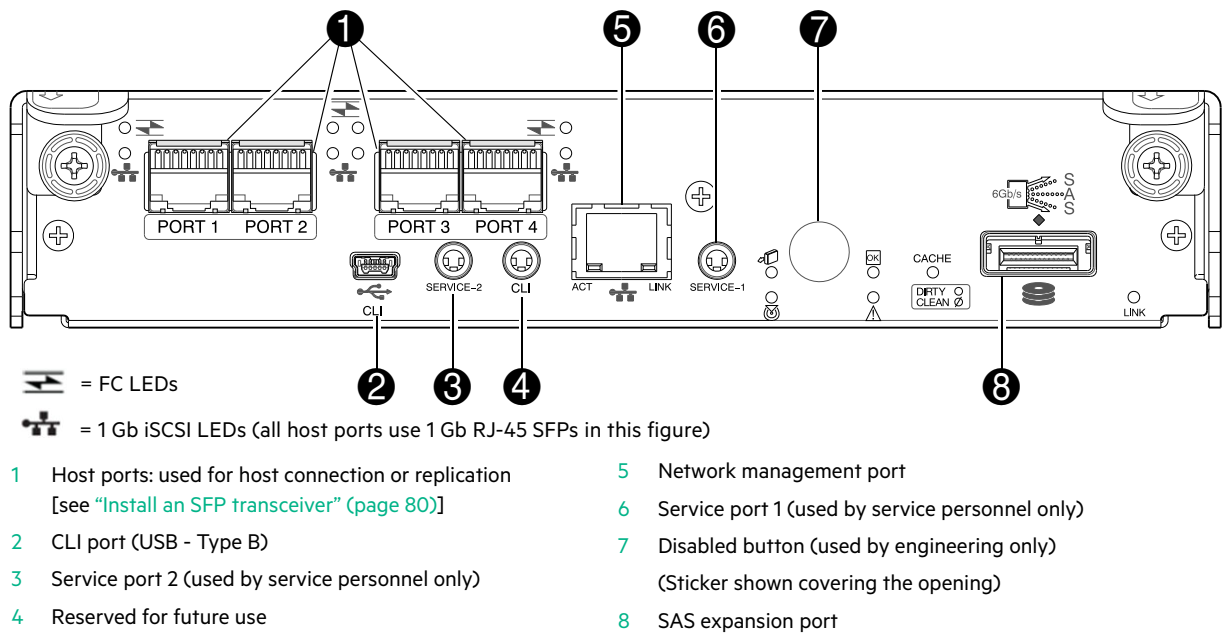


Figure 6 MSA 2050 SAN controller module face plate (1 Gb RJ-45)

NOTE: For more information about host port configuration, see the topic about configuring host ports within the SMU Reference Guide or online help.


⚠ **IMPORTANT:** See [Connecting to the controller CLI port](#) for information about enabling the controller enclosure USB Type - B CLI port for accessing the CLI to perform initial configuration tasks.

Drive enclosures

Drive enclosure expansion modules attach to MSA 2050 controller modules via the mini-SAS expansion port, allowing addition of disk drives to the system. MSA 2050 controller enclosures support adding the 6 Gb drive enclosures described below.

LFF and SFF drive enclosure — rear panel layout

MSA 2050 controllers support the MSA 2050 LFF Disk Enclosure and the MSA 2050 SFF Disk Enclosure, which share the same rear panel layout, as shown below.



- | | | | |
|---|--|---|---|
| 1 | Power supplies (AC shown) | 5 | Service port (used by service personnel only) |
| 2 | I/O module A | 6 | SAS In port |
| 3 | I/O module B | 7 | SAS Out port |
| 4 | Disabled button (used by engineering only) | | |

Figure 7 Supported drive enclosures: SFF/LFF rear panel

Cache

To enable faster data access from disk storage, the following types of caching are performed:

- Write-back caching. The controller writes user data in the cache memory on the module rather than directly to the drives. Later, when the storage system is either idle or aging—and continuing to receive new I/O data—the controller writes the data to the drive array.
- Read-ahead caching. The controller detects sequential array access, reads ahead into the next sequence of data, and stores the data in the read-ahead cache. Then, if the next read access is for cached data, the controller immediately loads the data into the system memory, avoiding the latency of a disk access.

NOTE: See the SMU Reference Guide for more information about volume cache options.

Transportable CompactFlash

During a power loss or array controller failure, data stored in cache is saved off to non-volatile memory (CompactFlash). The data is then written to disk after the issue is corrected. To protect against writing incomplete data to disk, the image stored on the CompactFlash is verified before committing to disk.

The CompactFlash memory card is located at the midplane-facing end of the controller module as shown below.

Controller module pictorial
(Midplane-facing rear view)




Figure 8 MSA 2050 CompactFlash memory card

If one controller fails, then later another controller fails or does not start, and the Cache Status LED is on or blinking, the CompactFlash will need to be transported to a replacement controller to recover data not flushed to disk (see [“Controller failure”](#) (page 51) for more information).

⚠ CAUTION: The CompactFlash memory card should only be removed for transportable purposes. To preserve the existing data stored in the CompactFlash, you must transport the CompactFlash from the failed controller to the replacement controller using a procedure outlined in the *HPE MSA Controller Module Replacement Instructions* shipped with the replacement controller module. Failure to use this procedure will result in the loss of data stored in the cache module.

📌 IMPORTANT: In dual controller configurations featuring one healthy partner controller, there is no need to transport failed controller cache to a replacement controller because the cache is duplicated between the controllers, provided that volume cache is set to standard on all volumes in the pool owned by the failed controller.

Supercapacitor pack

To protect RAID controller cache in case of power failure, MSA 2050 controllers are equipped with supercapacitor technology, in conjunction with CompactFlash memory, built into each controller module to provide extended cache memory backup time. The supercapacitor pack provides energy for backing up unwritten data in the write cache to the CompactFlash in the event of a power failure. Unwritten data in CompactFlash memory is automatically committed to disk media when power is restored. While the cache is being maintained by the supercapacitor, the Cache Status LED flashes at a rate of 1/10 second on and 9/10 second off.

Upgrading to the MSA 2050

For information about upgrading components for use with MSA controllers, see *Upgrading to the HPE MSA 1050/2050/2052*.

3 Installing the enclosures

Installation checklist

The following table outlines the steps required to install the enclosures and initially configure the system. To ensure a successful installation, perform the tasks in the order they are presented.

Table 2 Installation checklist

Step	Task	Where to find procedure
1.	Install the controller enclosure and optional drive enclosures in the rack, and attach the bezel or ear caps.	See the racking instructions poster.
2.	Connect the controller enclosure and LFF/SFF drive enclosures.	See “Connecting controller and drive enclosures” (page 18) .
3.	Connect power cords.	See the quick start instructions.
4.	Test enclosure connections.	See “Testing enclosure connections” (page 23) .
5.	Install required host software.	See “Host system requirements” (page 27) .
6.	Connect data hosts.	See “Connecting the enclosure to data hosts” (page 27) . If using the optional Remote Snap feature, also see “Connecting two storage systems to replicate volumes” (page 32) .
7.	Connect remote management hosts.	See “Connecting remote management hosts” (page 32) .
8.	Obtain IP values and set management port IP properties on the controller enclosure.	See “Obtaining IP values” (page 39) . See “Connecting to the controller CLI port” (page 38) ; with Linux and Windows topics.
9.	Perform initial configuration tasks ¹ : <ul style="list-style-type: none">• Sign in to the web-based Storage Management Utility (SMU).• Initially configure and provision the storage system using the SMU.	Topics below correspond to bullets at left: See “Getting Started” in the HPE MSA 1050/2050 SMU Reference Guide . See “Configuring the System” and “Provisioning the System” topics (SMU Reference Guide or online help) .

¹The SMU is introduced in [“Accessing the SMU” \(page 44\)](#). See the SMU Reference Guide or online help for additional information.

FDE considerations

The Full Disk Encryption feature available via the management interfaces requires use of self-encrypting drives (SED) which are also referred to as FDE-capable disk drive modules. When installing FDE-capable disk drive modules, follow the same procedures for installing disks that do not support FDE.

The procedures for using the FDE feature, such as securing the system, viewing disk FDE status, and clearing and importing keys are performed using the SMU or CLI commands (see the SMU Reference Guide or CLI Reference Guide for more information).

NOTE: When moving FDE-capable disk drive modules for a disk group, stop I/O to any volumes in the disk group before removing the disk drive modules. Follow the [“Removing the failed drive”](#) and [“Installing the replacement drive”](#) procedures within the *HPE MSA Drive Module Replacement Instructions*. Import the keys for the disks so that the disk content becomes available.

While replacing or installing FDE-capable disk drive modules, consider the following:

- If you are installing FDE-capable disk drive modules that do not have keys into a secure system, the system will automatically secure the disks after installation. Your system will associate its existing key with the disks, and you can transparently use the newly-secured disks.
- If the FDE-capable disk drive modules originate from another secure system, and contain that system's key, the new disks will have the Secure, Locked status. The data will be unavailable until you enter the passphrase for the other system to import its key. Your system will then recognize the metadata of the disk groups and incorporate it. The disks will have the status of Secure, Unlocked and their contents will be available.
 - To view the FDE status of disks, use the SMU or the `show fde-state` CLI command.
 - To import a key and incorporate the foreign disks, use the SMU or the `set fde-import-key` CLI command.

NOTE: If the FDE-capable disks contain multiple keys, you will need to perform the key importing process for each key to make the content associated with each key become available.

- If you do not want to retain the disks' data, you can repurpose the disks. Repurposing disks deletes all disk data, including lock keys, and associates the current system's lock key with the disks.
To repurpose disks, use the SMU or the `set disk` CLI command.
- You need not secure your system to use FDE-capable disks. If you install all FDE-capable disks into a system that is not secure, they will function exactly like disks that do not support FDE. As such, the data they contain will not be encrypted. If you decide later that you want to secure the system, all of the disks must be FDE-capable.
- If you install a disk module that does not support FDE into a secure system, the disk will have the Unusable status and will be unavailable for use.

If you are re-installing your FDE-capable disk drive modules as part of the process to replace the chassis FRU, you must insert the original disks and re-enter their FDE passphrase.

① **IMPORTANT:** The Fault/UID disk LED displays amber under the following conditions. See also [Figure 30 \(page 67\)](#).

- If an FDE disk is inserted into the storage enclosure in a secured locked state.
The disk is unusable by the system, and must either be unlocked or repurposed.
 - If a non-FDE disk is installed into an FDE-secured storage system.
The disk is unusable by the system, and must either be replaced with an FDE disk, or FDE must be turned off.
-

Connecting controller and drive enclosures

MSA 2050 controller enclosures support up to eight enclosures (including the controller enclosure). You can cable drive enclosures of the same type or of mixed LFF/SFF model type.

The firmware supports both straight-through and fault-tolerant SAS cabling. Fault-tolerant cabling allows any drive enclosure to fail—or be removed—while maintaining access to other enclosures. Straight-through cabling does not provide the same level of fault-tolerance as fault-tolerant cabling, but does provide some performance benefits as well as ensuring that all disks are visible to the array. Fault tolerance and performance requirements determine whether to optimize the configuration for high availability or high performance when cabling. MSA 2050 controller enclosures support 12 Gb/s disk drives downshifted to 6 Gb/s. Each enclosure has an expansion port using 6 Gb/s SAS lanes. When connecting multiple drive enclosures, use fault-tolerant cabling to ensure the highest level of fault tolerance.

For example, the illustration on the left in [Figure 10 \(page 20\)](#) shows controller module 1A connected to expansion module 2A, with a chain of connections cascading down (blue). Controller module 1B is connected to the lower expansion module (5B) of the last drive enclosure, with connections moving in the opposite direction (green).

Connecting the MSA 2050 controller to the LFF or SFF drive enclosure

The MSA 2050 LFF Disk Enclosure and the MSA 2050 SFF Disk Enclosure can be attached to an MSA 2050 controller enclosure using supported mini-SAS to mini-SAS cables of 0.5 m (1.64') to 2 m (6.56') length [see [Figure 9 \(page 20\)](#)]. Each drive enclosure provides two 0.5 m (1.64') mini-SAS to mini-SAS cables. Longer cables may be desired or required, and can be purchased separately.

Cable requirements for MSA 2050 enclosures

ⓘ IMPORTANT:

- When installing SAS cables to expansion modules, use only supported mini-SAS x4 cables with SFF-8088 connectors supporting your 6 Gb application.
 - See the QuickSpecs for information about which cables are provided with your MSA 2050 products.
<https://www.hpe.com/support/MSA2050QuickSpecs>
(If a website location has changed, an Internet search for “HPE MSA 2050 quickspecs” will provide a link.)
 - The maximum expansion cable length allowed in any configuration is 2 m (6.56').
 - When adding more than two drive enclosures, you may need to purchase additional 1 m or 2 m cables, depending upon number of enclosures and cabling method used (see QuickSpecs for supported cables):
 - Spanning 3, 4, or 5 drive enclosures requires 1 m (3.28') cables.
 - Spanning 6 or 7 drive enclosures requires 2 m (6.56') cables.
 - See the QuickSpecs (link provided above) regarding information about cables supported for host connection:
 - Qualified Fibre Channel SFP and cable options
 - Qualified 10GbE iSCSI SFP and cable options or qualified 10GbE Direct Attach Copper (DAC) cables
 - Qualified 1 Gb RJ-45 SFP and cable options
 - Qualified HD mini-SAS cable options
-

For additional information concerning cabling of MSA 2050 controllers, visit:

<https://www.hpe.com/support/MSA2050QuickSpecs>

<https://www.hpe.com/support/MSA2050BestPractices>

Browse for the following reference documents:

- HPE MSA 2050 Cable Configuration Guide
 - HPE Remote Snap technical white paper: [MSA Remote Snap Software](#)
-

NOTE: For clarity, the schematic illustrations of controller and expansion modules shown in this section provide only relevant details such as expansion ports within the module face plate outline. For detailed illustrations showing all components, see “[Controller enclosure—rear panel layout](#)” (page 13).

- ⓘ IMPORTANT:** MSA 2050 controller enclosures support dual-controller only. If a partner controller fails, the array will fail over and run on a single controller until the redundancy is restored. A controller module must be installed in each IOM slot to ensure sufficient airflow through the enclosure during operation.
-

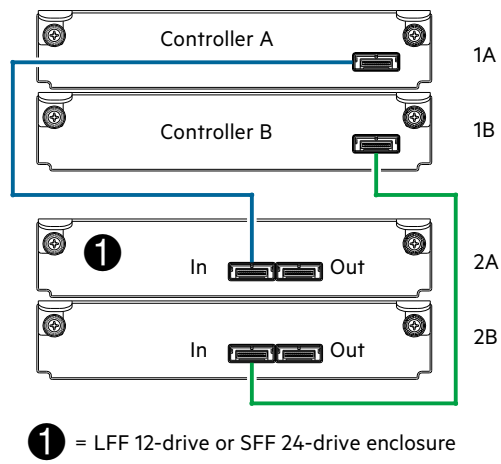


Figure 9 Cabling connections between the MSA 2050 controller and a single drive enclosure

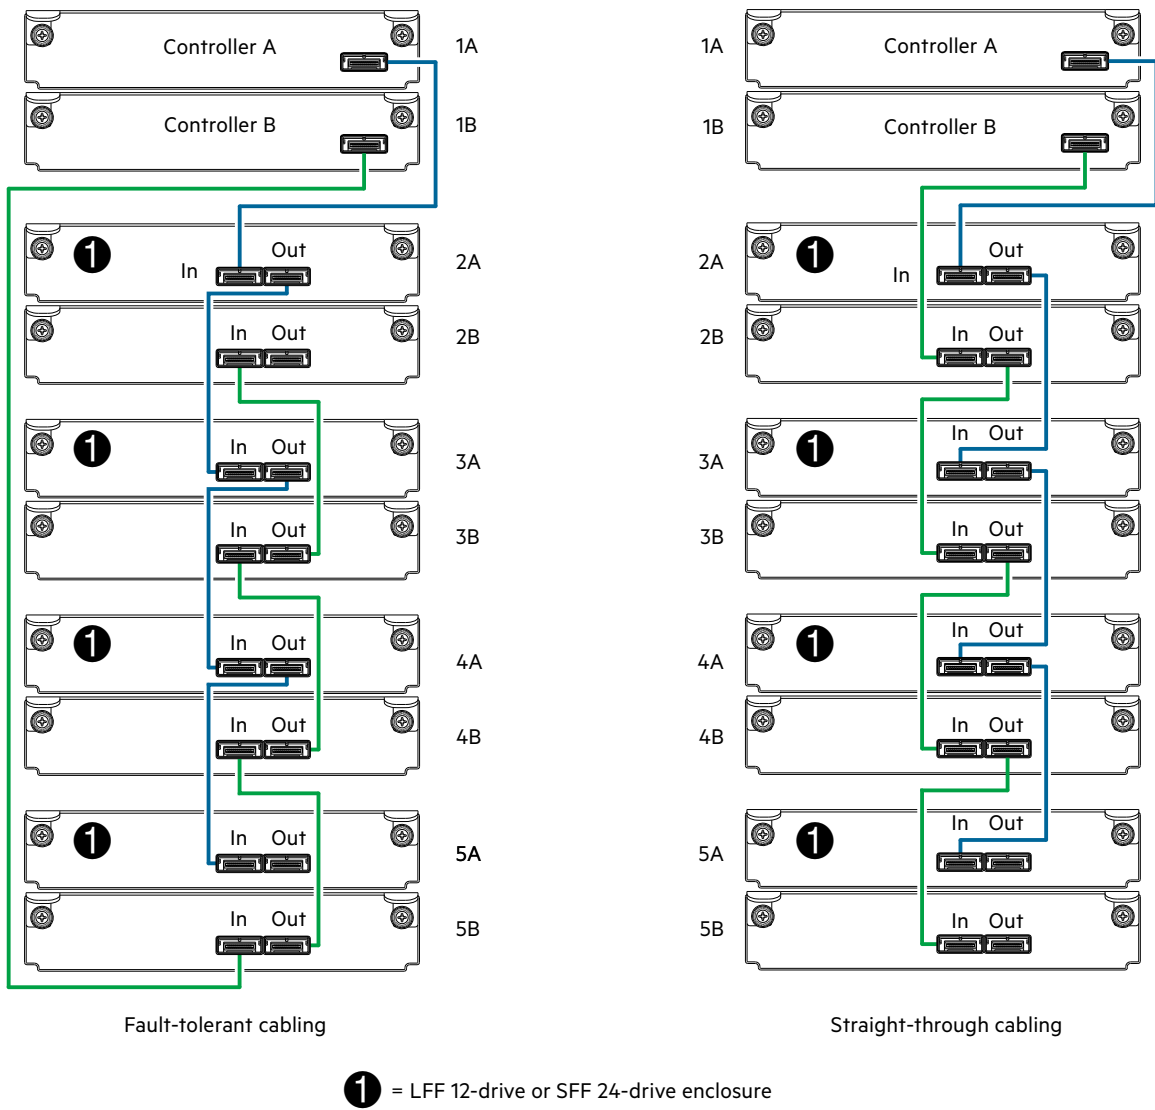



Figure 10 Cabling connections between MSA 2050 controllers and LFF and SFF drive enclosures

The diagram at left (above) shows fault-tolerant cabling of a dual-controller enclosure cabled to either the MSA 2050 LFF Disk Enclosure or the MSA 2050 SFF Disk Enclosure featuring dual-expansion modules. Controller module 1A is connected to expansion module 2A, with a chain of connections cascading down (blue). Controller module 1B is connected to the lower expansion module (5B), of the last drive enclosure, with connections moving in the opposite direction (green). Fault-tolerant cabling allows any drive enclosure to fail—or be removed—while maintaining access to other enclosures.

The diagram at right (above) shows the same storage components connected using straight-through cabling. Using this method, if a drive enclosure fails, the enclosures that follow the failed enclosure in the chain are no longer accessible until the failed enclosure is repaired or replaced.

[Figure 11 \(page 22\)](#) provides a sample diagram reflecting fault-tolerant cabling of a maximum number of supported MSA 2050 enclosures.



Note:
 The maximum number of supported drive enclosures (7) may require purchase of additional longer cables.

Drive enclosure IOM face plate key:

1 = LFF 12-drive enclosure or
 SFF 24-drive enclosure

Figure 11 Fault-tolerant cabling connections showing maximum number of enclosures

❗ **IMPORTANT:** For comprehensive configuration options and associated illustrations, refer to the *HPE MSA 2050 Cable Configuration Guide*.

Testing enclosure connections

NOTE: Once the power-on sequence for enclosures succeeds, the storage system is ready to be connected to hosts, as described in [“Connecting the enclosure to data hosts” \(page 27\)](#).

Powering on/powering off

Before powering on the enclosure for the first time:

- Install all disk drives in the enclosure so the controller can identify and configure them at power-up.
 - Connect the cables and power cords to the enclosures as explained in the quick start instructions.
-

NOTE: Power supplies used in MSA 2050 enclosures

- The MSA 2050 controller enclosures and drive enclosures equipped with AC power supplies do not have power switches (they are switchless). They power on when connected to a power source, and they power off when disconnected.
 - MSA 2050 controller enclosures and drive enclosures equipped with DC power supplies feature power switches.
-

- When powering up, make sure to power up the enclosures and associated host in the following order:

- Drive enclosures first

This ensures that disks in each drive enclosure have enough time to completely spin up before being scanned by the controller modules within the controller enclosure.

While enclosures power up, their LEDs blink. After the LEDs stop blinking—if no LEDs on the front and back of the enclosure are amber—the power-on sequence is complete, and no faults have been detected. See [“LED descriptions” \(page 64\)](#) for descriptions of LED behavior.

- Controller enclosure next

Depending upon the number and type of disks in the system, it may take several minutes for the system to become ready.

- Hosts last (if powered down for maintenance purposes)
-

💡 **TIP:** When powering off, you will reverse the order of steps used for powering on.

Power cycling procedures vary according to the type of power supply unit included with the enclosure. For controller and drive enclosures configured with the switchless AC power supplies, refer to the procedure described under [AC power supply](#) below. For procedures pertaining to controller and drive enclosures configured with DC power supplies, see [“DC power supply” \(page 25\)](#).

❗ **IMPORTANT:** See [“Power cord requirements” \(page 78\)](#) and the QuickSpecs for more information about power cords supported by MSA 2050 enclosures.

AC power supply

Enclosures equipped with switchless power supplies rely on the power cord for power cycling. Connecting the cord from the power supply power cord connector to the appropriate power source facilitates power on, whereas disconnecting the cord from the power source facilitates power off.




Figure 12 AC power supply

AC power cycle

To power on the system:

1. Obtain a suitable AC power cord for each AC power supply that will connect to a power source.
2. Plug the power cord into the power cord connector on the back of the drive enclosure (see [Figure 12](#)). Plug the other end of the power cord into the rack power source. Wait several seconds to allow the disks to spin up.
Repeat this sequence for each power supply within each drive enclosure.
3. Plug the power cord into the power cord connector on the back of the controller enclosure (see [Figure 12](#)). Plug the other end of the power cord into the rack power source.
Repeat the sequence for the controller enclosure's other switchless power supply.

To power off the system:

1. Stop all I/O from hosts to the system [see [“Stopping I/O” \(page 47\)](#)].
2. Shut down both controllers using *either* method described below:
 - o Use the SMU to shut down both controllers, as described in the online help and web-posted *HPE MSA 1050/2050 SMU Reference Guide*.
Proceed to [step 3](#).
 - o Use the CLI to shut down both controllers, as described in the *HPE MSA 1050/2050 CLI Reference Guide*.
3. Disconnect the power cord female plug from the power cord connector on the power supply module.
Perform this step for each power supply module (controller enclosure first, followed by drive enclosures).

NOTE: Power cycling for enclosures equipped with a power switch is described below.

DC power supply

DC power supplies are equipped with a power switch, as shown in the figure below.




Figure 13 DC power supply

Connect power cable to DC power supply

Locate two DC power cables that are compatible with your controller enclosure.





Figure 14 DC power cable featuring sectioned D-shell and lug connectors

See [Figure 13](#) and [Figure 14](#) when performing the following steps:

1. Verify that the enclosure power switches are in the **Off** position.
2. Connect a DC power cable to each DC power supply using the D-shell connector. Use the **UP>** arrow on the connector shell to ensure proper positioning (see adjacent left side view of D-shell connector).
3. Tighten the screws at the top and bottom of the shell, applying a torque between 1.7 N-m (15 in-lb) and 2.3 N-m (20 in-lb), to securely attach the cable to the DC power supply module.
4. To complete the DC connection, secure the other end of each cable wire component of the DC power cable to the target DC power source.



Check the three individual DC cable wire labels before connecting each cable wire lug to its power source. One cable wire is labeled *ground* (GND) and the other two wires are labeled *positive* (+L) and *negative* (-L), respectively (shown in [Figure 14](#) above).

CAUTION: Connecting to a DC power source outside the designated -48V DC nominal range (-36V DC to -72V DC) may damage the enclosure.

DC power cycle

To power on the system:

1. Power up drive enclosure(s).

Press the power switches at the back of each drive enclosure to the **On** position. Allow several seconds for the disks to spin up. See also [Figure 13 \(page 25\)](#).

2. Power up the controller enclosure next.

Press the power switches at the back of the controller enclosure to the **On** position. Allow several seconds for the disks to spin up. See also [Figure 13 \(page 25\)](#).

To power off the system:

1. Stop all I/O from hosts to the system [see [“Stopping I/O” \(page 47\)](#)].
2. Shut down both controllers using either method described below:
 - o Use the SMU to shut down both controllers, as described in the online help and *HPE MSA 1050/2050 SMU Reference Guide*.
Proceed to [step 3](#).
 - o Use the CLI to shut down both controllers, as described in the *HPE MSA 1050/2050 CLI Reference Guide*.
3. Press the power switches at the back of the controller enclosure to the **Off** position. See also [Figure 13 \(page 25\)](#).
4. Press the power switches at the back of each drive enclosure to the **Off** position. See also [Figure 13 \(page 25\)](#).

4 Connecting hosts

Host system requirements

Data hosts connected to HPE MSA 2050 arrays must meet requirements described herein. Depending on your system configuration, data host operating systems may require that multi-pathing is supported.

If fault-tolerance is required, then multi-pathing software may be required. Host-based multi-path software should be used in any configuration where two logical paths between the host and any storage volume may exist at the same time. This would include most configurations where there are multiple connections to the host or multiple connections between a switch and the storage.

- Use native Microsoft MPIO DSM support with **Windows Server 2016** and **Windows Server 2012**. Use either the Server Manager or the command-line interface (mpclaim CLI tool) to perform the installation. Refer to the following web sites for information about using Windows native MPIO DSM:
<https://www.support.microsoft.com>
<http://technet.microsoft.com> (search the site for “multipath I/O overview”)
- Use the HPE Multi-path Device Mapper for Linux Software with Linux servers. To download the appropriate device mapper multi-path enablement kit for your specific enterprise Linux operating system, go to <https://www.hpe.com/storage/spock>.

Connecting the enclosure to data hosts

A *host* identifies an external port to which the storage system is attached. The external port may be a port in an I/O adapter (such as an FC HBA) in a server. Cable connections vary depending on configuration. Common cable configurations are shown in this section. A list of supported configurations is available on the Hewlett Packard Enterprise site at: <https://www.hpe.com/support/msa2050>:

- HPE MSA 2050 Quick Start Instructions
- HPE MSA 2050 Cable Configuration Guide

These documents provide installation details and describe supported direct attach, switch-connect, and storage expansion configuration options for MSA 2050 products. For specific information about qualified host cabling options, see “[Cable requirements for MSA 2050 enclosures](#)” (page 19).

MSA 2050 SAN

MSA 2050 SAN models use Converged Network Controller technology, allowing you to select the desired host interface protocol(s) from the available FC or iSCSI host interface protocols supported by the system. The small form-factor pluggable (SFP transceiver or SFP) connectors used in host ports are further described in the subsections below. Also see “[MSA 2050 SAN](#)” (page 9) for more information concerning use of these host ports.

❗ **IMPORTANT:** Controller modules are not shipped with pre-installed SFPs. Within your product kit, locate the qualified SFP options, and install them into the host ports. See “[Install an SFP transceiver](#)” (page 80).

❗ **IMPORTANT:** Use the `set host-port-mode` CLI command to set the host interface protocol for MSA 2050 SAN host ports using qualified protocol options. MSA 2050 SAN models ship with host ports configured for FC. When connecting host ports to iSCSI hosts, you must use the CLI (not the SMU) to specify which ports will use iSCSI. It is best to do this before inserting the iSCSI SFPs into the host ports.

NOTE: MSA 2050 SAN controllers support the optionally-licensed Remote Snap replication feature. Remote Snap supports FC and iSCSI host interface protocols for replication. Replication sets can be created and viewed using either the SMU or CLI commands.

Fibre Channel protocol


The MSA 2050 SAN controller enclosures support two controller modules using the Fibre Channel interface protocol for host connection. Each controller module provides four host ports designed for use with an FC SFP supporting data rates up to 16 Gb/s. When configured with FC SFPs, MSA 2050 SAN controller enclosures can also be cabled to support the optionally-licensed Remote Snap replication feature via the FC ports.

The MSA 2050 SAN controller supports Fibre Channel Arbitrated Loop (public or private) or point-to-point topologies. Loop protocol can be used in a physical loop or in a direct connection between two devices. Point-to-point protocol is used to connect to a fabric switch. Point-to-point protocol can also be used for direct connection, and it is the only option supporting direct connection at 16 Gb/s. See the `set host-parameters` command within the CLI Reference Guide for command syntax and details about connection mode parameter settings relative to supported link speeds.

Fibre Channel ports are used in either of two capacities:


- To connect two storage systems through a Fibre Channel switch for use of Remote Snap replication.
- For attachment to FC hosts directly, or through a switch used for the FC traffic.


The first usage option requires valid licensing for the Remote Snap replication feature, whereas the second option requires that the host computer supports FC and optionally, multipath I/O.

 **TIP:** Use the SMU Configuration Wizard to set FC port speed. Within the SMU Reference Guide, see “Using the Configuration Wizard” and scroll to FC port options. Use the `set host-parameters` CLI command to set FC port options, and use the `show ports` CLI command to view information about host ports.

10GbE iSCSI protocol

The MSA 2050 SAN controller enclosures support two controller modules using the Internet SCSI interface protocol for host connection. Each controller module provides four host ports designed for use with a 10GbE iSCSI SFP or approved DAC cable supporting data rates up to 10 Gb/s, using either one-way or mutual CHAP (Challenge-Handshake Authentication Protocol).

 **TIP:** See the topics about configuring CHAP, and CHAP and replication in the SMU Reference Guide.

 **TIP:** Use the SMU Configuration Wizard to set iSCSI port options. Within the SMU Reference Guide, see “Using the Configuration Wizard” and scroll to iSCSI port options. Use the `set host-parameters` CLI command to set iSCSI port options, and use the `show ports` CLI command to view information about host ports.

The 10GbE iSCSI ports are used in either of two capacities:


- To connect two storage systems through a switch for use of Remote Snap replication.
- For attachment to 10GbE iSCSI hosts directly, or through a switch used for the 10GbE iSCSI traffic.

The first usage option requires valid licensing for the Remote Snap replication feature, whereas the second option requires that the host computer supports Ethernet, iSCSI, and optionally, multipath I/O.

1 Gb iSCSI protocol

The MSA 2050 SAN controller enclosures support two controller modules using the Internet SCSI interface protocol for host port connection. Each controller module provides four iSCSI host ports configured with an RJ-45 SFP supporting data rates up to 1 Gb/s, using either one-way or mutual CHAP.

 **TIP:** See the topics about configuring CHAP, and CHAP and replication in the SMU Reference Guide.

 **TIP:** Use the SMU Configuration Wizard to set iSCSI port options. Within the SMU Reference Guide, see “Using the Configuration Wizard” and scroll to iSCSI port options. Use the `set host-parameters` CLI command to set iSCSI port options, and use the `show ports` CLI command to view information about host ports.

The 1 Gb iSCSI ports are used in either of two capacities:

- To connect two storage systems through a switch for use of Remote Snap replication.
- For attachment to 1 Gb iSCSI hosts directly, or through a switch used for the 1 Gb iSCSI traffic.

The first usage option requires valid licensing for the Remote Snap replication feature, whereas the second option requires that the host computer supports Ethernet, iSCSI, and optionally, multipath I/O.

12 Gb SAS protocol

Host connection configurations

The MSA 2050 controller enclosures support up to eight direct-connect server connections, four per controller module. Connect appropriate cables from the server HBAs to the controller host ports as described below, and shown in the following illustrations.

NOTE: Not all operating systems support direct-connect. For more information, see the Single Point of Connectivity Knowledge (SPOCK) Storage compatibility matrix: <https://www.hpe.com/storage/spock>.

To connect the MSA 2050 SAN controller to a server or switch—using FC SFPs in controller ports—select Fibre Channel cables supporting 16 Gb data rates, that are compatible with the host port SFP connector (see the QuickSpecs). Such cables are also used for connecting a local storage system to a remote storage system via a switch, to facilitate use of the optional Remote Snap replication feature.

To connect the MSA 2050 SAN controller to a server or switch—using 10GbE iSCSI SFPs or approved DAC cables in controller ports—select the appropriate qualified 10GbE SFP option (see the QuickSpecs). Such cables are also used for connecting a local storage system to a remote storage system via a switch, to facilitate use of the optional Remote Snap replication feature.

To connect the MSA 2050 SAN controller to a server or switch—using the 1 Gb SFPs in controller ports—select the appropriate qualified RJ-45 SFP option (see the QuickSpecs). Such cables are also used for connecting a local storage system to a remote storage system via a switch, to facilitate use of the optional Remote Snap replication feature.


Connecting direct attach configurations

MSA 2050 controller enclosures support dual-controller only. If a partner controller fails, the array will fail over and run on a single controller until the redundancy is restored. A controller module must be installed in each IOM slot to ensure sufficient airflow through the enclosure during operation.

NOTE: The MSA 2050 SAN diagrams that follow use a single representation for each cabling example. This is due to the fact that the port locations and labeling are identical for each of the three possible interchangeable SFPs supported by the system.

One server/one HBA/dual path

MSA 2050 FC or iSCSI



MSA 2050 SAS





Figure 15 Connecting hosts: direct attach—one server/one HBA/dual path

Two servers/one HBA per server/dual path

MSA 2050 FC or iSCSI



MSA 2050 SAS





Figure 16 Connecting hosts: direct attach—two servers/one HBA per server/dual path

Four servers/one HBA per server/dual path

MSA 2050 FC or iSCSI



MSA 2050 SAS




Figure 17 Connecting hosts: direct attach—four servers/one HBA per server/dual path

Connecting switch attach configurations

Two servers/two switches

MSA 2050 FC or iSCSI




Figure 18 Connecting hosts: switch attach—two servers/two switches

Four servers/multiple switches/SAN fabric

MSA 2050 FC or iSCSI




Figure 19 Connecting hosts: switch attach—four servers/multiple switches/SAN fabric

Connecting remote management hosts

The management host directly manages systems out-of-band over an Ethernet network.

1. Connect an RJ-45 Ethernet cable to the network management port on each MSA 2050 controller.
2. Connect the other end of each Ethernet cable to a network that your management host can access (preferably on the same subnet).

NOTE: Connections to this device must be made with shielded cables—grounded at both ends—with metallic RFI/EMI connector hoods, in order to maintain compliance with NEBS and FCC Rules and Regulations.

NOTE: Access via HTTPS and SSH is enabled by default, and access via HTTP and Telnet is disabled by default.

Connecting two storage systems to replicate volumes

Remote Snap replication is a licensed feature for disaster-recovery. This feature performs asynchronous replication of block-level data from a volume in a primary system to a volume in a secondary system by creating an internal snapshot of the primary volume, and copying the changes to the data since the last replication to the secondary system via FC or iSCSI links. The two associated volumes form a replication set, and only the primary volume (source of data) can be mapped for access by a server. Both systems must be licensed to use Remote Snap, and must be connected through switches to the same fabric or network (no direct attach). The server accessing the replication set need only be connected to the primary system. If the primary system goes offline, a connected server can access the replicated data from the secondary system. Replication configuration possibilities are many, and can be cabled—in switch attach fashion—to support MSA 2050 SAN systems on the same network, or on different networks. As you consider the physical connections of your system—specifically connections for replication—keep several important points in mind:

- Ensure that controllers have connectivity between systems, whether the destination system is co-located or remotely located.
- Qualified Converged Network Controller options can be used for host I/O or replication, or both.

- The storage system does not provide for specific assignment of ports for replication. However, this can be accomplished using virtual LANs for iSCSI and zones for FC, or by using physically separate infrastructure. See also paragraph above [Figure 20 \(page 34\)](#).
- For remote replication, ensure that all ports assigned for replication are able to communicate appropriately with the remote replication system (see the CLI Reference Guide for more information) by using the `query peer-connection` CLI command.
- Allow a sufficient number of ports to perform replication. This permits the system to balance the load across those ports as I/O demands rise and fall. If some of the volumes replicated are owned by controller A and others are owned by controller B, then allow at least one port for replication on each controller module—and possibly more than one port per controller module—depending on replication traffic load.
- For the sake of system security, do not unnecessarily expose the controller module network port to an external network connection.

Conceptual cabling examples address cabling on the same network and cabling relative to different networks.

❗ **IMPORTANT:** Remote Snap must be licensed on all systems configured for replication, and the controller module firmware version must be compatible on all systems licensed for replication.

NOTE: Systems must be correctly cabled before performing replication. See the following documents for more information about using Remote Snap to perform replication tasks:

- HPE Remote Snap technical white paper: [MSA Remote Snap Software](#)
- HPE MSA 2050 Best Practices
- HPE MSA 1050/2050 SMU Reference Guide
- HPE MSA 1050/2050 CLI Reference Guide
- HPE MSA Event Descriptions Reference Guide
- HPE MSA 2050 Cable Configuration Guide

To access MSA 2050 documentation, see the Hewlett Packard Enterprise Information Library:

<https://www.hpe.com/support/msa2050>

Cabling for replication

This section shows example replication configurations for MSA 2050 SAN controller enclosures. The following illustrations provide conceptual examples of cabling to support Remote Snap replication.

NOTE: Simplified versions of controller enclosures are used in cabling illustrations to show host ports used for I/O or replication, given that only the external connectors used in the host interface ports differ.

- Virtual replication supports FC and iSCSI host interface protocols.
- The 2U enclosure rear panel represents MSA 2050 SAN models.
- Host ports used for replication must use the same protocol (either FC or iSCSI)
- Blue cables show I/O traffic and green cables show replication traffic.

Once the MSA 2050 systems are physically cabled, see the SMU Reference Guide or online help for information about configuring, provisioning, and using the optional Remote Snap feature.

Host ports and replication

MSA 2050 SAN controller modules can use qualified SFP options of the same type, or they can use a combination of qualified SFP options supporting different interface protocols. If you use a combination of different protocols, then host

ports 1 and 2 are set to FC (either both 16 Gb/s or both 8 Gb/s), and host ports 3 and 4 must be set to iSCSI (either both 10GbE or both 1 Gb). FC and iSCSI ports can either be used to perform I/O or replication.

-
- ① **IMPORTANT:** MSA 2050 controller enclosures support dual-controller only. If a partner controller fails, the array will fail over and run on a single controller until the redundancy is restored. A controller module must be installed in each IOM slot to ensure sufficient airflow through the enclosure during operation.
-

Each of the following diagrams show the rear panel of two MSA 2050 SAN controller enclosures equipped with dual-controller modules. The controller modules can use qualified SFP options of the same type, or they can use a combination of qualified SFP options supporting different interface protocols.

-
- ① **IMPORTANT:** MSA 2050 SAN controllers support FC and iSCSI host interface protocols for host connection or for performing replications.
-

Multiple servers/single network

The diagram below shows the rear panel of two MSA 2050 SAN controller enclosures with both I/O and replication occurring on the same physical network. With the replication configuration shown below, Virtual Local Area Network (VLAN) and zoning could be employed to provide separate networks for iSCSI and FC, respectively. Create a VLAN or zone for I/O and a VLAN or zone for replication to isolate I/O traffic from replication traffic. The configuration would appear physically as a single network, while logically, it would function as multiple networks.




Figure 20 Connecting two storage systems for Remote Snap: multiple servers/one switch/one location

The diagram below shows the rear panel of two MSA 2050 SAN controller enclosures with I/O and replication occurring on different physical networks. Use three switches to enable host I/O and replication. Connect two ports from each controller module in the left storage enclosure to the left switch. Connect two ports from each controller module in the right storage enclosure to the right switch. Connect two ports from each controller module in each enclosure to the middle switch. Use multiple switches to avoid a single point of failure inherent to using a single switch, and to physically isolate replication traffic from I/O traffic.




Figure 21 Connecting two storage systems for Remote Snap: multiple servers/switches/one location

Multiple servers/different networks/multiple switches

The diagram below shows the rear panel of two MSA 2050 SAN controller enclosures with both I/O and replication occurring on different networks.

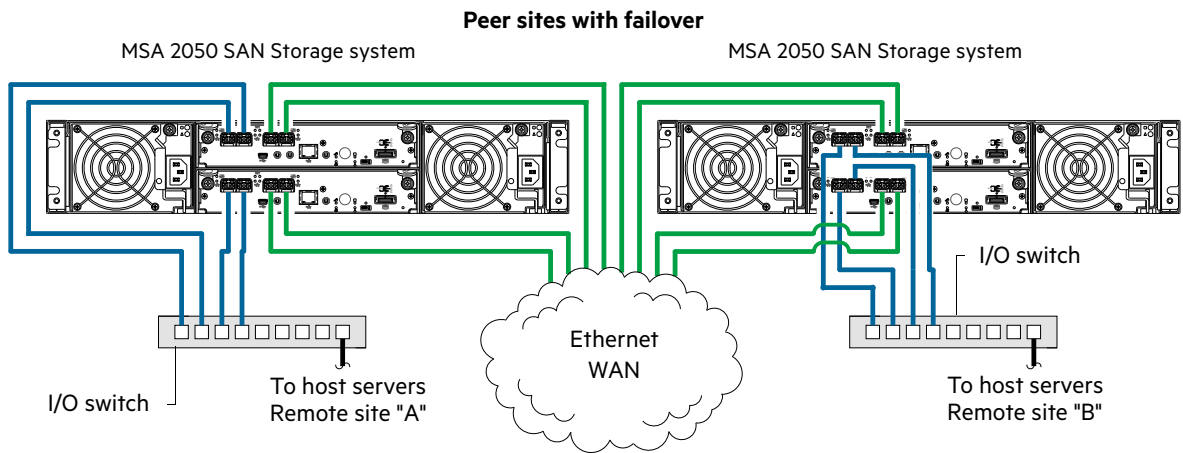


Figure 22 Connecting two storage systems for Remote Snap: multiple servers/switches/two locations

The diagram below also shows the rear-panel of two MSA 2050 SAN controller enclosures with I/O occurring on the FC SAN and replication occurring on the Ethernet WAN. This diagram represents two branch offices cabled to enable disaster recovery and backup. In case of failure at either the local site or the remote site, you can fail over the application to the available site.

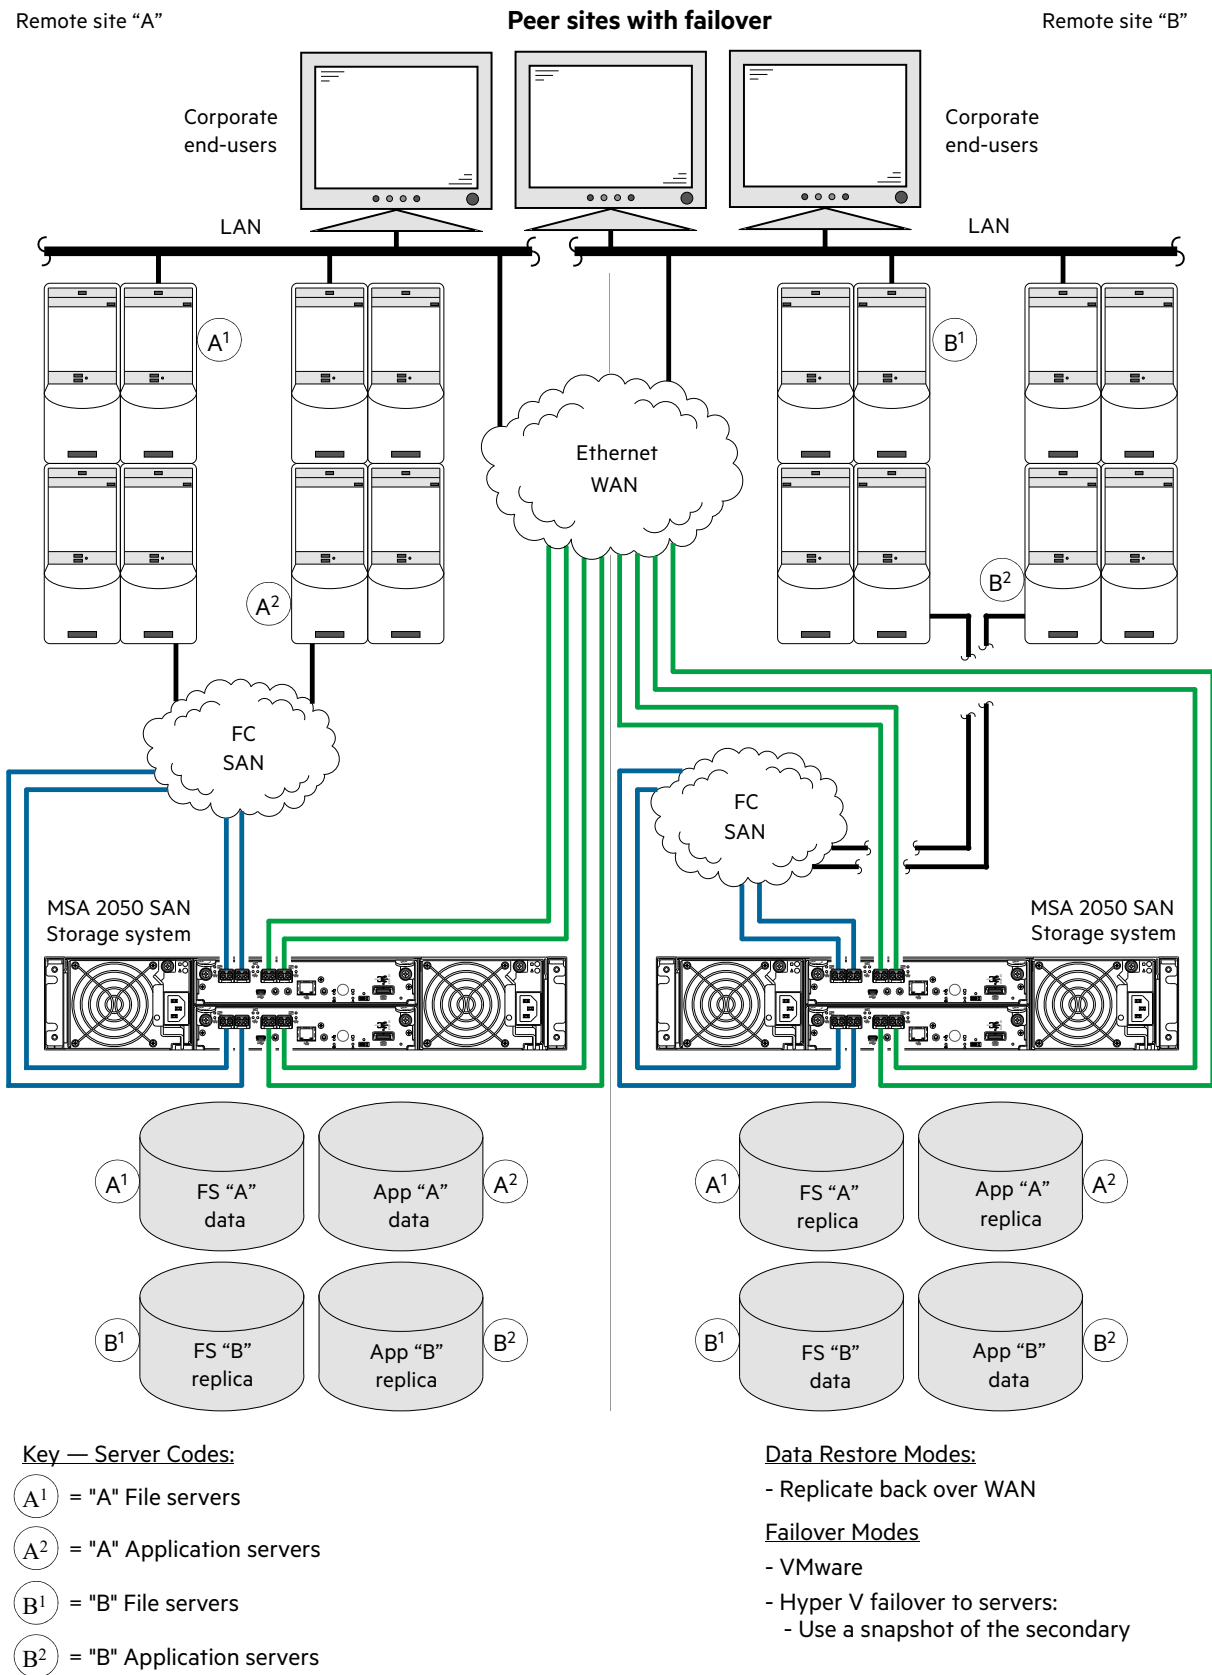


Figure 23 Connecting two storage systems for Remote Snap: multiple servers/SAN fabric/two locations

Although not shown in the preceding cabling examples, you can cable replication-enabled MSA 2050 SAN systems and compatible MSA 1040/2040 systems—via switch attach—for performing replication tasks limited to the Remote Snap functionality of the MSA 1040/2040 storage system.

Updating firmware

After installing the hardware and powering on the storage system components for the first time, verify that the controller modules, expansion modules, and disk drives are using the current firmware release.

NOTE: Update component firmware by installing a firmware file obtained from the HPE web download site as described in “[Accessing updates](#)” (page 61). To install an HPE ROM Flash Component or firmware Smart Component, follow the instructions on the HPE website.

Otherwise, to install a firmware binary file, follow the steps below.

Using the SMU, in the **System** topic, select **Action > Update Firmware**.

The Update Firmware panel opens. The Update Controller Modules tab shows versions of firmware components currently installed in each controller.

NOTE: Partner Firmware Update using management interfaces:

- The SMU provides an option for enabling or disabling Partner Firmware Update for the partner controller.
 - To enable or disable the setting via the CLI, use the `set advanced-settings` command, and set the `partner-firmware-upgrade` parameter. See the CLI Reference Guide for more information about command parameter syntax.
 - HPE recommends that Partner Firmware Update is enabled (the default setting).
-

Optionally, you can update firmware using SFTP or FTP as described in the SMU Reference Guide.

ⓘ **IMPORTANT:** See the topics about updating firmware within the SMU Reference Guide before performing a firmware update.

NOTE: To locate and download the latest software and firmware updates for your product, go to <https://www.hpe.com/support/downloads>.

5 Connecting to the controller CLI port

Device description

The MSA 2050 controllers feature a command-line interface port used to cable directly to the controller and initially set IP addresses, or perform other configuration tasks. This port employs a mini-USB Type B form factor, requiring a cable that is supplied with the controller, and additional support described herein, so that a server or other computer running a Linux or Windows operating system can recognize the controller enclosure as a connected device. Without this support, the computer might not recognize that a new device is connected, or might not be able to communicate with it. The USB device driver is implemented using the abstract control model (ACM) to ensure broad support.

For Linux computers, no new driver files are needed, but depending on the version of operating system, a Linux configuration file may need to be created or modified.

For Windows computers, if you are not using Windows 10/Server 2016, the Windows USB device driver must be downloaded from the HPE website, and installed on the computer that will be cabled directly to the controller command-line interface port (see also <https://www.hpe.com/support/downloads>).

NOTE: Directly cabling to the CLI port is an out-of-band connection because it communicates outside the data paths used to transfer information from a computer or network to the controller enclosure.

Emulated serial port

Once attached to the controller module as shown in [Figure 24 \(page 40\)](#), the management computer should detect a new USB device. Using the Emulated Serial Port interface, the controller presents a single serial port using a *vendor ID* and *product ID*. Effective presentation of the emulated serial port assumes the management controller previously had a terminal emulator installed (see [Table 3](#)). MSA 2050 controllers support the following applications to facilitate connection.

Table 3 Supported terminal emulator applications

Application	Operating system
HyperTerminal, TeraTerm, PuTTY	Microsoft Windows (all versions)
Minicom	Linux (all versions)
	Solaris
	HP-UX

Certain operating systems require a device driver or special mode of operation. Vendor and product identification are provided in [Table 4](#).

Table 4 Terminal emulator display settings

USB identification code type	Code
USB vendor identification	0x210c
USB product identification	0xa4a7

Preparing a Linux computer for cabling to the CLI port

You can determine if the operating system recognizes the USB (ACM) device by entering a command:

```
cat /proc/devices/ |grep -i "ttyACM"
```

If a device driver is discovered, the output will display:

```
ttyACM (and a device number)
```

You can query information about USB buses and the devices connected to them by entering a command:

```
lsusb
```

If a USB device driver is discovered, the output will display:

```
ID 210c:a4a7
```

The ID above is comprised of vendor ID and product ID terms as shown in [Table 4 \(page 38\)](#).

- ❗ **IMPORTANT:** Although Linux systems do not require installation of a device driver, on some operating system versions, certain parameters must be provided during driver loading to enable recognition of the MSA 2050 controllers. To load the Linux device driver with the correct parameters on these systems, the following command is required:

```
modprobe usbserial vendor=0x210c product=0xa4a7 use_acm=1
```

Optionally, the information can be incorporated into the `/etc/modules.conf` file.

Preparing a Windows computer for cabling to the CLI port

A Windows USB device driver is used for communicating directly with the controller command-line interface port using a USB cable to connect the controller enclosure and the computer.

- ❗ **IMPORTANT:** If using Windows 10/Server 2016, the operating system provides a native USB serial driver that supports the controller module's USB CLI port. However, if using an older version of Windows, you should download and install the USB device driver from your HPE MSA support page at <https://www.hpe.com/support/downloads>.
-

Obtaining IP values

One method of obtaining IP values for your system is to use a network management utility to discover “HPE MSA Storage” devices on the local LAN through SNMP. Alternative methods for obtaining IP values for your system are described in the following subsections.

Setting network port IP addresses using DHCP

In DHCP mode, network port IP address, subnet mask, and gateway values are obtained from a DHCP server if one is available. If a DHCP server is unavailable, current addressing is unchanged.

1. Look in the DHCP server's pool of leased addresses for two IP addresses assigned to “HPE MSA Storage.”
2. Use a ping broadcast to try to identify the device through the ARP table of the host.

If you do not have a DHCP server, you will need to ask your system administrator to allocate two IP addresses, and set them using the command-line interface during initial configuration (described below).

NOTE: For more information, see the Configuration Wizard topic about network configuration within the SMU Reference Guide.

Setting network port IP addresses using the CLI port and cable

You can set network port IP addresses manually using the command-line interface port and cable. If you have not done so already, you need to enable your system for using the command-line interface port [also see [“Using the CLI port and cable—known issues on Windows” \(page 43\)](#)].

NOTE: For Linux systems, see “Preparing a Linux computer for cabling to the CLI port” (page 38). For Windows systems see “Preparing a Windows computer for cabling to the CLI port” (page 39).

Network ports on controller module A and controller module B are configured with the following factory-default IP settings:

- **Management Port IP Address:** 10.0.0.2 (controller A), 10.0.0.3 (controller B)
- **IP Subnet Mask:** 255.255.255.0
- **Gateway IP Address:** 10.0.0.1

If the default IP addresses are not compatible with your network, you must set an IP address for each network port using the command-line interface.

Use the CLI commands described in the steps below to set the IP address for the network port on each controller module. Once new IP addresses are set, you can change them as needed using the SMU.

NOTE: Changing IP settings can cause management hosts to lose access to the storage system.

1. From your network administrator, obtain an IP address, subnet mask, and gateway address for controller A, and another for controller B.
Record these IP addresses so that you can specify them whenever you manage the controllers using the SMU or the CLI.
2. Use the provided USB cable to connect controller A to a USB port on a host computer. The USB mini 5 male connector plugs into the CLI port as shown in Figure 24 (generic controller module is shown).




Figure 24 Connecting a USB cable to the CLI port

3. Enable the CLI port for subsequent communication.
If the USB device is supported natively by the operating system, proceed to [step 4](#).
 - o Linux customers should enter the command syntax provided in “Preparing a Linux computer for cabling to the CLI port” (page 38).
 - o Windows customers should locate the downloaded device driver described in “Preparing a Windows computer for cabling to the CLI port” (page 39), and follow the instructions provided for proper installation.

4. Start and configure a terminal emulator, such as HyperTerminal or VT-100, using the display settings in [Table 5](#) and the connection settings in [Table 6](#) (also, see the note following this procedure).

Table 5 Terminal emulator display settings

Parameter	Value
Terminal emulation mode	VT-100 or ANSI (for color support)
Font	Terminal
Translations	None
Columns	80

Table 6 Terminal emulator connection settings

Parameter	Value
Connector	COM3 (for example) ^{1,2}
Baud rate	115,200
Data bits	8
Parity	None
Stop bits	1
Flow control	None

¹ Your server or laptop configuration determines which COM port is used for Disk Array USB Port.

² Verify the appropriate COM port for use with the CLI.

5. In the terminal emulator, connect to controller A.

6. Press **Enter** to display the CLI prompt (#).

The CLI displays the system version, MC version, and login prompt:

- a. At the login prompt, enter your user name.
- b. Enter your password.

ⓘ IMPORTANT: As of MSA firmware version VE270P003 (MSA 1050) and VL270P003 (MSA 2050 and MSA 2052), the MSA systems no longer support the default user name and password for logging into the system. After upgrading to this firmware version, you will be prompted to create a new user name and password combination using the setup user account.

7. At the prompt, type the following command to set the values you obtained in [step 1](#) for each network port, first for controller A and then for controller B:

```
set network-parameters ip address netmask netmask gateway gateway controller a|b
```

where:

- o *address* is the IP address of the controller
- o *netmask* is the subnet mask
- o *gateway* is the IP address of the subnet router
- o *a|b* specifies the controller whose network parameters you are setting

For example:

```
# set network-parameters ip 192.168.0.10 netmask 255.255.255.0 gateway 192.168.0.1 controller a
```

```
# set network-parameters ip 192.168.0.11 netmask 255.255.255.0 gateway 192.168.0.1 controller b
```

8. Type the following command to verify the new IP addresses:

```
show network-parameters
```

Network parameters, including the IP address, subnet mask, and gateway address are displayed for each controller.

9. Use the `ping` command to verify network connectivity.

For example, to ping the gateway in the examples above:

```
# ping 192.168.0.1
```

```
Info: Pinging 192.168.0.1 with 4 packets.
```

```
Success: Command completed successfully. - The remote computer responded with 4 packets.
```

10. In the host computer's command window, type the following command to verify connectivity, first for controller A and then for controller B:

```
ping controller-IP-address
```

If you cannot access your system for at least three minutes after changing the IP address, your network might require you to restart the Management Controller(s) using the CLI. When you restart a Management Controller, communication with it is temporarily lost until it successfully restarts.

Type the following command to restart the management controller on both controllers:

```
restart mc both
```

11. When you are done using the CLI, exit the emulator.
12. Retain the new IP addresses to access and manage the controllers, using either the SMU or the CLI.

NOTE: Using HyperTerminal with the CLI on a **Microsoft Windows** host:

On a host computer connected to a controller module's mini-USB CLI port, incorrect command syntax in a HyperTerminal session can cause the CLI to hang. To avoid this problem, use correct syntax, use a different terminal emulator, or connect to the CLI using SSH rather than the mini-USB cable.

Be sure to close the HyperTerminal session before shutting down the controller or restarting its Management Controller. Otherwise, the host's CPU cycles may rise unacceptably.

If communication with the CLI is disrupted when using an out-of-band cable connection, communication can sometimes be restored by disconnecting and reattaching the mini-USB cable as described in [step 2](#) on [page 40](#).

NOTE: If using a Windows operating system version older than Windows 10/Server 2016, access the USB device driver download from the HPE MSA support website at <https://www.hpe.com/support/downloads>.

Using the CLI port and cable—known issues on Windows

When using the CLI port and cable for setting controller IP addresses and other operations, be aware of the following known issues on Microsoft Windows platforms.

Problem

On Windows operating systems, the USB CLI port may encounter issues preventing the terminal emulator from reconnecting to storage after the Management Controller (MC) restarts or the USB cable is unplugged and reconnected.

Workaround

Follow these steps when using the mini-USB cable and USB Type B CLI port to communicate out-of-band between the host and controller module for setting network port IP addresses.

To restore a hung connection when the MC is restarted (any supported terminal emulator):

1. If the connection hangs, disconnect and quit the terminal emulator program.
 - a. Using Device Manager, locate the COM n port assigned to the Disk Array Port.
 - b. Right-click on the hung **Disk Array USB Port (COM n)**, and select **Disable**.
 - c. Wait for the port to disable.
2. Right-click on the previously hung—now disabled—**Disk Array USB Port (COM n)**, and select **Enable**.
3. Start the terminal emulator and connect to the COM port.
4. Set network port IP addresses using the CLI (see procedure on [page 39](#)).

NOTE: When using Windows 10/Server 2016 with PuTTY, the XON/XOFF setting must be disabled, or the COM port will not open.

6 Basic operation

Verify that you have completed the sequential “Installation Checklist” instructions in [Table 2 \(page 17\)](#). Once you have successfully completed steps 1 through 8 therein, you can access the management interface using your web browser to complete the system setup.

Accessing the SMU

Upon completing the hardware installation, you can access the web-based management interface—SMU (Storage Management Utility)—from the controller module to monitor and manage the storage system. Invoke your web browser, and enter the `https://IP-address` of the controller module’s network port in the address field (obtained during completion of “Installation Checklist” step 8), then press **Enter**. To sign in to the SMU for the first time, create a new user with the **manage** capability and a password. This brief sign in process assumes proper web browser setup.

❗ **IMPORTANT:** As of MSA firmware version VE270P003 (MSA 1050) and VL270P003 (MSA 2050 and MSA 2052), the MSA systems no longer support the default user name and password for logging into the system. After upgrading to this firmware version, you will be prompted to create a new user name and password combination using the setup user account.

❗ **IMPORTANT:** For detailed information about accessing and using the SMU, see the topic about getting started in the SMU Reference Guide.

The Getting Started section provides instructions for signing-in to the SMU, introduces key concepts, addresses browser setup, and provides tips for using the main window and the help window.

💡 **TIP:** After signing in to the SMU, you can use online help as an alternative to consulting the reference guide.

Configuring and provisioning the storage system

Once you have familiarized yourself with the SMU, use it to configure and provision the storage system. If you are licensed to use the optional Remote Snap feature, you may also need to set up storage systems for replication. Refer to the following topics within the SMU Reference Guide or online help:

- Configuring the system
 - Provisioning the system
 - Using Remote Snap to replicate volumes
-

❗ **IMPORTANT:** Some features within the storage system require a license. The license is specific to the controller enclosure and firmware version. See the topic about installing a license within the SMU Reference Guide for instructions about viewing the status of licensed features and installing a license.

❗ **IMPORTANT:** If the system is used in a VMware environment, set the system Missing LUN Response option to use its Illegal Request setting. To do so, see either the topic about changing the missing LUN response in the SMU Reference Guide, or the topic about the `set-advanced-settings` command in the CLI Reference Guide.

7 Troubleshooting

These procedures are intended to be used only during initial configuration, for the purpose of verifying that hardware setup is successful. They are not intended to be used as troubleshooting procedures for configured systems using production data and I/O.

USB CLI port connection

MSA 2050 controllers feature a CLI port employing a mini-USB Type B form factor. If you encounter problems communicating with the port after cabling your computer to the USB device, you may need to either download a device driver (Windows), or set appropriate parameters via an operating system command (Linux). See [“Connecting to the controller CLI port” \(page 38\)](#) for more information.

Fault isolation methodology

MSA 2050 controllers provide many ways to isolate faults. This section presents the basic methodology used to locate faults within a storage system, and to identify the associated Field Replaceable Units (FRUs) affected.

As noted in [“Basic operation” \(page 44\)](#), use the SMU to configure and provision the system upon completing the hardware installation. As part of this process, configure and enable event notification so the system will notify you when a problem occurs that is at or above the configured severity (see [“Using the Configuration Wizard > Configuring event notification”](#) within the SMU Reference Guide). With event notification configured and enabled, you can follow the recommended actions in the notification message to resolve the problem, as further discussed in the options presented below.

Basic steps

The basic fault isolation steps are listed below:

- Gather fault information, including using system LEDs [see [“Gather fault information” \(page 46\)](#)].
- Determine where in the system the fault is occurring [see [“Determine where the fault is occurring” \(page 46\)](#)].
- Review event logs [see [“Review the event logs” \(page 46\)](#)].
- If required, isolate the fault to a data path component or configuration [see [“Isolate the fault” \(page 47\)](#)].

Cabling systems to enable use of the licensed Remote Snap feature—to replicate volumes—is another important fault isolation consideration pertaining to initial system installation. See [“Isolating Remote Snap replication faults” \(page 55\)](#) for more information about troubleshooting during initial setup.

Options available for performing basic steps

When performing fault isolation and troubleshooting steps, select the option or options that best suit your site environment. Use of any option (four options are described below) is not mutually-exclusive to the use of another option. You can use the SMU to check the health icons/values for the system and its components to ensure that everything is okay, or to drill down to a problem component. If you discover a problem, both the SMU and the CLI provide recommended action text online. Options for performing basic steps are listed according to frequency of use:

- Use the SMU.
- Use the CLI.
- Monitor event notification.
- View the enclosure LEDs.

Use the SMU

The SMU uses health icons to show OK, Degraded, Fault, or Unknown status for the system and its components. The SMU enables you to monitor the health of the system and its components. If any component has a problem, the system health will be Degraded, Fault, or Unknown. Use the SMU to drill down to find each component that has a problem, and follow actions in the Recommendation field for the component to resolve the problem.

Use the CLI

As an alternative to using the SMU, you can run the `show system` command in the CLI to view the health of the system and its components. If any component has a problem, the system health will be Degraded, Fault, or Unknown, and those components will be listed as Unhealthy Components. Follow the recommended actions in the component Health Recommendation field to resolve the problem.

Monitor event notification

With event notification configured and enabled, you can view event logs to monitor the health of the system and its components. If a message tells you to check whether an event has been logged, or to view information about an event in the log, you can do so using either the SMU or the CLI. Using the SMU, you would view the event log and then click on the event message to see detail about that event. Using the CLI, you would run the `show events detail` command (with additional parameters to filter the output) to see the detail for an event.

View the enclosure LEDs

You can view the LEDs on the hardware (while referring to [LED descriptions](#) for your enclosure model) to identify component status. If a problem prevents access to either the SMU or the CLI, this is the only option available. However, monitoring/management is often done at a management console using storage management interfaces, rather than relying on line-of-sight to LEDs of racked hardware components.

Performing basic steps

You can use any of the available options in performing the basic steps comprising the fault isolation methodology.

Gather fault information

When a fault occurs, it is important to gather as much information as possible. Doing so will help you determine the correct action needed to remedy the fault.

Begin by reviewing the reported fault:

- Is the fault related to an internal data path or an external data path?
- Is the fault related to a hardware component such as a disk drive module, controller module, or power supply?

By isolating the fault to *one* of the components within the storage system, you will be able to determine the necessary action more quickly.

Determine where the fault is occurring

Once you have an understanding of the reported fault, review the enclosure LEDs. The enclosure LEDs are designed to alert users of any system faults, and might be what alerted the user to a fault in the first place.

When a fault occurs, the Fault ID status LED on the enclosure right ear [see [“Front panel components” \(page 11\)](#)] illuminates. Check the LEDs on the back of the enclosure to narrow the fault to a FRU, connection, or both. The LEDs also help you identify the location of a FRU reporting a fault.

Use the SMU to verify any faults found while viewing the LEDs. The SMU is also a good tool to use in determining where the fault is occurring if the LEDs cannot be viewed due to the location of the system. The SMU provides you with a visual representation of the system and where the fault is occurring. It can also provide more detailed information about FRUs, data, and faults.

Review the event logs

The event logs record all system events. Each event has a numeric code that identifies the type of event that occurred, and has one of the following severities:

- **Critical.** A failure occurred that may cause a controller to shut down. Correct the problem *immediately*.
- **Error.** A failure occurred that may affect data integrity or system stability. Correct the problem as soon as possible.
- **Warning.** A problem occurred that may affect system stability, but not data integrity. Evaluate the problem and correct it if necessary.

- Informational. A configuration or state change occurred, or a problem occurred that the system corrected. No immediate action is required.

For information about specific events, see the Event Descriptions Reference Guide, located on the Hewlett Packard Enterprise Information Library at: <https://www.hpe.com/support/msa2050>.

The event logs record all system events. It is very important to review the logs, not only to identify the fault, but also to search for events that might have caused the fault to occur. For example, a host could lose connectivity to a disk group if a user changes channel settings without taking the storage resources assigned to it into consideration. In addition, the type of fault can help you isolate the problem to either hardware or software.

Isolate the fault

Occasionally it might become necessary to isolate a fault. This is particularly true with data paths, due to the number of components comprising the data path. For example, if a host-side data error occurs, it could be caused by any of the components in the data path: controller module, cable, connectors, switch, or data host.

If the enclosure does not initialize

It may take up to two minutes for the enclosures to initialize. If the enclosure does not initialize:

- Perform a rescan.
- Power cycle the system.
- Make sure the power cord is properly connected, and check the power source that it is connected to.
- Check the event log for errors.

Correcting enclosure IDs

When installing a system with drive enclosures attached, the enclosure IDs might not agree with the physical cabling order. This is because the controller might have been previously attached to some of the same enclosures during factory testing, and it attempts to preserve the previous enclosure IDs if possible. To correct this condition, make sure that both controllers are up, and perform a rescan using the SMU or the CLI. This will reorder the enclosures, but can take up to two minutes for the enclosure IDs to be corrected.

To perform a rescan using the CLI, type the following command:

```
rescan
```

To rescan using the SMU:

1. Verify that both controllers are operating normally.
2. Do one of the following:
 - o Point to the **System** tab and select **Rescan Disk Channels**.
 - o In the **System** topic, select **Action > Rescan Disk Channels**.
3. Click **Rescan**.

Stopping I/O

When troubleshooting disk drive and connectivity faults, stop I/O to the affected disk groups from all hosts and remote systems as a data protection precaution. As an additional data protection precaution, it is recommended to conduct regularly scheduled backups of your data.

❗ **IMPORTANT:** Stopping I/O to a disk group is a host-side task, and falls outside the scope of this document.

When on-site, you can verify there is no I/O activity by briefly monitoring the system LEDs. When accessing the storage system remotely, this is not possible. Remotely, you can use the `show disk-group-statistics` CLI command to determine if input and output has stopped. Perform these steps:

1. Using the CLI, run the `show disk-group-statistics` command.
The Reads and Writes outputs show the number of these operations that have occurred since the statistic was last reset, or since the controller was restarted. Record the numbers displayed.
2. Run the `show disk-group-statistics` command a second time.
This provides you a specific window of time (the interval between requesting the statistics) to determine if data is being written to or read from the disk group. Record the numbers displayed.
3. To determine if any reads or writes occur during this interval, subtract the set of numbers you recorded in [step 1](#) from the numbers you recorded in [step 2](#).
 - o If the resulting difference is zero, then I/O has stopped.
 - o If the resulting difference is not zero, a host is still reading from or writing to this disk group. Continue to stop I/O from hosts, and repeat [step 1](#) and [step 2](#) until the difference in [step 3](#) is zero.

See the CLI Reference Guide for additional information on the Hewlett Packard Enterprise Information Library at: <https://www.hpe.com/support/msa2050>.

Diagnostic steps

This section describes possible reasons and actions to take when an LED indicates a fault condition during initial system setup. See “[LED descriptions](#)” ([page 64](#)) for descriptions of all LED statuses.

NOTE: Once event notification is configured and enabled using the SMU, you can view event logs to monitor the health of the system and its components using the GUI.

In addition to monitoring LEDs via line-of-sight observation of racked hardware components when performing diagnostic steps, you can also monitor the health of the system and its components using the management interfaces. Be mindful of this when reviewing the **Actions** column in the diagnostics tables, and when reviewing the step procedures provided in this chapter.

Is the enclosure front panel Fault/Service Required LED amber?

Table 7 Diagnostics LED status: Front panel “Fault/Service Required”

Answer	Possible reasons	Actions
No	System functioning properly.	No action required.
Yes	A fault condition exists/occurred. If installing an I/O module FRU, the module has not gone online and likely failed its self-test.	<ul style="list-style-type: none"> • Check the LEDs on the back of the controller enclosure to narrow the fault to a FRU, connection, or both. • Check 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 above actions do not resolve the fault, isolate the fault, and contact an authorized service provider for assistance. Replacement may be necessary.

Is the enclosure rear panel FRU OK LED off?

Table 8 Diagnostics LED status: Rear panel “FRU OK”

Answer	Possible reasons	Actions
No (blinking)	System functioning properly. System is booting.	No action required. Wait for system to boot.
Yes	The controller module is not powered on. The controller module has failed.	<ul style="list-style-type: none"> Check that the controller module is fully inserted and latched in place, and that the enclosure is powered on. Check the event log for specific information regarding the failure.

Is the enclosure rear panel Fault/Service Required LED amber?

Table 9 Diagnostics LED status: Rear panel “Fault/Service Required”

Answer	Possible reasons	Actions
No	System functioning properly.	No action required.
Yes (blinking)	One of the following errors occurred: <ul style="list-style-type: none"> Hardware-controlled power-up error Cache flush error Cache self-refresh error 	<ul style="list-style-type: none"> Restart this controller from the other controller using the SMU or the CLI. If the above action does not resolve the fault, remove the controller and reinsert it. If the above action does not resolve the fault, contact an authorized service provider for assistance. It may be necessary to replace the controller.

Are both disk drive module LEDs off (Online/Activity and Fault/UID)?

Table 10 Diagnostics LED status: Front panel disks “Online/Activity” and “Fault/UID”

Answer	Possible reasons	Actions
Yes	<ul style="list-style-type: none"> There is no power. The disk is offline. The disk is not configured. 	<ul style="list-style-type: none"> Check that the disk drive is fully inserted and latched in place, and that the enclosure is powered on.

NOTE: See “Disk drives used in MSA 2050 enclosures” (page 12).

Is the disk drive module Fault/UID LED blinking amber?

Table 11 Diagnostics LED status: Front panel disks “Fault/UID”

Answer	Possible reasons	Actions
No, but the Online/Activity LED is blinking.	The disk drive is rebuilding.	No action required. <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p>⚠ CAUTION: Do not remove a disk drive that is reconstructing. Removing a reconstructing disk drive might terminate the current operation and cause data loss.</p> </div>

Table 11 Diagnostics LED status: Front panel disks “Fault/UID” (continued)

Answer	Possible reasons	Actions
Yes, and the Online/Activity LED is off.	The disk drive is offline. A predictive failure alert may have been received for this device.	<ul style="list-style-type: none"> Check the event log for specific information regarding the fault. Isolate the fault. Contact an authorized service provider for assistance.
Yes, and the Online/Activity LED is blinking.	The disk drive is active, but a predictive failure alert may have been received for this device.	<ul style="list-style-type: none"> Check the event log for specific information regarding the fault. Isolate the fault. Contact an authorized service provider for assistance.

NOTE: See “FDE considerations” (page 17) and Figure 30 (page 67).

Is a connected host port Host Link Status LED off?

Table 12 Diagnostics LED status: Rear panel “Host Link Status”

Answer	Possible reasons	Actions
No	System functioning properly.	No action required (see Link LED note: page 71).
Yes	The link is down.	<ul style="list-style-type: none"> Check cable connections and reseal if necessary. Inspect cables for damage. Replace cable if necessary. Swap cables to determine if fault is caused by a defective cable. 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 data path component. Follow any recommended actions. Contact an authorized service provider for assistance. See “Isolating a host-side connection fault” (page 52).

Is a connected port Expansion Port Status LED off?

Table 13 Diagnostics LED status: Rear panel “Expansion Port Status”

Answer	Possible reasons	Actions
No	System functioning properly.	No action required.
Yes	The link is down.	<ul style="list-style-type: none"> Check cable connections and reseal if necessary. Inspect cable for damage. Replace cable if necessary. Swap cables to determine if fault is caused by a defective cable. Replace cable if necessary. In the SMU, review event logs for indicators of a specific fault in a host data path component. Follow any recommended actions. Contact an authorized service provider for assistance. See “Isolating a controller module expansion port connection fault” (page 55).

Is a connected port Network Port Link Status LED off?

Table 14 Diagnostics LED status: Rear panel “Network Port Link Status”

Answer	Possible reasons	Actions
No	System functioning properly.	No action required.
Yes	The link is down.	Use standard networking troubleshooting procedures to isolate faults on the network.

Is the power supply Input Power Source LED off?

Table 15 Diagnostics LED status: Rear panel power supply “Input Power Source”

Answer	Possible reasons	Actions
No	System functioning properly.	No action required.
Yes	The power supply is not receiving adequate power.	<ul style="list-style-type: none">• Verify that the power cord 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 above action does not resolve the fault, isolate the fault, and contact an authorized service provider for assistance.

Is the power supply Voltage/Current/Temperature/Fan Fault/Service Required LED amber?

Table 16 Diagnostics LED status: Rear panel power supply: “Voltage/Current/Temperature/Fan Fault/Service Required”

Answer	Possible reasons	Actions
No	System functioning properly.	No action required.
Yes	The power supply unit or a fan is operating at an unacceptable voltage/RPM level, or has failed.	<p>When isolating faults in the power supply, remember that the fans in both modules receive power through a common bus on the midplane, so if a power supply unit fails, the fans continue to operate normally.</p> <ul style="list-style-type: none">• Check that the power supply FRU is firmly locked into position.• Check that the power cable is connected to a power source.• Check that the power cable is connected to the power supply module.• Check proper ventilation (Refer to “Ventilation requirements” (page 76)).

Controller failure

Cache memory is flushed to CompactFlash in the case of a controller failure or power loss. During the write to CompactFlash process, only the components needed to write the cache to the CompactFlash are powered by the supercapacitor. This process typically takes 60 seconds per 1 Gbyte of cache. After the cache is copied to CompactFlash, the remaining power left in the supercapacitor is used to refresh the cache memory. While the cache is being maintained by the supercapacitor, the Cache Status LED flashes at a rate of 1/10 second on and 9/10 second off.

-
- ⓘ **IMPORTANT:** Transportable cache only applies to single-controller configurations. In dual controller configurations, there is no need to transport cache from a failed controller to a replacement controller because the cache is duplicated between the peer controllers (subject to volume cache optimization setting).
-

If the controller has failed or does not start, is the Cache Status LED on/blinking?

Table 17 Diagnostics LED status: Rear panel “Cache Status”

Answer	Actions
No, the Cache LED status is off, and the controller does not boot.	If valid data is thought to be in Flash, see Transporting cache ; otherwise, replace the controller module.
No, the Cache Status LED is off, and the controller boots.	The system has flushed data to disks. If the problem persists, replace the controller module.
Yes, at a strobe 1:10 rate - 1 Hz, and the controller does not boot.	See Transporting cache .
Yes, at a strobe 1:10 rate - 1 Hz, and the controller boots.	The system is flushing data to CompactFlash. If the problem persists, replace the controller module.
Yes, at a blink 1:1 rate - 1 Hz, and the controller does not boot.	See Transporting cache .
Yes, at a blink 1:1 rate - 1 Hz, and the controller boots.	The system is in self-refresh mode. If the problem persists, replace the controller module.

NOTE: See also “[Cache Status LED details](#)” (page 72).

Transporting cache

To preserve the existing data stored in the CompactFlash, you must transport the CompactFlash from the failed controller to a replacement controller using the procedure outlined in *HPE MSA Controller Module Replacement Instructions* shipped with the replacement controller module. Failure to use this procedure will result in the loss of data stored in the cache module.

CAUTION: Remove the controller module only after the copy process is complete, which is indicated by the Cache Status LED being off, or blinking at 1:10 rate.

Isolating a host-side connection fault

During normal operation, when a controller module host port is connected to a data host, the port’s host link status/link activity LED is green. If there is I/O activity, the LED blinks green. If data hosts are having trouble accessing the storage system, and you cannot locate a specific fault or cannot access the event logs, use the following procedure. This procedure requires scheduled downtime.

IMPORTANT: Do not perform more than one step at a time. Changing more than one variable at a time can complicate the troubleshooting process.

Host-side connection troubleshooting featuring host ports with SFPs

The procedure below applies to MSA 2050 SAN controller enclosures employing small form factor pluggable (SFP) transceiver connectors in 8/16 Gb FC, 10GbE iSCSI, or 1 Gb iSCSI host interface ports. In the following procedure, “SFP and host cable” is used to refer to any of the qualified SFP options supporting Converged Network Controller ports used for I/O or replication.

NOTE: When experiencing difficulty diagnosing performance problems, consider swapping out one SFP at a time to see if performance improves.

1. Halt all I/O to the storage system as described in [“Stopping I/O” \(page 47\)](#).
2. Check the host link status/link activity LED.
If there is activity, halt all applications that access the storage system.
3. Check the Cache Status LED to verify that the controller cached data is flushed to the disk drives.
 - o Solid – Cache contains data yet to be written to the disk.
 - o Blinking – Cache data is being written to CompactFlash.
 - o Flashing at 1/10 second on and 9/10 second off – Cache is being refreshed by the supercapacitor.
 - o Off – Cache is clean (no unwritten data).
4. Remove the SFP and host cable and inspect for damage.
5. Reseat the SFP and host cable.
Is the host link status/link activity LED on?
 - o Yes – Monitor the status to ensure that there is no intermittent error present. If the fault occurs again, clean the connections to ensure that a dirty connector is not interfering with the data path.
 - o No – Proceed to the next step.
6. Move the SFP and host cable to a port with a known good link status.
This step isolates the problem to the external data path (SFP, host cable, and host-side devices) or to the controller module port.
Is the host link status/link activity LED on?
 - o Yes – You now know that the SFP, host cable, and host-side devices are functioning properly. Return the SFP and cable to the original port. If the link status/link activity LED remains off, you have isolated the fault to the controller module port. Replace the controller module.
 - o No – Proceed to the next step.
7. Swap the SFP with the known good one.
Is the host link status/link activity LED on?
 - o Yes – You have isolated the fault to the SFP. Replace the SFP.
 - o No – Proceed to the next step.
8. Re-insert the original SFP and swap the cable with a known good one.
Is the host link status/link activity LED on?
 - o Yes – You have isolated the fault to the cable. Replace the cable.
 - o No – Proceed to the next step.
9. Verify that the switch, if any, is operating properly. If possible, test with another port.
10. Verify that the HBA is fully seated, and that the PCI slot is powered on and operational.
11. Replace the HBA with a known good HBA, or move the host side cable and SFP to a known good HBA.
Is the host link status/link activity LED on?
 - o Yes – You have isolated the fault to the HBA. Replace the HBA.
 - o No – It is likely that the controller module needs to be replaced.
12. Move the cable and SFP back to its original port.
Is the host link status/link activity LED on?
 - o No – The controller module port has failed. Replace the controller module.
 - o Yes – Monitor the connection for a period of time. It may be an intermittent problem, which can occur with damaged SFPs, cables, and HBAs.

Host-side connection troubleshooting featuring SAS host ports

The procedure below applies to MSA 2050 SAS controller enclosures employing 12 Gb SFF-8644 connectors in the HD mini-SAS host ports used for I/O.

1. Halt all I/O to the storage system as described in [“Stopping I/O” \(page 47\)](#).
2. Check the host activity LED.

If there is activity, halt all applications that access the storage system.

3. Check the Cache Status LED to verify that the controller cached data is flushed to the disk drives.
 - o Solid – Cache contains data yet to be written to the disk.
 - o Blinking – Cache data is being written to CompactFlash.
 - o Flashing at 1/10 second on and 9/10 second off – Cache is being refreshed by the supercapacitor.
 - o Off – Cache is clean (no unwritten data).
4. Reseat the host cable and inspect for damage.

Is the host link status LED on?

- o Yes – Monitor the status to ensure that there is no intermittent error present. If the fault occurs again, clean the connections to ensure that a dirty connector is not interfering with the data path.
 - o No – Proceed to the next step.
5. Move the host cable to a port with a known good link status.

This step isolates the problem to the external data path (host cable and host-side devices) or to the controller module port.

Is the host link status LED on?

- o Yes – You now know that the host cable and host-side devices are functioning properly. Return the cable to the original port. If the link status LED remains off, you have isolated the fault to the controller module port. Replace the controller module.
 - o No – Proceed to the next step.
6. Verify that the switch, if any, is operating properly. If possible, test with another port.
 7. Verify that the HBA is fully seated, and that the PCI slot is powered on and operational.
 8. Replace the HBA with a known good HBA, or move the host side cable to a known good HBA.

Is the host link status LED on?

- o Yes – You have isolated the fault to the HBA. Replace the HBA.
 - o No – It is likely that the controller module needs to be replaced.
9. Move the host cable back to its original port.

Is the host link status LED on?

- o No – The controller module port has failed. Replace the controller module.
- o Yes – Monitor the connection for a period of time. It may be an intermittent problem, which can occur with damaged cables and HBAs.

Isolating a controller module expansion port connection fault

During normal operation, when a controller module expansion port is connected to a drive enclosure, the expansion port status LED is green. If the connected port's expansion port LED is off, the link is down. Use the following procedure to isolate the fault.

This procedure requires scheduled downtime.

NOTE: Do not perform more than one step at a time. Changing more than one variable at a time can complicate the troubleshooting process.

1. Halt all I/O to the storage system as described in [“Stopping I/O” \(page 47\)](#).
2. Check the host activity LED.
If there is activity, halt all applications that access the storage system.
3. Check the Cache Status LED to verify that the controller cached data is flushed to the disk drives.
 - o Solid – Cache contains data yet to be written to the disk.
 - o Blinking – Cache data is being written to CompactFlash.
 - o Flashing at 1/10 second on and 9/10 second off – Cache is being refreshed by the supercapacitor.
 - o Off – Cache is clean (no unwritten data).
4. Reseat the expansion cable, and inspect it for damage.
Is the expansion port status LED on?
 - o Yes – Monitor the status to ensure there is no intermittent error present. If the fault occurs again, clean the connections to ensure that a dirty connector is not interfering with the data path.
 - o No – Proceed to the next step.
5. Move the expansion cable to a port on the controller enclosure with a known good link status.
This step isolates the problem to the expansion cable or to the controller module expansion port.
Is the expansion port status LED on?
 - o Yes – You now know that the expansion cable is good. Return the cable to the original port. If the expansion port status LED remains off, you have isolated the fault to the controller module expansion port. Replace the controller module.
 - o No – Proceed to the next step.
6. Move the expansion cable back to the original port on the controller enclosure.
7. Move the expansion cable on the drive enclosure to a known good expansion port on the drive enclosure.
Is the expansion port status LED on?
 - o Yes – You have isolated the problem to the drive enclosure port. Replace the expansion module.
 - o No – Proceed to the next step.
8. Replace the cable with a known good cable, ensuring the cable is attached to the original ports used by the previous cable.
Is the host link status LED on?
 - o Yes – Replace the original cable. The fault has been isolated.
 - o No – It is likely that the controller module must be replaced.

Isolating Remote Snap replication faults

Remote Snap replication is a licensed disaster-recovery feature that performs asynchronous replication of block-level data from a volume in a primary storage system to a volume in a secondary system. Remote Snap creates an internal snapshot of the primary volume, and copies changes to the data since the last replication to the secondary system via iSCSI or FC links. The primary volume exists in a primary pool in the primary storage system. Replication can be completed using either the SMU or CLI. See [“Connecting two storage systems to replicate volumes” \(page 32\)](#) for host connection information concerning Remote Snap.

Replication setup and verification

After storage systems and hosts are cabled for replication, you can use the SMU to prepare to use the Remote Snap feature. Optionally, you can use SSH to access the IP address of the controller module and access the Remote Snap feature using the CLI.

NOTE: Refer to the following manuals for more information about replication setup:

- See HPE Remote Snap technical white paper for replication best practices: : [MSA Remote Snap Software](#)
 - See *HPE MSA 1050/2050 SMU Reference Guide* for procedures to setup and manage replications
 - See *HPE MSA 1050/2050 CLI Reference Guide* for replication commands and syntax
 - See *HPE MSA Event Descriptions Reference Guide* for replication event reporting
-

Basic information for enabling the MSA 2050 SAN controller enclosures for replication supplements the troubleshooting procedures that follow.

- Familiarize yourself with replication content provided in the SMU Reference Guide.
- For best practices concerning replication-related tasks, see the technical white paper.
- In order to replicate an existing volume to a pool on the peer—ensure both the primary and secondary systems have the Remote Snap license installed—and follow these steps:
 - Find the port address on the secondary system:
Using the CLI, run the `show ports` command on the secondary system.
 - Verify that ports on the secondary system can be reached from the primary system using either method below:
 - Run the `query peer-connection` CLI command on the primary system, using a port address obtained from the output of the `show ports` command above.
 - In the SMU **Replications** topic, select **Action > Query Peer Connection**.
Ensure a pool exists on the secondary system.
 - Create a peer connection.
To create a peer connection, use the `create peer-connection` CLI command or in the SMU **Replications** topic, select **Action > Create Peer Connection**.
 - Create a replication set.
To create a replication set, use the `create replication-set` CLI command or in the SMU **Replications** topic, select **Action > Create Replication Set**.
 - Replicate.
To initiate replication, use the `replicate` CLI command or in the SMU **Replications** topic, select **Action > Replicate**.
- For descriptions of replication-related events, see the Event Descriptions Reference Guide.

Diagnostic steps for replication setup

The tables in this subsection show menu navigation for replication using the SMU.

-
- ❗ **IMPORTANT:** Remote Snap must be licensed on all systems configured for replication, and the controller module firmware must be compatible on all systems licensed for replication.
-

Can you successfully use the Remote Snap feature?

Answer	Possible reasons	Actions
Yes	System functioning properly.	No action required.
No	Remote Snap is not licensed on each controller enclosure used for replication.	<p>Verify licensing of the optional feature per system:</p> <ul style="list-style-type: none"> In the Home topic in the SMU, select Action > Install License. The License Settings panel opens and displays information about each licensed feature. If the Replication feature is not enabled, obtain and install a valid license for Remote Snap. For more information on licensing, see the “Installing a license” chapter in the SMU Reference Guide.
No	Compatible firmware revision supporting Remote Snap is not running on each system used for replication.	<p>Perform the following actions on each system used for replication:</p> <ul style="list-style-type: none"> In the System topic, select Action > Update Firmware. The Update Firmware panel opens. The Update Controller Modules tab shows firmware versions installed in each controller. If necessary, update the controller module firmware to ensure compatibility with other systems. For more information on compatible firmware, see the “Updating firmware” chapter in the SMU Reference Guide.
No	Invalid cabling connection. (If multiple controller enclosures are used, check the cabling for each system)	<p>Verify controller enclosure cabling.</p> <ul style="list-style-type: none"> Verify use of proper cables. Verify proper cabling paths for host connections. Verify cabling paths between replication ports and switches are visible to one another. Verify that cable connections are securely fastened. Inspect cables for damage and replace if necessary.
No	A system does not have a pool configured.	Configure each system to have a storage pool.

Table 18 Diagnostics for replication setup: Using Remote Snap feature

Can you create a replication set?

ⓘ **IMPORTANT:** Remote Snap must be licensed on all systems configured for replication, and the controller module firmware must be compatible on all systems licensed for replication.

Answer	Possible reasons	Actions
Yes	System functioning properly.	No action required.
No	On controller enclosures with iSCSI host interface ports, replication set creation fails due to use of CHAP.	If using CHAP (Challenge-Handshake Authentication Protocol), see the topics about configuring CHAP and working in replications within the SMU Reference Guide.
No	Unable to create the secondary volume (the destination volume on the pool to which you will replicate data from the primary volume)? ¹	<ul style="list-style-type: none"> Review event logs (in the footer, click the events panel and select Show Event List) for indicators of a specific fault in a replication data path component. Follow any recommended actions. Verify valid specification of the secondary volume according to either of the following criteria: <ul style="list-style-type: none"> A conflicting volume does not already exist Available free space in the pool

Table 19 Diagnostics for replication setup: Creating a replication set

Answer	Possible reasons	Actions
No	Communication link is down.	Review event logs for indicators of a specific fault in a host or replication data path component.
¹ After ensuring valid licensing, valid cabling connections, and network availability, create the replication set using the Replications topic; select Action > Create Replication Set .		

Table 19 Diagnostics for replication setup: Creating a replication set (continued)

Can you replicate a volume?

- ❗ **IMPORTANT:** Remote Snap must be licensed on all systems configured for replication, and the controller module firmware must be compatible on all systems licensed for replication.

Answer	Possible reasons	Actions
Yes	System functioning properly.	No action required.
No	Remote Snap is not licensed on each controller enclosure used for replication.	See actions described in “Can you successfully use the Remote Snap feature?” (page 57) .
No	Nonexistent replication set.	<ul style="list-style-type: none"> Determine existence of primary or secondary volumes. If a replication set has not been successfully created, use the Replications topic: select Action > Create Replication Set to create one. Review event logs (in the footer, click the events panel and select Show Event List) for indicators of a specific fault in a replication data path component. Follow any recommended actions.
No	Network error occurred during in-progress replication.	<ul style="list-style-type: none"> Review event logs for indicators of a specific fault in a replication data path component. Follow any recommended actions. Click in the Volumes topic, then click on a volume name in the volumes list. Click the Replication Sets tab to display replications and associated metadata. Replications that enter the suspended state can be resumed manually (see the SMU Reference Guide for additional information).
No	Communication link is down.	Review event logs for indicators of a specific fault in a host or replication data path component.

Table 20 Diagnostics for replication setup: Replicating a volume

Has a replication run successfully?

Answer	Possible reasons	Actions
Yes	System functioning properly.	No action required.
No	Last Successful Run shows N/A.	<ul style="list-style-type: none"> In the Volumes topic, click on the volume that is a member of the replication set. <ul style="list-style-type: none"> Select the Replication Sets table. Check the Last Successful Run information. If a replication has not run successfully, use the SMU to replicate as described in the section about working in the Replications topic within the SMU Reference Guide.
No	Communication link is down.	Review event logs for indicators of a specific fault in a host or replication data path component.

Table 21 Diagnostics for replication setup: Checking for a successful replication

Resolving voltage and temperature warnings

1. Check that all of the fans are working by making sure the Voltage/Fan Fault/Service Required LED on each power supply is off, or by using the SMU to check enclosure health status.
 - o In the lower corner of the footer, overall health status of the enclosure is indicated by a health status icon. For more information, point to the **System** tab and select **View System** to see the System panel. You can select from **Front**, **Rear**, and **Table** views on the System panel. If you point to a component, its associated metadata and health status displays onscreen.
See “[Options available for performing basic steps](#)” (page 45) for a description of health status icons and alternatives for monitoring enclosure health.
2. Make sure that all modules are fully seated in their slots with latches locked.
3. Make sure that no slots are left open for more than two minutes.
If you need to replace a module, leave the old module in place until you have the replacement or use a blank module to fill the slot. Leaving a slot open negatively affects the airflow and can cause the enclosure to overheat.
4. Make sure there is proper air flow, and no cables or other obstructions are blocking the front or rear of the array.
5. Try replacing each power supply module one at a time.
6. Replace the controller modules one at a time.
7. Replace SFPs one at a time (MSA 2050 SAN).

Sensor locations

The storage system monitors conditions at different points within each enclosure to alert you to problems. Power, cooling fan, temperature, and voltage sensors are located at key points in the enclosure. In each controller module and expansion module, the enclosure management processor (EMP) monitors the status of these sensors to perform SCSI enclosure services (SES) functions.

The following sections describe each element and its sensors.

Power supply sensors

Each enclosure has two fully redundant power supplies with load-sharing capabilities. The power supply sensors described in the following table monitor the voltage, current, temperature, and fans in each power supply. If the power supply sensors report a voltage that is under or over the threshold, check the input voltage.

Table 22 Power supply sensor descriptions

Description	Event/Fault ID LED condition
Power supply 1	Voltage, current, temperature, or fan fault
Power supply 2	Voltage, current, temperature, or fan fault

Cooling fan sensors

Each power supply includes two fans. The normal range for fan speed is 4,000 to 6,000 RPM. When a fan speed drops below 4,000 RPM, the EMP considers it a failure and posts an alarm in the storage system event log. The following table lists the description, location, and alarm condition for each fan. If the fan speed remains under the 4,000 RPM threshold, the internal enclosure temperature may continue to rise. Replace the power supply reporting the fault.

Table 23 Cooling fan sensor descriptions

Description	Location	Event/Fault ID LED condition
Fan 1	Power supply 1	< 4,000 RPM
Fan 2	Power supply 1	< 4,000 RPM
Fan 3	Power supply 2	< 4,000 RPM
Fan 4	Power supply 2	< 4,000 RPM

During a shutdown, the cooling fans do not shut off. This allows the enclosure to continue cooling.

Temperature sensors

Extreme high and low temperatures can cause significant damage if they go unnoticed. When a temperature fault is reported, it must be remedied as quickly as possible to avoid system damage. This can be done by warming or cooling the installation location.

Table 24 Controller platform temperature sensor descriptions

Description	Normal operating range	Warning operating range	Critical operating range	Shutdown values
CPU temperature (internal digital thermal sensor)	2°C–98°C	0°C–1°C, 99°C–104°C	None	≤ 0°C ≥ 104°C
SAS2008 internal digital sensor	3°C–112°C	0°C–2°C, 113°C–115°C	None	≤ 0°C ≥ 115°C
Supercapacitor pack thermistor	0°C–50°C	None	None	None
Onboard temperature 1	0°C–70°C	None	None	None
Onboard temperature 2	0°C–70°C	None	None	None
Onboard temperature 3	0°C–70°C	None	None	None

When a power supply sensor goes out of range, the Fault/ID LED illuminates amber and an event is logged.

Table 25 Power supply temperature sensor descriptions

Description	Normal operating range
Power Supply 1 temperature	–10°C–80°C
Power Supply 2 temperature	–10°C–80°C

Power supply module voltage sensors

Power supply voltage sensors ensure that the enclosure power supply voltage is within normal ranges. There are three voltage sensors per power supply.

Table 26 Voltage sensor descriptions

Sensor	Event/Fault LED condition
Power supply 1 voltage, 12V	< 11.00V > 13.00V
Power supply 1 voltage, 5V	< 4.00V > 6.00V
Power supply 1 voltage, 3.3V	< 3.00V > 3.80V

8 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/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>
 - Software Depot**
<https://www.hpe.com/support/softwaredepot>
- 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 HP 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 for your product, see the links provided below.

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

Additional warranty information

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.

A LED descriptions

Front panel LEDs

HPE MSA 2050 models support small form factor (SFF) and large form factor (LFF) enclosures. The SFF chassis, configured with 24 2.5" SFF disks, is used as a controller enclosure or drive enclosure. The LFF chassis, configured with 12 3.5" LFF disks, is used as either a controller enclosure or drive enclosure.

Enclosure bezel

The MSA 2050 enclosures are equipped with a removable bezel designed to cover the front panel during enclosure operation. The bezel assembly consists of a main body subassembly and two ear flange subassemblies, which attach the bezel to the left and right ear flanges of the 2U enclosure. See [Figure 1 \(page 11\)](#) for front view display of the bezel.

Enclosure bezel attachment




Figure 25 Partial exploded view showing bezel alignment with 2U chassis

Orient the enclosure bezel to align its back side with the front face of the enclosure as shown in [Figure 25](#). Face the front of the enclosure, and while supporting the base of the bezel—while grasping the left and right ear covers—position the bezel such that the mounting sleeves within the integrated ear covers align with the ball studs on the ear flanges (see [Figure 26](#)). Gently push-fit the bezel onto the ball studs to securely attach it to the front of the enclosure.




Figure 26 Detail views of enclosure ear cover mounting sleeves

Enclosure bezel removal


TIP: Please refer to [Figure 25](#) (bezel front) and [Figure 26](#) (bezel back) on [page 64](#) before removing the bezel from the enclosure front panel.

You may need to remove the bezel to access front panel components such as disk drives and ear kits. Although disk drive LEDs are not visible when the bezel is attached, you can monitor disk behavior from the management interfaces (see [“Fault isolation methodology”](#) ([page 45](#)) for more information about using LEDs together with event notification, the CLI, and the SMU for managing the storage system).

While facing the front of the enclosure, grasp the left and right ear covers, such that your fingers cup the bottom of each ear cover, with thumb at the top of each cover. Gently pull the top of the bezel while applying slight inward pressure below, to release the bezel from the ball studs.

NOTE: The bezel should be attached to the enclosure during operation to protect ear circuitry. To reattach the bezel to the enclosure front panel, follow the instructions provided in [“Enclosure bezel attachment”](#) ([page 64](#)).

MSA 2050 Array SFF or supported 24-drive expansion enclosure



Notes:

Integers on disks indicate drive slot numbering sequence.


The enlarged detail view at right shows LED icons from the bezel that correspond to chassis LEDs.

The detail view locator circle (above right) identifies the ear kit that connects to LED light pipes in the bezel (or ear cover).

LED	Description	Definition
1	Enclosure ID	Green — On 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	Disk drive Online/Activity	See “Disk drive LEDs” (page 67).
3	Disk drive Fault/UID	See “Disk drive LEDs” (page 67).
4	Unit Identification (UID)	Blue — Identified. Off — Identity LED off.
5	Heartbeat	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.
6	Fault ID	Amber — Fault condition exists. The event has been identified, but the problem needs attention. Off — No fault condition exists.

Figure 27 LEDs: MSA 2050 Array SFF or supported 24-drive expansion enclosure: front panel

MSA 2050 Array LFF or supported 12-drive expansion enclosure



Notes:

Integers on disks indicate drive slot numbering sequence.

The enlarged detail view at right shows LED icons from the bezel that correspond to chassis LEDs.


The detail view locator circle (above right) identifies the ear kit that connects to LED light pipes in the bezel (or ear cover).

LED	Description	Definition
1	Enclosure ID	Green — On 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	Disk drive Online/Activity	See “Disk drive LEDs” (page 67).
3	Disk drive Fault/UID	See “Disk drive LEDs” (page 67).
4	Unit Identification (UID)	Blue — Identified. Off — Identity LED off.
5	Heartbeat	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.
6	Fault ID	Amber — Fault condition exists. The event has been identified, but the problem needs attention. Off — No fault condition exists.

Figure 28 LEDs: MSA 2050 Array LFF or supported 12-drive expansion enclosure: front panel

Ear covers


Ear covers can be used instead of the enclosure bezel. Figure 29 callouts apply to the table for Figure 28. See Figure 26 (page 64) for detail views of mounting sleeve attachment to ball studs located on left and right ears.



Callout numbers pertain to chassis LED descriptions in the table above.

Figure 29 Ear covers option to enclosure bezel

Disk drive LEDs



Online/Activity (green)	Fault/UID (amber/blue)	Description
On	Off	Normal operation. The disk drive is online, but it is not currently active.
Blinking irregularly	Off	The disk drive is active and operating normally.
Off	Amber; blinking regularly (1 Hz)	Offline: the disk is not being accessed. A predictive failure alert may have been received for this device. Further investigation is required.
On	Amber; blinking regularly (1 Hz)	Online: possible I/O activity. A predictive failure alert may have been received for this device. Further investigation is required.
Blinking irregularly	Amber; blinking regularly (1 Hz)	The disk drive is active, but a predictive failure alert may have been received for this disk. Further investigation is required.
Off	Amber; solid ^{1,2}	Offline: no activity. A failure or critical fault condition has been identified for this disk.
Off	Blue; solid	Offline: the disk drive has been selected by a management application such as the SMU.
On or blinking	Blue; solid	The controller is driving I/O to the disk, and it has been selected by a management application such as the SMU.
Blinking regularly (1 Hz)	Off	<p>⚠ CAUTION: Do not remove the disk drive. Removing a disk may terminate the current operation and cause data loss. The disk is reconstructing.</p>
Off	Off	Either there is no power, the drive is offline, or the drive is not configured.

¹This Fault/UID state can indicate that the disk is a leftover. The fault may involve metadata on the disk rather than the disk itself. See the Clearing disk metadata topic in the SMU Reference Guide or online help.

²This Fault/UID state can indicate an FDE-related issue. The fault may indicate an inserted FDE disk that is in a locked state, or it may indicate insertion of a non-FDE disk into an FDE-secured system. See also "FDE considerations" (page 17) and the **Important** statement at the end of that section.


Figure 30 LEDs: Disk drive combinations — enclosure front panel

ⓘ IMPORTANT: For information about self-encrypting disk (SED) drives, see "FDE considerations" (page 17) and the SMU Reference Guide or online help.

Rear panel LEDs

Controller enclosure—rear panel layout

The diagram and table below display and identify important component items comprising the rear panel layout of the MSA 2050 controller enclosure (MSA 2050 SAN configured with SFPs is shown in the example). Diagrams and tables on the following pages further describe rear panel LED behavior for component field-replaceable units.



MSA 2050 controller enclosure (rear panel locator illustration)

- | | |
|--|--|
| 1 AC Power supplies [see Figure 35 (page 73)] | 8 Network management port |
| 2 Controller module A [see Figure 32 (page 69)] | 9 Service port 1 (used by service personnel only) |
| 3 Controller module B [see Figure 32 (page 69)] | 10 Disabled button (used by engineering only) (Stickers shown covering the openings) |
| 4 Host ports: used for host connection or replication | 11 SAS expansion port |
| 5 CLI port (USB - Type B) | 12 DC Power supply (2) — (DC model only) |
| 6 Service port 2 (used by service personnel only) | 13 DC Power switch [see Figure 35 (page 73)] |
| 7 Reserved for future use | |

Figure 31 MSA 2050 SAN Array: rear panel


A controller enclosure accommodates two power supply FRUs of the same type—either both AC or both DC—within the two power supply slots (see two instances of callout 1 above). The controller enclosure accommodates two controller module FRUs of the same type within the I/O module slots (see callouts 2 and 3 above).

IMPORTANT: MSA 2050 controller enclosures support dual-controller only. If a partner controller fails, the array will fail over and run on a single controller until the redundancy is restored. A controller module must be installed in each IOM slot to ensure sufficient airflow through the enclosure during operation.

The diagrams with tables that immediately follow provide descriptions of the different controller modules and power supply modules that can be installed into the rear panel of an MSA 2050 controller enclosure. Showing controller modules and power supply modules separately from the enclosure provides improved clarity in identifying the component items called out in the diagrams and described in the tables.

Descriptions are also provided for optional drive enclosures supported by MSA 2050 controller enclosures for expanding storage capacity.

MSA 2050 SAN controller module—rear panel LEDs



LED	Description	Definition
1	Host 8/16 Gb FC ¹ Link Status/ Link Activity	Off — No link detected. Green — The port is connected and the link is up. Blinking green — The link has I/O or replication activity.
2	Host 10GbE iSCSI ^{2,3} Link Status/ Link Activity	Off — No link detected. Green — The port is connected and the link is up. Blinking green — The link has I/O or replication activity.
3	Network Port Link Active Status ⁴	Off — The Ethernet link is not established, or the link is down. Green — The Ethernet link is up (applies to all negotiated link speeds).
4	Network Port Link Speed ⁴	Off — Link is up at 10/100base-T negotiated speeds. Amber — Link is up and negotiated at 1000base-T.
5	OK to Remove	Off — The controller module is not prepared for removal. Blue — The controller module is prepared for removal.
6	Unit Locator	Off — Normal operation. Blinking white — Physically identifies the controller module.
7	FRU OK	Off — Controller module is not OK. Blinking green — System is booting. Green — Controller module is operating normally.
8	Fault/Service Required	Amber — A fault has been detected or a service action is required. Blinking amber — Hardware-controlled power-up or a cache flush or restore error.
9	Cache Status	Green — Cache contains unwritten data and operation is normal. The unwritten information can be log or debug data that remains in the cache, so a Green cache status LED does not, by itself, indicate that any user data is at risk or that any action is necessary. Off — In a working controller, cache is clean (contains no unwritten data). This is an occasional condition that occurs while the system is booting. Blinking green — A CompactFlash flush or cache self-refresh is in progress, indicating cache activity. See also Cache Status LED details .
10	Expansion Port Status	Off — The port is empty or the link is down. On — The port is connected and the link is up.

¹When in FC mode, the SFPs must be a qualified 8 Gb or 16 Gb option as described in the QuickSpecs. A 16 Gb/s SFP can run at 16 Gb/s, 8 Gb/s, 4 Gb/s, or auto-negotiate its link speed. An 8 Gb/s SFP can run at 8 Gb/s, 4 Gb/s, or auto-negotiate its link speed.


²When in 10GbE iSCSI mode, the SFPs must be qualified 10GbE iSCSI options. These, and qualified 10GbE DAC cables are described in the QuickSpecs.

³When powering up and booting, iSCSI LEDs will be on/blinking momentarily, then they will switch to the mode of operation.

⁴When port is down, both LEDs are off.

Figure 32 LEDs: MSA 2050 SAN controller module (FC and 10GbE SFPs)

NOTE: See “MSA 2050 SAN” (page 9) for information about supported combinations of host interface protocols using Converged Network Controller ports.



LED	Description	Definition
1	Not used in example ¹	The FC SFP is not show in this example [see Figure 32 (page 69)].
2	Host 1 Gb iSCSI ^{2,3} Link Status/ Link Activity	Off — No link detected. Green — The port is connected and the link is up; or the link has I/O or replication activity.
3	Network Port Link Active Status ⁴	Off — The Ethernet link is not established, or the link is down. Green — The Ethernet link is up (applies to all negotiated link speeds).
4	Network Port Link Speed ⁴	Off — Link is up at 10/100base-T negotiated speeds. Amber — Link is up and negotiated at 1000base-T.
5	OK to Remove	Off — The controller module is not prepared for removal. Blue — The controller module is prepared for removal.
6	Unit Locator	Off — Normal operation. Blinking white — Physically identifies the controller module.
7	FRU OK	Off — Controller module is not OK. Blinking green — System is booting. Green — Controller module is operating normally.
8	Fault/Service Required	Amber — A fault has been detected or a service action is required. Blinking amber — Hardware-controlled power-up or a cache flush or restore error.
9	Cache Status	Green — Cache contains unwritten data and operation is normal. The unwritten information can be log or debug data that remains in the cache, so a Green cache status LED does not, by itself, indicate that any user data is at risk or that any action is necessary. Off — In a working controller, cache is clean (contains no unwritten data). This is an occasional condition that occurs while the system is booting. Blinking green — A CompactFlash flush or cache self-refresh is in progress, indicating cache activity. See also Cache Status LED details .
10	Expansion Port Status	Off — The port is empty or the link is down. On — The port is connected and the link is up.

¹When in FC mode, the SFPs must be a qualified 8 Gb or 16 Gb option as described in the QuickSpecs.


²When in 1 Gb iSCSI mode, the SFPs must be a qualified RJ-45 iSCSI option as described in the QuickSpecs. The 1 Gb iSCSI mode does not support an iSCSI optic option.

³When powering up and booting, iSCSI LEDs will be on/blinking momentarily, then they will switch to the mode of operation.

⁴When port is down, both LEDs are off.

Figure 33 LEDs: MSA 2050 SAN controller module (1 Gb RJ-45 SFPs)

MSA 2050 SAS controller module—rear panel LEDs



LED	Description	Definition
1	Host SAS ¹ Link Status	Off — No link detected. Green — The port is connected and the link is up.
2	Host SAS ¹ Link Activity	Off — The link is idle. Blinking green — The link has I/O activity.
3	Network Port Link Active Status ²	Off — The Ethernet link is not established, or the link is down. Green — The Ethernet link is up (applies to all negotiated link speeds).
4	Network Port Link Speed ²	Off — Link is up at 10/100base-T negotiated speeds. Amber — Link is up and negotiated at 1000base-T.
5	OK to Remove	Off — The controller module is not prepared for removal. Blue — The controller module is prepared for removal.
6	Unit Locator	Off — Normal operation. Blinking white — Physically identifies the controller module.
7	FRU OK	Off — Controller module is not OK. Blinking green — System is booting. Green — Controller module is operating normally.
8	Fault/Service Required	Amber — A fault has been detected or a service action is required. Blinking amber — Hardware-controlled power-up or a cache flush or restore error.
9	Cache Status	Green — Cache contains unwritten data and operation is normal. The unwritten information can be log or debug data that remains in the cache, so a Green cache status LED does not, by itself, indicate that any user data is at risk or that any action is necessary. Off — In a working controller, cache is clean (contains no unwritten data). This is an occasional condition that occurs while the system is booting. Blinking green — A CompactFlash flush or cache self-refresh is in progress, indicating cache activity. See also Cache Status LED details .
10	Expansion Port Status	Off — The port is empty or the link is down. On — The port is connected and the link is up.

¹See the qualified HD mini-SAS host cable options described in the QuickSpecs.

²When port is down, both LEDs are off.

Figure 34 LEDs: MSA 2050 SAS controller module (SFF-8644 12 Gb SAS connector)

NOTE: Once a Link Status LED is lit, it remains so, even if the controller is shutdown via the SMU or CLI.

When a controller is shutdown or otherwise rendered inactive—its Link Status LED remains illuminated—falsely indicating that the controller can communicate with the host. Though a link exists between the host and the chip on the

controller, the controller is not communicating with the chip. To reset the LED, the controller must be properly power-cycled [see “Powering on/powering off” (page 23)].

Cache Status LED details

Power on/off behavior

The storage enclosure's unified CPLD provides integrated Power Reset Management (PRM) functions. During power on, discrete sequencing for power on display states of internal components is reflected by blinking patterns displayed by the Cache Status LED (see Table 27).

Table 27 Cache Status LED – power on behavior

Item	Display states reported by Cache Status LED during power on sequence							
	0	1	2	3	4	5	6	7
Display state	0	1	2	3	4	5	6	7
Component	VP	SC	SAS BE	ASIC	Host	Boot	Normal	Reset
Blink pattern	On 1/Off 7	On 2/Off 6	On 3/Off 5	On 4/Off 4	On 5/Off 3	On 6/Off 2	Solid/On	Steady

Once the enclosure has completed the power on sequence, the Cache Status LED displays Solid/On (Normal), before assuming the operating state for cache purposes.

Cache status behavior


If the LED is blinking evenly, a cache flush is in progress. When a controller module loses power and write cache contains data that has not been written to disk, the supercapacitor pack provides backup power to flush (copy) data from write cache to CompactFlash memory. When cache flush is complete, the cache transitions into self-refresh mode.

If the LED is blinking momentarily slowly, the cache is in a self-refresh mode. In self-refresh mode, if primary power is restored before the backup power is depleted (3–30 minutes, depending on various factors), the system boots, finds data preserved in cache, and writes it to disk. This means the system can be operational within 30 seconds, and before the typical host I/O time-out of 60 seconds, at which point system failure would cause host-application failure. If primary power is restored after the backup power is depleted, the system boots and restores data to cache from CompactFlash, which can take about 90 seconds. The cache flush and self-refresh mechanism is an important data protection feature; essentially four copies of user data are preserved: one in controller cache and one in CompactFlash of each controller. The Cache Status LED illuminates solid green during the boot-up process. This behavior indicates the cache is logging all POSTs, which will be flushed to the CompactFlash the next time the controller shuts down.

CAUTION: If the Cache Status LED illuminates solid green—and you wish to shut-down the controller—do so from the user interface, so unwritten data can be flushed to CompactFlash.

Power supply LEDs

Power redundancy is achieved through two independent load-sharing power supplies. In the event of a power supply failure, or the failure of the power source, the storage system can operate continuously on a single power supply. Greater redundancy can be achieved by connecting the power supplies to separate circuits. DC power supplies are equipped with a power switch. AC power supplies do not have a power switch. Whether a power supply has a power switch is significant to powering on/off. Power supplies are used by controller and drive enclosures.




LED	Description	Definition
1	Input Source Power Good	Green — Power is on and input voltage is normal. Off — Power is off or input voltage is below the minimum threshold.
2	Voltage/Current/Temperature/Fan Fault/Service Required	Amber — Voltage (input or output), Current, Temperature is out of range or fan failure. Off — Output voltage is normal.

Figure 35 LEDs: MSA 2050 Storage system enclosure power supply modules

NOTE: See “Powering on/powering off” (page 23) for information about power-cycling enclosures.

MSA 2050 LFF and SFF drive enclosures—rear panel layout

MSA 2050 controllers support the 3.5" 12-drive enclosure and the 2.5" 24-drive enclosure for adding storage. The front panel of the 12-drive enclosure looks identical to the MSA 2050 Array LFF front panel. The front panel of the 24-drive enclosure looks identical to the MSA 2050 Array SFF front panel. The rear panel of the MSA 2050 LFF Disk Enclosure (12-drive) and the MSA 2050 SFF Disk Enclosure (24-drive) enclosures are identical, as shown below.



LED	Description	Definition
1	Power supply LEDs	See “Power supply LEDs” (page 72) .
2	Unit Locator	Off — Normal operation. Blinking white— Physically identifies the expansion module.
3	OK to Remove	Not implemented.
4	Fault/Service Required	Amber — A fault has been detected or a service action is required. Blinking amber — Hardware-controlled powerup or a cache flush or restore error.
5	FRU OK	Green — Expansion module is operating normally. Blinking green — System is booting. Off — Expansion module is not OK.
6	SAS In Port Status	Green — Port link is up and connected. Off — Port is empty or link is down.
7	SAS Out Port Status	Green — Port link is up and connected. Off — Port is empty or link is down.

Figure 36 LEDs: MSA 2050 3.5" 12-drive or 2.5" 24-drive enclosure rear panel

B Specifications and requirements

Safety requirements

Install the system in accordance with the local safety codes and regulations at the facility site. Follow all cautions and instructions marked on the equipment. Also, refer to the documentation included with your product ship kit.

Site requirements and guidelines

The following sections provide requirements and guidelines that you must address when preparing your site for the installation.

When selecting an installation site for the system, choose a location not subject to excessive heat, direct sunlight, dust, or chemical exposure. These conditions greatly reduce the system's longevity and might void your warranty.

Site wiring and AC power requirements

The following are required for all installations using AC power supplies:

- All AC mains and supply conductors to power distribution boxes for the rack-mounted system must be enclosed in a metal conduit or raceway when specified by local, national, or other applicable government codes and regulations.
- Ensure that the voltage and frequency of your power source match the voltage and frequency inscribed on the equipment's electrical rating label.
- To ensure redundancy, provide two separate power sources for the enclosures. These power sources must be independent of each other, and each must be controlled by a separate circuit breaker at the power distribution point.
- The system requires voltages within minimum fluctuation. The customer-supplied facilities' voltage must maintain a voltage with not more than ± 5 percent fluctuation. The customer facilities must also provide suitable surge protection.
- Site wiring must include an earth ground connection to the AC power source. The supply conductors and power distribution boxes (or equivalent metal enclosure) must be grounded at both ends.
- Power circuits and associated circuit breakers must provide sufficient power and overload protection. To prevent possible damage to the AC power distribution boxes and other components in the rack, use an external, independent power source that is isolated from large switching loads (such as air conditioning motors, elevator motors, and factory loads).

NOTE: For power requirements, see the QuickSpecs: <https://www.hpe.com/support/MSA2050QuickSpecs>. If a website location has changed, an Internet search for "HPE MSA 2050 quickspecs" will provide a link.

Site wiring and DC power requirements

The following are required for all installations using DC power supplies:

- All DC mains and supply conductors to power distribution boxes for the rack-mounted system must comply with local, national, or other applicable government codes and regulations.
- Ensure that the voltage of your power source matches the voltage inscribed on the equipment's electrical label.
- To ensure redundancy, provide two separate power sources for the enclosures. These power sources must be independent of each other, and each must be controlled by a separate circuit breaker at the power distribution point.
- The system requires voltages within minimum fluctuation. The customer-supplied facilities' voltage must maintain a voltage within the range specified on the equipment's electrical rating label. The customer facilities must also provide suitable surge protection.
- Site wiring must include an earth ground connection to the DC power source. Grounding must comply with local, national, or other applicable government codes and regulations.
- Power circuits and associated circuit breakers must provide sufficient power and overload protection.

Weight and placement guidelines

Refer to “[Physical requirements](#)” (page 77) for detailed size and weight specifications.

- The weight of an enclosure depends on the number and type of modules installed.
- Ideally, use two people to lift an enclosure. However, one person can safely lift an enclosure if its weight is reduced by removing the power supply modules and disk drive modules.
- Do not place enclosures in a vertical position. Always install and operate the enclosures in a horizontal/level orientation.
- When installing enclosures in a rack, make sure that any surfaces over which you might move the rack can support the weight. To prevent accidents when moving equipment, especially on sloped loading docks and up ramps to raised floors, ensure you have a sufficient number of helpers. Remove obstacles such as cables and other objects from the floor.
- To prevent the rack from tipping, and to minimize personnel injury in the event of a seismic occurrence, securely anchor the rack to a wall or other rigid structure that is attached to both the floor and to the ceiling of the room.

Electrical guidelines

- These enclosures work with single-phase power systems having an earth ground connection. To reduce the risk of electric shock, do not plug an enclosure into any other type of power system. Contact your facilities manager or a qualified electrician if you are not sure what type of power is supplied to your building.
- Enclosures are shipped with a grounding-type (three-wire) power cord. To reduce the risk of electric shock, always plug the cord into a grounded power outlet.
- Do not use household extension cords with the enclosures. Not all power cords have the same current ratings. Household extension cords do not have overload protection and are not meant for use with computer systems.

Ventilation requirements

Refer to “[Environmental requirements](#)” (page 78) for detailed environmental requirements.

- Do not block or cover ventilation openings at the front and rear of an enclosure. Never place an enclosure near a radiator or heating vent. Failure to follow these guidelines can cause overheating and affect the reliability and warranty of your enclosure.
- Leave a minimum of 15 cm (6 inches) at the front and back of each enclosure to ensure adequate airflow for cooling. No cooling clearance is required on the sides, top, or bottom of enclosures.
- Leave enough space in front and in back of an enclosure to allow access to enclosure components for servicing. Removing a component requires a clearance of at least 37 cm (15 inches) in front of and behind the enclosure.

Cabling requirements

- Keep power and interface cables clear of foot traffic. Route cables in locations that protect the cables from damage.
- Route interface cables away from motors and other sources of magnetic or radio frequency interference.
- Stay within the cable length limitations.

Management host requirements

A local management host with at least one USB Type B port connection is recommended for the initial installation and configuration of a controller enclosure. After you configure one or both of the controller modules with an Internet Protocol (IP) address, you then use a remote management host on an Ethernet network to configure, manage, and monitor.

NOTE: Connections to this device must be made with shielded cables—grounded at both ends—with metallic RFI/EMI connector hoods, in order to maintain compliance with NEBS and FCC Rules and Regulations.

Physical requirements

The floor space at the installation site must be strong enough to support the combined weight of the rack, controller enclosures, drive enclosures (expansion), and any additional equipment. The site also requires sufficient space for installation, operation, and servicing of the enclosures, together with sufficient ventilation to allow a free flow of air to all enclosures.

Table 28 and Table 29 list enclosure dimensions and weights. Weights are based on an enclosure having a full complement of disk drives, two controller or expansion modules, and two power supplies installed. “2U12” denotes the LFF enclosure (12 disks) and “2U24” denotes the SFF enclosure (24 disks).

Table 29 provides weight data for MSA 2050 controller enclosures and drive enclosures. For information about other HPE MSA drive enclosures that may be cabled to these systems, check the QuickSpecs at:

<https://www.hpe.com/support/MSA2050QuickSpecs>. If a website location has changed, an Internet search for “HPE MSA 2050 quickspecs” will provide a link.

Table 28 Rackmount enclosure dimensions

Specifications	Rackmount
2U Height (y-axis)	8.9 cm (3.5 inches)
Width (x-axis):	
• Chassis only	44.7 cm (17.6 inches)
• Chassis with bezel ear caps	47.9 cm (18.9 inches)
Depth (z-axis):	
SFF drive enclosure (2U24)	
• Back of chassis ear to controller latch	50.5 cm (19.9 inches)
• Front of chassis ear to back of cable bend	57.9 cm (22.8 inches)
LFF drive enclosure (2U12)	
• Back of chassis ear to controller latch	60.2 cm (23.7 inches)
• Front of chassis ear to back of cable bend	67.1 cm (26.4 inches)

Table 29 Rackmount enclosure weights

Specifications	Rackmount
MSA 2050 SAN Array SFF Enclosure	8.6 kg (19.0 lb) [chassis]
• Chassis with FRUs (no disks) ^{1,2}	19.9 kg (44.0 lb)
• Chassis with FRUs (including disk) ^{1,3}	25.4 kg (56.0 lb)
MSA 2050 SAN Array LFF Enclosure	9.9 kg (22.0 lb) [chassis]
• Chassis with FRUs (no disks) ^{1,2}	21.3 kg (47.0 lb)
• Chassis with FRUs (including disks) ^{1,3}	30.8 kg (68.0 lb)
MSA 2050 SFF Disk Enclosure	8.6 kg (19.0 lb) [chassis]
• Chassis with FRUs (no disks) ^{1,2}	19.9 kg (44.0 lb)
• Chassis with FRUs (including disks) ^{1,3}	25.4 kg (56.0 lb)
MSA 2050 LFF Disk Enclosure	9.9 kg (22.0 lb) [chassis]
• Chassis with FRUs (no disks) ^{1,2}	21.3 kg (47.0 lb)
• Chassis with FRUs (including disks) ^{1,3}	30.8 kg (68.0 lb)

¹Weights shown are nominal, and subject to variances.

²Weights may vary due to different power supplies, IOMs, and differing calibrations between scales.

³Weights may vary due to actual number and type of disk drives (SAS or SSD) installed.

Environmental requirements

NOTE: For operating and non-operating environmental technical specifications, see the QuickSpecs at: <https://www.hpe.com/support/MSA2050QuickSpecs>. If a website location has changed, an Internet search for “HPE MSA 2050 quickspecs” will provide a link.

Electrical requirements

Site wiring and power requirements

Each enclosure has two power supply modules for redundancy. If full redundancy is required, use a separate power source for each module. The AC power supply unit in each power supply module is auto-ranging and is automatically configured to an input voltage range from 100–240 VAC with an input frequency of 50–60 Hz. The power supply modules meet standard voltage requirements for both U.S. and international operation. The power supply modules use standard industrial wiring with line-to-neutral or line-to-line power connections.

Power cord requirements

Each enclosure is equipped with two power supplies of the same type (both AC or both DC). For enclosures equipped with AC power supply modules, use two power cords that are appropriate for use in a typical outlet in the destination country. Whether using AC or DC power supplies, each power cable connects one of the power supplies to an independent, external power source. To ensure power redundancy, connect the two suitable power cords to two separate circuits: for example, to one commercial circuit and one uninterruptible power source (UPS).

① **IMPORTANT:** See the QuickSpecs for information about power cables provided with your MSA 2050 Storage product. If a website location has changed, an Internet search for “HPE MSA 2050 quickspecs” will provide a link.

C Electrostatic discharge

Preventing electrostatic discharge

To prevent damaging the system, be aware of the precautions you need to follow when setting up the system or handling parts. A discharge of static electricity from a finger or other conductor may damage system boards or other static-sensitive devices. This type of damage may reduce the life expectancy of the device.

To prevent electrostatic damage:

- Avoid hand contact by transporting and storing products in static-safe containers.
- Keep electrostatic-sensitive parts in their containers until they arrive at static-protected workstations.
- Place parts in a static-protected area before removing them from their containers.
- Avoid touching pins, leads, or circuitry.
- Always be properly grounded when touching a static-sensitive component or assembly.

Grounding methods to prevent electrostatic discharge

Several methods are used for grounding. Use one or more of the following methods when handling or installing electrostatic-sensitive parts:

- Use a wrist strap connected by a ground cord to a grounded workstation or computer chassis. Wrist straps are flexible straps with a minimum of 1 megohm (± 10 percent) resistance in the ground cords. To provide proper ground, wear the strap snug against the skin.
- Use heel straps, toe straps, or boot straps at standing workstations. Wear the straps on both feet when standing on conductive floors or dissipating floor mats.
- Use conductive field service tools.
- Use a portable field service kit with a folding static-dissipating work mat.

If you do not have any of the suggested equipment for proper grounding, have an authorized reseller install the part. For more information on static electricity or assistance with product installation, contact an authorized reseller.

D SFP option for host ports

Locate the SFP transceivers

Locate the qualified SFP options for your MSA 2050 SAN controller module within your product ship kit. You can also obtain the part numbers using the QuickSpecs.

The SFP transceiver (SFP) should look similar to the generic SFP shown in the figure below. Follow the guidelines provided in [“Electrostatic discharge” \(page 79\)](#) when installing an SFP.




Figure 37 Install a qualified SFP option

TIP: See the “Configuring host ports” topic within the SMU Reference Guide for information about configuring MSA 2050 SAN host ports with host interface protocols of the same type or a combination of types. Also see [“Using the CLI port and cable—known issues on Windows” \(page 43\)](#).

Install an SFP transceiver

For each target MSA 2050 SAN host port, perform the following procedure to install an SFP. Refer to the figure above when performing the steps.

1. Orient the SFP as shown above, and align it for insertion into the target host port.
The SFP should be positioned such that the actuator pivot-hinge is on top.
2. If the SFP has a plug, remove it before installing the transceiver. Retain the plug.
3. Flip the actuator open as shown in the figure (near the left detail view).
The actuator on your SFP option may look slightly different than the one shown, and it may not open to a sweep greater than 90° (as shown in the figure).
4. Slide the SFP into the target host port until it locks into place.
5. Flip the actuator down, as indicated by the down-arrow next to the open actuator in the figure.
The installed SFP should look similar to the position shown in the right detail view.

6. When ready to attach to the host, obtain and connect a qualified cable option to the duplex jack at the end of the SFP connector.

NOTE: To remove an SFP module, perform the above steps in *reverse* order.

Verify component operation

View the host port Link Status/Link Activity LED on the controller module face plate. A green LED indicates that the port is connected and the link is up (see [LED descriptions](#) for information about controller module LEDs).

Index

Numerics

- 2U12
 - large form factor (LFF) enclosure 77
- 2U24
 - small form factor (SFF) enclosure 77

A

- accessing
 - CLI (command-line interface) 39

C

- cables
 - 10GbE iSCSI 29
 - 1Gb iSCSI 29
 - Ethernet 32
 - FCC compliance statement 32, 76
 - Fibre Channel 29
 - routing requirements 76
 - shielded 32, 76
 - USB for CLI 40
- cabling
 - connecting controller and drive enclosures 18
 - direct attach configurations 29
 - switch attach configurations 31
 - to enable Remote Snap replication 32
- cache
 - read ahead 15
 - self-refresh mode 72
 - write-through 15
- clearance requirements
 - service 76
 - ventilation 76
- command-line interface (CLI)
 - connecting USB cable to CLI port 40
 - using to set controller IP addresses 40
- CompactFlash
 - memory card 16
 - transporting 52
- components
 - MSA 2050
 - enclosure front panel
 - LFF enclosure 12
 - SFF enclosure 11
 - enclosure rear panel
 - AC power supply 13
 - DC power supply 13
 - supported drive enclosures
 - LFF drive enclosure 15
 - MSA 2050 SAN
 - enclosure rear panel
 - AC power supply 68
 - CLI port (USB - Type B) 14, 68

- DC power supply 68
- DC power switch 68
- host ports 14, 68
- mini-SAS expansion port 14
- network management port 14
- network port 68
- reserved port 14, 68
- SAS expansion port 68
- service port 1 14, 68
- service port 2 14, 68

- configuring
 - direct attach configurations 29
 - switch attach configurations 31
- connections
 - verify 23
- console requirement 76
- controller enclosures
 - connecting to data hosts 27
 - connecting to remote management hosts 32

D

- data hosts
 - defined 27
 - optional software 27
 - system requirements 27
- DHCP
 - server 39
- disk drive
 - slot numbering
 - LFF enclosure 12
 - SFF enclosure 11

E

- electromagnetic compatibility (EMC) 75
- electrostatic discharge
 - grounding methods 79
 - precautions 79
- enclosure
 - cabling 18
 - dimensions 77
 - IDs, correcting 47
 - input frequency requirement 78
 - input voltage requirement 78
 - installation checklist 17
 - site requirements 77
 - troubleshooting 47
 - web-browser based configuring and provisioning 44
 - weight 77
- Ethernet cables
 - requirements 32

F

faults

isolating

- expansion port connection fault 55
- host-side connection 52
- methodology 45

H

host interface ports

FC host interface protocol

- loop topology 28
- point-to-point protocol 28

iSCSI host interface protocol

- 1 Gb 29
- 10GbE 28
- mutual CHAP 28, 29

SFP transceivers 27

hosts

- defined 27
- stopping I/O 47

I

IDs, correcting for enclosure 47

installing enclosures

- installation checklist 17

IP addresses

- setting using CLI 39
- setting using DHCP 39

L

LEDs

- disk drives 67
- enclosure front panel
 - Enclosure ID 65, 66
 - Fault ID 65, 66
 - Heartbeat 65, 66
 - Unit Identification (UID) 65, 66
- enclosure rear panel
 - MSA 2050 SAN
 - 10GbE iSCSI Host Link Status/Link Activity 69
 - 1Gb iSCSI Host Link Status/Link Activity 70
 - Cache Status 69, 70
 - Expansion Port Status 69, 70
 - Fault/Service Required 69, 70
 - FC Host Link Status/Link Activity 69
 - FRU OK 69, 70
 - Network Port Link Active 69, 70
 - Network Port Link Speed 69, 70
 - OK to Remove 69, 70
 - Unit Locator 69, 70
 - MSA 2050 SAS
 - 6/12 Gb Host Link Activity 71
 - 6/12 Gb Host Link Status 71
 - Cache Status 71

Expansion Port Status 71

Fault/Service Required 71

FRU OK 71

Network Port Link Active 71

Network Port Link Speed 71

OK to Remove 71

Unit Locator 71

power supply unit

Input Source Power Good 73

Voltage/Fan Fault/Service Required 73

supported drive enclosures (expansion)

LFF enclosure rear panel

Fault/Service Required 74

FRU OK 74

OK to Remove 74

power supply 74

SAS In Port Status 74

SAS Out Port Status 74

Unit Locator 74

local management host requirement 76

P

physical requirements 77

power cord requirements 78

power cycle

- power off 24, 26
- power on 24, 26

power supply

- AC power requirements 75
- DC power requirements 75
- site wiring requirements 75

R

regulatory compliance

notices

- shielded cables 32, 76

requirements

- cabling 19
- clearance 76
- Ethernet cables 32
- host system 27
- physical 77
- ventilation 76

RFI/EMI connector hoods 32, 76

S

safety precautions 75

sensors

- locating 59
- power supply 59
- temperature 60
- voltage 60

- site planning
 - EMC 75
 - local management host requirement 76
 - physical requirements 77
 - safety precautions 75
- SMU 9
 - accessing web-based management interface 44
 - defined 44
 - getting started 44
 - Remote Snap replication 32, 55
 - storage system configuring and provisioning 44
- storage system setup
 - configuring 44
 - provisioning 44
 - replicating 44
- supercapacitor pack 16

T

- troubleshooting 45
 - controller failure, single controller configuration 51
 - correcting enclosure IDs 47
 - enclosure does not initialize 47
 - expansion port connection fault 55
 - host-side connection fault 52, 54
 - Remote Snap replication faults 55
 - using event notification 46
 - using system LEDs 48
 - using the CLI 46
 - using the SMU 45

V

- ventilation requirements 76

W

- warnings
 - voltage and temperature 59