SEAGATE

5005/4005/3005 Series
Hardware Installation and Maintenance Guide

Abstract

This document describes initial hardware setup for Seagate 5005/4005/3005 Series enclosures. It also describes removal and installation of customer-replaceable units for these enclosures. The document is intended for use by storage system administrators familiar with servers and computer networks, network administration, storage system administration and configurations, storage area network management, and relevant protocols.

Contents

Figures

Tables

About this guide

Introduction

This guide provides information about initial hardware setup, and removal and installation of customer-replaceable units (CRUs) for the Seagate 5005/4005/3005 Series controller enclosures, which can be used with J1212/J1224/J1284 expansion enclosures. In controller enclosures, 5005/4005 Series controller modules can be purchased with either of two types of host interface: SAS or FC/iSCSI. 3005 Series controller modules can be purchased only with an FC/iSCSI host interface. An FC/iSCSI controller module's host ports use converged network adapters, and is informally referred to as a converged network controller (CNC):

- FC/iSCSI controller enclosure:
    - Qualified Fibre Channel SFP option supporting (8/16Gb)
    - Qualified Internet SCSI (10GbE) SFP option
    - Qualified Internet SCSI (1Gb) Copper RJ-45 SFP option
- SAS (12Gb) controller enclosure

The 5005/4005/3005 Series controller enclosures are SBB-compliant (Storage Bridge Bay). These enclosures support large form factor disks or small form factor disks in 2U12 and 2U24 chassis. The 5005/4005 Series enclosures also support large form factor disks or small form factor disks in 5U84 chassis. The 2U12, 2U24, and 5U84 chassis form factors support controller enclosures and expansion enclosures.

The 5005/4005/3005 Series controller enclosures can optionally be cabled to supported expansion enclosures for adding storage. See also for maximum configuration information.

ⓘ **IMPORTANT:**   Product configuration characteristics for controller and optional expansion enclosures:

- 5005/4005 Series supports 2U and 5U chassis form factors that are configurable as either an RBOD, EBOD, or JBOD.
- 3005 Series supports 2U chassis form factors that are configurable as either an RBOD, EBOD, or JBOD.
- 5005/4005 Series provides four host ports per controller I/O module (SBB RAID canister).
- 3005 Series provides two host ports per controller I/O module.
- 5005/4005/3005 Series models are configured with two controller modules per enclosure.
- 5005/4005/3005 Series supports intermixing of 2U12 and 2U24 EBODs/JBODS; however, 2U EBODs cannot be intermixed with 5U84 EBODs within the same storage system.
- Optional J1224/J1212/J1284 models are configured with two expansion I/O modules (SBB expansion canisters) per enclosure. These models support EBOD and JBOD configurations.

These enclosures support virtual storage, which uses paged-storage technology. For virtual storage, a group of disks with an assigned RAID level is called a disk group.

Enclosure user interfaces

The 5005/4005/3005 Series enclosures support applications for configuring, monitoring, and managing the storage system. The web-based application GUI and the command-line interface are briefly described:

- RBOD/EBOD configuration: Storage Management Console (SMC) is the web interface for the enclosure, providing access to all common management functions for virtual storage.
- RBOD/EBOD configuration: The command-line interface (CLI) enables you to interact with the storage system using command syntax entered via the keyboard or scripting.
- JBOD configuration: See information about firmware support in .

NOTE: For more information about enclosure user interfaces, see the following:
- Storage Management Guide or online help - the guide describes the Storage Management Console GUI
- CLI Reference Guide

CNC ports used for host connection

Certain models use Converged Network Controller (CNC) 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 use the CLI to set all controller host ports to use either FC or iSCSI protocol using the `set host-port-mode` CLI command. 5005/4005/3005 controller enclosures support the following link speeds:

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

Alternatively, for 4-port 5005/4005 models you can use the CLI to set CNC ports to support a combination of host interface protocols. When configuring a combination of host interface protocols, host ports 0 and 1 are set to FC (either both 16Gb/s or both 8Gb/s), and host ports 2 and 3 must be set to iSCSI (either both 10GbE or both 1Gb/s), provided the CNC ports use qualified SFP connectors and cables required for supporting the selected host interface protocol.

The 2-port 3005 models do not support SFPs for multiple host interface protocols in combination. You must select a common host interface protocol and SFP for use in all CNC ports within the controller enclosure.

See "CNC technology" (page 66) and Figure 18 (page 29) for more information.

☼ **TIP:**   See the CLI Reference Guide for information about configuring CNC ports with host interface protocols of the same type or a combination of types.

CNC controller modules ship with CNC ports initially configured for FC. When connecting CNC ports to iSCSI hosts, you must use the CLI (not the SMC) to specify which ports will use iSCSI. It is best to do this before inserting the iSCSI SFPs into the CNC ports (see "Change the CNC port mode" (page 82) for instructions).

HD mini-SAS ports used for host connection

5005/4005 enclosures provide eight high-density mini-SAS ports (4-ports per controller module). The HD mini-SAS host interface protocol uses the SFF-8644 external connector interface defined for SAS3.0 to support a link rate of 12Gb/s using the qualified connectors and cable options. See Figure 19 (page 29) for more information.

Intended audience

This guide is intended for system administrators and storage administrators.

Prerequisites

Prerequisites for using this product include knowledge of:

- Server system administration
- Microsoft Windows servers
- Storage system configuration
- Storage area network (SAN) management and server-attached storage
- Fibre Channel (FC) protocol
- Serial Attached SCSI (SAS) protocol
- Internet SCSI (iSCSI) protocol

Related documentation

Table 1   Related documents

For information about	See
Enhancements, known issues, and late-breaking information not included in product documentation	Release Notes
Regulatory compliance and safety and disposal information	Product Regulatory Compliance and Safety*
Overview of hardware installation	Getting Started*
Obtaining and installing a license to use licensed features	Obtaining and Installing a License
Using the web interface to configure and manage the product	Storage Management Guide
Using the command-line interface (CLI) to configure and manage the product	CLI Reference Guide
Event codes and recommended actions	Event Descriptions Reference Guide

*Printed document included in product shipkit.

For additional information, see seagate.com/support.

Document conventions and symbols

Table 2   Document conventions

Convention	Element
Colored text	Cross-reference links
Black, underlined text	Email addresses
Colored, underlined text	Website addresses
Bold text	- Keys that are pressed - Text entered into a GUI element, such as a box - GUI elements that are clicked or selected, such as menu and list items, buttons, and check boxes
Italic text	Text emphasis
`Monospace` text	- File and directory names - System output - Code - Commands, their arguments, and argument values
`Monospace, italic` text	- Code variables - Command parameters
`Monospace, bold` text	Emphasis of file and directory names, system output, code, and text entered at the command line

△ **CAUTION:**   Indicates that failure to follow directions could result in damage to equipment or data.

ⓘ **IMPORTANT:**   Provides clarifying information or specific instructions.

NOTE: Provides additional information.

☼ **TIP:** Provides helpful hints and shortcuts.

1   Safety guidelines

Safe handling

⚠ **CAUTION:**   Use this equipment in a manner specified by the manufacturer: failure to do this may cancel the protection provided by the equipment.

- Permanently unplug the enclosure before you move it or if you think that it has become damaged in any way.
- A safe lifting height is 20U.
- Always remove the Power and Cooling Modules (PCMs) to minimize weight before you move the enclosure.
- Do not lift the enclosures by the handles on the PCMs—they are not designed to take the weight.

⚠ **CAUTION:**   Do not try to lift the enclosure by yourself:

- Fully configured 2U12 enclosures can weigh up to 32 kg (71 lb).
- Fully configured 2U24 enclosures can weigh up to 30 kg (66 lb).
- Fully configured 5U84 enclosures can weigh up to 135 kg (298 lb). An unpopulated enclosure weighs 46 kg (101 lb). Use appropriate lifting methods. Before lifting the enclosure:
  - Unplug all cables and label them for reconnection.
  - Remove DDICs from drawers and verify the drawers are closed and firmly locked. See also Figure 27 (page 34).
  - Use a minimum of three people to lift the enclosure using the lifting straps provided.
  - Avoid lifting the enclosure using the handles on any of the CRUs because they are not designed to take the weight.
  - Do not lift the enclosure higher than 20U. Use mechanical assistance to lift above this height.
- Observe the lifting hazard label affixed to the storage enclosure.

Operation

ⓘ **IMPORTANT:**   Operation of the enclosure with any CRU modules missing will disrupt the airflow, and the enclosure will not receive sufficient cooling. It is essential that all slots hold modules before the enclosure system is used. Empty drive slots (bays) in 2U enclosures must hold dummy drive carrier modules. See also Figure 43 (page 50).

- Observe the module bay caution label affixed to the module being replaced.
- Replace a defective PCM with a fully operational PCM as soon as possible. Do not remove a defective PCM unless you have a replacement model of the correct type ready for insertion.
- Before removal/replacement of a PCM or PSU, disconnect supply power from the module to be replaced. Please refer to "Replacing a power cooling module" (page 117).
- Observe the hazardous voltage warning label affixed to power cooling modules.

⚠ **CAUTION:**   5U84 enclosures only

- To prevent overturning, drawer interlocks stop users from opening both drawers at the same time. Do not attempt to force open a drawer when the other drawer in the enclosure is already open. In a rack containing more than one 5U84 enclosure, do not open more than one drawer per rack at a time.
- Observe the hot surface label affixed to the drawer. Operating temperatures inside enclosure drawers can reach 60ºC. Take care when opening drawers and removing DDICs.
- Due to product acoustics, ear protection should be worn during prolonged exposure to the product in operation.
- Observe the drawer caution label. Open drawers must not be used to support any other objects or equipment.

Electrical safety

- The 2U enclosure must only be operated from a power supply input voltage range of 100–240 VAC, 50–60 Hz.
- The 5U enclosure must only be operated from a power supply input voltage range of 200–240 VAC, 50–60 Hz.
- Provide a suitable power source with electrical overload protection to meet the technical specification requirements.
- The power cord must have a safe electrical earth connection. Check the connection to earth of the enclosure before you switch on the power supply.

ⓘ **IMPORTANT:**   The enclosure must be grounded before applying power.

- The plug on the power supply cord is used as the main disconnect device. Ensure that the socket outlets are located near the equipment and are easily accessible.
- 2U enclosures are intended to operate with two PCMs.
- 5U84 enclosures are intended to operate with two PSUs.
- Observe the power-supply disconnection caution label affixed to power cooling modules.

⚠ **CAUTION:**

- Do not remove the covers from the enclosure or any of the modules—there is a danger of electric shock inside. Return any damaged components to your supplier for repair.
- When bifurcated power cords (Y-leads) are used, these cords must only be connected to a power supply range of 200–240 VAC.

ⓘ **IMPORTANT:**   The RJ-45 socket on IOMs is for the Ethernet connection only and must not be connected to a telecommunications network.

Rack system safety precautions

The following safety requirements must be considered when the enclosure is mounted in a rack.

- The rack construction must be capable of supporting the total weight of the installed enclosures. The design should incorporate stabilizing features suitable to prevent the rack from tipping or being pushed over during installation or in normal use.
- When loading a rack with enclosures, fill the rack from the bottom up; and empty the rack from the top down.
- Always remove all power supply modules to minimize weight, before loading the enclosure into the rack.
- Do not try to lift the enclosure by yourself.

⚠ **CAUTION:**   To avoid danger of the rack falling over, under no circumstances should more than one enclosure be moved out of the cabinet at any one time.

- The system must be operated with low pressure rear exhaust installation. The back pressure created by rack doors and obstacles is not to exceed 5 pascals (0.5 mm water gauge).
- The rack design should take into consideration the maximum operating ambient temperature for the enclosure, which is 35ºC (95ºF) for RBODs and 40ºC (104ºF) for EBODs and JBODs.
- The rack should have a safe electrical distribution system. It must provide over-current protection for the enclosure and must not be overloaded by the total number of enclosures installed in the rack. When addressing these concerns, consideration should be given to the electrical power consumption rating shown on the nameplate.
- The electrical distribution system must provide a reliable earth connection for each enclosure in the rack.
- Each PSU or PCM in each enclosure has an earth leakage current of 1.0mA. The design of the electrical distribution system must take into consideration the total earth leakage current from all the PSUs/PCMs in all the enclosures. The rack will require labeling with "High Leakage Current. Earth connection essential before connecting supply."

- The rack, when configured with the enclosures, must meet the safety requirements of UL 60950-1 and IEC 60950-1.

2   System overview

Enclosure configurations

The storage system supports three controller enclosure configurations.

- 2U (rack space) controller enclosure – see Figure 1 (page 21) and Figure 2 (page 21): holds up to 12 low profile (1-inch high) 3.5" form factor disk drive modules in a horizontal orientation.
- 2U (rack space) controller enclosure – see Figure 3 (page 22) and Figure 4 (page 22): holds up to 24 low profile (5/8 inch high) 2.5" form factor disk drive modules in a vertical orientation.
- 5U (rack space) controller enclosure – see Figure 5 (page 23) and Figure 6 (page 23): holds up to 84 low profile (1-inch high) 3.5" form factor disk drive modules in a vertical orientation within the disk drawer. Two vertically stacked drawers each hold 42 disks. If used, 2.5" disks require 3.5" adapters.

These same chassis form factors are used for supported expansion enclosures; albeit with different I/O modules (IOMs). Each individual disk drive is hot pluggable and replaceable on site.

Expansion enclosures can be configured as either an EBOD or a JBOD. An EBOD is an expansion enclosure attached to a controller enclosure (RBOD). A JBOD is an expansion enclosure attached directly to a server that is equipped with a SAS HBA or SAS RAID adapter.

NOTE: Throughout this guide—and the management interfaces documents used with this guide—I/O module (IOM) is a general term denoting either a controller module (RAID canister) or an expansion module (expansion canister).

The enclosure configurations—including chassis and CRUs—are described on the following pages.

Cache

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

- Write-back caching. The controller writes user data into the cache memory in the controller module rather than directly to the disks. 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 disks.
- Read-ahead caching. The controller detects sequential data access, reads ahead into the next sequence of data—based upon settings—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.

☼ **TIP:**   See the Storage Management Guide for more information about cache options and settings for RBOD/EBOD configurations.

CompactFlash

During a power loss or controller failure, data stored in cache is saved off to non-volatile memory (CompactFlash). The data is restored to cache, and 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. Do not remove the card; it is used for cache recovery only.

NOTE: 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 controller module cache in case of power failure, each controller enclosure model is 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. In the event of power failure, while cache is maintained by the supercapacitor pack, the Cache Status LED blinks at a rate of 1/10 second on and 9/10 second off. See also "Cache Status LED details" (page 46).

NOTE: Please refer to "Enclosure variants" (page 23) for details about various 2U12 enclosure options.
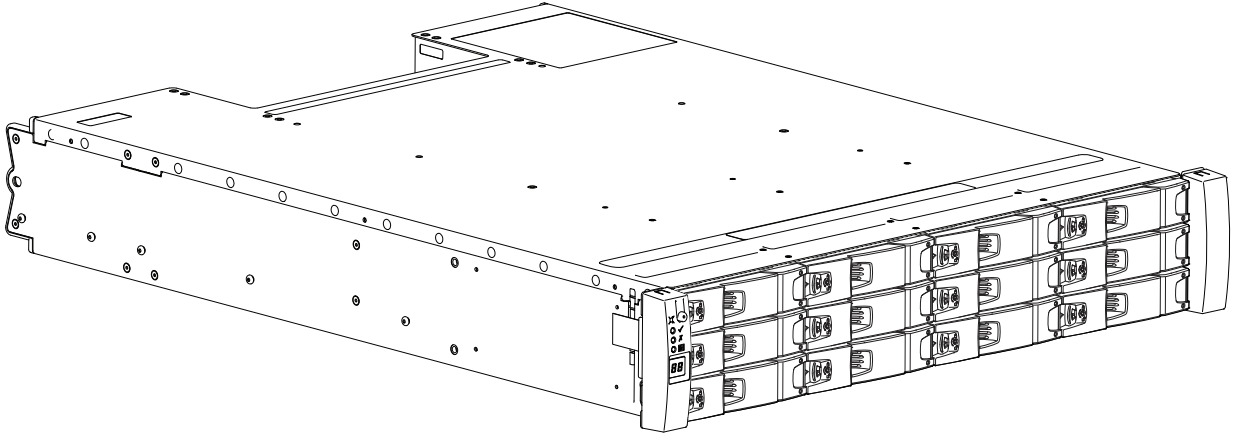
Figure 1 2U12 enclosure system – dimetric front orientation
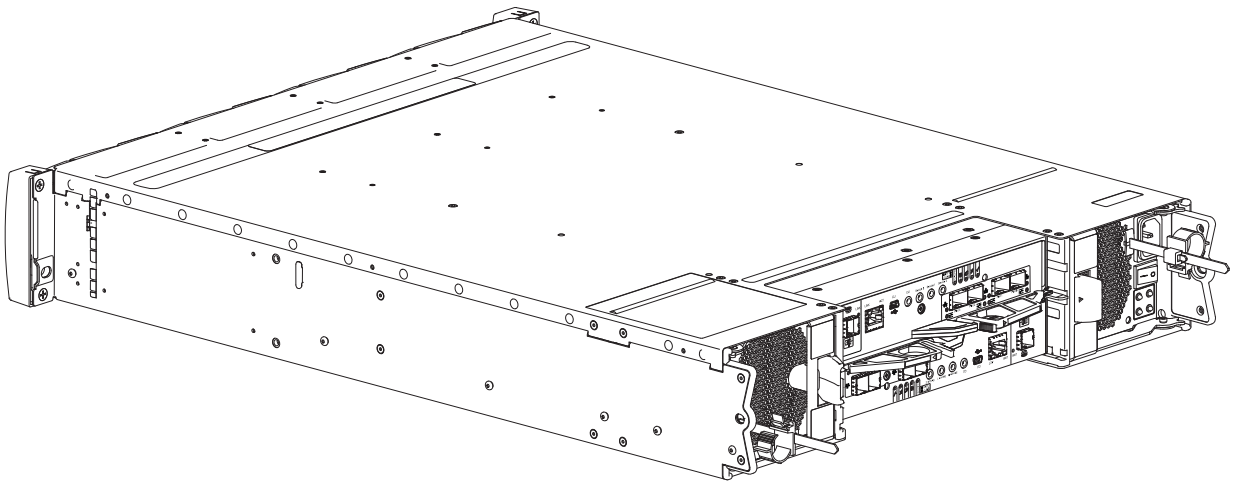
Figure 2 2U12 enclosure system – dimetric rear orientation

The 2U12 controller enclosure above is equipped with dual-controllers (4-port FC/iSCSI model shown).

NOTE: Please refer to "Enclosure variants" (page 23) for details about various 2U24 enclosure options.
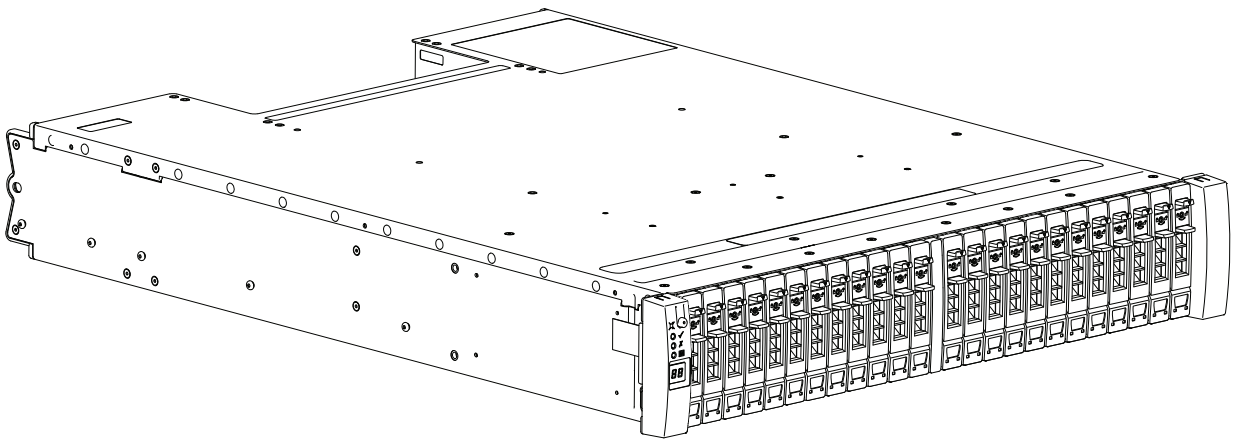
Figure 3 2U24 enclosure system – dimetric front orientation
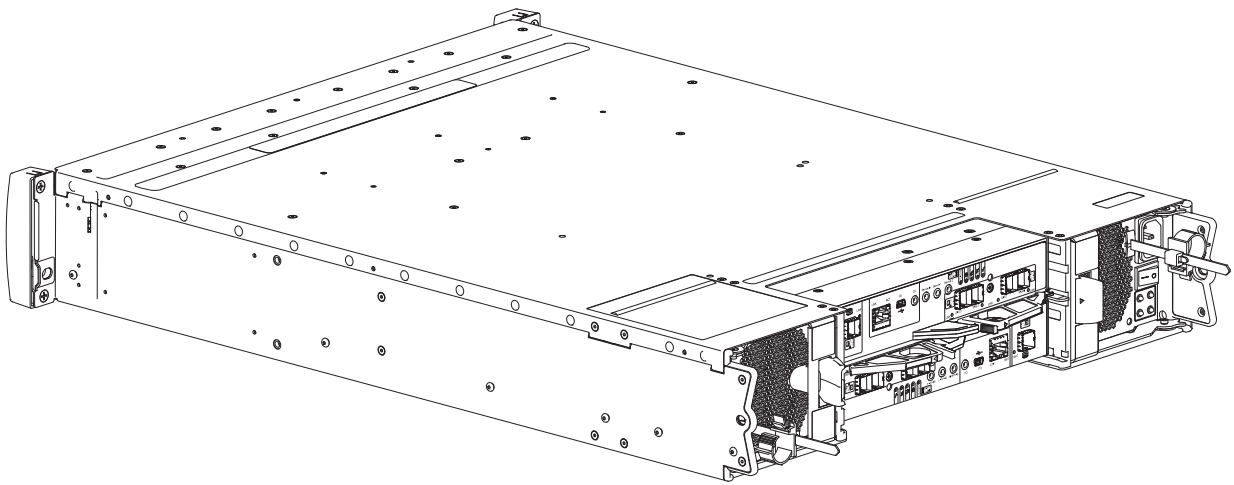
Figure 4 2U24 enclosure system – dimetric rear orientation

The 2U24 controller enclosure above is equipped with dual-controllers (4-port SAS model shown).

NOTE: Please refer to "Enclosure variants" (page 23) for details about various 5U84 enclosure options.
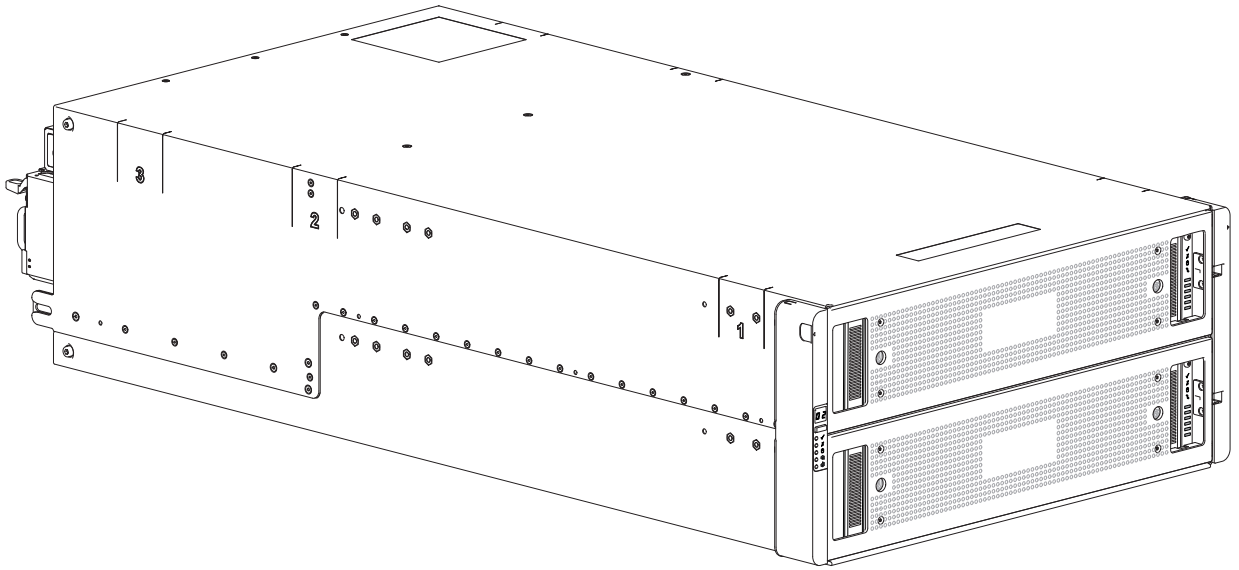
Figure 5 5U84 enclosure system – dimetric front orientation
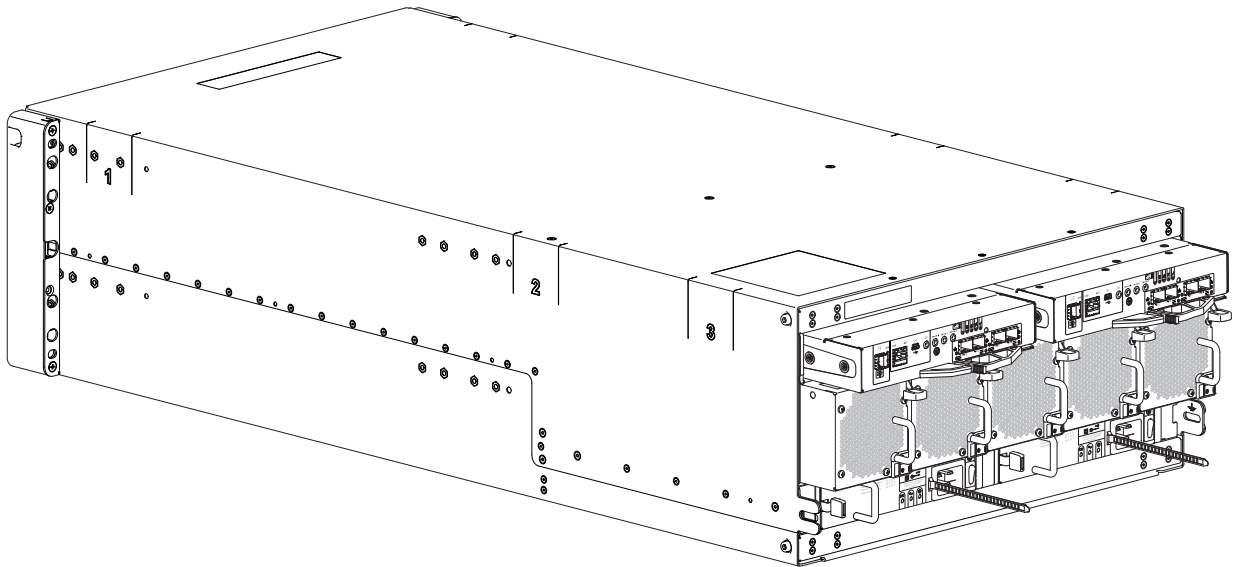
Figure 6 5U84 enclosure system – dimetric rear orientation

The 5U84 controller enclosure above is equipped with dual-controllers (4-port FC/iSCSI model shown).

Enclosure variants

The 2U chassis can be configured as a controller enclosure (5005, 4005, 3005) or an expansion enclosure (J1212/J1224) as shown in Table 3 and Table 4. The 5U chassis can be configured as a controller enclosure (5005, 4005) or an expansion enclosure (J1284) as shown in Table 5.

ⓘ **IMPORTANT:**   The 2U and 5U core products–including key components and CRUs–are described herein. Although many CRUs differ between the form factors, the IOMs are common to 2U12, 2U24, and 5U84 chassis. The IOMs are introduced in "2U enclosure core product" and cross-referenced from "5U enclosure core product".

2U12

12 x LFF (Large Form Factor) disk drives and 12 x HFF (Hybrid Form Factor) disk drives.

Table 3   2U12 enclosure variants

Product	Configuration	PCMs[1]	IOMs[2]
5005	12Gb/s direct dock LFF SAS	2	2
4005	12Gb/s direct dock LFF SAS	2	2
3005	12Gb/s direct dock LFF SAS	2	2
J1212	12Gb/s direct dock LFF SAS	2	2

1-Redundant PCMs must be compatible modules of the same type (both AC).

2-Supported controller module IOMs include 2-port and 4-port FC/iSCSI, and 4-port HD mini-SAS.
  Supported expansion module IOMs are used in expansion enclosures for adding storage.

2U24

24 x SFF (Small Form Factor) disk drives.

Table 4   2U24 enclosure variants

Product	Configuration	PCMs[1]	IOMs[2]
5005	12Gb/s direct dock SFF SAS	2	2
4005	12Gb/s direct dock SFF SAS	2	2
3005	12Gb/s direct dock LFF SAS	2	2
J1224	12Gb/s direct dock SFF SAS	2	2

1-Redundant PCMs must be compatible modules of the same type (both AC).

2-Supported controller module IOMs include 2-port and 4-port FC/iSCSI, and 4-port HD mini-SAS.
  Supported expansion module IOMs are used in expansion enclosures for adding storage.

5U84

84 x LFF or SFF disk drives (not interchangeable) held in two 42-slot vertically stacked drawers.

Table 5   5U84 enclosure variants

Product	Configuration	PSUs[1]	FCMs[2]	IOMs[3]
5005	12Gb/s direct dock SFF SAS	2	5	2
4005	12Gb/s direct dock SFF SAS	2	5	2
J1284	12Gb/s direct dock SFF SAS	2	5	2

1-Redundant PSUs must be compatible modules of the same type (both AC).

2-The fan control module (FCM) is a separate CRU (not integrated into a PCM).

3-Supported controller module IOMs include 4-port FC/iSCSI and 4-port HD mini-SAS.
  Supported expansion module IOMs are used in expansion enclosures for adding storage.

NOTE: The enclosure core product definitions follow.

- See "2U enclosure core product" for the 2U12 and 2U24 form factor.
- See "5U enclosure core product" for the 5U84 form factor.
- IOMs are used in 2U and 5U chassis and are described in "Controller modules" and "Expansion module" (page 28).

ⓘ **IMPORTANT:**   The enclosures support dual-controller configuration only. If a partner controller fails, the storage system will fail over and run on a single controller module until the redundancy is restored. For RBODs, EBODs, and JBODs, an IOM must be installed in each IOM slot to ensure sufficient air flow through the enclosure during operation.

2U enclosure core product

The design concept is based on an enclosure subsystem together with a set of plug-in modules. A typical enclosure system—as supplied—includes the following:

- An enclosure chassis which includes the midplane PCB and an integral operator's (Ops) panel that is mounted on the left ear flange at the front of the enclosure.
- Two 580W, 100–240V AC power cooling modules. See also Figure 30 (page 37).
- Two IOMs: 2 x SBB-compliant interface slots.
- Up to 24 disk drive modules. Where appropriate the disk drive carriers will include an Interposer card. See also "Enclosure variants" (page 23). Dummy drive carrier modules must be installed in all empty drive slots.
- A rail kit for rack mounting.

NOTE: The module quantities quoted above are the maximum that a 2U24 enclosure can support. The following figures show component locations–together with CRU slot indexing–relative to 2U enclosure front and rear panels.

2U enclosure front panel

Integers on disks indicate drive slot numbering sequence.
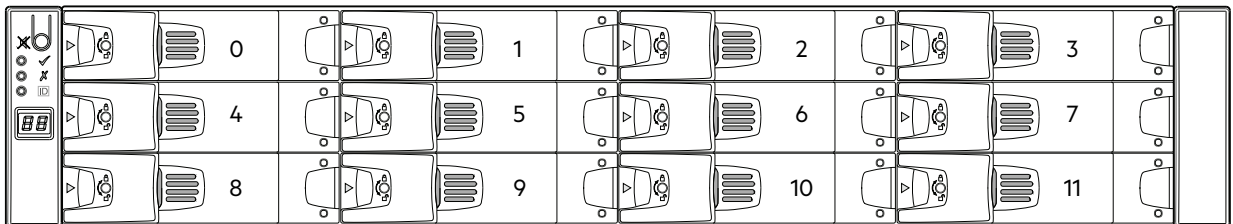
Figure 7   2U12 enclosure system – front panel components
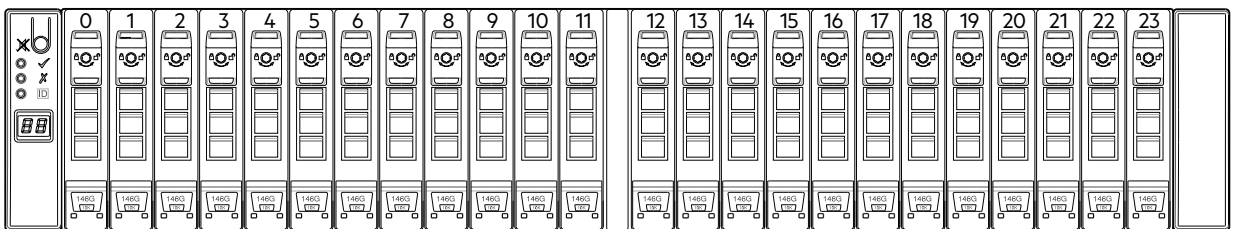
Figure 8   2U24 enclosure system – front panel components

2U enclosure rear panel

Numeric designators on PCMs and alphabetic designators on IOMs indicate slot sequencing for modules used in 2U enclosures. PCM and IOM modules are available as CRUs. The 5005 and 4005 RBODs use 4-port controller modules, whereas the 3005 RBODs use 2-port controller modules. The 5005/4005/3005 Series RBODs support the J1212/J1224/J1284 EBODs for optionally adding storage. The J1212/J1224/J1284 expansion enclosures can also be attached to server HBAs as JBODs. RAID controls reside on the server.
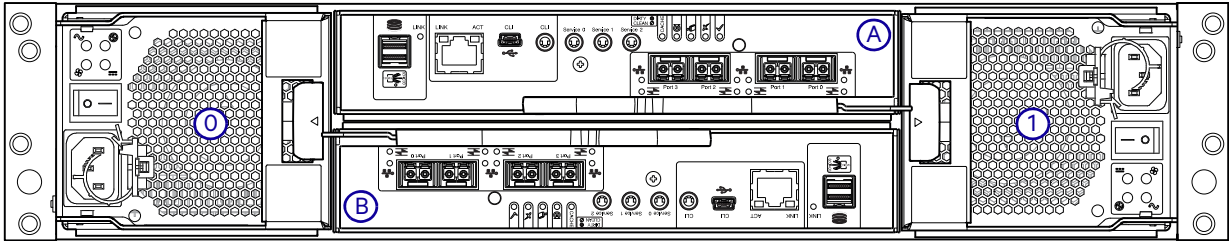
Figure 9 2U controller enclosure – rear panel components (4-port FC/iSCSI)
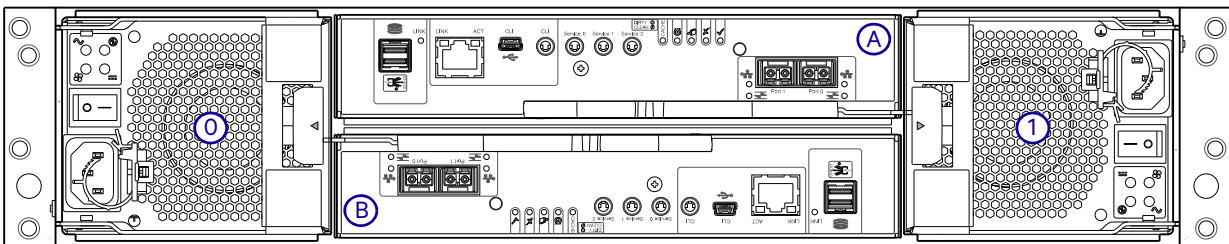
Figure 10 2U controller enclosure – rear panel components (2-port FC/iSCSI)
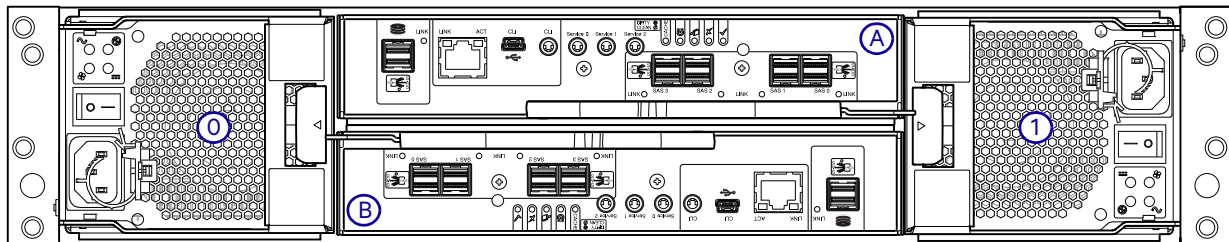
Figure 11 2U controller enclosure – rear panel components (4-port SAS)
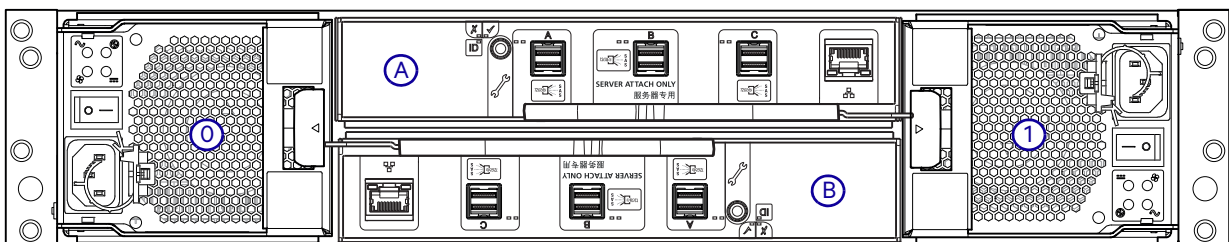
Figure 12 2U expansion enclosure – rear panel components

2U rear panel components

Controller modules

The top slot for holding IOMs is A and the bottom slot is B. The face plate details of the IOMs show the modules aligned for use in A. In this orientation, the IOM latch appears at the bottom of the module. It is shown in closed/locked position. The diagrams below identify ports. See "12Gb/s controller module LEDs" (page 40) for LED identification. These diagrams also show proper orientation for canister insertion into IOM slots located on the 5U84 enclosure rear panel.

Figure 13 (page 27) and Figure 14 (page 27) show CNC host interface ports that can be configured with 8/16Gb/s FC SFPs; 10GbE iSCSI SFPs; or 1Gb/s RJ-45 SFPs. See "Install an SFP transceiver" (page 179) for installing qualified SFP options in CNC ports.
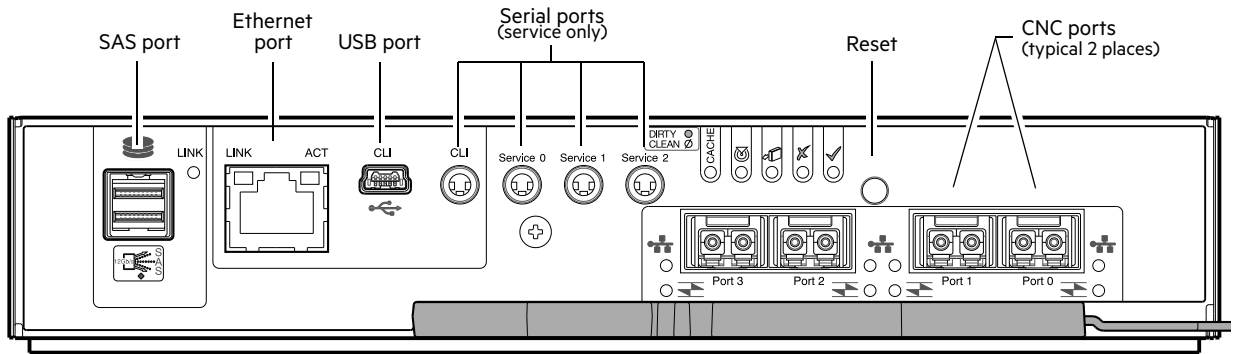
Figure 13   4-port FC/iSCSI controller module detail (5005/4005 Series)
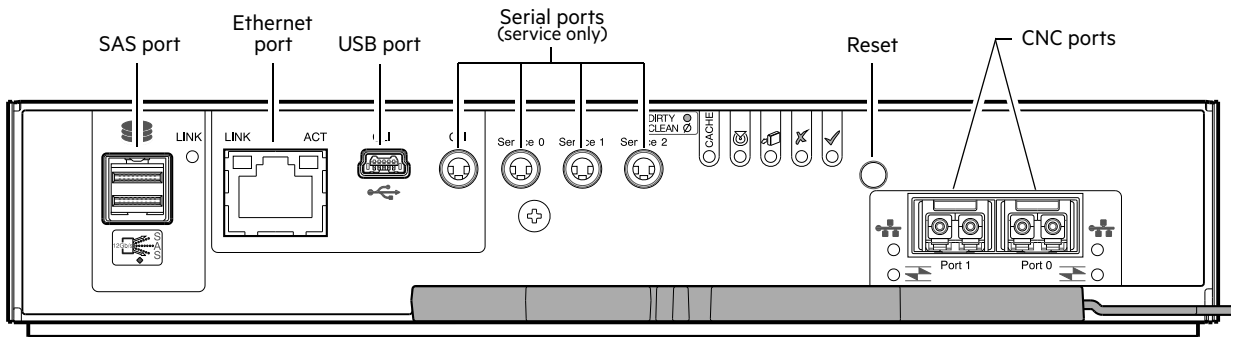
Figure 14   2-port FC/iSCSI controller module detail (3005 Series)

Figure 15 (page 27) shows SAS host interface ports that ship configured with 12Gb/s HD mini-SAS (SFF-8644) external connectors.
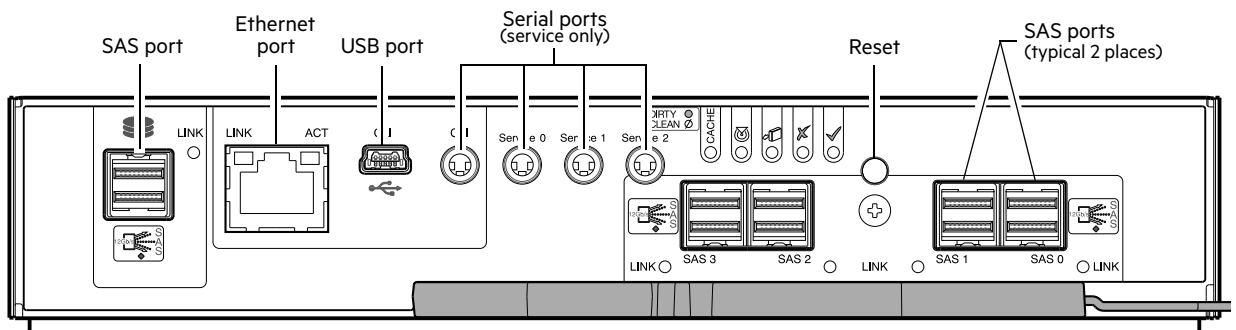
Figure 15   4-port HD mini-SAS controller module detail (5005/4005 Series)

Expansion module

Figure 16 (page 28) shows the IOM used in supported expansion enclosures for adding storage. Ports A/B/C ship configured with 12Gb/s HD mini-SAS (SFF-8644) external connectors.
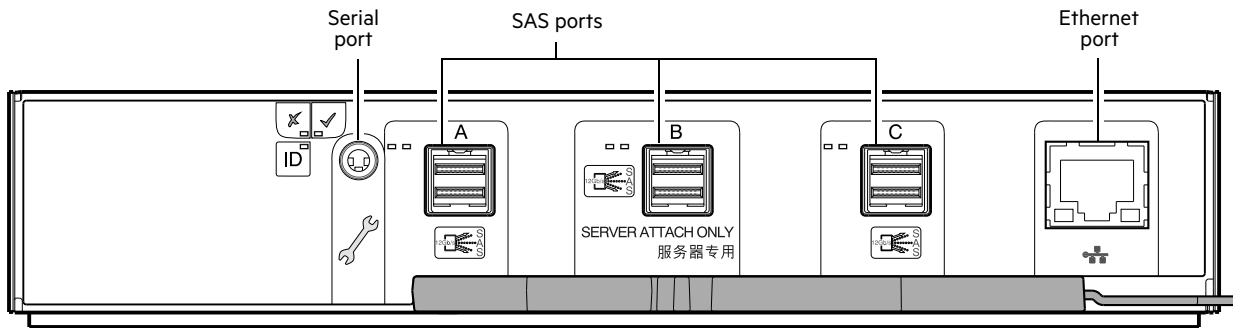
Figure 16 Expansion module detail (J1212/J1224/J1284)

ⓘ **IMPORTANT:**   RBOD/EBOD/JBOD configurations:

- When the expansion module shown above (Figure 16 (page 28)) is used with 5005/4005/3005 Series controller modules for adding storage, its middle HD mini-SAS expansion port ("B") is disabled by the firmware. See also Figure 51 (page 61).

- The Ethernet port on the expansion module is not used in RBOD/EBOD configurations and is disabled.

Power cooling module

Figure 17 (page 28) shows the power cooling module (PCM) used in controller enclosures and optional expansion enclosures. The example shows a PCM oriented for use in the left PCM slot of the enclosure rear panel.
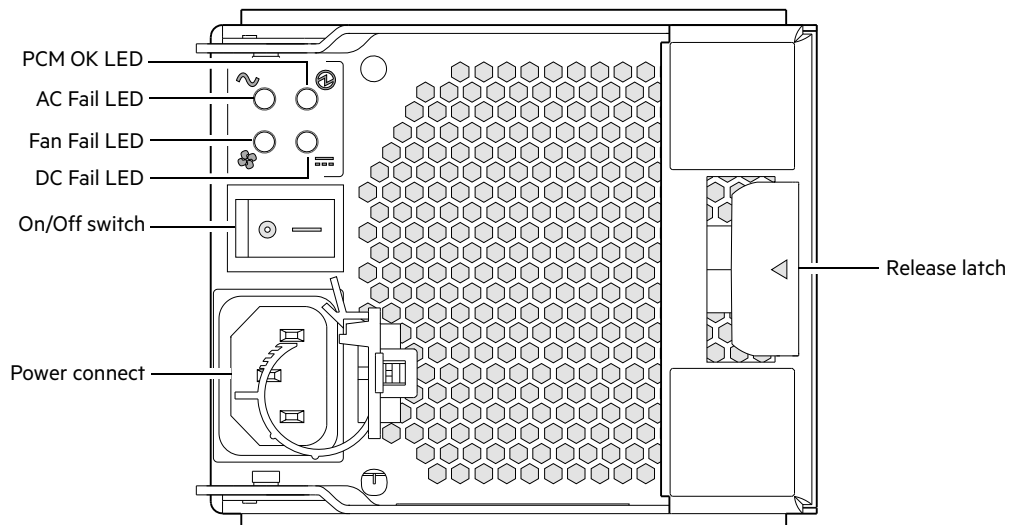
Figure 17 Power cooling module (PCM) detail

2U enclosure chassis

The 2U chassis consists of a sheet metal enclosure with an integrated midplane PCB and module runner system. The 5005/4005/3005 Series supports the 2U chassis form factor.

NOTE: Supported 2U chassis form factors used for configuring RBODs, EBODs, and JBODs:

- 2U12 chassis configured with 12 LFF disks: see Figure 1 (page 21).
- 2U24 chassis configured with 24 SFF disks: see Figure 3 (page 22).
- 2U12 empty chassis with midplane and module runner system: see Figure 18 (page 29).
- 2U24 empty chassis with midplane: see Figure 19 (page 29).

- The chassis has a 19-inch rack mounting that enables it to be installed onto standard 19-inch racks and uses two EIA units of rack space (3.5") for a 2U enclosure.
- The midplane PCB can support either 12 or 24 disk drive connections.
- There are either 12 or 24 drive slots at the front of the enclosure, in horizontal (12) or vertical (24) orientation, as defined by the enclosure variant. See also Figure 7 (page 25) and Figure 8 (page 25). Each drive slot holds a plug-in drive carrier module that can hold these drive types, dependent upon the enclosure type:
  - 2U12 enclosure: 12 low profile (1" high) 3.5" LFF disk drives, held horizontally.
  - 2U24 enclosure: 24 low profile (5/8" high) 2.5" LFF disk drives, held vertically.
- At the rear, the chassis assembly can hold a maximum of two PCMs and two SBB-compliant IOMs.
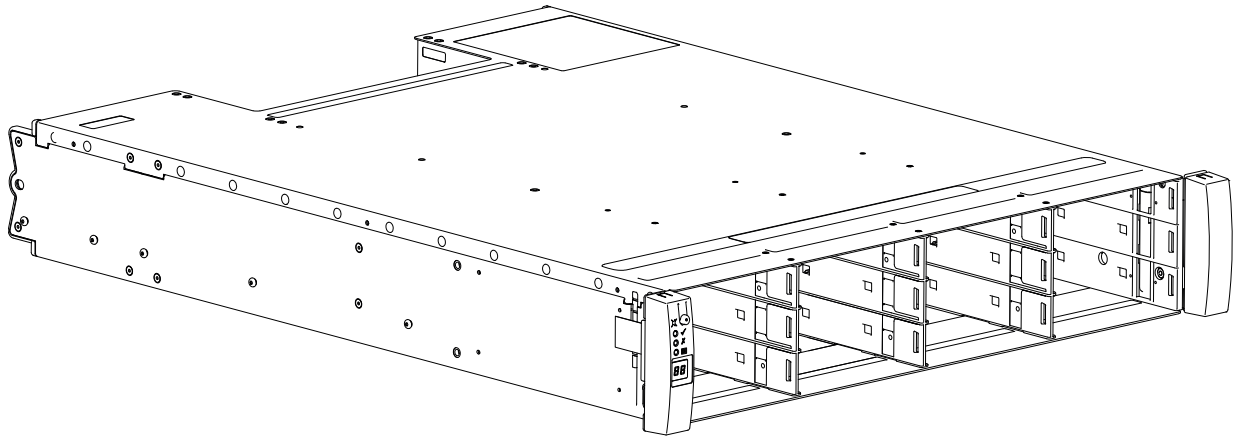
Figure 18 2U12 enclosure chassis – dimetric front orientation
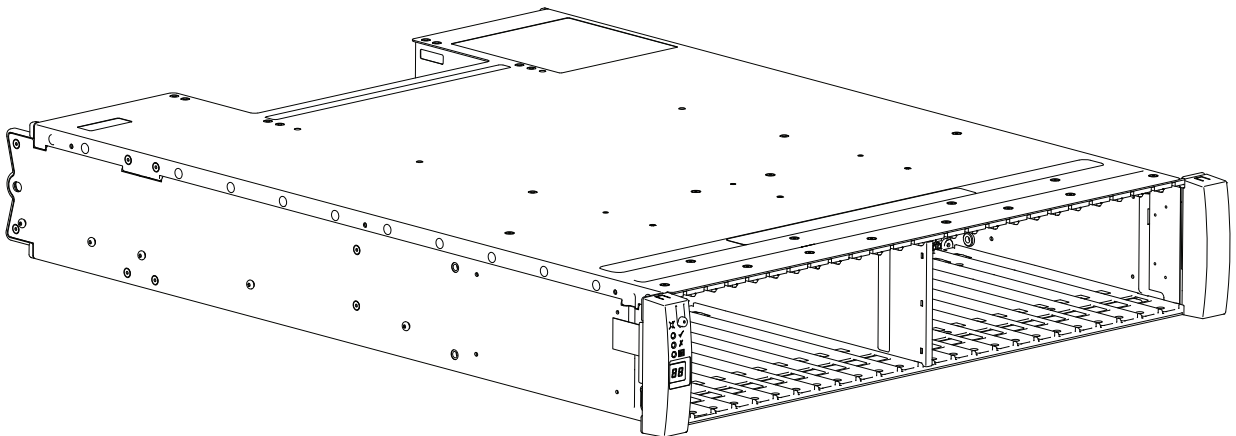
Figure 19 2U24 enclosure chassis – dimetric front orientation

---

NOTE: Either 2U chassis can be configured as a controller enclosure or as an *optional* expansion enclosure for adding storage. Expansion enclosures can be configured as an EBOD for use with an RBOD or as a server attached JBOD.

5U enclosure core product

The 5U chassis consists of a sheet metal enclosure, integrated PCBs, and a module runner system. The design concept is based on an enclosure subsystem together with a set of plug-in modules. A typical enclosure system—as supplied—includes the following:

- An enclosure chassis consisting of:
  - Two sliding drawers for holding Disk Drive in Carrier (DDIC) modules (42 disk slots per drawer).
  - An Operator's (Ops) panel installed on the left ear and a protective cover for the right ear.
  - An enclosure bezel (a bezel cover on each drawer face).
  - A midplane PCB into which other components and CRUs connect.
  - Four sideplanes and six baseplanes providing for connectivity of front-end components.
- Two power supply units.
- Five fan cooling modules.
- Two IOMs: 2 x SBB-compliant interface slots.
- Up to 84 DDIC modules populated within two drawers (42 DDICs per drawer; 14 DDICs per row).

  ☼ **TIP:** The 5U84 does not ship with DDICs installed. DDICs ship in a separate container, and must be installed into the enclosure drawers during product installation and setup. See the **Installation checklist** table on page 53.

- A rail kit for rack mounting.

ⓘ **IMPORTANT:** To ensure sufficient circulation and cooling throughout the enclosure, all PSU slots, cooling module slots, and IOM slots must contain a functioning CRU. Do not replace a faulty CRU until the replacement is available in hand.

5U enclosure front panel
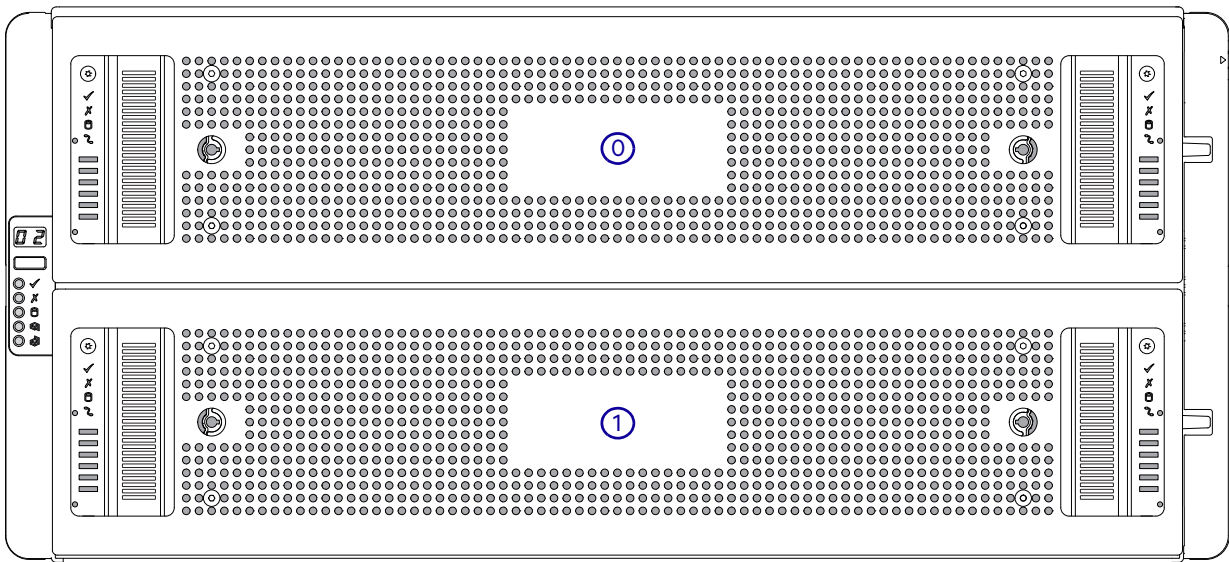
Figure 20 5U84 enclosure system – front panel components

Plan view of a drawer – showing installed DDICs and numeric indexing

Forward

Drawer front panel

Into slot

0/42	14/56	28/70
1/43	15/57	29/71
2/44	16/58	30/72
3/45	17/59	31/73
4/46	18/60	32/74
5/47	19/61	33/75
6/48	20/62	34/76
7/49	21/63	35/77
8/50	22/64	36/78
9/51	23/65	37/79
10/52	24/66	38/80
11/53	25/67	39/81
12/54	26/68	40/82
13/55	27/69	41/83

Forward
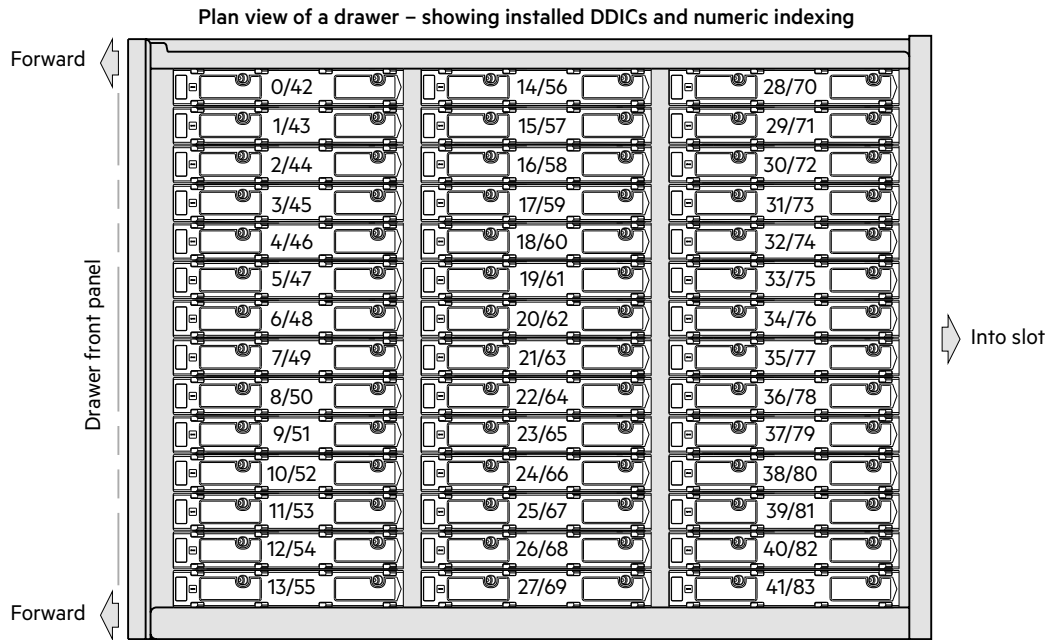
Figure 21 5U84 enclosure system – plan view of drawer accessed from front panel

shows a plan view of an enclosure drawer that is accessed from the enclosure front panel. The view is scaled and rotated -90° about the vertical axis to better fit the page, and the conceptual graphics are simplified for clarity.

ⓘ **IMPORTANT:**   Drawer sideplanes—also known as side cards—are *not* hot-swappable as field-replaceable units.

5U enclosure rear panel

Numeric designators on PSUs and FCMs, and alphabetic designators on IOMs indicate slot sequencing for modules used in 5U enclosures. PSU, FCM, and IOM modules are available as CRUs. 5005 and 4005 RBODs use 4-port controller modules and support the J1212/J1224/J1284 EBODs for optionally adding storage. The J1212/J1224/J1284 expansion enclosures can also be attached to server HBAs as JBODs. RAID controls reside on the server.
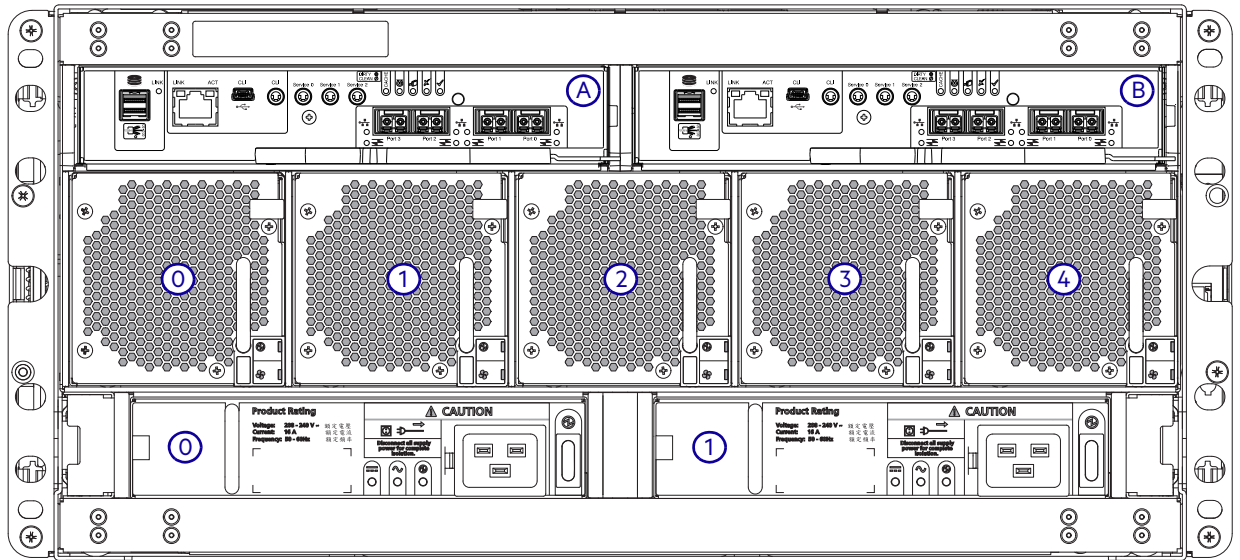
Figure 22 5U84 controller enclosure – rear panel components (4-port FC/iSCSI)
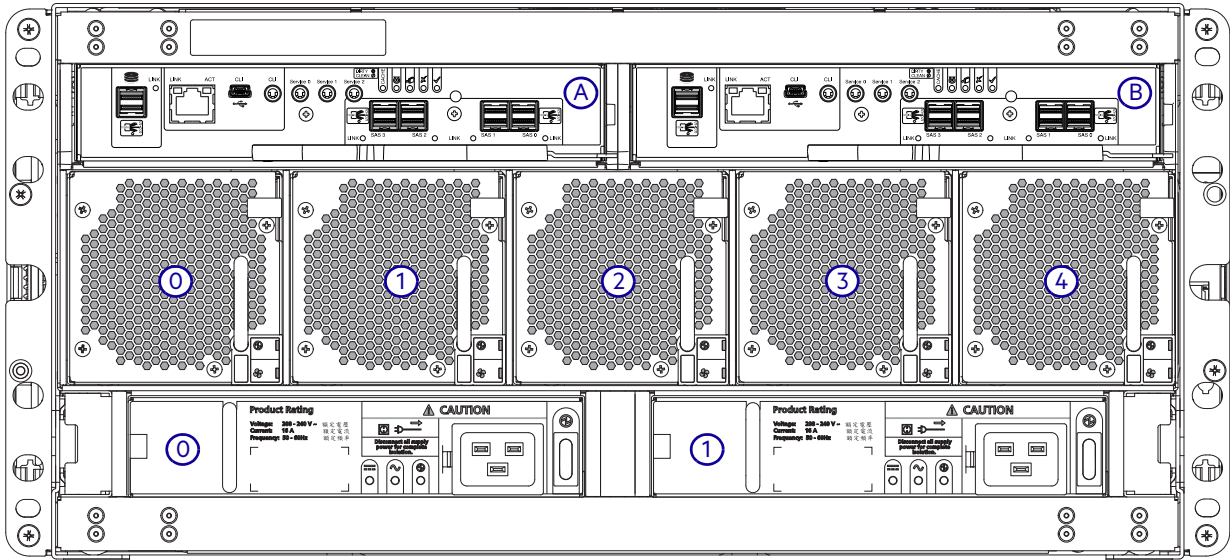
Figure 23 5U84 controller enclosure – rear panel components (4-port SAS)
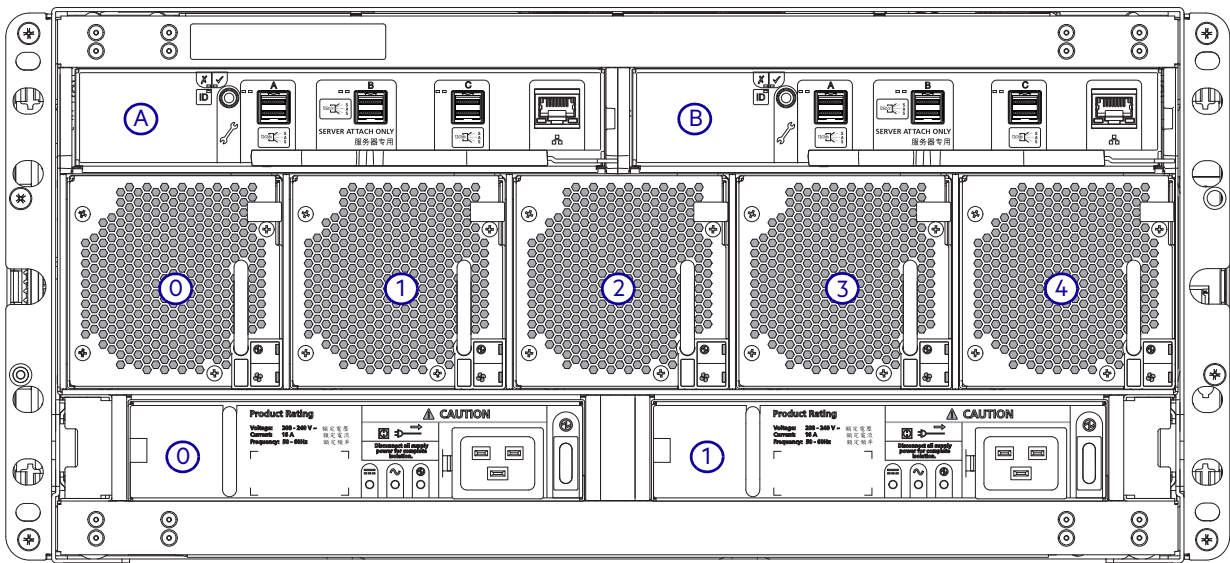
Figure 24 5U84 expansion enclosure – rear panel components

5U rear panel components

Controller modules

Except for the 2-port CNC module, the 5U84 controller enclosure uses the same IOMs used by 2U12 and 2U24 enclosures. See "Controller modules" (page 27). See also Figure 13 (page 27) and Figure 15 (page 27).

Expansion module

The 5U84 expansion enclosure uses the same IOM used by 2U12 and 2U24 enclosures. See "Expansion module" (page 28). See also Figure 16 (page 28) and the **IMPORTANT** entry on page 28.

Power supply module

shows the power supply unit used in 5U controller enclosures and optional 5U expansion enclosures.
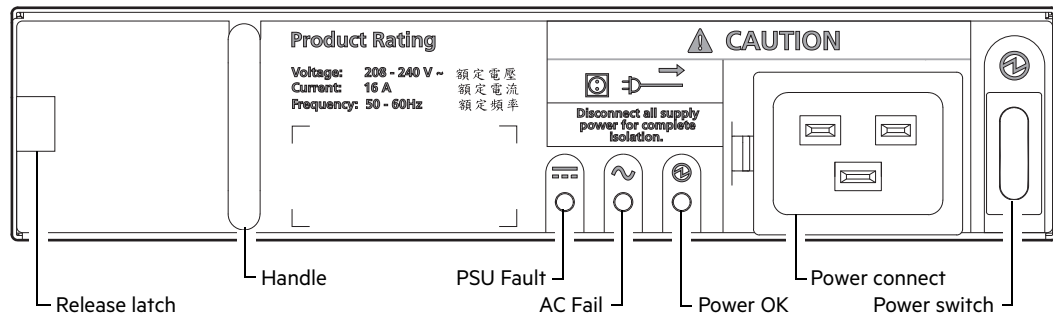
Figure 25 Power supply unit (PSU) detail

The 2U enclosures use a redundant PCM that integrates a power supply and cooling fan within the same CRU module, as shown in . Conversely, the 5U enclosures use separate CRU modules for power supply and cooling/circulation, respectively. shows the power supply module, which provides the enclosure with power connection and a power switch. shows the fan cooling module used in 5U enclosures. The FCM is smaller than the PCM, and five of them are used within the 5U enclosure to provide sufficient airflow throughout the enclosure.

Fan cooling module

shows the fan cooling module (FCM) used in 5U controller enclosures and optional 5U expansion enclosures.
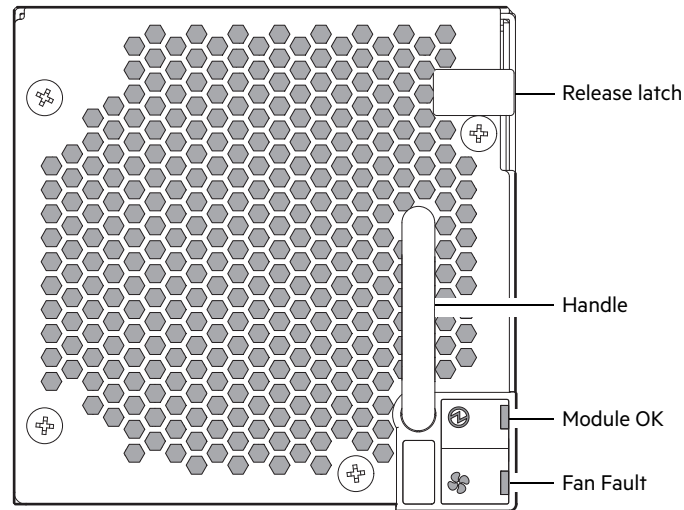
Figure 26 Fan cooling module (FCM) detail

5U enclosure chassis

The 5005/4005 Series support the 5U84, whereas the 3005 series does not. The 5U chassis consists of a sheet metal enclosure with an integrated midplane PCB, module runner system, and two drawers for drive modules.

NOTE: Supported 5U chassis form factors used for configuring RBODs, EBODs, and JBODs:

- 5U84 chassis configured with up to 84 LFF disks in DDICs: see .
- 5U84 chassis configured with SFF disks in 2.5" to 3.5" hybrid driver carrier adapter.

- 5U84 empty chassis with midplane, module runner system, and drawers.

- The chassis has a 19-inch rack mounting that enables it to be installed onto standard 19-inch racks and uses five EIA units of rack space (8.75") for a 5U84 enclosure.
- At the front of the enclosure two drawers can be opened and closed. Each drawer provides access to 42 slots for Disk Drive in Carrier (DDIC) modules. DDICs are top mounted into the drawers as shown in the drawer diagram on . The front of the enclosure also provides enclosure status LEDs and drawer status/activity LEDs.
- At the rear of the enclosure, access is provided to rear panel CRUs: two SBB-compliant IOMs; two PSUs, and five fan cooling modules.

5U enclosure drawers

Each enclosure drawer contains 42 slots, each of which can accept a single DDIC containing a 3.5" LFF disk drive or a 2.5" SFF disk drive with an adapter. See Figure 20 (page 30) and Figure 21 (page 31).

Opening a drawer does not interrupt the functioning of the storage system, and DDICs can be hot-swapped while the enclosure is in operation. However, drawers must not be left open for longer than 2 minutes, or airflow and cooling will be compromised.

⊙ **IMPORTANT:**   During normal operation, drawers should be closed to ensure proper airflow and cooling within the enclosure.

A drawer is designed to support its own weight, plus the weight of installed DDICs, when fully opened.

Safety features:

- To reduce the possibility of toppling, only one drawer can be open at a time.
- The drawer locks into place when fully opened and extended. To reduce pinching hazards, two latches must be released before the drawer can be pushed back into the drawer slot within the enclosure.

Power and data are sent via three baseplanes and two sideplanes. The direct-dock SATA version of the 5U enclosure has one active and one passive sideplane, and has dual power paths, but only a single data path.

Each drawer can be locked shut by turning both anti-tamper locks clockwise using a screwdriver with a Torx T20 bit. The anti-tamper locks are symmetrically placed on the left-hand and right-hand sides of the drawer bezel. Drawer status and activity LEDs can be monitored via two Drawer LEDs panels located next to the two drawer-pull pockets located on the left-hand and right-hand side of each drawer.

Left side   Right side
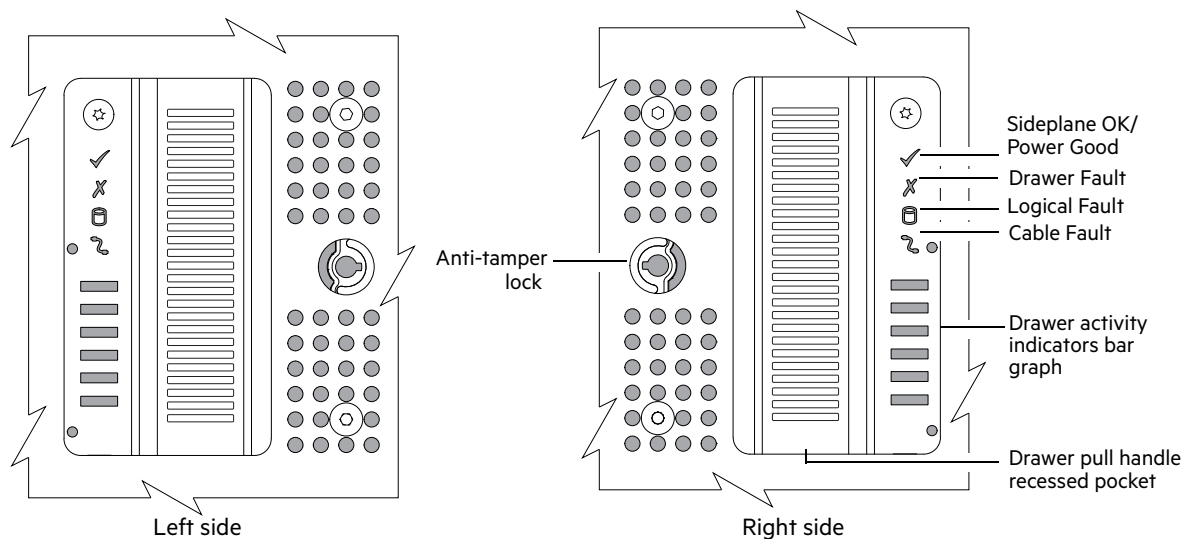
Figure 27   Drawer bezel details

Operator's (Ops) panel

Each of the enclosures supported by 5005/4005/3005 Series feature an Ops panel located on the chassis left ear flange. The Ops panel for 2U and 5U enclosures are described herein.

2U enclosure Ops panel

The enclosure front panel has an Operator's (Ops) panel mounted on the left ear flange. A flexible cable connects the Ops panel to the midplane. The Ops panel is a passive component: the midplane controls the panel, and the IOMs control all the panel's functions. An integral part of the enclosure chassis, the Ops panel is not replaceable on site. The Ops panel provides the functions shown in the illustration below and listed in the table. See also "Ops panel LEDs" (page 92).
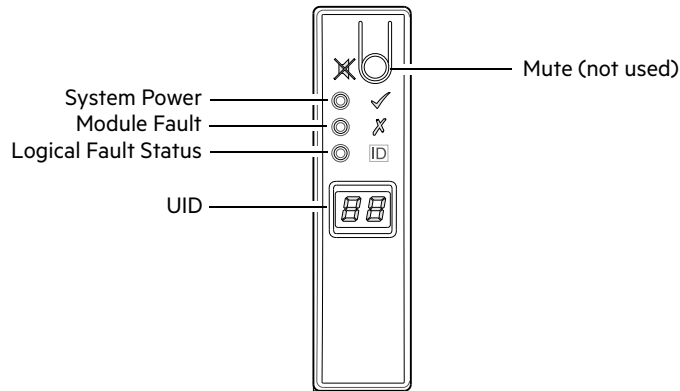
Ops panel functions (see left ear on 2U front panel)	
LED	**Status**
System Power On/Standby	Constant green: positive indication Constant amber: fault present
Module Fault	Constant or blinking amber: fault present
Logical Fault Status	Constant or blinking amber: fault present
Unit Identification Display	Green (seven segment display: enclosure sequence)
Thermal sensor	Located behind panel

Figure 28 LEDs: Ops panel – 2U enclosure front panel

System Power On/Standby LED (green/amber)

LED displays amber when only standby power is available. LED displays green when system power is available.

Module Fault LED (amber)

LED illuminates when experiencing a system hardware fault. It may be associated with a Fault LED on a PCM or IOM that helps the user to identify which component is causing the fault.

Logical Fault Status (amber)

This LED indicates a change of status or fault from something other than the enclosure management system. This may be initiated from the controller module or an external HBA.

Unit Identification Display

The UID is a dual seven-segment display that can be used to provide feedback to the user. Its primary function is to display an enclosure unit identification number to assist users in setting and maintaining multiple enclosure systems.

Thermal sensor

The thermal sensor is located on the outside of the enclosure, and it sends input to the enclosure about its external operating ambient temperature.

5U enclosure Ops panel

The front of the enclosure features an Operator's panel located on the left ear flange of the 5U chassis as shown in Figure 20 (page 30).

The Ops panel is an integral part of the chassis CRU, and is not replaceable on site. The 5U Ops panel provides the following:
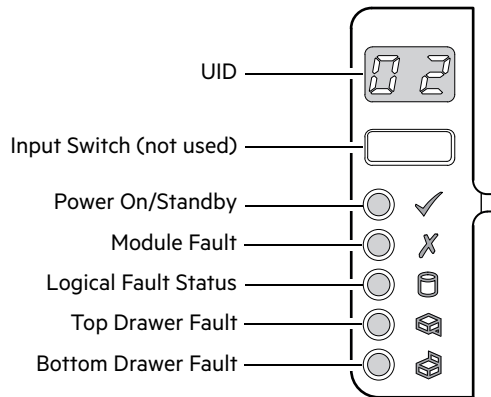
Ops panel functions (see left ear on 5U front panel)	
LED	**Status**
Unit Identification Display	Green (seven segment display: enclosure sequence)
System Power On/Standby	Constant green: positive indication Constant amber: system in standby (not operational)
Module Fault	Constant or blinking amber: fault present
Logical Fault Status	Constant or blinking amber: fault present
Top Drawer Fault	Constant or blinking amber: fault present in drive, cable, or sideplane.
Bottom Drawer Fault	Constant or blinking amber: fault present in drive, cable, or sideplane.

Figure 29 LEDs: Ops panel – 5U enclosure front panel

Unit Identification Display

The UID is a dual seven-segment display that can be used to provide feedback to the user. Its primary function is to display an enclosure unit identification number to assist users in setting and maintaining multiple enclosure systems.

System Power On/Standby LED (green/amber)

LED displays amber when only standby power is available (non-operational). LED displays green when system power is available (operational).

Module Fault LED (amber)

LED illuminates when experiencing a system hardware fault. It may be associated with a Fault LED on a PSU, FCM, IOM, DDIC, or Drawer that helps the user to identify which component is causing the fault.

Logical Status LED (amber)

This LED indicates a change of status or fault from something other than the enclosure management system. This may be initiated from the controller module or an external HBA. The indication is typically associated with a DDIC and LEDs at each disk position within the drawer, which help to identify the DDIC affected.

Drawer Fault LEDs (amber)

This LED indicates a disk, cable, or sideplane fault in the drawer indicated: Top (Drawer 0) or Bottom (Drawer 1).

Power cooling module – 2U enclosures

AC-DC power is provided by up to two auto-ranging power cooling modules (PCMs) with integrated axial cooling fans. The IOMs control fan speed. Also see "System airflow" (page 38) for optimal cooling within the enclosure(s).

580W PCM

The 580W PCM voltage operating range is nominally 100V–240V AC, and operates at 50–60 Hz input frequency. The dimetric rear orientation in Figure 30 (page 37) shows the PCM aligned for insertion into the right-hand PCM slot located on the enclosure rear panel.

PCM LEDs Inset
As shown on PCM in left-hand slot

PCM oriented for use in right-hand slot
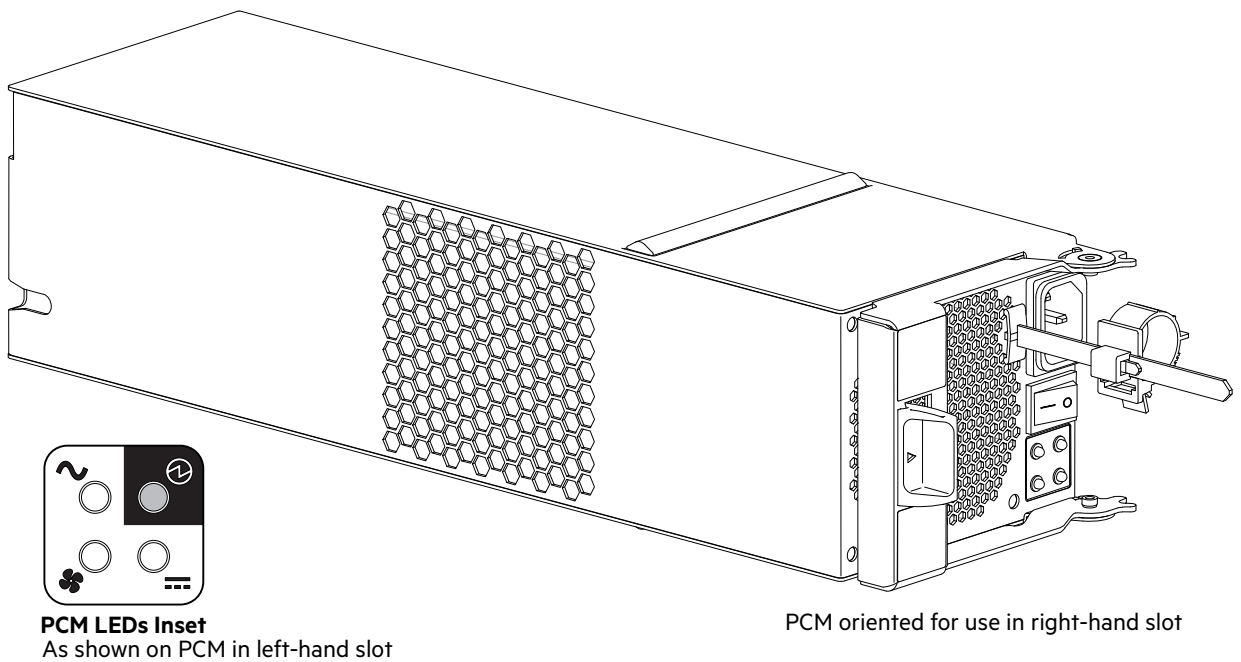
580W Power Cooling Module LED descriptions (see inset illustration detail above)	
PCM OK LED (Green)	Fan Fail LED (Amber/blinking amber)
AC Input Fail LED (Amber/blinking amber)	DC Fail LED (Amber/blinking amber)
LED behavior	
• If any of the PCM LEDs are illuminated amber, a module fault condition or failure has occurred. • For a detailed description of PCM LED behavior, see "580W PCM LEDs" (page 92).	

Figure 30 LEDs: 580W PCM – 2U enclosure rear panel

For a front view of a PCM panel located on the enclosure rear-panel, see Figure 17 (page 28).

Multiple power cooling modules

The 2U storage system includes two PCMs which provide redundant power control for the system so that if one PCM fails, the other maintains the power supply, and enclosure operation is not affected while you replace the faulty module.

PCMs are hot-pluggable, and replacement should only take a few seconds to do. Replacement must be completed as soon as possible after the removal of the defective PCM to avoid a thermal exception. The replacement procedure should be completed within an absolute maximum of 2 minutes.

> ⓘ **IMPORTANT:** Operation of the enclosure with any modules missing will disrupt the airflow, and the disks will not receive sufficient cooling. It is essential that all slots are fitted with PCMs prior to powering on the enclosure.

System airflow

The system must be operated with low pressure rear exhaust installation. Back pressure created by rack doors and obstacles must not exceed 5 pascals (0.5mm water gauge). The cooling system provides sufficient capacity to ensure that maximum temperatures are not exceeded.

> ⓘ **IMPORTANT:** The environment in which the enclosure operates must be dust-free to ensure adequate airflow.

Power supply unit – 5U enclosures

Power is provided by two 2,214W PSUs. The PSU voltage operating range is nominally 200V–240V AC, and operates at 50–60 Hz input frequency. The dimetric rear orientation in Figure 30 (page 37) shows the PSU aligned for insertion into its slot located on the enclosure rear panel.

PSU LEDs Inset

- Power OK
- AC Input Fail
- PSU Fail
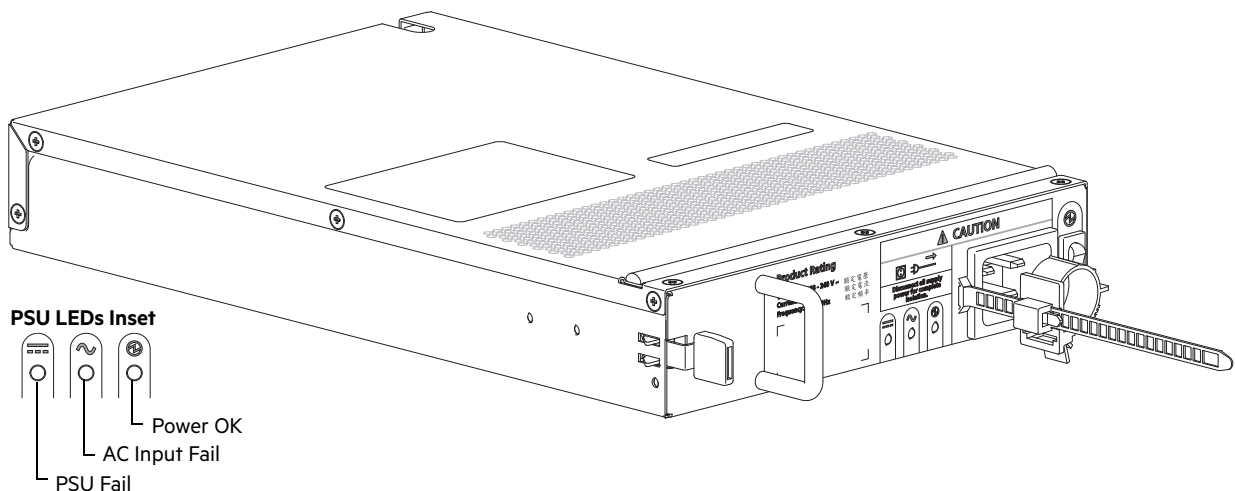
2,214W Power Supply Unit LED descriptions (see inset illustration detail above)	
Power OK LED (Green)	PSU Fail LED (Amber/blinking amber)
AC Input Fail LED (Amber/blinking amber)	
LED behavior	
• If any of the PSU LEDs are illuminated amber, a module fault condition or failure has occurred. • For a detailed description of PSU LED behavior, see "PSU LEDs" (page 95).	

Figure 31 LEDs: 2,214W PSU – 5U enclosure rear panel

Dual PSUs provide redundant power for the 5U system: if one PSU fails, the other keeps the system running while you replace the faulty module. PSUs are hot-swappable. Replacement of a PSU can be performed while the enclosure is running, but the procedure must be completed within two minutes of the removal of the defective module. Verify that you have a replacement PSU on hand before removing the defective module.

Fan cooling module – 5U enclosures

The five fan cooling modules (FCM) at the rear of the enclosure maintain all system components below their maximum temperature, assuming the ambient temperature is below 35ºC. Fan speed within the FCMs is governed by the controller modules.

FCMs are hot-swappable. Replacement of an FCM can be performed while the enclosure is running, but the procedure must be completed within two minutes of the removal of the defective module. Verify that you have a replacement module on hand before removing the defective FCM.

FCM LEDs Inset
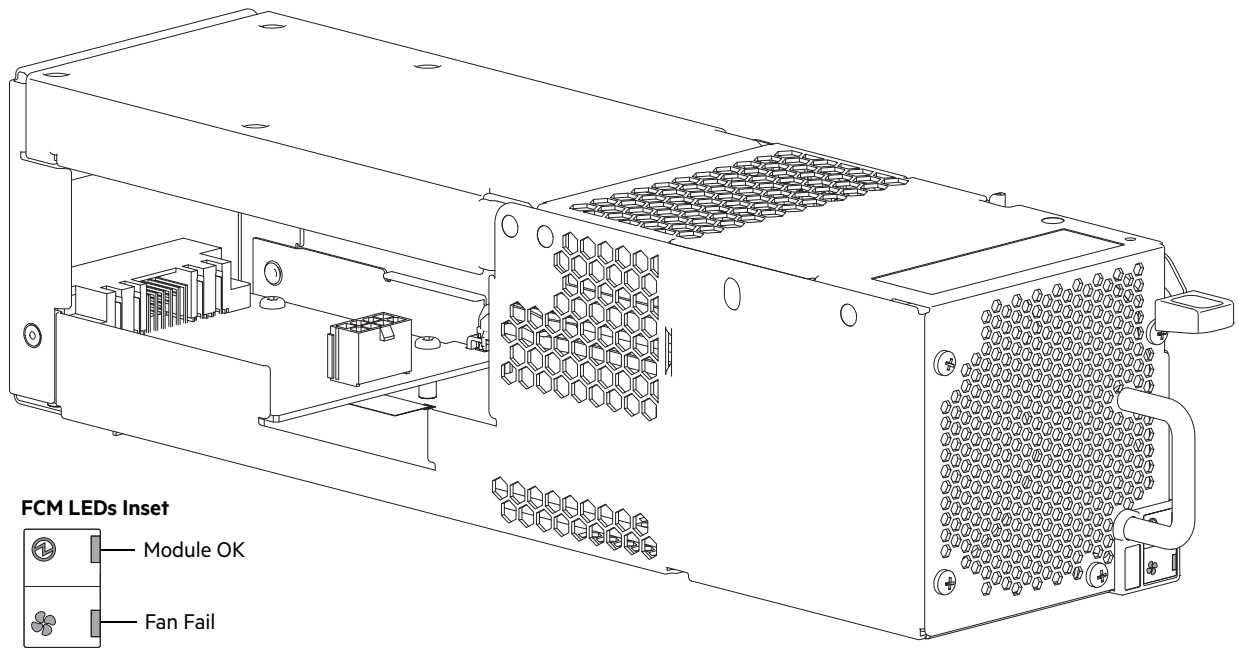
— Module OK

— Fan Fail

Fan Cooling Module LED descriptions (see inset illustration detail above)	
Module OK LED (Green)	Fan Fail LED (Amber/blinking amber)
LED behavior	
• If any of the FCM LEDs are illuminated amber, a module fault condition or failure has occurred.	
• For a detailed description of FCM LED behavior, see "Fan cooling module LEDs" (page 95).	

Figure 32 LEDs: FCM – 5U enclosure rear panel

NOTE: The "System airflow" (page 38) topic introduced for the 2U enclosures also applies to 5U enclosures.

Controller and expansion modules

This section describes the IOMs used in 12Gb/s storage enclosures. They are mechanically and electrically compliant to the latest SBB v2.1 specification.

The dimetric rear orientation in shows a 4-port FC/iSCSI controller module aligned for use in the top controller module slot located on the 2U enclosure rear panel. The controller module is also properly aligned for use in either IOM slot located on the 5U enclosure rear panel.
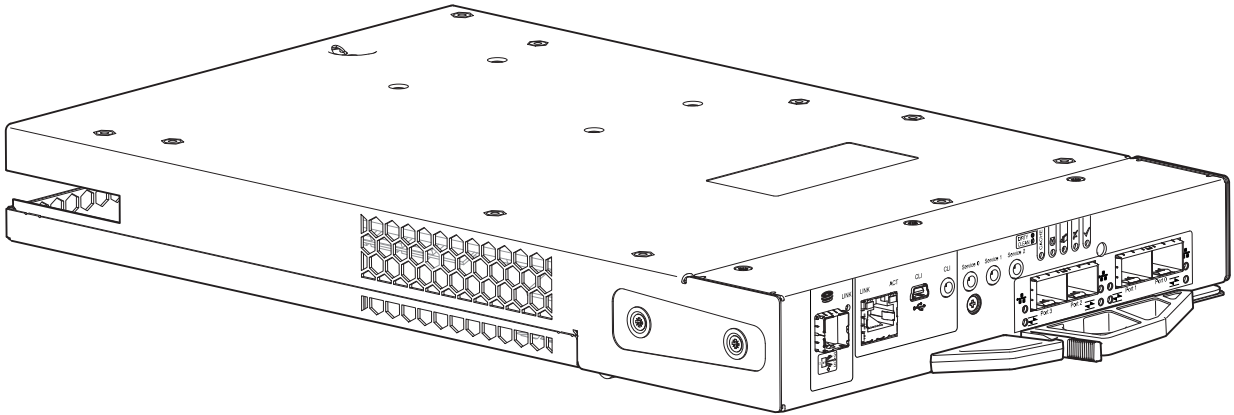
Figure 33 Controller module – dimetric rear orientation

Each controller module maintains VPD (Vital Product Data) in EEPROM devices, and are interconnected by SBB-defined I^{2}C buses on the midplane. In this way, the SBB modules can discover the type and capabilities of the partner SBB module(s), and vice versa, within the enclosure. An enclosure system alarm occurs when incompatible configurations are detected.

12Gb/s controller module LEDs

The diagrams with tables that immediately follow provide descriptions for the different controller modules that can be installed into the rear panel of a 5005/4005/3005 Series controller enclosure. Showing controller modules separately from the enclosure enables improved clarity in identifying the component items called out in the diagrams and described in the companion tables within the figure/table ensembles.

NOTE: Consider the following when viewing IOM diagrams appearing on the following pages

- In each IOM diagram, the canister is oriented for insertion into the top IOM slot (A) of 2U enclosures.

  When oriented for use in the bottom IOM slot (B) of 2U enclosures, the IOM labels appear upside down.

- In each IOM diagram, the canister is oriented for insertion into either IOM slot of 5U84 enclosures.

LED	Description	Definition
1	Host 4/8/16Gb 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 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	OK	Green — The controller is operating normally. Blinking green — System is booting. Off — The controller module is not OK or is powered off.
4	Fault	Off — The controller is operating normally. 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.
5	OK to Remove	Off — The controller is not prepared for removal. Blue — The controller module is prepared for removal.
6	Identify	White — The controller module is being identified.
7	Cache Status[4]	Green — Cache is dirty (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.
8	Network Port Link Active Status[5]	Off — The Ethernet link is not established, or the link is down. Green — The Ethernet link is up (applies to all negotiated link speeds).
9	Network Port Link Speed[5]	Off — Link is up at 10/100base-T negotiated speeds. Amber — Link is up and negotiated at 1000base-T.
10	Expansion Port Status	Off — The port is empty or the link is down. Green — The port is connected and the link is up.

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

2-When in 10GbE iSCSI mode, the SFPs must be a qualified 10GbE iSCSI optic option.

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

4-Cache Status LED supports power on behavior and operational (cache status) behavior. See also "Cache Status LED details" (page 46).

5-When port is down, both LEDs are off.

Figure 34 LEDs: 5005/4005 Series FC/iSCSI controller modules (FC and 10GbE SFPs)

LED	Description	Definition
1	Host 4/8/16Gb 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 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	OK	Green — The controller is operating normally. Blinking green — System is booting. Off — The controller module is not OK or is powered off.
4	Fault	Off — The controller is operating normally. 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.
5	OK to Remove	Off — The controller is not prepared for removal. Blue — The controller module is prepared for removal.
6	Identify	White — The controller module is being identified.
7	Cache Status	Green — Cache is dirty (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.
8	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).
9	Network Port Link Speed[4]	Off — Link is up at 10/100base-T negotiated speeds. Amber — Link is up and negotiated at 1000base-T.
10	Expansion Port Status	Off — The port is empty or the link is down. Green — The port is connected and the link is up.

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

2-When in 10GbE iSCSI mode, the SFPs must be a qualified 10GbE iSCSI optic option.

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

4-When port is down, both LEDs are off.

Figure 35 LEDs: 3005 Series FC/iSCSI controller modules (FC and 10GbE SFPs)

LED	Description	Definition
1	Not used in example[1]	The FC SFP is not shown in this example.
2	Host 1Gb 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	OK	Green — The controller is operating normally. Blinking green — System is booting. Off — The controller module is not OK or is powered off.
4	Fault	Off — The controller is operating normally. 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.
5	OK to Remove	Off — The controller is not prepared for removal. Blue — The controller module is prepared for removal.
6	Identify	White — The controller module is being identified.
7	Cache Status[4]	Green — Cache is dirty (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" (page 46).
8	Network Port Activity Status[5]	Off — The Ethernet link is not established, or the link is down. Green — The Ethernet link is up (applies to all negotiated link speeds).
9	Network Port Link Speed[5]	Off — Link is up at 10/100base-T negotiated speeds. Amber — Link is up and negotiated at 1000base-T.
10	Expansion Port Status	Off — The port is empty or the link is down. Green — The port is connected and the link is up.

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

2-When in 1GbE iSCSI mode, the SFPs must be a qualified 1GbE iSCSI optic option.

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

4-Cache Status LED supports power on behavior and operational (cache status) behavior. See also "Cache Status LED details" (page 46).

5-When port is down, both LEDs are off.

Figure 36 LEDs: 5005/4005 Series iSCSI controller module (1Gb RJ-45 SFPs)

LED	Description	Definition
1	Not used in example[1]	The FC SFP is not shown in this example.
2	Host 1Gb 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	OK	Green — The controller is operating normally. Blinking green — System is booting. Off — The controller module is not OK or is powered off.
4	Fault	Off — The controller is operating normally. 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.
5	OK to Remove	Off — The controller is not prepared for removal. Blue — The controller module is prepared for removal.
6	Identify	White — The controller module is being identified.
7	Cache Status	Green — Cache is dirty (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" (page 46).
8	Network Port Activity 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).
9	Network Port Link Speed[4]	Off — Link is up at 10/100base-T negotiated speeds. Amber — Link is up and negotiated at 1000base-T.
10	Expansion Port Status	Off — The port is empty or the link is down. Green — The port is connected and the link is up.

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

2-When in 1GbE iSCSI mode, the SFPs must be a qualified 1GbE iSCSI optic option.

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

4-When port is down, both LEDs are off.

Figure 37 LEDs: 3005 Series iSCSI controller module (1Gb RJ-45 SFPs)

LED	Description	Definition
1	Host 12Gb SAS[1-3] Link Status/ Link Activity	Green — The port is connected and the link is up. Amber — Partial link exists (one or more lanes down). Blinking green or amber — Host link activity is detected.
2	OK	Green — The controller is operating normally. Blinking green — System is booting. Off — The controller module is not OK or is powered off.
3	Fault	Off — The controller is operating normally. 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.
4	OK to Remove	Off — The controller is not prepared for removal. Blue — The controller module is prepared for removal.
5	Identify	White — The controller module is being identified.
6	Cache Status[4]	Green — Cache is dirty (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" (page 46).
7	Network Port Activity Status[5]	Off — The Ethernet link is not established, or the link is down. Green — The Ethernet link is up (applies to all negotiated link speeds).
8	Network Port Link Speed[5]	Off — Link is up at 10/100base-T negotiated speeds. Amber — Link is up and negotiated at 1000base-T.
9	Expansion Port Status	Off — The port is empty or the link is down. Green — The port is connected and the link is up.

1-Cables must be qualified HD mini-SAS cable options.

2-Use a qualified SFF-8644 to SFF-8644 cable option when connecting the controller to a 12Gb SAS HBA.

3-Use a qualified SFF-8644 to SFF-8088 cable option when connecting the controller to a 6Gb SAS HBA.

4-Cache Status LED supports power on behavior and operational (cache status) behavior. See also "Cache Status LED details" (page 46).

5-When port is down, both LEDs are off.

Figure 38 LEDs: 5005/4005 Series SAS controller module (HD mini-SAS)

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

When a controller is shut down 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 power-cycled.

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 6).

Table 6   Cache Status LED – power on behavior

Item	Display states reported by Cache Status LED during power on sequence							
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 is dirty (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.

> (!) **IMPORTANT:**   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.

Controller failure when a single controller is operational

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 1Gbyte 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 blinks at a rate of 1/10 second on and 9/10 second off.

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

Answer	Action
No, the Cache LED status is off, and the controller does not boot.	If the problem persists, 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.

Table 7   LEDs: Rear panel "Cache Status"

Answer	Action
Yes, at a strobe 1:10 rate - 1 Hz, and the controller does not boot.	You may need to replace the controller module.
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 - 2 Hz, and the controller does not boot.	You may need to replace the controller module.
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.

Table 7   LEDs: Rear panel "Cache Status" (continued)

12Gb/s expansion module LEDs

If optional expansion enclosures have been cabled to add storage, the supported J1212/J1224/J1284 enclosures are configured with dual expansion modules.
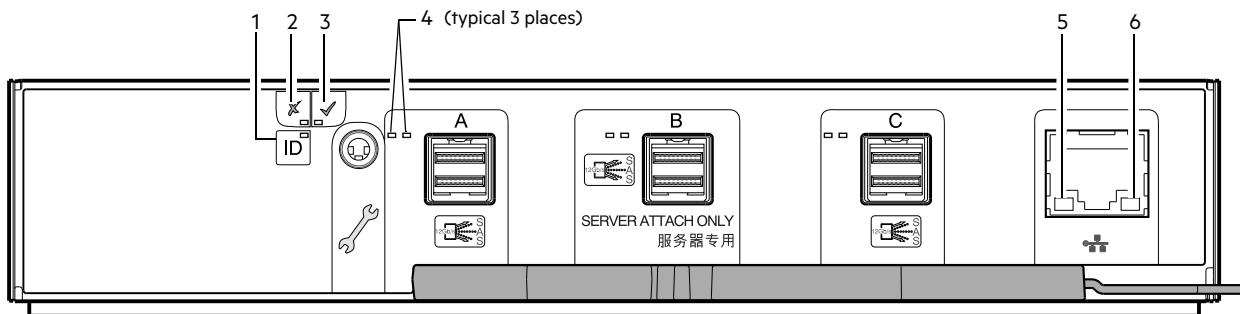
LED	Description	Definition
1	Identify	Blue — Expansion module is being identified. Blinking — 1 second on, 1 second off to identify canister.
2	Fault	Off — The expansion module is operating normally. Amber — A fault has been detected or a service action is required. Blinking — 1 second on, 1 second off for a non-critical cable fault.
3	OK	Green — The expansion module is operating normally. Blinking green — System is booting. Off — The expansion module is powered off.
4	HD mini-SAS connector LEDs (A/B/C)	See Table 8 for Activity (Green) and Fault (Amber) LED states.
5	Ethernet Port Link/Active Status (Left)	Not used in RBOD/EBOD configuration. JBOD configuration only. Green — Link Blinking green — Activity Off — No link
6	Ethernet Port Link Speed (Right)	Not used in RBOD/EBOD configuration. JBOD configuration only. Green — 100-Base-TX Off — 10-Base-TX

Figure 39 LEDs: J1212/J1224/J1284 expansion module

The following table provides companion data for the figure above relative to LED states for A/B/C SAS port expansion.

Table 8   LEDs: J1212/J1224/J1284 expansion activity states

Condition	Activity (Green)	Fault (Amber)
No cable present	Off	Off
Cable present: all links up/no activity	On	Off

Table 8 LEDs: J1212/J1224/J1284 expansion activity states (continued)

Condition	Activity (Green)	Fault (Amber)
Cable present: all links up/with aggregate port activity	Blinking	Off
Critical fault: Any fault causing operation of the cable to cease or fail to start (e.g., over current trip).	Off	On
Non-critical fault: any fault that does not cause the connection to cease operation (e.g., not all links are established; over temperature).	Blinking	Blinking 1s on/1s off

ⓘ **IMPORTANT:** RBOD/EBOD/JBOD configurations:

- When the expansion module shown above (Figure 39 (page 47)) is used with 5005/4005/3005 Series controller modules for adding storage, its middle HD mini-SAS expansion port ("B") is disabled by the firmware. See also Figure 51 (page 61)

- The Ethernet port on the expansion module is not used in RBOD/EBOD configurations, and is disabled.

Disk drive module

Different disk drive modules are used by the different chassis form factors (2U12, 2U24, 5U84). The disk modules used in 2U enclosures are referred to as drive carrier modules, whereas those used in 5U enclosures are called DDICs (Disk Drive in Carrier).

Drive carrier module in 2U chassis

The drive carrier module comprises a hard disk held by a carrier.

- Each 2U12 drive slot holds a single low profile 1.0-inch high, 3.5-inch form factor disk drive in its carrier. The disk drives are horizontal. A 2.5" to 3.5" carrier adapter is available to accommodate 2.5" disk drives.

- Each 2U24 drive slot holds a single low profile 5/8-inch high, 2.5-inch form factor disk drive in its carrier. The disk drives are vertical.

The carriers have mounting locations for:

- Direct dock SAS drives.

A sheet steel carrier holds each drive, which provides thermal conduction, radio frequency, and electro-magnetic induction protection, and physically protects the drive.

The front cap also has an ergonomic handle which gives the following functions:

- Secure location of the carrier into and out of drive slots.
- Positive spring-loading of the drive/midplane connector.

The carrier can use this interface:

- Dual path direct dock Serial Attached SCSI.
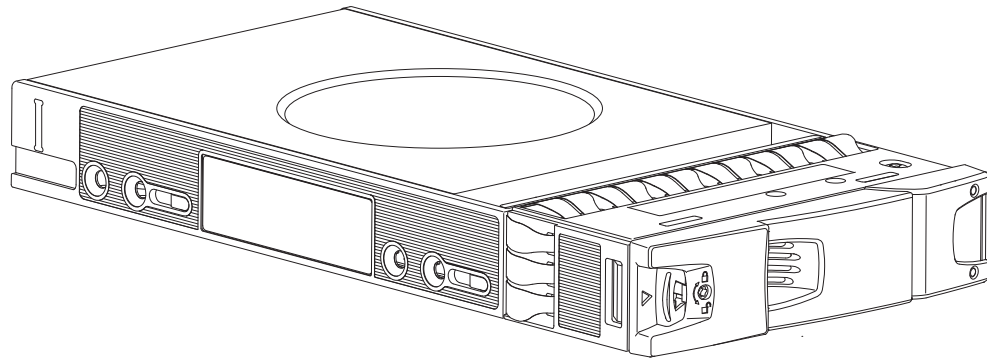
3.5" SAS drive

Figure 40 Dual path LFF 3.5" disk drive carrier modules
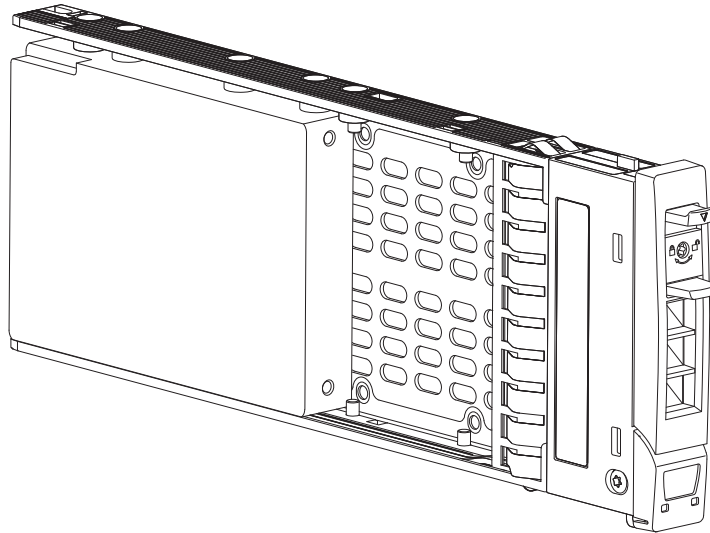
Figure 41 Dual path SFF 2.5" drive carrier module with no Interposer card
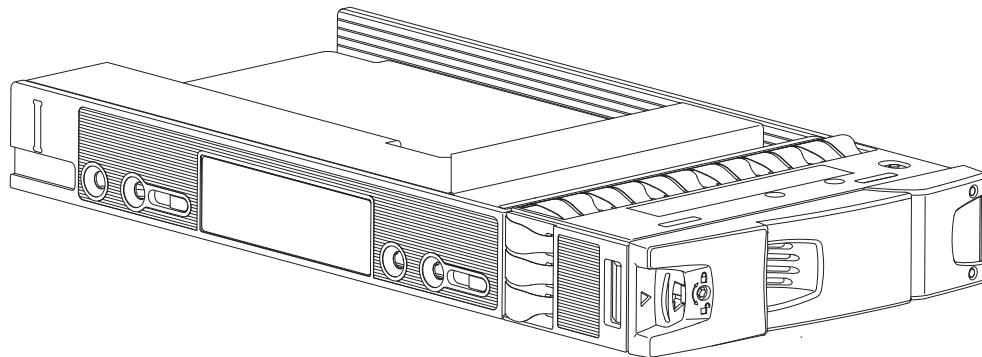
Figure 42 2.5" to 3.5" hybrid drive carrier adapter

Drive status indicators

Green and amber LEDs on the front of each drive carrier module indicate disk drive status. The SEP controls these LEDs. "Disk drive carrier module LEDs" (page 93) describes the LED states.

Dummy drive carrier modules

Dummy drive carrier modules, also known as drive blanks, are provided in 3.5" (2U12) and 2.5" (2U24) form factors. They must be installed in empty disk drive slots to create a balanced air flow.

(3.5" drive slot)

(2.5" drive slot)
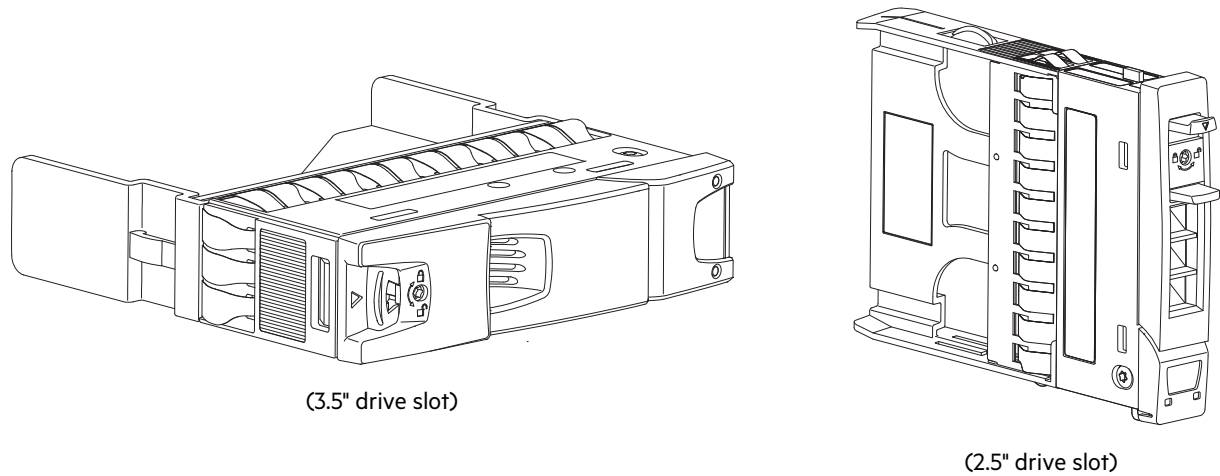
Figure 43 Dummy drive carrier module

DDIC in 5U chassis

Each disk drive is housed in a carrier that enables secure insertion of the disk drive into the drawer with the appropriate SAS carrier transition card. Figure 44 (page 50) shows a DDIC with a 3.5" disk.
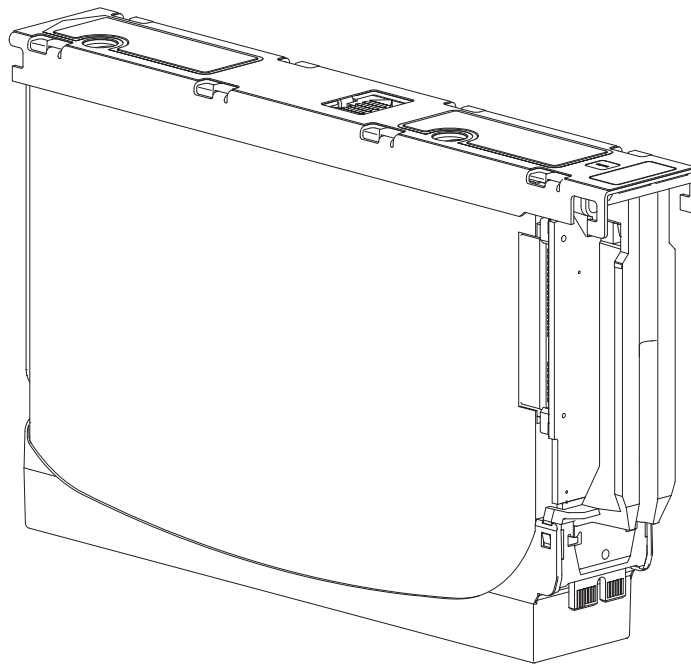
Figure 44 DDIC with 3.5" disk

Figure 45 (page 51) shows a DDIC with an adapter and a 2.5" disk.
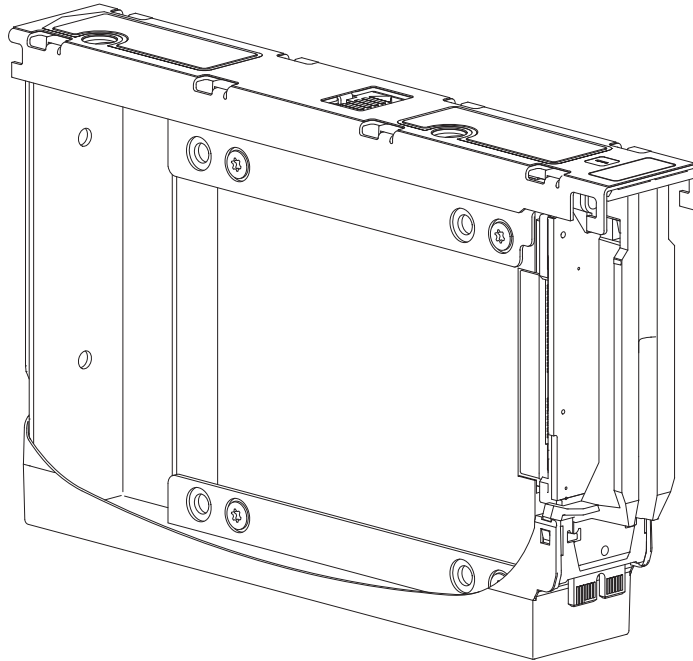
Figure 45 DDIC with adapter and 2.5" disk

The DDIC features a latch button (center of top view) and a slide latch with arrow label (left of latch button). These features allow you to install and secure the DDIC into the disk slot within the drawer. They also allow you to disengage the DDIC from its slot, and remove it from the drawer. The DDIC features a single Drive Fault LED (top view on right), which illuminates amber when the disk has a fault.

Populating drawers with DDICs

The 5U84 enclosure does not ship with DDICs installed. Please refer to Figure 21 (page 31) when reviewing these rules:

- The minimum number of disks supported by the enclosure is 14.
- DDICs must be added to disk slots in rows (14 disks at a time).
- Beginning at the front of the drawer(s), install DDICs consecutively by number, and alternately between the top drawer and the bottom drawer. Namely, install first at slots 0 through 13 in the top drawer, and then 42 through 55 in the bottom drawer. After that, install slots 14 through 27, and so on.
- The number of populated rows must not differ by more than one row between the top and bottom drawers.
- Hard disk drives (HDD) and solid state drives (SDD) can be mixed in the same drawer.
- HDDs installed in the same row should have the same rotational speed.
- Although DDICs holding 3.5" disks can be intermixed with DDICs holding 2.5" disks within the enclosure, each row should be populated with disks of the same form factor (all LFF or all SFF).

NOTE: For additional information, see "Populating drawers" (page 139) within "Module removal and replacement".

Figure 103 (page 137) shows an empty drawer, which is representative of how the 5U84 enclosure is delivered. Figure 46 (page 52) shows a drawer that is fully populated with DDICs:

- See Figure 44 (page 50) for the DDIC holding the 3.5" disk
- See Figure 45 (page 51) for the DDIC holding the 2.5" disk with 3.5" adapter
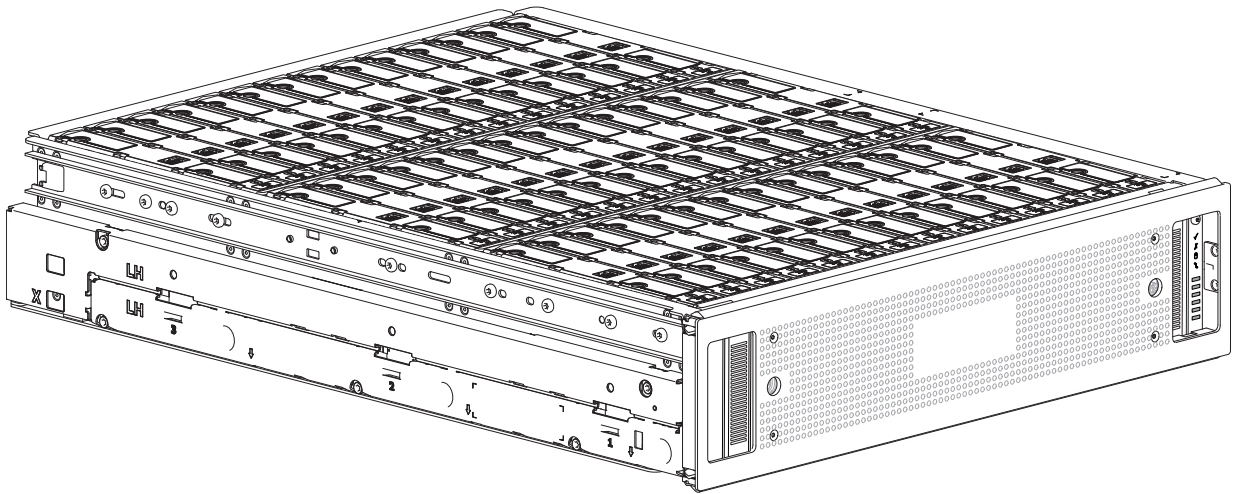
Figure 46 5U84 enclosure drawer fully populated with DDICs

Enclosure management

SBB IOMs actively manage the enclosure. Each module has a SAS expander with its own storage enclosure processor (SEP) that provides a SES target for a host to interface to through the ANSI SES Standard. If one of these modules fails, the other module will continue to operate.

Refer to a module's specification or the SES Interface specification for definitions of the module's functions and its SES control. See also "Management interfaces".

Management interfaces

RBOD/EBOD configuration

Upon completing the hardware installation, you can access the controller module's web-based management interface—Storage Management Console (SMC)—to configure, monitor, and manage the storage system. See also "Accessing the SMC" (page 90).

The controller module also provides a CLI in support of command entry and scripting.

JBOD configuration

Management of the enclosure is also performed by Generic Enclosure Management (GEM) firmware through the Ethernet port (10/100Base-TX) of each expansion module. The GEM firmware provides a built-in command line interface (GEM CLI) for interacting with the enclosure. See "Overview of GEM" (page 160) for more information.

3   Installation

Installation checklist

This chapter shows how to plan for and successfully install your enclosure system into an industry standard 19-inch rack cabinet.

⚠ **CAUTION:**   To install the system, use only the power cords supplied, or power cables that match the specification quoted in "AC power cords" (page 174).

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

ⓘ **IMPORTANT:**   Steps 1-7 in the following installation checklist are common to RBOD/EBOD/JBOD configurations. Steps 8-11 pertain only to RBOD/EBOD configurations. Step 12 pertains only to JBOD configurations.

Table 9   Installation checklist

Item	Step	Task	Where to find procedure
RBOD EBOD JBOD	1	Unpack the enclosure.	See "Unpacking the enclosure" (page 56).
	2	Install the controller enclosure and optional drive enclosures in the rack.[1]	See "Required tools" (page 57). See "Requirements for rackmount installation" (page 57). See "Installing the 2U enclosure" (page 58). See "Installing the 5U enclosure" (page 59).
	3	5U84 enclosure only: populate drawers with disks (DDIC–disk in drive carrier) 2U enclosures ship with disks installed.	See "FDE considerations" (page 60). See "Accessing drawers" (page 136). See "Populating drawers with DDICs" (page 51). See "Installing a DDIC" (page 139).
	4	Connect power cords.	See "Power cord connection" (page 64).
	5	Test enclosure connectivity.	See "Testing enclosure connections" (page 64).
	6	Install required host software.	See "Host system requirements" (page 65).
	7	Connect hosts.[2]	For RBOD/EBOD configuration, see "Connecting the controller enclosure to hosts" (page 65). For JBOD configuration, see "Storage system cabling" (page 149).

Table 9 Installation checklist (continued)

Item	Step	Task	Where to find procedure
RBOD EBOD	8	Connect remote management hosts.	See "Connecting a management host on the network" (page 74)
	9	Obtain IP values and set network port IP properties on the controller enclosure.	See "New user setup" (page 75). For USB CLI port and cable see "USB device connection".
	10	For CNC models, verify the host interface protocol setting (not necessary for SAS models).	See "CNC technology" (page 66), The CNC IOMs allow for setting the host interface protocol for qualified SFP options. See "Change the CNC port mode" (page 82).
	11	Perform initial configuration tasks:[3]	Topics below correspond to bullets at left:
		• Sign-in to the web-browser interface to access the application GUI.	See the "Getting Started" chapter in the Storage Management Guide.
		• Verify firmware revisions and update if necessary.	See "Updating firmware for RBOD/EBOD configurations" (page 75). Also see the same topic in the Storage Management Guide.
		• Initially configure and provision the system using the SMC.	See the topics about configuring the system and provisioning the system in the Storage Management Guide.
JBOD	12	Perform initial configuration tasks:[4]	Topics below correspond to bullets at left:
		• Using the IOM's Ethernet port, access the GEM CLI interface via Telnet.	See "Connecting to the GEM CLI" (page 161).
		• Verify firmware revisions and update if necessary.	See "Obtaining the JBOD firmware download" (page 156) and "JBOD firmware update procedure" (page 157).
		• Initially configure and provision the system using the host's RAID application.	See the documentation provided for host-based RAID controls.

1-The environment in which the enclosure operates must be dust-free to ensure adequate airflow.

2-For more information about hosts, see the About hosts topic in the Storage Management Guide.

3-The Storage Management Console is introduced in "Accessing the SMC" (page 90). See the Storage Management Guide or online help for additional information.

4-GEM is introduced in "Overview of GEM" (page 160).

Planning for installation

Before beginning the enclosure installation, familiarize yourself with the system configuration requirements. The figures listed below show the locations for each plug-in module:

- 2U12 front panel: see Figure 7 (page 25)
- 2U24 front panel: see Figure 8 (page 25)
- 2U controller enclosure rear panel: see Figure 9 and Figure 11 (page 26)
- 2U expansion enclosure rear panel: see Figure 12 (page 26) (product option)
- 5U front panel: see Figure 20 (eBezel) and Figure 20 (page 30) (drawer with slot index)
- 5U controller enclosure rear panel: see Figure 22 and Figure 23 (page 32)
- 5U expansion enclosure rear panel: see Figure 24 (page 32) (product option)

⊙ **IMPORTANT:**   Installation work should be performed by qualified service personnel.

Table 10 Storage system configuration

Module type	Location	Description
Drive carrier modules	2U Front panel	All drive slots must hold either a drive carrier or dummy drive carrier module. Empty slots are not allowed. At least one disk must be installed.
DDIC	5U Front panel drawers	Maximum 84 disks are installed (42 disks per drawer). Minimum 14 disks are required. Follow drawer population rules on page 51.
Power cooling modules	2U Rear panel	Two PCMs provide full power redundancy, allowing the system to continue to operate while a faulty PCM is replaced.
Power supply unit modules	5U Rear panel	Two PSUs provide full power redundancy, allowing the system to continue to operate while a faulty PSU is replaced.
Fan cooling modules	5U Rear panel	Five FCMs provide airflow circulation, maintaining all system components below the maximum temperature allowed.
I/O modules	Rear panel	Two IOMs must be installed for this configuration (RBOD, EBOD, JBOD).

NOTE: Although different disk drive modules and rear panel CRUs are used by the different 5005/4005/3005 Series enclosure form factors, whether used for RBOD, EBOD, or JBOD configuration, the IOMs are common across 2U and 5U enclosures. The only exception is the 2-port CNC module which is not used in the 5U84.

Preparing for installation

NOTE: Enclosure configurations

- 2U enclosures are delivered with CRUs and all drive carrier modules installed.
- 5U enclosures are delivered with many CRUs installed; however, DDICs must be installed during system setup.

△ **CAUTION:** Lifting enclosures

- A 2U enclosure—together will all its component parts—is too heavy for one person to lift and install into the rack cabinet. Two people are required to safely move a 2U enclosure.
- A 5U enclosure—delivered without DDICs installed—requires four people to lift it from the box. A suitable mechanical lift is required to hoist the enclosure for positioning in the rack.

Make sure you wear an effective anti-static wrist or ankle strap and obey conventional ESD precautions when touching modules and components. Do not touch midplane, motherboard, or module connectors. See also "ESD precautions" (page 116). This section provides important preparation requirements and handling procedures for use during product installation.

Preparing the site and host server

Before beginning the enclosure installation, verify that the site where you will install your storage system has the following:

- A standard AC power supply from a independent source or a rack power distribution unit with an Uninterruptible Power Supply (UPS).
- A host computer configured with the appropriate software, BIOS, and drives. Contact your supplier for the correct software configurations.

Before installing the enclosure, verify the existence of the following:

- Depending upon the controller module: SAS, FC, or iSCSI HBA and appropriate switches (if used)
- Qualified cable options for host connection
- One power cord per PCM or PSU
- Rail kit (for rack installation)

Please refer to your supplier for a list of qualified accessories for use with the enclosure. The accessories box contains the power cords and other accessories.

Unpacking the enclosure

1. Examine the packaging for crushes, cuts, water damage, or any other evidence of mishandling during transit.

    If you suspect that damage has happened, photograph the package before opening, for possible future reference. Retain original packaging materials for use with returns.

2. The unpacking sequence pertaining to 2U enclosures is shown in Figure 47.
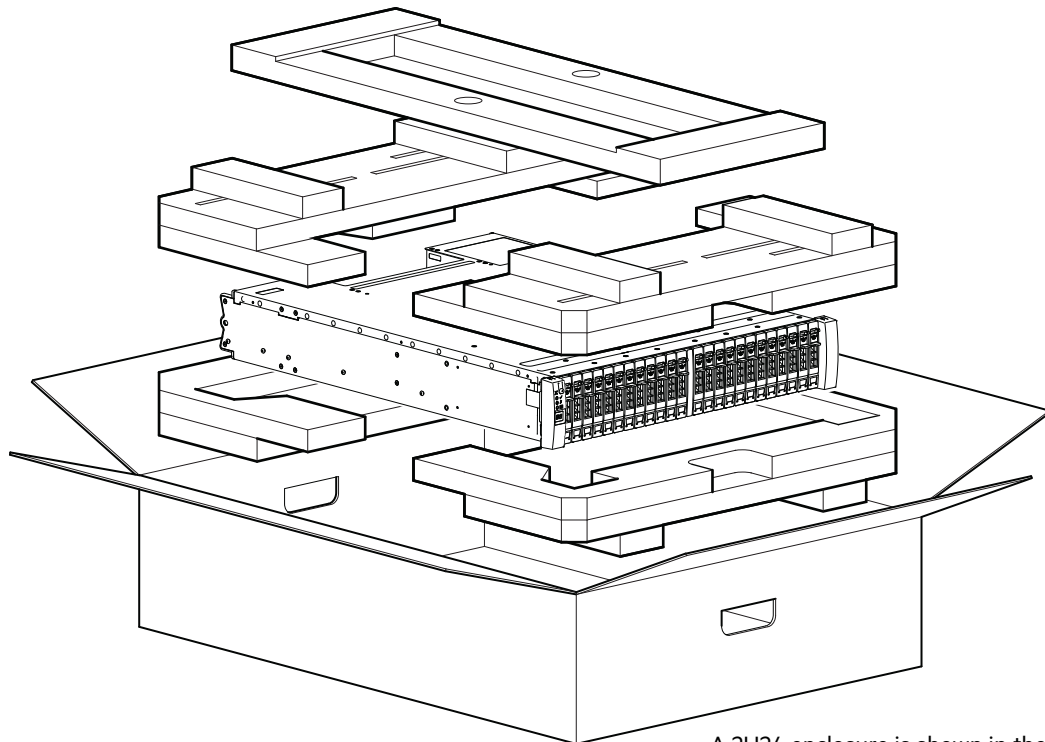
A 2U24 enclosure is shown in the example

Figure 47 Unpacking the 2U12 and 2U24 enclosures

2U enclosures are supplied with the midplane PCB and all plug-in modules installed. For information about plug-in module replacement, see "Module removal and replacement". Dummy drive carriers must be installed in unused drive slots.

3. The unpacking procedure pertaining to the 5U84 is shown in Figure 48.

⚠ **CAUTION:** The enclosure does not include DDICs, but all rear panel modules are installed. This partially populated enclosure is quite heavy: 64 kg (142 lb). Verify that each strap is securely wrapped and buckled. Four people are required to lift the 5U84 from the box.
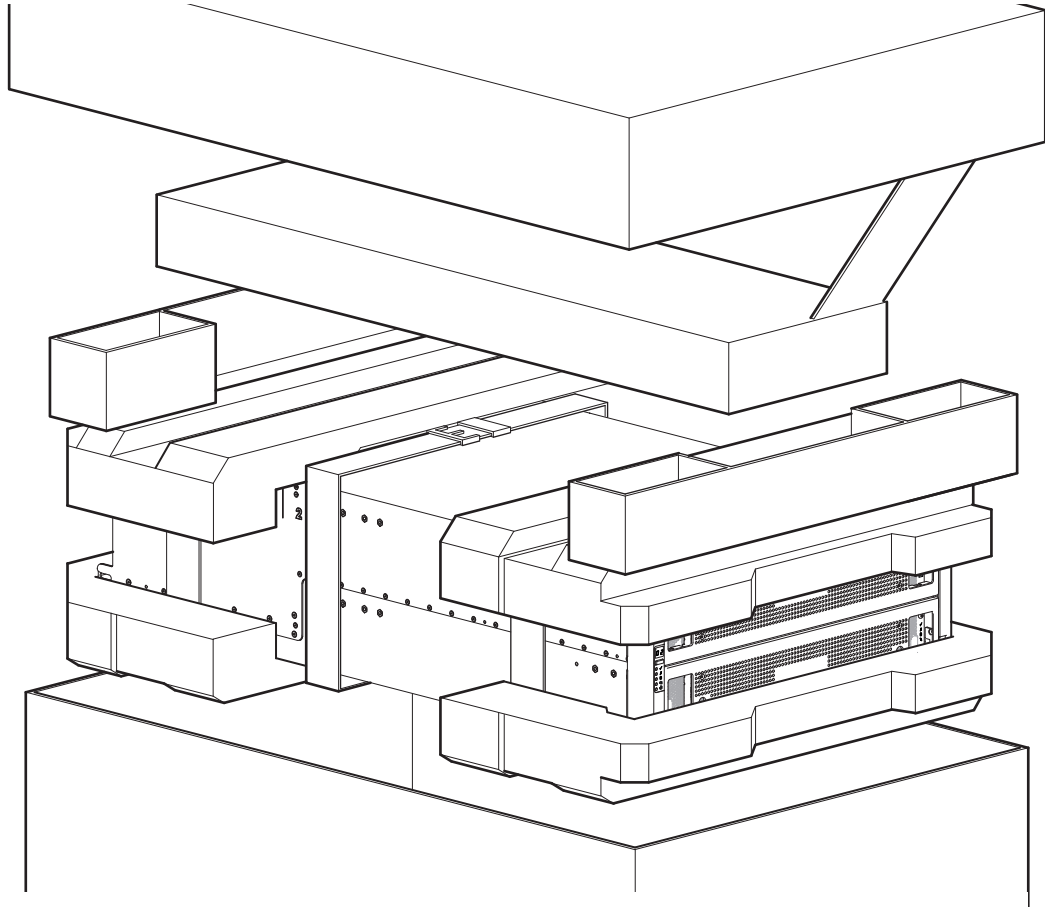
Figure 48 Unpacking the 5U84 enclosure

The railkit and accessories box is located immediately below the box lid in the illustration above.

△ **CAUTION:**   With four persons—positioned one at each corner of the enclosure—grip the straps securely by the loops, and lift the enclosure out of the box, using appropriate lifting technique. Place the enclosure in a static-protected area.

NOTE: If your product model uses CNC ports for FC or iSCSI, you must locate and install the SFPs. See also "Locate the SFP transceivers" (page 178).

Required tools

Flat blade screwdriver. Torx T10/T20 bits for locks and select CRU replacement.

Requirements for rackmount installation

You can install the enclosure in an industry standard 19-inch cabinet capable of holding 2U form factors.

- Minimum depth: 707 mm (27.83") from rack posts to maximum extremity of enclosure (includes rear panel cabling and cable bend radii).
- Weight:
  - ○ Up to 32 kg (71 lb), dependent upon configuration, per 2U enclosure.
  - ○ Up to 128 kg (282 lb), dependent upon configuration, per 5U enclosure.
- The rack should cause a maximum back pressure of 5 pascals (0.5 mm water gauge).
- Before you begin, ensure that you have adequate clearance in front of the rack for installing the rails.

Rackmount rail kit

Various sets of rack mounting rails are available for use in 19-inch rack cabinets. These rails have been designed and tested for the maximum enclosure weight, and to make sure that multiple enclosures may be installed without loss of space within the rack. Use of other mounting hardware may cause some loss of rack space. Contact your supplier to make sure suitable mounting rails are available for the rack you plan to use.

Installing the 2U enclosure

1. Remove the rack mounting rail kit from the accessories box, and examine for damage.

2. Use the procedure below to attach the rail kit brackets to the rack post as shown in Figure 49.

   a. Set the location pin at the rear of the rail into a rear rack post hole. Attach the bracket to the rear rack post: use the washers and screws supplied. Leave the screws loose.

   b. Extend the rail to fit between the front and rear rack posts.

   c. Attach the bracket to the front rack post using the washers and screws supplied. Leave the screws loose.

   d. Tighten the two clamping screws located along the inside of the rear section of the rack bracket.

   e. Repeat the above sequence of steps for the companion rail.
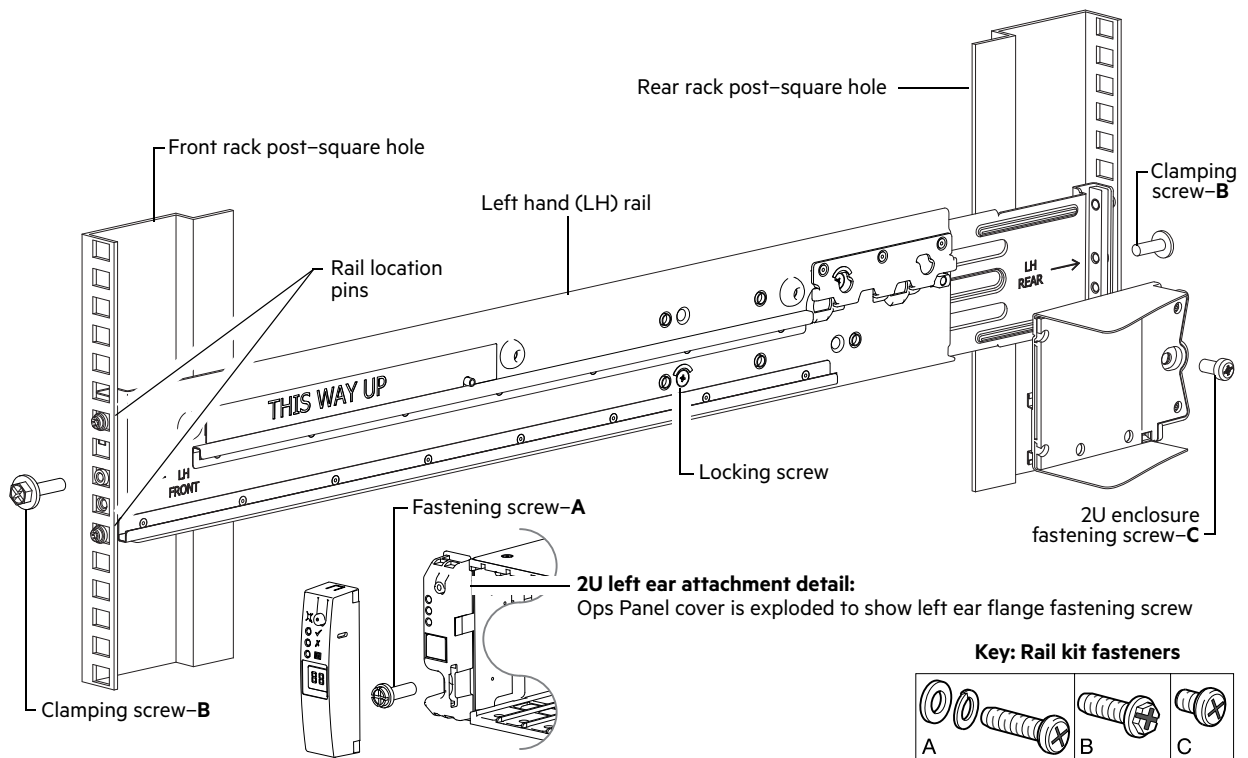
Figure 49 Securing brackets to the rail (left hand rail shown for 2U)

3. Install the enclosure into the rack:

   a. Lift the enclosure and align it with the installed rack rails, taking care to ensure that the enclosure remains level.

   b. Carefully insert the chassis slides into the rack rails and push fully in.

   c. Tighten the mounting screws in the rear rail kit brackets.

   d. Remove the enclosure until it reaches the hard stops—approximately 400 mm (15.75")—and tighten the mounting screws in the front rail kit bracket.

   e. Return the enclosure to the fully home position.

Installing the 5U enclosure

The 5U84 enclosure is delivered without the disks installed. Due to the weight of the enclosure, install it into the rack without DDICs installed, and remove the rear panel CRUs to lighten the enclosure weight.

The adjustment range of the railkit from the inside of the front post to the inside of the rear post is 660 mm – 840 mm. This range suits a one meter deep rack within Rack Specification IEC 60297.

1. To facilitate access, remove the door from the rack.

2. Ensure that the pre-assembled rails are at their shortest length (see also the reference label on the inside of the rail).

3. Locate the rail location pins inside the front of the rack, and extend the length of the rail assembly to enable the rear location pins to locate. Ensure the pins are fully located in the square or round holes in the rack posts. See also Figure 50.
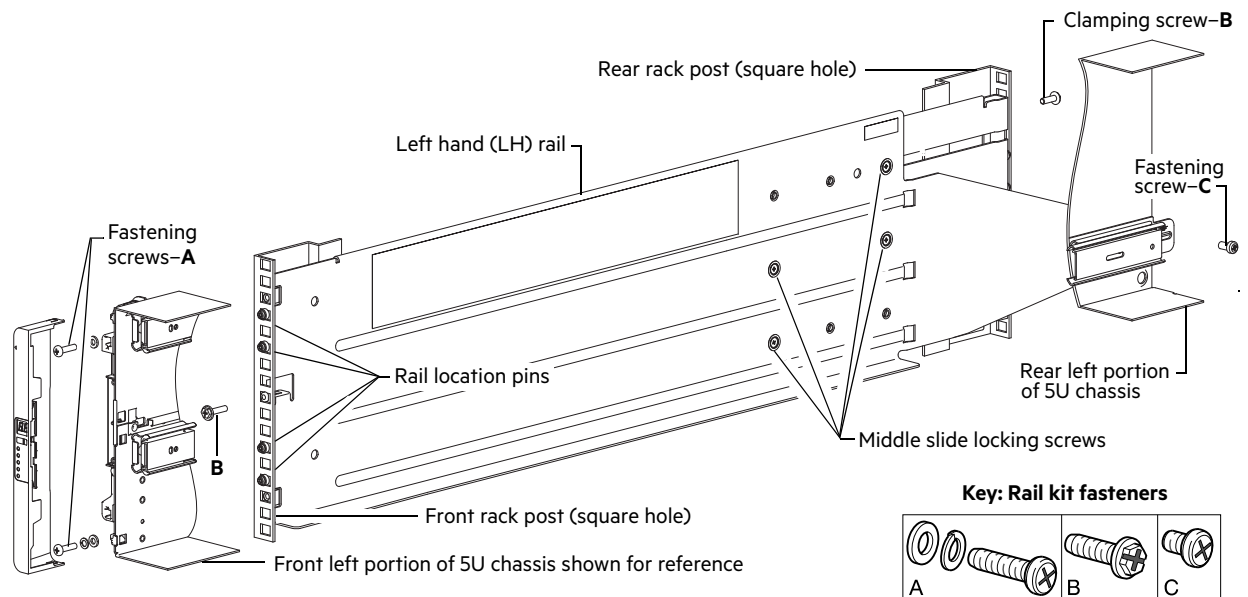
Figure 50 Securing brackets to the rail (left hand rail shown for 5U)

4. Fully tighten all clamping screws and middle slide locking screws.

5. Ensure the rear spacer clips (x4) are fitted to the edge of the rack post.

6. Slide the enclosure fully home on its rails.

7. Fasten the front of the enclosure using the enclosure fastening screws (x4) as shown in Figure 50.

8. Fix the rear of the enclosure to the sliding bracket with the rear enclosure fastening screws.

△ **CAUTION:**   Use only power cords supplied, or power cords that comply with the specifications in Table 44 (page 174).

△ **CAUTION:**   Once the enclosure is installed in the rack, dispose of the lifting straps. Due to the difficulty in attaching the straps once the enclosure is installed into the rack, the straps are not suitable for removing the enclosure from the rack.

Reinsert all of the rear panel modules, and install all of the DDICs into the drawers accessed from the front panel per the following instructions:

- Installing a fan cooling module "Installing an FCM" (page 143)
- Installing a PSU "Installing a PSU" (page 142)
- Installing an IOM "Installing an IOM" (page 145)
- Installing a DDIC "Installing a DDIC" (page 139)
  See also:
  - "Accessing drawers" (page 136)
  - "Populating drawers" (page 139)

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 SMC or CLI commands (see the Storage Management 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 procedures in "Module removal and replacement" for replacing disk drive modules relative to the enclosure type (2U or 5U). Import the keys for the disks so that the disk content becomes available. This feature applies only to RBOD/EBOD configurations.

Connecting the controller enclosure and optional expansion enclosures

This section describes cabling expansion enclosures to an RBOD. For information about cabling JBODs to host server HBAs, see "Connecting a JBOD" (page 149).

5005/4005/3005 Series controller enclosures—available in 24-drive (2.5") or 12-drive (3.5") chassis—support up to ten 2U enclosures (including the controller enclosure), or a maximum of 240 disk drives. 5005/4005 Series controller enclosures support the high-density 5U84 chassis, with a maximum of 84 disks in a carrier designed to hold either 2.5" or 3.5" disks. 5005/4005 Series supports up to four 5U enclosures (including the controller enclosure), or a maximum of 336 disk drives. The enclosures support both *straight-through* and *reverse* SAS cabling. Reverse cabling allows any drive enclosure to fail—or be removed—while maintaining access to other enclosures. Fault tolerance and performance requirements determine whether to optimize the configuration for high availability or high performance when cabling.

Cabling diagrams in this section show fault-tolerant cabling patterns. Controller and expansion modules are identified by enclosure ID and IOM ID, such as 0A and 0B for controller enclosures, 1A and 1B for the first expansion enclosure in a cascade, and so forth. When connecting multiple expansion enclosures, use reverse cabling to ensure the highest level of fault tolerance, enabling controllers to access remaining expansion enclosures if an expansion enclosure fails.

---

① **IMPORTANT:** Whereas 2U12 and 2U24 expansion enclosures can be intermixed, 2U expansion enclosures cannot be intermixed with 5U84 expansion enclosures in the same storage system.

Cable requirements for expansion enclosures used with a controller enclosure

When adding storage, use only Seagate or OEM-qualified cables, and observe the following guidelines:

- When installing SAS cables to expansion modules, use only supported HD mini-SAS x4 cables.
- Qualified HD mini-SAS to HD mini-SAS 0.5 m (1.64') cables are used to connect cascaded enclosures in the rack.

- 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 cables, depending upon number of enclosures and cabling method used.
- You may need to order additional or longer cables when reverse-cabling a fault-tolerant configuration.

The rear panel view of the 2U12 and 2U24 controller enclosures are nearly identical to one another. The rear panel views of the expansion enclosures are also nearly identical to one another.

Whether configured as an RBOD or EBOD, the rear panel view of the 5U84 is very different than the 2U12 or 2U24. Although the 5U84 uses the same SBB IOMs used in the 2U enclosures, the remaining CRUs accessible from the enclosure rear panel differ from those used in the 2U enclosures.

NOTE: For clarity, the schematic diagrams show only relevant details such as IOM face plate outlines and expansion ports. For detailed illustrations see "2U enclosure rear panel" (page 26) and "5U enclosure rear panel" (page 31).
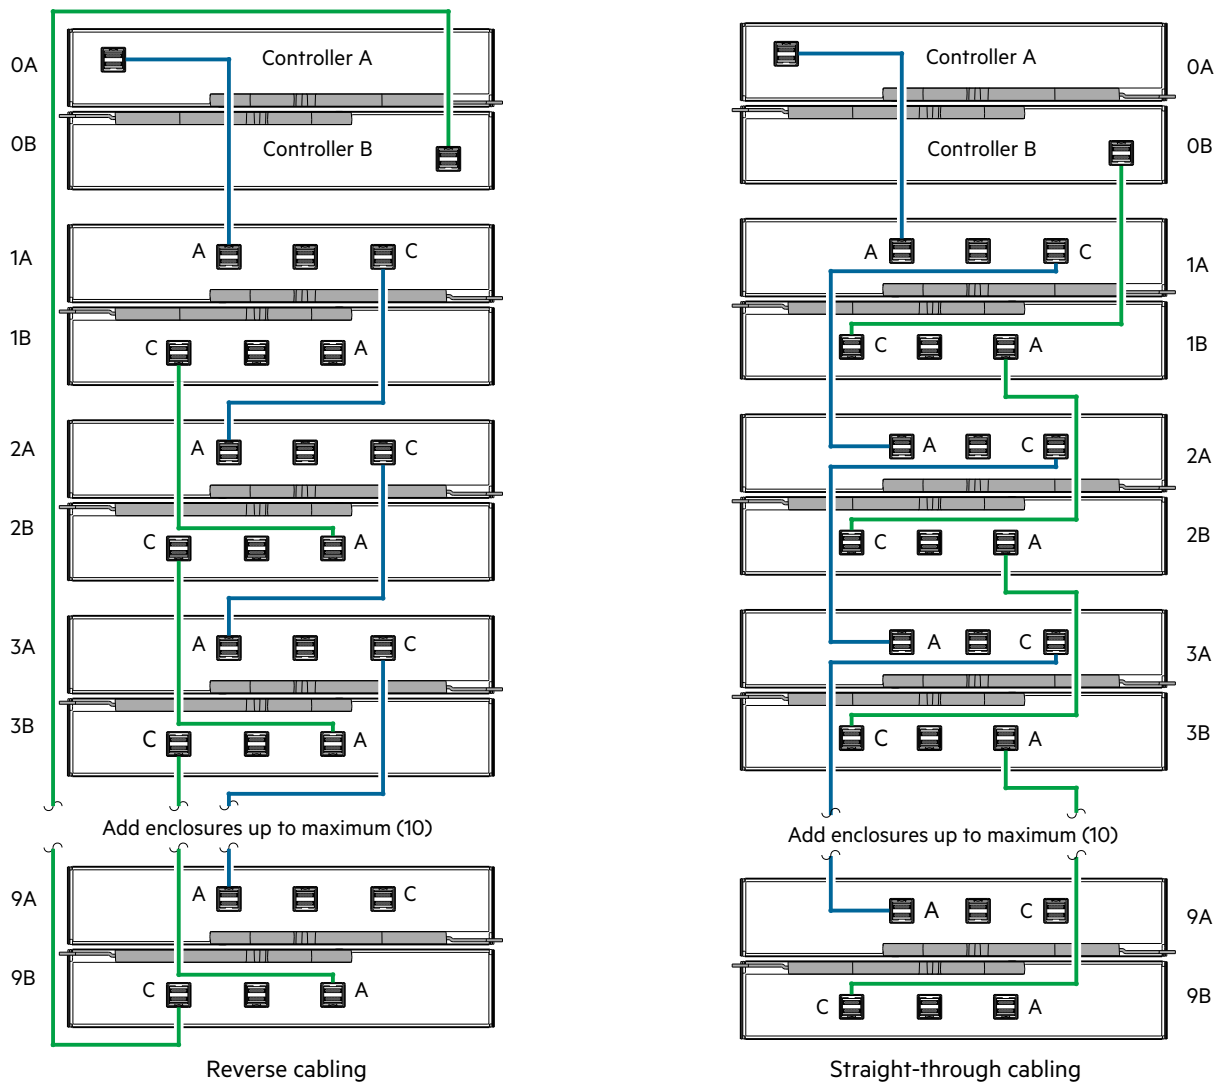
Figure 51 Cabling connections between a 2U controller enclosure and 2U expansion enclosures

The previous diagram (left) shows reverse cabling of a 5005/4005/3005 Series dual-controller 2U enclosure and supported 2U drive enclosures configured with dual expansion modules. Controller module 0A is connected to expansion module 1A, with a chain of connections cascading down (blue). Controller module 0B is connected to the lower expansion

module (9B), of the last expansion enclosure, with connections moving in the opposite direction (green). Reverse cabling allows any expansion enclosure to fail—or be removed—while maintaining access to other enclosures.

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

The 2U drive enclosures shown in the above figure can either be of the same type or they can be a mixture of the J1212 and J1224 models. Given that supported drive enclosure models use 12Gb/s SAS link-rate and SAS 3.0 expanders, they can be ordered in desired sequence within the system, following the controller enclosure. The middle SAS ports on expansion modules are not used. See also Figure 16 (page 28) and the **IMPORTANT** entry beneath that figure. Refer to these diagrams when cabling multiple expansion enclosures together with 5005/4005/3005 Series 2U controller enclosures.

5005/4005 Series does not support intermixing 2U and 5U84 enclosures when configuring a storage system. Representative examples showing supported configurations and configuration limits are provided above for 2U, and on the following pages for 5U84.
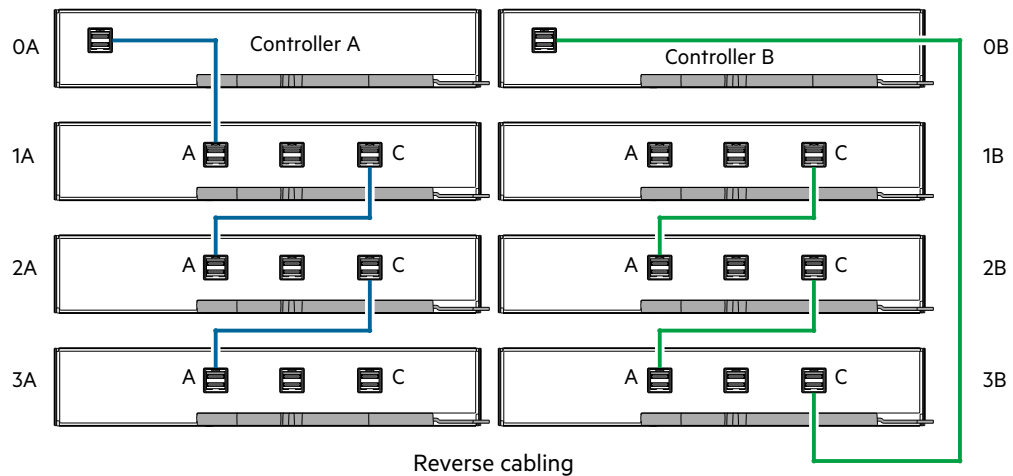
Reverse cabling

Figure 52 Cabling connections between a 5U controller enclosure and 5U expansion enclosures–reverse

Figure 52 and Figure 53 (page 63) show maximum configuration cabling for a 5005/4005 Series 5U controller enclosure with 5U expansion enclosures (4 enclosures including the controller enclosure).

NOTE: See also Figure 16 (page 28) and the **IMPORTANT** entry beneath that figure. Refer to these diagrams when cabling multiple 5U expansion enclosures together with 5005/4005 Series 5U controller enclosures.

Straight-through cabling
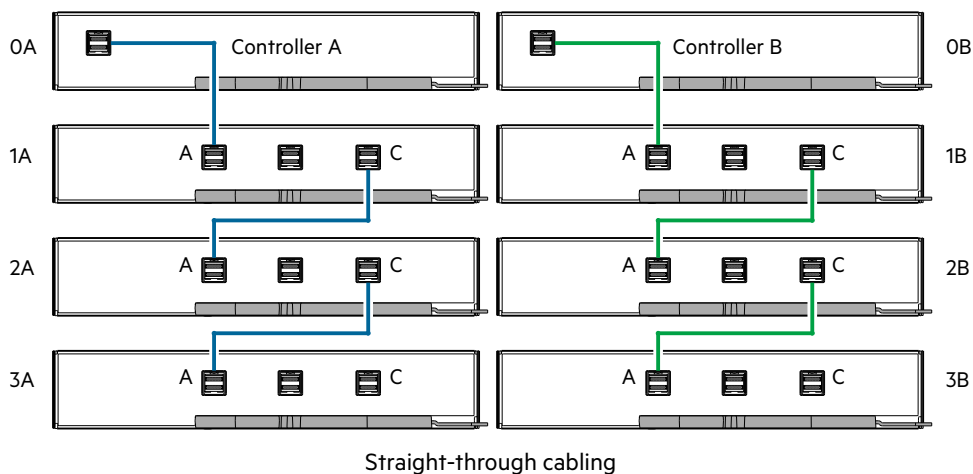
Figure 53 Cabling connections between a 5U controller enclosure and 5U expansion enclosures–straight-through

The comparative characteristics of using either reverse cabling or straight-through cabling methods are discussed beneath Figure 51 (page 61).
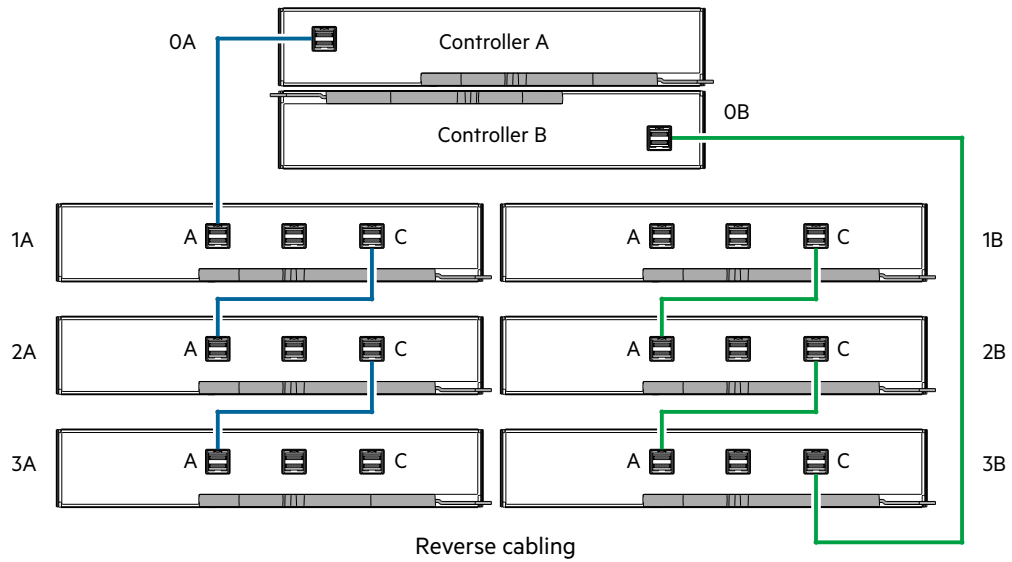
Reverse cabling

Figure 54 Cabling connections between a 2U controller enclosure and 5U expansion enclosures–reverse

Figure 54 and Figure 55 show maximum configuration cabling for a 5005/4005 Series 5U controller enclosure with 5U expansion enclosures (4 enclosures including the controller enclosure).

NOTE: See also Figure 16 (page 28) and the **IMPORTANT** entry beneath that figure. Refer to these diagrams when cabling multiple 5U expansion enclosures together with 5005/4005 Series 2U controller enclosures.

Straight-through cabling
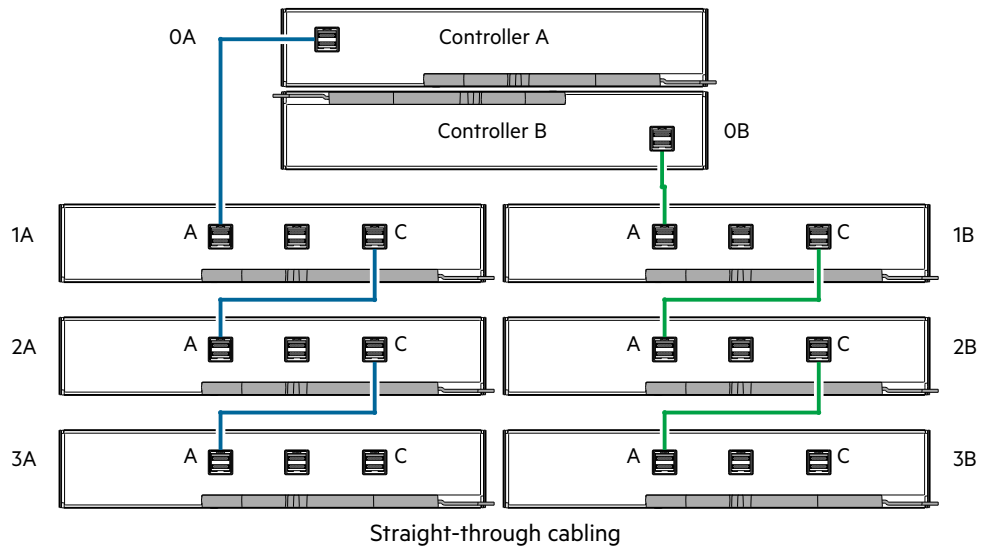
Figure 55 Cabling connections between a 2U controller enclosure and 5U expansion enclosures–straight-through

The comparative characteristics of using either reverse cabling or straight-through cabling methods are discussed beneath Figure 51 (page 61).

Power cord connection

Connect a power cord from each PCM or PSU on the enclosure rear panel to the PDU (power distribution unit) as shown in the illustrations below (4-port SAS models are shown in the examples).
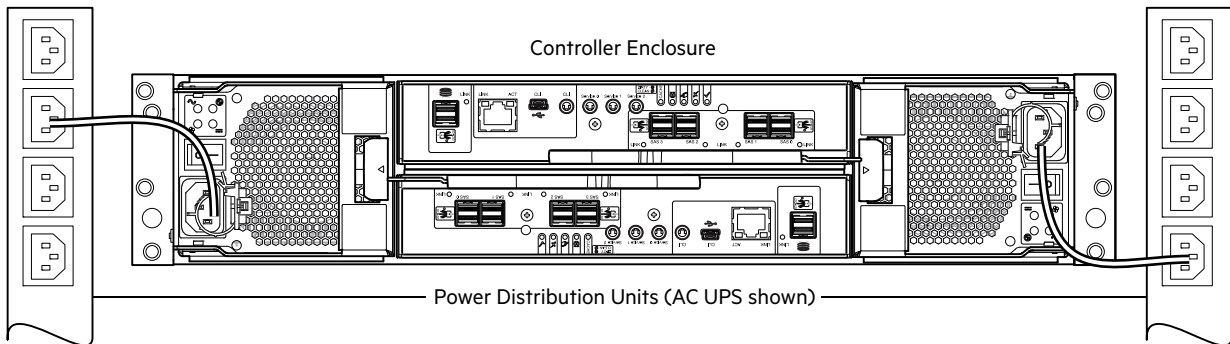
Figure 56 Typical AC power cord connection from PDU to PCM (2U)
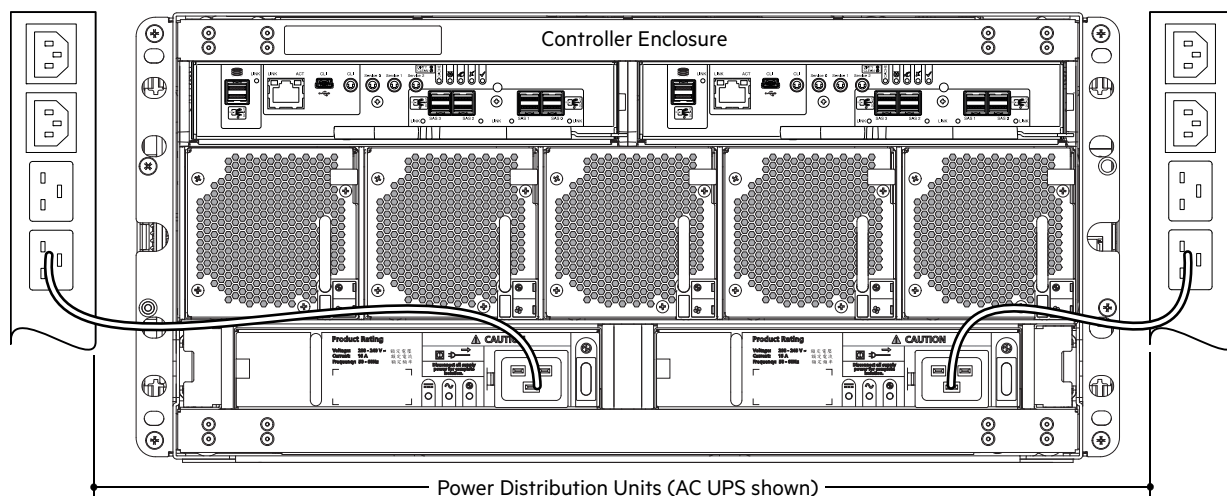
Figure 57 Typical AC power cord connection from PDU to PSU (5U)

ⓘ **IMPORTANT:** When more than one PCM (2U) or PSU (5U) is fitted, all power cords must be connected to at least two separate and independent power supplies to ensure redundancy. When the storage system is ready for operation, ensure that each PCM or PSU power switch is set to the **On** position. See also "Powering on/powering off" (page 88).

△ **CAUTION:** Power connection concerns:

- Always remove the power connections *before* you remove the PCM (2U) or PSU (5U) from the enclosure.
- When bifurcated power cords (Y leads) are used, these cords must only be connected to a supply range of 200–240V AC.

Testing enclosure connections

See "Powering on/powering off" (page 88). Once the power-on sequence succeeds, the storage system is ready to be connected as described in "Connecting the controller enclosure to hosts" (page 65) for RBOD/EBOD configurations, and "Storage system cabling" (page 149) for JBOD configurations.

Grounding checks

The product must only be connected to a power source that has a safety electrical earth connection.

△ **CAUTION:** If more than one enclosure goes in a rack, the importance of the earth connection to the rack increases because the rack will have a larger Earth Leakage Current (Touch Current).

Examine the earth connection to the rack before power on. An electrical engineer who is qualified to the appropriate local and national standards must do the examination.

Host system requirements

Hosts connected to a 5005/4005/3005 Series controller enclosure or JBOD must meet the following requirements:

Depending on your system configuration, host operating systems may require that multipathing is supported.

If fault tolerance is required, then multipathing software may be required. Host-based multipath 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 2008, Windows Server 2012, and Windows Server 2016. Use either the Server Manager or the *mpclaim* CLI tool to perform the installation.

See the following web sites for information about using native Microsoft MPIO DSM:

https://support.microsoft.com

https://technet.microsoft.com (search the site for "multipath I/O overview")

NOTE: For information about connecting a JBOD to a host server's HBAs, see "Storage system cabling" (page 149).

Cabling considerations for RBOD/EBOD configurations

Common cabling configurations address hosts, controller enclosures, expansion enclosures, and switches. Host interface ports on 5005/4005/3005 Series controller enclosures can connect to respective hosts via direct-attach or switch-attach. Cabling systems to enable use of the optional replication feature—to replicate volumes—is yet another important cabling consideration. See "Connecting a management host on the network" (page 74) for RBOD/EBOD configurations, and "Storage system cabling" (page 149) for JBOD configurations. The FC and iSCSI product models can be licensed to support replication, whereas the HD mini-SAS models cannot.

Use only Seagate or OEM-qualified cables for host connection:

- Qualified Fibre Channel SFP and cable options
- Qualified 10GbE iSCSI SFP and cable options
- Qualified 1Gb RJ-45 SFP and cable options
- Qualified HD mini-SAS cable options

Connecting the controller enclosure to 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. This section describes host interface protocols supported by 5005/4005/3005 Series controller enclosures, while showing a few common cabling configurations.

NOTE: 5005/4005/3005 Series controllers use Unified LUN Presentation (ULP), which enables a host to access mapped volumes through any controller host port.

ULP can show all LUNs through all host ports on both controllers, and the interconnect information is managed by the controller firmware. ULP appears to the host as an active-active storage system, allowing the host to select any available path to access the LUN, regardless of disk group ownership.

> ☼ **TIP:** See the topic about configuring system settings in the Storage Management Guide to initially configure the system, or change system configuration settings (such as Configuring host ports).

CNC technology

The 5005/4005/3005 Series FC/iSCSI models use Converged Network Controller (CNC) 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 CNC ports are further described in the subsections below. Also see "CNC ports used for host connection" (page 14) for more information concerning use of CNC ports.

> **NOTE:** Controller modules are *not* shipped with pre-installed SFPs. Within your product kit, you will need to locate the qualified SFP options and install them into the CNC ports. See "Install an SFP transceiver" (page 179).

> ⓘ **IMPORTANT:** Use the `set host-port-mode` CLI command to set the host interface protocol for CNC ports using qualified SFP options. 5005/4005/3005 Series models ship with CNC ports configured for FC. When connecting CNC ports to iSCSI hosts, you must use the CLI (not the SMC) to specify which ports will use iSCSI. It is best to do this before inserting the iSCSI SFPs into the CNC ports (see "Change the CNC port mode" (page 82) for instructions).

Fibre Channel protocol ⇄

5005/4005/3005 Series FC controller enclosures support two controller modules using the Fibre Channel interface protocol for host connection. Each 5005/4005 FC controller module provides four host ports. CNC ports are designed for use with an FC SFP supporting data rates up to 16Gb/s.

The controllers support Fibre Channel Arbitrated Loop (public or private) or point-to-point topologies. Loop protocol can be used in a physical loop or for 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 16Gb/s. See the `set host-parameters` command within the CLI Reference Guide for command syntax and details about parameter settings relative to supported link speeds. Fibre Channel ports are used for attachment to FC hosts directly, or through a switch used for the FC traffic. The host computer must support FC and optionally, multipath I/O.

The Fibre Channel ports are used in either of two capacities:

- To connect two storage systems through a switch for use of 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 replication feature, whereas the second option requires that the host computer supports Ethernet, FC, and optionally, multipath I/O.

> ☼ **TIP:** Use the SMC to set FC port speed. Within the Storage Management Guide, see the topic about configuring host ports. 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.

Fibre Channel zoning

The host should be cabled to two different Fibre Channel switches for redundancy. Alternatively, a single Fibre Channel switch is supported, but it introduces a single point of failure. When using Fibre Channel for front-end connectivity, zones must be established to ensure that storage is visible to the servers,

> **CAUTION:** There are two types of FC zoning: port zoning and WWN zoning. Port zoning uses the switch port number as the zone member, while WWN zoning uses each device's WWN as the zone member. WWN zoning is recommended since a device may reside on any port, or change physical ports and still be visible.

It is recommended to create zones using a single initiator host port and multiple system ports. This means one zone for each server HBA port, that includes the HBA's WWN and all the storage WWNs connected to that switch.

iSCSI settings

The host should be cabled to two different Ethernet switches for redundancy. If using switches with mixed traffic (LAN/iSCSI), then VLAN should be used to isolate iSCSI traffic from the rest of the switches.

10GbE iSCSI protocol ⬚

5005/4005/3005 Series 10GbE iSCSI controller enclosures support two controller modules using the Internet SCSI interface protocol for host connection. Each 5005/4005 10GbE iSCSI controller module provides four host ports. CNC ports are designed for use with a 10GbE iSCSI SFP supporting data rates up to 10Gb/s, using either one-way or mutual CHAP (Challenge-Handshake Authentication Protocol).

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

- To connect two storage systems through a switch for use of 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 replication feature, whereas the second option requires that the host computer supports Ethernet, iSCSI, and optionally, multipath I/O.

> **TIP:** See the topic about configuring CHAP in the Storage Management Guide.

> **TIP:** Use the SMC to set iSCSI port options. Within the Storage Management Guide, see the topic about configuring host ports. 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.

1Gb iSCSI protocol ⬚

5005/4005/3005 Series 1Gb iSCSI controller enclosures support two controller modules using the Internet SCSI interface protocol for host port connection. Each 5005/4005 1Gb iSCSI controller module provides four host ports. The CNC ports are designed for use with an RJ-45 SFP supporting data rates up to 1Gb/s, using either one-way or mutual CHAP.

> **TIP:** See the topic about configuring CHAP in the Storage Management Guide.

> **TIP:** Use the SMC to set iSCSI port options. Within the Storage Management Guide, see the topic about configuring host ports. 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 1Gb iSCSI ports are used in either of two capacities:

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

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

SAS protocol

5005/4005 SAS models use 12Gb/s host interface protocol and qualified cable options for host connection.

12Gb HD mini-SAS host ports

5005/4005 12Gb SAS controller enclosures support two controller modules. The 5005/4005 12Gb/s SAS controller module provides four SFF-8644 HD mini-SAS host ports. These host ports support data rates up to 12Gb/s. HD mini-SAS host ports are used for attachment to SAS hosts directly. The host computer must support SAS and optionally, multipath I/O. Use a qualified cable option when connecting to a host.

Host connection for RBOD/EBOD configurations

5005/4005 controller enclosures support up to eight direct-connect server connections, four per controller module. Connect appropriate cables from the server's HBAs to the controller module's host ports as described below, and shown in the following illustrations.

Fibre Channel host connection

To connect controller modules supporting (4/8/16Gb) FC host interface ports to a server HBA or switch, using the controller's CNC ports, select a qualified FC SFP option.

Use the cabling diagrams to connect the host servers to the switches. See the Support Matrix for supported Fibre Channel HBAs.

Install and connect each FC HBA to a switch that is connected to the host ports on the two controllers as shown in the "Switch attach" section of this chapter.

In hybrid examples, one server and switch manages FC traffic, and the other server and switch manages iSCSI traffic.

For FC, each initiator must be zoned with a single host port or multiple host ports only (single initiator, multi-target of the same kind).

Connecting host servers directly to the storage system is also supported.

Qualified options support cable lengths of 1 m (3.28'), 2 m (6.56'), 5 m (16.40'), 15 m (49.21'), 30 m (98.43'), and 50 m (164.04') for OM4 multimode optical cables and OM3 multimode FC cables, respectively. A 0.5 m (1.64') cable length is also supported for OM3. In addition to providing host connection, these cables are used for connecting two storage systems via a switch, to facilitate use of the optional replication feature.

10GbE iSCSI host connection

To connect controller modules supporting 10GbE iSCSI host interface ports to a server HBA or switch, using the controller's CNC ports, select a qualified 10GbE SFP option.

Qualified options support cable lengths of 0.5 m (1.64'), 1 m (3.28'), 3 m (9.84'), 5 m (16.40'), and 7 m (22.97') for copper cables; and cable lengths of 0.65 m (2.13'), 1 m (3.28'), 1.2 m (3.94'), 3 m (9.84'), 5 m (16.40'), and 7 m (22.97') for direct attach copper (DAC) cables. In addition to providing host connection, these cables are used for connecting two storage systems via a switch, to facilitate use of the optional replication feature.

1Gb iSCSI host connection

To connect controller modules supporting 1Gb iSCSI host interface ports to a server HBA or switch, using the controller's CNC ports, select a qualified 1Gb RJ-45 copper SFP option supporting (CAT5-E minimum) Ethernet cables of the same

lengths specified for 10GbE iSCSI above. In addition to providing host connection, these cables are used for connecting two storage systems via a switch, to facilitate use of the optional replication feature.

HD mini-SAS host connection

To connect controller modules supporting HD mini-SAS host interface ports to a server HBA, using the controller's SFF-8644 dual HD mini-SAS host ports, select a qualified HD mini-SAS cable option.

A qualified SFF-8644 to SFF-8644 cable option is used for connecting to a 12Gb/s enabled host; whereas a qualified SFF-8644 to SFF-8088 cable option is used for connecting to a 6Gb/s host. Qualified SFF-8644 to SFF-8644 options support cable lengths of 0.5 m (1.64'), 1 m (3.28'), 2 m (6.56'), and 4 m (13.12'). Qualified SFF-8644 to SFF-8088 options support cable lengths of 1 m (3.28'), 2 m (6.56'), 3 m (9.84'), and 4 m (13.12').

Connecting direct attach RBOD/EBOD configurations

A dual-controller configuration improves application availability because in the event of a controller failure, the affected controller fails over to the healthy partner controller with little interruption to data flow. A failed controller can be replaced without the need to shut down the storage system. The 5005/4005/3005 Series enclosures are configured with dual controller modules.

NOTE: In the examples that follow, a single diagram represents CNC and SAS host connections for controller enclosures. The location and sizes of the host ports are very similar. Blue cables show controller A paths and green cables show controller B paths for host connection.

5005/4005 Series

Server

Controller enclosure

0A

0B
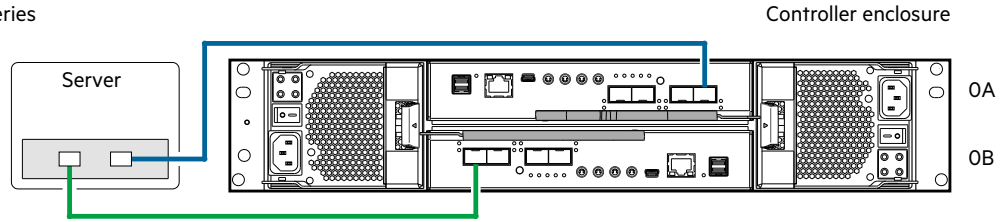
Figure 58 Connecting hosts: 2U direct attach – one server/ one HBA/ dual path

3005 Series

Server

Controller enclosure

0A

0B
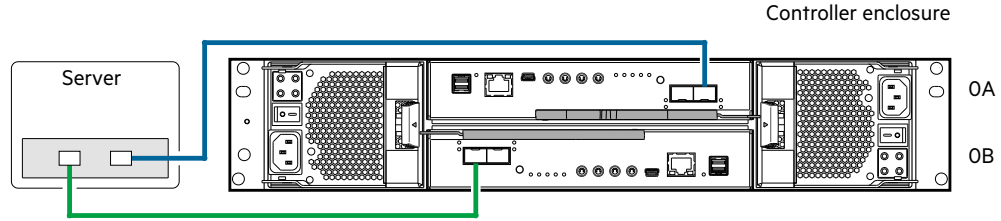
Figure 59 Connecting hosts: 2U direct attach – one server/ one HBA/ dual path

5005/4005 Series

0A   0B

Server

Controller enclosure
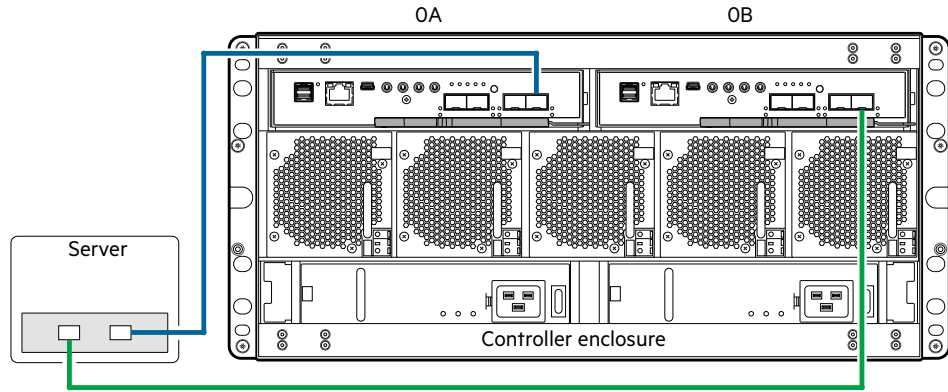
Figure 60 Connecting hosts: 5U direct attach – one server/ one HBA per server/ dual path

5005/4005 Series   Controller enclosure

Server 1   0A   Server 2

0B
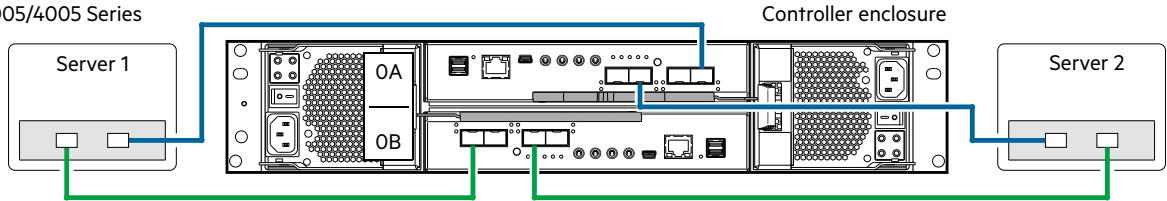
Figure 61 Connecting hosts: 2U direct attach – two servers/ one HBA per server/ dual path

3005 Series   Controller enclosure

Server 1   0A   Server 2

0B
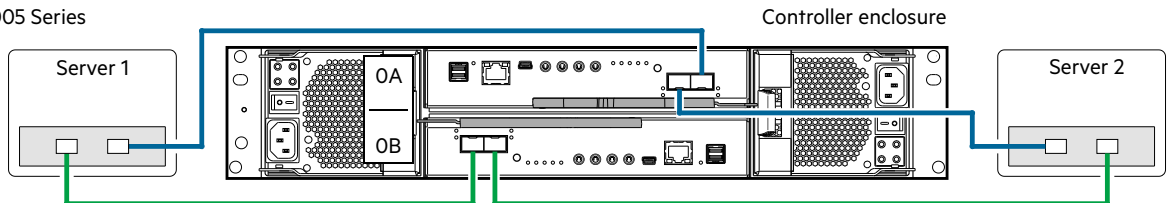
Figure 62 Connecting hosts: 2U direct attach – two servers/ one HBA per server/ dual path
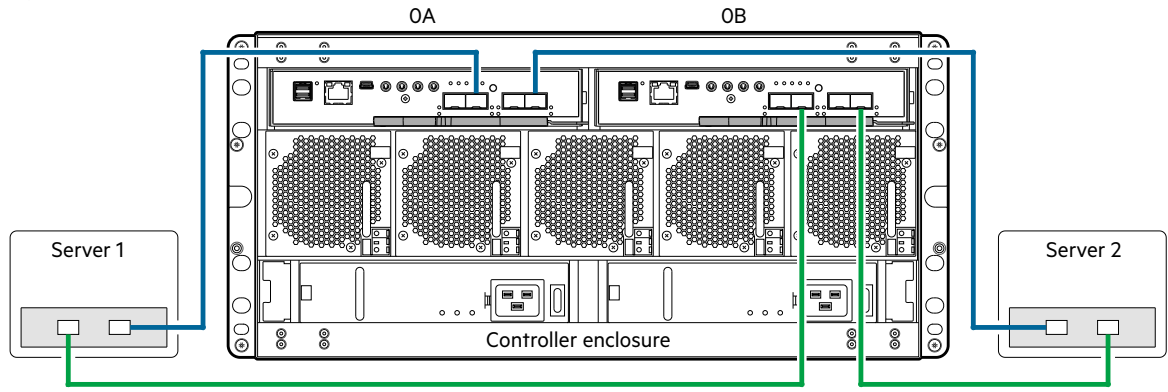
5005/4005 Series

Figure 63 Connecting hosts: 5U direct attach – two servers/ one HBA per server/ dual path
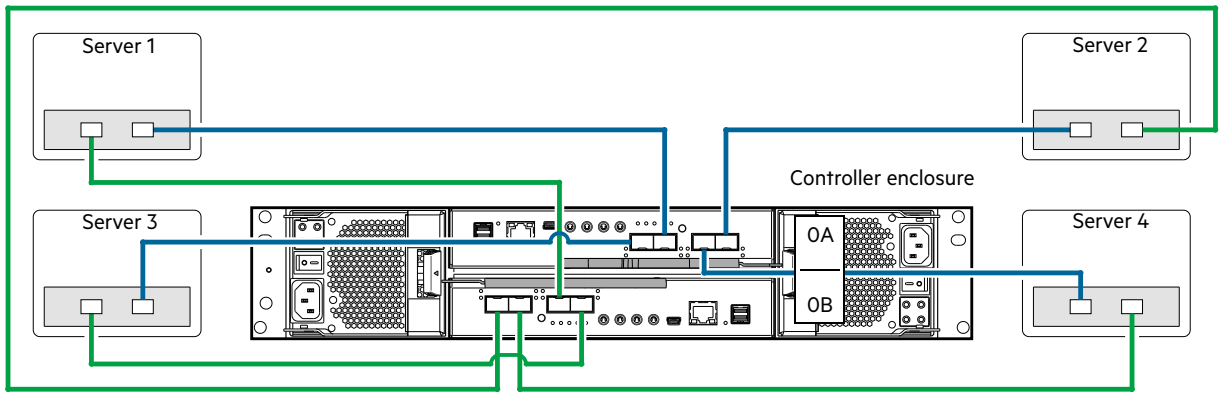
5005/4005 Series

Figure 64 Connecting hosts: 2U direct attach – four servers/ one HBA per server/ dual path

5005/4005 Series
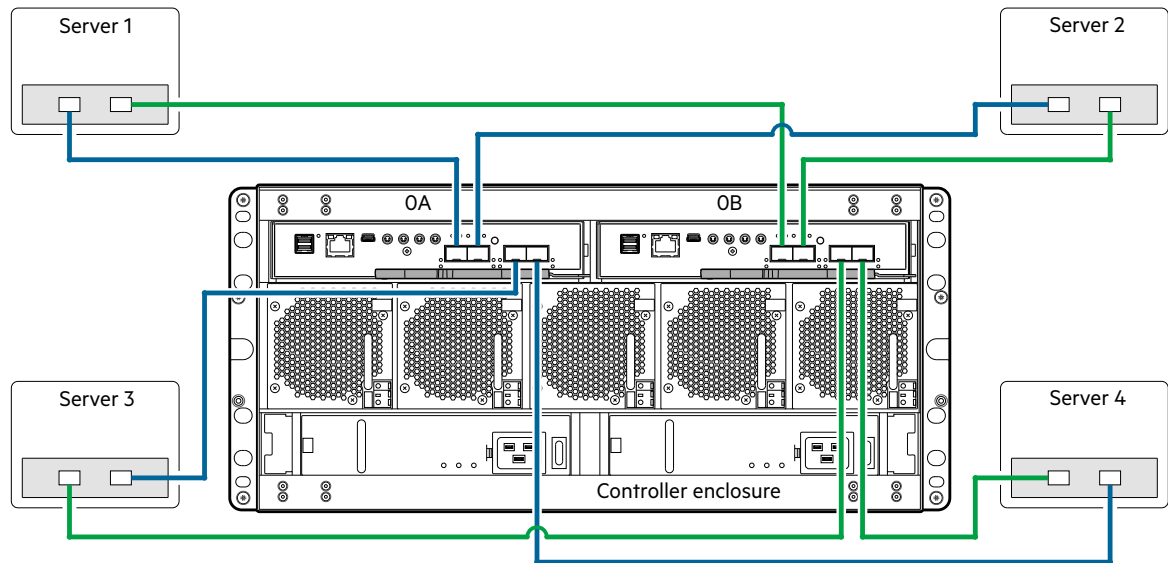
Figure 65 Connecting hosts: 5U direct attach – four servers/ one HBA per server/ dual path

Switch attach

A switch attach solution—or SAN—places a switch between the servers and the controller enclosures within the storage system. Using switches, a SAN shares a storage system among multiple servers, reducing the number of storage systems required for a particular environment. Using switches increases the number of servers that can be connected to the storage system.

5005/4005 Series
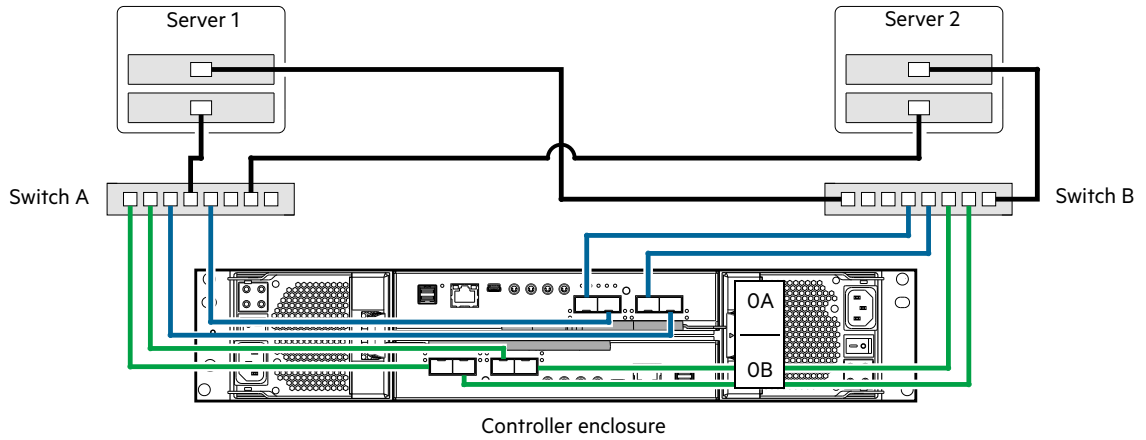
Figure 66 Connecting hosts: 2U switch attach – two servers/ two switches

3005 Series
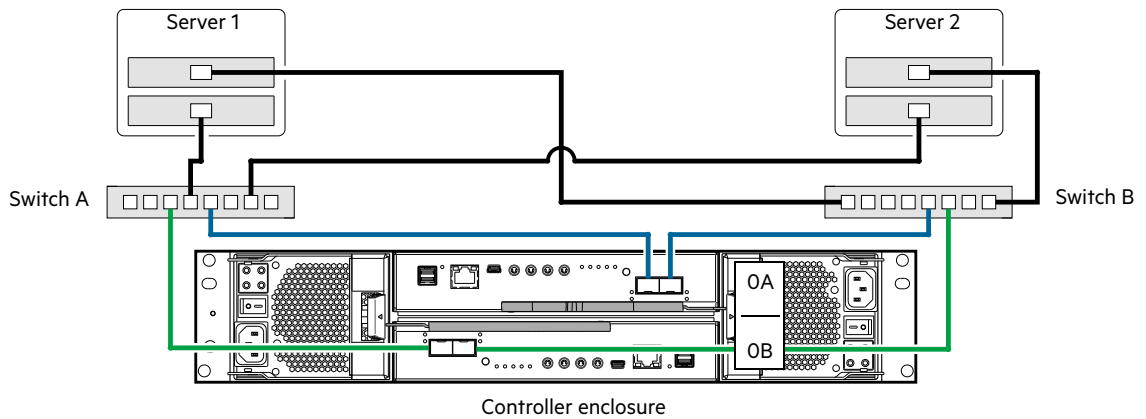
Figure 67 Connecting hosts: 2U switch attach – two servers/ two switches

5005/4005 Series
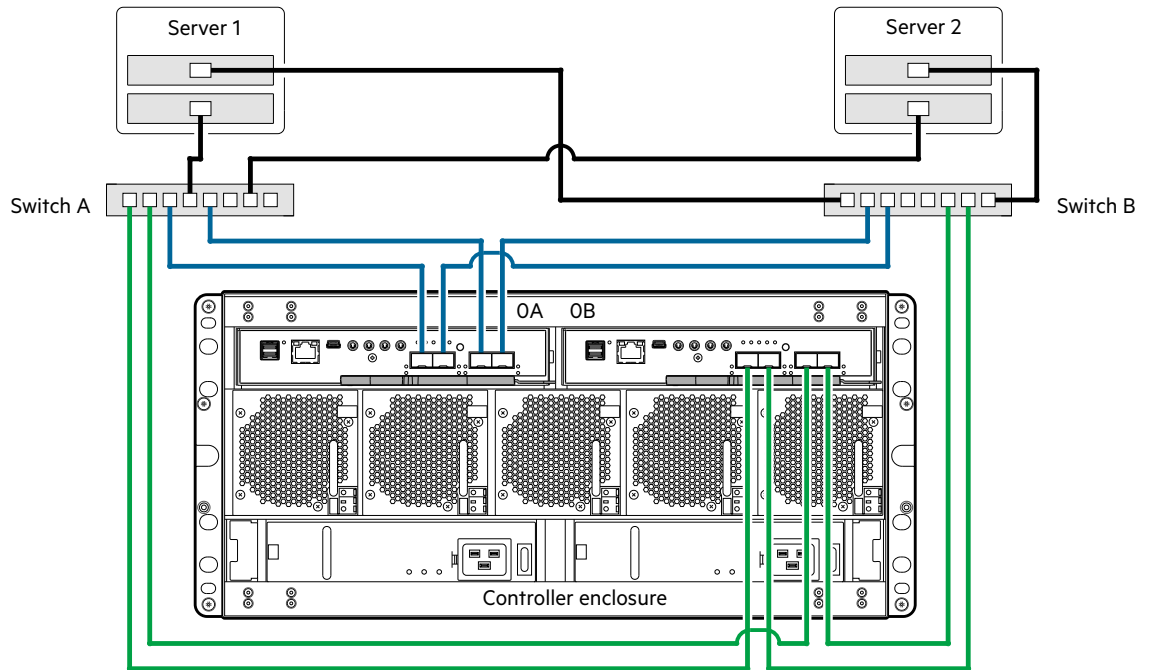
Figure 68 Connecting hosts: 5U switch attach – two servers/ two switches

5005/4005 Series
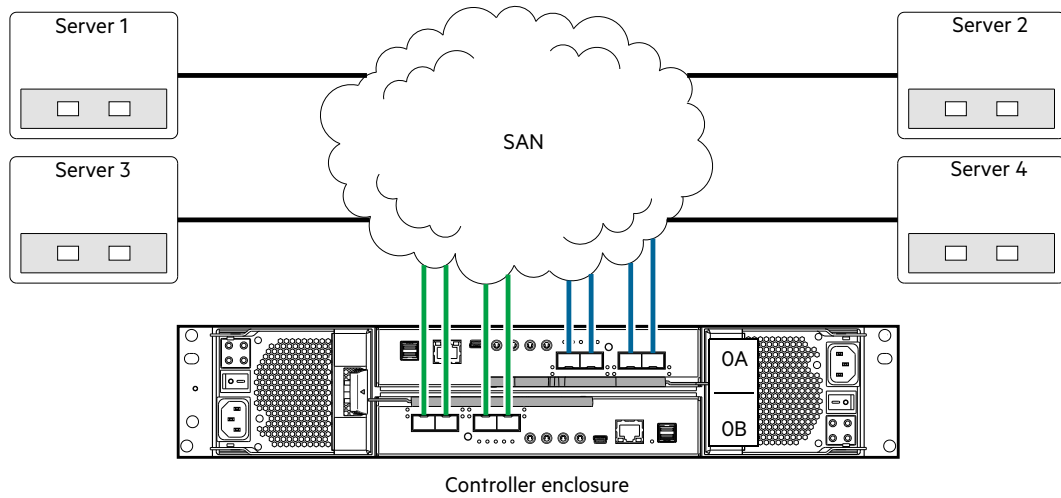
Figure 69 Connecting hosts: 2U switch attach – four servers/ multiple switches/ SAN fabric
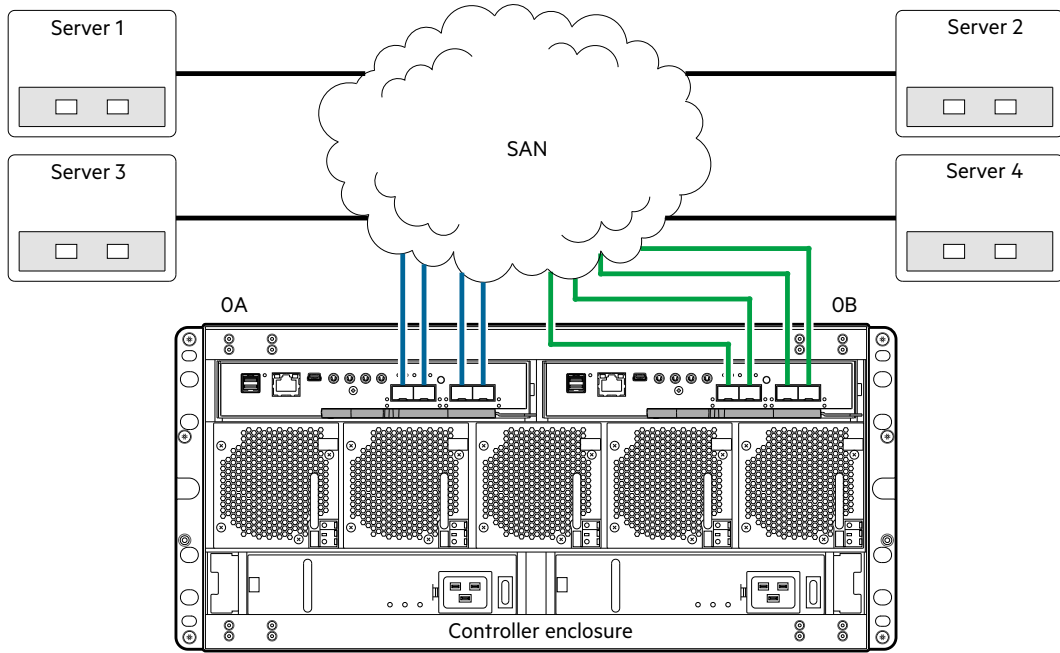
Figure 70 Connecting hosts: 5U switch attach – four servers/ multiple switches/ SAN fabric

Connecting a management host on the network

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

1. Connect an RJ-45 Ethernet cable to the network port on each controller.

2. Connect the other end of each Ethernet cable to a network that your management host can access (preferably on the same subnet). Do not interconnect iSCSI and management Ethernet on the same network.

Controller enclosure
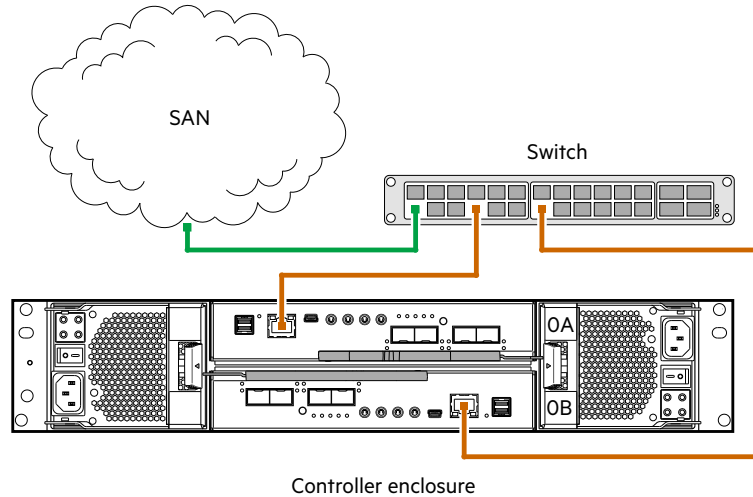
Figure 71 Connect a 2U controller enclosure to the management network

Controller enclosure (partial view)
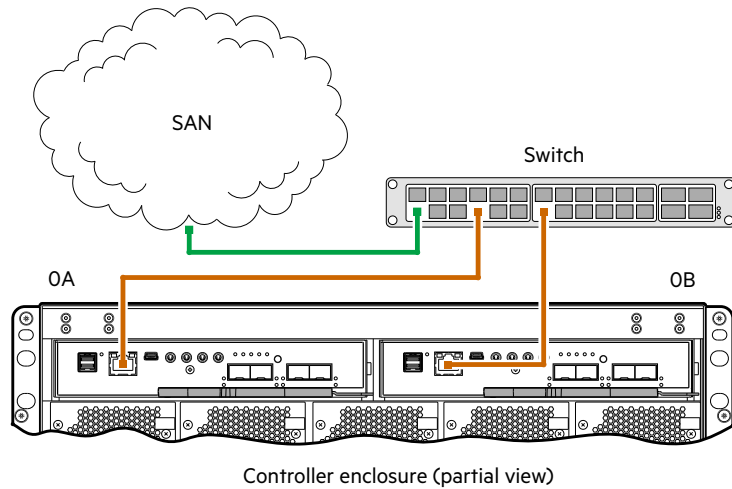
Figure 72 Connect a 5U84 controller enclosure to the management network

---

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 FCC Rules and Regulations.
- If you connect the iSCSI and management ports to the same physical switches, separate VLANs are recommended.
- See also the topic about configuring controller network ports in the Storage Management Guide.

Updating firmware for RBOD/EBOD configurations

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.

For RBOD/EBOD configurations, using the Storage Management Console, in the **System** topic, select **Action > Update Firmware**. The Update Firmware panel opens. The Update Controller Module tab shows versions of firmware components currently installed in each controller.

For JBOD configurations, see .

NOTE: The SMC provides an option for enabling or disabling Partner Firmware Update for the partner controller. To enable or disable the setting, use the `set advanced-settings` CLI command, and set the `partner-firmware-upgrade` parameter. See the CLI Reference Guide for more information about command parameter syntax.

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

ⓘ **IMPORTANT:**   See the topic about updating firmware in the Storage Management Guide before performing a firmware update. Partner Firmware Update (PFU) is enabled by default on 5005/4005/3005 Series systems.

New user setup

This product does not provide default users. When you first connect to the system, you will be prompted to set up a new user with *manage* capability.

> **NOTE:** At start up, the CLI (Serial/SSH) and WBI (HTTPS) are the only interfaces through which the system can be configured until the new user is created.

If you connect to the CLI console via the USB port, the following interactive shell output will be presented for creating a new user.

```
Welcome to Storage Manager

Use the following procedure to complete the initial setup for using this system.

Do you want to continue? [y|n]   y

Step -1: You must create a new user to manage this system. You need to specify a username and password.

The new user will be created with default preferences. Do you want to continue? [y/n]   y

Username input rules:

- The value is case sensitive.

- The value can have a maximum of 29 characters.

- The value can include printable UTF-8 characters except a space or: \" \' , < > \\

Password input rules:

- The value is case sensitive.

- The value can have 8-32 characters.

- The value can include printable UTF-8 characters except a space or: \" \' , < > \\

Enter username > <new username>

Enter password > <password>

Confirm Password > <password>

Thank you. New user has been created successfully.

The setup process has been completed.

You will be disconnected now. Please login again with your new user account.
```

Obtaining IP values

> **NOTE:** You can manually set static IP values (default method) for each controller, or you can specify that IP values should be set automatically for both controllers through communication with a Dynamic Host Configuration Protocol (DHCP) server.

> **TIP:** See the topic about configuring network ports in the Storage Management Guide.

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. You must have some means of determining what addresses have been assigned, such as the list of bindings on the DHCP server.

Setting network port IP addresses using the CLI port and cable

If you did not use DHCP to set network port IP values, set them manually as described below.

- You can use the 3.5mm stereo plug CLI port and the supplied 3.5mm/DB9 serial cable.

  –or–

- You can use the generic mini-USB cable (not included) and the USB CLI port. If you plan on using a mini-USB cable, you will need to enable the port for communication using DB9 serial cable. See also "Using the CLI port and cable—known issues on Windows" (page 177).

Network ports on controller module A and controller module B are configured with the following default values:

- **Network 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 CLI embedded in each controller module. The CLI enables you to access the system using the USB (Universal Serial Bus) communication interface and terminal emulation software.

NOTE:

- If you are using the mini USB CLI port and cable, see Appendix D - "USB device connection":
- Unless using Windows 10 or Server 2016, Windows customers should download and install the device driver as described in "Obtaining the software download" (page 176).
- Linux customers should prepare the USB port as described in "Setting parameters for the device driver" (page 177).

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 SMC. Be sure to change the IP address before changing the network configuration. See "Accessing the SMC" (page 90) for more information concerning the web-based storage management application.

1. From your network administrator, obtain an IP address, subnet mask, and gateway address for controller A and another for controller B.
2. Record these IP addresses so you can specify them whenever you manage the controllers using the SMC or the CLI.

   –or–
3. 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 73 (generic controller module is shown).

Within the illustration below, the IOM is aligned for use in the bottom (0B) controller module slot in a 2U enclosure.
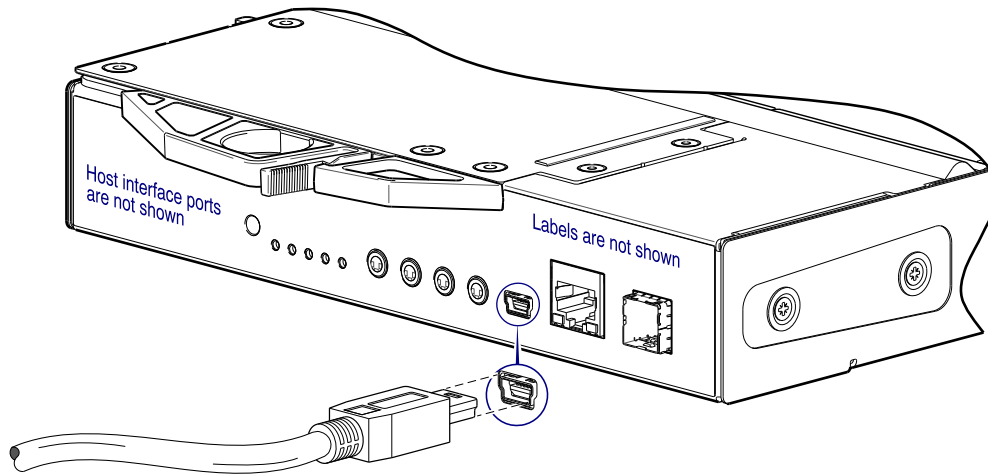
Figure 73 Connecting a USB cable to the CLI port

4. If using a mini-USB cable, enable the CLI port for subsequent communication:

    Note: Skip this step if using the 3.5mm/DB9 serial cable.

    o   Linux customers should enter the command syntax provided in "Setting parameters for the device driver" (page 177).

    o   Windows customers should locate the downloaded device driver described in "Obtaining the software download" (page 176), and follow the instructions provided for proper installation.

5. Start and configure a terminal emulator, such as HyperTerminal or VT-100, using the display settings used in Table 11 (page 78), and the connection settings in Table 12 (page 78) (also, see the note following this procedure).

Table 11   Terminal emulator display settings

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

Table 12   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.

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

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

    The CLI displays the system version, Management Controller version, and login prompt.

NOTE: Follow the new user setup process to create the first user on the system. See "New user setup" (page 75)

8. To use DHCP to set network port IP values, enter the following command at the prompt:

   `# set network-parameters dhcp`

   –or–

9. To use custom static IP addresses, enter the following CLI command at the prompt 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`

10. Enter the following CLI 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.

11. Use the ping command to verify connectivity to the gateway address.

    For example:

    `# ping 192.168.0.1`

    A success message will say that the remote computer responded with four packets.

12. 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, you might need to restart the Management Controller(s) using the serial CLI.

    When you restart a Management Controller, communication with it is temporarily lost until it successfully restarts.

    Enter the following CLI command to restart the Management Controller in both controllers:

    `restart mc both`

    > ⓘ **IMPORTANT:** When configuring an iSCSI system or a system using a combination of FC and iSCSI SFPs, do *not* restart the Management Controller or exit the terminal emulator session until configuring the CNC ports as described in "Change the CNC port mode" (page 82).

13. When you are done using the CLI, exit the emulator.
14. Retain the IP addresses (recorded in step 1) for accessing and managing the controllers using the SMC 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 telnet 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 CLI cable as described in step 3 and Figure 73 (page 78).

Configure controller network ports

IMPORTANT: If you used the default 10.0.0.2/10.0.0.3 addresses to access the guided setup, it is important to consider changing those IPv4 addresses to avoid an IP conflict should you have more than one ME4 Series array on your network.

You can manually set static IP addressing parameters for network ports or you can specify that IP values be set automatically using DHCP for IPv4 or Auto for IPv6, which uses DHCPv6 and/or SLAAC. When setting IP values, you can choose either IPv4 or IPv6 formatting for each controller. Additionally, you can set the addressing mode and IP version differently for each controller and use them concurrently. For example, you could set IPv4 on controller A to Manual to enable static IP addressing, and IPv6 on controller B to Auto to enable automatic IP addressing.

When using DHCP mode, the system obtains values for the network port IP address, subnet mask, and gateway from a DHCP server if one is available. If a DHCP server is unavailable, current addressing is unchanged. You must have some means of determining what addresses have been assigned, such as the list of bindings on the DHCP server. When using Auto mode, addresses are retrieved from both DHCP and Stateless address auto-configuration (SLAAC). DNS settings are also automatically retrieved from the network.

Each controller has the following factory-default IP settings:

- IP address source: DHCP
- Controller A IP address: 10.0.0.2
- Controller B IP address: 10.0.0.3
- IP subnet mask: 255.255.255.0
- Gateway IP address: 10.0.0.1

When DHCP is enabled in the storage system, the following initial values are set and remain set until the system is able to contact a DHCP server for new addresses:

- Controller IP addresses: 169.254.x.x (where the value of x.x is the lowest 16 bits of the controller serial number)
- IP subnet mask: 255.255.0.0
- Gateway IP address: 10.0.0.0

169.254.x.x addresses (including gateway 169.254.0.1) are on a private subnet that is reserved for unconfigured systems and the addresses are not routable. This prevents the DHCP server from reassigning the addresses and possibly causing a conflict where two controllers have the same IP address. As soon as possible, change these IP values to proper values for your network.

For IPv6, when Manual mode is enabled you can enter up to four static IP addresses for each controller. When Auto is enabled, the following initial values are set and remain set until the system is able to contact a DHCPv6 and/or SLAAC server for new addresses:

- Controller A IP address: fd6e:23ce:fed3:19d1::1
- Controller B IP address: fd6e:23ce:fed3:19d1::2
- Gateway IP address: fd6e:23ce:fed3:19d1::3

CAUTION: Changing IP settings can cause management hosts to lose access to the storage system after the changes are applied in the confirmation step.

Set IPv4 values for network ports

1. Perform one of the following to access Network options:
   o In the **Home** topic, select **Action** > **System Settings**, and then click the **Network** tab.
   o In the **System** topic, select **Action** > **System Settings**, and then click the **Network** tab.
2. Select the **IPv4** tab. IPv4 uses 32-bit addresses.
3. Select the type of IP address to use for each controller. Choose **Source** > **manual** to enter static IP addresses or choose **Source** > **DHCP** to have the system automatically obtain values from a DHCP server.
4. If you chose manual, enter the unique IP address, IP mask, and gateway values for each controller, then record the IP values you assign.

NOTE: The following IP addresses are reserved for internal use by the storage system: 169.254.255.1, 169.254.255.2, 169.254.255.3, 169.254.255.4, and 127.0.0.1. Because these addresses are routable, do not use them anywhere in your network.

5. Perform one of the following:
   o To save your settings and continue configuring your system, click **Apply**.
   o To save your settings and close the panel, click **Apply and Close**.

   A confirmation panel appears.
6. Click **OK** to continue. If you chose DHCP and the controllers successfully obtained IP values from the DHCP server, the new IP values appear. Record the new addresses and sign out to use the new IP address to access the MESM.

Set IPv6 values for network ports1

1. Perform one of the following to access network options:
   o In the **Home** topic, select **Action** > **System Settings**, and then click the **Network** tab.
   o In the **System** topic, select **Action** > **System Settings**, and then click the **Network** tab.
2. Select the **IPv6** tab. IPv6 uses 128-bit addresses.
3. Select the type of IP address to use for each controller. Choose **Source** > **manual** to enter up to four static IP addresses for each controller, or choose **Source** > **auto** to have the system automatically obtain values.
4. If you chose manual, perform the following:
   a. Enter the unique IP address, gateway value, and address for each controller.
   b. Record the IP values you assign.
   c. Click **Add**.
5. Click **Add Address** to continue adding up to four IP addresses.

NOTE: The following IP addresses are reserved for internal use by the storage system: 169.254.255.1, 169.254.255.2, 169.254.255.3, 169.254.255.4, and 127.0.0.1. Because these addresses are routable, do not use them anywhere in your network.

6. Perform one of the following:

- o To save your settings and continue configuring your system, click **Apply**.
- o To save your settings and close the panel, click **Apply and Close.**

  A confirmation panel appears.

7. Click **Yes** to save your changes. Otherwise, click **No**.

Sign out to use the new IP addresses to access the MESM.

Change the CNC port mode

This subsection applies to 5005/4005/3005 Series FC/iSCSI models only. While the USB cable is still connected and the terminal emulator session remains active, perform the following steps to change the CNC port mode from the default setting (FC), to either iSCSI or FC-and-iSCSI used in combination.

When using FC SFPs and iSCSI SFPs in combination, host ports 0 and 1 are set to FC (either both 16Gb/s or both 8Gb/s), and host ports 2 and 3 must be set to iSCSI (either both 10GbE or both 1Gb/s).

Set CNC port mode to iSCSI

To set the CNC port mode for use with iSCSI SFPs, run the following CLI command at the command prompt:

`set host-port-mode iSCSI`

The command notifies you that it will change host port configuration, stop I/O, and restart both controllers. When asked if you want to continue, enter **y** to change the host port mode to use iSCSI SFPs.

Once the `set host-port-mode` CLI command completes, it will notify you that the specified system host port mode was set, and that the command completed successfully.

Continue with .

Set CNC port mode to FC and iSCSI

To set the CNC port mode for use with FC SFPs and iSCSI SFPs in combination, run the following CLI command at the command prompt:

`set host-port-mode FC-and-iSCSI`

The command notifies you that it will change host port configuration, stop I/O, and restart both controllers. When asked if you want to continue, enter **y** to change the host port mode to use FC and iSCSI SFPs.

Once the `set host-port-mode` CLI command completes, it will notify you that the specified system host port mode was set, and that the command completed successfully.

Continue with .

Connecting two controller enclosures to replicate volumes

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 standard 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 the replication feature, 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 the CNC-based systems (SAS systems do not support replication) 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 the paragraph above Figure 74 (page 85).
- For remote replication, ensure that all ports assigned for replication are able to communicate appropriately with the 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 are provided addressing cabling on the same network and cabling relative to different networks.

ⓘ **IMPORTANT:** The replication feature must be licensed on all systems configured for replication, and the controller module firmware must be compatible on all systems used for replication. See the topic about licensed features in the Storage Management Guide.

Cabling for replication

This section shows example replication configurations for CNC-based controller enclosures.

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.

- Replication supports FC and iSCSI host interface protocols.
- The 2U enclosure rear panel represents 5005/4005 (4-port) models, and 3005 (2-port) models.
- The 5U enclosure rear panel represents 5005/4005 Series (4-port) 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 CNC-based systems are physically cabled, see the Storage Management Guide or online help for information about configuring, provisioning, and using the optional replication feature with out-of-band management.

Host ports and replication

5005/4005/3005 Series 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 0 and 1 are set to FC (either both 16Gb/s or 8Gb/s), and ports 2 and 3 must be set to iSCSI (either both 10GbE or both 1Gb). FC and iSCSI ports can be used to perform I/O or replication.

IMPORTANT: 5005/4005/3005 Series enclosures support dual-controller configuration only. If a partner controller fails, the storage system will fail over and run on a single controller module 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 5005/4005/3005 Series controller enclosures equipped with dual-controller modules. Some of the diagrams show 2U enclosures, while others show 5U enclosures. 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.

IMPORTANT: 5005/4005/3005 Series controller enclosures support FC and iSCSI host interface protocols for host connection or for performing replications.

Multiple servers/single network

Figure 74 and Figure 75 (page 85) shows the rear panel of two 2U enclosures with I/O and replication occurring on the same physical network. Figure 76 (page 85) shows the rear panel of two 5U enclosures with I/O and replication occurring on the same physical network.

With the replication configurations shown in each of these figures, 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. Either or these configurations would appear physically as a single network, while logically, either configuration would function as multiple networks.
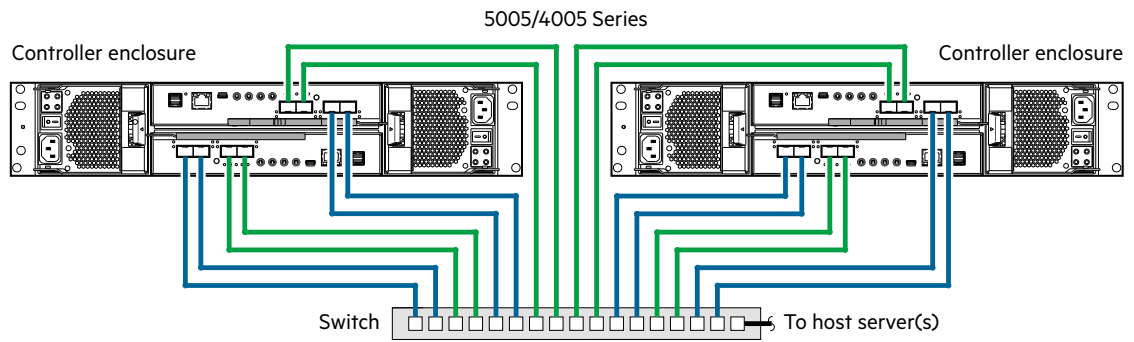
Figure 74 Connecting two 2U storage systems for replication: multiple servers/one switch/one location
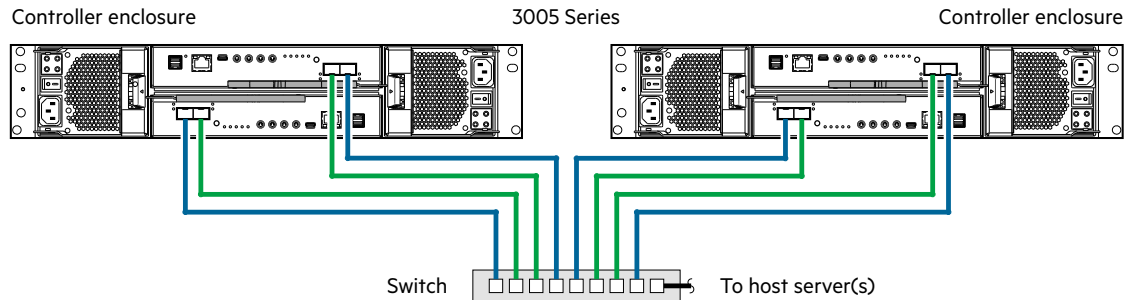
Figure 75 Connecting two 2U storage systems for replication: multiple servers/one switch/one location
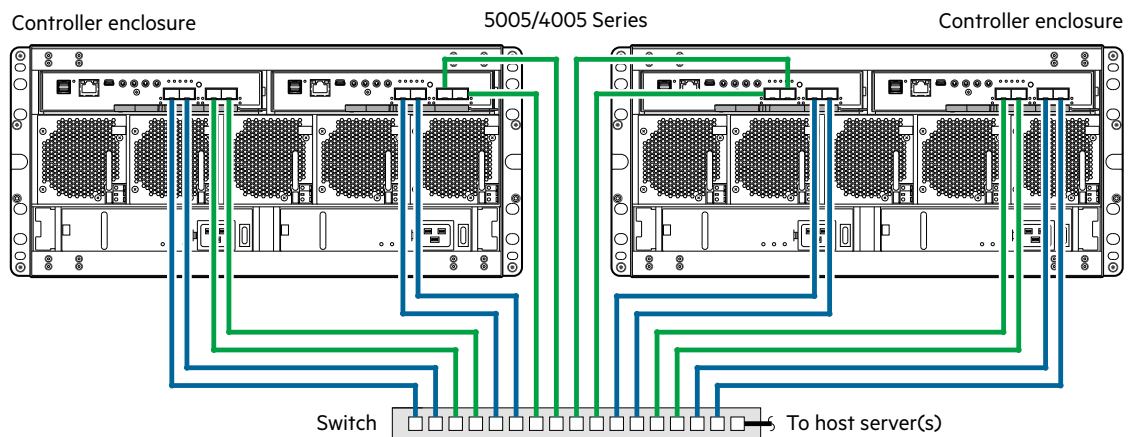
Figure 76 Connecting two 5U storage systems for replication: multiple servers/one switch/one location

Figure 77 (2U), Figure 78 (page 86) (2U) and Figure 79 (page 86) (5U) each show the rear panel of two enclosures with I/O and replication occurring on different physical networks. For optimal protection, 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.
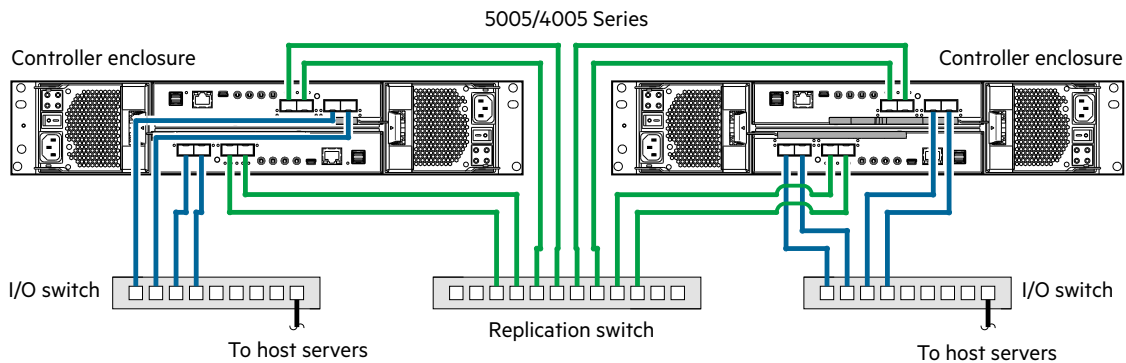
Figure 77 Connecting two 2U storage systems for replication: multiple servers/switches/one location
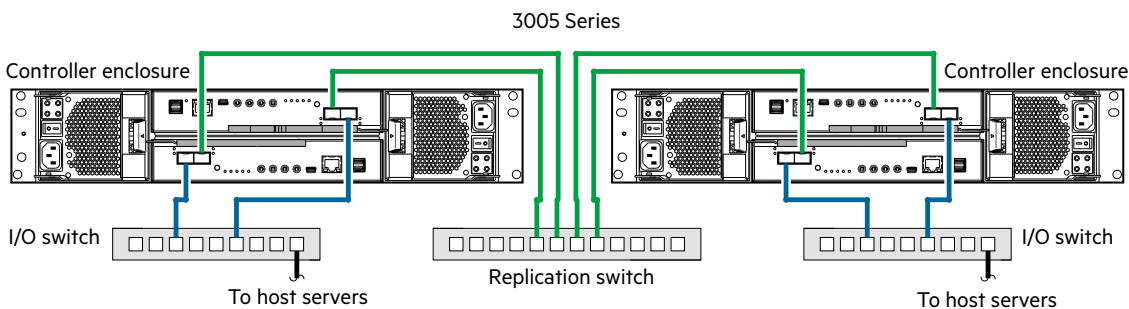
Figure 78 Connecting two 2U storage systems for replication: multiple servers/switches/one location
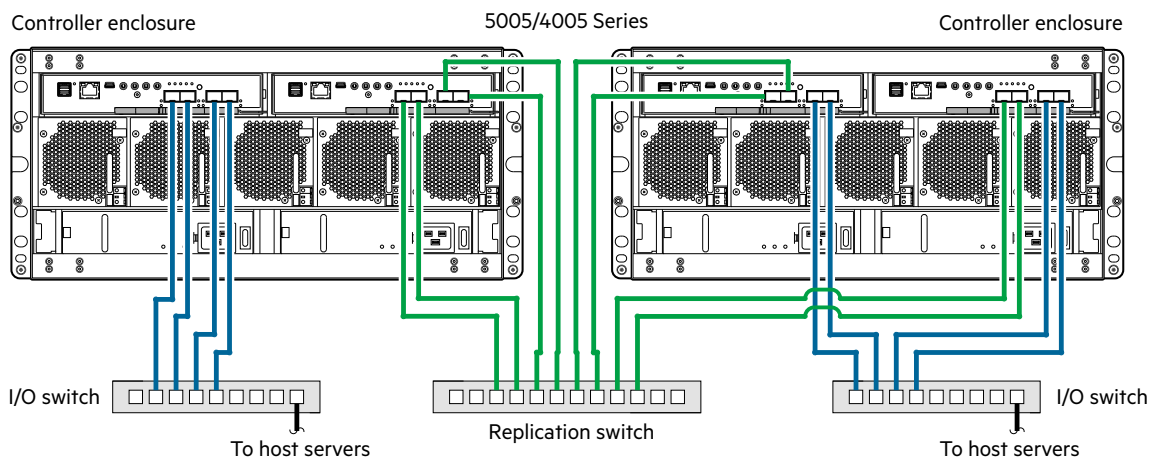
Figure 79 Connecting two 5U storage systems for replication: multiple servers/switches/one location

Multiple servers/different networks/multiple switches

Figure 80 (2U), Figure 81 (2U) and Figure 82 (5U) each show the rear panel of two enclosures with I/O and replication occurring on different networks. The I/O traffic occurs on the SAN, and the replication traffic occurs on the Ethernet WAN.
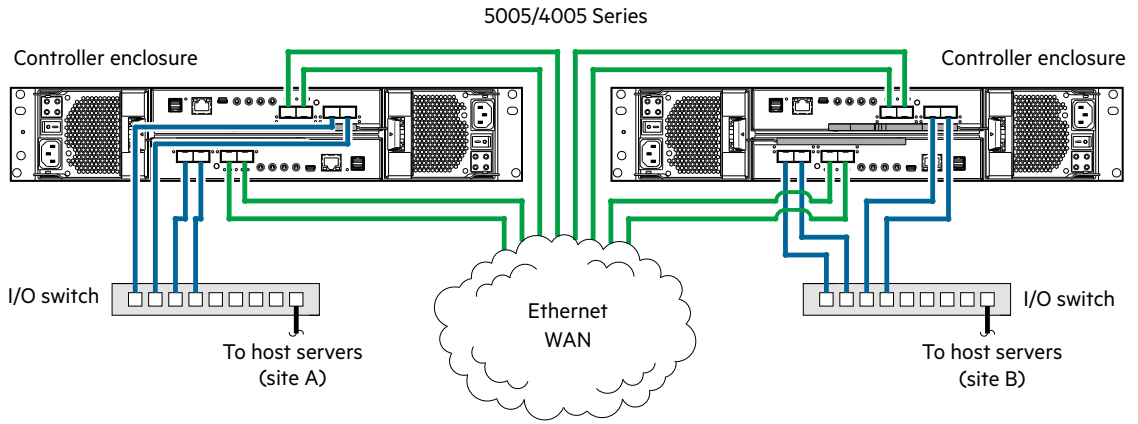
Figure 80 Connecting two 2U storage systems for replication: multiple servers/switches/two locations
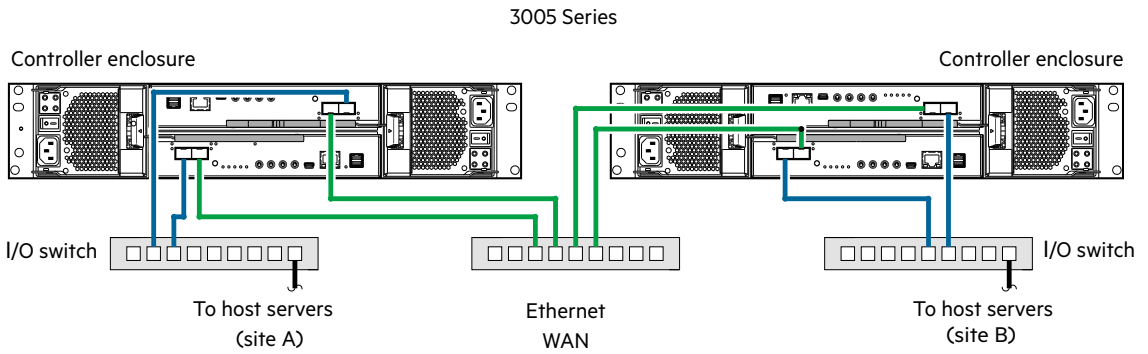
Figure 81 Connecting two 2U storage systems for replication: multiple servers/switches/two locations
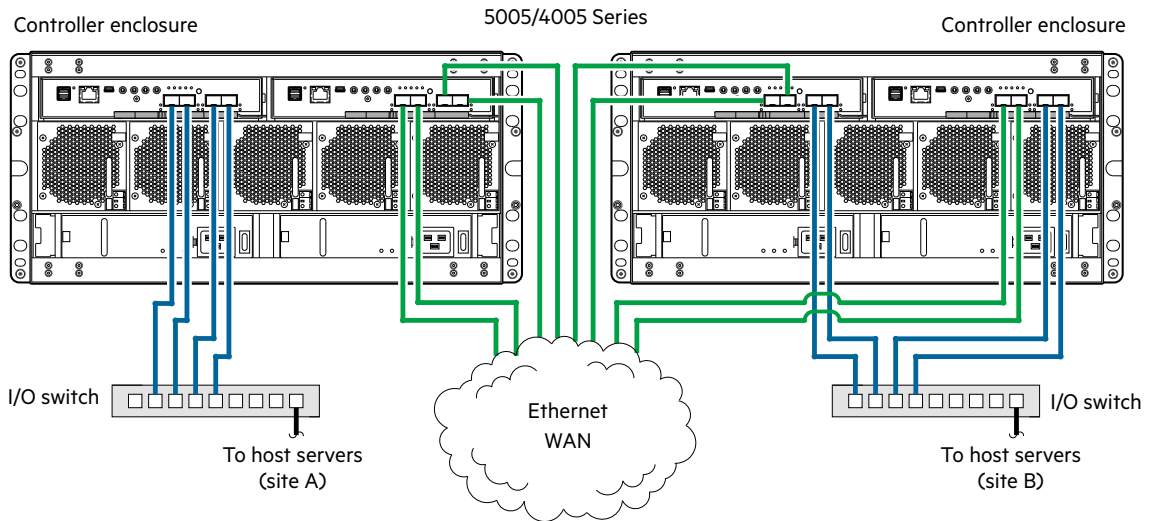
Figure 82 Connecting two 5U storage systems for replication: multiple servers/switches/two locations

The diagrams shown in Figure 80 and Figure 82 represent 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.

4 Operation

Before you begin

Before powering on the enclosure system, make sure that all modules are firmly seated in their correct slots. Verify that you have successfully completed the sequential "Installation Checklist" instructions in Table 9 (page 53). Once you have completed all of the necessary steps, you can access the management interfaces using your web-browser to complete the system setup.

Powering on/powering off

⚠ **CAUTION:** Do not operate the enclosure system until the ambient temperature is within the specified operating range described in "Environmental requirements" (page 170). If the drive modules have been recently installed, make sure they have had time to adjust to the environmental conditions before they are used with production data for I/O.

1. With 2U enclosures, power on the powered down system by connecting the power cables from the PCMs to the PDU, and moving the power switch on each PCM to the **on** position. See Figure 56 (page 64).

   The System Power LED on the 2U Ops panel should be lit green when the enclosure power is activated.

   With 5U enclosures, power on the powered down system by connecting the power cables from the PSUs to the PDU, and moving the power switch on each PSU to the **on** position. See Figure 57 (page 64).

   The Power on/Standby LED on the 5U Ops panel should be lit green when the enclosure power is activated.

2. Power the system down by moving the power switch on the PCM or PSU to the **off** position.

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

- Drive enclosures *first*

  This ensures that the disks in the 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.

- 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.

- Data host *last* (if powered down for maintenance purposes).

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

⊙ **IMPORTANT:** If main power is lost for any reason, upon restoration of power, the system will restart automatically.

NOTE: Enclosure Ops panels
- See Table 14 (page 92) for details pertaining to 2U Ops panel LEDs and related fault conditions.
- See Table 19 (page 96) for details pertaining to 5U84 Ops panel LEDs and related fault conditions.

NOTE: Guidelines for consideration when powering enclosures on and off

- Remove the AC cord before inserting or removing a PCM (2U) or PSU (5U).
- Move the PCM or PSU switch to the **off** position *before* connecting or disconnecting the AC power cord.
- Allow 15 seconds between powering off and powering on the PCM or PSU.
- Allow 15 seconds before powering on one PSU or PCM in the system, and powering off another PCM or PSU.
- Never power off a PCM or PSU while any amber LED is lit on the partner PCM or PSU.
- The enclosure must be left in a power on state for 30 seconds following resumption from standby before the enclosure can be placed into standby again (applies to 5U84 only).
- Although the enclosure supports standby, the expansion module shuts off completely during standby and cannot receive a user command to power back on. An AC power cycle is the only method to return the 5U84 to full power from standby.

Operator's (Ops) panel LEDs

Once the enclosure has successfully powered on, you can observe its Ops panel–located on the left ear of the front panel–for LED behavior reflecting enclosure status. Refer to the appropriate Ops panel figure/table from the options below for descriptions of enclosure status LEDs and related fault conditions.

- For 2U12 and 2U24 enclosures, see Figure 28 (page 35).
- For 5U84 enclosures, see Figure 29 (page 36).

Unit Identification Number

The Unit Identification Display (UID) is a dual seven segment display that is used to provide feedback to the user. Its primary purpose is to display an Enclosure UID number to assist in setting up, monitoring, and managing storage systems comprised of multiple enclosures.

The UID is stored in the enclosure VPD and is used by management interfaces. The UID can drive all seven of the segments, plus the dot/decimal point in each character of the display.

Software/SES

The enclosure UID number can be read and set through the management interfaces and SES.

Disk drive LEDs

Disk modules used in 2U chassis

Each disk drive module includes two LEDs (green and amber) as shown in "Disk drive carrier module LEDs" (page 93).

- In normal operation, the green LED will be on and will blink as the drive operates.
- In normal operation, the amber LED state will be:
    - Off if a disk is not present
    - Off as the disk successfully operates
    - On if there is a disk fault

DDICs used in 5U chassis

Each disk has a single amber drive fault LED as shown in "DDIC LED" (page 97). If the LED is illuminated, a disk failure condition has occurred. The DDIC should be replaced as soon as possible, using the procedure described in "Replacing a DDIC" (page 137), while adhering to guidelines provided in "Populating drawers" (page 139).

Accessing the SMC

Upon completing the hardware installation, you can access the controller module's web-based management interface—Storage Management Console (SMC)—to configure, monitor, and manage the storage system. Invoke your web browser and enter the IP address of the controller module's network port in the address field, then press **Enter**. To sign-in to the SMC, use the default user name `manage` and password `!manage`. If the default user name or password—or both—have been changed for security reasons, enter the secure login credentials instead of the defaults shown above. This brief sign-in discussion assumes proper web browser setup.

ⓘ **IMPORTANT:**   For detailed information about accessing and using the SMC, see the topic about getting started in the Storage Management Guide.

In addition to summarizing the processes to configure and provision a new system for the first time, the getting started topics provide instructions for signing in to the SMC, introduce key system concepts, address browser setup, and provide tips for using the main window and the help window.

☼ **TIP:**   After signing-in to the SMC, you can use online help as an alternative to consulting the Storage Management Guide.

Configuring and provisioning the storage system

RBOD/EBOD configuration

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

- Getting started
- Configuring the system
- Provisioning the system
- Using the optional replication feature

NOTE: See the topic about licensed features in the Storage Management Guide for instructions about creating a temporary license, or installing a permanent license.

ⓘ **IMPORTANT:**   If the system is used in a VMware environment, set the system's 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 Storage Management Guide, or see the `set advanced-settings` command topic in the CLI Reference Guide.

JBOD configuration

Use the host application RAID controls to configure and provision the storage system.

5 Troubleshooting and problem solving

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.

NOTE: For further troubleshooting help after setup, and when data is present, see seagate.com/support.

Overview

The enclosure system includes a Storage Enclosure Processor (SEP) and associated monitoring and control logic to enable it to diagnose problems with the enclosure's power, cooling, and drive systems. Management interfaces allow for provisioning, monitoring, and managing the storage system.

ⓘ **IMPORTANT:** See "Fault isolation methodology" (page 100) when conducting system diagnostics.

Initial start-up problems

Faulty power cords

Check that you have correctly cabled the system. Contact your supplier for replacements if:

- Power cables are missing or damaged.
- Plugs are incorrect.
- Power cables are too short.

Computer does not recognize the enclosure system

1. Verify that the interface cables from the enclosure to the host computer are fitted correctly.
2. Verify that the drive carrier modules have been correctly installed.
3. Check any visible SAS indicators (RBOD, EBOD, JBOD, and HBA).
4. Check HBA BIOS for SAS target visibility.
5. Verify that the operating system device driver has been installed correctly.

NOTE: If the enclosure fails initialization, see "If the enclosure does not initialize" (page 102).

LEDs

LED colors are used consistently throughout the enclosure and its components for indicating status:

- Green: good or positive indication
- Blinking green/amber: non-critical condition
- Amber: critical fault

2U enclosure LEDs

580W PCM LEDs

Under normal conditions, the PCM OK LEDs will be a constant green. See also Figure 30 (page 37). When a fault occurs, the colors of the LEDs will display as shown in the following table.

Table 13   PCM LED states

PCM OK (Green)	Fan Fail (Amber)	AC Fail (Amber)	DC Fail (Amber)	Status
Off	Off	Off	Off	No AC power on any PCM
Off	Off	On	On	No AC power on this PCM only
On	Off	Off	Off	AC present; PCM working correctly
On	Off	Off	On	PCM fan speed is outside acceptable limits
Off	On	Off	Off	PCM fan has failed
Off	On	On	On	PCM fault (over temperature, over voltage, over current)
Off	Blinking	Blinking	Blinking	PCM firmware download is in progress

Ops panel LEDs

The Ops panel displays the aggregated status of all the modules. See also Figure 28 (page 35). The Ops panel LEDs are defined in the following table.

Table 14   Ops panel LED states

System Power (Green/Amber)	Module Fault (Amber)	LED display	Associated LEDs/alarms	Status
On	Off	X		5V standby power present, overall power failed or switched off
On	On	On		Ops panel power on (5s) test state
On	Off	X		Power on, all functions good
On	On	X	PCM fault LEDs, fan fault LEDs	Any PCM fault, fan fault, over or under temperature
On	On	X	SBB module LEDs	Any SBB module fault
On	On	X	No module LEDs	Enclosure logical fault
On	Blink	X	Module status LED on SBB module	Unknown (invalid or mixed) SBB module type installed, I^{2}C bus failure (inter-SBB communications). EBOD/JBOD VPD configuration error
On	Blink	X	PCM fault LEDs, fan fault LEDs	Unknown (invalid or mixed) PCM type installed or I^{2}C bus failure (PCM communications)
		Blink		Enclosure identification or invalid ID selected
X = Disregard				

Actions for RBOD/EBOD configurations:

- If the Ops panel Module Fault LED is on, check the module LEDs on the enclosure rear panel to narrow the fault to a CRU, a connection, or both.
- Check the event log for specific information regarding the fault, and follow any Recommended Actions.
- If installing an IOM CRU:
  - Remove and reinstall the IOM per "Replacing an IOM" (page 125).
  - Check the event log for errors.
- If the CRU Fault LED is on, a fault condition is detected.
  - Restart this controller from the partner controller using the SMC or CLI.
  - If the restart does not resolve the fault, remove the IOM and reinsert it.
- If the above actions do not resolve the fault, contact Seagate for assistance.

NOTE: For JBOD-related actions, see "Basic steps" (page 155).

Disk drive carrier module LEDs

Disk drive status is monitored by a green LED and an amber LED mounted on the front of each drive carrier module, as shown in Figure 83. The drive module LEDs are identified in the figure, and the LED behavior is described in the table following the figure.

- In normal operation the green LED will be on, and will flicker as the drive operates.
- In normal operation the amber LED will be:
  - Off if there is no drive present.
  - Off as the drive operates.
  - On if there is a drive fault.
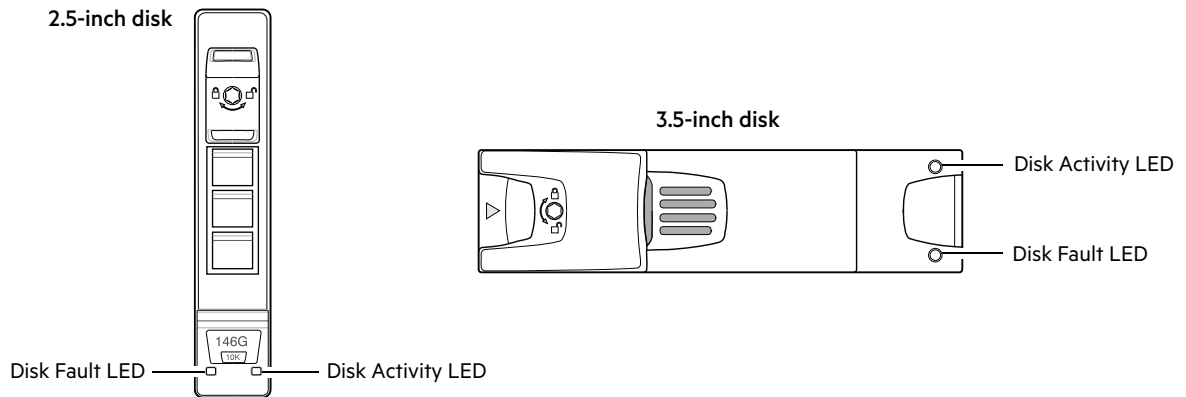
Activity LED (Green)	Fault LED (Amber)	Status/condition*
Off	Off	Off (disk module/enclosure)
Off	Off	Not present
Blink off with activity	Blinking: 1s on/1s off	Identify
1 down: Blink with activity 2 down: Off	On	Drive link (PHY lane) down
On	On	Fault (leftover/failed/locked-out)
Blink off with activity	Off	Available
Blink off with activity	Off	Storage system: Initializing

Activity LED (Green)	Fault LED (Amber)	Status/condition*
Blink off with activity	Off	Storage system: Fault-tolerant
Blink off with activity	Off	Storage system: Degraded (non-critical)
Blink off with activity	Blinking: 3s on/1s off	Storage system: Degraded (critical)
On	Off	Storage system: Quarantined
Blink off with activity	Blinking: 3s on/1s off	Storage system: Offline (dequarantined)
Blink off with activity	Off	Storage system: Reconstruction
Blink off with activity	Off	Processing I/O (whether from host or internal activity)

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

Figure 83 LEDs: Drive carrier LEDs (SFF and LFF modules)

IOM LEDs

IOM LEDs pertain to controller modules and expansion modules, respectively. Expansion enclosures are supported for optionally adding storage. See also "12Gb/s controller module LEDs" (page 40).

Controller module LEDs

Table 15   Controller module LED states

CRU OK (Green)	CRU Fault (Amber)	External host port activity (Green)	Status
On	Off		Controller module OK
Off	On		Controller module fault – see "Replacing an IOM" (page 125)
		Off	No external host port connection
		On	External host port connection – no activity
		Blinking	External host port connection – activity
Blinking			System is booting

Actions:

- If the CRU OK LED is blinking, wait for the system to boot.
- If the CRU OK LED is off, and the IOM is powered on, the module has failed.
    - Check that the IOM is fully inserted and latched in place, and that the enclosure is powered on.
    - Check the event log for specific information regarding the failure.
- If the CRU Fault LED is on, a fault condition is detected.
    - Restart this controller from the partner controller using the SMC or CLI.
    - If the restart does not resolve the fault, remove the IOM and reinsert it.
- If the above actions do not resolve the fault, contact your supplier for assistance. IOM replacement may be necessary.

Expansion module LEDs

Expansion IOM status is monitored by the LEDs located on the face plate. See also Figure 12 (page 26). LED behaviors for expansion enclosures are described in Table 16. For actions pertaining to Table 15 or Table 16, see **Actions** above. See also "12Gb/s expansion module LEDs" (page 47).

Table 16   Expansion module LED states

CRU OK (Green)	CRU Fault (Amber)	SAS port activity (Green)	Status
On	Off		Expansion module OK
Off	On		Expansion module fault – see "Replacing an IOM" (page 125)
		Off	No external port connection
		On	HD mini-SAS port connection – no activity
		Blinking	HD mini-SAS port connection – activity
Blinking			EBOD/JBOD VPD error

5U enclosure LEDs

NOTE: When the 5U84 enclosure is powered on, all LEDs are lit for a short period to ensure they are working. This behavior does not indicate a fault unless LEDs remain lit after several seconds.

PSU LEDs

See Figure 25 (page 33) for a visual description of the Power Supply Unit (PSU) module faceplate.

Table 17   PSU LED states

CRU Fail (Amber)	AC Missing (Amber)	Power (Green)	Status
On	Off	Off	No AC power to either PSU.
On	On	Off	PSU present, but not supplying power or PSU alert state. (usually due to critical temperature)
Off	Off	On	Mains AC present, switch on. This PSU is providing power.
Off	Off	Blinking	AC power present, PSU in standby (other PSU is providing power).
Blinking	Blinking	Off	PSU firmware download in progress.
Off	On	Off	AC power missing, PSU in standby (other PSU is providing power).
On	On	On	Firmware has lost communication with the PSU module.
On	—	Off	PSU has failed. Follow procedure in "Replacing a PSU" (page 140).

Fan cooling module LEDs

See Figure 26 (page 33) for a visual description of the Fan Cooling Module (FCM) faceplate.

Table 18   FCM LED descriptions

LED	Status/description
Module OK	Constant green indicates that the FCM is working correctly. Off indicates the fan module has failed. Follow procedure in "Replacing an FCM" (page 142).
Fan Fault	Amber indicates the fan module has failed. Follow procedure in "Replacing an FCM" (page 142).

Ops panel LEDs

The Ops panel displays the aggregated status of all the modules. See also Figure 29 (page 36).

Table 19   Ops panel LED descriptions

LED	Status/description
Unit ID Display	Usually shows the ID number for the enclosure, but can be used for other purposes. (e.g., Blinking to locate enclosure)
Power On/Standby	Amber if the system is in standby. Green if the system has full power.
Module Fault	Amber indicates a fault in a PSU, FCM, or IOM. Check the drawer LEDs for indication of a disk fault. See also "Drawer LEDs" (page 96).
Logical Status	Amber indicates a fault from something other than firmware (usually a disk, an HBA, or an internal or external RAID controller). Check the drawer LEDs for indication of a disk fault. See also "Drawer LEDs" (page 96).
Drawer 0 Fault	Amber is there is a disk, cable, or sideplane fault in drawer 0. Open the drawer and check DDICs for faults.
Drawer 1 Fault	Amber is there is a disk, cable, or sideplane fault in drawer 1. Open the drawer and check DDICs for faults.

Drawer LEDs

See Figure 27 (page 34) for a visual description of the Drawer LED inserts located on each drawer bezel.

Table 20   Drawer LED descriptions

LED	Status/description
Sideplane OK/ Power Good	Green if the sideplane card is working and there are no power problems.
Drawer Fault	Amber if a drawer component has failed. If the failed component is a disk, the LED on the failed DDIC will light amber. Follow procedure in "Replacing a DDIC" (page 137). If the disks are OK, contact your service provider to identify the cause of the failure, and resolve the problem.
Logical Fault	Amber (solid) indicates a disk fault. Amber (blinking) indicates that one or more storage systems are in an impacted state.
Cable Fault	Amber indicates the cabling between the drawer and the back of the enclosure has failed. Contact your service provider to resolve the problem.
Activity Bar Graph	Displays the amount of data I/O from zero segments lit (no I/O) to all six segments lit (maximum I/O).

DDIC LED

The DDIC supports LFF 3.5" and SFF 2.5" disks as shown in Figure 44 and Figure 45 (page 51). The illustration below shows the top panel of the DDIC as viewed when the disk is aligned for insertion into a drawer slot. See also Figure 20 and Figure 21 (page 31).

DDIC
(longitudinal view–top face of carrier)

Slide latch   Latch button (locked)   Disk Fault LED
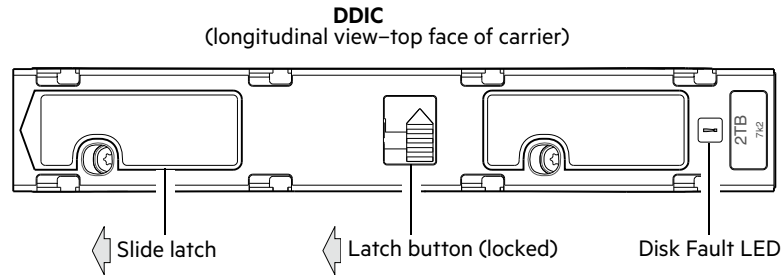
Fault LED (Amber)	Status/condition*
Off	Off (disk module/enclosure)
Off	Not present
Blinking: 1s on/1s off	Identify
Any links down: On	Drive link (PHY lane) down
On	Fault (leftover/failed/locked-out)
Off	Available
Off	Storage system: Initializing
Off	Storage system: Fault-tolerant
Off	Storage system: Degraded (non-critical)
Blinking: 3s on/1s off	Storage system: Degraded (critical)
Off	Storage system: Quarantined
Blinking: 3s on/1s off	Storage system: Offline (dequarantined)
Off	Storage system: Reconstruction
Off	Processing I/O (whether from host or internal activity)
*If multiple conditions occur simultaneously, the LED state will behave as indicated by the condition listed earliest in the table, as rows are read from top to bottom.	

Figure 84 LEDs: DDIC – 5U enclosure disk slot in drawer

Each DDIC has a single Drive Fault LED. A disk drive fault is indicated if the Drive Fault LED is lit amber. In the event of a disk failure, follow the procedure in "Replacing a DDIC" (page 137).

IOM LEDs

NOTE: IOM CRUs are common to the 2U and 5U enclosures.

- For information about controller module LEDs see "Controller module LEDs" (page 94)
- For information about expansion module LEDs see "Expansion module LEDs" (page 95)

Temperature sensors

Temperature sensors throughout the enclosure and its components monitor the thermal health of the storage system. Exceeding the limits of critical values will cause a notification to occur.

Troubleshooting 2U enclosures

The following sections describe common problems that can occur with your enclosure system, and some possible solutions. For all of the problems listed in Table 21, the Module Fault LED on the Ops panel will light amber to indicate a fault. See also "2U enclosure Ops panel" (page 35). All alarms will also report via SES.

Table 21   2U alarm conditions

Status	Severity	Alarm
PCM alert - loss of DC power from a single PCM	Fault - loss of redundancy	S1
PCM fan fail	Fault - loss of redundancy	S1
SBB module detected PCM fault	Fault	S1
PCM removed	Configuration error	None
Enclosure configuration error (VPD)	Fault – critical	S1
Low warning temperature alert	Warning	S1
High warning temperature alert	Warning	S1
Over temperature alarm	Fault – critical	S4
I^2C bus failure	Fault – loss of redundancy	S1
Ops panel communication error (I^2C)	Fault – critical	S1
RAID error	Fault – critical	S1
SBB interface module fault	Fault – critical	S1
SBB interface module removed	Warning	None
Drive power control fault	Warning – no loss of disk power	S1
Drive power control fault	Fault – critical–loss of disk power	S1
Drive removed	Warning	None
Insufficient power available	Warning	None

For details about replacing modules, see "Module removal and replacement" (page 116).

NOTE: See the Event Descriptions Reference Guide for more information about enclosure-related events and recommended actions.

PCM Faults

Symptom	Cause	Recommended action
Ops panel Module Fault LED is amber[1]	Any power fault	Verify AC mains connections to PCM are live
Fan Fail LED is illuminated on PCM[2]	Fan failure	Replace PCM

1-See Figure 28 (page 35) for visual reference of Ops panel LEDs.
2-See Figure 30 (page 37) for visual reference of PCM LEDs.

Thermal monitoring and control

The storage enclosure system uses extensive thermal monitoring and takes a number of actions to ensure component temperatures are kept low, and to also minimize acoustic noise. Air flow is from the front to back of the enclosure.

Symptom	Cause	Recommended action
If the ambient air is below 25ºC (77ºF), and the fans are observed to increase in speed, then some restriction on airflow may be causing additional internal temperature rise. **NOTE:** This is not a fault condition.	The first stage in the thermal control process is for the fans to automatically increase in speed when a thermal threshold is reached. This may be caused by higher ambient temperatures in the local environment, and may be perfectly normal. **NOTE:** This threshold changes according to the number of disks and power supplies fitted.	1. Check the installation for any airflow restrictions at either the front or back of the enclosure. A minimum gap of 25 mm (1") at the front and 50 mm (2") at the rear is recommended. 2. Check for restrictions due to dust build-up. Clean as appropriate. 3. Check for excessive re-circulation of heated air from rear to front. Use of the enclosure in a fully enclosed rack is not recommended. 4. Verify that all blank modules are in place. 5. Reduce the ambient temperature.

Thermal alarm

Symptom	Cause	Recommended action
1. Ops panel Module Fault LED is amber. 2. Fan Fail LED is illuminated on one or more PCMs.	Internal temperature exceeds a preset threshold for the enclosure.	1. Verify that the local ambient environment temperature is within the acceptable range. See also "Environmental requirements" (page 170). 2. Check the installation for any airflow restrictions at either the front or back of the enclosure. A minimum gap of 25 mm (1") at the front and 50 mm (2") at the rear is recommended. 3. Check for restrictions due to dust build-up. Clean as appropriate. 4. Check for excessive re-circulation of heated air from rear to front. Use of the enclosure in a fully enclosed rack is not recommended. 5. If possible, shut down the enclosure and investigate the problem before continuing.

Troubleshooting 5U enclosures

The table describes common problems that can occur with your enclosure system, together with possible solutions. For all of the problems listed in Table 22, the Module Fault LED on the Ops panel will light amber to indicate a fault. See also "5U enclosure Ops panel" (page 36). All alarms will also report via SES.

Table 22 5U alarm conditions

Status	Severity
PSU alert–loss of DC power from a single PSU	Fault–loss of redundancy
Cooling module fan failure	Fault–loss of redundancy
SBB I/O module detected PSU fault	Fault

Table 22   5U alarm conditions (continued)

Status	Severity
PSU removed	Configuration error
Enclosure configuration error (VPD)	Fault–critical
Low temperature warning	Warning
High temperature warning	Warning
Over-temperature alarm	Fault–critical
Under-temperature alarm	Fault–critical
I^{2}C bus failure	Fault–loss of redundancy
Ops panel communication error (I^{2}C)	Fault–critical
RAID error	Fault–critical
SBB I/O module fault	Fault–critical
SBB I/O module removed	Warning
Drive power control fault	Warning–no loss of drive power
Drive power control fault	Fault–critical: loss of drive power
Insufficient power available	Warning

For information about replacing modules, see "Module removal and replacement" (page 116).

Thermal considerations

NOTE: Thermal sensors in the 5U enclosure and its components monitor the thermal health of the storage system.

- Exceeding the limits of critical values will activate the over-temperature alarm.
- For information about 5U enclosure alarm notification, see Table 22 (page 99).

USB CLI port connection

5005/4005/3005 Series 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 "USB device connection" (page 175) for more information.

Fault isolation methodology

Seagate 5005/4005/3005 Series storage systems provide many ways to isolate faults. This section presents the basic methodology used to locate faults within a storage system, and to identify the pertinent CRUs affected.

NOTE: The subtopics in this section pertain to RBOD/EBOD configurations. See "Fault isolation methodology" (page 155) for troubleshooting server-attached JBOD faults.

As noted in "Accessing the SMC" (page 90), use the Storage Management Console 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 the Configuring event notification topic within the Storage Management 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

- Gather fault information, including using system LEDs as described in "Gather fault information" (page 102).
- Determine where in the system the fault is occurring as described in "Determine where the fault is occurring" (page 102).
- Review event logs as described in "Review the event logs" (page 102).
- If required, isolate the fault to a data path component or configuration as described in "Isolate the fault" (page 102).

Cabling systems to enable use of the licensed replication feature—to replicate volumes—is another important fault isolation consideration pertaining to initial system installation. See "Isolating replication faults" (page 106) 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 SMC 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, either the SMC or the CLI provide recommended-action text online. Options for performing basic steps are listed according to frequency of use:

- Use the SMC
- Use the CLI
- Monitor event notification
- View the enclosure LEDs

Use the SMC

The SMC uses health icons to show OK, Degraded, Fault, or Unknown status for the system and its components. The SMC 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 web application's GUI 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 SMC, you can run the `show system` CLI command 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 the SMC or the CLI. Using the SMC, 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 the SMC 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 described above 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 unit?*

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

Determine where the fault is occurring

When a fault occurs, the Module Fault LED—located on the Ops panel on an enclosure's left ear—illuminates. Check the LEDs on the back of the enclosure to narrow the fault to a CRU, connection, or both. The LEDs also help you identify the location of a CRU reporting a fault. See also "2U enclosure rear panel" (page 26) and "5U enclosure rear panel" (page 31).

Use the SMC to verify any faults found while viewing the LEDs. The SMC 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. This web-application provides you with a visual representation of the system and where the fault is occurring. The SMC also provides more detailed information about CRUs, 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.

NOTE: Some events also have a Resolved severity that indicates that a previously logged non-Informational condition has been resolved. See the Event Descriptions Reference Guide for information about specific events.

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, or data host.

If the enclosure does not initialize

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

- Perform a rescan
- Power cycle the system
- Make sure the power cord is properly connected, and check the power source to which it is connected
- 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 enclosures in a different configuration, 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 SMC 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 perform a rescan using the SMC:

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**.

Host I/O

When troubleshooting disk drive and connectivity faults, stop I/O to the affected disk groups from all hosts as a data protection precaution. As an additional data protection precaution, it is helpful to conduct regularly scheduled backups of your data. See also "Stopping I/O for RBOD/EBOD configurations" (page 127).

Dealing with hardware faults

Ensure that you have obtained a replacement module of the same type before removing any faulty module as described in "Module removal and replacement" (page 116).

---

① **IMPORTANT:**   If the enclosure system is powered up and you remove any module, replace it immediately. If the system is used with any modules missing for more than a few seconds, the enclosure(s) can overheat, causing power failure and potential data loss. Such action can invalidate the product warranty.

---

① **IMPORTANT:**   Observe applicable/conventional ESD precautions when handling modules and components, as described in "ESD precautions" (page 116). Avoid contact with midplane components, module connectors, leads, pins, and exposed circuitry.

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 host activity 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 procedures. This procedure requires scheduled downtime.

---

① **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 CNC ports

The procedure below applies to controller enclosures employing small form factor pluggable (SFP) transceiver connectors in 4/8/16Gb/s FC, 10GbE iSCSI, or 1Gb/s iSCSI host interface ports. In the following procedure, "SFP and host cable" is used to refer to any qualified SFP option supporting CNC 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. See also "Stopping I/O for RBOD/EBOD configurations" (page 127).
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.
    - Solid – Cache contains data yet to be written to the disk.
    - Blinking – Cache data is being written to CompactFlash.
    - Flashing at 1/10 second on and 9/10 second off – Cache is being refreshed by the supercapacitor.
    - 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?
    - 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.
    - 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?
    - Yes – You now know that the SFP, 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's port. Replace the controller module.
    - No – Proceed to the next step.
7. Swap the SFP with the known good one.

    Is the host link status/link activity LED on?
    - Yes – You have isolated the fault to the SFP. Replace the SFP.
    - 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?
    - Yes – You have isolated the fault to the cable. Replace the cable.
    - 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?
    - Yes – You have isolated the fault to the HBA. Replace the HBA.
    - 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 5005/4005 Series controller enclosures employing 12Gb SFF-8644 connectors in the HD mini-SAS host ports used for I/O. These models do not support replication.

1. Halt all I/O to the storage system .

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's 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 "Stopping I/O for RBOD/EBOD configurations" (page 127).
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's 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's 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's 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 replication faults

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. The replication feature 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 SMC or the CLI. See "Connecting two controller enclosures to replicate volumes" (page 82) for host connection information concerning the replication feature.

Replication setup and verification

After storage systems are cabled for replication, you can use the SMC to prepare for using the replication feature. Alternatively, you can use SSH or telnet to access the IP address of the controller module and access the replication feature using the CLI.

Basic information for enabling the 5005/4005/3005 Series controller enclosures for replication supplements the troubleshooting procedures that follow.

- Familiarize yourself with replication content provided in the Storage Management Guide.
- For virtual replication, in order to replicate an existing volume to a pool on the peer in the primary system or secondary system, follow these steps:
  - o Find the port address on the secondary system:

    Using the CLI, run the `show ports` command on the secondary system.
  - o 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 SMC **Replications** topic, select **Action > Query Peer Connection**.
  - o Create a peer connection.

    To create a peer connection, use the `create peer-connection` CLI command or in the SMC **Replications** topic, select **Action > Create Peer Connection**.
  - o Create a virtual replication set.

    To create a replication set, use the `create replication-set` CLI command or in the SMC **Replications** topic, select **Action > Create Replication Set**.
  - o Replicate.
  - o To initiate replication, use the `replicate` CLI command or in the SMC **Replications** topic, select **Action > Replicate**.
- For descriptions of replication-related events, see the Event Descriptions Reference Guide.

NOTE: These steps are a general outline of the replication setup. Refer to the following manuals for more information about replication setup:

- See the Storage Management Guide for procedures to setup and manage replications.
- See the CLI Reference Guide for replication commands and syntax.
- See the Event Descriptions Reference Guide for replication event reporting.

ⓘ **IMPORTANT:** Controller module firmware must be compatible on all systems used for replication. For license information, see the Storage Management Guide.

Diagnostic steps for replication setup

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

Can you successfully use the replication feature?

Answer	Possible reasons	Action
Yes	System functioning properly.	No action required.
No	The replication feature is not licensed on each controller enclosure used for replication.	Verify licensing of the optional feature per system: • In the **Home** topic in the SMC, 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 this feature. See the Storage Management Guide for license information. **NOTE:** Replication is not supported by SAS controller enclosures.
	Compatible firmware revision supporting the replication feature is not running on each system used for replication.	Perform the following actions on each system used for virtual replication: • 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 the other systems. • For more information on compatible firmware, see the topic about updating firmware in the Storage Management Guide.
	Invalid cabling connection. (If multiple enclosures are used, check the cabling for each system.)	Verify controller enclosure cabling: • 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 23   Diagnostics for replication setup: Using the replication feature

Can you create a replication set?

Answer	Possible reasons	Action
Yes	System functioning properly.	No action required.
No	On controller enclosures equipped 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 Storage Management Guide.

	Unable to create the secondary volume (the destination volume on the pool to which you will replicate data from the primary volume)?[1]	• 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:
		   o A conflicting volume does not already exist
		   o Available free space in the pool

	Communication link is down.	Review event logs for indicators of a specific fault in a host or replication data path component.
After ensuring valid licensing, valid cabling, and network availability, create the replication set using the **Replications** topic: select **Action > Create Replication Set**.		

Table 24 Diagnostics for replication setup: Creating a replication set

Can you replicate a volume?

Answer	Possible reasons	Action
Yes	System functioning properly.	No action required.
No	The replication feature is not licensed on each controller enclosure used for replication.	See actions described in "Can you successfully use the replication feature?" (page 108).
No	Nonexistent replication set.	• 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.
	Network error occurred during in-progress replication.	• 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 Storage Management Guide for additional information).
	Communication link is down.	Review event logs for indicators of a specific fault in a host or replication data path component.

Table 25 Diagnostics for replication setup: Replicating a volume

Has a replication run successfully?

Answer	Possible reasons	Action
Yes	System functioning properly.	No action required.

No	Last Successful Run shows N/A.	• In the **Volumes** topic, click on the volume that is a member of the replication set.
		
		    o Select the **Replication Sets** table.
		
		    o Check the Last Successful Run information.
		
		• If the replication has not run successfully, use the SMC to replicate as described in the topic about working in replications in the Storage Management Guide.

	Communication link is down.	Review event logs for indicators of a specific fault in a host or replication data path component.

Table 26   Diagnostics for replication setup: Checking for a successful replication

Continuous operation during replacement

Your hardware or software enclosure management application determines the capability for replacing a failed disk without the loss of access to any file system on the enclosure. Enclosure access and use during this period is uninterrupted. If an enclosure is equipped with redundant PCMs or PSU, sufficient power is provided to the system while the faulty module is replaced.

NOTE: 5005/4005/3005 Series enclosures support hot-plug replacement of redundant controller modules, power supplies, and expansion modules. Hot-add replacement of expansion enclosures is also supported.

Firmware updates

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. Periodically, you should ensure that the firmware versions used in enclosure modules are compatible. Also see "Updating firmware for RBOD/EBOD configurations" (page 75) for RBOD/EBOD configurations, and "Obtaining the JBOD firmware download" (page 156) for JBOD configurations.

Customer-replaceable units

NOTE: The following tables share the table endnotes located beneath Table 30 (page 111). See Figure 85 (page 115) for pictorial views of IOM CRUs used in the different enclosure form factors supported by 5005/4005/3005 Series.

Table 27   5005 Series enclosure models

5005 Series: 4-port controller enclosure matrix–2U					
2.5" (SFF) 24-drive controller enclosures			3.5" (LFF) 12-drive controller enclosures		
Model	Description	Form	Model	Description	Form
5005	Fibre Channel (8/16Gb/s) SFP[1,4]	2U24	5005	Fibre Channel (8/16Gb/s) SFP[1,4]	2U12
5005	Internet SCSI (10GbE) SFP[2,4]	2U24	5005	Internet SCSI (10GbE) SFP[2,4]	2U12
5005	Internet SCSI (1Gb/s) SFP[3,4]	2U24	5005	Internet SCSI (1Gb/s) SFP[3,4]	2U12
5005	HD mini-SAS (12Gb/s)[5]	2U24	5005	HD mini-SAS (12Gb/s)[5]	2U12

Table 28   4005 Series enclosure models

4005 Series: 4-port controller enclosure matrix–2U					
2.5" (SFF) 24-drive controller enclosures			3.5" (LFF) 12-drive controller enclosures		
Model	Description	Form	Model	Description	Form
4005	Fibre Channel (8/16Gb/s) SFP[1,4]	2U24	4005	Fibre Channel (8/16Gb/s) SFP[1,4]	2U12
4005	Internet SCSI (10GbE) SFP[2,4]	2U24	4005	Internet SCSI (10GbE) SFP[2,4]	2U12
4005	Internet SCSI (1Gb/s) SFP[3,4]	2U24	4005	Internet SCSI (1Gb/s) SFP[3,4]	2U12
4005	HD mini-SAS (12Gb/s)[5]	2U24	4005	HD mini-SAS (12Gb/s)[5]	2U12

Table 29   3005 Series enclosure models

3005 Series: 2-port controller enclosure matrix–2U					
2.5" (SFF) 24-drive controller enclosures			3.5" (LFF) 12-drive controller enclosures		
Model	Description	Form	Model	Description	Form
3005	Fibre Channel (8/16Gb/s) SFP[1]	2U24	3005	Fibre Channel (8/16Gb/s) SFP[1]	2U12
3005	Internet SCSI (10GbE) SFP[2]	2U24	3005	Internet SCSI (10GbE) SFP[2]	2U12
3005	Internet SCSI (1Gb/s) SFP[3]	2U24	3005	Internet SCSI (1Gb/s) SFP[3]	2U12

Table 30   High-density 5005/4005 Series enclosure models

High-density 4-port controller enclosure matrix–5U					
5005 Series (2.5"/3.5") 84-drive controller enclosures			4005 Series (2.5"/3.5") 84-drive controller enclosures		
Model	Description	Form	Model	Description	Form
5005	Fibre Channel (8/16Gb/s) SFP[1,4]	5U84	4005	Fibre Channel (8/16Gb/s) SFP[1,4]	5U84
5005	Internet SCSI (10GbE) SFP[2,4]	5U84	4005	Internet SCSI (10GbE) SFP[2,4]	5U84
5005	Internet SCSI (1Gb/s) SFP[3,4]	5U84	4005	Internet SCSI (1Gb/s) SFP[3,4]	5U84
5005	HD mini-SAS (12Gb/s)[5]	5U84	4005	HD mini-SAS (12Gb/s)[5]	5U84

1-This model uses a qualified FC SFP option within the CNC ports (used for host connection). When in FC mode, the SFPs must be a qualified 8Gb or 16Gb fiber-optic option. A 16Gb/s SFP can run at 16Gb/s, 8Gb/s, 4Gb/s, or auto-negotiate its link speed. An 8Gb/s SFP can run at 8Gb/s, 4Gb/s, or auto-negotiate its link speed.

2-This model uses a qualified 10GbE iSCSI option within the controller module CNC ports (used for host connection).

3-This model uses a qualified 1Gb iSCSI SFP option within the controller module CNC ports (used for host connection).

4-CNC ports support same-type or mixed-type SFPs in combination as described in "CNC ports used for host connection" (page 14).

5-This model uses SFF-8644 connectors and qualified cable options for host connection.

CRUs addressing 2U 12-drive chassis

Table 31   5005/4005/3005 Series product components for 2U 12-drive chassis

Item	Enclosure component description
1	Disk drive (LFF)
	a) 3.5" disk drive module (disks of differing type/speed and storage capacity: SAS, SSD)
	b) Dummy drive carrier module (blank to maintain optimum air flow within enclosure)
2	Ear components (not customer replaceable)
	a) Left ear assembly
	b) Right ear assembly
3	Chassis (sheet metal enclosure that is configurable as an RBOD, EBOD, or JBOD)
4	Midplane PCB (included with chassis; not available separately)
5	Power cooling module (PCM) available as AC unit (chassis uses two PCMs of same model type)
6	Expansion module for storage expansion I/O (used when chassis is configured as an EBOD or JBOD)
	a) 3-port 12Gb/s SBB expansion module
7	Controller module host interfaces for storage system (used when chassis is configured as an RBOD)
	a) 5005 4-port FC/iSCSI supports (8/16Gb/s FC; 10GbE iSCSI; or 1Gb/s iSCSI) qualified SFP options*
	b) 4005 4-port FC/iSCSI supports (8/16Gb/s FC; 10GbE iSCSI; or 1Gb/s iSCSI) qualified SFP options*
	c) 3005 2-port FC/iSCSI supports (8/16Gb/s FC; 10GbE iSCSI; or 1Gb/s iSCSI) qualified SFP options*
	d) 5005 4-port mini-SAS (12Gb/s) supports various qualified SFF-8644 to SFF-8644 or SFF-8088 options
	e) 4005 4-port mini-SAS (12Gb/s) supports various qualified SFF-8644 to SFF-8644 or SFF-8088 options
8	Small form-pluggable SFP connectors
	a) SFP transceiver: 8/16Gb/s FC; 10GbE iSCSI; 1Gb/s iSCSI
9	Rail kit (variable attachment options)
	a) 2U rack mount kit, shelf, short, all HW
10	Cable kit [Cable package: standard HD mini-SAS (SFF-8644) to HD mini-SAS (SFF-8644)] Cable kit [Cable package: standard HD mini-SAS (SFF-8644) to mini-SAS (SFF-8088)] Cable kit [Cable package: USB Type B; CLI (USB)]
11	AC power cord compatible with AC PCM
*5005 and 4005 models support FC and iSCSI SFPs used in combination.	

NOTE: Figure 85 (page 115) shows dimetric pictorial representations of controller modules and expansion modules used in 2U 12-drive 5005/4005/3005 Series enclosures. Modules are shown oriented for insertion into IOM slot A located on the 2U enclosure rear panel. This orientation also applies to both IOM slots in 5U84 enclosures.

CRUs addressing 2U 24-drive chassis

Table 32   5005/4005/3005 Series product components for 2U 24-drive chassis

Item	Enclosure component description
1	Disk drive (SFF)
	a) 2.5" disk drive module (disks of differing type/speed and storage capacity: SAS, SSD)
	b) Dummy drive carrier module (blank to maintain optimum air flow within enclosure)
2	Ear components (not customer replaceable)
	a) Left ear assembly
	b) Right ear assembly
3	Chassis (sheet metal enclosure that is configurable as an RBOD, EBOD, or JBOD)
4	Midplane PCB (included with chassis; not available separately)
5	Power cooling module (PCM) available as AC unit (chassis uses two PCMs of same model type)
6	Expansion module for storage expansion I/O (used when chassis is configured as an EBOD/JBOD)
	a) 3-port 12Gb/s SBB expansion module
7	Controller module host interfaces for storage system (used when chassis is configured as an RBOD)
	a) 5005 4-port FC/iSCSI supports (8/16Gb/s FC; 10GbE iSCSI; or 1Gb/s iSCSI) qualified SFP options*
	b) 4005 4-port FC/iSCSI supports (8/16Gb/s FC; 10GbE iSCSI; or 1Gb/s iSCSI) qualified SFP options*
	c) 3005 2-port FC/iSCSI (12Gb/s) supports (8/16Gb/s FC; 10GbE iSCSI; or 1Gb/s iSCSI) qualified SFP options*
	d) 5005 4-port mini-SAS (12Gb/s) supports various qualified SFF-8644 to SFF-8644 or SFF-8088 options
	e) 4005 4-port mini-SAS (12Gb/s) supports various qualified SFF-8644 to SFF-8644 or SFF-8088 options
8	Small form-pluggable SFP connectors
	a) SFP transceiver: 8/16Gb/s FC; 10GbE iSCSI; 1Gb/s iSCSI
9	Rail kit (variable attachment options)
	a) 2U rack mount kit, shelf, short, all HW
10	Cable kit [Cable package: standard HD mini-SAS (SFF-8644) to HD mini-SAS (SFF-8644)] Cable kit [Cable package: standard HD mini-SAS (SFF-8644) to mini-SAS (SFF-8088)] Cable kit [Cable package: USB Type B; CLI (USB)]
11	AC power cord compatible with AC PCM
*5005 and 4005 models support FC and iSCSI SFPs used in combination.	

NOTE: Figure 85 (page 115) shows dimetric pictorial representations of controller modules and expansion modules used in 2U 24-drive 5005/4005/3005 Series enclosures. Modules are shown oriented for insertion into IOM slot A located on the 2U enclosure rear panel. This orientation also applies to both IOM slots in 5U84 enclosures.

CRUs addressing high-density 5U 84-drive chassis

Table 33   5005/4005 Series product components for 5U 84-drive chassis

Item	Enclosure component description
1	Disk drive in carrier (DDIC) LFF (DDICs must be installed into drawers after delivery)
	a) 2.5" disk with 3.5" adapter (disks of differing type/speed and storage capacity: SAS, SSD)
	b) 3.5" disk (disks of differing type/speed and storage capacity: SAS, SSD)
2	Ear components (not customer replaceable)
	a) Left ear assembly
	b) Right ear assembly
3	Drawers for holding DDICs (included with chassis; not available separately)
	Top drawer accessible from enclosure front panel (empty–42 DDIC slots)
	Bottom drawer accessible from enclosure front panel (empty–42 DDIC slots)
4	Chassis (sheet metal enclosure that is configurable as an RBOD, EBOD, or JBOD)
5	PCBs –midplane sideplane, baseplanes (included with chassis; not available separately)
6	Power supply unit (PSU) available as AC module (chassis uses two PSUs of same model type)
7	Fan cooling module (FCM) (chassis uses five FCMs of same model type)
8	Expansion module for storage expansion I/O (used when chassis is configured as an EBOD/JBOD)
	a) 3-port 12Gb/s SBB expansion module
9	Controller module host interfaces for storage system (used when chassis is configured as an RBOD)
	a) 5005 4-port FC/iSCSI supports (8/16Gb/s FC; 10GbE iSCSI; or 1Gb/s iSCSI) qualified SFP options*
	b) 4005 4-port FC/iSCSI supports (8/16Gb/s FC; 10GbE iSCSI; or 1Gb/s iSCSI) qualified SFP options*
	c) 5005 4-port mini-SAS (12Gb/s) supports various qualified SFF-8644 to SFF-8644 or SFF-8088 options
	d) 4005 4-port mini-SAS (12Gb/s) supports various qualified SFF-8644 to SFF-8644 or SFF-8088 options
10	Small form-pluggable SFP connectors
	a) SFP transceiver: 8/16Gb/s FC; 10GbE iSCSI; 1Gb/s iSCSI
11	Rail kit (variable attachment options)
	a) 5U rack mount kit, shelf, long, all HW
12	Cable kit [Cable package: standard HD mini-SAS (SFF-8644) to HD mini-SAS (SFF-8644)] Cable kit [Cable package: standard HD mini-SAS (SFF-8644) to mini-SAS (SFF-8088)] Cable kit [Cable package: USB Type B; CLI (USB)]
13	AC power cord compatible with AC PSU (each PSU uses a power cord)
*5005 and 4005 models support FC and iSCSI SFPs used in combination.	

NOTE: Figure 85 (page 115) shows dimetric pictorial representations of controller modules and expansion modules used in 5U 84-drive 5005/4005 Series enclosures. Modules are shown oriented for insertion into IOM slots located on the 2U enclosure rear panel. This orientation also applies to both IOM slots in 5U84 enclosures.

4-port FC/iSCSI controller module
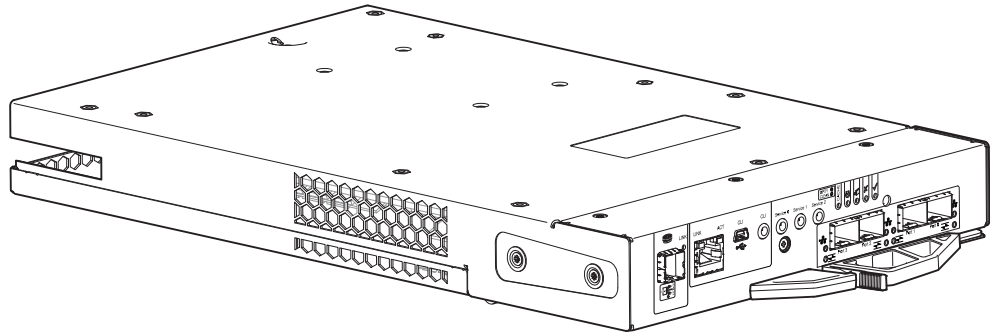
2-port SAS controller module (3005 Series only)
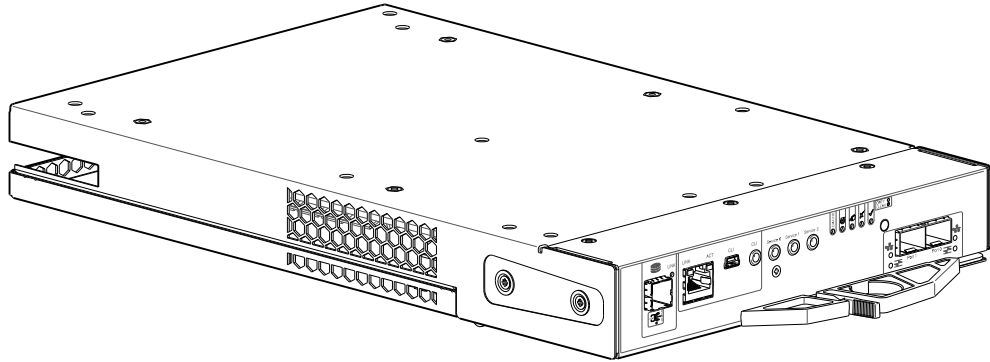
4-port SAS controller module
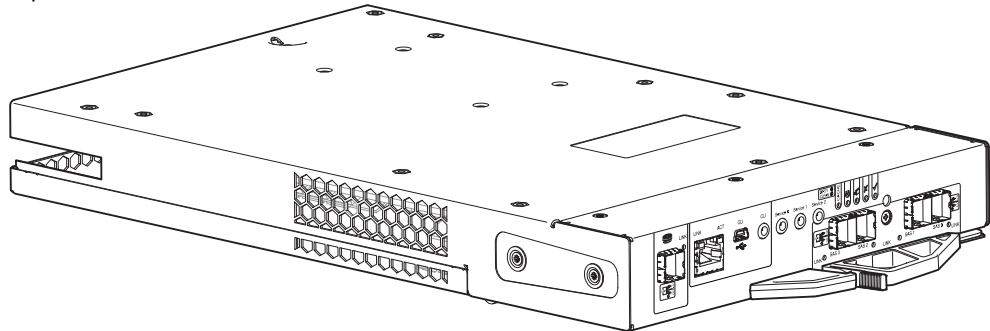
3-port expansion module
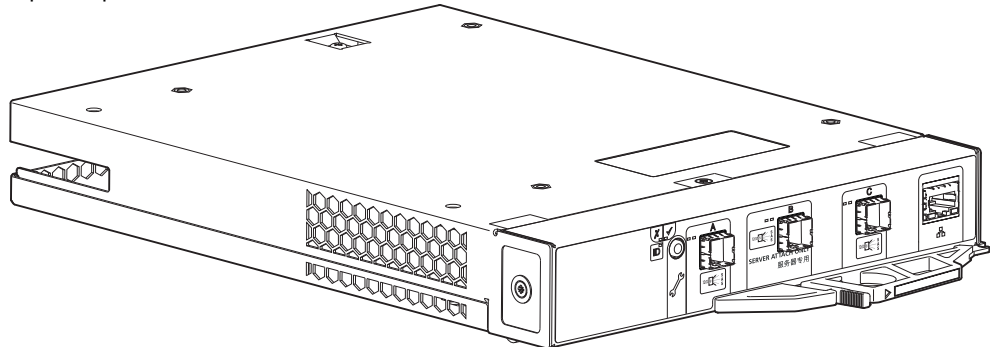
Figure 85 SBB IOMs used in 5005/4005/3005 Series storage enclosures

6   Module removal and replacement

Overview

This chapter provides procedures for replacing CRUs (customer-replaceable units), including precautions, removal instructions, installation instructions, and verification of successful installation. Each procedure addresses a specific task. Certain procedures refer to related documentation. See "Related documentation" (page 15) for a list of these documents and where to find them online.

NOTE: CRU replacement procedures are grouped by chassis form factor.

- For 2U12 and 2U24 CRU replacement, see "CRU replacement for 2U chassis" (page 117)
- For 5U84 CRU replacement, see "CRU replacement for 5U chassis" (page 135)

ESD precautions

Before you begin *any* of the procedures, consider the following precautions and preventive measures.

Preventing electrostatic discharge

To prevent electrostatic discharge (ESD) from damaging the system, be aware of the precautions to consider 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.

⚠ **CAUTION:**   Parts can be damaged by electrostatic discharge. Follow these precautions:

- 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.
- Remove clutter (plastic, vinyl, foam) from the static-protected workstation.

Grounding methods to prevent electrostatic discharge

Several methods are used for grounding. Adhere to the following precautions when handling or installing electrostatic-sensitive parts.

⚠ **CAUTION:**   Parts can be damaged by electrostatic discharge. Use proper anti-static protection:

- Keep the replacement CRU in the ESD bag until needed; and when removing a CRU from the enclosure, immediately place it in the ESD bag and anti-static packaging.
- Wear an ESD wrist strap connected by a ground cord to a grounded workstation or unpainted surface of the 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.
- If an ESD wrist strap is unavailable, touch an unpainted surface of the chassis before handling the component.
- 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 technician install the part. For more information about static electricity or assistance with product installation, contact customer support. For additional information, see seagate.com/support.

CRU replacement for 2U chassis

NOTE: Unless noted otherwise within a passage pertaining to a particular CRU, the replacement procedure should be completed within two minutes of the removal of a defective module.

Replacing a power cooling module

This section provides procedures for replacing a failed power cooling module (PCM). Illustrations in PCM replacement procedures show rear panel views of the enclosure, with the PCM properly oriented for insertion into the rear panel of the enclosure.

A single PCM is sufficient to maintain operation of the enclosure. You need not halt operations and completely power-off the enclosure when replacing only one PCM; however, a complete orderly shutdown is required if replacing both units simultaneously.

△ **CAUTION:**   Do not remove the cover from the PCM due to danger from electric shock inside. Return the PCM to your supplier for repair.

See CAUTION bullets regarding *electrostatic discharge* and *anti-static protection* on page 116.

☼ **TIP:**   The illustrations show PCM module replacement within the right slot as you view the enclosure rear panel. To replace a PCM in the left slot, you would first rotate the module 180º about its longitudinal axis, so that it properly aligns with its connectors on the back of the midplane. See also Figure 12 (page 26).

Removing a PCM

△ **CAUTION:**   Removing a power supply unit significantly disrupts the enclosure's airflow. Do not remove the PCM until you have received the replacement module. It is important that all slots are filled when the enclosure is in operation.

Before removing the PCM, disconnect the power from the PCM by either the mains switch (where present) or by physically removing the power source in order to ensure your system has warning of imminent power shutdown. A faulty PCM must be replaced by a fully operational PCM within 24 hours. Ensure that you correctly identify the faulty PCM before beginning the step procedure.

1. Stop all I/O from hosts to the enclosure. See also "Stopping I/O for RBOD/EBOD configurations" (page 127).

    ☼ **TIP:**   This step is not required for hot-swapping. However, it is required when replacing both PCMs at once.

2. Use management software to shut down any other system components necessary.

    ☼ **TIP:**   This step is not required for hot-swapping. However, it is required when replacing both PCMs at once.

3. Switch off the faulty PCM, and disconnect the power supply cable.
4. If replacing a single PCM via hot-swap, proceed to step 6.

5. If replacing both PCMs, verify that the enclosure was shut down using management interfaces, and that the enclosure is powered off.

6. Verify that the power cord is disconnected.

7. Grasp the latch and the side of the PCM handle between thumb and forefinger, squeeze together and open the handle to cam the PCM out of the enclosure as shown in Figure 86.
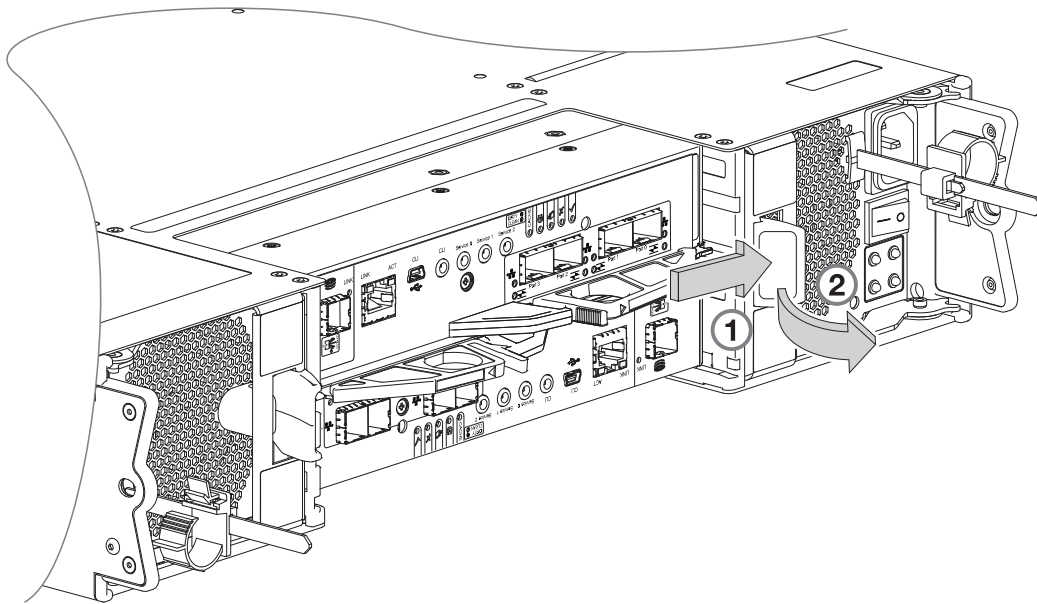
Figure 86 Removing a PCM (1 of 2)

8. Grip the handle and withdraw the PCM, taking care to support the base of the module with both hands as you remove it from the enclosure as shown in Figure 87.
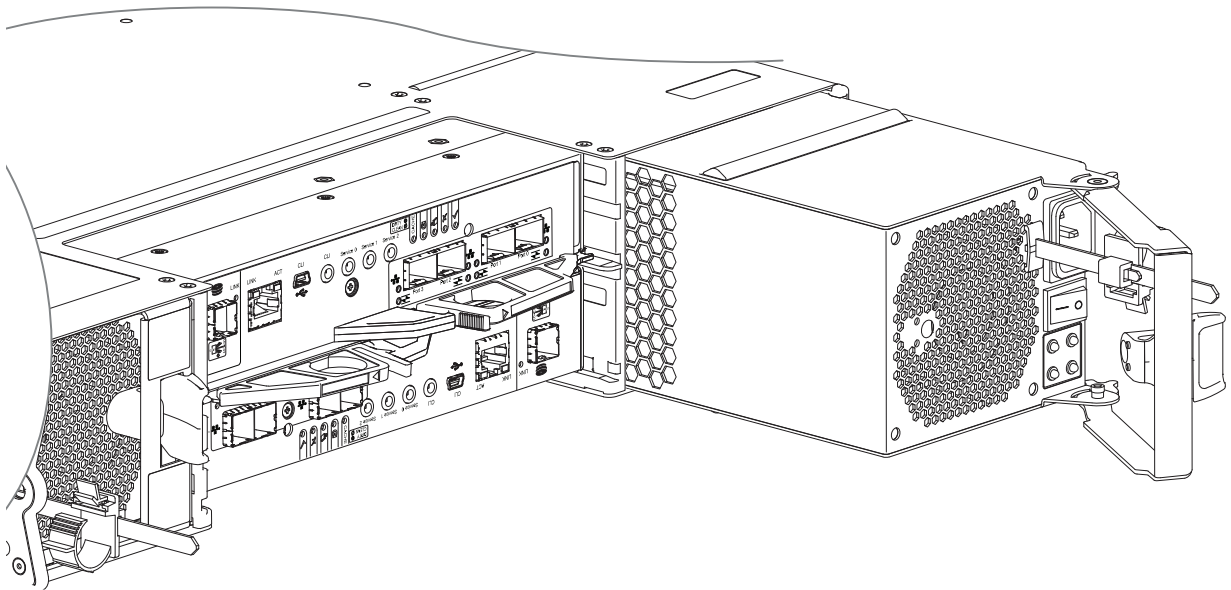
Figure 87 Removing a PCM (2 of 2)

> **NOTE:** The remove PCM illustrations show a chassis configured as a 4-port FC/iSCSI controller enclosure. The procedure applies to all 2U controller enclosures and expansion enclosures.

9. If replacing two PCMs, repeat steps 5 through 8, being mindful of the illustrations TIP on page 117 (this is the first TIP appearing on the page).

Installing a PCM

Refer to Figure 86 (page 118) and Figure 87 (page 118) when performing this procedure, but ignore the directional arrows—since you will insert the module into the slot—rather than extracting it.

ⓘ **IMPORTANT:** Handle the PCM carefully, and avoid damaging the connector pins. Do not install the PCM if any pins appear to be bent.

---

1. Check for damage, especially to all module connectors.
2. With the PCM handle in the open position, slide the module into the enclosure, taking care to support the base and weight of the module with both hands.
3. Cam the module home by manually closing the PCM handle. You should hear a click as the latch handle engages and secures the PCM to its connector on the back of the midplane.
4. Connect the power cable to the power source and the PCM.
5. Secure the strain relief bales.
6. Using the management interfaces (the SMC or CLI), verify whether the health of the new PCM is OK. Verify that the green PCM OK LED is on/blinking per Table 13 (page 92). Verify that cooling fans are spinning with no fail states. Verify that Ops panel states show no amber module faults.
7. If replacing two PCMs, repeat steps 1 through 5, being mindful of the illustrations TIP on page 117 this is the first TIP appearing on the page).

Replacing a drive carrier module

A disk drive module consists of a disk in a carrier or sled. Disk drive modules are hot-swappable, which means they can be replaced without halting I/O to the disk groups, or powering off the enclosure. The new disk must be of the same type, and possess capacity equal to or greater than the one being replaced. Otherwise, the storage system cannot use the new disk to reconstruct the disk group.

△ **CAUTION:** Removing a disk drive module impacts the airflow and cooling ability of the enclosure. If the internal temperature exceeds acceptable limits, the enclosure may overheat and automatically shut down or restart. To avoid potential overheating, wait 20 seconds to allow the internal disks to stop spinning, then insert the new disk drive module.

ⓘ **IMPORTANT:** Familiarize yourself with FDE considerations relative to disk module installation and replacement. See also "FDE considerations" (page 60).

NOTE: When moving FDE-capable disk drive modules for a disk group, stop I/O to the disk group before removing the disk drive modules. Import the keys for the disks so that the disk content becomes available. See the Storage Management Guide or CLI Reference Guide for more information.

See CAUTION bullets regarding *electrostatic discharge* and *anti-static protection* on page 116.

☼ **TIP:** The illustrations show disk module replacement within the drive slots as you view the enclosure front panel.

Removing a LFF drive carrier module

1. Press the latch in the carrier handle towards the handle hinge to release the carrier handle as shown below.
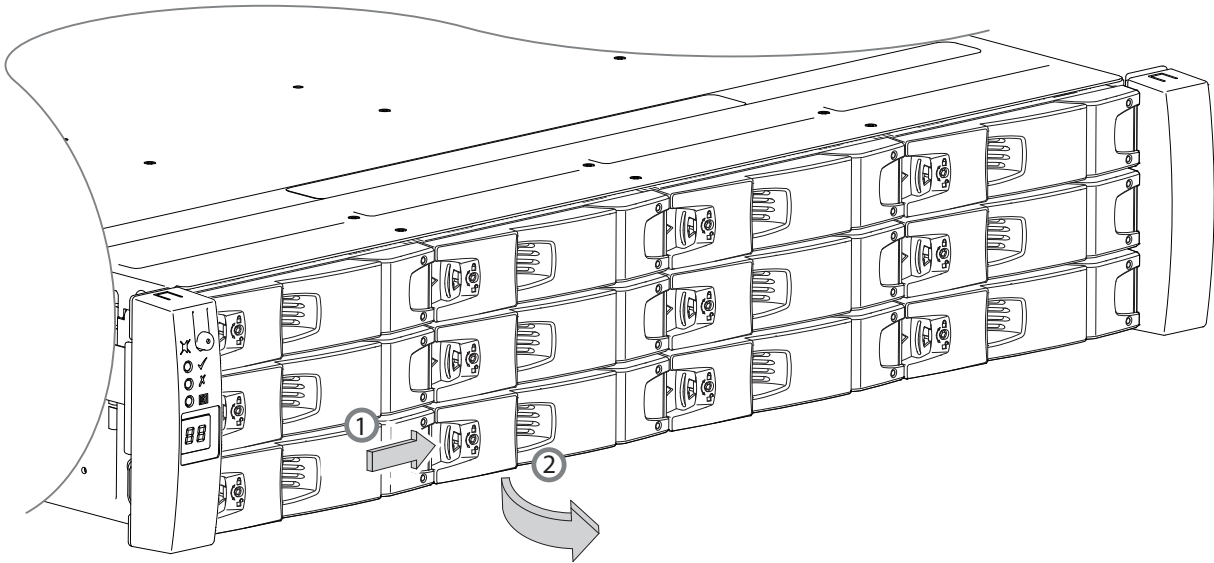
Figure 88 Removing a LFF disk drive module (1 of 2)

2. Gently move the drive carrier module approximately 25 mm (1-inch), then wait 30 seconds.
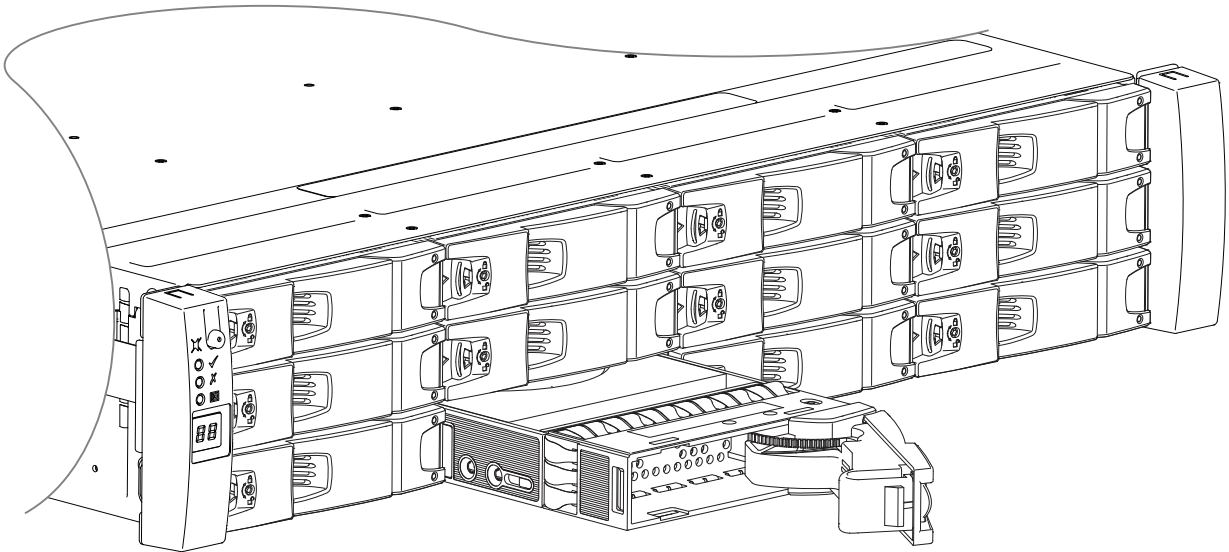
Figure 89 Removing a LFF disk drive module (2 of 2)

3. Remove the module fully from the drive slot.

△ **CAUTION:**   To ensure optimal cooling throughout the enclosure, dummy drive carrier modules must be fitted to all unused drive slots. See also Figure 43 (page 50).

Installing a LFF drive carrier module

⚠ **CAUTION:** A drive carrier module cannot be installed if its anti-tamper lock is activated outside the enclosure. See step 1 of the procedure described on page 122.

1. Release the drive carrier handle by depressing the latch in the handle.
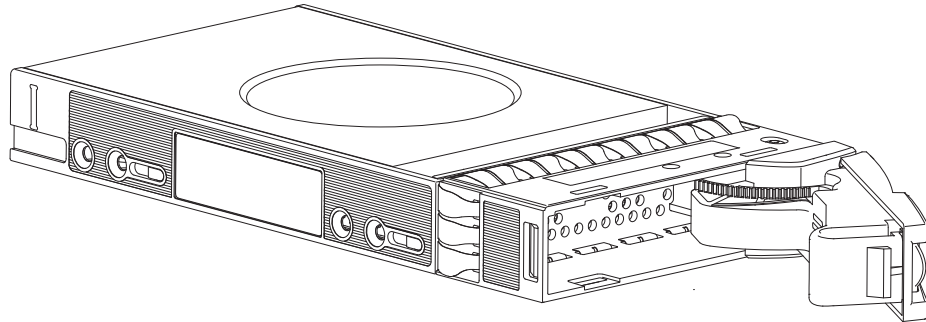
Figure 90 LFF drive carrier module in open position

2. Insert the drive carrier module into the enclosure. Make sure that the drive carrier is positioned such that the top of the disk is facing up, and the handle opens from the left as you face the enclosure front panel.
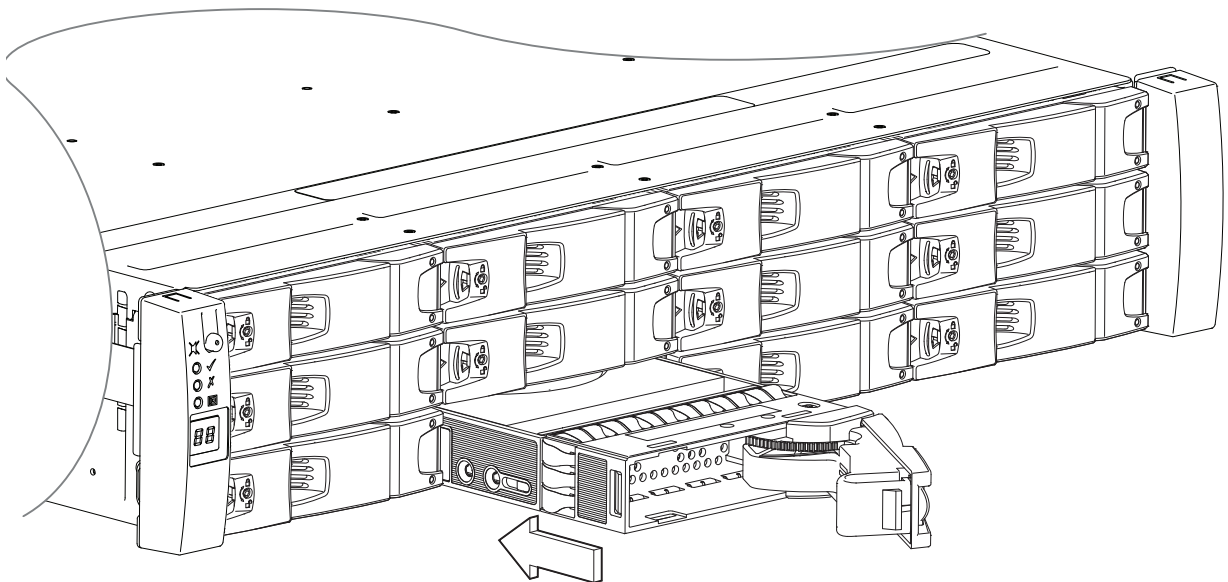
Figure 91 Installing a LFF drive carrier module (1 of 2)

3. Slide the drive carrier fully into the enclosure.

4. Cam the drive carrier home. The camming foot on the carrier will engage into a slot in the enclosure. Continue to push firmly until the handle fully engages. You should hear a click as the latch handle engages and holds the handle closed.
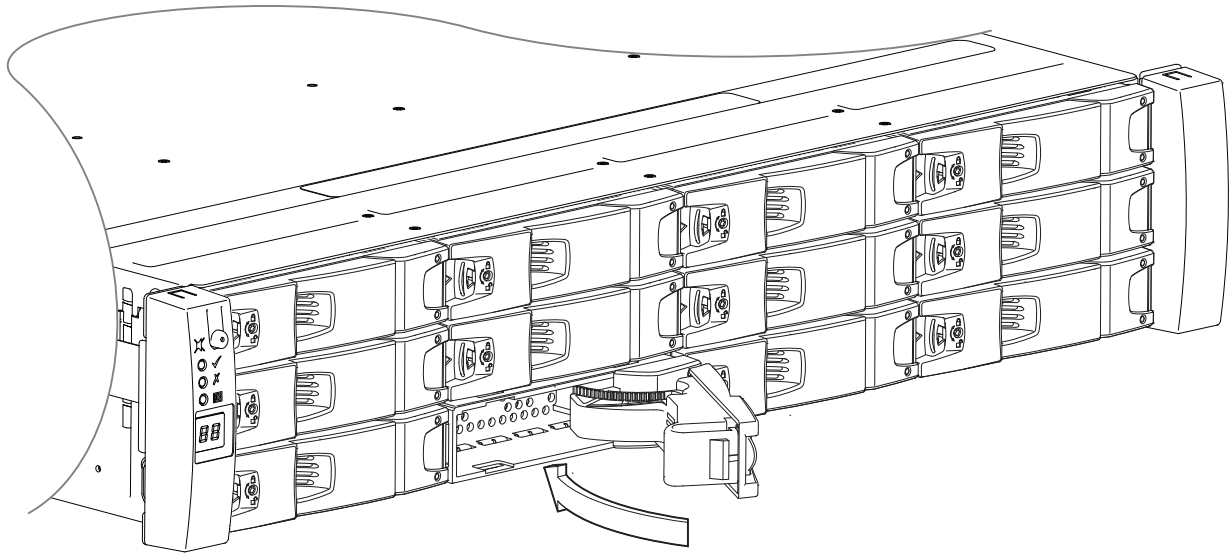
Figure 92 Installing a LFF drive carrier module (2 of 2)

5. Using the management interfaces (the SMC or CLI), verify whether the health of the new disk is OK. Verify that the Green Drive LED is on/blinking per Figure 83 (page 94). Verify that Ops panel states show no amber module faults.

Using the anti-tamper locks

The LFF and SFF disk drive modules provide anti-tamper locks. Please refer to Figure 93 when using this procedure.

1. Carefully put the T10 lock key into the cutout in the handle.

   The LFF disk module is shown in the example below. On SFF disk modules, the socket is located between the latch and the carrier handle as shown in Figure 83 (page 94). The lock socket mechanism is identical on both modules.

2. Position the key into its socket.

3. Perform one of the following actions to either activate or deactivate the anti-tamper lock:

   o Lock: Rotate the key in a clockwise direction until the indicator is visible in the aperture beside the key.

   o Unlock: Rotate the key in a counterclockwise direction until the indicator is visible in the aperture beside the key.

4. Remove the T10 lock key.
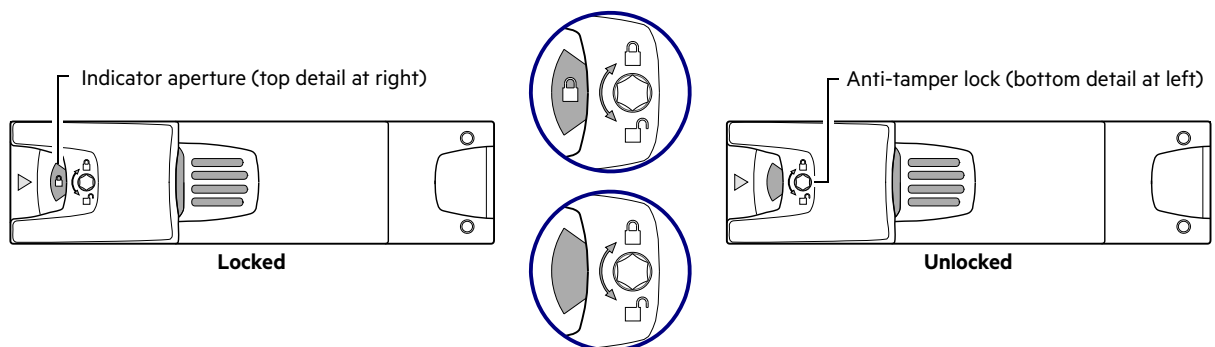
Figure 93 Activating the anti-tamper lock

Removing a SFF drive carrier module

The removal/replacement procedure for SFF drive carrier modules is basically the same as for LFF models, except that the SFF carriers are mounted vertically.

1. Press the latch in the carrier handle downward to release the carrier handle so that it can revolve outward as shown below.
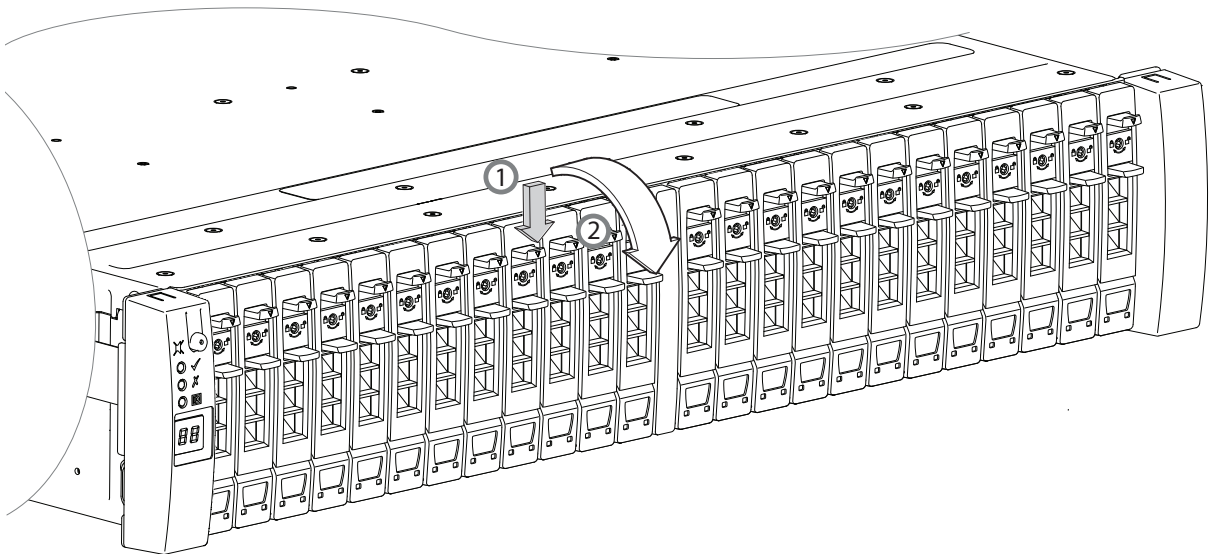
Figure 94 Removing a SFF disk drive module (1 of 2)

2. Gently move the drive carrier module outward from the drive slot.
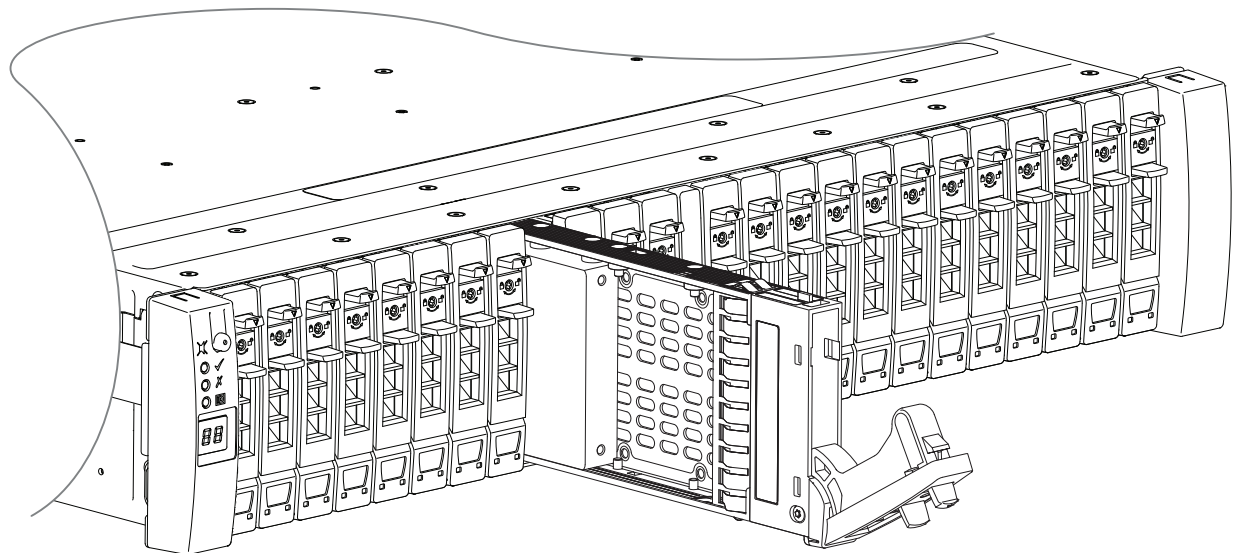
Figure 95 Removing a SFF disk drive module (2 of 2)

3. Remove the module fully from the drive slot.

△ **CAUTION:**   To ensure optimal cooling throughout the enclosure, dummy drive carrier modules must be fitted to all unused drive slots. See also Figure 43 (page 50).

Installing a SFF drive carrier module

1. If the anti-tamper lock is engaged, unlock it per the instructions provided in "Using the anti-tamper locks" (page 122).

2. Release the carrier handle by pressing the latch in the handle downwards, and opening the hinged handle as shown in Figure 96 (page 124). Insert the carrier into the enclosure in a vertical position.

ⓘ **IMPORTANT:** Make sure the carrier is positioned such that the disk is on its left side and the handle opens from the top.
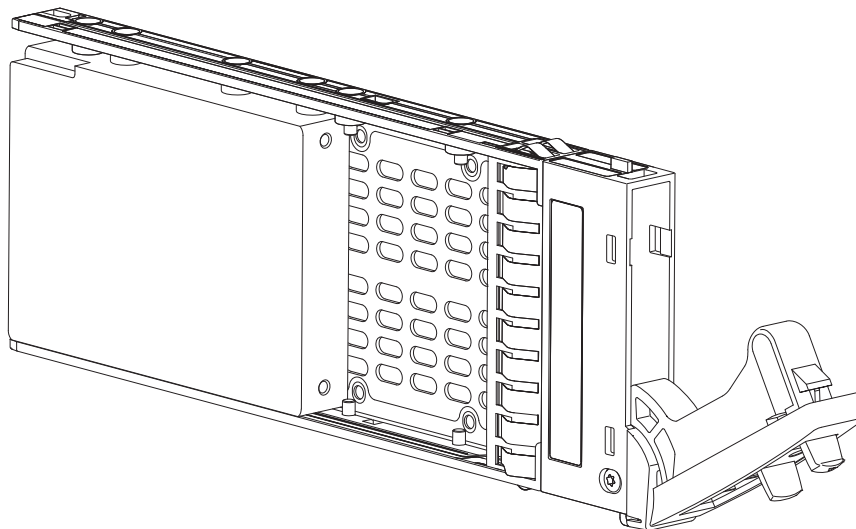
Figure 96 SFF drive carrier module in open position

3. Slide the carrier fully into the enclosure until it is stopped by the camming lever on the bottom of the carrier.
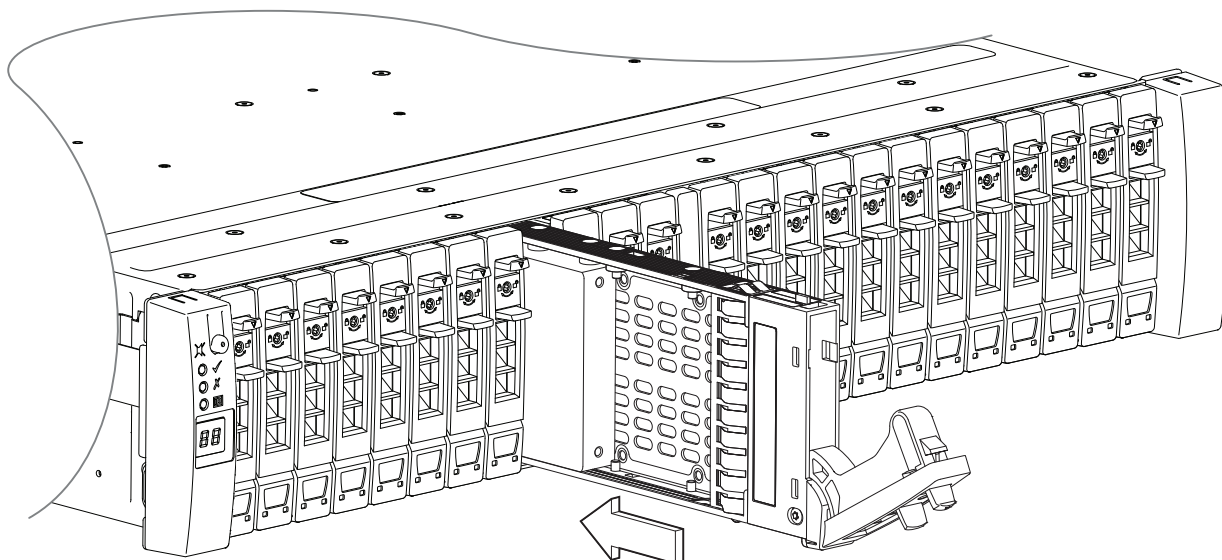
Figure 97 Installing a SFF drive carrier module (1 of 2)

4. Cam the carrier home. The camming lever on the carrier will engage into the a slot in the enclosure. Continue to push firmly until the handle fully engages. You should hear a click as the latch handle engages and holds the handle closed.
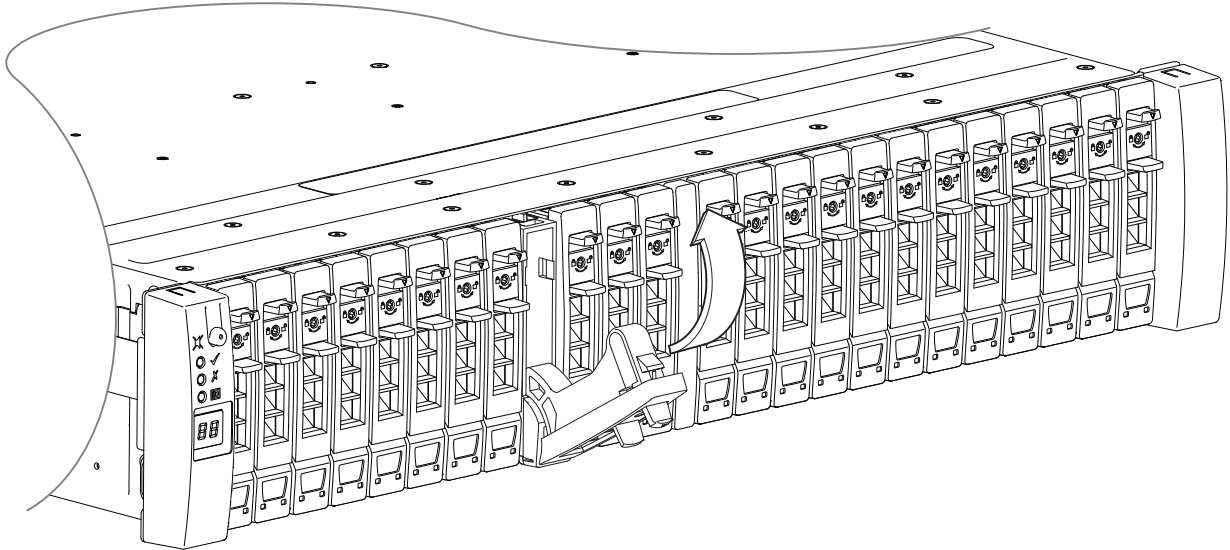
Figure 98 Installing a SFF drive carrier module (2 of 2)

5. Using the management interfaces (the SMC or CLI), verify whether the health of the new disk is OK. Verify that the Green Drive LED is on/blinking per Figure (page 93). Verify that Ops panel states show no amber module faults.

Replacing a dummy drive carrier

A dummy drive carrier module is removed from the enclosure simply by pulling the module out of the drive slot. A dummy drive carrier module is installed in the enclosure by properly aligning it and inserting it into the drive slot, followed by pushing it securely into place.

△ **CAUTION:**   To ensure optimal cooling throughout the enclosure, dummy drive carrier modules must be fitted to all unused drive slots. See also Figure 43 (page 50).

Replacing an IOM

ⓘ **IMPORTANT:**   The 5005/4005/3005 Series storage enclosures support dual-controller configuration only. If a partner controller fails, the storage system will fail over and run on a single controller module until the redundancy is restored. A controller module must be installed in each IOM slot to ensure sufficient air flow through the enclosure during operation.

I/O module (IOM) denotes either a controller module (RAID canister) or an expansion module (expansion canister). In a dual-controller configuration, controller and expansion modules are hot-swappable, which means you can replace one module without halting I/O to disk groups, or powering off the enclosure. In this case, the second module takes over operation of the storage system until you install the new module.

You may need to replace a controller module or an expansion module when:

- The Fault LED is illuminated
- Health status reporting in the SMC indicates a problem with the module
- Events in the SMC indicate a problem with the module
- Troubleshooting indicates a problem with the module

☼ **TIP:** The illustrations show IOM module replacement within the top slot (A) as you view the enclosure rear panel. To replace an IOM in the bottom slot (B), you would first rotate the module 180º about its longitudinal axis, so that it properly aligns with its connectors on the back of the midplane. See also Figure 12 (page 26).

Before you begin

Removing a controller or expansion module from an operational enclosure significantly changes air flow within the enclosure. Openings must be populated for the enclosure to cool properly. Leave modules in the enclosure until ready to install a replacement. If replacing both controllers in a dual-controller enclosure, use the SMC to record configuration settings before installing the new controller modules. See "Removing an IOM" (page 129), and "Installing an IOM" (page 130) for instructions on installing an additional controller module.

△ **CAUTION:** When replacing a controller module, ensure that less than 10 seconds elapse between inserting it into a slot and fully latching it in place. Not doing so might cause the controller to fail. If it is not latched within 10 seconds, remove the controller module from the slot, and repeat the process.

When two controller modules are installed in the enclosure, they must be of the same model type. When replacing both controller modules in an operational enclosure, follow these guidelines:

1. Replace one controller as described in these instructions.
2. Wait 30 minutes: this pause ensures that the controller and its ownership of disk groups has sufficient time to stabilize. See also "Verifying component operation" (page 131).
3. Check the system status and event logs to verify that the system is stable.
4. Replace the partner controller as described in these instructions.

Configuring partner firmware update for RBOD/EBOD configurations

In a dual-controller system in which PFU is enabled, when you update firmware on one controller, the system automatically updates the partner controller. Disable partner firmware update *only* if requested by a service technician.

Use the SMC or CLI to change the PFU setting.

ⓘ **IMPORTANT:** See the "Updating firmware" topic within the Storage Management Guide before performing a firmware update.

NOTE: The SMC and CLI provide an option for enabling or disabling Partner Firmware Update for the partner controller as described in the Storage Management Guide. 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.

Firmware update for JBOD configurations

Refer to the JBOD firmware update topics in the "Connecting a JBOD" appendix.

Verifying component failure

For RBOD/EBOD configurations, select from the following methods to verify component failure:

- Use the SMC to check the health icons/values of the system and its components to either ensure that everything is okay, or to drill down to a problem component. The SMC uses health icons to show OK, Degraded, Fault, or Unknown status for the system and its components. If you discover a problem component, follow the actions in its Recommendation field to resolve the problem.

- As an alternative to using the SMC, you can run the CLI show system command 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. If you discover a problem component, follow the actions in its Health Recommendations field to resolve the problem.

- Monitor event notification — With event notification configured and enabled, use the SMC to view the event log, or run the CLI show events detail command to see details for events.

- Check Fault LED (back of enclosure on IOM face plate): Amber = Fault condition.

- Check that the OK LED (back of enclosure) is off.

NOTE: For JBOD-related actions, see "Determine where the fault is occurring" (page 155).

Stopping I/O for RBOD/EBOD configurations

When troubleshooting disk drive and connectivity faults, stop I/O to the affected disk groups from all hosts as a data protection precaution. As an additional data protection precaution, it is helpful 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 that there is no I/O activity by briefly monitoring the system LEDs; however, when accessing the storage system remotely, this is not possible. Remotely, you can use the show disk-group-statistics command to determine if input and output has stopped on RBOD/EBOD systems. Perform these steps:

1. Using the CLI, run the show disk-group-statistics command.

   The Reads and Writes fields 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 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.

NOTE: See the CLI Reference Guide for additional information. Optionally, you can use the SMC to monitor IOPs and MB/s.

Shutting down a controller module

Shutting down the Storage Controller in a controller module ensures that a proper failover sequence is used, which includes stopping all I/O operations and writing any data in write cache to disk. If the Storage Controller in both controller modules is shut down, hosts cannot access the system's data. Perform a shut down before you remove a controller module from an enclosure, or before you power off its enclosure for maintenance, repair, or a move.

Use the SMC or the CLI to perform a shut down.

Using the SMC

1. Sign-in to the SMC
2. In the System panel in the banner, select **Restart System**.

   The Controller Restart and Shut Down panel opens.
3. Select the **Shut Down** operation, which automatically selects the controller type **Storage**.
4. Select the controller module to shut down: A, B, or both.
5. Select **OK**. A confirmation panel appears.
6. Select **Yes** to continue; otherwise, select **No**. If you selected Yes, a message describes shutdown activity.

   > **NOTE:** If an iSCSI port is connected to a Microsoft Windows host, the following event is recorded in the Windows event log: `Initiator failed to connect to the target.`

NOTE: See the Storage Management Guide for additional information.

Using the CLI

1. Log-in to the CLI.
2. In your dual-controller system, verify that the partner controller is online by running the command;

   `show controllers`
3. Shut down the failed controller—A or B—by running the command:

   `shutdown a` *or* `shutdown b`

   The blue OK to Remove LED (back of enclosure) illuminates to indicate that the controller module can be safely removed.
4. Illuminate the white Identify LED of the enclosure that contains the controller module to remove by running the command:

   `set led enclosure 0 on`

   The Display LED on the Ops panel located on the enclosure left ear will be blinking green when the above command is invoked.

NOTE: See the CLI Reference Guide for additional information.

Removing an IOM

> (!) **IMPORTANT:** Considerations for removing controller modules:
>
> - In a dual-controller environment, you may hot-swap a single controller module in an operational enclosure, provided you first shut down the faulty controller using the SMC or the CLI.
> - In a dual-controller environment—if replacing both controller modules—you must adhere to the instructions provided in "Before you begin" (page 126).
> - Do not remove a faulty module unless its replacement is on-hand. All modules must be in place when the system is in operation.

See CAUTION bullets regarding *electrostatic discharge* and *anti-static protection* on page 116.

Illustrations in the controller module replacement procedures show rear panel views of the 4-port enclosure, but also apply to the 2-port enclosure; IOMs are properly aligned for insertion into the IOM slots.

1. Verify that you have successfully shut down the controller module using the SMC or the CLI.

2. Locate the enclosure whose UID LED (Ops panel on enclosure front left ear) is illuminated, and within the enclosure, locate the controller module whose OK to Remove LED is blue (rear panel).

3. Disconnect any cables connected to the controller.

   Label each cable to facilitate re-connection to the replacement IOM.

4. Grasp the module latch between the thumb and forefinger, and squeeze the flange and handle together to release the latch handle from its docking member (detail No.1), and swing the latch out to release the IOM from its seated position (detail No.2) as shown in Figure 99 and Figure 100 (page 130).
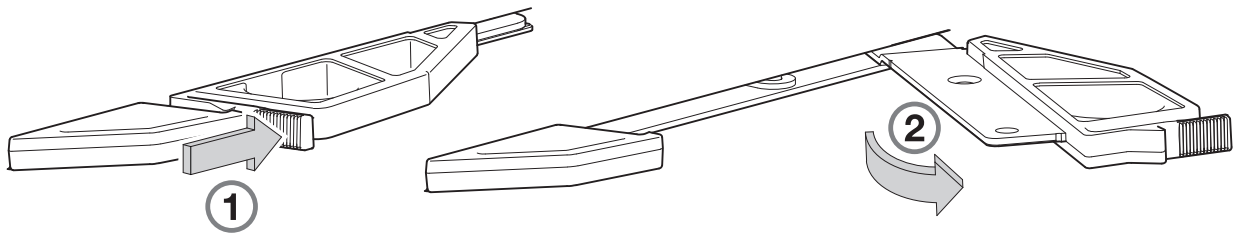
Figure 99 IOM latch operation detail

5. Swing the latch handle open as shown in detail No.2 within Figure 99 and Figure 100.
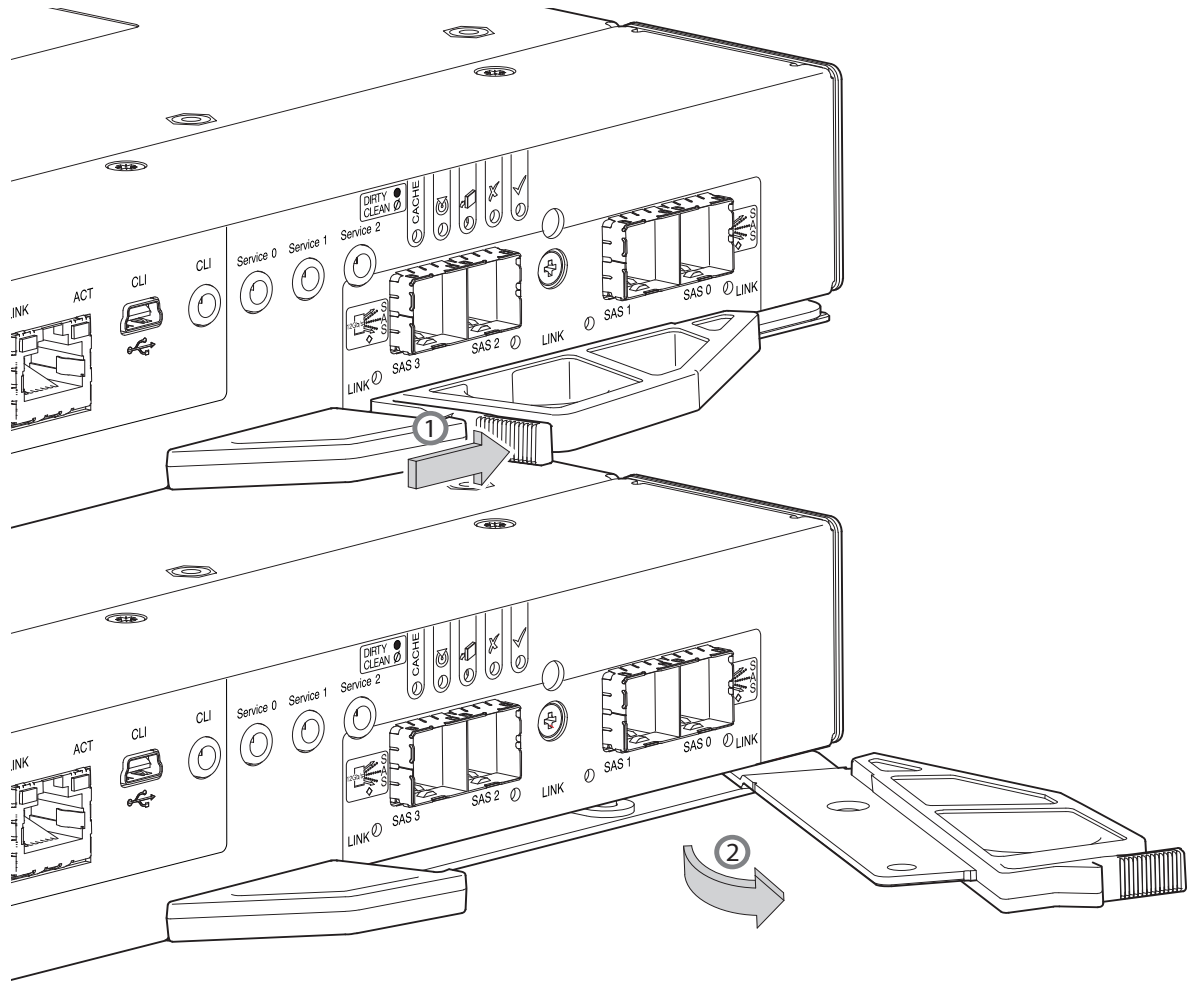
4-port FC/iSCSI controller module shown

Figure 100 Removing an IOM

6. Grip the latch handle and ease the IOM forward from the slot as shown in detail No.2 within Figure 100.

NOTE: The illustration above shows a 4-port controller module of model type SAS. However, the procedure applies to all IOMs discussed herein, including the 2-port model. They all use the same latch mechanism, but feature different face plate geometry.

---

7. Place both hands on the canister body, and pull it straight out of the enclosure such that the controller module remains level during removal.

Installing an IOM

See CAUTION bullets regarding *electrostatic discharge* and *anti-static protection* on page 116.

△ **CAUTION:** If passive copper cables are connected, the cable must not have a connection to a common ground/earth point.

---

1. Examine the IOM for damage, and closely inspect the interface connector. Do not install if the pins are bent.
2. Grasp the IOM using both hands, and with the latch in the open position, orient the module and align it for insertion into the target IOM slot.
3. Ensuring that the IOM is level, slide it into the enclosure as far as it will go.

   A controller module that is only partially seated will prevent optimal performance of the controller enclosure. Verify that the controller module is fully seated before continuing.
4. Set the module in position by manually closing the latch.

   You should hear a click as the latch handle engages and secures the IOM to its connector on the back of the midplane.
5. Reconnect the cables.

NOTE: In a dual-controller system in which PFU is enabled, when you update the firmware on one controller, the system automatically updates the partner controller.

Verifying component operation

Controller module

After replacing the controller module, verify that the CRU OK LED (rear panel) illuminates green, indicating that the controller has completed initializing, and is online/operating normally. It may take two to five minutes for the replacement controller to become ready. If you are replacing either controller module, and PFU is enabled, you may need to wait 30 minutes to ensure that the two controllers—with their respective ownership of the disk groups—have enough time to fully stabilize.

ⓘ **IMPORTANT:** Use the SMC or CLI to perform a restart *only* if necessary. See the topic about restarting controllers in the Storage Management Guide for more information.

Expansion module

If the storage system is configured with expansion enclosures, the replacement expansion module may take up to one minute to initialize after the cables are connected. Verify that firmware on both expansion modules is compatible and current.

NOTE: See the topic about updating expansion module firmware within the Storage Management Guide.

Replacing a 2U storage enclosure chassis

The controller enclosure or expansion enclosure chassis replacement procedure replaces a damaged chassis CRU. The procedure includes removing all CRU modules from a damaged chassis and installing them into a replacement chassis.

ⓘ **IMPORTANT:** The 5005/4005/3005 Series enclosures using 2U12 or 2U24 chassis are described in "5005 Series enclosure models" (page 111), "4005 Series enclosure models" (page 111) and "3005 Series enclosure models" (page 111).

Whether your controller enclosure product is a 2U12 or 2U24 model, a fully functional replacement chassis requires the successful removal and installation of the following components:

- All power and data cables
- All disk drive modules
- Two power cooling modules (both AC)
- Two IOMs of the same model type

This procedure references the CRU component procedures described elsewhere in this chapter.

Before you begin

⚠ **CAUTION:** Do not remove the enclosure until you have received the replacement enclosure.

See CAUTION bullets regarding *electrostatic discharge* and *anti-static protection* on page 116.

1. Schedule down time that will allow for shutdown; sixty minutes of replacement work; and restart.
2. Verify the existence of a known/good backup of the system.
3. Record system settings for future use and label all cables.
4. Prepare a suitable static-protected work environment to accommodate chassis replacement.

Verifying component failure

The controller enclosure CRU includes the enclosure's metal housing, the module runner system, the integrated Ops panel, and the assembled/installed midplane PCB that connects controller modules, disk drive modules, and power and cooling modules. See the empty enclosure chassis pictorials: 2U12–Figure 18 (page 29) and 2U24–Figure 19 (page 29). This CRU replaces an enclosure that has been damaged, or whose midplane has been damaged. Often, a damaged midplane will appear as though a RAID canister has failed. If you replace a controller module and it does not remedy the fault, you may need to replace the enclosure chassis.

You can observe enclosure health (front panel and rear panel) using management interfaces to verify enclosure/component failure or enclosure/component operation. See also "Using RBOD/EBOD management interfaces" (page 135), and "Using JBOD management interfaces" (page 135).

Preparing to remove a damaged 2U storage enclosure chassis

Because you are removing and replacing an entire storage enclosure chassis, neither the hot-swap capability that applies to replacing individual redundant CRUs in an operational storage enclosure, nor the hot-add of a drive enclosure to an operational storage system, apply to this procedure.

1. Stop all I/O from hosts to the system. See also "Stopping I/O for RBOD/EBOD configurations" (page 127).
2. Shut down the controllers. See also "Shutting down a controller module" (page 128).
3. Power off the system – controller enclosure *first*; expansion enclosures *next*: see "Powering on/powering off" (page 88).

Table 34   Replacing a 2U storage enclosure chassis and installing its CRUs

To accomplish this sequential process		See the following procedures	
1.	Facing the front of the 2U enclosure, remove the disk modules from the damaged chassis and place them in a static-protected work area nearby.[1]	a.	"Removing a LFF drive carrier module" (page 120), or
		b.	"Removing a SFF drive carrier module" (page 123)
2.	Facing the rear of the 2U enclosure, remove the power and data cables.	"Removing a damaged 2U storage enclosure chassis from the rack" (page 133) (step 1)	
3.	Remove the damaged storage enclosure chassis from the rack.	"Removing a damaged 2U storage enclosure chassis from the rack" (page 133)	

Table 34 Replacing a 2U storage enclosure chassis and installing its CRUs (continued)

To accomplish this sequential process	See the following procedures
4. Remove the two power and cooling modules (PCMs) from the damaged chassis rear panel, and install them in the replacement chassis.	**a.** See PCM related TIP on page 117
	b. "Removing a PCM" (page 117)
	c. "Installing a PCM" (page 119)
5. Remove the two IOMs from the damaged chassis rear panel, and install them in the replacement chassis.[2]	**a.** "Replacing an IOM" (page 125) (preparation subsections)
	b. "Removing an IOM" (page 129)
	c. "Installing an IOM" (page 130)
	d. "Verifying component operation" (page 131).
6. Install the replacement storage enclosure in the rack.	"Installing the replacement 2U storage enclosure in the rack" (page 134)
7. Facing the front of the 2U enclosure, install disk drive modules into the replacement enclosure.[3]	**a.** "FDE considerations" (page 60)
	b. "Installing a LFF drive carrier module" (page 121), or
	c. "Installing a SFF drive carrier module" (page 124)
8. Complete the chassis replacement process.	"Completing the process" (page 134)

1-If you temporarily stack the disks before installing them, insert static dissipative foam between the drive modules.

2-Within the replacement enclosure, each IOM must be reinstalled into the same IOM slot from which it was extracted from the damaged enclosure (0A > 0A and 0B > 0B, respectively). For FC/iSCSI IOMs, the SFPs installed in the CNC ports need not be removed. For SFP replacement, see "SFP option for CNC ports" (page 178).

3-Within the replacement enclosure, reinstall each disk drive or disk drive blank into the same disk slot from which it was removed from the damaged enclosure. To ensure optimal cooling throughout the enclosure, dummy drive carriers must be fitted to all unused drive slots.

Removing a damaged 2U storage enclosure chassis from the rack

This section provides a procedure for removing a damaged storage enclosure chassis from its rack location.

△ **CAUTION:**   It is recommended that all disk drive modules be removed before removing the enclosure. Two people are required to move the 2U enclosure.

---

1. Disconnect the power cables and data cables between devices as needed:
    a. Between the cascaded enclosures.
    b. Between the controller and peripheral SAN devices.
    c. Between the controller and the host.

    NOTE: Label the cables to facilitate reconnection of the storage devices.

    ---

2. Remove the retaining screws that secure the front and rear of the chassis to the rack and rails.

    NOTE: Do not remove the ear components from the failed chassis: integrated ear components and covers are provided with the replacement chassis CRU.

    ---

3. Maintaining a level position, carefully slide the enclosure chassis from the rack.
4. Place the chassis on a static-protected work surface near the replacement enclosure chassis, with the removed disk drive modules and screws.

Installing the replacement 2U storage enclosure in the rack

This section provides a procedure for installing the replacement enclosure in its rack location.

ⓘ **IMPORTANT:**   Refer to "Installing the 2U enclosure" (page 58) when installing the enclosure into the rack.

ⓘ **IMPORTANT:**   Install the enclosure into the rack before re-inserting the disk drive modules. Two people are required to move the enclosure.

---

1. Support the bottom of the enclosure chassis. Carefully lift/align the chassis and while maintaining a level position for the enclosure, slide it into the rack.
2. Using the appropriate mounting hardware, secure the enclosure chassis to the rack.
3. Using the applicable retaining screws, secure the front and rear of the enclosure chassis to the rack and rails.

Completing the process

This section provides a procedure for ensuring that the CRU components installed in the replacement controller enclosure chassis function properly.

1. Reconnect data cables between devices, as needed, to return to the original cabling configuration:
   o Between cascaded storage enclosures.
   o Between the controller and peripheral or SAN devices.
   o Between the controller enclosure and the host.
2. Reconnect power cables to the storage enclosures. See "Powering on/powering off" (page 88).

Verifying component operation

1. Restart system devices by moving the power switch on the power supply to the **On** position in the following sequence:
   a. Expansion enclosures *first*.
   b. Controller enclosure *next* (RBOD/EBOD configurations).
   c. Data host *last* (if powered down for maintenance purposes).

   Allow time for each device to complete its Power On Self Tests (POST) before proceeding.
2. For RBOD/EBOD configurations, perform a rescan to force a fresh discovery of all expansion enclosures connected to the controller enclosure. This step clears the internal SAS layout information, reassigns enclosure IDs, and ensures the enclosures are displayed in the proper order. Use the CLI or SMC to perform the rescan:

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

   ```
   rescan
   ```

   To perform a rescan using the SMC:
   a. Verify that both controllers are operating normally.
   b. In the **System** topic, select **Action > Rescan Disk Channels**.
   c. Select **Rescan**.

Using LEDs

View LEDs on the enclosure front and rear panels.

Verify front panel LEDs

Front panel LEDs reside on the Ops panel located on the left ear flange. Disk LEDs are located on the carrier modules.

- Verify that the System Power On/Standby LED is illuminated green, and that the Module Fault LED is not illuminated.
- Verify that the enclosure ID LED located on the left ear is illuminated green.
- Verify that the disk module's Green LED is illuminated green or blinking green, and that the disk module's Amber LED is not illuminated.

Verify rear panel LEDs

Rear panel LEDs reside on PCM and IOM face plates.

- Verify that each power cooling module's PCM OK LED is illuminated green.
- For IOMs, verify that the OK LED is illuminated green, indicating that the module has completed initializing, and is online.

Using RBOD/EBOD management interfaces

In addition to viewing LEDs as described above, you can use management interfaces to monitor the health status of the system and its components, provided you have configured and provisioned the system, and enabled event notification.

Select from the following methods to verify component operation:

- Use the SMC to check the health icons/values of the system and its components to either ensure that everything is okay, or to drill down to a problem component. The SMC uses health icons to show OK, Degraded, Fault, or Unknown status for the system and its components. If you discover a problem component, follow the actions in its Recommendation field to resolve the problem.
- As an alternative to using the SMC, 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. If you discover a problem component, follow the actions in its Health Recommendations 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 the SMC or the CLI. Using the SMC, you would view the event log and then hover over 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 (see the CLI Reference Guide for more information about command parameters and syntax).

Using JBOD management interfaces

For server-attached JBOD configurations, you should be able to reconnect the storage system cables, and then power on the enclosure(s). The host application and enclosure management interfaces should interact to discover storage components and bring the storage system online. See also "JBOD firmware update procedure" (page 157) if a firmware upgrade is necessary.

CRU replacement for 5U chassis

NOTE: Unless noted otherwise within a passage pertaining to a particular CRU, the replacement procedure should be completed within two minutes of the removal of a defective module.

Accessing drawers

To observe or replace a DDIC, you must open the drawer in which it resides. The top drawer (Drawer 0) and the bottom drawer (Drawer 1) are accessed from the enclosure front panel. See also Figure 20 (page 30).

Opening a drawer

1. Verify that the anti-tamper locks are not engaged. The red arrows on the locks point inwards if the locks are disengaged as shown in Figure 101. If necessary, unlock them by rotating counter-clockwise using a Torx T20 bit.

Left side   Right side
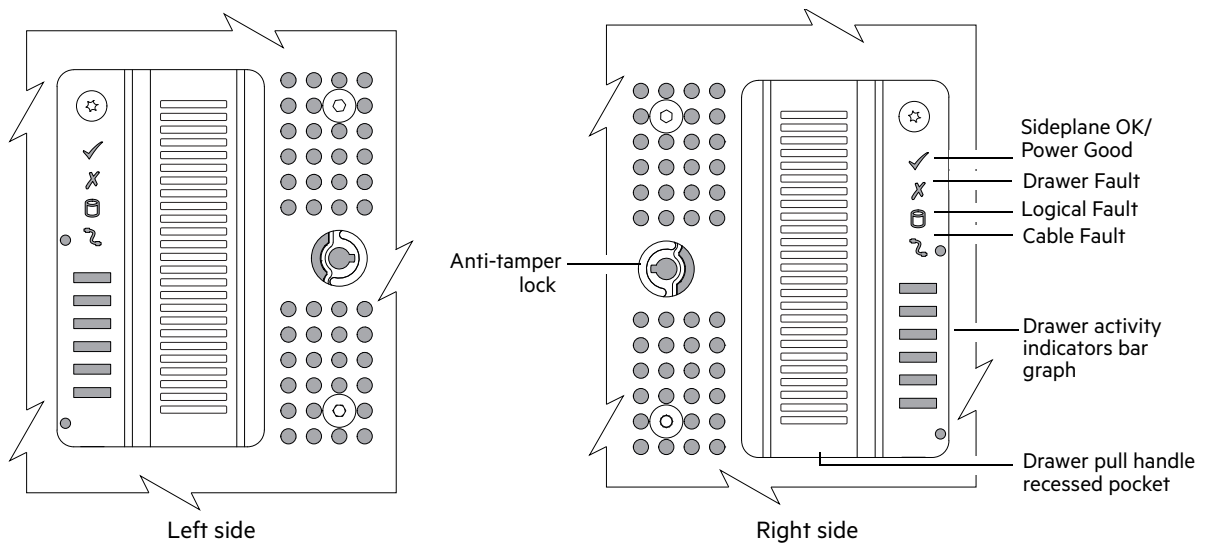
Figure 101 Drawer front panel details

2. Push the drawer latches inward and hold them as shown in Figure 102.
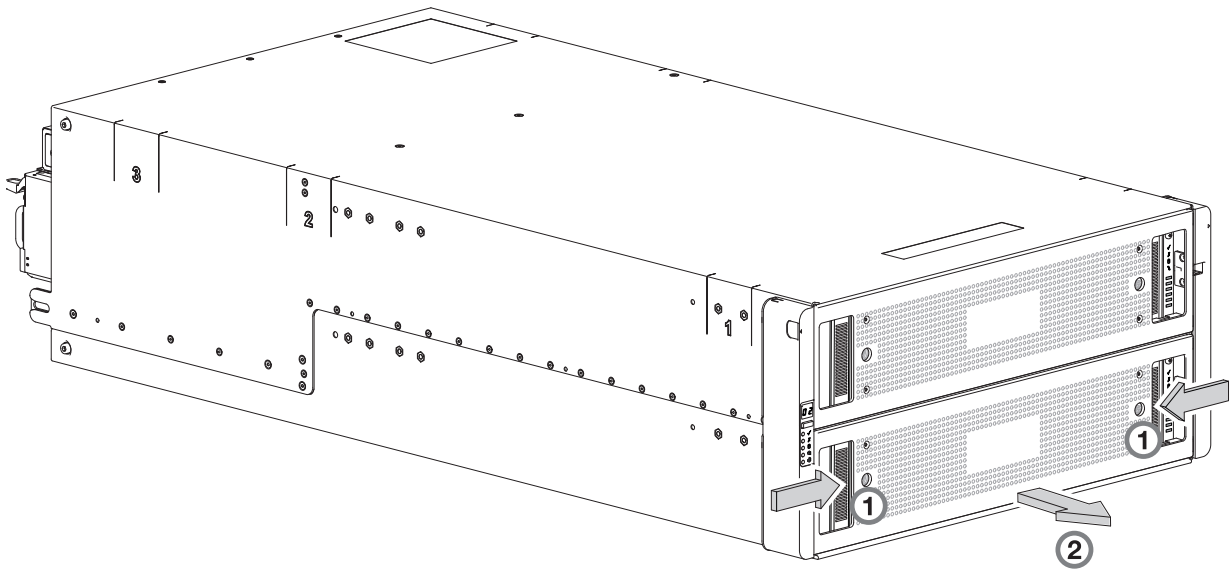
Figure 102 Opening a drawer (1 of 2)

3. Pull the drawer outward until it locks at the drawer stops as shown in Figure 103.

The drawer is shown empty, which is how the enclosure is delivered.

Slide rail latch inset
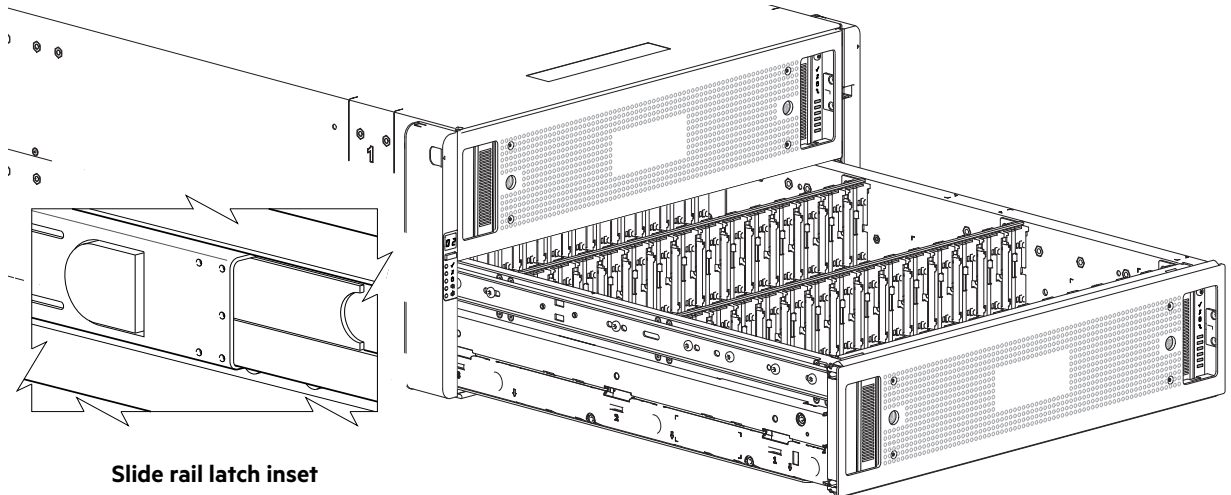
Figure 103 Opening a drawer (2 of 2)

ⓘ **IMPORTANT:**   The drawer must not remain open for more than two minutes whilst the enclosure is powered on.

Closing a drawer

1. Press and hold the black latches on the sides of the open drawer in each extended top rail.

    Figure 103 shows a magnified detail of a slide latch, which resides on the left and right drawer rails.

2. Push the drawer in slightly.

3. Release the drawer latches.

4. Push the drawer all the way into the enclosure, making sure that it clicks home.

Replacing a DDIC

This procedure describes removal and installation of a Disk Drive in Carrier (DDIC). Illustrations in the DDIC replacement procedures show drawer-centric views of the enclosure, and DDICs are properly aligned for insertion into the disk slots.

See CAUTION bullets regarding *electrostatic discharge* and *anti-static protection* on page 116.

Removing a DDIC

NOTE: Remove a DDIC only if a replacement is available. Closing a drawer with one or more drives missing can potentially cause cooling problems. See also "Populating drawers" (page 139).

---

1. Determine which drawer contains the disk to be replaced.

    o   If the drive number is known, use the information contained in Figure 21 (page 31), which provides a single plan view of a drawer that is dual-indexed with *top drawer* (left integer) and *bottom drawer* (right integer) slot numbering.

    o   If the disk has failed, a fault LED is lit on the front panel of the affected drawer. The illuminated LED will either be the Drawer LED or the Logical LED.

    o   If the disk has failed, the Drive Fault LED on the DDIC cover is lit amber.

2. Open the relevant drawer per the instructions provided in "Opening a drawer" (page 136).

3. Locate the DDIC to be replaced using any of the methods listed in step 1 above.

4. On the face of the DDIC, push the latch button in the direction shown in Figure 104 (page 138) to unlock the DDIC from its seated position in the slot.

Figure 104 Removing a DDIC (1 of 2)

5. Pull the DDIC upwards and out of the drawer slot.
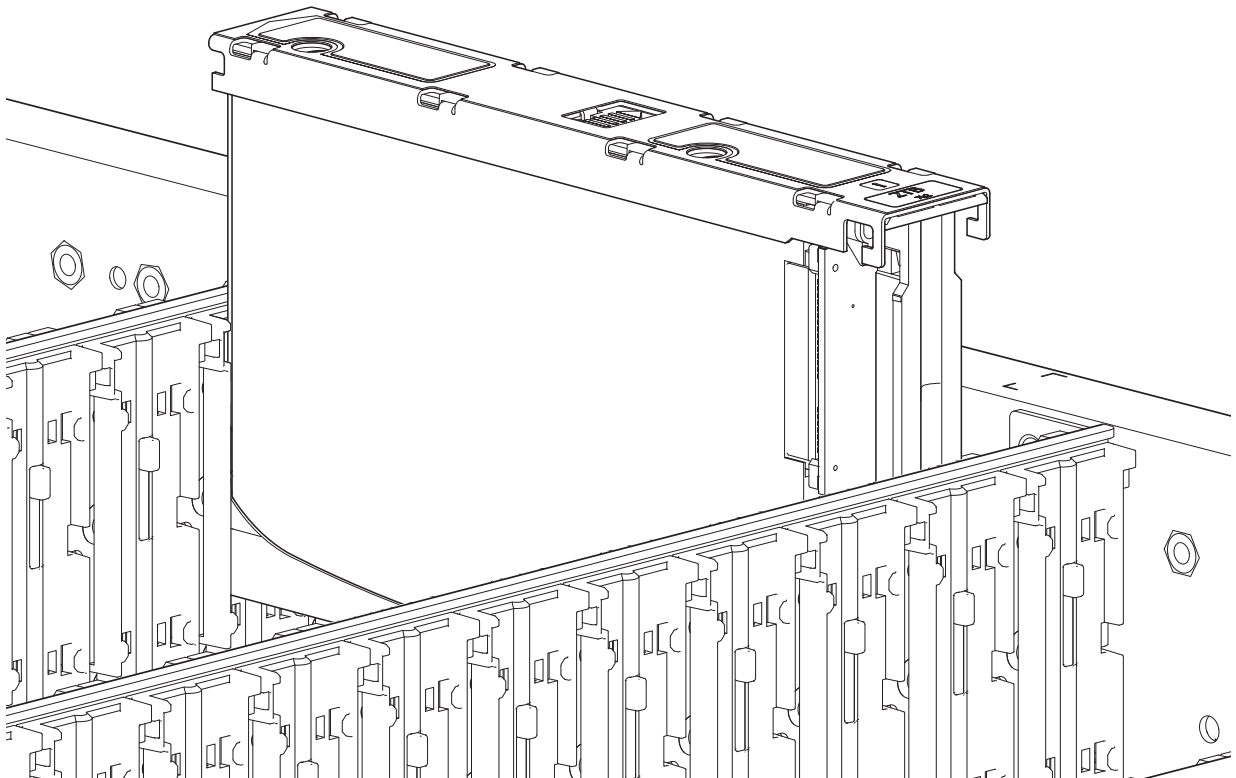
Figure 105 Removing a DDIC (2 of 2)

Installing a DDIC

> ⊕ **IMPORTANT:** Failed disks must be replaced with approved disks. Contact your service provider for details.

1. Open the relevant drawer per the instructions provided in "Opening a drawer" (page 136).
2. Align the DDIC with the target disk slot as shown in Figure 105 (page 138) and insert it into the disk slot.
3. Lower the DDIC into the disk slot.
   a. Push the DDIC downwards and hold it down.
   b. Move the slide latch in the direction shown in Figure 106.

DDIC
(longitudinal view–top face of carrier)

Slide latch   Latch button   Drive Fault LED
                     (locked)
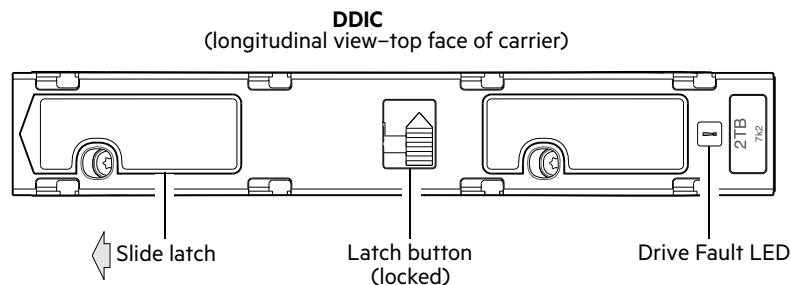
Figure 106 Installing a DDIC

4. Verify the following:
   a. The latch button is in the locked position, as shown in Figure 106.
   b. The Drive Fault LED is not lit.
5. Close the drawer per the instructions provided in "Closing a drawer" (page 137).

Populating drawers

General guidelines for populating a drawer with DDICs are provided at "Populating drawers with DDICs" (page 51). Additional guidelines are provided for replacing disks in previously populated drawers, or populating enclosures delivered with the half-populated enclosure configuration option.

Preparation

Customers with multiple enclosures may spread the 42 disks of an expansion package across those enclosures, provided the DDICs are installed 14 at a time to completely fill empty rows. The installation pattern providing the best airflow and thermal performance is described in this section.

The drawers must be populated with DDICs in whole rows. Each drawer contains 3 rows of 14 DDICs. Rules and assumptions are listed:

- The minimum number of disks in an enclosure is 14.
- The number of rows must not differ by more than 1 between the top and bottom drawers.
- The rows should be populated from front to rear of drawer.
- The disks of the expansion package must match the disks originally shipped with the 5U84 enclosure. Both groups of disks must share the same model type and capacity.

  > **NOTE:** Part numbers for expansion packages are not listed because they change over time when disks ship with new firmware, or new disk models become available. Contact your account manager for part numbers.

- If the two groups of disks have different firmware, all disks must be updated with current/compatible firmware. See the Storage Management Guide or online help for additional information about updating firmware.

Installation guidelines

The recommended order for partially populating drives in the 5U84 enclosure optimizes the airflow through the chassis. Please reference the following illustrations:

- Figure 20 (page 30) shows location and indexing of drawers accessed from the enclosure front panel.
- Figure 21 (page 31) shows dual-indexing of disk slots across front, middle, and back rows of a drawer.

The 5U84 ships with drawers installed in the chassis, However, to avoid shock and vibration issues during transit, the enclosure does not ship with DDICs installed in the drawers. An enclosure is configured with either 42 disks (half-populated) or 84 disks (fully populated) for customer delivery. If half-populated, the rows containing disks should be populated with a full complement of DDICs (no blank slots in the row). The bullet-list below identifies rows in drawers that should contain DDICs when the enclosure is configured as half-populated:

- Top drawer–front row
- Top drawer–middle row
- Bottom drawer–front row

If additional disks are incrementally installed into a half-populated enclosure, the DDICs must be added one row at a time (no blank slots in row) in the sequence listed:

- Bottom drawer–middle row
- Top drawer–back row
- Bottom drawer–back row

Replacing a PSU

Illustrations in the power supply unit module replacement procedures show rear panel views of the enclosure, and PSUs are properly aligned for insertion into the PSU slots.

See CAUTION bullets regarding *electrostatic discharge* and *anti-static protection* on page 116.
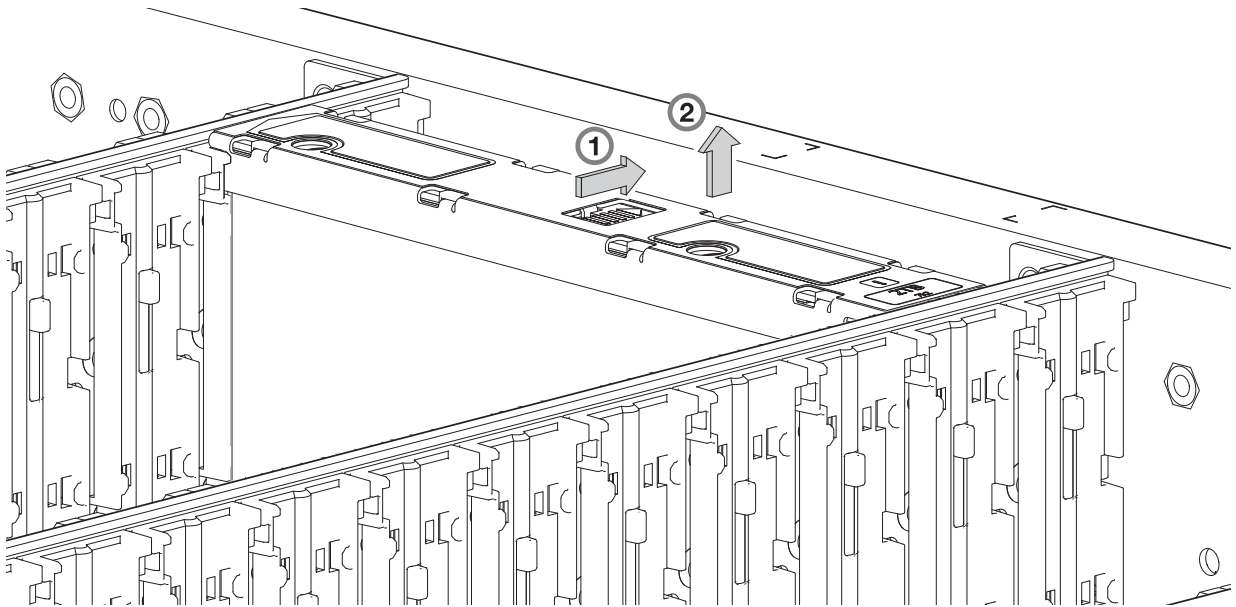
Removing a PSU

⚠ **CAUTION:**   Removing a power supply unit significantly disrupts the enclosure's airflow. Do not remove the PSU until you have received the replacement module. It is important that all slots are filled when the enclosure is in operation.

Before removing the PSU, disconnect the power from the PSU by either the mains switch (where present) or by physically removing the power source in order to ensure your system has warning of imminent power shutdown. Ensure that you correctly identify the faulty PSU before beginning the step procedure.

1. Stop all I/O from hosts to the enclosure. See also "Stopping I/O for RBOD/EBOD configurations" (page 127).

   ☼ **TIP:**   This step is not required for hot-swapping. However, it is required when replacing both PSUs at once.

---

2. Use management software to shut down any other system components necessary.

   ☼ **TIP:**   This step is not required for hot-swapping. However, it is required when replacing both PSUs at once.

---

3. Per Figure 31 (page 38), verify the Power OK LED is lit, then switch off the faulty PSU, and disconnect the power supply cable.
4. If replacing a single PSU via hot-swap, proceed to step 6.

5. If replacing both PSUs, verify that the enclosure was shut down using management interfaces, and that the enclosure is powered off.

6. Verify that the power cord is disconnected.

7. Refer to Figure 107 (page 141) when performing this step:

  a. Push the release latch to the right and hold it in place (detail No.1).

  b. With your other hand, grasp the handle and pull the PSU outward (detail No.2). Supporting the PSU with both hands, remove it from the enclosure. See also Figure 108 (page 141).
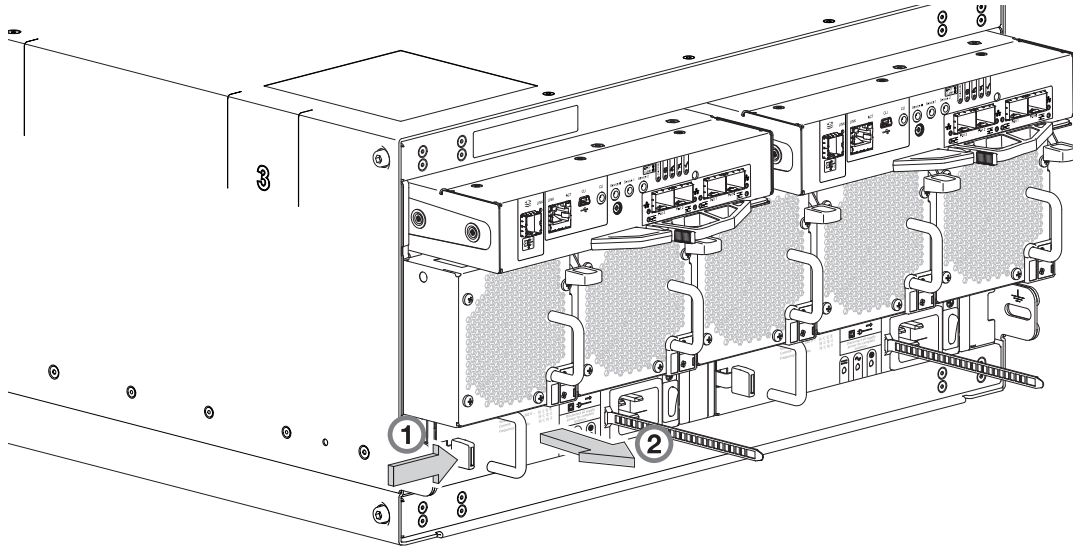
Figure 107 Removing a PSU (1 of 2)
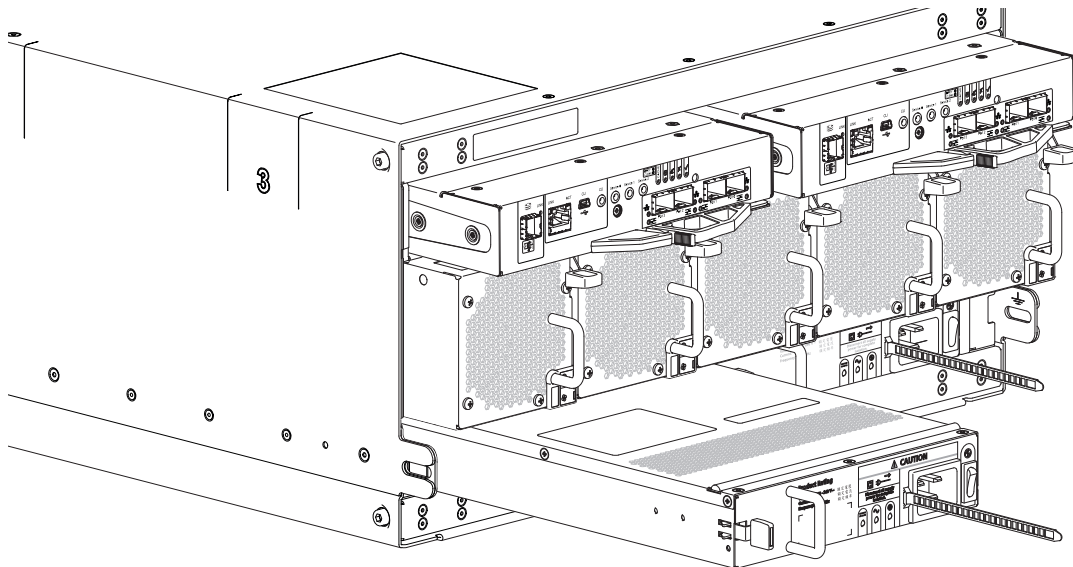
Figure 108 Removing a PSU (2 of 2)

8. If replacing both PSUs. repeat steps 5 through 7.

ⓘ **IMPORTANT:** The PSU slot must not be empty for more than 2 minutes while the enclosure is powered.

Installing a PSU

1. Ensure that the PSU is switched off.

   If replacing both PSUs, the enclosure must be powered off via an orderly shutdown using the management interfaces.

2. Orient the PSU for insertion into the target slot on the enclosure rear panel, as shown in Figure 108.

3. Slide the PSU into the slot until the latch clicks home.

4. Connect the AC power cord.

5. Move the PSU power switch to the **On** position.

6. Wait for the Power OK LED on the newly inserted PSU to illuminate green. See also Figure 31 (page 38).

   o If the Power OK LED does not illuminate, verify that the PSU is properly inserted and seated in the slot.

   o If properly seated, the module may be defective. Check the SMC and the event logs for more information.

   o Using the management interfaces (the SMC or CLI), determine if the health of the new PSU is OK. Verify that the Power OK LED is green, and that the Ops panel states show no amber module faults.

7. If replacing both PSUs, repeat steps 1 through 6.

Replacing an FCM

Illustrations in the fan cooling module replacement procedures show rear panel views of the enclosure, and FCMs are properly aligned for insertion into the FCM slots.

See CAUTION bullets regarding *electrostatic discharge* and *anti-static protection* on page 116.

Removing an FCM

⚠ **CAUTION:**   Removing an FCM significantly disrupts the enclosure's airflow. Do not remove the FCM until you have received the replacement module. It is important that all slots are filled when the enclosure is in operation.

---

1. Identify the fan cooling (FCM) module to be removed. If the FCM module has failed, the Fan Fault LED will illuminate amber. See also Figure 26 (page 33).

2. Refer to Figure 109 when performing this step:

   a. Push the release latch down and hold it in place (detail No.1).

   b. With your other hand, grasp the handle and pull the FCM outward (detail No.2). Supporting the FCM with both hands, remove it from the enclosure. See also Figure 110 (page 143).
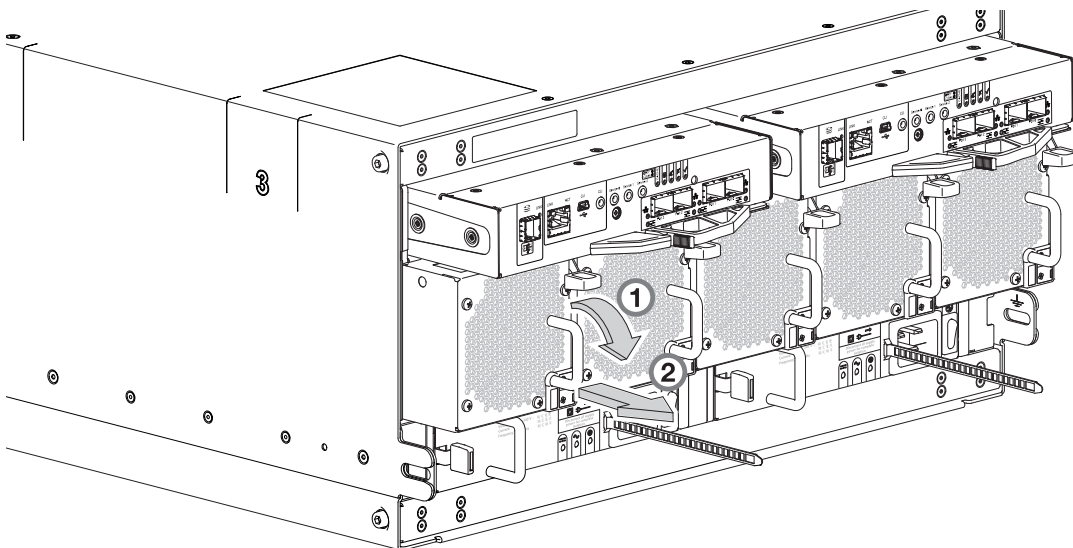
Figure 109 Removing an FCM (1 of 2)
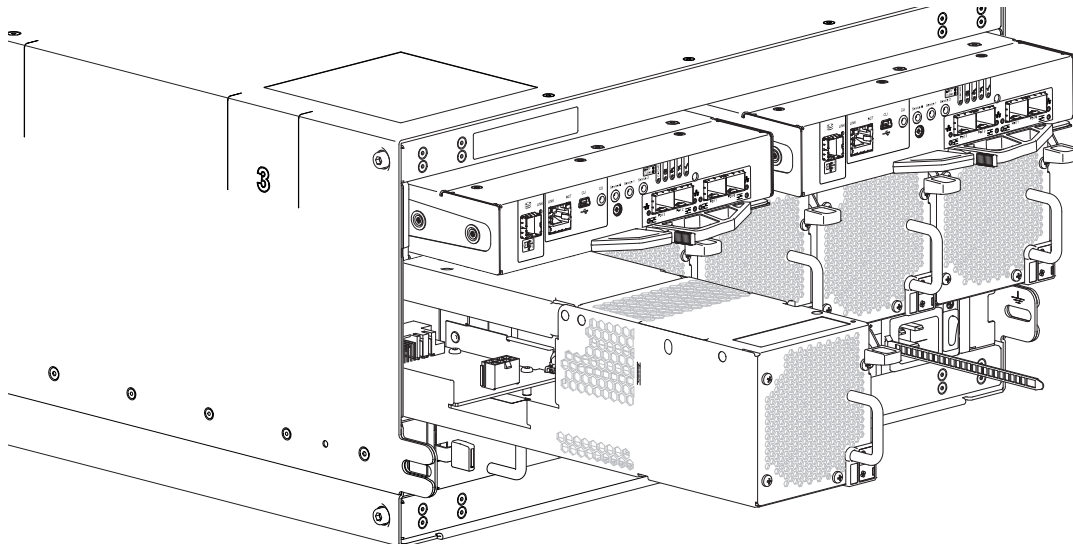
Figure 110 Removing an FCM (2 of 2)

---

① **IMPORTANT:**   The FCM slot must not be empty for more than 2 minutes while the enclosure is powered.

Installing an FCM

1.   You can hotswap the replacement of a single FCM; however, if replacing multiple FCMs, the enclosure must be powered off via an orderly shutdown using the management interfaces.

2.   Orient the FCM for insertion into the target slot on the enclosure rear panel, as shown in Figure 110.

3.   Slide the FCM into the slot until the latch clicks home.

     The enclosure should automatically detect and make use of the new module.

4.   Wait for the Module OK LED on the newly inserted FCM to illuminate green. See also Figure 32 (page 39).

     o   If the Module OK LED does not illuminate, verify that the FCM is properly inserted and seated in the slot.

     o   If properly seated, the module may be defective. Check the SMC and the event logs for more information.

     o   Using the management interfaces (the SMC or CLI), determine if the health of the new FCM is OK. Verify that the Module OK LED is green, and that the Ops panel states show no amber module faults.

5.   If replacing multiple FCMs, repeat steps 1 through 4.

Replacing an IOM

5005/4005/3005 Series supports a common set of controller modules in RBODs, and a common expansion module in EBODs and JBODS, within supported 2U and 5U enclosures as shown in Figure 85 (page 115). Be mindful of the intricacies associated with enclosure IOMs before engaging in replacement. Key considerations pertaining to IOM replacement in 2U enclosures apply equally to 5U enclosures. Please familiarize yourself with these topics before replacing an IOM in the 5U84 enclosure:

- **IMPORTANT** and ensuing discussion on page 125

- "Before you begin" (page 126)

- "Configuring partner firmware update for RBOD/EBOD configurations" (page 126)

- "Verifying component failure" (page 127)

- "Stopping I/O for RBOD/EBOD configurations" (page 127)

- "Shutting down a controller module" (page 128)

Removing an IOM

ⓘ **IMPORTANT:**   Considerations for removing controller modules:

- In a dual-controller environment, you may hot-swap a single controller module in an operational enclosure, provided you first shut down the faulty controller using the SMC or the CLI.
- In a dual-controller environment—if replacing both controller modules—you must adhere to the instructions provided in "Before you begin" (page 126).
- Do not remove a faulty module unless its replacement is on-hand. All modules must be in place when the system is in operation.

See CAUTION bullets regarding *electrostatic discharge* and *anti-static protection* on page 116.

Illustrations in the controller module replacement procedures show rear panel views of the enclosure, and IOMs are properly aligned for insertion into the IOM slots.

1. Verify that you have successfully shut down the controller module using the SMC or the CLI.
2. Locate the enclosure whose UID LED (Ops panel on enclosure left front ear) is illuminated, and within the enclosure, locate the controller module whose OK to Remove LED is blue (rear panel).
3. Disconnect any cables connected to the IOM.

   Label each cable to facilitate re-connection to the replacement IOM.
4. Grasp the module latch between the thumb and forefinger, and squeeze the flange and handle together to release the latch handle from its docking member (detail No.1), and swing the latch out to release the IOM from its seated position (detail No.2) as shown in Figure 111 and Figure 112.
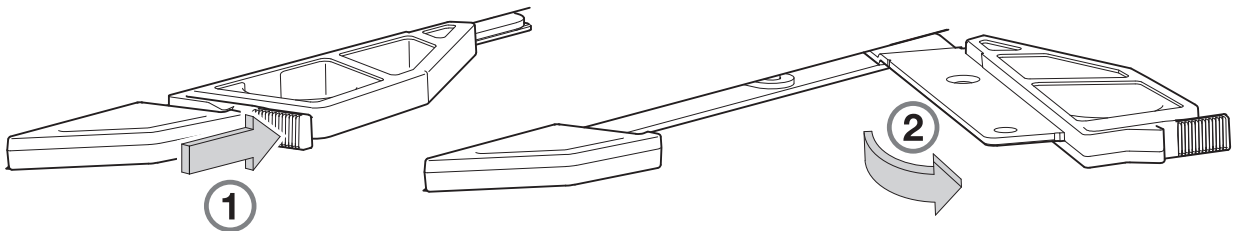
Figure 111 IOM latch operation detail

5. Grip the latch handle and ease the IOM outward from its installed position as shown in detail No.3 in Figure 112.
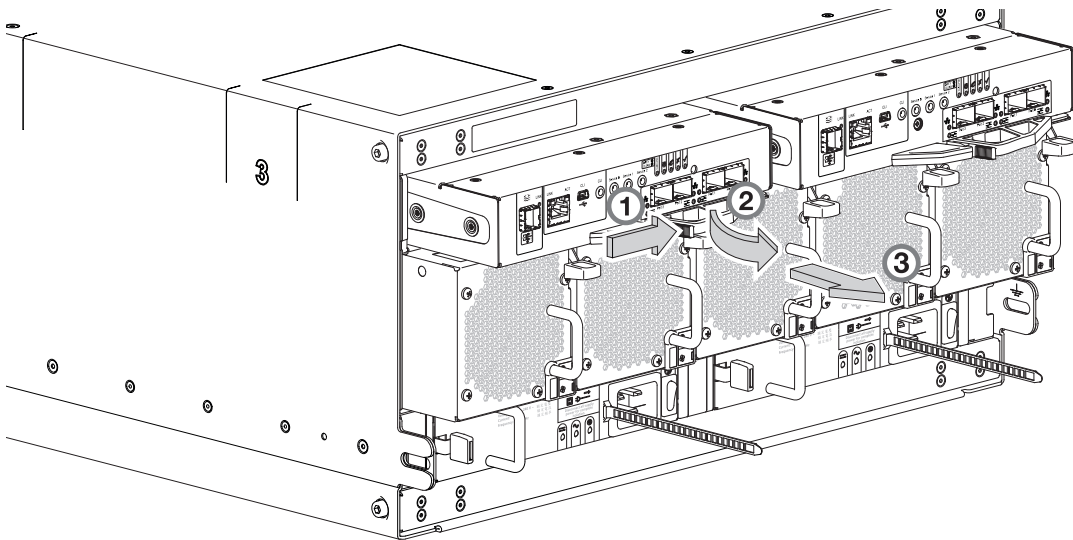
Figure 112 Removing an IOM (1 of 2)

NOTE: The illustrations above and below show a 4-port controller module of type FC/iSCSI. However, the procedure applies to all IOMs used in 5U84 enclosures. They all use the same latch mechanism, but feature different face plate geometry as shown in Figure 85 (page 115).
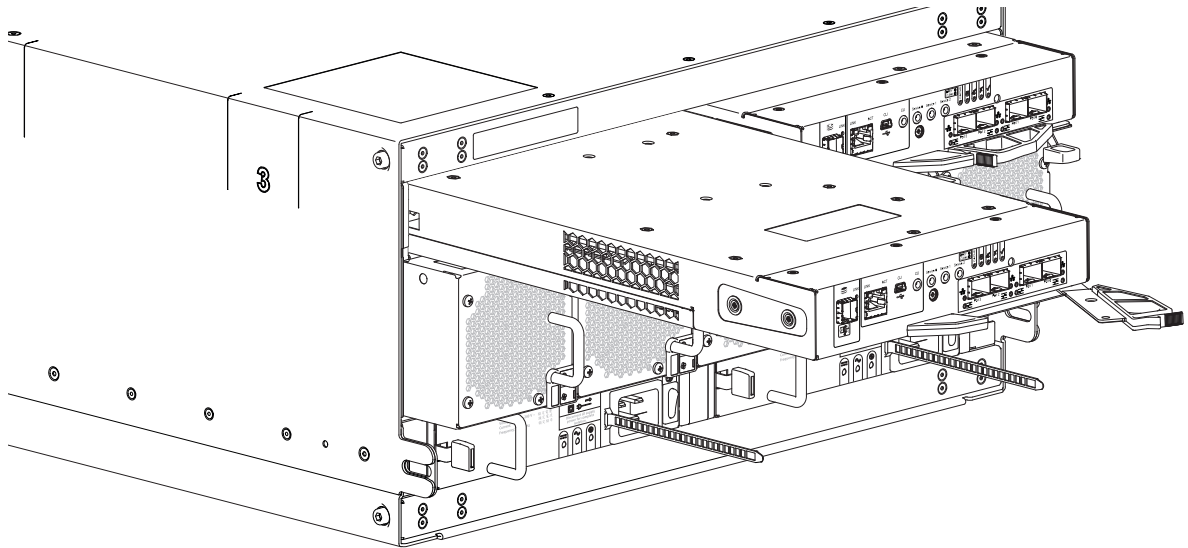
Figure 113 Removing an IOM (2 of 2)

6. Place both hands on the canister body, and pull it straight out of the enclosure.

(!) **IMPORTANT:**   The SBB IOM slot must not be empty for more than two minutes while the enclosure is powered.

Installing an IOM

See CAUTION bullets regarding *electrostatic discharge* and *anti-static protection* on page 116.

△ **CAUTION:**   If passive copper cables are connected, the cable must not have a connection to a common ground/earth point.

NOTE: When performing the following procedure, refer to Figure 113.

---

1. Examine the IOM for damage, and closely inspect the interface connector. Do not install if the pins are bent.
2. Grasp the IOM using both hands, and with the latch in the open position, orient the module and align it for insertion into the target IOM slot.
3. Ensuring that the IOM is level, slide it into the enclosure as far as it will go.

   A controller module that is only partially seated will prevent optimal performance of the controller enclosure. Verify that the controller module is fully seated before continuing.
4. Set the module in position by manually closing the latch. See detail No.2, followed by No.1 in Figure 111 (page 144). You should hear a click as the latch handle engages, securing the IOM to its connector on the back of the midplane.
5. Reconnect the cables.
6. See "Verifying component operation" (page 131).

NOTE: In a dual-controller system in which PFU is enabled, when you update the firmware on one controller, the system automatically updates the partner controller.

Replacing a 5U storage enclosure chassis

The controller enclosure or expansion enclosure chassis replacement procedure replaces a damaged 5U chassis CRU. The procedure includes removing all CRU modules from a damaged chassis and installing them into a replacement chassis.

ⓘ **IMPORTANT:**   The 5005/4005 Series enclosures using 5U84 chassis are described in Table 30 (page 111) and Table 33 (page 114).

Whether your controller enclosure product is configured as an RBOD, an EBOD, or a JBOD, a fully functional replacement chassis requires the successful removal and installation of the following components:

- All power and data cables
- All DDICs from enclosure drawers
- Two power supply unit modules (both AC)
- Five fan cooling modules
- Two IOMs of the same model type

This procedure references the CRU component procedures described elsewhere in this chapter.

Before you begin

△ **CAUTION:**   Do not remove the enclosure until you have received the replacement enclosure.

See CAUTION bullets regarding *electrostatic discharge* and *anti-static protection* on page 116.

1. Schedule down time that will allow for shutdown; sixty minutes of replacement work; and restart.
2. Verify the existence of a known/good backup of the system.
3. Record system settings for future use and label all cables.
4. Prepare a suitable static-protected work environment to accommodate chassis replacement.

Verifying component failure

The controller enclosure CRU includes the enclosure's metal housing, the module runner system, the integrated Ops panel, enclosure drawers for holding DDICs, and the assembled/installed midplane, sideplane, and baseplane PCBs that connect various components within the enclosure. See Figure 5 (page 23) for a pictorial view of the 5U84.

The 5U84 chassis CRU replaces an enclosure that has been damaged, or whose midplane has been damaged. Often, a damaged midplane will appear as though a RAID canister has failed. If you replace a controller module and it does not remedy the fault, you may need to replace the enclosure chassis.

You can observe enclosure health (front panel and rear panel) using management interfaces to verify enclosure/component failure or enclosure/component operation. See also "Using RBOD/EBOD management interfaces" (page 135).

Preparing to remove a damaged 5U storage enclosure chassis

Because you are removing and replacing an entire storage enclosure chassis, neither the hot-swap capability that applies to replacing individual redundant CRUs in an operational storage enclosure, nor the hot-add of a drive enclosure to an operational storage system, apply to this procedure.

1. Stop all I/O from hosts to the system. See also "Stopping I/O for RBOD/EBOD configurations" (page 127).
2. Shut down the controllers. See also "Shutting down a controller module" (page 128).
3. Power off the system – controller enclosure *first*; expansion enclosures *next*: see "Powering on/powering off" (page 88).

Table 35   Replacing a 5U storage enclosure chassis and installing its CRUs

To accomplish this sequential process		See the following procedures
1.	Facing the front of the 5U enclosure, access the drawers and remove DDICs from the damaged chassis. Place them in a static-protected work area nearby.[1]	a. "Opening a drawer" (page 136)
		b. "Removing a DDIC" (page 137)
		c. "Closing a drawer" (page 137)
2.	Facing the rear of the 5U enclosure, remove the power and data cables.	"Removing a damaged 5U storage enclosure chassis from the rack" (page 148) (step 1)
3.	Remove the damaged storage enclosure chassis from the rack.[2]	"Removing a damaged 5U storage enclosure chassis from the rack" (page 148)
4.	Remove the two power supply unit (PSU) modules from the damaged chassis rear panel, and install them in the replacement chassis.	a. See PSU-related TIP on page 140
		b. "Removing a PSU" (page 140)
		c. "Installing a PSU" (page 142)
5.	Remove the five fan cooling modules (FCMs) from the damaged chassis rear panel, and install them in the replacement chassis.	a. "Removing an FCM" (page 142)
		b. "Installing an FCM" (page 143)
6.	Remove the two IOMs from the damaged chassis rear panel, and install them in the replacement chassis.[3]	a. "Replacing an IOM" (page 143) (preparation subsections)
		b. "Removing an IOM" (page 144)
		c. "Installing an IOM" (page 145)
		d. "Verifying component operation" (page 131)
7.	Install the replacement storage enclosure in the rack.	"Installing the replacement 5U storage enclosure in the rack" (page 148)
8.	Facing the front of the 5U enclosure, access the drawers and install DDICs retrieved from the damaged enclosure into the replacement enclosure.[4]	a. "FDE considerations" (page 60)
		b. "Accessing drawers" (page 136)
		c. "Populating drawers" (page 139)
		d. "Installing a DDIC" (page 139)
9.	Complete the chassis replacement process	a. "Completing the process" (page 134)
		b. "Verifying component operation" (page 134)

1-If you temporarily stack the disks before installing them, insert static dissipative foam between the drive modules.

2-Fully configured 5U84 enclosures can weigh up to 128 kg (282 lb) and require a suitable mechanical lift.

3-Within the replacement enclosure, each IOM must be reinstalled into the same IOM slot from which it was extracted from the damaged enclosure (0A > 0A and 0B > 0B, respectively). For FC/iSCSI IOMs, the SFPs installed in the CNC ports need not be removed. For SFP replacement, see "SFP option for CNC ports" (page 178).

4-Within the replacement enclosure, reinstall each DDIC into the same disk slot and drawer from which it was removed from the damaged enclosure. To ensure optimal cooling throughout the enclosure follow drawer population guidelines.

Removing a damaged 5U storage enclosure chassis from the rack

This section provides a procedure for removing a damaged 5U storage enclosure chassis from its rack location.

⚠ **CAUTION:**   It is recommended that all DDICs be removed from the drawers before removing the enclosure from the rack. A suitable mechanical lift is required to move a high-density 5U84 enclosure.

1. Disconnect the power cables and data cables between devices as needed:
   a. Between the cascaded enclosures.
   b. Between the controller and peripheral SAN devices.
   c. Between the controller and the host.

   NOTE: Label the cables to facilitate reconnection of the storage devices.

2. Position a suitable mechanical lift to receive the enclosure.
3. Remove the retaining screws that secure the front and rear of the chassis to the rack and rails.

   NOTE: Do not remove the ear components from the failed chassis: integrated ear components and covers are provided with the replacement chassis CRU.

4. Maintaining a level position, carefully slide the enclosure chassis from the rack.
5. Place the chassis on a static-protected work surface near the replacement enclosure chassis, with the removed DDICs and rackmount screws.

Installing the replacement 5U storage enclosure in the rack

This section provides a procedure for installing the replacement enclosure in its rack location.

�घ **IMPORTANT:**   Refer to "Installing the 5U enclosure" (page 59) when installing the enclosure into the rack.

ⓘ **IMPORTANT:**   Install the enclosure into the rack before re-inserting the DDICs. A suitable mechanical lift is required to move the enclosure.

1. Support the bottom of the enclosure chassis. Carefully lift/align the chassis and while maintaining a level position for the enclosure, slide it into the rack.
2. Using the appropriate mounting hardware, secure the enclosure chassis to the rack.
3. Using the applicable retaining screws, secure the front and rear of the enclosure chassis to the rack and rails.

A   Connecting a JBOD

This appendix provides information about connecting JBODs once they have been secured into the rack. When used as a JBOD, the enclosure assumes the RAID functionality—if any—resides on the host server to which it is attached.

Storage system cabling

Collectively, the data path comprised of the server-based HBA(s) (RAID adapter/software), the expansion enclosures /firmware, and the cables connecting these enclosures to the host server(s) is referred to as the storage system.

The firmware supports both straight-through and reverse SAS cabling. Reverse cabling allows any enclosure to fail—or be removed—while maintaining access to the other enclosures. Straight-through cabling does not provide the same level of fault-tolerance as reverse cabling. Fault tolerance and performance requirements determine whether to optimize the configuration for high availability (reverse) or high performance (straight-through) when cabling.

Maximum configuration for a server-attached JBOD storage system is contingent on the total number of devices (HBAs or RAID adapters, disks, expanders) supported by the host. Consult your vendors' server and HBA documentation for configuration limits. If HBAs selected by the customer do not support enough devices, it may not be possible to achieve the system sizes shown in the example illustrations.

ⓘ **IMPORTANT:**   Cabling rules and system configuration limits

- Host server must be configured with compatible HBAs and any applicable RAID controls. See also "Host system requirements" (page 65).
- Only 12Gb expansion enclosures (J1212/J1224/J1284) are supported.
- When connecting enclosures, avoid creating invalid loops between SAS ports (A/B/C):
  - Figure 114 (page 150) shows an IOM incorrectly looped to itself
  - Figure 115 (page 150) shows an IOM incorrectly looped to its partner
  - Figure 116 (page 150) shows an IOM incorrectly looped to another JBOD instead of the server HBA
  - Figure 117 (page 151) shows an IOM incorrectly looped to another JBOD instead of the server HBA
- In a single SAS chain, the firmware versions used for IOMs and enclosure CRUs must be compatible.
- Intermixing of 2U12, 2U24, and 5U84 enclosures for use as SAS-chained JBODs is permitted.
  Conversely, RBOD/EBOD configurations do not support intermixing of 2U/5U expansion enclosures.
- Use reverse cabling to optimize the storage system for high availability/redundancy.
- Valid example configurations by enclosure type are shown on the following pages:
  - 2U enclosure cascade: supports up to 8 enclosures or 192 disks
  - 5U84 enclosure cascade: supports up to 4 enclosures or 336 disks
  - Mixed 2U/5U84 cascade: supports two 2U12s and two 5U84s; one 2U24 and two 5U84s; or a maximum of 192 disks

---

When adding storage, use only Seagate or OEM-qualified cables, and observe the following guidelines:

- When installing SAS cables to expansion modules, use only supported SFF-8644 cables.
- Qualified HD mini-SAS to HD mini-SAS 0.5 m (1.64') cables are used to connect cascaded enclosures in the rack.
- The maximum expansion cable length allowed in any server-attached JBOD configuration is 4 m (13.12').
- When adding more than two expansion enclosures, you may need to purchase additional cables, depending upon number of enclosures and cabling method used.
- You may need to order additional or longer cables when reverse-cabling a fault-tolerant configuration. Both passive and active optical cables are supported. Optical cables span a greater distance than passive cables.

See the rear panel views of the 2U enclosures Figure 12 (page 26) and 5U84 enclosure Figure 24 (page 32) for the location and orientation of the IOMs and other replaceable CRU modules.
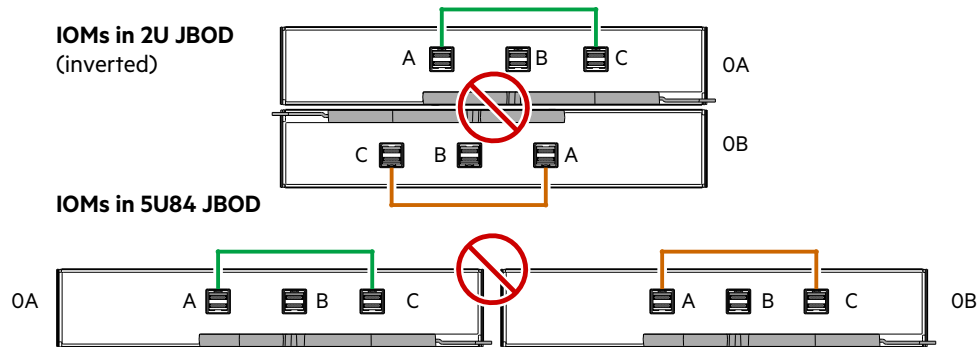
Figure 114 SAJ cabling: invalid configuration–IOM looped to self
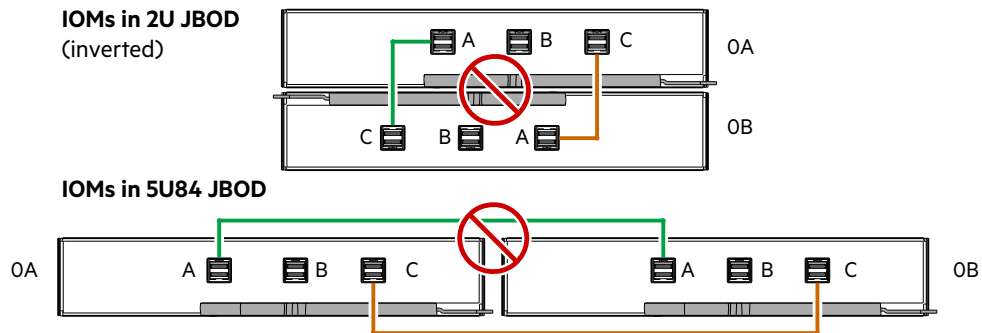
Figure 115 SAJ cabling: invalid configuration–IOM looped to partner
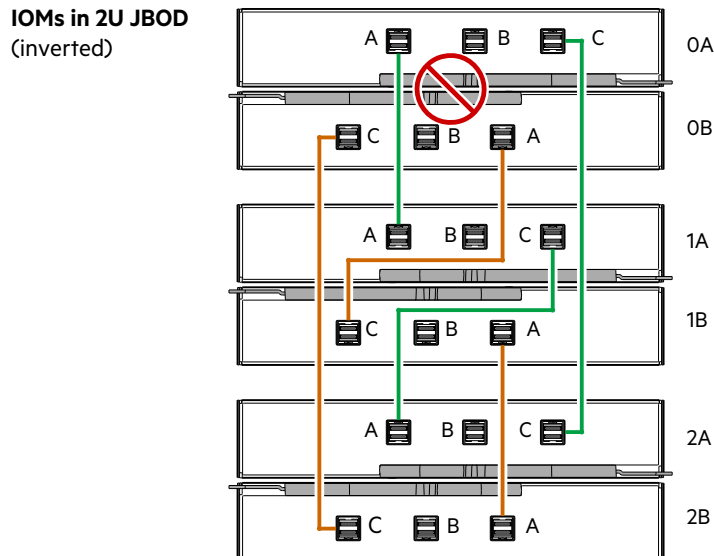
Figure 116 SAJ cabling: invalid configuration–IOM looped to JBOD instead of server

IOMs in 5U84 JBOD

Figure 117 SAJ cabling: invalid configuration–IOM looped to JBOD instead of server
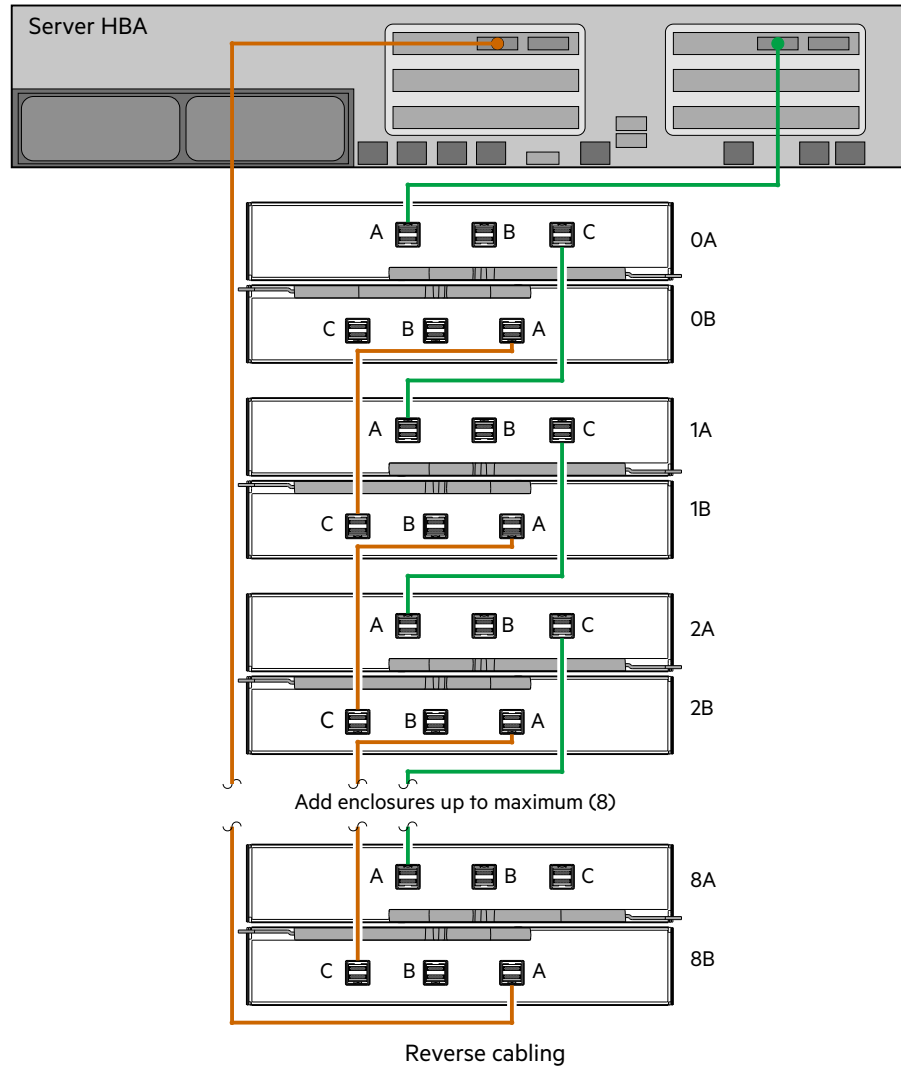
Figure 118 SAJ cabling: single host/multiple 2U enclosures

The diagram above shows cabling of a single HBA host server and supported 2U JBODs equipped with dual expansion modules, using the reverse cabling method.

The HBA connection on the right is cabled to expansion port 1A, with a chain of SAS connections cascading down (green). The HBA connection on the left is cabled to the lower expansion module (8B) of the last expansion enclosure, with connections moving in the opposite direction (tan).
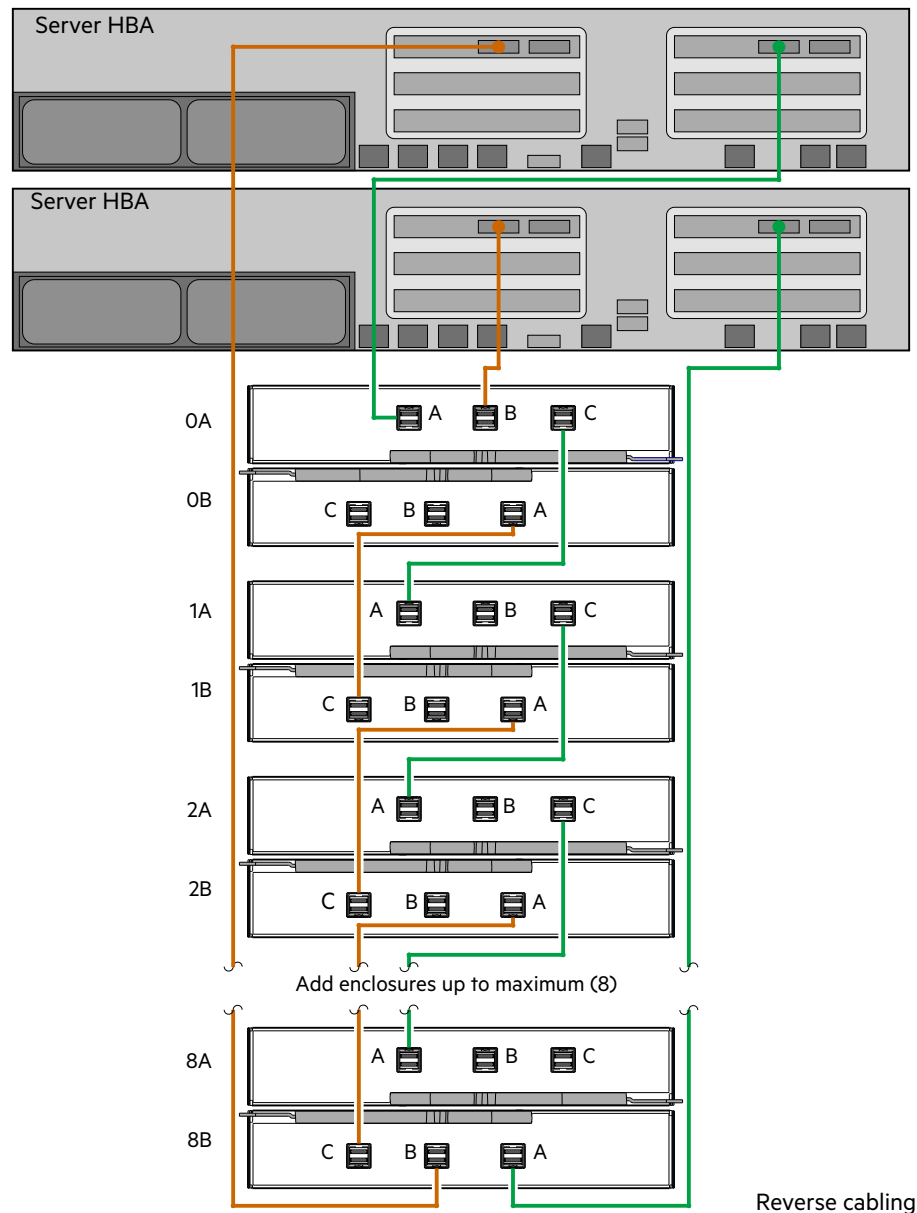
Figure 119 SAJ cabling: dual host/multiple 2U enclosures

The diagram above is similar to the previous diagram, although it uses dual hosts, together with two server connections from IOM 1A and IOM 8B, respectively.

Figure 120 (page 153) provides two additional examples of cabling 2U expansion enclosures. These examples use the middle SAS port. The diagram on the left shows reverse cabling. The diagram on the right shows the same storage components connected using straight-through cabling. Using this method, if an expansion enclosure fails, the enclosures that follow the failed controller in the SAS chain are no longer accessible until the failed enclosure is repaired or replaced.

Figure 121 (page 153) provides an example of cabling 5U84 expansion enclosures to a single host using the reverse cabling method.

Figure 122 (page 154) provides an example of cabling 5U84 expansion enclosures to dual hosts using the reverse cabling method.

Figure 123 (page 154) provides an example of cabling intermixed 2U and 5U84 enclosures using the reverse cabling method.
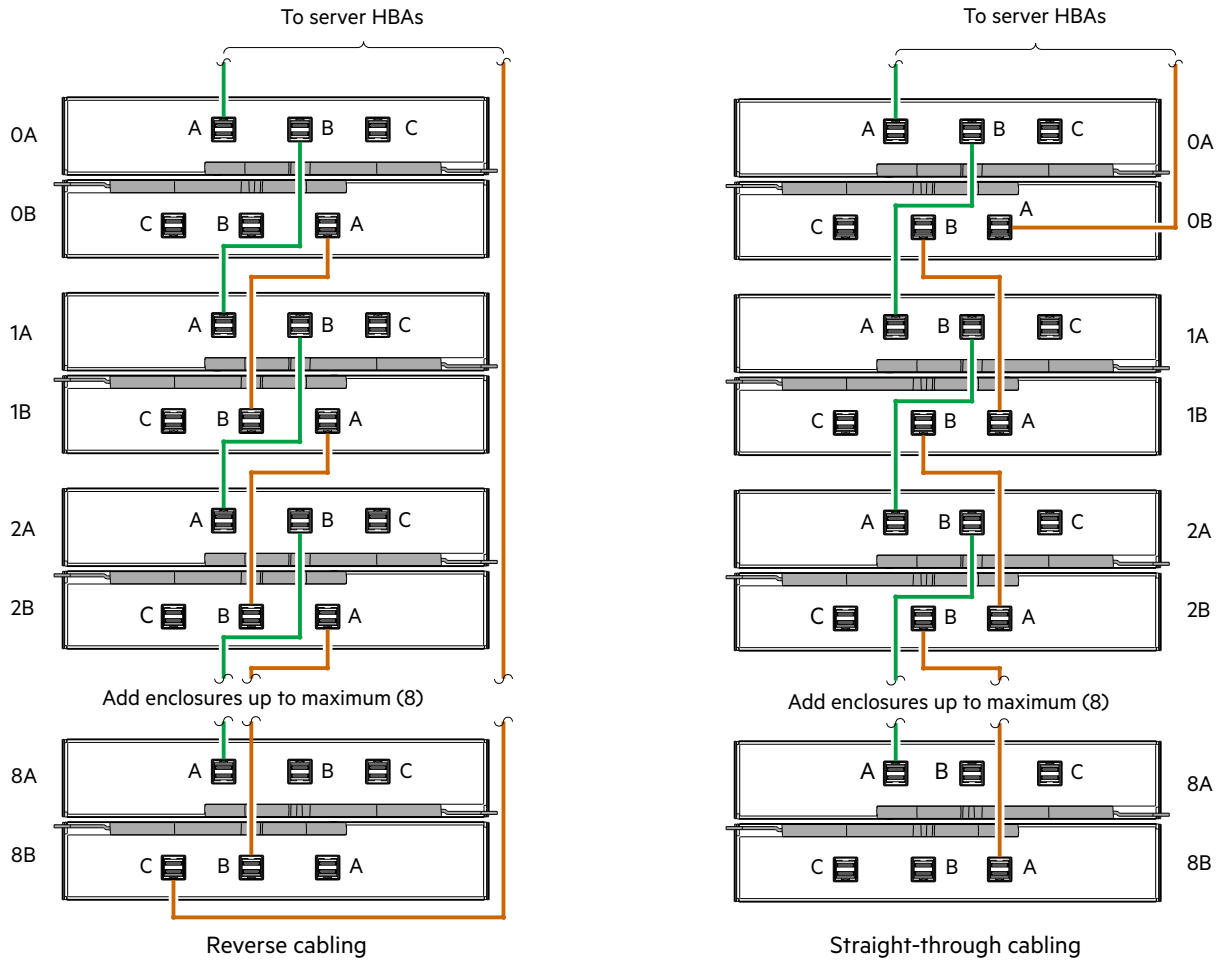
Figure 120 SAJ cabling: dual host/multiple 2U enclosures
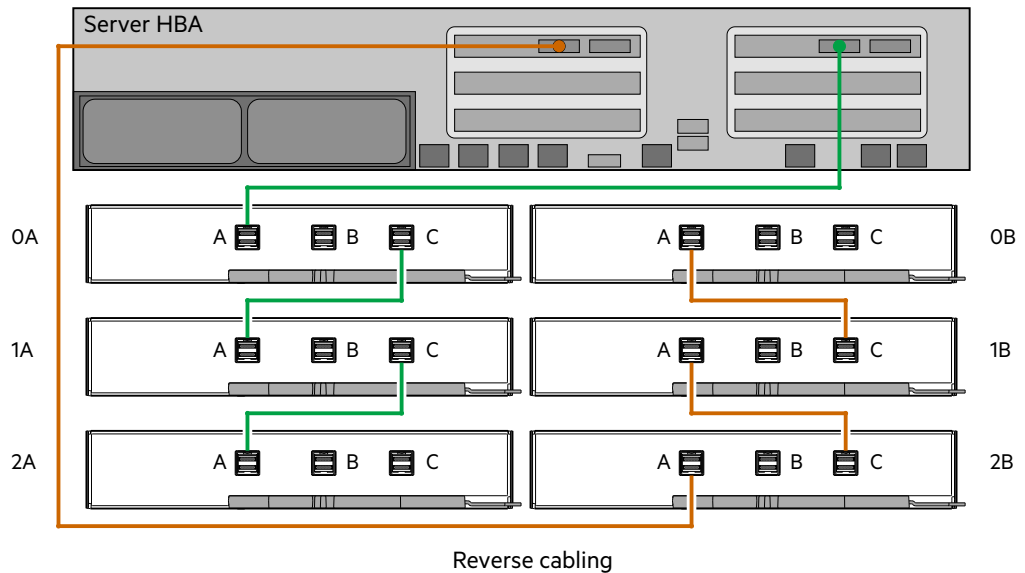
Figure 121 SAJ cabling: single host/multiple 5U84 enclosures
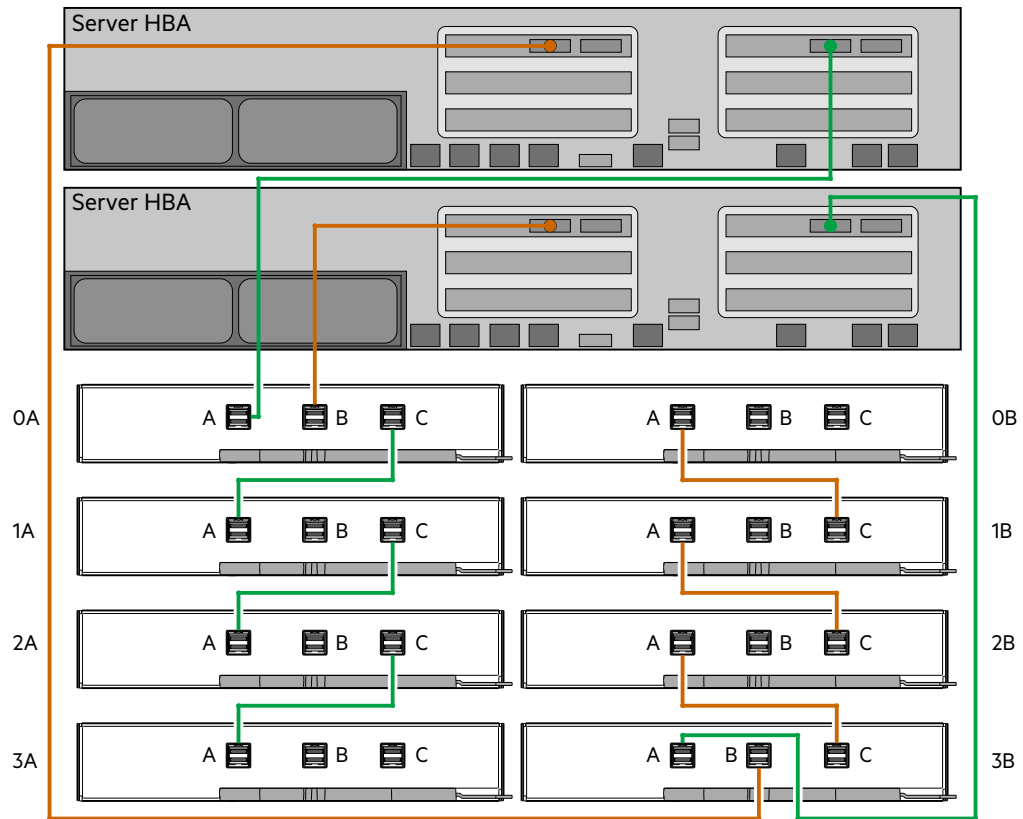
Figure 122 SAJ cabling: dual host/multiple 5U84 enclosures
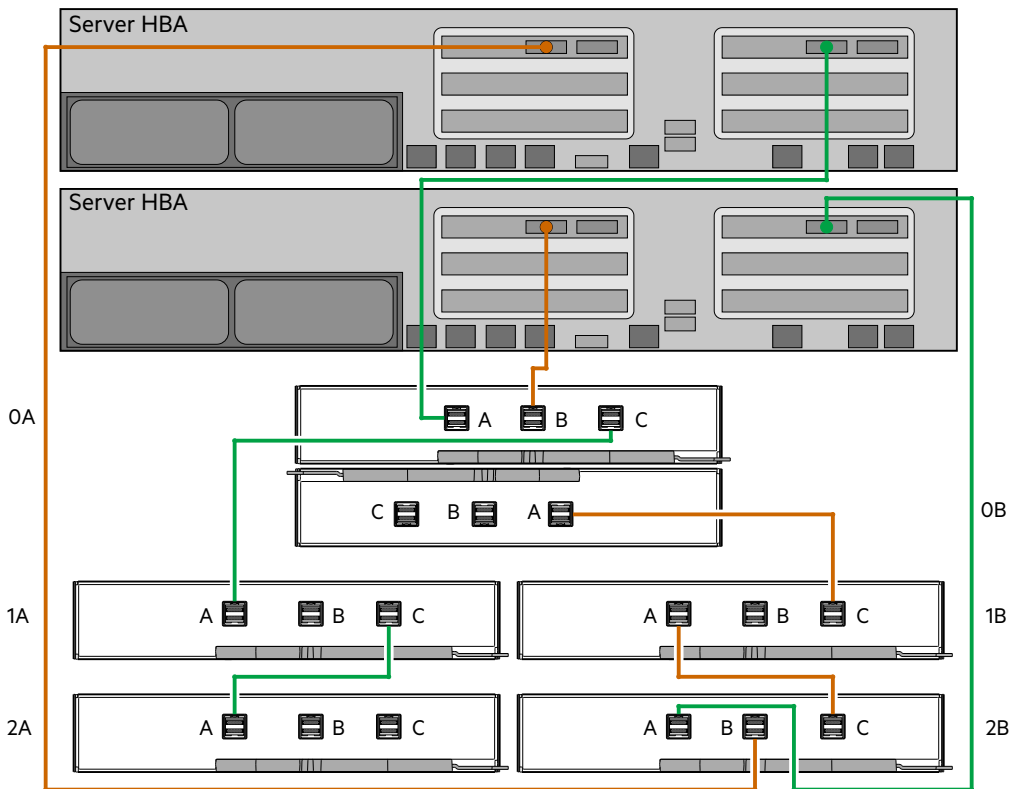
Figure 123 SAJ cabling: dual host/intermixed 2U and 5U84 enclosures

Fault isolation methodology

This section presents the basic methodology used to locate faults within a server-attached JBOD storage system, and to identify the pertinent CRUs affected.

Basic steps

- Gather fault information, including using system LEDs.
- Determine where in the system the fault is occurring.
- Review logs from the host application/operating system.
- Review `ddump` output from *both* controllers:
  - See "ddump" (page 163) for a description of the `ddump` GEM CLI command.
  - `Ddumps` can be captured via the following interfaces: serial, Telnet, WCLI/SES.
  - Invoke `help ddump` for additional information about capturing `ddump` output.

  See also "GEM CLI General Commands" (page 162) for related subset of the command table.
- If required, isolate the fault to a data path component or configuration as described in "Isolate the fault" below.

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 unit?*

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

Determine where the fault is occurring

When a fault occurs, the Module Fault LED—located on the left ear of the enclosure front panel (Ops panel)—illuminates. See also Figure 28 (page 35) for 2U enclosures, and Figure 29 (page 36) for 5U84 enclosures. Check the status of the other front panel LEDs. Also check the LEDs on the drawer bezel (5U84 only) and the rear panel of the enclosure to narrow the fault to a CRU, connection, or both.

- A system crash or unexpected reset of GEM should include `statesave` information.

  A firmware fault results in the automatic capture of state and trace data known as `statesave` information. The existence of this information is flagged in the corresponding Enclosure Electronics Diagnostics Element in SES. It is important to collect this information as soon as a `statesave` capture event is detected, to improve the opportunity for effective fault diagnosis.
- If a `statesave` capture signal is detected, the application should collect the diagnostic information by running the `ddump` CLI command on both IOMs, capturing the command output to a text file.

The LEDs help you to identify the location of a CRU reporting a fault. Each CRU reports failure if it detects a fault. It does so via LED and SES notification. See "Module removal and replacement" (page 116) for information about removing a faulty CRU and installing a replacement.

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 involved. For example, if a host-side data error occurs, it could be caused by a cable, connection, or faulty RAID adapter. If an expansion-side data error occurs, it could be caused by a faulty IOM, cable, connection, or other failed CRU or chassis component, to include a midplane or other PCBA.

If the enclosure is not recognized or fails to initialize

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

- Verify that each IOM are securely seated within its slot.
- Verify that the interface cables from the enclosure to the host computer are fitted correctly.
- Check initialization of the disk drive modules:
  - o For 2U enclosures, verify the Disk Activity LED on each installed module is green, and the Disk Fault LED is off.
  - o For the 5U84, verify no illuminated amber states for Module Fault or Drawer Fault LEDs on the Ops panel or drawers. Also, when viewing the DDICs installed in disk slots within each drawer, verify that none of the DDICs show an illuminated amber state for the Drive Fault LED located on the top of the module.
- Check that the disk drive modules have been correctly installed.
- Check any visible SAS indicators (JBOD and server adapters).

  These are the host port activity and fault indicators.
- Check server adapter setup utilities for SAS target visibility.
- Verify that the operating system driver supporting the HBA/RAID adapter on the host has been installed correctly.

Obtaining the JBOD firmware download

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. Periodically, you should ensure that the firmware versions used in enclosure modules are compatible.

1. Verify that the management computer has Internet access.
2. For customer support, see seagate.com/support
   a. Select **Products > Systems > General Support**
   b. From **Downloads**, select the firmware update that applies to your JBOD model and operating system.

   The firmware download bundle should appear in your Downloads directory.
3. Follow instructions provided in the firmware download bundle.

   See also Firmware installation instructions (below).

Obtaining JBOD IP values

You can manually set static IP values for each IOM, or you can specify that IP values should be set automatically for both IOMs through communication with a DHCP server. 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. You must have some means of determining what addresses have been assigned, such as the list of bindings on the DHCP server.

NOTE: DHCP-assigned IP addresses can be queried from each IOM using the `ipconfig` CLI command. See also "ipconfig" (page 165).

JBOD firmware installation instructions

The downloaded firmware bundle includes the Seagate JBOD Firmware Update Utility known as `jbodupdate`. It is a cross-platform tool combining the enclosure firmware content and the logic required to detect and update any attached enclosures.

NOTE: An enclosure model/firmware release (.txt) file in the GEM firmware bundle provides release note information and overview installation instructions for Windows and Linux. Given that `jbodupdate` is not available within VMware, the text file provides guidance for scheduled system shutdown prior to doing a firmware update on these systems.

JBOD firmware update procedure

The `jbodupdate` utility uses SES to upgrade the enclosure. As such, the enclosure must be connected to a host server containing an HBA. `Jbodupdate` starts by first discovering the SES targets available in the SCSI topology detected by the server. From the visible SES targets, an enclosure map is constructed, allowing each SES device to be mapped to a data path within the enclosure. Multiple paths to the same physical device are recorded so they can be used to access the enclosure in the event of a path loss during the firmware update.

After discovery, the enclosure components are queried to determine the steps required to update them to the firmware levels contained in the update package. Provided an update operation is specified in the command line arguments for `jbodupdate`, the tool will commence with the update on a per enclosure basis, with each enclosure being updated in the order discovered. The tool will typically update each of the enclosure CRUs in the following order:

1. Update path A IOM
2. Activate path A IOM
3. Update path A sideplanes (high-density enclosures only)
4. Activate path a sideplanes (high-density enclosures only)
5. Update path B IOM
6. Activate path B IOM
7. Update path B sideplanes (high-density only)
8. Activate path B sideplanes (high-density only)
9. Update power supplies and fan modules
10. Update midplane
11. Activate midplane

NOTE: The above sequence of steps may be subject to change on a per release basis if the specific USM (Unified System Management) requires an alternative procedure.

To minimize update duration, `jbodupdate` compares the component versions running on the enclosure against the versions of the firmware in the bundle. If the versions match, the update is skipped. `Jbodupdate` also tracks the activation action required for each component downloaded. This enables `jbodupdate` to determine the level of disruption that will occur to the data path when an activation is performed, and select the least disruptive option.

For high-availability (HA) enclosures, where disruptive path activations cannot be avoided, the enclosure staggers the activation of the separate data paths so that I/O may continue during the download.

The multipath logic within `jbodupdate` endeavors to pick the most appropriate SCSI pass-through device for each enclosure component. If this device is not available, it will pick the next most appropriate path. If direct SCSI access to an enclosure component is not available, `jbodupdate` will use the remote update facility within the enclosure to complete the update.

NOTE: provides a high-level flow diagram of the procedure performed by `jbodupdate` when a single enclosure is specified for firmware update using the `--enc` option.
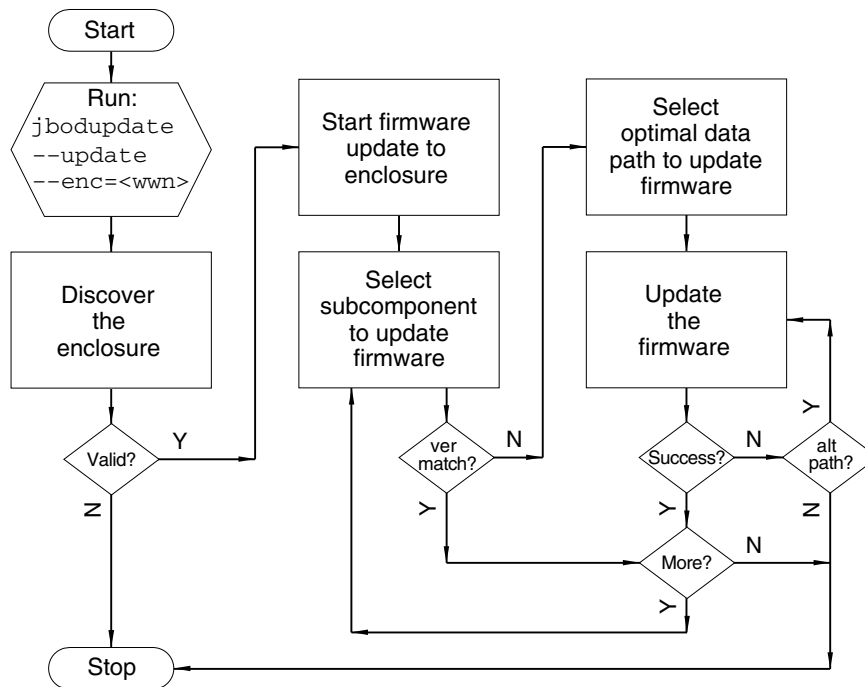
Figure 124 Update sequence for enclosure firmware

Command-line arguments

Usage: `jbodupdate [options]`

`options`: Options for specifying command behavior. Note that option arguments are only displayed against the long option below. Where a long option argument is specified, the same argument applies to the short option.

`-c` | `--continue`	Instructs the utility to continue with the update if an error occurs. If this option is not specified, the utility will exit on error.
`-d` | `--dev=<devid>`	Specifies a SES device on which the operations will be performed. The <devid> argument is the path for the SCSI pass-through device for the SES target. This option may not be used with the `--enc` option.
`-e` | `--enc=<encl>`	Specifies the enclosure on which the operation will be performed. The <encl> argument for this option takes the form of the device WWN or may be set to `all` to specify all discovered enclosures. This option may not be used together with the `--dev` option.
`-f` | `--force`	Forces the update of all enclosure components regardless of whether they are detected to be at the same level as the firmware bundle content. **NOTE**: Use this option *only* if instructed to do so by customer support.
`-p` | `--path=<path>`	Optional argument that when combined with `--update`, selects the data path that will be updated. Regardless of the path setting, common enclosure components will always be updated. Path may be set to:

`all`	The tool will update all data path components (default option)
`1`	Update data path A components
`2`	Update data path B components
`local`	Update local data path components (applicable with `--dev` only)
`remote`	Update partner data path components (applicable with `--dev` only)

`-s` | `--show`	Instructs the utility to display firmware version and inventory information for the devices specified by the `--enc` and `--dev` options. Specifying no `--dev` or `--enc` option is treated the same as specifying `--enc=all`. No update will be performed if this option is specified.
`-u` | `--update`	Instructs the utility to update the devices specified by the `--enc` or `--dev` options. If `--show` is not specified, this option is selected by default.
`-v` | `--ver`	Displays the bundle version and exits.

NOTE: If `--help`, `--show`, or `--ver` are not specified, the `--update` operation is applied by default. Either `--dev` or `--enc` are mandatory for enclosure update operations.

Example command use cases

Acquire an enclosure inventory list

The `jbodupdate` utility supports the ability to discover available Seagate JBODs and present an inventory list of all enclosure CRUs and their versions. This is achieved using the `--show` option as demonstrated below. The inventory output is structured in a YAML-like format, and can be machine parsed if necessary.

An update summary is provided as part of the output from the `show` command option. The summary indicates if the enclosure firmware is currently down-level, and whether the update would result in any data path disruption.

Command: **jbodupdate --show**

Auto-update all enclosures

The `jbodupdate` utility supports the ability to discover available Seagate JBODs and automatically apply the necessary steps to update them to the firmware levels contained in the bundle. The automated update procedures endeavor to limit disruption to a single data path where possible.

Command: **jbodupdate --update --enc=all**

Query a specific enclosure

The `jbodupdate` utility supports the ability to restrict the output of the `--show` option to include information from a specific enclosure. This is accomplished using the `--enc` option together with the enclosure's WWN.

Command: **jbodupdate --show --enc=**<wwn>
Command: **jbodupdate --show --enc=500c0ff012345678** (example WWN)

Query a specific SES target

The `jbodupdate` utility supports the ability to restrict the output of the `--show` option to include information from a specific SES device ID. This is accomplished using the SES device's SCSI pass-through device ID on Linux.

Command: **jbodupdate --show --dev=**<devid>
Command: **jbodupdate --show --dev=SCSI1:0,1,0** (Windows)
Command: **jbodupdate --show --dev=/dev/sg0** (Linux)

Auto-update a specific enclosure

The `jbodupdate` utility supports the ability to specify the enclosure to be updated using the `--enc` or `--dev` options. When this option is selected, the utility will update the specified enclosure, and then exit.

Command: **jbodupdate --update --enc=**<wwn>
Command: **jbodupdate --update --dev=**<devid>

Auto-update a specific data path

The `jbodupdate` utility supports the ability to specify which data path is affected by the firmware update procedure. This is useful for synchronizing I/O disruptive enclosure updates with a higher-level I/O stack. When a data path is specified, the update procedure will update components in that data path, as well as non-data path components such as power supplies and fan modules. Components within the alternate data path are not updated.

Command: **jbodupdate --update --dev=**<devid> **--path=**<path>

Overview of GEM

GEM (Generic Enclosure Management) is the firmware responsible for actively managing the expansion module and enclosure electronics. Additional software is required for other components in the storage system. GEM resides in flash memory.

In a dual-IOM enclosure, GEM runs on both expansion modules in an active-active redundant configuration. This setup allows for the collection of consistent enclosure status from either channel, as well as providing dual path control capabilities. Because there are two instances of GEM, enclosure operation is maintained in the event that one of the expansion modules fails.

Table 36 describes the enclosure management tasks performed by GEM.

Table 36   Enclosure management tasks performed by GEM

No.	Task	Description
1	Thermal management	Monitor and control temperatures and enclosure cooling.
2	Power management	Intelligent power and fault tracking to ensure that the enclosure power supplies are never oversubscribed. GEM allows you to power cycle individual IOMs as well as the enclosure on whole.
3	Drive management	Spin up drives using input from power management.
4	SAS expander management	Control of access to drives through SAS, and provision of status information about link quality and device mapping within the enclosure.
5	High availability	Dual-IOMs only: ensure the management environment for both IOMs is synchronized, and detecting/managing scenarios wherein an IOM has failed.
6	Diagnostics	Collection of fault information gathered from the enclosure and its display via the various fault indicators, as well as the electronic interfaces provided.

GEM CLI

The GEM firmware provides a command-line interface (GEM CLI) for interacting with the enclosure. The GEM CLI is not intended as a general means of controlling the system (like SES or IPMI), and is most often used for gathering diagnostic information in the event that a fault has occurred, or for setting TCP/IP options.

NOTE: The GEM CLI undergoes ongoing development, and its behavior is subject to change.

Table 2 describes document conventions. For GEM CLI commands listed in this appendix, the following conventions are used to describe command syntax.

Table 37   Symbols used in GEM command parameter syntax

Symbol	Name	Description
<>	Angle bracket	Used to delimit mandatory parameters. These terms must be substituted with a suitable argument. Available options are described for each command. Parameters without angle brackets are literal text. For example, "all" is sometimes used to specify all available options.
[]	Square bracket	Used to show optional parameters.
		Vertical bar	Used to show logical OR, meaning that one option or the other can be selected.

Example:

```
help [all]|[<cmd>]
```

This means the `help` command can be invoked with the `all` parameter or the name of a specific command.

```
help all
help ddump
```

Numeric parameters can be specified in decimal, hexadecimal, or octal.

- Decimal—use the numeric value for example: "14"
- Octal—use a leading zero, for example: "016"
- Hexadecimal—use a leading zero and x: for example "0xE"

String parameters must be enclosed in quotes if the parameter includes spaces or tabs, or if the case must be preserved. By default, strings are converted to lowercase.

To send a command to the partner IOM, precede the command with a minus sign ("-"). For example:

`ver` gives software version numbers for the current IOM; whereas `-ver` gives the corresponding information for the partner IOM.

A single command cannot exceed 120 characters in length.

Serial port settings

The following serial port settings are required when interfacing a terminal emulator client with GEM.

Baud rate (bits/second):	115,200
Data bits:	8
Parity	None
Stop bits:	1
Flow control:	None (see **Note**)

NOTE: The IOM serial port does not use flow control. Therefore, transmitted or received characters may be dropped during large data transfers.

- When collecting `ddump` output using the serial port, it is recommended that the output be inspected for signs of dropped characters before submitting it for further analysis.
- Telnet and the in-band CLI protocols are not susceptible to dropped characters.

Connecting to the GEM CLI

The TCP/IP settings must be configured appropriately before connecting an Ethernet cable to the enclosure. Ensure the network is configured to use DHCP, and verify that the DHCP server is assigned to the correct IP addresses.

> ⓘ **IMPORTANT:** Do not connect the Ethernet port to the public network. All IOM Ethernet ports should reside in the local area network.

Telnet CLI settings

The JBOD IOMs offer GEM CLI access using the Telnet over the standard TCP/IP port 23. The TCP/IP and Ethernet settings may be changed using the `ipconfig` command. See also "ipconfig" (page 165).

> **NOTE:** The default user name is **manage**. If the default user name has been changed for security reasons, enter the secure login credentials instead of the default.

GEM CLI status codes

The GEM CLI commands can be configured to complete with either a human readable or a hexadecimal machine readable status code. By default, GEM does not render a status code upon command completion. When command return codes are enabled, the status codes below will be used. The table provides both the human readable strings and corresponding machine readable codes.

No.	Status	Code
1	SUCCESS	0
2	INSUFFICIENT_RESOURCES	1
3	NOT_SUPPORTED	2
4	NOT_IMPLEMENTED	3
5	INVALID_PARAMETER	4
6	IO_TIMEOUT	5
7	CANCELLED	6
8	UNSUCCESSFUL	7
9	NO_MEMORY	8
10	DEVICE_REMOVED	9
11	REVISION_MISMATCH	a
12	INVALID_HANDLE	b
13	INVALID_DEVICE_REQUEST	c
14	INTERRUPTED	d
15	DEVICE_NOT_READY	e
16	FILE_NOT_FOUND	f
17	UNKNOWN	10

GEM CLI General Commands

The GEM CLI commands described in this appendix are intended for general use, and can be issued without any possibility of causing data loss or other serious problems. The commands that follow are organized alphabetically.

change_password

Description

Changes the password for the user currently logged in. The password is changed at the enclosure level, so all IOMs use the updated login credentials on completion.

The command is interactive, and the current password must be specified before a new password can be specified. The new password must be reentered.

Syntax

```
change_password
```

Example

```
# change_password
Enter Old Password:********
Enter New Password:********
Reenter New Password:********
password change success
```

cleardriveindicator

Command name	`cleardriveindicator`
Command synopsis	Clears the locate/identify LED.
Command description	Invoking this command sets the identify flag off, and stops the Fault LED blinking for the specified disk drive module.
Command arguments	`cleardriveindicator <drivebay>` `<drivebay>:` drive to flag for visual identification, indexed from 0
Command type	Control: Drive manager command
Access level	Testing

ddump

Command name	`ddump`
Command synopsis	Provides system-wide diagnostic dump for storage system.
Command description	Call all the commands of type diagnostic which do not demand an argument. Provides a single shot diagnostic dump.
Command arguments	`ddump [canister]` `[canister]:` provides diagnostic information for canister only **Note:** If no argument is provided, the command provides diagnostic information for all enclosure components.
Command type	Diagnostic: GEM CLI in-built command
Access level	General

ddumpo

Command name	`ddumpo`
Command synopsis	Provides ordered system-wide diagnostic dump.
Command description	Call all the commands of type diagnostic which do not demand an argument. Provides `ddump` output in a command type categorized order.

Command arguments	ddump [canister]
	[canister]: provides diagnostic information for canister only
	Note: If no argument is provided, the command provides diagnostic information for all items that are type diagnostic.
Command type	Diagnostic: Engineering command
Access level	Engineering

exit

Command name	exit
Command synopsis	Exits the Telnet session.
Command description	Ends the active Telnet session.
Command arguments	None
Command type	Control: CLI in-built command
Access level	General

getboardid

Command name	getboardid
Command synopsis	Get the PCB ID and mode.
Command description	Get the IOM slot ID and operating mode.
Command arguments	getboardid [hex]
	[hex]: returns slot ID (byte 1) and HA mode (byte 2) in hexadecimal form
	Note: If the canister is the master, the HA mode is set to 0x01, and if the canister is the slave, the HA mode is set to 0x00.
Command type	Debug: Canister manager command
Access level	General

getvpd

Command name	getvpd
Command synopsis	Get the VPD information for the enclosure.
Command description	Report the VPD information for the enclosure.
	Note: Show information for enclosure, canister, midplane, and power supply.
Command arguments	None
Command type	Debug: Canister manager command
Access level	General

help

Command name	help
Command synopsis	Display information about commands.
Command description	Provide mechanism to discover available commands and display command usage information.

Command arguments	`help [none|all|testing]`
	none (default): lists the synopsis for commands with "general" access level
	`[all]`: lists synopsis for all commands, regardless of access level
	`[testing]`: lists synopsis for commands with "testing" access level
	`[cmd name]`: lists detailed help for specified command (e.g., `help ddump`)
Command type	Status: CLI in-built command
Access level	General

hid_set_ident

Command name	`hid_set_ident`
Command synopsis	Set enclosure identify state.
Command description	Enable user to set/clear the identity state of the enclosure. This state is signified by flashing the Ops panel identify LED.
Command arguments	`hid_set_ident <0|1>`
	`[0]`: cancel the identify state
	`[1]`: enable the identify state
Command type	Debug: Human Interface Device (HID) command
Access level	Engineering

ipconfig

Command name	`ipconfig`
Command synopsis	Perform TCP/IP configuration and save in VPD.
Command description	Enable user to view/configure TCP/IP settings and save changes in enclosure VPD.
Command arguments	`ipconfig` [arguments are listed below]
	none: lists current active `ipconfig` settings
	`[dhcp <enable>|dhcp <disable>]`: enable or disable DHCP
	`[ip ip_address netmask gateway]`: set IP configuration
	`[dns dns_server_1 <dns_server_2>]`: set DNS server
	`[telnet <enable>|<disable>]`: enable or disable the Telnet server
	`[ether <enable>|<disable>]`: enable or disable Ethernet
Command type	Diagnostic: Hardware Abstraction Layer (HAL) command
Access level	General

logcmt

Command name	`logcmt`
Command synopsis	Insert a comment into a logdump.
Command description	Enable user to insert a custom entry into the enclosure diagnostic log.
Command arguments	`logcmt [<comment>]`
	<comment>: alpha-numeric string (maximum 72 characters)
Command type	Diagnostic: Logging command
Access level	General

logdump

Command name	`logdump`
Command synopsis	Display logged messages.
Command description	Enable user to access and view logged messages.
Command arguments	`logdump` [arguments are listed below]

none: Displays all logged messages from RAM; newest to oldest, with all fields enabled

`[r|n]`: retrieve log messages from RAM (`r`) or non-volatile (`n`) memory

`[old|new]`: retrieve log messages in ascending (`new`) or descending (`old`) order

`[last n|all]`: display last specified number (n) of messages or all (default if omitted)

`[enable_timestamp]`: controls timestamp field (`1`=enable; `0`=disable)

`[enable_subsystem_name]`: controls subsystem name field (`1`=enable; `0`=disable)

`[enable_service_name]`: controls service name field (`1`=enable; `0`=disable)

Command type	Diagnostic: Logging command
Access level	General

logdumpnv

Command name	`logdumpnv`
Command synopsis	Display non-volatile logged messages.
Command description	This command is an alias for `logdump n`.
Command arguments	`logdumpnv` [arguments are listed below]

`[old|new]`: retrieve log messages in ascending (`new`) or descending (`old`) order

`[last n|all]`: display last specified number (n) of messages or all (default if omitted)

`[enable_timestamp]`: controls timestamp field (`1`=enable; `0`=disable)

`[enable_subsystem_name]`: controls subsystem name field (`1`=enable; `0`=disable)

`[enable_service_name]`: controls service name field (`1`=enable; `0`=disable)

Note: Minimum number of arguments is 1 and maximum is 6 (e.g., `logdumpnv new 10`)

Command type	Diagnostic: Logging command
Access level	General

poweroffdrive

Command name	`poweroffdrive`
Command synopsis	Power off a drive, including spin-down.
Command description	Move disk into the "off" state form any other state.
Command arguments	`poweroffdrive` [arguments are listed below]

`[all]`: all enclosure drive slots (default if omitted)

`[n]`: drive slot to operate on, indexed from 0

Command type	Control: Drive manager command
Access level	Engineering

powerondrive

Command name	`powerondrive`
Command synopsis	Power on a drive, including spin-up.
Command description	Move disk into the "online" state form any other state.
Command arguments	`powerondrive` [arguments are listed below]
	`[all]`: all enclosure drive slots (default if omitted)
	`[n]`: drive slot to operate on, indexed from 0
Command type	Control: Drive manager command
Access level	Engineering

report_faults

Command name	`report_faults`
Command synopsis	Report all system-wide faults.
Command description	Output all known faults collected from each GEM service.
Command arguments	None
Command type	Diagnostic: CLI in-built command
Access level	General

reset_saved_zone_info

Command name	`reset_saved_zone_info`
Command synopsis	Reset the saved zoning information to system default.
Command description	Erase the saved zone partition in flash, then update the current, shadow, saved context in SDK, and saved zoning partition in flash to the default configuration.
Command arguments	None
Command type	Control: Flash update command
Access level	General

set_encl_id

Command name	`set_encl_id`
Command synopsis	Set the enclosure ID (if supported).
Command description	Allow the user to set the Ops panel enclosure ID (if supported).
Command arguments	None—the current ID will display. Use integer between 1 and 99 for ID.
Command type	Status: Human Interface Device (HID) command
Access level	Engineering

setcanisterindicator

Command name	`setcanisterindicator`
Command synopsis	Clear, then set the local canister (IOM) indicator statuses.
Command description	Set the local IOM identify indicator LED state. See also Figure 39 (page 47)
	Note: The blue identify LED will blink 1 second on, 1 second off if enabled.

Command arguments	`setcanisterindicator [0|1]`
	[0]: clear the canister Identify LED
	[1]: set the canister Identify LED to On
	Note: Invoking this command without an argument returns the current state of Identify LED.
Command type	Test: Canister manager service command
Access level	Testing

setdriveindicator

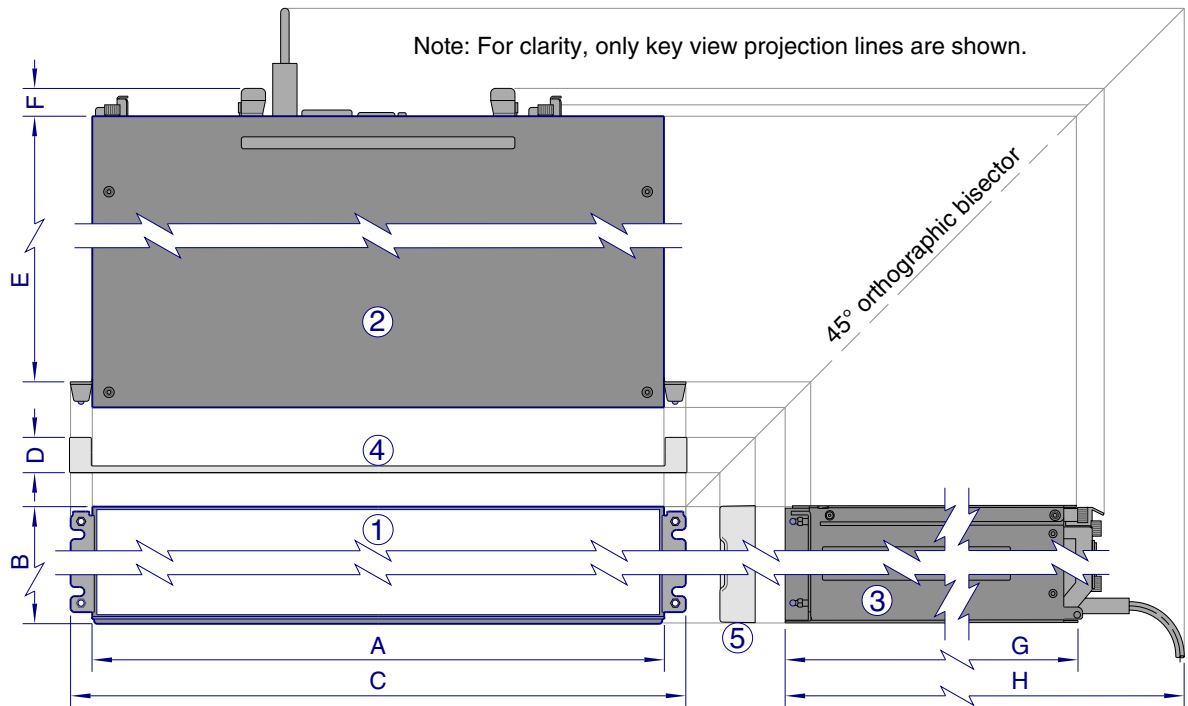|	
Command name	`setdriveindicator`
Command synopsis	Set the disk drive indicator LED.
Command description	Set the drive Fault LED to blink (provides easier identification of the drive).
Command arguments	`setdriveindicator <drivebay>`
	`[drivebay]`: Slot in which the disk resides (Integer)
Command type	Control: Canister manager command
Access level	Testing

settime

Command name	`settime`
Command synopsis	Set unit time.
Command description	Set the unit time using either Universal or Epoch measure.
	Universal (-u) time specified in days (days), hours (hh), minutes (mm), seconds (ss)
	Note: Setting Universal time is moot if Epoch time was previously set.
	Epoch (-c) time specified in YYYY-MM-DD, hours (hh), minutes (mm), seconds (ss)
Command arguments	`settime [-u|-c]` (see additional arguments below)
	`[-u|days|hh|mm|ss]`: (Universal parameter string)
	`[-c|YYYY|MM|DD|hh|mm|ss]`: (Epoch parameter string)
Command type	Control: Canister manager command
Access level	General

B   Technical specifications

Enclosure dimensions

Figure 125 Rackmount enclosure dimensions

Note: For clarity, only key view projection lines are shown.

45° orthographic bisector

Key for generic enclosure multiview diagram:

① − Front view: simplified (section removed)   ④ − Enclosure bezel/ear cap top view (reference only)
② − Top view: simplified (section removed)   ⑤ − Enclosure bezel/ear cap side view (reference only)
③ − Side view: simplified (section removed)   A–H: Critical–fit dimensions (see table)

	A		B		C		D		E		F		G		H	
Form	**cm**	**in**	**cm**	**in**	**cm**	**in**	**cm**	**in**	**cm**	**in**	**cm**	**in**	**cm**	**in**	**cm**	**in**
2U12[1]	44.5	17.5	8.9	3.5	48.3	19.0	2.5	1.0	58.1	22.9	2.5	1.0	61.2	24.1	66.9	26.4
2U24[2]	44.5	17.5	8.9	3.5	48.3	19.0	2.5	1.0	58.1	22.9	2.5	1.0	61.2	24.1	66.9	26.4
5U84[3]	44.5	17.5	22.2	8.8	48.3	19.0	4.4	1.8	85.1	33.5	6.5	2.5	89.5	35.3	98.4	38.8

1-The 2U12 enclosure uses 3.5" LFF disks. It can also use 2.5" SFF disks with 3.5" adapter. Dummy LFF drive carriers are fitted to empty slots.
2-The 2U24 enclosure uses 2.5" SFF disks. Dummy SFF drive carriers are fitted to empty slots.
3-The 5U84 uses 3.5" LFF disks in the DDIC carrier. It can also use 2.5" SFF disks with 3.5" adapter in the DDIC. Follow slot population rules.

Enclosure weights

Table 38   Enclosure weights

CRU/component	2U12 (kg/lb)	2U24 (kg/lb)	5U84 (kg/lb)
Storage enclosure (empty)	4.8/10.56	4.8/10.56	64/141
Disk drive carrier	0.9/1.98	0.3/0.66	0.8/1.8
Dummy disk drive carrier (air management sled)	0.05/0.11	0.05/0.11	—
Power Cooling Module (PCM)	3.5/7.7	3.5/7.7	—

Table 38 Enclosure weights (continued)

CRU/component	2U12 (kg/lb)	2U24 (kg/lb)	5U84 (kg/lb)
Power Supply Unit (PSU)	—	—	2.7/6
Fan Cooling Module (FCM)	—	—	1.4/3
SBB controller module (maximum weight)	2.6/5.8	2.6/5.8	2.6/5.8
SBB expansion module	1.5/3.3	1.5/3.3	1.5/3.3
RBOD enclosure (fully populated with modules: maximum weight)	32/71	30/66	135/298
EBOD/JBOD enclosure (fully populated with modules: maximum weight)	28/62	25/55	130/287

1-Weights shown are nominal, and subject to variances.

2-2U rail kits add between 2.8 kg (6.2 lb) and 3.4 kg (7.4 lb) to the aggregate enclosure weight. 5U rail kits add significantly more weight.

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

4-Weights may vary due to actual number and type of disk drives (SAS or SSD) and air management modules installed.

Environmental requirements

Table 39 Ambient temperature and humidity

Specification	Temperature range	Relative humidity	Max. Wet Bulb
Operating	RBOD: 5ºC to 35ºC (41ºF to 95ºF) EBOD/JBOD: 5ºC to 40ºC (41ºF to 104ºF)	20% to 80% non-condensing	28ºC
Non-operating (shipping)	-40ºC to +70ºC (-40ºF to +158ºF)	5% to 100% non-precipitating	29ºC

Specification	Measurement/description
Airflow	System must be operated with low pressure rear exhaust installation. Back pressure created by rack doors and obstacles not to exceed 5Pa (0.5 mm H_2O)
Altitude, operating	2U enclosures: 0 to 3,000 meters (0 to 10,000 feet) Maximum operating temperature is de-rated by 5ºC above 2,133 meters (7,000 feet)
	5U enclosures: -100 to 3,000 meters (-330 to 10,000 feet) Maximum operating temperature is de-rated by 1ºC above 900 meters (3,000 feet)
Altitude, non-operating	-100 to 12,192 meters (-330 to 40,000 feet)
Shock, operating	5.0 g, 10 ms, ½ sine pulses, Y-axis
Shock, non-operating	2U enclosures: 30.0 g, 10 ms, ½ sine pulses
	5U enclosures: 30.0 g, 10 ms, ½ sine pulses (Z-axis); 20.0 g, 10 ms, ½ sine pulses (X- and Y-axes)
Vibration, operating	0.21 G_{rms} 5 Hz to 500 Hz random
Vibration, non-operating	1.04 G_{rms} 2 Hz to 200 Hz random
Vibration, relocation	0.3 G_{rms} 2 Hz to 200 Hz 0.4 decades per minute
Acoustics	Operating sound power 2U enclosures: $\leq L_{WAd}$ 6.6 Bels (re 1 pW) @ 23ºC 5U enclosures: $\leq L_{WAd}$ 8.0 Bels (re 1 pW) @ 23ºC
Orientation and mounting:	19" rack mount (2 EIA units; 5 EIA units)
• Rack rails	To fit 800 mm depth racks compliant with the SSI server rack specification
• Rack characteristics	Back pressure not exceeding 5Pa (~0.5 mm H_2O)

Power cooling module

Specifications for the HE580W PCM are provided in the table.

Table 40   2U Power cooling module specifications

Specification	Measurement/description	
Dimensions (size)	84.3 mm high x 104.5 mm wide x 340.8 mm long: • X-axis length: 104.5 mm (4.11 in) • Y-axis length: 84.3 mm (3.32 in) • Z-axis length: 340.8 mm (37.03)	
Maximum output power	580 W	
Voltage range	100–200 VAC rated	
Frequency	50–60 Hz	
Voltage range selection	Auto-ranging: 90–264 VAC, 47–63 Hz	
Maximum inrush current	20A	
Power factor correction	≥ 95% @ nominal input voltage	
Efficiency	115 VAC/60 Hz	230 VAC/50 Hz
	> 80% @ 10% load > 87% @ 20% load > 90% @ 50% load > 87% @ 100% load > 85% @ surge	> 80% @ 10% load > 88% @ 20% load > 92% @ 50% load > 88% @ 100% load > 85% @ surge
Harmonics	Meets EN61000-3-2	
Output	+5 V @ 42A, +12 V @ 38A, +5 V standby voltage @ 2.7A	
Operating temperature	0 to 57ºC (32ºF to +135ºF)	
Hot pluggable	Yes	
Switches and LEDs	AC mains switch and four status indicator LEDs	
Enclosure cooling	Dual axial cooling fans with variable fan speed control	

Power supply unit

Table 41   5U Power supply unit specifications

Specification	Measurement/description
Maximum output power	2,214 W maximum continuous output power at high line voltage
Voltage	+12 V at 183 A (2,196 W) +5 V standby voltage at 2.7 A
Voltage range	200–240 VAC rated
Frequency	50–60 Hz
Power factor correction	≥ 95% @ 100% load
Efficiency	82% @ 10% load 90% @ 20% load 94% @ 50% load 91% @ 100% load
Holdup time	5 ms from ACOKn high to rails out of regulation (see SBB v2 specification)
Mains inlet connector	IEC60320 C20 with cable retention

Table 41   5U Power supply unit specifications (continued)

Specification	Measurement/description
Weight	3 kg (6.6 lb)
Cooling fans	2 stacked fans: 80x80x38mm

C Standards and regulations

International standards

The enclosure system complies with the requirements of the following agencies and standards:

- CE to EN 60950-1
- CB report to IEC 60950-1
- UL & cUL to UL 60950-1 second edition

Potential for radio frequency interference

USA Federal Communications Commission (FCC)

Notice

This equipment has been tested and found to comply with the limits for a class A digital device, pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy, and if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his or her expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. The supplier is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

European regulations

This equipment complies with European Regulations EN 55022 Class A: Limits and Methods of Measurement of Radio Disturbance Characteristics of Information Technology Equipment and EN50082-1: Generic Immunity.

Safety compliance

Table 42   Safety compliance specifications

System product type approval	UL/cUL/CE
Safety compliance	UL 60950-1 second edition
	IEC 60950-1
	EN 60950-1

EMC compliance

Table 43   EMC compliance specifications

Conducted emissions limit levels	CFR47 Part 15B Class A
	EN55022 Class A
	CISPR Class A

Table 43 EMC compliance specifications (continued)

Radiated emissions limit levels	CFR47 Part 15B Class A
	EN55022 Class A
	CISPR Class A
Harmonics and flicker	EN61000-3-2/3
Immunity limit levels	EN55024

AC power cords

Table 44 AC power cord specifications

1	United States of America–Must be NRTL Listed (National Recognized Test Laboratory – e.g., UL):		
	Chassis form factor	**2U12/2U24**	**5U84**
	Cord type	SV or SVT, 18 AWG minimum, 3 conductor, 2.0M max length	SJT or SVT, 12 AWG minimum, 3 conductor
	Plug (AC source)	NEMA 5–15P grounding-type attachment plug rated 120V 10A *or* IEC 320, C-14, 250V, 10A	IEC 320, C-20, 250V, 20A *or* A suitable plug rated 250V, 20A
	Socket	IEC 320, C-13, 250V, 10A	IEC 320, C-19, 250V, 20A
2	Europe and others–General requirements:		
	Chassis form factor	**2U12/2U24**	**5U84**
	Cord type	Harmonized, H05-WF-3G1.0	Harmonized, H05-WF-3G2.5
	Plug (AC source)		IEC 320, C-20, 250V, 16A *or* A suitable plug rated 250V, 16A
	Socket	IEC 320, C-13, 250V, 10A	IEC 320, C-19, 250V, 16A

ⓘ **IMPORTANT:**   The plug and the complete power cable assembly must meet the standards appropriate to the country, and must have safety approvals acceptable in that country.

Recycling of Waste Electrical and Electronic Equipment (WEEE)

At the end of the product's life, all scrap/waste electrical and electronic equipment should be recycled in accordance with national regulations applicable to the handling of hazardous/toxic electrical and electronic waste materials.

Please contact your supplier for a copy of the Recycling Procedures applicable to your country.

ⓘ **IMPORTANT:**   Observe all applicable safety precautions detailed in the preceding chapters (weight restrictions, handling batteries and lasers, etc.) when dismantling and disposing of this equipment.

D   USB device connection

Rear panel USB ports

The 5005/4005/3005 Series controller modules contain a USB (Universal Serial Bus) Device interface pertaining to the Management Controller (MC). The Device interface is accessed via a port on the controller module face plate. This appendix describes the USB Type B port labeled CLI, which enables direct connection between a management computer and the RBOD controller module, using the command-line interface and appropriate cable.

USB CLI port

Connect USB cable to CLI port
on the canister face plate

Host interface ports
are not shown

Labels are not shown
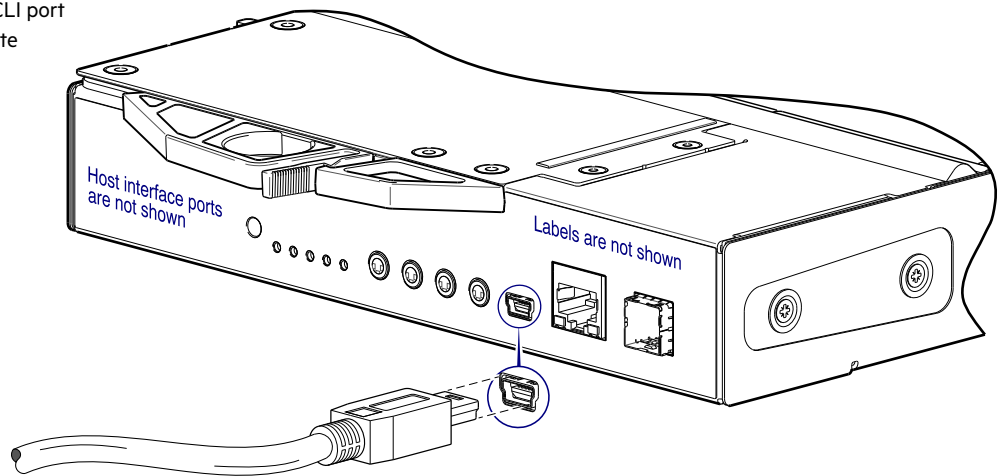
Figure 126 USB device connection—CLI port

5005/4005/3005 Series controller modules feature a USB CLI port used to cable directly to the controller and initially set IP addresses, or perform other configuration tasks. The USB CLI port employs a mini-USB Type B form factor, and requires a specific cable and additional support, 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.

For Linux computers, no new driver files are needed, but a Linux configuration file must be created or modified. See also "Setting parameters for the device driver" (page 177). For Windows computers a special device driver file, `gserial.inf`, must be downloaded from a CD or web site, and installed on the computer that will be cabled directly to the controller's CLI port. See also "Microsoft Windows" (page 176).

Emulated serial port

Once attached to the controller module, the management computer should detect a new USB device. Using the Emulated Serial Port interface, the 5005/4005/3005 Series controller presents a single serial port using a *customer vendor ID* and *product ID*. Effective presentation of the emulated serial port assumes the management computer previously had terminal emulator installed. See also "Supported host applications" (page 176). Serial port configuration is unnecessary.

ⓘ **IMPORTANT:**   Certain operating systems require a device driver or special mode of operation to enable proper functioning of the USB CLI port. See also "Device driver/special operation mode" (page 176).

Supported host applications

5005/4005/3005 Series controllers support the following applications to facilitate connection.

Table 45 Supported terminal emulator applications

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

Command-line interface

Once the management computer detects connection to the USB-capable device, the Management Controller awaits input of characters from the host computer via the command-line. To see the command-line prompt, you must press **Enter**. The MC provides direct access to the CLI.

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

Device driver/special operation mode

Certain operating systems require a device driver or special mode of operation. Product and vendor identification information required for such setup is provided below.

Table 46 USB vendor and product identification codes

USB identification code type	Code
USB Vendor ID	0x210c
USB Product ID	0xa4a7

Microsoft Windows

Microsoft Windows operating systems provide a USB serial port driver. However, the USB driver requires details for connecting to 5005/4005/3005 Series controller enclosures. The 5005/4005/3005 Series provides a device driver for use in the Windows environment. The USB device driver and installation instructions are available via a download.

Obtaining the software download

⊙ **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, using the procedure below.

---

1. Verify that the management computer has Internet access.
   a. For customer support, see seagate.com/support.
2. Select **Products > Systems > General Support**.
   a. From **Downloads**, select the DH USB Driver.

      The DH USB Driver zip file will appear in your Downloads directory.
   b. Follow the instructions accompanying the device driver topic for Microsoft Windows.

Although Linux operating systems do not require installation of a device driver, certain parameters must be provided during driver loading to enable recognition of the 5005/4005/3005 Series controller enclosures.

Setting parameters for the device driver

1. Enter the following command:

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

2. Press **Enter** to execute the command.

   The Linux device driver is loaded with the parameters required to recognize the controllers.

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

Using the CLI port and cable—known issues on Windows

When using the CLI port and cable for setting network port IP addresses, 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 create a new connection or open an existing connection (HyperTerminal):

1. From the Windows Control Panel, select Device Manager.

2. Connect using the USB COM port and Detect Carrier Loss option.

   a. Select **Connect To > Connect using: >** pick a COM port from the list.

   b. Select the **Detect Carrier Loss** check box.

   The Device Manager page should show "Ports (COM & LPT)" with an entry entitled "Disk Array USB Port (COM*n*)"—where *n* is your system's COM port number.

3. Set network port IP addresses using the CLI (see procedure on).

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).

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

E SFP option for CNC ports

Locate the SFP transceivers

Locate the qualified SFP options for your FC/iSCSI controller canister within your product ship kit. The SFP transceiver (SFP) should look similar to the generic SFP shown in Figure 127 and Figure 128 below. When installing an SFP, refer to this figure.
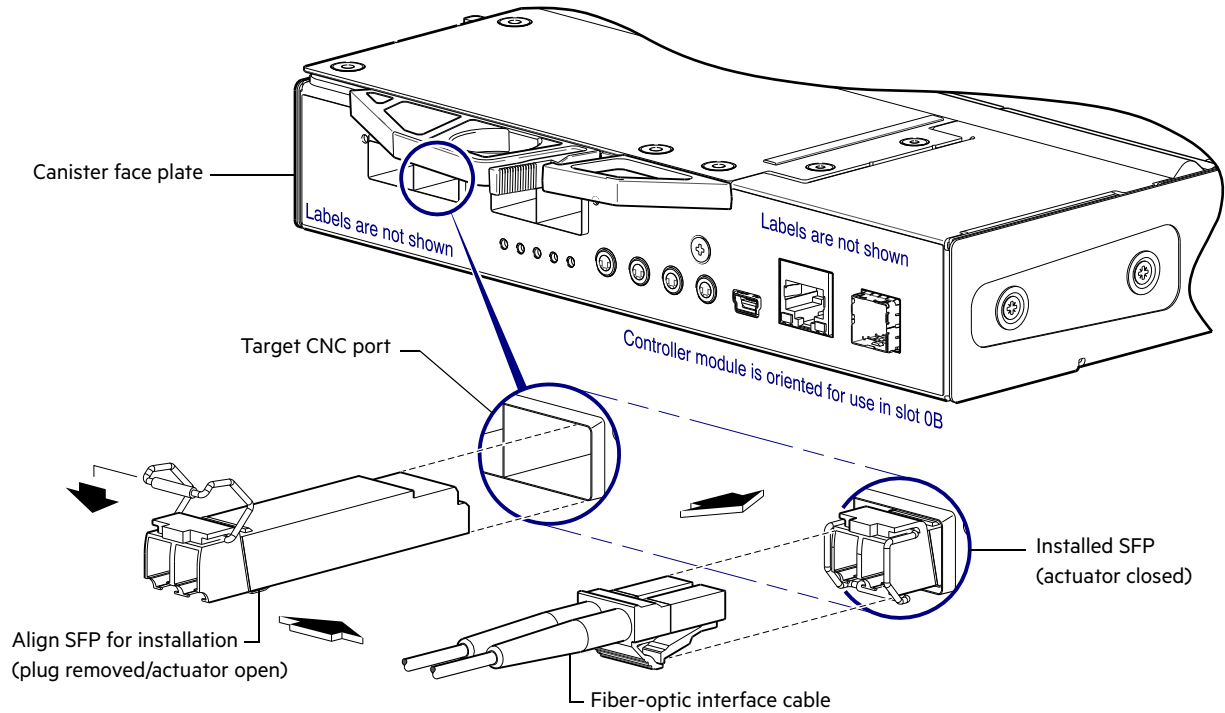
Figure 127 Install a qualified SFP option into a 5005 or 4005 Series FC/iSCSI controller module
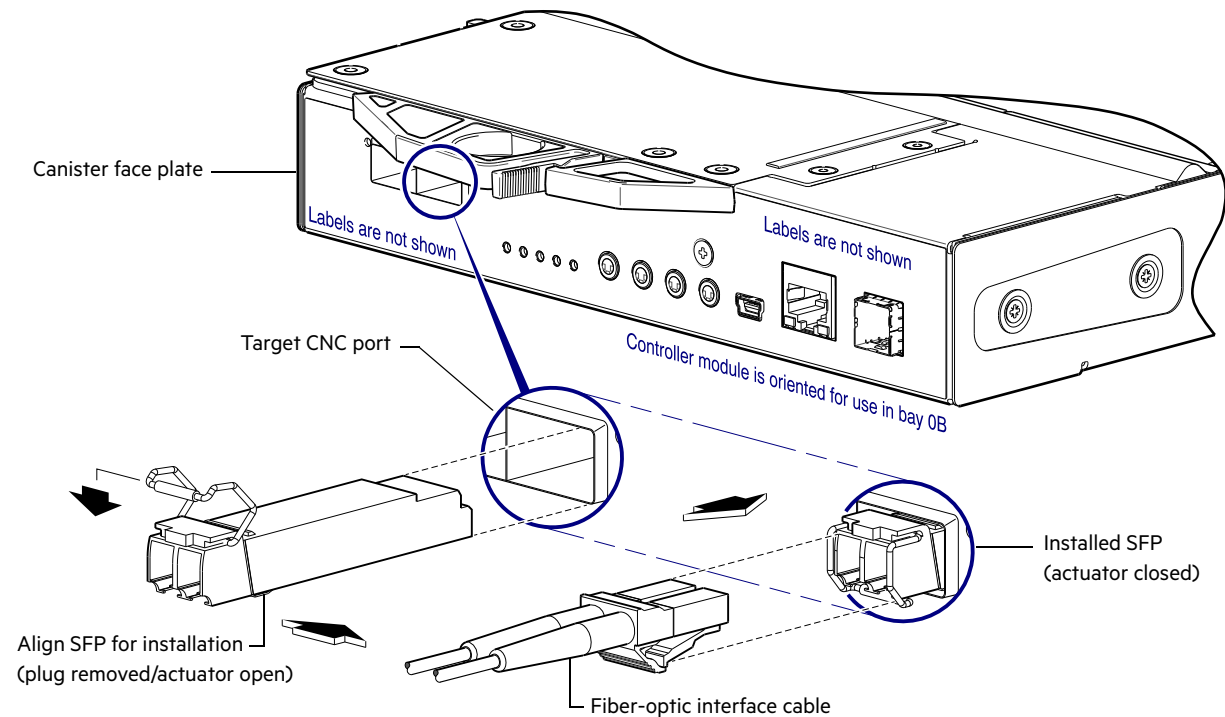
Figure 128 Install a qualified SFP option into a 3005 Series FC/iSCSI controller module

Install an SFP transceiver

For each target CNC port, perform the following procedure to install an SFP. Refer to the figure above when performing the steps. Follow the guidelines provided in "ESD precautions" (page 116) when installing an SFP.

1. Orient the SFP for the target CNC port and canister position, and align it for insertion.

   Depending upon whether the SFP is installed into controller A or B, it will either install right-side up, or upside down in 2U enclosures.

2. If the SFP has a plug, remove it before installing the transceiver. Retain the plug.

3. Flip the actuator open (sweep up or down) according to the SFP position and alignment for canister slot 0A (top) or 0B (bottom).

   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 above).

4. Slide the SFP into the target CNC port until it locks securely into place.

5. Flip the actuator closed (sweep up or down) according to its position in canister slot 0A or 0B.

   The installed SFP should look similar to the position shown in the right detail view above.

6. When ready to attach to the host, obtain and connect a qualified fiber-optic interface cable into the duplex jack at the end of the SFP connector.

Verify component operation

View the CNC port Link Status/Link Activity LED on the controller canister face plate. A green LED indicates that the port is connected and the link is up:

- For FC SFPs, see Figure 34 (page 41) and Figure 35 (page 42) and refer to table entry No.1
- For 10GbE iSCSI SFPs, see Figure 34 (page 41) and Figure 35 (page 42) and refer to table entry No.2
- For 1Gb iSCSI SFPs, see Figure 36 (page 43) and Figure 37 (page 44) and refer to table entry No.2

NOTE: To remove an external SFP connector, perform the installation steps in *reverse* order relative to what is described in "Install an SFP transceiver" (page 179).

Index