



## **Cisco HX C240 M6 Node Installation and Service Guide**

**First Published:** 2021-07-06

**Last Modified:** 2022-10-20

### **Americas Headquarters**

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
<http://www.cisco.com>  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The following information is for FCC compliance of Class A devices: 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 users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If the equipment causes interference to radio or television reception, which can be determined by turning the equipment off and on, users are encouraged to try to correct the interference by using one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio/TV technician for help.

Modifications to this product not authorized by Cisco could void the FCC approval and negate your authority to operate the product.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2021–2022 Cisco Systems, Inc. All rights reserved.



## CONTENTS

### [Introduction](#) ?

---

#### PREFACE

---

### [Communications, Services, Bias-free Language, and Additional Information](#) ix

#### CHAPTER 1

### [Overview](#) 1

[Overview](#) 1

[External Features](#) 4

[PCIe Risers](#) 23

[Summary of Node Features](#) 29

[Serviceable Component Locations](#) 40

---

#### CHAPTER 2

### [Installing the Node](#) 43

[Preparing for Installation](#) 43

[Installation Warnings and Guidelines](#) 43

[Rack Requirements](#) 45

[Installing the Node in a Rack](#) 46

[Installing the Cable Management Arm \(Optional\)](#) 49

[Reversing the Cable Management Arm \(Optional\)](#) 50

[Initial Server Setup](#) 51

[Connecting to the Server Locally For Setup](#) 52

[Connecting to the Server Remotely For Setup](#) 53

[Setting Up the System With the Cisco IMC Configuration Utility](#) 54

[NIC Mode and NIC Redundancy Settings](#) 56

[Updating the BIOS and Cisco IMC Firmware](#) 57

[Accessing the System BIOS](#) 58

[Smart Access Serial](#) 58

Smart Access USB 59

---

**CHAPTER 3**

**Maintaining the Node 61**

Status LEDs and Buttons 61

Front-Panel LEDs 62

Rear-Panel LEDs 65

Internal Diagnostic LEDs 67

Preparing For Component Installation 68

Required Equipment For Service Procedures 69

Shutting Down and Removing Power From the Node 69

Shutting Down Using the Power Button 69

Shutting Down Using The Cisco IMC GUI 70

Shutting Down Using The Cisco IMC CLI 70

Removing the Node Top Cover 70

Serial Number Location 72

Hot Swap vs Hot Plug 72

Replacing the Air Duct 72

Removing the Air Duct 72

Installing the Air Duct 74

Removing and Replacing Components 75

Serviceable Component Locations 76

Replacing Front-Loading SAS/SATA Drives 78

Front-Loading SAS/SATA Drive Population Guidelines 78

4K Sector Format SAS/SATA Drives Considerations 79

Replacing a Front-Loading SAS/SATA Drive 80

Replacing Rear-Loading SAS/SATA Drives 81

Rear-Loading SAS/SATA Drive Population Guidelines 81

Replacing a Rear-Loading SAS/SATA Drive 82

Replacing Mid-Mounted SAS/SATA Drives (LFF Server) 83

Basic Troubleshooting: Reseating a SAS/SATA Drive 85

Reseating a SAS/SATA Drive 86

Replacing Front-Loading NVMe SSDs 87

Front-Loading NVMe SSD Population Guidelines 87

Front-Loading NVMe SSD Requirements and Restrictions 87

Enabling Hot-Plug Support in the System BIOS	88
Replacing a Front-Loading NVMe SSD	88
Cabling NVMe Drives 1 Through 4 (UCS C240 M6 24 SFF Drives Only)	90
Replacing Rear-Loading NVMe SSDs	94
Rear-Loading NVMe SSD Population Guidelines	95
Rear-Loading NVMe SSD Requirements and Restrictions	95
Replacing a Rear-Loading NVMe SSD	96
Basic Troubleshooting: Reseating an NVMe Drive	97
Reseating an NVMe Drive	98
Replacing Fan Modules	99
Removing the Fan Tray	100
Installing the Fan Tray	101
Replacing CPUs and Heatsinks	102
CPU Configuration Rules	102
Tools Required For CPU Replacement	104
Removing CPUs and Heat Sinks	104
Installing the CPUs and Heatsinks	109
Additional CPU-Related Parts to Order with RMA Replacement CPUs	113
Additional CPU-Related Parts to Order with RMA Replacement System Chassis	114
Replacing Memory DIMMs	115
DIMM Population Rules and Memory Performance Guidelines	115
Replacing DIMMs	117
Replacing Intel Optane DC Persistent Memory Modules	118
Intel Optane DC Persistent Memory Module Population Rules and Performance Guidelines	118
Installing Intel Optane DC Persistent Memory Modules	119
Node BIOS Setup Utility Menu for DCPMM	120
Replacing a Mini-Storage Module	122
Replacing a Mini-Storage Module Carrier	122
Replacing an M.2 SSD in a Mini-Storage Carrier For M.2	123
Replacing the RTC Battery	123
Replacing Power Supplies	124
Supported Power Supplies	124
Replacing AC Power Supplies	125
Replacing DC Power Supplies	126

Installing DC Power Supplies (First Time Installation)	127
Grounding for DC Power Supplies	129
Replacing a PCIe Riser	129
Replacing a PCIe Card	130
PCIe Slot Specifications	130
Replacing a PCIe Card	131
Cisco Virtual Interface Card (VIC) Considerations	133
Replacing an mLOM Card	134
Replacing a SAS Storage Controller Card (RAID or HBA)	135
Storage Controller Card Firmware Compatibility	135
Removing the Dual Storage Controller Cards	135
Installing the Dual Storage Controller Cards	139
Removing the Storage Controller Card	142
Installing the Storage Controller Card	146
Verify Cabling	149
Replacing a SATA Interposer Card (12-Drive SFF Server Only)	151
Replacing the Supercap (RAID Backup)	155
Replacing a Boot-Optimized M.2 RAID Controller Module	158
Cisco Boot-Optimized M.2 RAID Controller Considerations	158
Replacing a Cisco Boot-Optimized M.2 RAID Controller	160
Replacing a Chassis Intrusion Switch	162
Installing a Trusted Platform Module (TPM)	163
TPM Considerations	163
Installing and Enabling a TPM	163
Installing TPM Hardware	164
Enabling the TPM in the BIOS	164
Enabling the Intel TXT Feature in the BIOS	165
Removing the PCB Assembly (PCBA)	166
Service Headers and Jumpers	167
Using the Clear CMOS Switch (SW12, Switch 9)	169
Using the BIOS Recovery Switch (SW12, Switch 5)	170
Procedure 1: Reboot With bios.cap Recovery File	170
Using the BIOS Recovery Switch (SW12, Switch 5) and bios.cap File	171
Using the Clear BIOS Password Switch (SW12, Switch 6)	172

Using the Boot Alternate Cisco IMC Image Header (CN3, Pins 1-2) 172

Using the System Firmware Secure Erase Header (CN3, Pins 3-4) 173

---

**APPENDIX A**

**Node Specifications 175**

Physical Specifications 175

Environmental Specifications 175

Power Specifications 176

1050 W AC Power Supply 176

1600 W AC Power Supply 177

1050 W DC Power Supply 178

Power Cord Specifications 179

---

**APPENDIX B**

**Storage Controller Considerations 183**

Storage Controller Card Firmware Compatibility 183

RAID Backup (Supercap) 183

Mixing Drive Types in RAID Groups 184

Storage Controller Cable Connectors and Backplanes 184

For More RAID Utility Information 188

---

**APPENDIX C**

**GPU Card Installation 189**

Node Firmware Requirements 189

GPU Card Configuration Rules 189

Requirement For All GPUs: Memory-Mapped I/O Greater Than 4 GB 190

Installing a Single-Wide GPU Card 191

Installing a Double-Wide GPU Card 194

Replacing a Heatsink 196

Removing a Heat Sink 197

Installing a Heatsink 200

Installing GRID Licenses From the NVIDIA Licensing Portal to the License Server 202

Reading Your License Server's MAC Address 203

Installing Licenses From the Licensing Portal 203

Viewing Available GRID Licenses 203

Viewing Current License Usage 204

Managing GRID Licenses 204

Acquiring a GRID License on Windows 204

Acquiring a GRID License on Linux 204

Using gpumodeswitch 205

Installing Drivers to Support the GPU Cards 206

- 1. Updating the Server BIOS 206
- 2. Updating the GPU Card Drivers 206





## Communications, Services, Bias-free Language, and Additional Information

---

- To receive timely, relevant information from Cisco, sign up at [Cisco Profile Manager](#).
- To get the business impact you're looking for with the technologies that matter, visit [Cisco Services](#).
- To submit a service request, visit [Cisco Support](#).
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit [Cisco Marketplace](#).
- To obtain general networking, training, and certification titles, visit [Cisco Press](#).
- To find warranty information for a specific product or product family, access [Cisco Warranty Finder](#).

### Documentation Feedback

To provide feedback about Cisco technical documentation, use the feedback form available in the right pane of every online document.

### Cisco Bug Search Tool

[Cisco Bug Search Tool](#) (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.

### Bias-Free Language

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.





# CHAPTER 1

## Overview

---

- [Overview, on page 1](#)
- [External Features, on page 4](#)
- [Summary of Node Features, on page 29](#)
- [Serviceable Component Locations, on page 40](#)

## Overview

The Cisco HyperFlex (HX) C240 M6 is a stand-alone 2U rack server chassis that can operate in both standalone environments and as part of the Cisco Unified Computing System (Cisco UCS).

The Cisco HX C240 M6 servers support a maximum of two 3rd Gen Intel® Xeon® Scalable Processors, in either one or two CPU configurations.

The servers support:

- 16 DIMM slots per CPU for 3200-MHz DDR4 DIMMs in capacities up to 128 GB DIMMs.
- A maximum of 8 or 12 TB of memory is supported for a dual CPU configuration populated with either:
  - DIMM memory configurations of either 32 128 GB DDR DIMMs, or 16 128 GB DDR4 DIMMs plus 16 512 GB Intel® Optane™ Persistent Memory Modules (DCPMMs).
- The servers have different supported drive configurations depending on whether they are configured with large form factor (LFF) or small form factor (SFF) front-loading drives.
- The C240 M6 12 LFF supports midplane mounted storage through a maximum of r4 LFF HDDs.
- Up to 2 M.2 SATA RAID cards for server boot.
- Rear Storage risers (2 slots each)
- One rear PCIe riser (3 slots)
- Internal slot for a 12 G SAS RAID controller with SuperCap for write-cache backup, or for a SAS HBA.
- Network connectivity through either a dedicated modular LAN over motherboard card (mLOM) that accepts a series 14xxx Cisco virtual interface card (VIC) or a third-party NIC. These options are in addition to Intel x550 10Gbase-T mLOM ports built into the server motherboard.
- One mLOM/VIC card provides 10/25/40/50/100 Gbps. The following mLOMs are supported:

- Cisco HX VIC 15428 Quad Port CNA MLOM (UCSC-M-V5Q50G) supports:
  - a x16 PCIe Gen4 Host Interface to the rack server
  - four 10G/25G/50G SFP56 ports
  - 4GB DDR4 Memory, 3200 MHz
  - Integrated blower for optimal ventilation
- Cisco HX VIC 1467 Quad Port 10/25G SFP28 mLOM (UCSC-M-V25-04) supports:
  - a x16 PCIe Gen3 Host Interface to the rack server
  - four 10G/25G QSFP28 ports
  - 2GB DDR3 Memory, 1866 MHz
- Cisco HX VIC 1477 Dual Port 40/100G QSFP28 (UCSC-M-V100-04)
  - a x16 PCIe Gen3 Host Interface to the rack server
  - two 40G/100G QSFP28 ports
  - 2GB DDR3 Memory, 1866 MHz

These options are in addition to Intel x550 10Gbase-T mLOM ports built into the server motherboard.

- The following virtual interface cards (VICs) are supported in addition to some third-party VICs:
  - Cisco HX VIC 1455 quad port 10/25G SFP28 PCIe (UCSC-PCIE-C25Q-04=)
  - Cisco HX VIC 1495 Dual Port 40/100G QSFP28 CNA PCIe (UCSC-PCIE-C100-042)
- Two power supplies (PSUs) that support N+1 power configuration.
- Six modular, hot swappable fans.

### Server Configurations, LFF

The server is orderable with the following configuration for LFF drives.

- Cisco HX C240 M6 LFF 12 (HX-C240-M6L)—Large form-factor (LFF) drives, with a 12-drive backplane.
  - Front-loading drive bays 1—12 support 3.5-inch SAS/SATA drives.
  - The midplane drive cage supports four 3.5-inch SAS-only drives.
  - Optionally, rear-loading drive bays support either two or four SAS/SATA or NVMe drives.

### Server Configurations, SFF 12 SAS/SATA

The SFF 12 SAS/SATA configuration (HX-C240-M6-S) can be configured with 12 SFF drives and an optional optical drive. Also, The SFF configurations can be ordered as either an I/O-centric configuration or a storage centric configuration. This server supports the following:

- A maximum of 12 Small form-factor (SFF) drives, with a 12-drive backplane.

- Front-loading drive bays 1—12 support a maximum of 12 2.5-inch SAS/SATA drives as SSDs or HDDs.
- Optionally, drive bays 1—4 can support 2.5-inch NVMe SSDs. In this configuration, any number of NVMe drives can be installed up to the maximum of 4.



---

**Note** NVMe drives are supported only on a dual CPU server.

---

- The server can be configured with a SATA Interposer card. If your server uses a SATA Interposer card, up to a maximum of 8 SATA-only drives can be configured. These drives can be installed only in slots 1-8.
- Drive bays 5 —12 support SAS/SATA SSDs or HDDs only; no NVMe.
- Optionally, the rear-loading drive bays support four 4 2.5-inch SAS/SATA or NVMe drives.

### Server Configurations, 24 SFF SAS/SATA

The SFF 24 SAS/SATA configuration (HX-C240-M6SX) can be ordered as either an I/O-centric configuration or a storage centric configuration. This server supports the following:

- A maximum of 24 small form-factor (SFF) drives, with a 24-drive backplane.
  - Front-loading drive bays 1—24 support 2.5-inch SAS/SATA drives as SSDs or HDDs.
  - Optionally, drive bays 1—4 can support 2.5-inch NVMe SSDs. In this configuration, any number of NVMe drives can be installed up to the maximum of 4.



---

**Note** NVMe drives are supported only on a dual CPU server.

---

- Drive bays 5 —24 support SAS/SATA SSDs or HDDs only; no NVMe.
- Optionally, the rear-loading drive bays support four 4 2.5-inch SAS/SATA or NVMe drives.
- As an option, this server can be ordered with "GPU ready" configuration. This option supports adding GPUs at a later date even though the GPU is not purchased at the time the server is initially ordered.



---

**Note** To order the GPU Ready configuration through the Cisco online ordering and configuration tool, you must select the GPU air duct PID to enable GPU ready configuration. Follow the additional rules displayed in the tool. For additional information, see [GPU Card Configuration Rules, on page 189](#).

---

### Server Configurations, 12 NVMe

The SFF 12 NVMe configuration (HX-C240-M6N) can be ordered as an NVMe-only server. The NVMe-optimized server requires two CPUs. This server supports the following:

- A maximum of 12 SFF NVMe drives as SSDs with a 12-drive backplane, NVMe-optimized.
  - Front-loading drive bays 1—12 support 2.5-inch NVMe PCIe SSDs *only*.
  - The two rear-loading drive bays support two 2.5-inch NVMe SSDs *only*. These drive bays are the top and middle slot on the left of the rear panel.

### Server Configurations, 24 NVMe

The SFF 24 NVMe configuration (HX-C240-M6SN) can be ordered as an NVMe-only server. The NVMe-optimized server requires two CPUs. This server supports the following:

- A maximum of 24 SFF NVMe drives as SSDs with a 24-drive backplane, NVMe-optimized.
  - Front-loading drive bays 1—24 support 2.5-inch NVMe PCIe SSDs *only*.
  - The two rear-loading drive bays support two 2.5-inch NVMe SSDs *only*. These drive bays are the top and middle slot on the left of the rear panel.
  - As an option, this server can be ordered with "GPU ready" configuration. This option supports adding GPUs at a later date even though the GPU is not purchased at the time the server is initially ordered.




---

**Note** To order the GPU Ready configuration through the Cisco online ordering and configuration tool, you must select the GPU air duct PID to enable GPU ready configuration. Follow the additional rules displayed in the tool. For additional information, see [GPU Card Configuration Rules, on page 189](#).

---

## External Features

This topic shows the external features of the different configurations of the node.

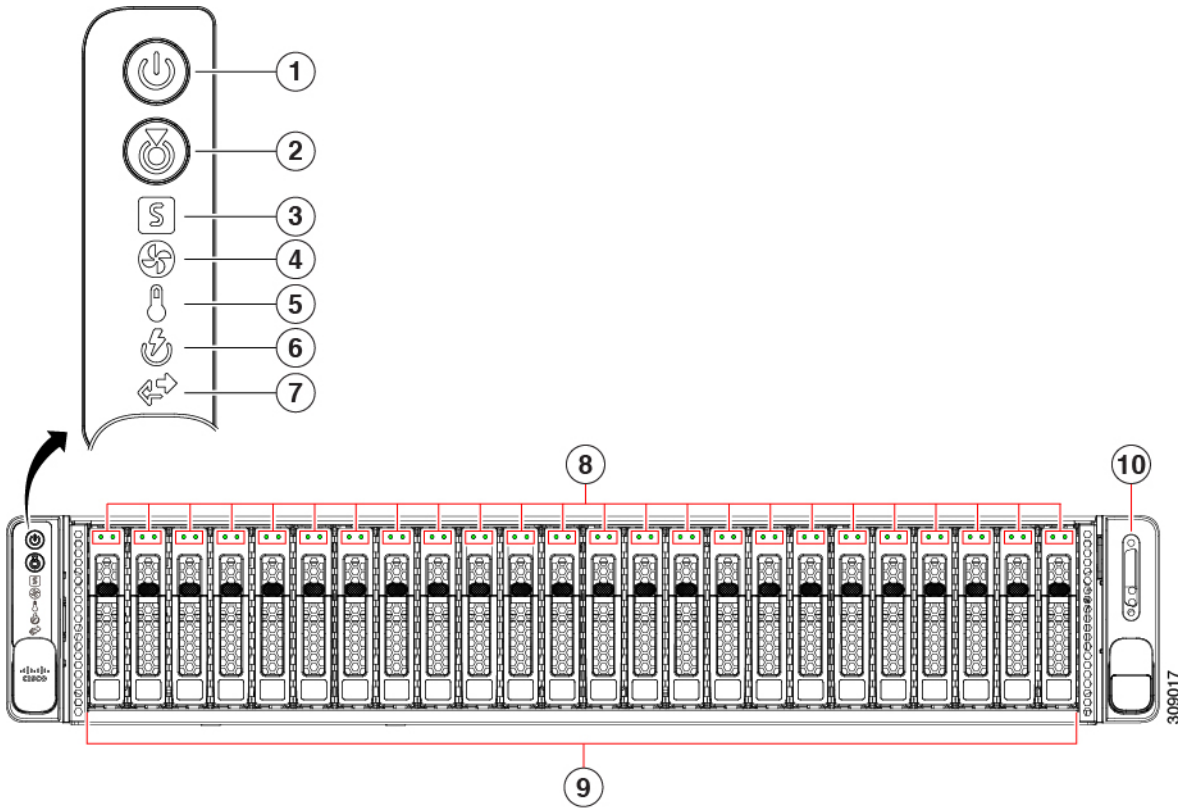
For definitions of LED states, see [Front-Panel LEDs, on page 62](#).

### Cisco HX C240 M6 Node 24 SAS/SATA Front Panel Features

The following figure shows the front panel features of Cisco HX C240 M6SX, which is the small form-factor (SFF), 24 SAS/SATA drive version of the node. Front-loading drives can be mix and match in slots 1 through 4 to support up to four SFF NVMe or SFF SAS/SATA drives. HX C240 M6 servers with any number of NVMe drives must be dual CPU systems.

This configuration can support up to 4 optional SAS/SATA drives in the rear PCIe slots.

Figure 1: Cisco HX C240 M6 node 24 SAS/SATA Front Panel



1	Power Button/Power Status LED	2	Unit Identification LED
3	System Status LEDs	4	Fan Status LED
5	Temperature Status LED	6	Power Supply Status LED
7	Network Link Activity LED	8	Drive Status LEDs
9	NVMe Drive Bays, front loading Drive bays 1—24 support front-loading SFF SAS/SATA drives.  Drive bays 1 through 4 can support SAS/SATA hard drives and solid-state drives (SSDs) or NVMe PCIe drives. Any number of NVMe drives up to 4 can reside in these slots.  Drive bays 5 - 24 support SAS/SATA hard drives and solid-state drives (SSDs) only.  Drive bays are numbered 1 through 24 with bay 1 as the leftmost bay.	10	KVM connector (used with KVM cable that provides one DB-15 VGA, one DB-9 serial, and two USB 2.0 connectors)





<p><b>9</b></p>	<p>Drive Bays front loading</p> <p>Drive bays 1—12 support front-loading SFF SAS/SATA drives.</p> <p>Drive bays 1 thorough 4 can support SAS/SATA hard drives and solid-state drives (SSDs) as well as NVMe PCIe drives. Any number of NVMe drives up to 4 can reside in these slots.</p> <p>Drive bays 5 - 12 support SAS/SATA hard drives and solid-state drives (SSDs) only.</p> <p>Drive bays are numbered 1 through 24 with bay 1 as the leftmost bay.</p> <p><b>Note</b> If the node has a SATA Interposer card, a maximum of 8 SATA drives is supported in slots 1 through 8.</p>	<p><b>10</b></p>	<p>Drive Bays 13 through 24 are blocked off with sheet metal.</p>
<p><b>11</b></p>	<p>KVM connector (used with KVM cable that provides one DB-15 VGA, one DB-9 serial, and two USB 2.0 connectors)</p>	<p>-</p>	

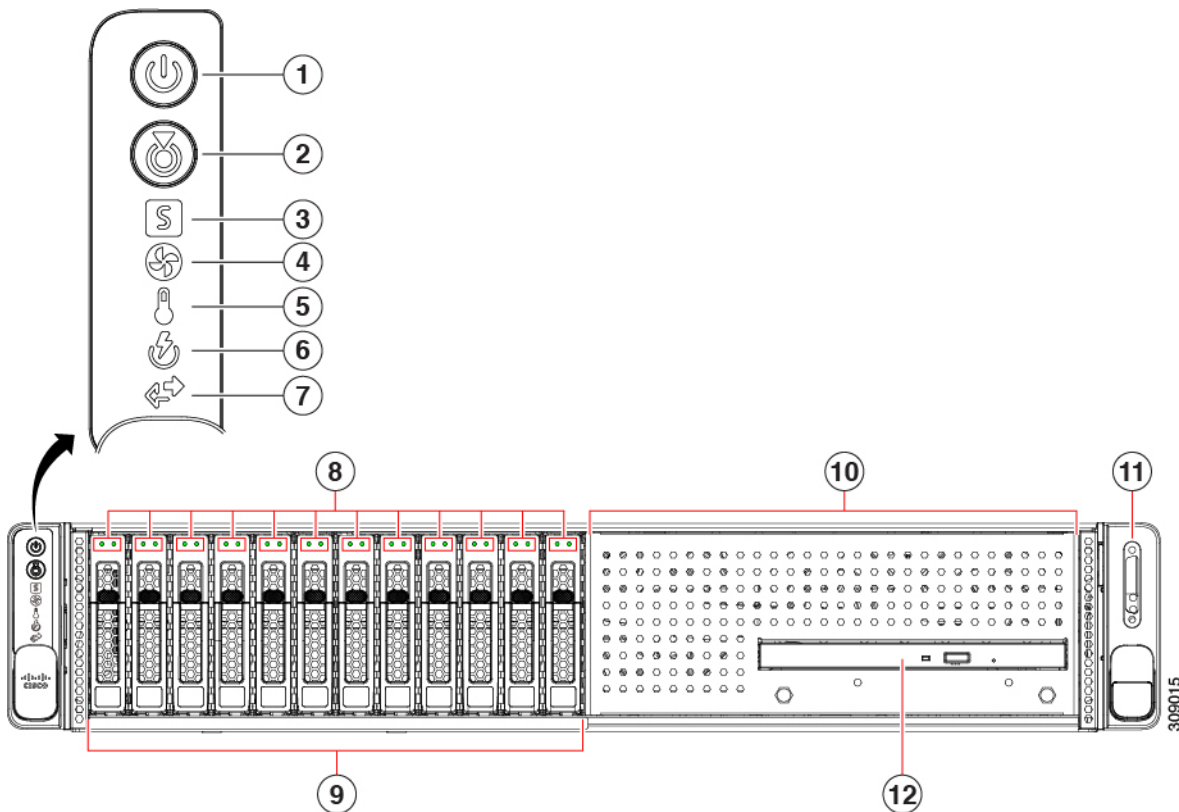
#### Cisco HX C240 M6 Node 12 SAS/SATA Drives (Plus Optical) Front Panel Features

The following figure shows the front panel features of Cisco HX C240 M6S, which is the small form-factor (SFF) drive, 12-drive version of the node. Front-loading drives can be mix and match in slots 1 through 4 to support up to four SFF NVMe or SFF SAS/SATA drives. HX C240 M6 servers with any number of NVMe drives must be dual CPU systems.

This configuration can support up to 4 optional SAS/SATA drives in the rear PCIe slots.

For definitions of LED states, see [Front-Panel LEDs, on page 62](#).

Figure 3: Cisco HX C240 M6 node 12 SAS/SATA Plus Optical Drive, Front Panel Features



1	Power Button/Power Status LED	2	Unit Identification LED
3	System Status LEDs	4	Fan Status LED
5	Temperature Status LED	6	Power Supply Status LED
7	Network Link Activity LED	8	Drive Status LEDs
9	<p>Drive Status LEDs</p> <p>Drive bays 1—12 support front-loading SFF drives.</p> <p>Drive bays 1 through 4 can support SAS/SATA hard drives and solid-state drives (SSDs) as well as NVMe PCIe drives. Any number of NVMe drives up to 4 can reside in these slots.</p> <p><b>Note</b> If the node has a SATA Interposer card, a maximum of 8 SATA drives is supported in slots 1 through 8.</p>	10	Drive bays 13-24 are blocked off with sheet metal.
11	KVM connector (used with KVM cable that provides one DB-15 VGA, one DB-9 serial, and two USB 2.0 connectors)	12	Optional optical DVD drive is installed horizontally.

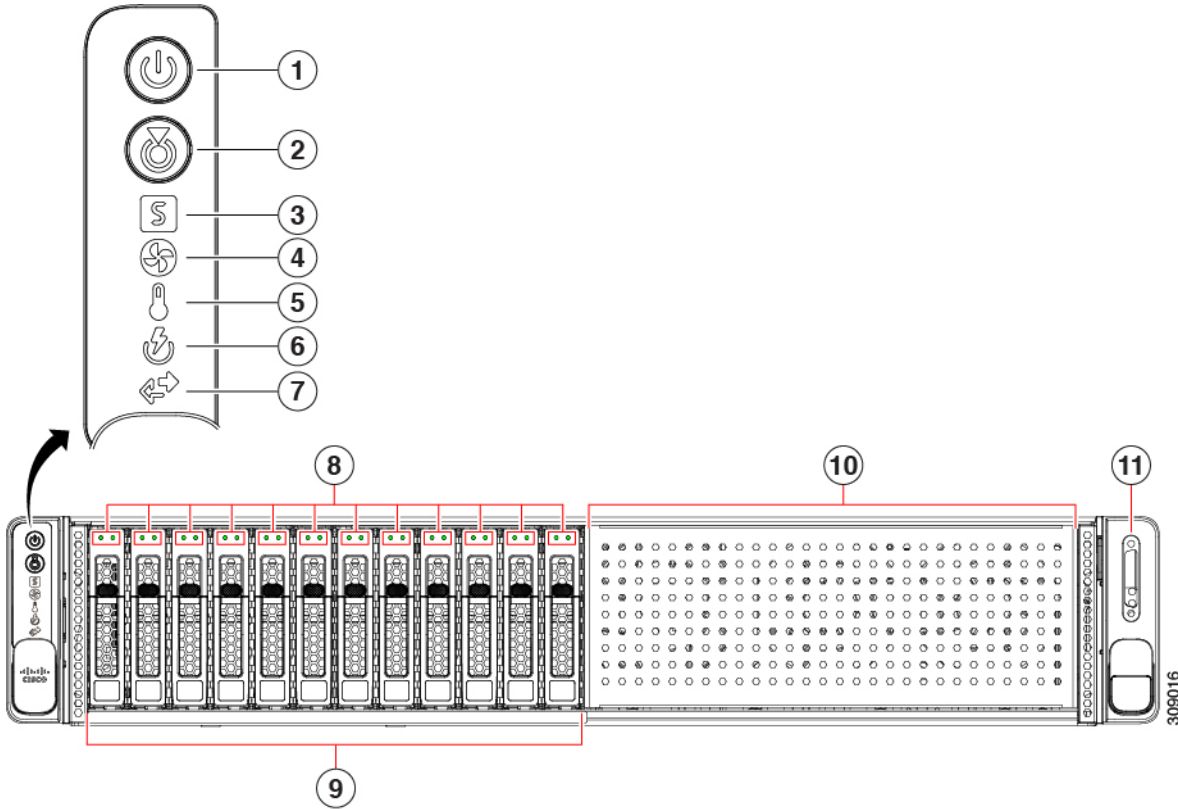
### Cisco HX C240 M6 Node 12 NVMe Drive Front Panel Features

The following figure shows the front panel features of Cisco HX C240 M6N, which is the small form-factor (SFF) drive, 12 NVMe drive version of the node. Front-loading drives are all NVMe only. HX C240 M6 servers with any number of NVMe drives must be dual CPU systems.

This configuration can support up to 2 optional NVMe drives in the rear PCIe slots.

For definitions of LED states, see [Front-Panel LEDs, on page 62](#).

Figure 4: Cisco HX C240 M6 Server 12 NVMe Front Panel



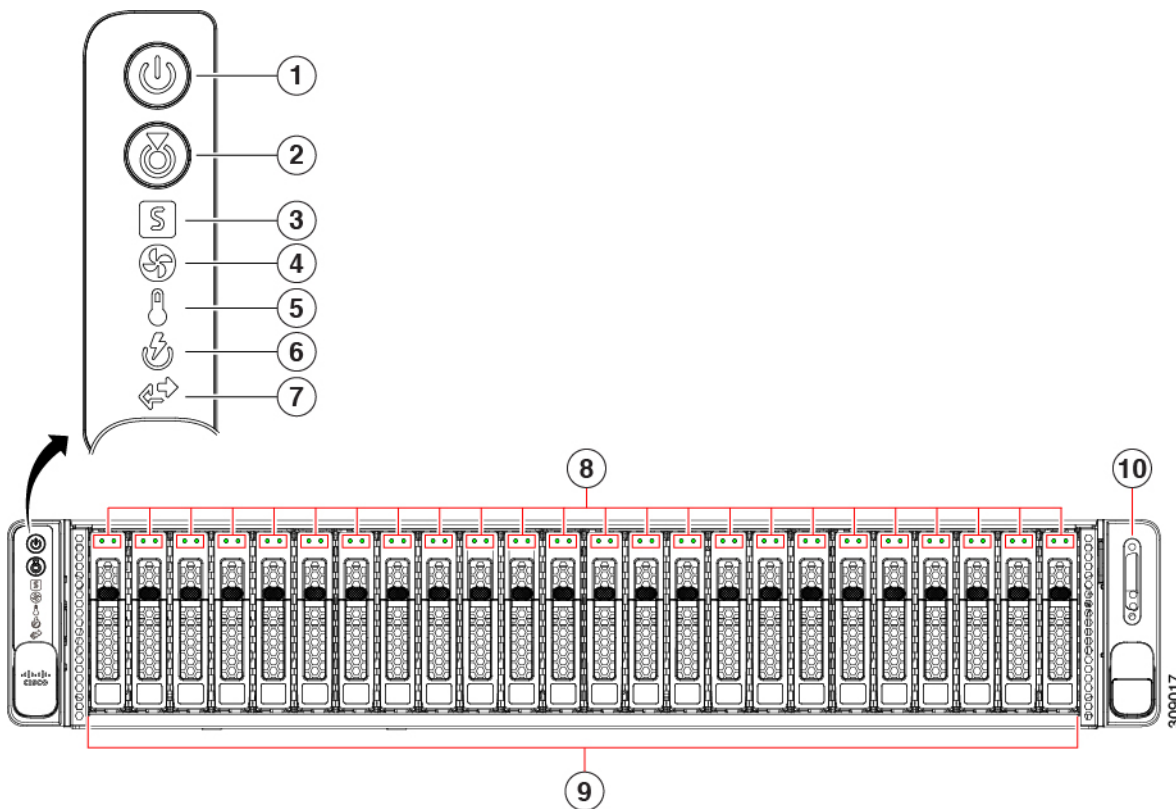
1	Power Button/Power Status LED	2	Unit Identification LED
3	System Status LEDs	4	Fan Status LED
5	Temperature Status LED	6	Power Supply Status LED
7	Network Link Activity LED	8	Drive Status LEDs
9	Drive bays 1—12 support front-loading SFF NVMe drives only. Drive bays are numbered 1 through 12 with bay 1 as the leftmost bay.	10	Drive bays 13 —24 are blocked off with sheet metal.
11	KVM connector (used with KVM cable that provides one DB-15 VGA, one DB-9 serial, and two USB 2.0 connectors)	-	

### Cisco HX C240 M6 Server 24 NVMe Drives Front Panel Features

The following figure shows the front panel features of Cisco HX C240 M6SN, which is the small form-factor (SFF) drive, 24 NVMe drive version of the node. Front-loading drives are all NVMe; SAS/SATA drives are not supported. HX C240 M6 servers with any number of NVMe drives must be dual CPU systems.

This configuration can support up to 2 optional NVMe drives in the rear PCIe slots.

Figure 5: Cisco HX C240 M6 Server 24 NVMe Front Panel




1	Power Button/Power Status LED	2	Unit Identification LED
3	System Status LEDs	4	Fan Status LED
5	Temperature Status LED	6	Power Supply Status LED
7	Network Link Activity LED	8	Drive Status LEDs
9	Drive Status LEDs	10	Drive bays 1—24 support front-loading SFF NVMe drives. Drive bays are numbered 1 through 24 with bay 1 as the leftmost bay.
11	KVM connector (used with KVM cable that provides one DB-15 VGA, one DB-9 serial, and two USB 2.0 connectors)	-	

### Cisco HX C240 M6 Server 12 LFF Drive Front Panel Features

The following figure shows the front panel features of the large form-factor (LFF) configuration of the node. This version of the node supports 12 3.5-inch LFF SAS-only front-loading hard disk drives (HDDs) plus up to 4 3.5-inch LFF mid-plane mounted HDDs. As an option, the node can also support up to 4 SFF drives as SAS, SATA, or NVME in the rear PCIe slots.

For definitions of LED states, see [Front-Panel LEDs](#), on page 62.


**Figure 6: Cisco UCS C240 M6 Server 12 LFF Drive Front Panel**



<b>1</b>	Power Button/Power Status LED	<b>2</b>	Unit Identification LED
<b>3</b>	System Status LEDs	<b>4</b>	Fan Status LED
<b>5</b>	Temperature Status LED	<b>6</b>	Power Supply Status LED
<b>7</b>	Network Link Activity LED	<b>8</b>	Drive Status LEDs
<b>9</b>	KVM connector (used with KVM cable that provides one DB-15 VGA, one DB-9 serial, and two USB 2.0 connectors)	<b>10</b>	Drive bays 1—12 support front-loading LFF SAS-only drives.  Drive bays are numbered 1 through 12 with bay 1 as the leftmost bay, 12 as the rightmost bottom bay.

### Common Rear Panel Features

The following illustration shows the rear panel hardware features that are common across all models of the node.




<p><b>1</b></p>	<p>Rear hardware configuration options:</p> <ul style="list-style-type: none"> <li>• For I/O-Centric, these are PCIe slots.</li> <li>• For Storage-Centric, these are storage drives bays.</li> </ul> <p>This illustration shows the slots unpopulated</p>	<p><b>2</b></p>	<p>Power supplies (two, redundant as 1+1)</p> <p>See <a href="#">Power Specifications, on page 176</a> for specifications and supported options.</p>
<p><b>3</b></p>	<p>VGA video port (DB-15 connector)</p>	<p><b>4</b></p>	<p>Serial port (RJ-45 connector)</p>
<p><b>5</b></p>	<p>1-Gb/10-Gb auto negotiating Ethernet ports, 2 which are ports 1 and 2 in the cluster.</p> <p>These LAN ports (LAN1 and LAN2) support 1 Gbps and 10 Gbps and auto negotiate to the optimal speed based on the link partner capability.</p> <p>The third port in the cluster is the dedicated 1 Gb management port.</p>	<p><b>6</b></p>	<p>USB 3.0 ports, 2</p>
<p><b>7</b></p>	<p>Rear unit identification button/LED</p>	<p><b>8</b></p>	<p>Modular LAN-on-motherboard (mLOM) card slot (x16)</p>

### Cisco HX C240 M6 Server 24 Drive Rear Panel, I/O Centric

The Cisco HX C240 M6 24 SAS/SATA SFF version has a rear configuration option, for either I/O (I/O Centric) or Storage (Storage Centric) with the I/O Centric version of the node offering PCIe slots and the Storage Centric version of the node offering drive bays.

The following illustration shows the rear panel features for the I/O Centric version of the Cisco HX C240 M6SX.

- For features common to all versions of the node, see Common Rear Panel Features.
- For definitions of LED states, see [Rear-Panel LEDs, on page 65](#).



1	Riser 1A	2	Riser 2A
3	Riser 3A or 3C	-	-

The following table shows the riser options for this version of the node.

**Table 1: Cisco HX C240 M6 24 SFF SAS/SATA/NVMe (UCSC-C240-M6SX)**

Riser	Options
<p>Riser 1</p> <p>This riser is I/O-centric and controlled by CPU 1 or CPU 2.</p>	<p>Riser 1A supports three PCIe slots numbered bottom to top:</p> <ul style="list-style-type: none"> <li>• Slot 1 is full-height, 3/4 length, x8, NCSI</li> <li>• Slot 2 is full-height, full-length, x16, NCSI</li> <li>• Slot 3 is full-height, full-length, x8, no NCSI</li> </ul>
<p>Riser 2</p> <p>This riser is I/O-centric and controlled by CPU 2.</p>	<p>Riser 2A supports three PCIe slots:</p> <ul style="list-style-type: none"> <li>• Slot 1 is full-height, 3/4 length, x8, NCSI</li> <li>• Slot 2 is full-height, full-length, x16, NCSI</li> <li>• Slot 3 is full-height, full-length, x8, no NCSI</li> </ul>
<p>Riser 3</p> <p>This riser is I/O-centric and controlled by CPU 2.</p>	<p>Riser 3A supports two PCIe slots:</p> <ul style="list-style-type: none"> <li>• Slot 7 is full-height, full-length, x8</li> <li>• Slot 8 is full-height, full-length, x8</li> </ul> <p>Riser 3C supports a GPU only.</p> <ul style="list-style-type: none"> <li>• Supports one full-height, full-length, double-wide GPU (PCIe slot 7 only), x16</li> <li>• Slot 8 is blocked by double-wide GPU</li> </ul>

### Cisco HX C240 M6 Server 12 SAS/SATA Drive Rear Panel, I/O Centric


The Cisco HX C240 M6 12 SAS/SATA SFF version has a rear configuration option, for either I/O (I/O Centric) or Storage (Storage Centric) with the I/O Centric version of the node offering PCIe slots and the Storage Centric version of the node offering drive bays.



**Note** This version of node has an option for a DVD drive on the front of the node. The rear panel shown here is the same for both the standard node and the DVD drive version of the node.

The following illustration shows the rear panel features for the I/O Centric version of the Cisco HX C240 M6S.

- For features common to all versions of the node, see Common Rear Panel Features.
- For definitions of LED states, see [Rear-Panel LEDs, on page 65](#).



1	Riser 1A	2	Riser 2A
3	Riser 3A or 3C	-	

The following table shows the riser options for this version of the node.

**Table 2: Cisco HX C240 M6 12 SFF SAS/SATA (USC-C240M6-S)**

Riser	Options
Riser 1 This riser is I/O-centric and controlled by CPU 1.	Riser 1A supports three PCIe slots numbered bottom to top: <ul style="list-style-type: none"> <li>• Slot 1 is full-height, 3/4 length, x8, NCSI</li> <li>• Slot 2 is full-height, full-length, x16, NCSI</li> <li>• Slot 3 is full-height, full-length, x8, no NCSI</li> </ul>
Riser 2 This riser is I/O-centric and controlled by CPU 2.	Riser 2A supports three PCIe slots: <ul style="list-style-type: none"> <li>• Slot 4 is full-height, 3/4 length, x8, NCSI</li> <li>• Slot 5 is full-height, full-length, x16, NCSI</li> <li>• Slot 6 is full-height, full-length, x8, no NCSI</li> </ul>




Riser	Options
<p>Riser 3</p> <p>This riser is I/O-centric and controlled by CPU 2.</p>	<p>Riser 3A supports two PCIe slots</p> <ul style="list-style-type: none"> <li>Slot 7 is full-height, full-length, x8, no NCSI</li> <li>Slot 8 is full-height, full-length, x8, no NCSI</li> </ul> <p>Riser 3C supports a GPU only.</p> <ul style="list-style-type: none"> <li>Supports one full-height, full-length, double-wide GPU (PCIe slot 7 only), x16</li> <li>Slot 8 is blocked by double-wide GPU</li> </ul>

**Cisco HX C240 M6 Server 24 NVMe Drive Rear Panel, I/O Centric**

The Cisco HX C240 M6 24 NVMe version has a rear configuration option, for either I/O (I/O Centric) or Storage (Storage Centric) with the I/O Centric version of the node offering PCIe slots and the Storage Centric version of the node offering drive bays.

The following illustration shows the rear panel features for the I/O Centric version of the Cisco HX C240 M6SN.

- For features common to all versions of the node, see Common Rear Panel Features.
- For definitions of LED states, see [Rear-Panel LEDs, on page 65](#).



The following table shows the riser options for this version of the node.

1	Riser 1A	2	Riser 2A
3	Riser 3A, 3B, or 3C (not supported)	-	

**Table 3: Cisco HX C240 M6 24 SFF NVMe (UCSC-C240M6-SN)**


Riser	Options
Riser 1 This riser is I/O-centric and controlled by CPU 1 or CPU 2.	Riser 1A supports three PCIe slots: <ul style="list-style-type: none"> <li>• Slot 1 is full-height, 3/4 length, x8, NCSI</li> <li>• Slot 2 is full-height, full-length, x16, NCSI</li> <li>• Slot 3 is full-height, full-length, x8, no NCSI</li> </ul>
Riser 2A This riser is I/O-centric and controlled by CPU 2.	Riser 2A supports three PCIe slots: <ul style="list-style-type: none"> <li>• Slot 4 is full-height, 3/4 length, x8</li> <li>• Slot 5 is full-height, full-length, x16</li> <li>• Slot 6 is full-height, full-length, x8</li> </ul>
Riser 3	Riser 3A, 3B, and 3C are not supported.

**Cisco HX C240 M6 Server 12 NVMe Drive Rear Panel, I/O Centric**

The Cisco HX C240 M6 24 NVMe version has a rear configuration option, for either I/O (I/O Centric) or Storage (Storage Centric) with the I/O Centric version of the node offering PCIe slots and the Storage Centric version of the node offering drive bays.

The following illustration shows the rear panel features for the I/O Centric version of the Cisco HX C240 M6N.

- For features common to all versions of the node, see Common Rear Panel Features.
- For definitions of LED states, see [Rear-Panel LEDs, on page 65](#).



The following table shows the riser options for this version of the node.

<b>1</b>	Riser 1A	<b>2</b>	Riser 2A
<b>3</b>	Riser 3A, 3B, or 3C (Not Supported)	-	

Table 4: Cisco HX C240 M6 24 SFF NVMe (UCSC-C240M6-SN)


Riser	Options
Riser 1 This riser is I/O-centric and controlled by CPU 1 or CPU 2.	Riser 1A supports three PCIe slots: <ul style="list-style-type: none"> <li>• Slot 1 is full-height, 3/4 length, x8, NCSI</li> <li>• Slot 2 is full-height, full-length, x16, NCSI</li> <li>• Slot 3 is full-height, full-length, x8, no NCSI</li> </ul>
Riser 2A This riser is I/O-centric and controlled by CPU 2.	Riser 2A supports three PCIe slots: <ul style="list-style-type: none"> <li>• Slot 4 is full-height, 3/4 length, x8, NCSI</li> <li>• Slot 5 is full-height, full-length, x16, NCSI</li> <li>• Slot 6 is full-height, full-length, x8</li> </ul>
Riser 3	Riser 3A, 3B, and 3C are not supported.

**Cisco HX C240 M6 Server 24 Drive Rear Panel, Storage Centric**

The Cisco HX C240 M6 24 SAS/SATA SFF version has a rear configuration option, for either I/O (I/O Centric) or Storage (Storage Centric) with the I/O Centric version of the node offering PCIe slots and the Storage Centric version of the node offering drive bays.

The following illustration shows the rear panel features for the Storage Centric version of the Cisco HX C240 M6SX.

- For features common to all versions of the node, see Common Rear Panel Features.
- For definitions of LED states, see [Rear-Panel LEDs, on page 65](#).



The following table shows the riser options for this version of the node.

1	Riser 1B	2	Riser 2A (Not supported)
3	Riser 3B, 3C	-	

Table 5: Cisco HX C240 M6 24 SFF SAS/SATA/NVMe (UCSC-C240-M6SX)


Riser	Options
<p>Riser 1</p> <p>This riser is Storage-centric and controlled by CPU 2.</p>	<p>Riser 1B supports two SFF SAS/SATA/NVMe drives</p> <ul style="list-style-type: none"> <li>• Slot 1 is reserved</li> <li>• Slot 2 (drive bay 102), x4</li> <li>• Slot 3 (drive bay 101), x4</li> </ul> <p>When the node uses a hardware RAID controller card, SAS/SATA HDDs or SSDs, or NVMe PCIe SSDs are supported in the rear bays.</p>
<p>Riser 2</p> <p>This riser is I/O-centric and controlled by CPU 2.</p>	<p>Riser 2A is not supported for a Storage-centric version of the node.</p>
<p>Riser 3</p> <p>This riser is controlled by CPU 2.</p>	<p>Riser 3B has two PCIe slots that can support two SFF drives (NVMe).</p> <ul style="list-style-type: none"> <li>• Slot 7 (drive bay 104), x4</li> <li>• Slot 8 (drive bay 103), x4</li> </ul> <p>When the node uses a hardware RAID controller card, SAS/SATA HDDs or SSDs, or NVMe PCIe SSDs, are supported in the rear bays.</p> <p>Riser 3C has two PCIe slots that can support a GPU.</p> <ul style="list-style-type: none"> <li>• Slot 7 supports one full-height, full-length, double-wide GPU, x16</li> <li>• Slot 8 is blocked when a double-wide GPU is installed</li> </ul>

### Cisco HX C240 M6 Server 12 Drive Rear Panel, Storage Centric

The Cisco HX C240 M6 12 SAS/SATA SFF version has a rear configuration option, for either I/O (I/O Centric) or Storage (Storage Centric) with the I/O Centric version of the node offering PCIe slots and the Storage Centric version of the node offering drive bays.

The following illustration shows the rear panel features for the Storage Centric version of the Cisco HX C240 M6S.

- For features common to all versions of the node, see Common Rear Panel Features.
- For definitions of LED states, see [Rear-Panel LEDs, on page 65](#).



The following table shows the riser options for this version of the node.

1	Riser 1A	2	Riser 2A
3	Riser 3B, 3C	-	

**Table 6: Cisco HX C240 M6 12 SFF SAS/SATA (CSC-C240M6-S)**


Riser	Options
Riser 1	Riser 1A is not supported for a Storage Centric version of the node.
Riser 2	Riser 2A is not supported for a Storage Centric version of the node.
Riser 3 This riser is Storage-centric and controlled by CPU 2.	<p>Riser 3B can support two SFF drives (SAS/SATA/NVMe):</p> <ul style="list-style-type: none"> <li>Slot 7 (drive bay 104), x4</li> <li>Slot 8 (drive bay 103), x4</li> </ul> <p>When you are using a HWRAID controller, SAS/SATA HDDs or SSDs, or NVME PCIe SSDs, are supported in the rear bays.</p> <p>Riser 3C has two PCIe slots that can support a GPU.</p> <ul style="list-style-type: none"> <li>Slot 7 supports one full-height, full-length, double-wide GPU, x16</li> <li>Slot 8 is blocked when a double-wide GPU is installed</li> </ul>

**Cisco HX C240 M6 Server 24 NVMe Drive Rear Panel, Storage Centric**

The Cisco HX C240 M6 24 NVMe SFF version has a rear configuration option, for either I/O (I/O Centric) or Storage (Storage Centric) with the I/O Centric version of the node using PCIe slots and the Storage Centric version of the node offering drive bays.

The following illustration shows the rear panel features for the Storage Centric version of the Cisco HX C240 M6SN.

- For features common to all versions of the node, see Common Rear Panel Features.
- For definitions of LED states, see [Rear-Panel LEDs, on page 65](#).



The following table shows the riser options for this version of the node.

<b>1</b>	Riser 1B	<b>2</b>	Riser 2A (Not Supported)
<b>3</b>	Riser 3A, 3B, or 3C (Not Supported)	-	

**Table 7: Cisco HX C240 M6 24 SFF NVMe (UCSC-C240M6-SN)**


Riser	Options
Riser 1B This riser is Storage centric and controlled by CPU 2.	Riser 1B supports two NVMe drives: <ul style="list-style-type: none"> <li>• Slot 1 is reserved</li> <li>• Slot 2 (drive bay 102), x4</li> <li>• Slot 3 (drive bay 101), x4</li> </ul> When the node uses a hardware RAID controller card, NVMe PCIe SSDs are supported in the rear bays.
Riser 2	Riser 2A does not support storage devices.
Riser 3	Riser 3A, 3B, and 3C are not supported.

**Cisco HX C240 M6 Server 12 NVMe Drive Rear Panel, Storage Centric**

The Cisco HX C240 M6 12 NVMe SFF version has a rear configuration option, for either I/O (I/O Centric) or Storage (Storage Centric) with the I/O Centric version of the node using PCIe slots and the Storage Centric version of the node offering drive bays.

The following illustration shows the rear panel features for the Storage Centric version of the Cisco HX C240 M6N.

- For features common to all versions of the node, see Common Rear Panel Features.
- For definitions of LED states, see [Rear-Panel LEDs, on page 65](#).



The following table shows the riser options for this version of the node.

<b>1</b>	Riser 1B	<b>2</b>	Riser 2A (Not Supported)
<b>3</b>	Riser 3A, 3B, or 3C (Not Supported)	-	

**Table 8: Cisco HX C240 M6 12 SFF NVMe (UCSC-C240-M6N)**


Riser	Options
Riser 1B This riser is Storage-centric and controlled by CPU 2.	Riser 1B supports two NVME drives: <ul style="list-style-type: none"> <li>Slot 1 is reserved</li> <li>Slot 2 (drive bay 102), x4</li> <li>Slot 3 (drive bay 101), x4</li> </ul> When the node uses a hardware RAID controller card, NVMe PCIe SSDs are supported in the rear bays.
Riser 2	Riser 2A does not support storage devices.
Riser 3	Riser 3A, 3B, and 3C are not supported.

**Cisco HX C240 M6 Server 12 LFF Drive Rear Panel**

Unlike the SFF versions of the node, the Cisco HX C240 M6 LFF has one supported hardware configuration on the rear panel. The rear panel hardware configuration offers both PCIe slots and drive bays.

The following illustration shows the rear panel features for the I/O Centric version of the Cisco HX C240 LFF.

- For features common to all versions of the node, see Common Rear Panel Features.
- For definitions of LED states, see [Rear-Panel LEDs, on page 65](#).



The following table shows the riser options for this version of the node.

1	Riser 1B	2	Riser 2A
3	Riser 3B	-	

Table 9: Cisco HX C240 M6 12 LFF (HX-C240-LFF)

Riser	Options
<p>Riser 1</p> <p>This riser is controlled by CPU 1.</p>	<p>Riser 1B supports three PCIe slots numbered bottom to top.</p> <ul style="list-style-type: none"> <li>Slot 1 is reserved for a drive controller.</li> <li>Slot 2 supports one drive (drive bay 102), x4</li> <li>Slot 3 supports one drive (drive bay 101), x4</li> </ul> <p>When using a hardware RAID controller card or SAS</p> <p>When using a hardware RAID controller card or SAS HBA in the node, SAS/SATA HDDs or SSDs are supported in the rear bays.</p> <p>NVMe PCIe SSDs are supported in the rear bays without need for a RAID controller.</p>
<p>Riser 2</p> <p>This riser is controlled by CPU 2.</p>	<p>Riser 2A supports three PCIe slots:</p> <ul style="list-style-type: none"> <li>Slot 4 is full-height, 3/4 length, x8</li> <li>Slot 5 is full-height, full-length, x16</li> <li>Slot 6 is full-height, full-length, x8</li> </ul>



Riser	Options
Riser 3 This riser is controlled by CPU 2.	Riser 3B supports two drives: <ul style="list-style-type: none"><li>• Slot 7 (drive bay 104), x4</li><li>• Slot 8 (drive bay 103), x4</li></ul> When the node uses a hardware RAID controller card or SAS HBA, SAS/SATA HDDs or SSDs are supported in the rear bays. NVMe PCIe SSDs are supported in the rear bays without need for a RAID controller


## PCIe Risers

The following different PCIe riser options are available.


### Riser 1 Options

This riser supports two options, Riser 1A and 1B.

PCIe Riser 1A (outside)




PCIe Riser 1A (inside)




1	PCIe slot 1, Full height, $\frac{3}{4}$ length, x8, NCSI	2	PCIe slot 2, full height, full length, x16, NCSI, GPU capable
3	PCIe slot 3, full height, full length, x8, no NCSI	4	Edge Connectors

PCIe Riser 1B (outside)



PCIe Riser 1B (inside)




309593

1	PCIe slot 1, Reserved for drive controller	2	Drive Bay 102, x4
3	Drive 103, x4	4	Edge Connectors


**Riser 2**

This riser supports one option, Riser 2A.

PCIe Riser 2A (outside)



PCIe Riser 2A (inside)




1	PCIe slot 4, Full height, $\frac{3}{4}$ length, x8	2	PCIe slot 5, full height, full length, x16
---	--	---	--

3	PCIe slot 6, full height, full length, x16	4	Edge Connectors
---	--	---	-----------------


**Riser 3**

This riser supports three options, 3A, 3B, and 3C.

PCIe Riser 3A (outside)




PCIe Riser 3A (inside)




300596

<b>1</b>	PCIe slot 7, Full height, full length x8	<b>2</b>	PCIe slot 8, full height, full length, x16
<b>3</b>	Edge Connectors		

PCIe Riser 3B (outside)




PCIe Riser 3B (inside)




<b>1</b>	PCIe Slot 7, Drive Bay 103, x4	<b>2</b>	PCIe Slot 8, Drive Bay 103, x4
<b>3</b>	Edge Connectors		

PCIe Riser 3C (outside)



PCIe Riser 3C (inside)



1	PCIe Slot 7, supports one full height, full length, double-wide GPU (slot 7 only), x16	2	Edge Connectors
---	--	---	-----------------

## Summary of Node Features

The following tables list a summary of the node features for the LFF version and SFF versions of the server.

Table 10: Server Features, SFF

Feature	Description
Chassis	Two rack-unit (2RU) chassis
Central Processor	One or two 3rd Generation Intel Xeon processors.
Chipset	Intel® C621 series chipset
Memory	32 slots for registered DIMMs (RDIMMs) or load-reduced DIMMs (LR DIMMs) and support for Intel® Optane™ Persistent Memory Modules (DCPMMs)
Multi-bit error protection	Multi-bit error protection is supported
Video	<p>The Cisco Integrated Management Controller (CIMC) provides video using the Matrox G200e video/graphics controller:</p> <ul style="list-style-type: none"> <li>• Integrated 2D graphics core with hardware acceleration</li> <li>• Embedded DDR memory interface supports up to 512 MB of addressable memory (8 MB is allocated by default to video memory)</li> <li>• Supports display resolutions up to 1920 x 1200 16bpp @ 60Hz</li> <li>• High-speed integrated 24-bit RAMDAC</li> <li>• Single lane PCI-Express host interface running at Gen 1 speed</li> </ul>
Network and management I/O	<p>Rear panel:</p> <ul style="list-style-type: none"> <li>• One 1-Gb Ethernet dedicated management port (RJ-45 connector)</li> <li>• Two 1-Gb/10-Gb BASE-T Ethernet LAN ports (RJ-45 connectors)</li> </ul> <p>The dual LAN ports can support 1 Gbps and 10 Gbps, depending on the link partner capability.</p> <ul style="list-style-type: none"> <li>• One RS-232 serial port (RJ-45 connector)</li> <li>• One VGA video connector port (DB-15 connector)</li> <li>• Two USB 3.0 ports</li> </ul> <p>Front panel:</p> <ul style="list-style-type: none"> <li>• One front-panel keyboard/video/mouse (KVM) connector that is used with the KVM breakout cable. The breakout cable provides two USB 2.0, one VGA, and one DB-9 serial connector.</li> </ul>



Feature	Description
Power	<p>Up to two of the following hot-swappable power supplies:</p> <ul style="list-style-type: none"> <li>• 1050 W (AC)</li> <li>• 1050 W (DC)</li> <li>• 1600 W (AC)</li> <li>• 2300 W (AC)</li> </ul> <p>One power supply is mandatory; one more can be added for 1 + 1 redundancy as long power supplies are the same.</p> <p>For additional information, see <a href="#">Supported Power Supplies, on page 124</a></p>
ACPI	<p>The advanced configuration and power interface (ACPI) 4.0 standard is supported.</p>
Front Panel	<p>The front panel controller provides status indications and control buttons</p>
Cooling	<p>Six hot-swappable fan modules for front-to-rear cooling.</p>
InfiniBand	<p>The PCIe bus slots in this server support the InfiniBand architecture.</p>
Expansion Slots	<p>For the SFF versions of the server, three half-height riser slots are supported:</p> <ul style="list-style-type: none"> <li>• Riser 1A (3 PCIe slots)</li> <li>• Riser 1B (2 drive bays)</li> <li>• Riser 2A (3 PCIe slots)</li> <li>• Riser 3A (2 PCIe slots)</li> <li>• Riser 3B (2 drive bays)</li> <li>• Riser 3C (1 full-length, double-wide GPU)</li> </ul> <p><b>Note</b> Not all risers are available in every server configuration option.</p> <p>One or two dedicated slots (depending on the server type) for a SATA interposer or storage controller(s).</p>

Feature	Description
Interfaces	<p>Rear panel:</p> <ul style="list-style-type: none"><li>• One 1Gbase-T RJ-45 management port</li><li>• Two 10Gbase-T LOM ports</li><li>• One RS-232 serial port (RJ45 connector)</li><li>• One DB15 VGA connector</li><li>• Two USB 3.0 port connectors</li><li>• One flexible modular LAN on motherboard (mLOM) slot that can accommodate various interface cards</li></ul> <p>Front panel supports one KVM console connector that supplies:</p> <ul style="list-style-type: none"><li>• two USB 2.0 connectors,</li><li>• one VGA DB15 video connector</li><li>• one serial port (RS232) RJ45 connector</li></ul>

Feature	Description
Internal Storage Devices	

Feature	Description
	<ul style="list-style-type: none"> <li>• UCSC-C240-M6S:               <ul style="list-style-type: none"> <li>• Up to 12 SFF SAS/SATA hard drives (HDDs) or SAS/SATA solid state drives (SSDs)</li> <li>• Optionally, up to four SFF NVMe PCIe SSDs. These drives must be placed in front drive bays 1, 2, 3, and 4 only, can be mixed with SAS/SATA drives. The rest of the bays (5 - 12) can be populated with SAS/SATA SSDs or HDDs. Two CPUs are required in a server that has any number of NVMe drives.</li> <li>• Optionally, one front-facing DVD drive</li> <li>• Optionally, up to two SFF rear-facing SAS/SATA/NVMe drives</li> <li>• If using a SATA Interposer, up to 8 SATA-only drives can be installed (slots 1-8 only).</li> </ul> </li> <li>• UCSC-C240-M6SX:               <ul style="list-style-type: none"> <li>• Up to 24 front SFF SAS/SATA hard drives (HDDs) or SAS/SATA solid state drives (SSDs).</li> <li>• Optionally, up to four front SFF NVMe PCIe SSDs. These drives must be placed in front drive bays 1, 2, 3, and 4 only. The rest of the bays (5 - 24) can be populated with SAS/SATA SSDs or HDDs. Two CPUs are required in a server that has any number of NVMe drives.</li> <li>• Optionally, up to four SFF rear-facing SAS/SATA/NVMe drives</li> </ul> </li> <li>• UCSC-C240-M6N:               <ul style="list-style-type: none"> <li>• Up to 12 front NVMe (only) drives</li> <li>• Optionally, up to 2 rear NVMe (only) drives</li> <li>• Two CPUs are required in a server that has any number of NVMe drives.</li> </ul> </li> <li>• UCSC-C240-M6SN:               <ul style="list-style-type: none"> <li>• Up to 24 front NVMe drives (only).</li> <li>• Optionally, up to 2 rear NVMe drives (only)</li> <li>• Two CPUs are required when choosing NVMe SSDs</li> </ul> </li> <li>• Other Storage:               <ul style="list-style-type: none"> <li>• A mini-storage module connector on the motherboard</li> </ul> </li> </ul>

Feature	Description
	<p>supports a boot-optimized RAID controller carrier that holds up two SATA M.2 SSDs. Mixing different capacity SATA M.2 SSDs is not supported.</p> <ul style="list-style-type: none"> <li>• Optional 2 M.2 RAID cards for use as boot volumes.</li> </ul>
Integrated Management Processor	<p>Baseboard Management Controller (BMC) running Cisco Integrated Management Controller (CIMC) firmware.</p> <p>Depending on your CIMC settings, the CIMC can be accessed through the 1GE dedicated management port, the 1GE/10GE LOM ports, or a Cisco virtual interface card (VIC).</p> <p>CIMC manages certain components within the server, such as the Cisco 12G SAS HBA.</p>
Storage Controllers	<p>One SATA Interposer board, 12G RAID HBA, or one or two 12G SAS HBAs plug into a dedicated slot.</p> <ul style="list-style-type: none"> <li>• SATA Interposer board: <ul style="list-style-type: none"> <li>• AHCI support for up to eight SATA-only drives (slots 1- 8)</li> <li>• Supported only on the UCSC-C240M6-S server</li> </ul> </li> <li>• Cisco 12G RAID controller with 4GB FBWC (for UCSC-240-M6S server) <ul style="list-style-type: none"> <li>• RAID support (RAID 0, 1, 5, 6, 10, 50 and 60) and SW RAID0</li> <li>• Supports up to 14 internal SAS/SATA drives</li> </ul> </li> <li>• Cisco M6 12G SAS RAID controller with 4GB FBWC (for UCSC-240-M6SX server) <ul style="list-style-type: none"> <li>• RAID support (RAID 0, 1, 5, 6, 10, 50, and 60) and SRAID0</li> <li>• Supports up to 28 internal SAS/SATA drives</li> </ul> </li> <li>• Cisco M6 12G SAS HBA (for UCSC-240-M6S and UCSC-240-M6SX servers) <ul style="list-style-type: none"> <li>• RAID Support 0, 1, and 10</li> <li>• JBOD/Pass-through Mode support</li> <li>• Each HBA supports up to 14 SAS/SATA internal drives</li> </ul> </li> </ul>

Feature	Description
Modular LAN over Motherboard (mLOM) slot	<p>The dedicated mLOM slot on the motherboard can flexibly accommodate the following cards:</p> <ul style="list-style-type: none"> <li>• Cisco Virtual Interface Cards (VICs)</li> <li>• Quad Port Intel i350 1GbE RJ45 Network Interface Card (NIC)</li> </ul> <p><b>Note</b> The four Intel i350 ports are provided on an optional card that plugs into the mLOM slot. These ports are separate from the two embedded LAN ports on the motherboard.</p>
Node Management	Cisco Intersight provides server management.
CIMC	Cisco Integrated Management Controller (CIMC) 4.2(1) or later is required for the server.

Table 11: Node Features, LFF

Feature	Description
Chassis	Two rack-unit (2RU) chassis
Central Processor	One or two 3rd Generation Intel Xeon processors.
Chipset	Intel® C621 series chipset
Memory	32 slots for registered DIMMs (RDIMMs) or load-reduced DIMMs (LR DIMMs) and support for Intel® Optane™ Persistent Memory Modules (DCPMMs)
Multi-bit error protection	Multi-bit error protection is supported
Video	<p>The Cisco Integrated Management Controller (CIMC) provides video using the Matrox G200e video/graphics controller:</p> <ul style="list-style-type: none"> <li>• Integrated 2D graphics core with hardware acceleration</li> <li>• DDR2/3 memory interface supports up to 512 MB of addressable memory (8 MB is allocated by default to video memory)</li> <li>• Supports display resolutions up to 1920 x 1200 16bpp @ 60Hz</li> <li>• High-speed integrated 24-bit RAMDAC</li> <li>• Single lane PCI-Express host interface running at Gen 1 speed</li> </ul>

Feature	Description
Network and management I/O	<p>Rear panel:</p> <ul style="list-style-type: none"> <li>• One 1Gbase-T RJ-45 management port</li> <li>• Two 10Gbase-T LOM ports</li> <li>• One RS-232 serial port (RJ-45 connector)</li> <li>• One DB15 VGA connector</li> <li>• Two USB 3.0 ports</li> <li>• One flexible modular LAN on motherboard (mLOM) slot that can accommodate various interface cards</li> </ul> <p>Front panel supports one KVM console connector that supplies:</p> <ul style="list-style-type: none"> <li>• two USB 2.0 connectors,</li> <li>• one VGA DB15 video connector</li> <li>• one serial port (RS232) RJ45 connector</li> </ul>
Power	<p>Up to two of the following hot-swappable power supplies:</p> <ul style="list-style-type: none"> <li>• 1050 W (AC)</li> <li>• 1050 W (DC)</li> <li>• 1600 W (AC)</li> <li>• 2300 W (AC)</li> </ul> <p>One power supply is mandatory; one more can be added for 1 + 1 redundancy. With two power supplies, both must be the same type and wattage.</p>
ACPI	The advanced configuration and power interface (ACPI) 4.0 standard is supported.
Front Panel	The front panel controller provides status indications and control buttons
Cooling	Six hot-swappable fan modules for front-to-rear cooling.
InfiniBand	The PCIe bus slots in this server support the InfiniBand architecture.
Expansion Slots	<ul style="list-style-type: none"> <li>• Riser 1B (1 PCIe slot reserved for a drive controller and 2 HDD slots)</li> <li>• Riser 2A (3 PCIe slots)</li> <li>• Riser 3B (2 HDD slots)</li> </ul>

Feature	Description
Interfaces	<p>Rear panel:</p> <ul style="list-style-type: none"> <li>• One 1Gbase-T RJ-45 management port</li> <li>• Two 10Gbase-T LOM ports</li> <li>• One RS-232 serial port (RJ45 connector)</li> <li>• One DB15 VGA connector</li> <li>• Two USB 3.0 port connectors</li> <li>• One flexible modular LAN on motherboard (mLOM) slot that can accommodate various interface cards</li> </ul> <p>Front panel: Supports one KVM console connector that supplies two USB 2.0 connectors, one VGA DB15 video connector, and one serial port (RS232) RJ45 connector</p>
Internal Storage Devices	<ul style="list-style-type: none"> <li>• Large Form Factor (LFF) drives with a 12-drive backplane. The server can hold a maximum of: <ul style="list-style-type: none"> <li>• 12 LFF 3.5-inch front-loading SAS-only hard drives (HDDs, 4 mid-plane LFF drives,</li> <li>• As an option, up to four 3.5-inch mid-plane SAS-only LFF HDDs</li> <li>• Optionally, up to four rear-facing SAS/SATA HDDs/SSDs or up to four rear-facing NVMe PCIe SSDs</li> </ul> </li> <li>• A mini-storage module connector on the motherboard supports a boot-optimized RAID controller carrier that holds up to two SATA M.2 SSDs. Mixing different capacity SATA M.2 SSDs is not supported.</li> </ul>




Feature	Description
Storage Controllers	<p>The 12G RAID HBA or 12G SAS HBA plugs into slot 1 (bottom slot) of riser 1B.</p> <ul style="list-style-type: none"> <li>• Cisco M6 12G SAS RAID Controller with 4GB FBWC <ul style="list-style-type: none"> <li>• RAID support (RAID 0, 1, 5, 6, 10, 50 and 60) and SRAID</li> <li>• Supports up to 32 internal SAS/SATA drives</li> <li>• Plugs into drive slot 1 of riser 1B</li> </ul> </li> <li>• Cisco M6 12G SAS HBA <ul style="list-style-type: none"> <li>• RAID 0, 1, and 10 support</li> <li>• JBOD/Pass-through Mode support</li> <li>• Supports up to 32 SAS/SATA internal drives</li> <li>• Plugs into slot 1 of riser 1B</li> </ul> </li> </ul>
Integrated Management Processor	<p>Baseboard Management Controller (BMC) running Cisco Integrated Management Controller (CIMC) firmware.</p> <p>Depending on your CIMC settings, the CIMC can be accessed through the 1GE dedicated management port, the 1GE/10GE LOM ports, or a Cisco virtual interface card (VIC).</p> <p>CIMC manages certain components within the server, such as the Cisco 12G SAS HBA.</p>
Modular LAN over Motherboard (mLOM) slot	<p>The dedicated mLOM slot on the motherboard can flexibly accommodate the following cards:</p> <ul style="list-style-type: none"> <li>• Cisco Virtual Interface Cards (VICs)</li> <li>• Quad Port Intel i350 1GbE RJ45 Network Interface Card (NIC)</li> </ul> <p><b>Note</b> The four Intel i350 ports are provided on an optional card that plugs into the mLOM slot. These ports are separate from the two LAN ports embedded on the motherboard.</p>
Node Management	Cisco Intersight provides node management.
CIMC	Cisco Integrated Management Controller 4.2(1) or later is required for this node.

# Serviceable Component Locations

This topic shows the locations of the field-replaceable components and service-related items. The view in the following figure shows the server with the top cover removed.

Figure 7: Cisco HX C240 M6 Server, Serviceable Component Locations



1	Front-loading drive bays.	2	Cooling fan modules (six, hot-swappable)
3	DIMM sockets on motherboard (16 per CPU) See <a href="#">DIMM Population Rules and Memory Performance Guidelines</a> , on page 115 for DIMM slot numbering.  <b>Note</b> An air baffle rests on top of the DIMM and CPUs when the server is operating. The air baffle is not displayed in this illustration.	4	CPU socket 1
5	CPU socket 2	6	M.2 RAID Controller

7	<p>PCIe riser 3 (PCIe slots 7 and 8 numbered from bottom to top), with the following options:</p> <ul style="list-style-type: none"> <li>• 3A (Default Option)—Slots 7 (x16 mechanical, x8 electrical), and 8 (x16 mechanical, x8 electrical). Both slots can accept a full height, full length GPU card.</li> <li>• 3B (Storage Option)—Slots 7 (x24 mechanical, x4 electrical) and 8 (x24 mechanical, x4 electrical). Both slots can accept 2.5-inch SFF universal HDDs.</li> <li>• 3C (GPU Option)—Slots 7 (x16 mechanical, x16 electrical) and 8 empty (NCSI support limited to one slot at a time). Slot 7 can support a full height, full length GPU card.</li> </ul>	8	<p>PCIe riser 2 (PCIe slots 4, 5, 6 numbered from bottom to top), with the following options:</p> <ul style="list-style-type: none"> <li>• 2A (Default Option)—Slot 4 (x24 mechanical, x8 electrical) supports full height, <math>\frac{3}{4}</math> length card; Slot 5 (x24 mechanical, x16 electrical) supports full height, full length GPU card; Slot 6 (x16 mechanical, x8 electrical) supports full height, full length card.</li> </ul>
9	<p>PCIe riser 1 (PCIe slot 1, 2, 3 numbered bottom to top), with the following options:</p> <ul style="list-style-type: none"> <li>• 1A (Default Option)—Slot 1 (x24 mechanical, x8 electrical) supports full height, <math>\frac{3}{4}</math> length card; Slot 2 (x24 mechanical, x16 electrical) supports full height, full length GPU card; Slot 3 (x16 mechanical, x8 electrical) supports full height, full length card.</li> <li>• 1B (Storage Option)—Slot 1 (x24 mechanical, x8 electrical) supports full height, <math>\frac{3}{4}</math> length card; Slot 2 (x4 electrical), supports 2.5-inch SFF universal HDD; Slot 3 (x4 electrical), supports 2.5-inch SFF universal HDD</li> </ul>	-	





## CHAPTER 2

# Installing the Node

---

- [Preparing for Installation, on page 43](#)
- [Installing the Node in a Rack, on page 46](#)
- [Initial Server Setup, on page 51](#)
- [NIC Mode and NIC Redundancy Settings, on page 56](#)
- [Updating the BIOS and Cisco IMC Firmware, on page 57](#)
- [Accessing the System BIOS, on page 58](#)
- [Smart Access Serial, on page 58](#)
- [Smart Access USB, on page 59](#)

## Preparing for Installation

This section contains the following topics:

### Installation Warnings and Guidelines



**Note**

---

Before you install, operate, or service a node, review the [Regulatory Compliance and Safety Information for Cisco UCS C-Series Servers](#) for important safety information.

---



**Warning**

**IMPORTANT SAFETY INSTRUCTIONS**

**This warning symbol means danger. You are in a situation that could cause bodily injury. Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents. Use the statement number provided at the end of each warning to locate its translation in the translated safety warnings that accompanied this device.**

Statement 1071

---



---

**Warning** To prevent the system from overheating, do not operate it in an area that exceeds the maximum recommended ambient temperature of: 35° C (95° F).

Statement 1047

---



---

**Warning** The plug-socket combination must be accessible at all times, because it serves as the main disconnecting device.

Statement 1019

---



---

**Warning** This product relies on the building's installation for short-circuit (overcurrent) protection. Ensure that the protective device is rated not greater than: 250 V, 15 A.

Statement 1005

---



---

**Warning** Installation of the equipment must comply with local and national electrical codes.

Statement 1074

---



---

**Warning** This unit is intended for installation in restricted access areas. A restricted access area can be accessed only through the use of a special tool, lock, and key, or other means of security.

Statement 1017

---



---

**Caution** To ensure proper airflow it is necessary to rack the servers using rail kits. Physically placing the units on top of one another or “stacking” without the use of the rail kits blocks the air vents on top of the servers, which could result in overheating, higher fan speeds, and higher power consumption. We recommend that you mount your servers on rail kits when you are installing them into the rack because these rails provide the minimal spacing required between the servers. No additional spacing between the servers is required when you mount the units using rail kits.

---



---

**Caution** Avoid uninterruptible power supply (UPS) types that use ferroresonant technology. These UPS types can become unstable with systems such as the Cisco UCS, which can have substantial current draw fluctuations from fluctuating data traffic patterns.

---

When you are installing a node, use the following guidelines:

- Plan your site configuration and prepare the site before installing the node. See the Cisco Site Preparation Guide for the recommended site planning tasks.

- Ensure that there is adequate space around the node to allow for accessing the node and for adequate airflow. The airflow in this node is from front to back.
- Ensure that the air-conditioning meets the thermal requirements listed in the [Environmental Specifications, on page 175](#).
- Ensure that the cabinet or rack meets the requirements listed in the [Rack Requirements, on page 45](#).
- Ensure that the site power meets the power requirements listed in the [Power Specifications, on page 176](#). If available, you can use an uninterruptible power supply (UPS) to protect against power failures.

## Rack Requirements

The rack must be of the following type:

- A standard 19-in. (48.3-cm) wide, four-post EIA rack, with mounting posts that conform to English universal hole spacing, per section 1 of ANSI/EIA-310-D-1992.
- The rack-post holes can be square 0.38-inch (9.6 mm), round 0.28-inch (7.1 mm), #12-24 UNC, or #10-32 UNC when you use the Cisco-supplied slide rails.
- The minimum vertical rack space per node must be two rack units (RUs), equal to 3.5 in. (88.9 mm).

### Supported Cisco Slide Rail Kits

The node supports the following rail kit options:

- Cisco part UCSC-RAIL-M6= (Ball bearing rail kit for C220 & C240 M6 rack servers)
- Cisco part UCSC-CMA-C240M6= (Reversible CMA for ball bearing rail kit)

### Rack Installation Tools Required

The slide rails sold by Cisco Systems for this node do not require tools for installation.

### Slide Rail and Cable Management Arm Dimensions

The slide rails for this node have an adjustment range of 24 to 36 inches (610 to 914 mm).

The optional cable management arm (CMA) adds additional length requirements:

- The additional distance from the rear of the node to the rear of the CMA is 5.4 inches (137.4 mm).
- The total length of the node including the CMA is 35.2 inches (894 mm).

# Installing the Node in a Rack



**Warning** To prevent bodily injury when mounting or servicing this unit in a rack, you must take special precautions to ensure that the system remains stable. The following guidelines are provided to ensure your safety:

This unit should be mounted at the bottom of the rack if it is the only unit in the rack.


When mounting this unit in a partially filled rack, load the rack from the bottom to the top with the heaviest component at the bottom of the rack.

If the rack is provided with stabilizing devices, install the stabilizers before mounting or servicing the unit in the rack.

Statement 1006

**Step 1** Attach the inner rails to the sides of the node:

- a) Align an inner rail with one side of the node so that the three keyed slots in the rail align with the three pegs on the side of the node.
- b) Set the keyed slots over the pegs, and then slide the rail toward the front to lock it in place on the pegs. The front slot has a metal clip that locks over the front peg.



1	Keyed installation slots	2	Front of node
---	--------------------------	---	---------------


- c) Install the second inner rail to the opposite side of the node.

**Step 2** Open the front securing plate on both slide-rail assemblies. The front end of the slide-rail assembly has a spring-loaded securing plate that must be open before you can insert the mounting pegs into the rack-post holes.

On the *outside* of the assembly, push the green-arrow button toward the rear to open the securing plate.



Figure 8: Front Securing Mechanism, Inside of Front End



1	Front mounting pegs	3	Securing plate shown pulled back to the open position
2	Rack post between mounting pegs and opened securing plate	-	

**Step 3** Install the outer slide rails into the rack:

- a) Align one slide-rail assembly front end with the front rack-post holes that you want to use.

The slide rail front-end wraps around the outside of the rack post and the mounting pegs enter the rack-post holes from the outside-front. See [Figure 8: Front Securing Mechanism, Inside of Front End](#), on page 47.

**Note** The rack post must be between the mounting pegs and the *open* securing plate.


- b) Push the mounting pegs into the rack-post holes from the outside-front.
- c) Press the securing plate release button, marked PUSH. The spring-loaded securing plate closes to lock the pegs in place.
- d) Adjust the slide-rail length, and then push the rear mounting pegs into the corresponding rear rack-post holes. The slide rail must be level front-to-rear.
- The rear mounting pegs enter the rear rack-post holes from the *inside* of the rack post.
- e) Attach the second slide-rail assembly to the opposite side of the rack. Ensure that the two slide-rail assemblies are at the same height and are level front-to-back.
- f) Pull the inner slide rails on each assembly out toward the rack front until they hit the internal stops and lock in place.

**Step 4** Insert the node into the slide rails:

**Caution** This node can weigh up to 64 pounds (29 kilograms) when fully loaded with components. We recommend that you use a minimum of two people or a mechanical lift when lifting the node. Attempting this procedure alone could result in personal injury or equipment damage.


- a) Align the rear ends of the inner rails that are attached to the node sides with the front ends of the empty slide rails on the rack.

b) Push the inner rails into the slide rails on the rack until they stop at the internal stops



c) Slide the inner-rail release clip toward the rear on both inner rails, and then continue pushing the node into the rack until its front slam-latches engage with the rack posts.

Figure 9: Inner-Rail Release Clip




1	Inner-rail release clip	3	Outer slide rail attached to rack post
2	Inner rail attached to node and inserted into outer slide rail	-	

**Step 5**

(Optional) Secure the node in the rack more permanently by using the two screws that are provided with the slide rails. Perform this step if you plan to move the rack with servers installed.

With the node fully pushed into the slide rails, open a hinged slam latch lever on the front of the node and insert a screw through the hole that is under the lever. The screw threads into the static part of the rail on the rack post and prevents the node from being pulled out. Repeat for the opposite slam latch.




## Installing the Cable Management Arm (Optional)



**Note** The cable management arm (CMA) is reversible left-to-right. To reverse the CMA, see [Reversing the Cable Management Arm \(Optional\)](#), on page 50 before installation.

**Step 1** With the node pushed fully into the rack, slide the CMA tab of the CMA arm that is farthest from the node onto the end of the stationary slide rail that is attached to the rack post. Slide the tab over the end of the rail until it clicks and locks.

Figure 10: Attaching the CMA to the Rear Ends of the Slide Rails



1	CMA tab on arm farthest from node attaches to end of stationary outer slide rail.	3	CMA tab on width-adjustment slider attaches to end of stationary outer slide rail.
2	CMA tab on arm closest to the node attaches to end of inner slide rail attached to node.	4	Rear of node


- Step 2** Slide the CMA tab that is closest to the node over the end of the inner rail that is attached to the node. Slide the tab over the end of the rail until it clicks and locks
- Step 3** Pull out the width-adjustment slider that is at the opposite end of the CMA assembly until it matches the width of your rack.
- Step 4** Slide the CMA tab that is at the end of the width-adjustment slider onto the end of the stationary slide rail that is attached to the rack post. Slide the tab over the end of the rail until it clicks and locks.
- Step 5** Open the hinged flap at the top of each plastic cable guide and route your cables through the cable guides as desired.

## Reversing the Cable Management Arm (Optional)

- Step 1** Rotate the entire CMA assembly 180 degrees, left-to-right. The plastic cable guides must remain pointing upward.
- Step 2** Flip the tabs at the ends of the CMA arms so that they point toward the rear of the node.

- Step 3** Pivot the tab that is at the end of the width-adjustment slider. Depress and hold the metal button on the outside of the tab and pivot the tab 180 degrees so that it points toward the rear of the node.

**Figure 11: Reversing the CMA**



1	CMA tab on end of width-adjustment slider	2	Metal button on outside of tab
---	---	---	--------------------------------

## Initial Server Setup



**Note** This section describes how to power on the server, assign an IP address, and connect to server management when using the node in standalone mode.

### Server Default Settings

The node is shipped with these default settings:

- The NIC mode is *Shared LOM EXT*.

Shared LOM EXT mode enables the 1-Gb/10-Gb Ethernet ports *and* the ports on any installed Cisco virtual interface card (VIC) to access the Cisco Integrated Management Interface (Cisco IMC). If you want to use the 10/100/1000 dedicated management ports to access Cisco IMC, you can connect to the node and change the NIC mode as described in [Setting Up the System With the Cisco IMC Configuration Utility, on page 54](#).

- The NIC redundancy is *Active-Active*. All Ethernet ports are utilized simultaneously.
- DHCP is enabled.
- IPv4 is enabled.

### Connection Methods

There are two methods for connecting to the system for initial setup:

- Local setup—Use this procedure if you want to connect a keyboard and monitor directly to the system for setup. This procedure can use a KVM cable (Cisco PID N20-BKVM) or the ports on the rear of the node.
- Remote setup—Use this procedure if you want to perform setup through your dedicated management LAN.




---

**Note** To configure the system remotely, you must have a DHCP node on the same network as the system. Your DHCP node must be preconfigured with the range of MAC addresses for this node. The MAC address is printed on a label that is on the pull-out asset tag on the front panel. This node has a range of six MAC addresses assigned to the Cisco IMC. The MAC address printed on the label is the beginning of the range of six contiguous MAC addresses.

---

This section contains the following topics:

## Connecting to the Server Locally For Setup

This procedure requires the following equipment:

- VGA monitor
- USB keyboard
- Either the supported Cisco KVM cable (Cisco PID N20-BKVM); or a USB cable and VGA DB-15 cable

---

**Step 1** Attach a power cord to each power supply in your node, and then attach each power cord to a grounded AC power outlet.

If you are using DC power supplies, see [Installing DC Power Supplies \(First Time Installation\), on page 127](#).

Wait for approximately two minutes to let the node boot to standby power during the first bootup. You can verify system power status by looking at the system Power Status LED on the front panel. The system is in standby power mode when the LED is amber.

**Step 2** Connect a USB keyboard and VGA monitor to the node using one of the following methods:

- Connect an optional KVM cable (Cisco PID N20-BKVM) to the KVM connector on the front panel. Connect your USB keyboard and VGA monitor to the KVM cable.
- Connect a USB keyboard and VGA monitor to the corresponding connectors on the rear panel.

**Step 3** Open the Cisco IMC Configuration Utility:

- Press and hold the front panel power button for four seconds to boot the node.
- During bootup, press **F8** when prompted to open the Cisco IMC Configuration Utility.

**Note** The first time that you enter the Cisco IMC Configuration Utility, you are prompted to change the default password. The default password is *password*. The Strong Password feature is enabled.

The following are the requirements for Strong Password:

- The password can have a minimum of 8 characters and a maximum of 14 characters.
- The password must not contain the user's name.
- The password must contain characters from three of the following four categories:
  - English uppercase letters (A through Z)
  - English lowercase letters (a through z)
  - Base 10 digits (0 through 9)
  - Non-alphabetic characters !, @, #, \$, %, ^, &, \*, -, \_, =, “

**Step 4** Continue with [Setting Up the System With the Cisco IMC Configuration Utility, on page 54](#).

---

## Connecting to the Server Remotely For Setup

This procedure requires the following equipment:

- One RJ-45 Ethernet cable that is connected to your management LAN.

### Before you begin



**Note** To configure the system remotely, you must have a DHCP server on the same network as the system. Your DHCP server must be preconfigured with the range of MAC addresses for this server node. The MAC address is printed on a label that is on the pull-out asset tag on the front panel. This server node has a range of six MAC addresses assigned to the Cisco IMC. The MAC address printed on the label is the beginning of the range of six contiguous MAC addresses.

---

**Step 1** Attach a power cord to each power supply in your node, and then attach each power cord to a grounded AC power outlet.

If you are using DC power supplies, see [Installing DC Power Supplies \(First Time Installation\), on page 127](#).

Wait for approximately two minutes to let the node boot to standby power during the first bootup. You can verify system power status by looking at the system Power Status LED on the front panel. The system is in standby power mode when the LED is amber.

**Step 2** Plug your management Ethernet cable into the dedicated management port on the rear panel.

**Step 3** Allow your preconfigured DHCP node to assign an IP address to the node node.

**Step 4** Use the assigned IP address to access and log in to the Cisco IMC for the node node. Consult with your DHCP node administrator to determine the IP address.

**Note** The default username for the node is *admin*. The default password is *password*.

**Step 5** From the Cisco IMC Server Summary page, click **Launch KVM Console**. A separate KVM console window opens.

**Step 6** From the Cisco IMC Summary page, click **Power Cycle Server**. The system reboots.

**Step 7** Select the KVM console window.

**Note** The KVM console window must be the active window for the following keyboard actions to work.

**Step 8** When prompted, press **F8** to enter the Cisco IMC Configuration Utility. This utility opens in the KVM console window.

**Note** The first time that you enter the Cisco IMC Configuration Utility, you are prompted to change the default password. The default password is *password*. The Strong Password feature is enabled.

The following are the requirements for Strong Password:

- The password can have a minimum of 8 characters and a maximum of 14 characters.
- The password must not contain the user's name.
- The password must contain characters from three of the following four categories:
  - English uppercase letters (A through Z)
  - English lowercase letters (a through z)
  - Base 10 digits (0 through 9)
  - Non-alphabetic characters !, @, #, \$, %, ^, &, \*, -, \_, =, “

**Step 9** Continue with [Setting Up the System With the Cisco IMC Configuration Utility, on page 54](#).

## Setting Up the System With the Cisco IMC Configuration Utility

### Before you begin

The following procedure is performed after you connect to the system and open the Cisco IMC Configuration Utility.

**Step 1** Set the NIC mode to choose which ports to use to access Cisco IMC for node management:

- *Shared LOM EXT* (default)—This is the shared LOM extended mode, the factory-default setting. With this mode, the Shared LOM and Cisco Card interfaces are both enabled. You must select the default *Active-Active* NIC redundancy setting in the following step.

In this NIC mode, DHCP replies are returned to both the shared LOM ports and the Cisco card ports. If the system determines that the Cisco card connection is not getting its IP address from a Cisco UCS Manager system because the node is in standalone mode, further DHCP requests from the Cisco card are disabled. Use the Cisco Card NIC mode if you want to connect to Cisco IMC through a Cisco card in standalone mode.

- *Shared LOM*—The 1-Gb/10-Gb Ethernet ports are used to access Cisco IMC. You must select either the *Active-Active* or *Active-standby* NIC redundancy setting in the following step.
- *Dedicated*—The dedicated management port is used to access Cisco IMC. You must select the *None* NIC redundancy setting in the following step.
- *Cisco Card*—The ports on an installed Cisco UCS Virtual Interface Card (VIC) are used to access the Cisco IMC. You must select either the *Active-Active* or *Active-standby* NIC redundancy setting in the following step.

See also the required VIC Slot setting below.



- *VIC Slot*—Only if you use the Cisco Card NIC mode, you must select this setting to match where your VIC is installed. The choices are Riser1, Riser2, or Flex-LOM (the mLOM slot).
  - If you select Riser1, you must install the VIC in slot 2.
  - If you select Riser2, you must install the VIC in slot 5.
  - If you select Flex-LOM, you must install an mLOM-style VIC in the mLOM slot.

**Step 2** Set the NIC redundancy to your preference. This node has three possible NIC redundancy settings:

- *None*—The Ethernet ports operate independently and do not fail over if there is a problem. This setting can be used only with the Dedicated NIC mode.
- *Active-standby*—If an active Ethernet port fails, traffic fails over to a standby port. Shared LOM and Cisco Card modes can each use either Active-standby or Active-active settings.
- *Active-active* (default)—All Ethernet ports are utilized simultaneously. The Shared LOM EXT mode must use only this NIC redundancy setting. Shared LOM and Cisco Card modes can each use either Active-standby or Active-active settings.

**Step 3** Choose whether to enable DHCP for dynamic network settings, or to enter static network settings.

**Note** Before you enable DHCP, you must preconfigure your DHCP node with the range of MAC addresses for this node. The MAC address is printed on a label on the rear of the node. This node has a range of six MAC addresses assigned to Cisco IMC. The MAC address printed on the label is the beginning of the range of six contiguous MAC addresses.

The *static* IPv4 and IPv6 settings include the following:

- The Cisco IMC IP address.  
For IPv6, valid values are 1 - 127.
- The gateway.  
For IPv6, if you do not know the gateway, you can set it as none by entering :: (two colons).
- The preferred DNS node address.  
For IPv6, you can set this as none by entering :: (two colons).

**Step 4** (Optional) Make VLAN settings.

**Step 5** Press **F1** to go to the second settings window, then continue with the next step.

From the second window, you can press **F2** to switch back to the first window.

**Step 6** (Optional) Set a hostname for the node.

**Step 7** (Optional) Enable dynamic DNS and set a dynamic DNS (DDNS) domain.

**Step 8** (Optional) If you check the Factory Default check box, the node reverts to the factory defaults.

You can use this option to reset user credentials in future. For detailed steps, refer *Cisco UCS C-Series Integrated Management Controller GUI Configuration Guide* for your Cisco IMC release at [Configuration Guides](#).

**Step 9** (Optional) Set a default user password.

**Note** The factory default username for the node is *admin*. The default password is *password*.

**Step 10** (Optional) Enable auto-negotiation of port settings or set the port speed and duplex mode manually.

**Note** Auto-negotiation is applicable only when you use the Dedicated NIC mode. Auto-negotiation sets the port speed and duplex mode automatically based on the switch port to which the node is connected. If you disable auto-negotiation, you must set the port speed and duplex mode manually.

**Step 11** (Optional) Reset port profiles and the port name.

**Step 12** Press **F5** to refresh the settings that you made. You might have to wait about 45 seconds until the new settings appear and the message, “Network settings configured” is displayed before you reboot the node in the next step.

**Step 13** Press **F10** to save your settings and reboot the node.

**Note** If you chose to enable DHCP, the dynamically assigned IP and MAC addresses are displayed on the console screen during bootup.

### What to do next

Use a browser and the IP address of the Cisco IMC to connect to the Cisco IMC management interface. The IP address is based upon the settings that you made (either a static address or the address assigned by your DHCP node).



**Note** The factory default username for the node is *admin*. The default password is *password*.

To manage the node, see the *Cisco UCS C-Series Integrated Management Controller GUI Configuration Guide* or the *Cisco UCS C-Series Servers Integrated Management Controller CLI Configuration Guide* for instructions on using those interfaces for your Cisco IMC release. The links to the configuration guides are in the [Cisco Integrated Management Controller](#).

## NIC Mode and NIC Redundancy Settings

**Table 12: Valid NIC Redundancy Settings For Each NIC Mode**

NIC Mode	Valid NIC Redundancy Settings
Shared LOM EXT	Active-active
Dedicated	None
Shared LOM	Active-active Active-standby
Cisco Card	Active-active Active-standby

This node has the following NIC mode settings that you can choose from:

- *Shared LOM EXT* (default)—This is the shared LOM extended mode, the factory-default setting. With this mode, the Shared LOM and Cisco Card interfaces are both enabled. You must select the default *Active-Active* NIC redundancy setting in the following step.

In this NIC mode, DHCP replies are returned to both the shared LOM ports and the Cisco card ports. If the system determines that the Cisco card connection is not getting its IP address from a Cisco UCS Manager system because the node is in standalone mode, further DHCP requests from the Cisco card are disabled. Use the Cisco Card NIC mode if you want to connect to Cisco IMC through a Cisco card in standalone mode.

- *Shared LOM*—The 1-Gb/10-Gb Ethernet ports are used to access Cisco IMC. You must select either the *Active-Active* or *Active-standby* NIC redundancy setting in the following step.
- *Dedicated*—The dedicated management port is used to access Cisco IMC. You must select the *None* NIC redundancy setting in the following step.
- *Cisco Card*—The ports on an installed Cisco UCS Virtual Interface Card (VIC) are used to access the Cisco IMC. You must select either the *Active-Active* or *Active-standby* NIC redundancy setting in the following step.

See also the required VIC Slot setting below.

- *VIC Slot*—Only if you use the Cisco Card NIC mode, you must select this setting to match where your VIC is installed. The choices are Riser1, Riser2, or Flex-LOM (the mLOM slot).
  - If you select Riser1, you must install the VIC in slot 2.
  - If you select Riser2, you must install the VIC in slot 5.
  - If you select Flex-LOM, you must install an mLOM-style VIC in the mLOM slot.

This node has the following NIC redundancy settings that you can choose from:

- *None*—The Ethernet ports operate independently and do not fail over if there is a problem. This setting can be used only with the Dedicated NIC mode.
- *Active-standby*—If an active Ethernet port fails, traffic fails over to a standby port. Shared LOM and Cisco Card modes can each use either Active-standby or Active-active settings.
- *Active-active* (default)—All Ethernet ports are utilized simultaneously. The Shared LOM EXT mode must use only this NIC redundancy setting. Shared LOM and Cisco Card modes can each use either Active-standby or Active-active settings.

## Updating the BIOS and Cisco IMC Firmware



### Caution

When you upgrade the BIOS firmware, you must also upgrade the Cisco IMC firmware to the same version or the server does not boot. Do not power off the server until the BIOS and Cisco IMC firmware are matching or the server does not boot.

Cisco provides the *Cisco Host Upgrade Utility* to assist with simultaneously upgrading the BIOS, Cisco IMC, and other firmware to compatible levels.

The server uses firmware obtained from and certified by Cisco. Cisco provides release notes with each firmware image. There are several possible methods for updating the firmware:

- **Recommended method for firmware update:** Use the Cisco Host Upgrade Utility to simultaneously upgrade the Cisco IMC, BIOS, and component firmware to compatible levels.  
See the *Cisco Host Upgrade Utility Quick Reference Guide* for your firmware release at the documentation roadmap link below.
- You can upgrade the Cisco IMC and BIOS firmware by using the Cisco IMC GUI interface.  
See the *Cisco UCS C-Series Rack-Mount Server Configuration Guide*.
- You can upgrade the Cisco IMC and BIOS firmware by using the Cisco IMC CLI interface.  
See the *Cisco UCS C-Series Rack-Mount Server CLI Configuration Guide*.

For links to the documents listed above, see the [Cisco UCS C-Series Documentation Roadmap](#).

## Accessing the System BIOS

**Step 1** Enter the BIOS Setup Utility by pressing the **F2** key when prompted during bootup.

**Note** The version and build of the current BIOS are displayed on the Main page of the utility.

**Step 2** Use the arrow keys to select the BIOS menu page.

**Step 3** Highlight the field to be modified by using the arrow keys.

**Step 4** Press **Enter** to select the field that you want to change, and then modify the value in the field.

**Step 5** Press the right arrow key until the Exit menu screen is displayed.

**Step 6** Follow the instructions on the Exit menu screen to save your changes and exit the setup utility (or press **F10**). You can exit without saving changes by pressing **Esc**.

## Smart Access Serial

This server supports the Smart Access Serial feature. This feature allows you to switch between host serial and Cisco IMC CLI.

- This feature has the following requirements:
  - A serial cable connection, which can use either the RJ-45 serial connector on the server rear panel, or a DB-9 connection when using the KVM cable (Cisco PID N20-BKVM) on the front-panel KVM console connector.
  - Console redirection must be enabled in the server BIOS.
  - Terminal type must be set to VT100+ or VTUFT8.
  - Serial-over-LAN (SOL) must be disabled (SOL is disabled by default).
- To switch from host serial to Cisco IMC CLI, press **Esc+9**.

You must enter your Cisco IMC credentials to authenticate the connection.

- To switch from Cisco IMC CLI to host serial, press **Esc+8**.



---

**Note** You cannot switch to Cisco IMC CLI if the serial-over-LAN (SOL) feature is enabled.

---

- After a session is created, it is shown in the CLI or web GUI by the name `serial`.

## Smart Access USB

This node supports the Smart Access USB feature. The board management controller (BMC) in this node can accept a USB mass storage device and access the data on it. This feature allows you to use the front-panel USB device as a medium to transfer data between the BMC and the user without need for network connectivity. This can be useful, for example, when remote BMC interfaces are not yet available, or are not accessible due to network misconfiguration.

- This feature has the following requirements:
  - The KVM cable (Cisco PID N20-BKVM) connected to the front panel KVM console connector.
  - A USB storage device connected to one of the USB 2.0 connectors on the KVM cable. The USB device must draw less than 500 mA to avoid disconnect by the current-protection circuit.



---

**Note** Any mouse or keyboard that is connected to the KVM cable is disconnected when you enable Smart Access USB.

---

- You can use USB 3.0-based devices, but they will operate at USB 2.0 speed.
- We recommend that the USB device have only one partition.
- The file system formats supported are: FAT16, FAT32, MSDOS, EXT2, EXT3, and EXT4. NTFS is not supported.
- The front-panel KVM connector has been designed to switch the USB port between Host OS and BMC.
- Smart Access USB can be enabled or disabled using any of the BMC user interfaces. For example, you can use the Cisco IMC Configuration Utility that is accessed by pressing **F8** when prompted during bootup.
  - Enabled: the front-panel USB device is connected to the BMC.
  - Disabled: the front-panel USB device is connected to the host.
- In a case where no management network is available to connect remotely to Cisco IMC, a Device Firmware Update (DFU) shell over serial cable can be used to generate and download technical support files to the USB device that is attached to front panel USB port.





## CHAPTER 3

# Maintaining the Node

---

- [Status LEDs and Buttons](#), on page 61
- [Preparing For Component Installation](#), on page 68
- [Removing the Node Top Cover](#), on page 70
- [Serial Number Location](#), on page 72
- [Hot Swap vs Hot Plug](#), on page 72
- [Replacing the Air Duct](#), on page 72
- [Removing and Replacing Components](#), on page 75
- [Replacing the RTC Battery](#), on page 123
- [Replacing Power Supplies](#), on page 124
- [Replacing a PCIe Riser](#), on page 129
- [Replacing a PCIe Card](#), on page 130
- [Replacing an mLOM Card](#), on page 134
- [Replacing a SAS Storage Controller Card \(RAID or HBA\)](#), on page 135
- [Replacing a SATA Interposer Card \(12-Drive SFF Server Only\)](#), on page 151
- [Replacing the Supercap \(RAID Backup\)](#), on page 155
- [Replacing a Boot-Optimized M.2 RAID Controller Module](#), on page 158
- [Replacing a Chassis Intrusion Switch](#), on page 162
- [Installing a Trusted Platform Module \(TPM\)](#), on page 163
- [Removing the PCB Assembly \(PCBA\)](#), on page 166
- [Service Headers and Jumpers](#), on page 167

## Status LEDs and Buttons

This section contains information for interpreting LED states.

# Front-Panel LEDs

Figure 12: Front Panel LEDs




Table 13: Front Panel LEDs, Definition of States

	LED Name	States
1 SAS	SAS/SATA drive fault <b>Note</b> NVMe solid state drive (SSD) drive tray LEDs have different behavior than SAS/SATA drive trays.	<ul style="list-style-type: none"> <li>• Off—The hard drive is operating properly.</li> <li>• Amber—Drive fault detected.</li> <li>• Amber, blinking—The device is rebuilding.</li> <li>• Amber, blinking with one-second interval—Drive locate function activated in the software.</li> </ul>
2 SAS	SAS/SATA drive activity LED	<ul style="list-style-type: none"> <li>• Off—There is no hard drive in the hard drive tray (no access, no fault).</li> <li>• Green—The hard drive is ready.</li> <li>• Green, blinking—The hard drive is reading or writing data.</li> </ul>



<b>1</b> <b>NVMe</b>	NVMe SSD drive fault <b>Note</b> NVMe solid state drive (SSD) drive tray LEDs have different behavior than SAS/SATA drive trays.	<ul style="list-style-type: none"> <li>• Off—The drive is not in use and can be safely removed.</li> <li>• Green—The drive is in use and functioning properly.</li> <li>• Green, blinking—the driver is initializing following insertion, or the driver is unloading following an eject command.</li> <li>• Amber—The drive has failed.</li> <li>• Amber, blinking—A drive Locate command has been issued in the software.</li> </ul>
<b>2</b> <b>NVMe</b>	NVMe SSD activity	<ul style="list-style-type: none"> <li>• Off—No drive activity.</li> <li>• Green, blinking—There is drive activity.</li> </ul>
<b>3</b>	Power button/LED	<ul style="list-style-type: none"> <li>• Off—There is no AC power to the node.</li> <li>• Amber—The node is in standby power mode. Power is supplied only to the Cisco IMC and some motherboard functions.</li> <li>• Green—The node is in main power mode. Power is supplied to all node components.</li> </ul>
<b>4</b>	Unit identification	<ul style="list-style-type: none"> <li>• Off—The unit identification function is not in use.</li> <li>• Blue, blinking—The unit identification function is activated.</li> </ul>

5	System health	<ul style="list-style-type: none"> <li>• Green—The node is running in normal operating condition.</li> <li>• Green, blinking—The node is performing system initialization and memory check.</li> <li>• Amber, steady—The node is in a degraded operational state (minor fault). For example: <ul style="list-style-type: none"> <li>• Power supply redundancy is lost.</li> <li>• CPUs are mismatched.</li> <li>• At least one CPU is faulty.</li> <li>• At least one DIMM is faulty.</li> <li>• At least one drive in a RAID configuration failed.</li> </ul> </li> <li>• Amber, 2 blinks—There is a major fault with the system board.</li> <li>• Amber, 3 blinks—There is a major fault with the memory DIMMs.</li> <li>• Amber, 4 blinks—There is a major fault with the CPUs.</li> </ul>
6	Fan status	<ul style="list-style-type: none"> <li>• Green—All fan modules are operating properly.</li> <li>• Amber, blinking—One or more fan modules breached the non-recoverable threshold.</li> </ul>
7	Temperature status	<ul style="list-style-type: none"> <li>• Green—The node is operating at normal temperature.</li> <li>• Amber, steady—One or more temperature sensors breached the critical threshold.</li> <li>• Amber, blinking—One or more temperature sensors breached the non-recoverable threshold.</li> </ul>
8	Power supply status	<ul style="list-style-type: none"> <li>• Green—All power supplies are operating normally.</li> <li>• Amber, steady—One or more power supplies are in a degraded operational state.</li> <li>• Amber, blinking—One or more power supplies are in a critical fault state.</li> </ul>
9	Network link activity	<ul style="list-style-type: none"> <li>• Off—The Ethernet LOM port link is idle.</li> <li>• Green—One or more Ethernet LOM ports are link-active, but there is no activity.</li> <li>• Green, blinking—One or more Ethernet LOM ports are link-active, with activity.</li> </ul>

<p><b>10</b></p>	<p>DVD drive activity</p>	<ul style="list-style-type: none"> <li>• Off—The drive is idle.</li> <li>• Green, steady—The drive is spinning up a disk.</li> <li>• Green, blinking—The drive is accessing data.</li> </ul>
------------------	---------------------------	--

## Rear-Panel LEDs

Figure 13: Rear Panel LEDs




Table 14: Rear Panel LEDs, Definition of States

	LED Name	States
<p><b>1</b></p>	<p>Rear unit identification</p>	<ul style="list-style-type: none"> <li>• Off—The unit identification function is not in use.</li> <li>• Blue, blinking—The unit identification function is activated.</li> </ul>
<p><b>2</b></p>	<p>1-Gb/10-Gb Ethernet link speed (on both LAN1 and LAN2)</p>	<ul style="list-style-type: none"> <li>• Amber—Link speed is 100 Mbps.</li> <li>• Amber—Link speed is 1 Gbps.</li> <li>• Green—Link speed is 10 Gbps.</li> </ul>
<p><b>3</b></p>	<p>1-Gb/10-Gb Ethernet link status (on both LAN1 and LAN2)</p>	<ul style="list-style-type: none"> <li>• Off—No link is present.</li> <li>• Green—Link is active.</li> <li>• Green, blinking—Traffic is present on the active link.</li> </ul>
<p><b>4</b></p>	<p>1-Gb Ethernet dedicated management link speed</p>	<ul style="list-style-type: none"> <li>• Off—Link speed is 10 Mbps.</li> <li>• Amber—Link speed is 100 Mbps.</li> <li>• Green—Link speed is 1 Gbps.</li> </ul>

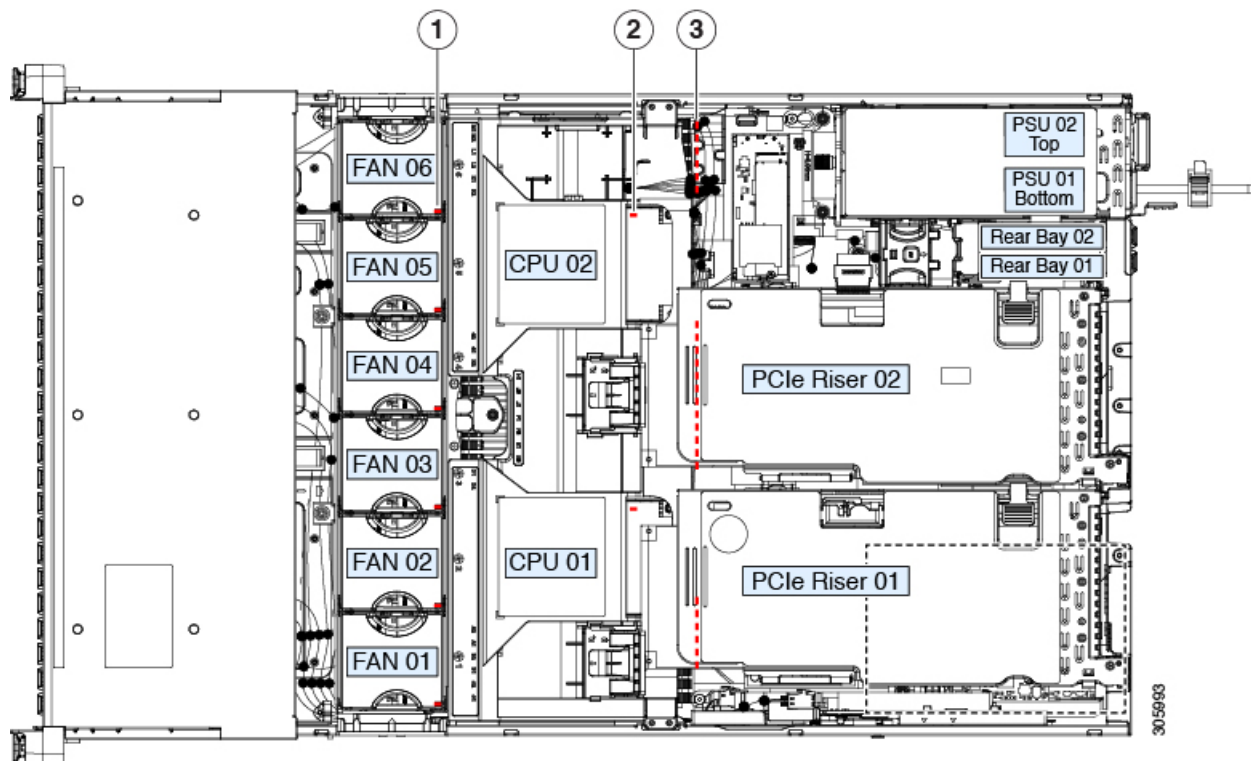
5	1-Gb Ethernet dedicated management link status	<ul style="list-style-type: none"> <li>• Off—No link is present.</li> <li>• Green—Link is active.</li> <li>• Green, blinking—Traffic is present on the active link.</li> </ul>
6	Power supply status (one LED each power supply unit)	<p><b>AC power supplies:</b></p> <ul style="list-style-type: none"> <li>• Off—No AC input (12 V main power off, 12 V standby power off).</li> <li>• Green, blinking—12 V main power off; 12 V standby power on.</li> <li>• Green, solid—12 V main power on; 12 V standby power on.</li> <li>• Amber, blinking—Warning threshold detected but 12 V main power on.</li> <li>• Amber, solid—Critical error detected; 12 V main power off (for example, over-current, over-voltage, or over-temperature failure).</li> </ul> <p><b>DC power supply (HX-PSUV2-1050DC):</b></p> <ul style="list-style-type: none"> <li>• Off—No DC input (12 V main power off, 12 V standby power off).</li> <li>• Green, blinking—12 V main power off; 12 V standby power on.</li> <li>• Green, solid—12 V main power on; 12 V standby power on.</li> <li>• Amber, blinking—Warning threshold detected but 12 V main power on.</li> <li>• Amber, solid—Critical error detected; 12 V main power off (for example, over-current, over-voltage, or over-temperature failure).</li> </ul>
7 SAS	<p>SAS/SATA drive fault</p> <p><b>Note</b> NVMe solid state drive (SSD) drive tray LEDs have different behavior than SAS/SATA drive trays.</p>	<ul style="list-style-type: none"> <li>• Off—The hard drive is operating properly.</li> <li>• Amber—Drive fault detected.</li> <li>• Amber, blinking—The device is rebuilding.</li> <li>• Amber, blinking with one-second interval—Drive locate function activated in the software.</li> </ul>

<b>8</b> <b>SAS</b>	SAS/SATA drive activity LED	<ul style="list-style-type: none"> <li>• Off—There is no hard drive in the hard drive tray (no access, no fault).</li> <li>• Green—The hard drive is ready.</li> <li>• Green, blinking—The hard drive is reading or writing data.</li> </ul>
<b>7</b> <b>NVMe</b>	NVMe SSD drive fault <b>Note</b> NVMe solid state drive (SSD) drive tray LEDs have different behavior than SAS/SATA drive trays.	<ul style="list-style-type: none"> <li>• Off—The drive is not in use and can be safely removed.</li> <li>• Green—The drive is in use and functioning properly.</li> <li>• Green, blinking—the driver is initializing following insertion, or the driver is unloading following an eject command.</li> <li>• Amber—The drive has failed.</li> <li>• Amber, blinking—A drive Locate command has been issued in the software.</li> </ul>
<b>8</b> <b>NVMe</b>	NVMe SSD activity	<ul style="list-style-type: none"> <li>• Off—No drive activity.</li> <li>• Green, blinking—There is drive activity.</li> </ul>

## Internal Diagnostic LEDs

The node has internal fault LEDs for CPUs, DIMMs, and fan modules.

Figure 14: Internal Diagnostic LED Locations



<p><b>1</b></p>	<p>Fan module fault LEDs (one on the top of each fan module)</p> <ul style="list-style-type: none"> <li>• Amber—Fan has a fault or is not fully seated.</li> <li>• Green—Fan is OK.</li> </ul>	<p><b>3</b></p>	<p>DIMM fault LEDs (one behind each DIMM socket on the motherboard)</p> <p>These LEDs operate only when the node is in standby power mode.</p> <ul style="list-style-type: none"> <li>• Amber—DIMM has a fault.</li> <li>• Off—DIMM is OK.</li> </ul>
<p><b>2</b></p>	<p>CPU fault LEDs (one behind each CPU socket on the motherboard).</p> <p>These LEDs operate only when the node is in standby power mode.</p> <ul style="list-style-type: none"> <li>• Amber—CPU has a fault.</li> <li>• Off—CPU is OK.</li> </ul>	<p>-</p>	

## Preparing For Component Installation

This section includes information and tasks that help prepare the node for component installation.

## Required Equipment For Service Procedures

The following tools and equipment are used to perform the procedures in this chapter:

- T-30 Torx driver (supplied with replacement CPUs for heatsink removal)
- #1 flat-head screwdriver (used during CPU or heatsink replacement)
- #1 Phillips-head screwdriver (for M.2 SSD and intrusion switch replacement)
- Electrostatic discharge (ESD) strap or other grounding equipment such as a grounded mat

## Shutting Down and Removing Power From the Node

The node can run in either of two power modes:

- Main power mode—Power is supplied to all node components and any operating system on your drives can run.
- Standby power mode—Power is supplied only to the service processor and certain components. It is safe for the operating system and data to remove power cords from the node in this mode.



---

**Caution**

After a node is shut down to standby power, electric current is still present in the node. To completely remove power, you must disconnect all power cords from the power supplies in the node, as directed in the service procedures.

You can shut down the node by using the front-panel power button or the software management interfaces.

---

### Shutting Down Using the Power Button

---

**Step 1** Check the color of the Power button/LED:

- Amber—The node is already in standby mode, and you can safely remove power.
- Green—The node is in main power mode and must be shut down before you can safely remove power.

**Step 2** Invoke either a graceful shutdown or a hard shutdown:

**Caution** To avoid data loss or damage to your operating system, you should always invoke a graceful shutdown of the operating system.

- Graceful shutdown—Press and release the **Power** button. The operating system performs a graceful shutdown, and the node goes to standby mode, which is indicated by an amber Power button/LED.
- Emergency shutdown—Press and hold the **Power** button for 4 seconds to force the main power off and immediately enter standby mode.

**Step 3** If a service procedure instructs you to completely remove power from the node, disconnect all power cords from the power supplies in the node.

---

## Shutting Down Using The Cisco IMC GUI

You must log in with user or admin privileges to perform this task.

---

- Step 1** In the Navigation pane, click the **Server** tab.
- Step 2** On the Server tab, click **Summary**.
- Step 3** In the Actions area, click **Power Off Server**.
- Step 4** Click **OK**.

The operating system performs a graceful shutdown, and the node goes to standby mode, which is indicated by an amber Power button/LED.

- Step 5** If a service procedure instructs you to completely remove power from the node, disconnect all power cords from the power supplies in the node.
- 

## Shutting Down Using The Cisco IMC CLI

You must log in with user or admin privileges to perform this task.

---

- Step 1** At the server prompt, enter:

**Example:**

```
server# scope chassis
```

- Step 2** At the chassis prompt, enter:

**Example:**

```
server/chassis# power shutdown
```

The operating system performs a graceful shutdown, and the node goes to standby mode, which is indicated by an amber Power button/LED.

- Step 3** If a service procedure instructs you to completely remove power from the node, disconnect all power cords from the power supplies in the node.
- 

## Removing the Node Top Cover

- Step 1** Remove the top cover:

- a) If the cover latch is locked, slide the lock sideways to unlock it.

When the latch is unlocked, the handle pops up so that you can grasp it.

- b) Lift on the end of the latch so that it pivots vertically to 90 degrees.
- c) Simultaneously, slide the cover back and lift the top cover straight up from the node and set it aside.



- Step 2** Replace the top cover:



- a) With the latch in the fully open position, place the cover on top of the node about one-half inch (1.27 cm) behind the lip of the front cover panel.
- b) Slide the cover forward until the latch makes contact.
- c) Press the latch down to the closed position. The cover is pushed forward to the closed position as you push down the latch.
- d) Lock the latch by sliding the lock button to sideways to the left.

Locking the latch ensures that the node latch handle does not protrude when you install the blade.

**Figure 15: Removing the Top Cover**



309157

1	Cover lock	2	Cover latch handle
---	------------	---	--------------------

---

## Serial Number Location

The serial number for the node is printed on a label on the top of the node, near the front. See [Removing the Node Top Cover, on page 70](#).

## Hot Swap vs Hot Plug

Some components can be removed and replaced without shutting down and removing power from the node. This type of replacement has two varieties: hot-swap and hot-plug.

- Hot-swap replacement—You do not have to shut down the component in the software or operating system. This applies to the following components:
  - SAS/SATA hard drives
  - SAS/SATA solid state drives
  - Cooling fan modules
  - Power supplies (when redundant as 1+1)
- Hot-plug replacement—You must take the component offline before removing it for the following component:
  - NVMe PCIe solid state drives

## Replacing the Air Duct

The node has an air duct under the top sheet metal cover. The air duct ensures proper cooling and air flow across the node from intake (the cool aisle of the data center) to exhaust (the hot aisle in the data center). The air duct is in the middle of the node and covers the CPU and DIMMs.

To replace the node's air duct, use the following procedures:

- [Removing the Air Duct, on page 72](#)
- [Installing the Air Duct, on page 74](#)

## Removing the Air Duct


Use this procedure to remove the air duct when needed.

---

**Step 1** Remove the server top cover.


**Step 2** Locate the detents for the air duct.

The following illustration highlights the middle detent for illustrative purposes only. When removing the air duct, always grasp the detents closest to the chassis sidewalls (left and right).



**Step 3** Grasp the left and right detent and lift it out of the chassis.

**Note** You might need to slide the air duct towards the back of the server while lifting the air duct up.



---


**What to do next**

When you are done servicing the node, install the air duct. See [Installing the Air Duct](#), on page 74.

## Installing the Air Duct

The air duct sits behind the front-loading drive cage and covers the CPU and DIMMs in the middle of the node.

- 
- Step 1** Orient the air duct as shown.
- Step 2** Lower the air duct into place and gently press down to ensure that all of its edges sit flush.  
If the air duct is not seated correctly, it can obstruct installing the node's top cover.
- Step 3** When the air duct is correctly seated, attach the node's top cover.  
The node top cover should sit flush so that the metal tabs on the top cover match the indents in the top edges of the air duct.



## Removing and Replacing Components



### Warning

Blank faceplates and cover panels serve three important functions: they prevent exposure to hazardous voltages and currents inside the chassis; they contain electromagnetic interference (EMI) that might disrupt other equipment; and they direct the flow of cooling air through the chassis. Do not operate the system unless all cards, faceplates, front covers, and rear covers are in place.

Statement 1029



**Caution** When handling node components, handle them only by carrier edges and use an electrostatic discharge (ESD) wrist-strap or other grounding device to avoid damage.




**Tip** You can press the unit identification button on the front panel or rear panel to turn on a flashing, blue unit identification LED on both the front and rear panels of the node. This button allows you to locate the specific node that you are servicing when you go to the opposite side of the rack. You can also activate these LEDs remotely by using the Cisco IMC interface.

This section describes how to install and replace node components.

## Serviceable Component Locations

This topic shows the locations of the field-replaceable components and service-related items. The view in the following figure shows the server with the top cover removed.

Figure 16: Cisco HX C240 M6 Server, Serviceable Component Locations



1	Front-loading drive bays.	2	Cooling fan modules (six, hot-swappable)
---	---------------------------	---	--

<p><b>3</b></p>	<p>DIMM sockets on motherboard (16 per CPU)</p> <p>See <a href="#">DIMM Population Rules and Memory Performance Guidelines</a> , on page 115 for DIMM slot numbering.</p> <p><b>Note</b> An air baffle rests on top of the DIMM and CPUs when the server is operating. The air baffle is not displayed in this illustration.</p>	<p><b>4</b></p>	<p>CPU socket 1</p>
<p><b>5</b></p>	<p>CPU socket 2</p>	<p><b>6</b></p>	<p>M.2 RAID Controller</p>
<p><b>7</b></p>	<p>PCIe riser 3 (PCIe slots 7 and 8 numbered from bottom to top), with the following options:</p> <ul style="list-style-type: none"> <li>• 3A (Default Option)—Slots 7 (x16 mechanical, x8 electrical), and 8 (x16 mechanical, x8 electrical). Both slots can accept a full height, full length GPU card.</li> <li>• 3B (Storage Option)—Slots 7 (x24 mechanical, x4 electrical) and 8 (x24 mechanical, x4 electrical). Both slots can accept 2.5-inch SFF universal HDDs.</li> <li>• 3C (GPU Option)—Slots 7 (x16 mechanical, x16 electrical) and 8 empty (NCSI support limited to one slot at a time). Slot 7 can support a full height, full length GPU card.</li> </ul>	<p><b>8</b></p>	<p>PCIe riser 2 (PCIe slots 4, 5, 6 numbered from bottom to top), with the following options:</p> <ul style="list-style-type: none"> <li>• 2A (Default Option)—Slot 4 (x24 mechanical, x8 electrical) supports full height, <sup>3</sup>/<sub>4</sub> length card; Slot 5 (x24 mechanical, x16 electrical) supports full height, full length GPU card; Slot 6 (x16 mechanical, x8 electrical) supports full height, full length card.</li> </ul>
<p><b>9</b></p>	<p>PCIe riser 1 (PCIe slot 1, 2, 3 numbered bottom to top), with the following options:</p> <ul style="list-style-type: none"> <li>• 1A (Default Option)—Slot 1 (x24 mechanical, x8 electrical) supports full height, <sup>3</sup>/<sub>4</sub> length card; Slot 2 (x24 mechanical, x16 electrical) supports full height, full length GPU card; Slot 3 (x16 mechanical, x8 electrical) supports full height, full length card.</li> <li>• 1B (Storage Option)—Slot 1 (x24 mechanical, x8 electrical) supports full height, <sup>3</sup>/<sub>4</sub> length card; Slot 2 (x4 electrical), supports 2.5-inch SFF universal HDD; Slot 3 (x4 electrical), supports 2.5-inch SFF universal HDD</li> </ul>	<p>-</p>	

## Replacing Front-Loading SAS/SATA Drives



**Note** You do not have to shut down the node or drive to replace SAS/SATA hard drives or SSDs because they are hot-swappable.

To replace rear-loading SAS/SATA drives, see [Replacing Rear-Loading SAS/SATA Drives, on page 81](#).

### Front-Loading SAS/SATA Drive Population Guidelines

The node is orderable in four different versions, each with a different front panel/drive-backplane configuration.

- Cisco HyperFlex C240 M6 24 SAS/SATA—Small form-factor (SFF) drives, with 24-drive backplane.
  - Front-loading drive bays 1—24 support 2.5-inch SAS/SATA drives.
  - Optionally, front-loading drive bays 1 through 4 can support 2.5-inch NVMe SSDs.
- Cisco HyperFlex C240 M6 24 NVMe—SFF drives, with 24-drive backplane.
  - Front-loading drive bays 1—24 support 2.5-inch NVMe PCIe SSDs *only*.
- Cisco HyperFlex C240 M6 12 SAS/SATA plus optical drive—SFF drives, with 12-drive backplane and DVD drive option.
  - Front-loading drive bays 1—12 support 2.5-inch SAS/SATA drives.
  - Optionally, front-loading drive bays 1 through 4 can support 2.5-inch NVMe SSDs.
- Cisco HyperFlex C240 M6 12 NVMe—SFF drives, with 24-drive backplane.
  - Front-loading drive bays 1—12 support 2.5-inch NVMe PCIe SSDs *only*
- Cisco HyperFlex C240 M6 12 LFF SAS/SATA—Large form-factor (LFF) drives, with 12-drive backplane.
  - Front-loading drive bays 1—12 support 3.5-inch SAS-only drives.
  - Optionally, up to 4 mid-plane mounted SAS-only HDDs can be supported.
  - Optionally, rear drive bays can support up to 4 SFF SAS/SATA or NVMe HDDs.

Drive bay numbering is shown in the following figures.

**Figure 17: Small Form-Factor Drive (24-Drive) Versions, Drive Bay Numbering**







Figure 18: Large Form-Factor Drive (12-Drive) Version, Drive Bay Numbering



Observe these drive population guidelines for optimum performance:

- When populating drives, add drives to the lowest-numbered bays first.
- Front-loading drives are hot pluggable, but each drive requires a 10 second delay between hot removal and hot insertion.
- Keep an empty drive blanking tray in any unused bays to ensure proper airflow.
- You can mix SAS/SATA hard drives and SAS/SATA SSDs in the same node. However, you cannot configure a logical volume (virtual drive) that contains a mix of hard drives and SSDs. That is, when you create a logical volume, it must contain all SAS/SATA hard drives or all SAS/SATA SSDs.

## 4K Sector Format SAS/SATA Drives Considerations

- You must boot 4K sector format drives in UEFI mode, not legacy mode. See the procedures in this section.
- Do not configure 4K sector format and 512-byte sector format drives as part of the same RAID volume.
- For operating system support on 4K sector drives, see the interoperability matrix tool for your node: [Hardware and Software Interoperability Matrix Tools](#)

### Setting Up UEFI Mode Booting in the BIOS Setup Utility

- Step 1** Enter the BIOS setup utility by pressing the **F2** key when prompted during bootup.
- Step 2** Go to the **Boot Options** tab.
- Step 3** Set **UEFI Boot Options** to **Enabled**.
- Step 4** Under **Boot Option Priorities**, set your OS installation media (such as a virtual DVD) as your **Boot Option #1**.
- Step 5** Go to the **Advanced** tab.
- Step 6** Select **LOM** and **PCIe Slot Configuration**.
- Step 7** Set the **PCIe Slot ID: HBA Option ROM** to **UEFI Only**.
- Step 8** Press **F10** to save changes and exit the BIOS setup utility. Allow the node to reboot.
- Step 9** After the OS installs, verify the installation:
  - a) Enter the BIOS setup utility by pressing the **F2** key when prompted during bootup.
  - b) Go to the Boot Options tab.

- c) Under **Boot Option Priorities**, verify that the OS you installed is listed as your **Boot Option #1**.

---

## Setting Up UEFI Mode Booting in the Cisco IMC GUI

---

- Step 1** Use a web browser and the IP address of the node to log into the Cisco IMC GUI management interface.
- Step 2** Navigate to **Server > BIOS**.
- Step 3** Under Actions, click **Configure BIOS**.
- Step 4** In the Configure BIOS Parameters dialog, select the **Advanced** tab.
- Step 5** Go to the **LOM** and **PCIe Slot Configuration** section.
- Step 6** Set the **PCIe Slot: HBA Option ROM** to **UEFI Only**.
- Step 7** Click **Save Changes**. The dialog closes.
- Step 8** Under BIOS Properties, set **Configured Boot Order** to **UEFI**.
- Step 9** Under Actions, click **Configure Boot Order**.
- Step 10** In the Configure Boot Order dialog, click **Add Local HDD**.
- Step 11** In the Add Local HDD dialog, enter the information for the 4K sector format drive and make it first in the boot order.
- Step 12** Save changes and reboot the node. The changes you made will be visible after the system reboots.
- 

## Replacing a Front-Loading SAS/SATA Drive





---

**Note** You do not have to shut down the node or drive to replace SAS/SATA hard drives or SSDs because they are hot-swappable.

---

- Step 1** Remove the drive that you are replacing or remove a blank drive tray from the bay:
- Press the release button on the face of the drive tray.
  - Grasp and open the ejector lever and then pull the drive tray out of the slot.
  - If you are replacing an existing drive, remove the four drive-tray screws that secure the drive to the tray and then lift the drive out of the tray.
- Step 2** Install a new drive:
- Place a new drive in the empty drive tray and install the four drive-tray screws.
  - With the ejector lever on the drive tray open, insert the drive tray into the empty drive bay.
  - Push the tray into the slot until it touches the backplane, and then close the ejector lever to lock the drive in place.

Figure 19: Replacing a Drive in a Drive Tray



1	Ejector lever	3	Drive tray screws (two on each side)
2	Release button	4	Drive removed from drive tray

## Replacing Rear-Loading SAS/SATA Drives



**Note** You do not have to shut down the server or drive to replace SAS/SATA hard drives or SSDs because they are hot-swappable.

### Rear-Loading SAS/SATA Drive Population Guidelines

The rear drive bay support differs by node PID and which type of RAID controller is used in the node:

- HX C240 M6 24 SAS/SATA—Small form-factor (SFF) drives, with 24-drive backplane.
  - Hardware RAID—Rear drive bays support SAS or NVMe drives
  - Intel™ Virtual RAID on CPU—Rear drive bays support NVMe drives only.
- HX C240 M6 24 SAS/SATA—SFF drives, with 24-drive backplane.
  - Rear drive bays support only NVMe SSDs.
- HX C240 M6 8 SAS/SATA plus optical—SFF drives, with 8-drive backplane and DVD drive option.
  - Hardware RAID—Rear drive bays support SAS or NVMe drives
  - Intel™ Virtual RAID on CPU—Rear drive bays support NVMe drives only.
- HX C240 M6 12 LFF—Large form-factor (LFF) drives, with 12-drive backplane.
  - Hardware RAID—Rear drive bays support SAS or NVMe drives
  - Intel™ Virtual RAID on CPU—Rear drive bays support NVMe drives only.
- The rear drive bay numbering follows the front-drive bay numbering in each node version:
  - 8-drive node—rear bays are numbered bays 9 and 10.
  - 12-drive node—rear bays are numbered bays 13 and 14.
  - 24-drive node—rear bays are numbered bays 25 and 26.
- When populating drives, add drives to the lowest-numbered bays first.
- Keep an empty drive blanking tray in any unused bays to ensure proper airflow.
- You can mix SAS/SATA hard drives and SAS/SATA SSDs in the same node. However, you cannot configure a logical volume (virtual drive) that contains a mix of hard drives and SSDs. That is, when you create a logical volume, it must contain all SAS/SATA hard drives or all SAS/SATA SSDs.

## Replacing a Rear-Loading SAS/SATA Drive




---


**Note** You do not have to shut down the node or drive to replace SAS/SATA hard drives or SSDs because they are hot-swappable.

---


- Step 1** Remove the drive that you are replacing or remove a blank drive tray from the bay:
- a) Press the release button on the face of the drive tray.
  - b) Grasp and open the ejector lever and then pull the drive tray out of the slot.
  - c) If you are replacing an existing drive, remove the four drive-tray screws that secure the drive to the tray and then lift the drive out of the tray.
- Step 2** Install a new drive:
- a) Place a new drive in the empty drive tray and install the four drive-tray screws.



- Step 1** Open the node top cover.
- Step 2** Grasp the handle for the mid-mount drive cage, and swing the cage cover open.  
When the cage cover is open, it will be pointing up at a 90-degree angle.
- Step 3** Grasping the cage cover handle, pull up on the drive cage until the bottom row of drives clears the top of the node.  
When pulling on the mid-mount drive cage, it will arc upward.
- Step 4** Grasp the drive handle and pull the drive out of the mid-mount drive cage.



- Step 5** Orient the drive so that the handle is at the bottom and align it with its drive bay.
- Step 6** Holding the drive level, slide it into the drive bay until it connects with the midplane.
- Step 7** Push down on the drive cage so that it seats into the node.
- Step 8** Grasp the handle and close the node cage cover.
- Note** Make sure that the node cage cover is completely closed, and the node cage is completely seated in the node.  
When the node cage is completely seated, its top is flush with the fans and rear PCI riser cages.



**Step 9** Install the node's top cover.

If the node's top cover does not close easily, check that the mid-mount drive cage is completely seated into the node.

## Basic Troubleshooting: Reseating a SAS/SATA Drive

Due to a drive manufacturing issue, it is possible that a false positive UBAD error can occur on SAS/SATA HDDs installed in the node.

- Only drives that are managed by the UCS MegaRAID controller are affected.
- Drives can be affected regardless where they are installed in the node (front-loaded, rear-loaded, and so on).
- Both SFF and LFF form factor can be affected.
- Drives installed in all Cisco UCS C-Series nodes with M3 processors and later can be affected.
- Drives can be affected regardless of whether they are configured for hotplug or not.
- The UBAD error is not always valid, so the drive is not always defective or in need of repair or replacement. However, it is also possible that the error is correct, or the drive will need replacement.

**Before submitting the drive to the RMA process**, it is a best practice to reseat the drive. If the false UBAD error exists, reseating the drive can clear it. If successful, reseating the drive reduces inconvenience, cost, and service interruption, and optimizes your node uptime.



**Note** Reseat the drive only if a UBAD error occurs. Other errors are transient, and you should not attempt diagnostics and troubleshooting without the assistance of Cisco personnel. Contact Cisco TAC for assistance with other drive errors.

To reseat the drive, see [Reseating a SAS/SATA Drive, on page 86](#).

## Reseating a SAS/SATA Drive

Sometimes, SAS/SATA drives can throw a false UBAD error and reseating the drive can clear the error.

Use the following procedure to reseat the drive.



---

**Caution** This procedure requires powering down the node, which will cause a service interruption.

---

### Before you begin

Before attempting this procedure, be aware of the following:

- Before reseating the drive, it is a best practice to back up any data on it.
- When reseating the drive, make sure to reuse the same drive bay.
  - Do not move the drive to a different slot.
  - Do not move the drive to a different node.
  - If you do not reuse the same slot, the Cisco management software (for example, Cisco IMM) might require a rescan/rediscovery of the node.
- Regardless of the status of hotplug/hot insertion on the drive, it is a best practice to gracefully power down the drive before physically reseating it. To gracefully power down the drive, use the appropriate Cisco management platform:
  - Cisco Intersight Managed Mode (Cisco IMM)
  - Cisco UCS Manager (UCSM)
  - Cisco Integrated Management Controller (CIMC)

---

**Step 1** If you have not powered down the node, do so now.

- a) Use your node management software to gracefully power down the node.

See the appropriate Cisco management software documentation.

- b) If node power down through software is not available, you can power down the node by pressing the power button.

See [Status LEDs and Buttons, on page 61](#).

**Step 2** Choose the appropriate option:

- a) For a front-loading drive, see [Replacing a Front-Loading SAS/SATA Drive, on page 80](#)
- b) For a rear-loading drive, see [Replacing a Rear-Loading SAS/SATA Drive, on page 82](#)
- c) For a mid-mount drive, see [Replacing Mid-Mounted SAS/SATA Drives \(LFF Server\), on page 83](#)

**Note** While the drive is removed, it is a best practice to perform a visual inspection. Check the drive bay to ensure that no dust or debris is present. Also, check the connector on the back of the drive and the connector on the inside of the node for any obstructions or damage.



- Step 3** When the drive is correctly reseated, restart the node.
- Step 4** During node boot up, watch the drive's LEDs to verify correct operation.  
See [Status LEDs and Buttons](#), on page 61.
- Step 5** If reseating the drive does not clear the UBAD error, choose the appropriate option:
- Contact Cisco Systems for assistance with troubleshooting.
  - Begin an RMA of the errored drive.

---

## Replacing Front-Loading NVMe SSDs

This section is for replacing 2.5-inch or 3.5-inch form-factor NVMe solid-state drives (SSDs) in front-panel drive bays.

### Front-Loading NVMe SSD Population Guidelines

The front drive bay support for 2.5-inch NVMe SSDs differs by node PID:

- HX C240 M6 SFF 24 SAS/SATA—Small form-factor (SFF) drives, with 24-drive backplane. Drive bays 1 and 4 support 2.5-inch NVMe SSDs.
- HX C240 M6 24 NVMe—SFF drives, with 24-drive backplane. Drive bay 1 - 24 support *only* 2.5-inch NVMe SSDs.
- HX C240 M6 12 SAS/SATA plus optical—SFF drives, with 8-drive backplane and DVD drive option. Drive bays 1 and 4 support 2.5-inch NVMe SSDs.
- HX C240 M6 12 NVMe—SFF drives, with 12-drive backplane. Drive bay 1 - 12 support *only* 2.5-inch NVMe SSDs.
- HX C240 M6 LFF—Large form-factor (LFF) drives, with 12-drive backplane. Drive bays 1 - 4 support 2.5-inch NVMe SSDs. If you use 2.5-inch NVMe SSDs, a size-converter drive tray (HX-LFF-SFF-SLED2) is required for this version of the node.

### Front-Loading NVMe SSD Requirements and Restrictions

Observe these requirements:

- The node must have two CPUs. PCIe riser 2 is not available in a single-CPU system.
- PCIe cable. This is the cable that carries the PCIe signal from the front-panel drive backplane to PCIe riser 1B or 3B. The cable differs by node version:
  - For small form factor (SFF) drive versions of the node: CBL-NVME-C240SFF
  - For the large form factor (LFF) drive version of the node: CBL-NVME-C240LFF
- Hot-plug support must be enabled in the system BIOS. If you ordered the system with NVMe drives, hot-plug support is enabled at the factory.

Observe these restrictions:

- NVMe 2.5 SSDs support booting only in UEFI mode. Legacy boot is not supported. For instructions on setting up UEFI boot, see [Setting Up UEFI Mode Booting in the BIOS Setup Utility, on page 79](#) or [Setting Up UEFI Mode Booting in the Cisco IMC GUI, on page 80](#).
- You cannot control NVMe PCIe SSDs with a SAS RAID controller because NVMe SSDs interface with the node via the PCIe bus.
- You can combine NVMe SSDs in the same system, but the same partner brand must be used. For example, two *Intel* NVMe SFF 2.5-inch SSDs and two *HGST* SSDs is an invalid configuration.
- UEFI boot is supported in all supported operating systems. Hot-insertion and hot-removal are supported in all supported operating systems except VMWare ESXi.

## Enabling Hot-Plug Support in the System BIOS

Hot-plug (OS-informed hot-insertion and hot-removal) is disabled in the system BIOS by default.

- If the system was ordered with NVMe PCIe SSDs, the setting was enabled at the factory. No action is required.
- If you are adding NVMe PCIe SSDs after-factory, you must enable hot-plug support in the BIOS. See the following procedures.

### Enabling Hot-Plug Support Using the BIOS Setup Utility

- 
- Step 1** Enter the BIOS setup utility by pressing the **F2** key when prompted during bootup.
  - Step 2** Navigate to **Advanced > PCI Subsystem Settings > NVMe SSD Hot-Plug Support**.
  - Step 3** Set the value to **Enabled**.
  - Step 4** Save your changes and exit the utility.
- 

### Enabling Hot-Plug Support Using the Cisco IMC GUI

- 
- Step 1** Use a browser to log in to the Cisco IMC GUI for the node.
  - Step 2** Navigate to **Compute > BIOS > Advanced > PCI Configuration**.
  - Step 3** Set NVME SSD Hot-Plug Support to **Enabled**.
  - Step 4** Save your changes.
- 

## Replacing a Front-Loading NVMe SSD

This topic describes how to replace 2.5- or form-factor NVMe SSDs in the front-panel drive bays.




---

**Note** OS-surprise removal is not supported. OS-informed hot-insertion and hot-removal are supported on all supported operating systems except VMware ESXi.

---



---

**Note** OS-informed hot-insertion and hot-removal must be enabled in the system BIOS. See [Enabling Hot-Plug Support in the System BIOS, on page 88](#).

---

**Step 1** Remove an existing front-loading NVMe SSD:

- a) Shut down the NVMe SSD to initiate an OS-informed removal. Use your operating system interface to shut down the drive, and then observe the drive-tray LED:
  - Green—The drive is in use and functioning properly. Do not remove.
  - Green, blinking—the driver is unloading following a shutdown command. Do not remove.
  - Off—The drive is not in use and can be safely removed.
- b) Press the release button on the face of the drive tray.
- c) Grasp and open the ejector lever and then pull the drive tray out of the slot.
- d) Remove the four drive tray screws that secure the SSD to the tray and then lift the SSD out of the tray.

**Note** If this is the first time that front-loading NVMe SSDs are being installed in the node, you must install a PCIe cable with PCIe riser 2C.


**Step 2** Install a new front-loading NVMe SSD:

- a) Place a new SSD in the empty drive tray and install the four drive-tray screws.
- b) With the ejector lever on the drive tray open, insert the drive tray into the empty drive bay.
- c) Push the tray into the slot until it touches the backplane, and then close the ejector lever to lock the drive in place.

**Step 3** Observe the drive-tray LED and wait until it returns to solid green before accessing the drive:

- Off—The drive is not in use.
- Green, blinking—the driver is initializing following hot-plug insertion.
- Green—The drive is in use and functioning properly.

Figure 21: Replacing a Drive in a Drive Tray



1	Ejector lever	3	Drive tray screws (two on each side)
2	Release button	4	Drive removed from drive tray

## Cabling NVMe Drives 1 Through 4 (UCS C240 M6 24 SFF Drives Only)

When adding or replacing NVMe drives, two specific cables are required and are available through CBL-FNVME-240M6=.

- One NVMe C cable (74-126742-01), which connects front-loading drives 1 and 2 to the motherboard.
- One NVMe D cable (74-124687-01), which connects front-loading drives 3 and 4 to the motherboard.


Connectors are keyed, and they are different at each end of the cable to prevent improper installation. The backplane connector IDs are silk screened onto the interior of the node.

For this task, you need the appropriate cables.

**Before you begin**

Specific cables are required to add or replace front-loading NVMe drives 1 through 4 in Cisco HX C240 M6 24-SFF drive nodes. This procedure is for Cisco HX C240 M6 24-drive nodes only.


- Step 1** Remove the node top cover.  
See [Removing the Node Top Cover, on page 70](#).
- Step 2** Remove the fan tray.  
See [Removing the Fan Tray, on page 100](#)
- Step 3** Locate the NVMe backplane connectors.



540835


1	Backplane Connector, NVMe D	2	Backplane Connector, NVMe C
3	Motherboard connector, NVMe D	-	Motherboard connector, NVMe C

- Step 4** Orient the NVMe D cable correctly, lower it into place, and attach both ends.




**Step 5** Attach both ends of the NVMe D cable.

**Note** This cable must be installed first to allow the NVMe C cable to lie on top of it.



**Step 6** Orient the NVMe D cable correctly, lower it into place, and attach both ends.

**Note** The NVMe D cable lies on top of the NVMe C cable.



- Step 7** If drives are installed in any of slots 1 through 4, look at the drive LEDs to verify correct operation. See [Front-Panel LEDs](#), on page 62.
- Step 8** When drives are successfully booted up to runtime, reinstall the fan tray. See [Installing the Fan Tray](#), on page 101.
- Step 9** Replace the top cover.

## Replacing Rear-Loading NVMe SSDs

This section is for replacing 2.5-inch form-factor NVMe solid-state drives (SSDs) in rear-panel drive bays.



## Rear-Loading NVMe SSD Population Guidelines

The rear drive bay support differs by node PID and which type of RAID controller is used in the node for non-NVMe drives:

- HX C240 M6 24 SAS/SATA—Small form-factor (SFF) drives, with 24-drive backplane.
  - Hardware RAID—Rear drive bays support SAS or NVMe drives
- HX C240 M6 12 SAS/SATA—SFF drives, with 12-drive backplane.
  - Rear drive bays support only NVMe SSDs.
- HX C240 M6 24 NVMe—SFF drives, with 24-drive backplane.
  - Hardware RAID—Rear drive bays support NVMe drives only
- HX C240 M6 12 NVMe—SFF drives, with 12-drive backplane.
  - Hardware RAID—Rear drive bays support NVMe drives only.
- HX C240 M6 12 LFF—Large form-factor (LFF) drives, with 12-drive backplane.
  - Hardware RAID—Rear drive bays support SAS or NVMe drives
- The rear drive bay numbering follows the front-drive bay numbering in each node version:
  - 12-drive node—rear bays are numbered bays 103 and 104.
  - 24-drive node—rear bays are numbered bays 101 through 104.
- When populating drives, add drives to the lowest-numbered bays first.
- Drives are hot pluggable, but each drive requires a 10-second delay between hot removal and hot insertion.
- Keep an empty drive blanking tray in any unused bays to ensure proper airflow.

## Rear-Loading NVMe SSD Requirements and Restrictions

Observe these requirements:

- The node must have two CPUs. PCIe riser 2 is not available in a single-CPU system.
- PCIe riser 1A and 3A support NVMe rear drives.
- Rear PCIe cable and rear drive backplane.
- Hot-plug support must be enabled in the system BIOS. If you ordered the system with NVMe drives, hot-plug support is enabled at the factory.

Observe these restrictions:

- NVMe SSDs support booting only in UEFI mode. Legacy boot is not supported. For instructions on setting up UEFI boot, see [Setting Up UEFI Mode Booting in the BIOS Setup Utility, on page 79](#) or [Setting Up UEFI Mode Booting in the Cisco IMC GUI, on page 80](#).
- You cannot control NVMe PCIe SSDs with a SAS RAID controller because NVMe SSDs interface with the node via the PCIe bus.

- You can combine NVMe 2.5-inch SSDs in the same system, but the same partner brand must be used. For example, two *Intel* NVMe SFF 2.5-inch SSDs and two *HGST* SSDs is an invalid configuration.
- UEFI boot is supported in all supported operating systems. Hot-insertion and hot-removal are supported in all supported operating systems except VMWare ESXi.

## Replacing a Rear-Loading NVMe SSD

This topic describes how to replace 2.5-inch form-factor NVMe SSDs in the rear-panel drive bays.




---

**Note** OS-surprise removal is not supported. OS-informed hot-insertion and hot-removal are supported on all supported operating systems except VMware ESXi.

---




---

**Note** OS-informed hot-insertion and hot-removal must be enabled in the system BIOS. See [Enabling Hot-Plug Support in the System BIOS, on page 88](#).

---

### Step 1 Remove an existing rear-loading NVMe SSD:

- Shut down the NVMe SSD to initiate an OS-informed removal. Use your operating system interface to shut down the drive, and then observe the drive-tray LED:
  - Green—The drive is in use and functioning properly. Do not remove.
  - Green, blinking—the driver is unloading following a shutdown command. Do not remove.
  - Off—The drive is not in use and can be safely removed.
- Press the release button on the face of the drive tray.
- Grasp and open the ejector lever and then pull the drive tray out of the slot.
- Remove the four drive tray screws that secure the SSD to the tray and then lift the SSD out of the tray.

**Note** If this is the first time that rear-loading NVMe SSDs are being installed in the node, you must install PCIe riser 2B or 2C and a rear NVMe cable kit.


### Step 2 Install a new front-loading NVMe SSD:

- Place a new SSD in the empty drive tray and install the four drive-tray screws.
- With the ejector lever on the drive tray open, insert the drive tray into the empty drive bay.
- Push the tray into the slot until it touches the backplane, and then close the ejector lever to lock the drive in place.

### Step 3 Observe the drive-tray LED and wait until it returns to solid green before accessing the drive:

- Off—The drive is not in use.
- Green, blinking—the driver is initializing following hot-plug insertion.
- Green—The drive is in use and functioning properly.

Figure 22: Replacing a Drive in a Drive Tray



1	Ejector lever	3	Drive tray screws (two on each side)
2	Release button	4	Drive removed from drive tray

## Basic Troubleshooting: Reseating an NVMe Drive

Due to a drive-manufacturing issue, it is possible that a false positive UBAD error can occur on NVMe SSDs installed in the node.

- Only drives that are managed by the UCS MegaRAID controller are affected.
- Drives can be affected regardless where they are installed in the node (front-loaded, rear-loaded, and so on).
- Drives installed in all Cisco UCS C-Series nodes with M3 processors and later can be affected.

- Drives can be affected regardless of whether they are configured for hotplug or not.
- The UBAD error is not always valid, so the drive is not always defective or in need of repair or replacement. However, it is also possible that the error is correct, or the drive will need replacement.

**Before submitting the drive to the RMA process**, it is a best practice to reseat the drive. If the false UBAD error exists, reseating the drive can clear it. If successful, reseating the drive reduces inconvenience, cost, and service interruption, and optimizes your node uptime.




---

**Note** Reseat the drive only if a UBAD error occurs. Other errors are transient. You should not attempt diagnostics and troubleshooting without the assistance of Cisco personnel. Contact Cisco TAC for assistance with other drive errors.

---

To reseat the drive, see [Reseating an NVMe Drive, on page 98](#).

## Reseating an NVMe Drive

Sometimes, NVMe drives can throw a false UBAD error and reseating the drive can clear the error.

Use the following procedure to reseat the drive.




---

**Caution** This procedure requires powering down the node, which will cause a service interruption.

---

### Before you begin

Before attempting this procedure, be aware of the following:

- Before reseating the drive, it is a best practice to back up any data on it.
- When reseating the drive, make sure to reuse the same drive bay.
  - Do not move the drive to a different slot.
  - Do not move the drive to a different node.
  - If you do not reuse the same slot, the Cisco management software (for example, Cisco IMM) might require a rescan/rediscovery of the node.
- Regardless of the status of hotplug/hot insertion on the drive, it is a best practice to gracefully power down the drive before physically reseating it. To gracefully power down the drive, use the appropriate Cisco management platform:
  - Cisco Intersight Managed Mode (Cisco IMM)
  - Cisco UCS Manager (UCSM)
  - Cisco Integrated Management Controller (CIMC)

---

**Step 1** If you have not powered down the node, do so now.

- a) Use your node management software to gracefully power down the node.

See the appropriate Cisco management software documentation.

- b) If node power down through software is not available, you can power down the node by pressing the power button.

See [Status LEDs and Buttons, on page 61](#).

**Step 2** Choose the appropriate option:

- a) For a front-loading drive, see [Replacing a Front-Loading NVMe SSD, on page 88](#)
- b) For a rear-loading drive, see [Replacing a Rear-Loading NVMe SSD, on page 96](#)

**Note** While the drive is removed, it is a best practice to perform a visual inspection. Check the drive bay to ensure that no dust or debris is present. Also, check the connector on the back of the drive and the connector on the inside of the node for any obstructions or damage.

**Step 3** When the drive is correctly reseated, restart the node.

**Step 4** During node boot up, watch the drive's LEDs to verify correct operation.

See [Status LEDs and Buttons, on page 61](#).

**Step 5** If reseating the drive does not clear the UBAD error, choose the appropriate option:

- a) Contact Cisco Systems for assistance with troubleshooting.
- b) Begin an RMA of the errored drive.

---

## Replacing Fan Modules



**Tip**

There is a fault LED on the top of each fan module. This LED lights green when the fan is correctly seated and is operating OK. The LED lights amber when the fan has a fault or is not correctly seated.



**Caution**

You do not have to shut down or remove power from the node to replace fan modules because they are hot-swappable. However, to maintain proper cooling, do not operate the node for more than one minute with any fan module removed.

---

**Step 1** Remove an existing fan module:

- a) Slide the node out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.

**Caution** If you cannot safely view and access the component, remove the node from the rack.


- b) Remove the top cover from the node as described in [Removing the Node Top Cover, on page 70](#).
- c) Grasp and squeeze the fan module release latches on its top. Lift straight up to disengage its connector from the motherboard.

**Step 2** Install a new fan module:

- a) Set the new fan module in place. The arrow printed on the top of the fan module should point toward the rear of the node.

- b) Press down gently on the fan module to fully engage it with the connector on the motherboard.
- c) Replace the top cover to the node.
- d) Replace the node in the rack, replace cables, and then fully power on the node by pressing the Power button.

**Figure 23: Top View of Fan Module**




<b>1</b>	Fan module release latches	<b>2</b>	Fan module fault LED
----------	----------------------------	----------	----------------------

## Removing the Fan Tray

The fan tray can be removed either with all fan modules in place, or when some, or all, of the fan modules have been removed.

- Step 1** Remove the screws that secure the fan tray to the chassis.
- a) Locate the screws that secure the fan tray to the node.
  - b) Using a #2 Phillips screwdriver, loosen the screws.
- Step 2** Disconnect the fan tray cable from the fan tray, leaving the motherboard connection in place.
- Step 3** Remove the fan tray from the node.
- a) Grasp the handles at the top of the fan tray.
  - b) Holding the fan tray level and making sure that the fan tray cable does not obstruct removal, lift the fan tray up until it is removed from the chassis.



### What to do next

Reinsert the fan tray into the chassis. See [Installing the Fan Tray, on page 101](#).


## Installing the Fan Tray

You can install the fan tray with or without fans installed. Use the following procedure to install the fan tray.

### Step 1

Install the fan tray.

- a) Align the fan tray with the guides on the inside of the chassis.
- b) Make sure that the fan tray cable is out of the way and will not obstruct installation.
- c) Holding the fan tray by the handles, slide it into place in the chassis.



- Step 2** Reconnect the fan tray cable.
- Step 3** Close the top cover, or perform additional procedures, if needed.

## Replacing CPUs and Heatsinks

This section contains the following topics:

- [CPU Configuration Rules, on page 102](#)
- [Tools Required For CPU Replacement, on page 104](#)
- [Removing CPUs and Heat Sinks, on page 104](#)
- [Installing the CPUs and Heatsinks, on page 109](#)
- [Additional CPU-Related Parts to Order with RMA Replacement CPUs, on page 113](#)

### CPU Configuration Rules

This node has two CPU sockets on the motherboard. Each CPU supports 8 DIMM channels (16 DIMM slots). See [DIMM Population Rules and Memory Performance Guidelines](#), on page 115.

- The node can operate with one CPU, or two identical CPUs installed.
- The minimum configuration is that the node must have at least CPU 1 installed. Install CPU 1 first, and then CPU 2.




- The following restrictions apply when using a single-CPU configuration:
  - Any unused CPU socket must have the socket dust cover from the factory in place.
  - The maximum number of DIMMs is 16 (only CPU 1 channels A through H).
- Two different form factors exist for heatsinks, a low profile and a high profile. The node can be ordered with either, but you cannot mix high- and low-profile CPUs and heatsinks in the same node. A single node must have all of one type.

The CPU and heatsink installation procedure is different depending on the type of heatsink used in your node.


- Low profile (HX-HSLP-M6), which has 4 T30 Torx screws on the main heatsink, and 2 Phillips-head screws on the extended heatsink.

This heat sink is required for nodes that contain one or more GPUs.

This heat sink is not supported on C240 M6 LFF nodes.



- High profile (HX-HSHP-240M6), which has 4 T30 Torx screws.



## Tools Required For CPU Replacement

You need the following tools and equipment for this procedure:

- T-30 Torx driver—Supplied with replacement CPU.
- #1 flat-head screwdriver—Supplied with replacement CPU.
- CPU assembly tool—Supplied with replacement CPU. Orderable separately as Cisco PID HX-CPUAT=.
- Heatsink cleaning kit—Supplied with replacement CPU. Orderable separately as Cisco PID HX-HSCK=. One cleaning kit can clean up to four CPUs.
- Thermal interface material (TIM)—Syringe supplied with replacement CPU. Use only if you are reusing your existing heatsink (new heatsinks have a pre-applied pad of TIM). Orderable separately as Cisco PID HX-CPU-TIM=.

One TIM kit covers one CPU.

See also [Additional CPU-Related Parts to Order with RMA Replacement CPUs, on page 113](#).


## Removing CPUs and Heat Sinks

Use the following procedure to remove an installed CPU and heatsink from the blade server node. With this procedure, you will remove the CPU from the motherboard, disassemble individual components, then place the CPU and heatsink into the fixture that came with the CPU.

**Step 1** Choose the appropriate method to loosen the securing screws, based on the whether the CPU has a high-profile or low-profile heatsink.

- For a CPU with a high-profile heatsink, proceed to step a.

- For a CPU with a low-profile heatsink, skip to step 2.
  - a) Using a T30 Torx driver, loosen all the securing nuts.
  - b) Push the rotating wires towards each other to move them to the unlocked position.
- Caution** Make sure that the rotating wires are as far inward as possible. When fully unlocked, the bottom of the rotating wire disengages and allows the removal of the CPU assembly. If the rotating wires are not fully in the unlocked position, you can feel resistance when attempting to remove the CPU assembly.
- c) Grasp the CPU and heatsink along the edge of the carrier and lift the CPU and heatsink off of the motherboard.
- Caution** While lifting the CPU assembly, make sure not to bend the heatsink fins. Also, if you feel any resistance when lifting the CPU assembly, verify that the rotating wires are completely in the unlocked position.



- d) Go to step 3.

## Step 2 Remove the CPU.

- a) Using a #2 Phillips screwdriver, loosen the two Phillips head screws for the extended heatsink.
- b) Using a T30 Torx driver, loosen the four Torx securing nuts.
- c) Push the rotating wires towards each other to move them to the unlocked position.

**Caution** Make sure that the rotating wires are as far inward as possible. When fully unlocked, the bottom of the rotating wire disengages and allows the removal of the CPU assembly. If the rotating wires are not fully in the unlocked position, you can feel resistance when attempting to remove the CPU assembly.

- d) Grasp the CPU and heatsink along the edge of the carrier and lift the CPU and heatsink off of the motherboard.

**Caution** While lifting the CPU assembly, make sure not to bend the heatsink fins. Also, if you feel any resistance when lifting the CPU assembly, verify that the rotating wires are completely in the unlocked position.

- e) Go to step 3.


**Step 3** Put the CPU assembly on a rubberized mat or other ESD-safe work surface.

When placing the CPU on the work surface, the heatsink label should be facing up. Do not rotate the CPU assembly upside down.

**Step 4** Attach a CPU dust cover (HX-CPU-M6-CVR⇒) to the CPU socket.

- a) Align the posts on the CPU bolstering plate with the cutouts at the corners of the dust cover.  
b) Lower the dust cover and simultaneously press down on the edges until it snaps into place over the CPU socket.

**Caution** Do not press down in the center of the dust cover!




**Step 5** Detach the CPU from the CPU carrier by disengaging CPU clips and using the TIM breaker.

- a) Turn the CPU assembly upside down, so that the heatsink is pointing down.

This step enables access to the CPU securing clips.

- b) Gently lift the TIM breaker in a 90-degree upward arc to partially disengage the CPU clips on this end of the CPU carrier.




- c) Lower the TIM breaker into the u-shaped securing clip to allow easier access to the CPU carrier.

**Note** Make sure that the TIM breaker is completely seated in the securing clip.

- d) Gently pull up on the outer edge of the CPU carrier (2) so that you can disengage the second pair of CPU clips near both ends of the TIM breaker.

**Caution** Be careful when flexing the CPU carrier! If you apply too much force you can damage the CPU carrier. Flex the carrier only enough to release the CPU clips. Make sure to watch the clips while performing this step so that you can see when they disengage from the CPU carrier.

- e) Gently pull up on the outer edge of the CPU carrier so that you can disengage the pair of CPU clips (3 in the following illustration) which are opposite the TIM breaker.
- f) Grasp the CPU carrier along the short edges and lift it straight up to remove it from the heatsink.




**Step 6** Transfer the CPU and carrier to the fixture.

- a) When all the CPU clips are disengaged, grasp the carrier, and lift it and the CPU to detach them from the heatsink.

**Note** If the carrier and CPU do not lift off of the heatsink, attempt to disengage the CPU clips again.

- b) Flip the CPU and carrier right-side up so that the words PRESS HERE are visible.  
 c) Align the posts on the fixture and the pin 1 locations on the CPU carrier and the fixture (1 in the following illustration).  
 d) Lower the CPU and CPU carrier onto the fixture.



**Step 7** Use the provided cleaning kit (HX-HSCK) to remove all of the thermal interface barrier (thermal grease) from the CPU, CPU carrier, and heatsink.

**Important** Make sure to use only the Cisco-provided cleaning kit, and make sure that no thermal grease is left on any surfaces, corners, or crevices. The CPU, CPU carrier, and heatsink must be completely clean.

### What to do next

Choose the appropriate option:


- If you will be installing a CPU, go to [Installing the CPUs and Heatsinks, on page 109](#).
- If you will not be installing a CPU, verify that a CPU socket cover is installed. This option is valid only for CPU socket 2 because CPU socket 1 must always be populated in a runtime deployment.

## Installing the CPUs and Heatsinks

Use this procedure to install a CPU if you have removed one, or if you are installing a CPU in an empty CPU socket. To install the CPU, you will move the CPU to the fixture, then attach the CPU assembly to the CPU socket on the node mother board.

**Step 1** Remove the CPU socket dust cover (HX-CPU-M6-CVR=) on the node motherboard.

- a) Push the two vertical tabs inward to disengage the dust cover.
- b) While holding the tabs in, lift the dust cover up to remove it.



- c) Store the dust cover for future use.

**Caution** Do not leave an empty CPU socket uncovered. If a CPU socket does not contain a CPU, you must install a CPU dust cover.

**Step 2** Grasp the CPU fixture on the edges labeled PRESS, lift it out of the tray, and place the CPU assembly on an ESD-safe work surface.


**Step 3** Apply new TIM.

**Note** The heatsink must have new TIM on the heatsink-to-CPU surface to ensure proper cooling and performance.

- If you are installing a new heatsink, it is shipped with a pre-applied pad of TIM. Go to step 4.
  - If you are reusing a heatsink, you must remove the old TIM from the heatsink and then apply new TIM to the CPU surface from the supplied syringe. Continue with step a below.
- a) Apply the Bottle #1 cleaning solution that is included with the heatsink cleaning kit (UCSX-HSCK=), as well as the spare CPU package, to the old TIM on the heatsink and let it soak for a least 15 seconds.
  - b) Wipe all of the TIM off the heatsink using the soft cloth that is included with the heatsink cleaning kit. Be careful to avoid scratching the heatsink surface.
  - c) Completely clean the bottom surface of the heatsink using Bottle #2 to prepare the heatsink for installation.
  - d) Using the syringe of TIM provided with the new CPU (UCS-CPU-TIM=), apply 1.5 cubic centimeters (1.5 ml) of thermal interface material to the top of the CPU. Use the pattern shown in the following figure to ensure even coverage.



**Figure 24: Thermal Interface Material Application Pattern**



**Caution** Use only the correct heatsink for your CPU. CPU 1 uses heatsink UCSB-HS-M6-R and CPU 2 uses heatsink HX-HS-M6-F.

**Step 4** Attach the heatsink to the CPU fixture.


- a) Make sure the rotating wires are in the unlocked position so that the feet of the wires do not impede installing the heatsink.
- b) Grasp the heatsink by the fins and align the pin 1 location of the heatsink with the pin 1 location on the CPU fixture, then lower the heatsink onto the CPU fixture.

**Step 5** Install the CPU assembly onto the CPU motherboard socket.


- a) Push the rotating wires (1 in the following image) to the unlocked position so that they do not obstruct installation.
- b) Grasp the heatsink by the fins, align the pin 1 location on the heatsink with the pin 1 location on the CPU socket (2 in the following image), then seat the heatsink onto the CPU socket.
- c) Holding the CPU assembly level, lower it onto the CPU socket.
- d) Push the rotating wires away from each other to lock the CPU assembly into the CPU socket.

**Caution** Make sure that you close the rotating wires completely before using the Torx driver to tighten the securing nuts.

- e) Choose the appropriate option to secure the CPU to the socket.
  - For a CPU with a high-profile heatsink, set the T30 Torx driver to 12 in-lb of torque and tighten the 4 securing nuts to secure the CPU to the motherboard (4).



- For a CPU with a low-profile heatsink, set the T30 Torx driver to 12 in-lb of torque and tighten the 4 securing nuts to secure the CPU to the motherboard (3) first. Then, set the torque driver to 6 in-lb of torque and tighten the two Phillips head screws for the extended heatsink (4).



## Additional CPU-Related Parts to Order with RMA Replacement CPUs

When a return material authorization (RMA) of the CPU is done on a Cisco HX C-Series node, additional parts might not be included with the CPU spare. The TAC engineer might need to add the additional parts to the RMA to help ensure a successful replacement.



**Note** The following items apply to CPU *replacement* scenarios. If you are replacing a system chassis and *moving* existing CPUs to the new chassis, you do not have to separate the heatsink from the CPU. See [Additional CPU-Related Parts to Order with RMA Replacement System Chassis](#), on page 114.

- Scenario 1—You are reusing the existing heatsinks:

- Heat sink cleaning kit (HX-HSCK=)  
One cleaning kit can clean up to four CPUs.
- Thermal interface material (TIM) kit for M6 nodes (HX-CPU-TIM=)  
One TIM kit covers one CPU.
- Scenario 2—You are replacing the existing heatsinks:




---

**Caution** Use only the correct heatsink for your CPUs to ensure proper cooling. There are two different heatsinks, a low profile (HX-HSLP-M6) which is used with GPUs, and a high-profile (HX-HSLP-M6).

---

- New heatsinks have a pre-applied pad of TIM.
- Heat sink cleaning kit (HX-HSCK=)  
One cleaning kit can clean up to four CPUs.
- Scenario 3—You have a damaged CPU carrier (the plastic frame around the CPU):
  - CPU Carrier: HX-M6-CPU-CAR=
  - #1 flat-head screwdriver (for separating the CPU from the heatsink)
  - Heatsink cleaning kit (HX-HSCK=)  
One cleaning kit can clean up to four CPUs.
  - Thermal interface material (TIM) kit for M6 nodes (HX-CPU-TIM=)  
One TIM kit covers one CPU.

A CPU heat sink cleaning kit is good for up to four CPU and heat sink cleanings. The cleaning kit contains two bottles of solution, one to clean the CPU and heat sink of old TIM and the other to prepare the surface of the heat sink.

New heat sink spares come with a pre-applied pad of TIM. It is important to clean any old TIM off of the CPU surface prior to installing the heat sinks. Therefore, even when you are ordering new heat sinks, you must order the heat sink cleaning kit.

## Additional CPU-Related Parts to Order with RMA Replacement System Chassis

When a return material authorization (RMA) of the system chassis is done on a Cisco HyperFlex C-Series server node, you move existing CPUs to the new chassis.




---

**Note** Unlike previous generation CPUs, the M6 node CPUs do not require you to separate the heatsink from the CPU when you *move* the CPU-heatsink assembly. Therefore, no additional heatsink cleaning kit or thermal-interface material items are required.

---

- The only tool required for moving a CPU/heatsink assembly is a T-30 Torx driver.

# Replacing Memory DIMMs



**Caution** DIMMs and their sockets are fragile and must be handled with care to avoid damage during installation.



**Caution** Cisco does not support third-party DIMMs. Using non-Cisco DIMMs in the node might result in system problems or damage to the motherboard.



**Note** To ensure the best node performance, it is important that you are familiar with memory performance guidelines and population rules before you install or replace DIMMs.


## DIMM Population Rules and Memory Performance Guidelines

The following sections provide partial information for memory usage, mixing, and population guidelines.

### DIMM Slot Numbering

The following figure shows the numbering of the DIMM slots on the motherboard.

*Figure 25: DIMM Slot Numbering*



309367

### DIMM Population Rules

Observe the following guidelines when installing or replacing DIMMs for maximum performance:

- Each CPU supports eight memory channels, A through H.

- CPU 1 supports channels P1 A1, P1 A2, P1 B1, P1 B2, P1 C1, P1 C2, P1 D1, P1 D2, P1 E1, P1 E2, P1 F1, P1 F2, P1 G1, P1 G2, P1 H1, and P1 H2.
- CPU 2 supports channels P2 A1, P2 A2, P2 B1, P2 B2, P2 C1, P2 C2, P2 D1, P2 D2, P2 E1, P2 E2, P2 F1, P2 F2, P2 G1, P2 G2, P2 H1, and P2 H2.
- Each channel has two DIMM sockets (for example, channel A = slots A1, A2).
- In a single-CPU configuration, populate the channels for CPU1 only (P1 A1 through P1 H2).
- For optimal performance, populate DIMMs in the order shown in the following table, depending on the number of CPUs and the number of DIMMs per CPU. If your node has two CPUs, balance DIMMs evenly across the two CPUs as shown in the table.



**Note** The sections below list recommended configurations. Using 5, 7, 9, 10, or 11 DIMMs per CPU is not recommended.

### Memory Population Order

The Cisco UCS C240 M6 node has two memory options, DIMMs only or DIMMs plus Intel Optane Persistent Memory 200 series memory.

Memory slots are color coded, blue and black. The color-coded channel population order is blue slots first, then black.

The following tables show the memory population order for each memory option.

**Table 15: DIMMs Population Order**

Number of DDR4 DIMMs per CPU (Recommended Configurations)	Populate CPU 1 Slot		Populate CPU2 Slots	
	P1 Blue #1 Slots	P1 Black #2 Slots	P2 Blue #1 Slots	P2 Black #2 Slots
1	A1	-	A1	
2	(A1, E1)	-	(A1, E1)	
4	(A1, C1); (E1, G1)	-	(A1, C1); (E1, G1)	
6	(A1, C1,D1); (E1, G1, H1)	-	(A1, C1,D1); (E1, G1, H1)	
8	(A1, B1,C1, D1, E1, F1, G1, H1)	-	(A1, B1,C1, D1, E1, F1, G1, H1)	
12	A1, C1, D1, E1, G1, H1	A2, C2, D2, E2, G2, H2	A1, C1, D1, E1, G1, H1	A2, C2, D2, E2, G2, H2
16	All populated (A1 through H1)	All populated (A2 through H2)	All populated (A1 through H1)	All populated (A2 through H2)

Table 16: DIMM Plus Intel Optane Persistent Memory 200 Series Memory Population Order

Total Number of DIMMs per CPU	DDR4 DIMM Slot	Intel Optane Persistent Memory 200 Series DIMM Slot
4+4 DIMM	A1, C1, E1, G1,	B1, D1, F1, H1
8+1 DIMMs	A1, B1, C1, D1, E1, F1, G1, H1	A1
8+4 DIMMs	A1, B1, C1, D1, E1, F1, G1, H1	A1, C1, E1, G1
8+8 DIMMs	A0, B0, C0, D0, E0, F0, G0, H0	A1, B1, C1, D1, E1, F1, G1, H1

### Memory Mirroring

The CPUs in the node support memory mirroring only when an even number of channels are populated with DIMMs. If one or three channels are populated with DIMMs, memory mirroring is automatically disabled.

Memory mirroring reduces the amount of memory available by 50 percent because only one of the two populated channels provides data. The second, duplicate channel provides redundancy.

## Replacing DIMMs

### Identifying a Faulty DIMM

Each DIMM socket has a corresponding DIMM fault LED, directly in front of the DIMM socket. See [Internal Diagnostic LEDs, on page 67](#) for the locations of these LEDs. When the node is in standby power mode, these LEDs light amber to indicate a faulty DIMM.

#### Step 1

Remove an existing DIMM:

- a) Shut down and remove power from the node as described in [Shutting Down and Removing Power From the Node, on page 69](#).
- b) Slide the node out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.
 

**Caution** If you cannot safely view and access the component, remove the node from the rack.
- c) Remove the top cover from the node as described in [Removing the Node Top Cover, on page 70](#).
- d) Remove the air baffle that covers the front ends of the DIMM slots to provide clearance.
- e) Locate the DIMM that you are removing, and then open the ejector levers at each end of its DIMM slot.

#### Step 2

Install a new DIMM:

- Note** Before installing DIMMs, see the memory population rules for this node: [DIMM Population Rules and Memory Performance Guidelines , on page 115](#).
- a) Align the new DIMM with the empty slot on the motherboard. Use the alignment feature in the DIMM slot to correctly orient the DIMM.
  - b) Push down evenly on the top corners of the DIMM until it is fully seated and the ejector levers on both ends lock into place.
  - c) Replace the top cover to the node.

- d) Replace the node in the rack, replace cables, and then fully power on the node by pressing the Power button.

---

## Replacing Intel Optane DC Persistent Memory Modules

This topic contains information for replacing Intel Optane Data Center Persistent memory modules (DCPMMs), including population rules and methods for verifying functionality. DCPMMs have the same form-factor as DDR4 DIMMs and they install to DIMM slots.



---

**Caution** DCPMMs and their sockets are fragile and must be handled with care to avoid damage during installation.

---



---

**Note** To ensure the best node performance, it is important that you are familiar with memory performance guidelines and population rules before you install or replace DCPMMs.

---

DCPMMs can be configured to operate in one of the following modes:

- **Memory Mode:** The module operates as 100% memory module. Data is volatile and DRAM acts as a cache for DCPMMs.
- **App Direct Mode:** The module operates as a solid-state disk storage device. Data is saved and is non-volatile.

## Intel Optane DC Persistent Memory Module Population Rules and Performance Guidelines


This topic describes the rules and guidelines for maximum memory performance when using Intel Optane DC persistent memory modules (DCPMMs) with DDR4 DRAM DIMMs.

### DIMM Slot Numbering

The following figure shows the numbering of the DIMM slots on the node motherboard.



Figure 26: DIMM Slot Numbering



309367

### Configuration Rules

Observe the following rules and guidelines:

- To use DCPMMs in this node, two CPUs must be installed.
- The DCPMMs run at 2666 MHz. If you have 2933 MHz RDIMMs or LRDIMMs in the node and you add DCPMMs, the main memory speed clocks down to 2666 MHz to match the speed of the DCPMMs.
- Each DCPMM draws 18 W sustained, with a 20 W peak.
- When using DCPMMs in a node:
  - The DDR4 DIMMs installed in the node must all be the same size.
  - The DCPMMs installed in the node must all be the same size and must have the same SKU.

## Installing Intel Optane DC Persistent Memory Modules



**Note** DCPMM configuration is always applied to all DCPMMs in a region, including a replacement DCPMM. You cannot provision a specific replacement DCPMM on a preconfigured server.

### Step 1

Remove an existing DCPMM:

- Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Node, on page 69](#).
- Slide the server out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.

**Caution** If you cannot safely view and access the component, remove the server from the rack.

- c) Remove the top cover from the server as described in [Removing the Node Top Cover, on page 70](#).
- d) Remove the air baffle that covers the front ends of the DIMM slots to provide clearance.

**Caution** If you are moving DCPMMs with active data (persistent memory) from one server to another as in an RMA situation, each DCPMM must be installed to the identical position in the new server. Note the positions of each DCPMM or temporarily label them when removing them from the old server.

- e) Locate the DCPMM that you are removing, and then open the ejector levers at each end of its DIMM slot.

## Step 2

Install a new DCPMM:

**Note** Before installing DCPMMs, see the population rules for this server: [Intel Optane DC Persistent Memory Module Population Rules and Performance Guidelines, on page 118](#).

- a) Align the new DCPMM with the empty slot on the motherboard. Use the alignment feature in the DIMM slot to correctly orient the DCPMM.
- b) Push down evenly on the top corners of the DCPMM until it is fully seated and the ejector levers on both ends lock into place.
- c) Reinstall the air baffle.
- d) Replace the top cover to the server.
- e) Replace the server in the rack, replace cables, and then fully power on the server by pressing the Power button.

## Step 3

Perform post-installation actions:

- If the existing configuration is in 100% Memory mode, and the new DCPMM is also in 100% Memory mode (the factory default), the only action is to ensure that all DCPMMs are at the latest, matching firmware level.
- If the existing configuration is fully or partly in App-Direct mode and new DCPMM is also in App-Direct mode, then ensure that all DCPMMs are at the latest matching firmware level and also re-provision the DCPMMs by creating a new goal.
- If the existing configuration and the new DCPMM are in different modes, then ensure that all DCPMMs are at the latest matching firmware level and also re-provision the DCPMMs by creating a new goal.

There are a number of tools for configuring goals, regions, and namespaces.

- To use the server's BIOS Setup Utility, see [Node BIOS Setup Utility Menu for DCPMM, on page 120](#).
- To use Cisco IMC or Cisco UCS Manager, see the [Cisco UCS: Configuring and Managing Intel Optane DC Persistent Memory Modules](#) guide.

## Node BIOS Setup Utility Menu for DCPMM



**Caution** Potential data loss: If you change the mode of a currently installed DCPMM from App Direct to Memory Mode, any data in persistent memory is deleted.

DCPMMs can be configured by using the node's BIOS Setup Utility, Cisco IMC, Cisco UCS Manager, or OS-related utilities.

- To use the BIOS Setup Utility, see the section below.
- To use Cisco IMC, see the configuration guides for Cisco IMC 4.0(4) or later: [Cisco IMC CLI and GUI Configuration Guides](#)
- To use Cisco UCS Manager, see the configuration guides for Cisco UCS Manager 4.0(4) or later: [Cisco UCS Manager CLI and GUI Configuration Guides](#)

The node BIOS Setup Utility includes menus for DCPMMs. They can be used to view or configure DCPMM regions, goals, and namespaces, and to update DCPMM firmware.

To open the BIOS Setup Utility, press **F2** when prompted onscreen during a system boot.

The DCPMM menu is on the Advanced tab of the utility:

**Advanced > Intel Optane DC Persistent Memory Configuration**

From this tab, you can access other menus:

- **DIMMs:** Displays the installed DCPMMs. From this page, you can update DCPMM firmware and configure other DCPMM parameters.
  - Monitor health
  - Update firmware
  - Configure security
 

You can enable security mode and set a password so that the DCPMM configuration is locked. When you set a password, it applies to all installed DCPMMs. Security mode is disabled by default.
  - Configure data policy

- **Regions:** Displays regions and their persistent memory types. When using App Direct mode with interleaving, the number of regions is equal to the number of CPU sockets in the node. When using App Direct mode without interleaving, the number of regions is equal to the number of DCPMMs in the node.

From the Regions page, you can configure memory goals that tell the DCPMM how to allocate resources.

- Create goal config
- **Namespaces:** Displays namespaces and allows you to create or delete them when persistent memory is used. Namespaces can also be created when creating goals. A namespace provisioning of persistent memory applies only to the selected region.
 

Existing namespace attributes such as the size cannot be modified. You can only add or delete namespaces.
- **Total capacity:** Displays the total DCPMM resource allocation across the node.

### Updating the DCPMM Firmware Using the BIOS Setup Utility

You can update the DCPMM firmware from the BIOS Setup Utility if you know the path to the .bin files. The firmware update is applied to all installed DCPMMs.

1. Navigate to **Advanced > Intel Optane DC Persistent Memory Configuration > DIMMs > Update firmware**
2. Under **File:**, provide the file path to the .bin file.
3. Select **Update**.

## Replacing a Mini-Storage Module

The mini-storage module plugs into a vertical riser card that attaches to motherboard by two captive screws to provide additional internal storage. The module is an M.2 SSD Carrier which provides two M.2 form-factor SSD sockets. See also [Replacing a Boot-Optimized M.2 RAID Controller Module, on page 158](#).



---

**Note** The Cisco IMC firmware does not include an out-of-band management interface for the M.2 drives installed in the M.2 version of this mini-storage module (HX-MSTOR-M2). The M.2 drives are not listed in Cisco IMC inventory, nor can they be managed by Cisco IMC. This is expected behavior.

---

## Replacing a Mini-Storage Module Carrier

This topic describes how to remove and replace a mini-storage module carrier. The carrier sits in an M.2 vertical riser card that is attached to the motherboard by two captive screws.

The carrier has one media socket on its top and one socket on its underside. Use the following procedure for an M.2 SSD mini-storage module carrier.

- 
- Step 1** Shut down and remove power from the node as described in [Shutting Down and Removing Power From the Node, on page 69](#).
- Step 2** Slide the node out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.
- Caution** If you cannot safely view and access the component, remove the node from the rack.
- Step 3** Remove the top cover from the node as described in [Removing the Node Top Cover, on page 70](#).
- Step 4** Locate the mini-storage module carrier in its socket between PCIe riser 2 and 3.
- Step 5** Using a Phillips screwdriver, loosen each of the captive screws and lift the M.2 riser out of the node.
- Step 6** Remove a carrier from its socket:
- Using a Phillips screwdriver, loosen the screw that holds the module to the carrier.
  - Push outward on the securing clips that holds each end of the carrier.
  - Lift both ends of the carrier to disengage it from the socket on the motherboard.
  - Set the carrier on an anti-static surface.
- Step 7** Install a carrier to its socket:
- Position carrier over socket, with the carrier's connector facing down. Two alignment pegs must match with two holes on the carrier.
  - Gently push down the socket end of the carrier so that the two pegs go through the two holes on the carrier.
  - Push down on the carrier so that the securing clips click over it at both ends.
- Step 8** Replace the top cover to the node.
- Step 9** Replace the node in the rack, replace cables, and then fully power on the node by pressing the Power button.
-

## Replacing an M.2 SSD in a Mini-Storage Carrier For M.2

This topic describes how to remove and replace an M.2 SATA or NVMe SSD in a mini-storage carrier for M.2 (PID HX-MSTOR-M2). The carrier has one M.2 SSD socket on its top and one socket on its underside.

### Population Rules For Mini-Storage M.2 SSDs

- Both M.2 SSDs must be either SATA or NVMe; do not mix types in the carrier.
- You can use one or two M.2 SSDs in the carrier.
- M.2 socket 1 is on the top side of the carrier; M.2 socket 2 is on the underside of the carrier (the same side as the carrier's motherboard connector).

- 
- Step 1** Power off the node and then remove the mini-storage module carrier from the node as described in [Replacing a Mini-Storage Module Carrier, on page 122](#).
- Step 2** Remove an M.2 SSD:
- a) Use a #1 Phillips-head screwdriver to remove the single screw that secures the M.2 SSD to the carrier.
  - b) Remove the M.2 SSD from its socket on the carrier.
- Step 3** Install a new M.2 SSD:
- a) Insert the new M.2 SSD connector-end into the socket on the carrier with its label side facing up.
  - b) Press the M.2 SSD flat against the carrier.
  - c) Install the single screw that secures the end of the M.2 SSD to the carrier.
- Step 4** Install the mini-storage module carrier back into the node and then power it on as described in [Replacing a Mini-Storage Module Carrier, on page 122](#).
- 

## Replacing the RTC Battery

**Warning**

There is danger of explosion if the battery is replaced incorrectly. Replace the battery only with the same or equivalent type recommended by the manufacturer. Dispose of used batteries according to the manufacturer's instructions.

[Statement 1015]

**Warning**

**Recyclers:** Do not shred the battery! Make sure you dispose of the battery according to appropriate regulations for your country or locale.

The real-time clock (RTC) battery retains system settings when the node is disconnected from power. The battery type is CR2032. Cisco supports the industry-standard CR2032 battery, which can be purchased from most electronic stores.

- 
- Step 1** Remove the RTC battery:

- a) Shut down and remove power from the node as described in [Shutting Down and Removing Power From the Node, on page 69](#).
- b) Slide the node out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.

**Caution** If you cannot safely view and access the component, remove the node from the rack.

- c) Remove the top cover from the node as described in [Removing the Node Top Cover, on page 70](#).
- d) Remove PCIe riser 1 from the node to provide clearance to the RTC battery socket that is on the motherboard. See [Replacing a PCIe Riser, on page 129](#).
- e) Locate the horizontal RTC battery socket.
- f) Remove the battery from the socket on the motherboard. Gently pry the securing clip to the side to provide clearance, then lift up on the battery.

**Step 2** Install a new RTC battery:

- a) Insert the battery into its socket and press down until it clicks in place under the clip.

**Note** The positive side of the battery marked “3V+” should face up.

- b) Replace PCIe riser 1 to the node. See [Replacing a PCIe Riser, on page 129](#).
- c) Replace the top cover to the node.
- d) Replace the node in the rack, replace cables, and then fully power on the node by pressing the Power button.

## Replacing Power Supplies

When two power supplies are installed they are redundant as 1+1 by default, but they also support cold redundancy mode. Cold redundancy (CR) suspends power delivery on one or more power supplies and forces the remainder of the load to be supplied by the active PSU(s). As a result, total power efficiency is improved by best utilizing the PSU efficiency when compared to load characteristics.

- See also [Power Specifications, on page 176](#)
- See also [Rear-Panel LEDs, on page 65](#) for information about the power supply LEDs.

This section includes procedures for replacing AC and DC power supply units.

## Supported Power Supplies

The Cisco HyperFlex C240 M6 supports the following power supplies.



**Caution** Do not mix PSU types in the same node. PSU must be the same type and wattage.

For detailed information, see [Power Specifications, on page 176](#).

PSU Type	Supported In	Notes
1050 W AC	All Cisco HyperFlex C240 M6 models	One power supply is mandatory; one more can be added for 1 + 1 redundancy as long power supplies are the same.

PSU Type	Supported In	Notes
1050 W DC	All Cisco HyperFlex C240 M6 models	One power supply is mandatory; one more can be added for 1 + 1 redundancy as long power supplies are the same.
1600 W AC	All Cisco HyperFlex C240 M6 models	One power supply is mandatory; one more can be added for 1 + 1 redundancy as long power supplies are the same.
2300 W AC	All Cisco HyperFlex C240 M6 models	One power supply is mandatory; one more can be added for 1 + 1 redundancy as long power supplies are the same.

## Replacing AC Power Supplies



**Note** If you have ordered a node with power supply redundancy (two power supplies), you do not have to power off the node to replace a power supply because they are redundant as 1+1.



**Note** Do not mix power supply types or wattages in the node. Both power supplies must be identical.



**Caution** DO NOT interchange power supplies of Cisco HyperFlex C240 M5 server nodes and Cisco HyperFlex C240 SD M5 server nodes with the power supplies of the Cisco HyperFlex C240 M6 node.

**Step 1** Remove the power supply that you are replacing or a blank panel from an empty bay:

- a) Perform one of the following actions:
  - If your node has only one power supply, shut down and remove power from the node as described in [Shutting Down and Removing Power From the Node, on page 69](#).
  - If your node has two power supplies, you do not have to shut down the node.
- b) Remove the power cord from the power supply that you are replacing.
- c) Grasp the power supply handle while pinching the release lever toward the handle.
- d) Pull the power supply out of the bay.

**Step 2** Install a new power supply:

- a) Grasp the power supply handle and insert the new power supply into the empty bay.
- b) Push the power supply into the bay until the release lever locks.
- c) Connect the power cord to the new power supply.
- d) Only if you shut down the node, press the Power button to boot the node to main power mode.

## Replacing DC Power Supplies



**Note** This procedure is for replacing DC power supplies in a node that already has DC power supplies installed. If you are installing DC power supplies to the node for the first time, see [Installing DC Power Supplies \(First Time Installation\)](#), on page 127.



**Warning** A readily accessible two-poled disconnect device must be incorporated in the fixed wiring.

Statement 1022



**Warning** This product requires short-circuit (overcurrent) protection, to be provided as part of the building installation. Install only in accordance with national and local wiring regulations.

Statement 1045



**Warning** Installation of the equipment must comply with local and national electrical codes.

Statement 1074



**Note** If you are replacing DC power supplies in a node with power supply redundancy (two power supplies), you do not have to power off the node to replace a power supply because they are redundant as 1+1.



**Note** Do not mix power supply types or wattages in the node. Both power supplies must be identical.

**Step 1** Remove the DC power supply that you are replacing or a blank panel from an empty bay:


- a) Perform one of the following actions:
  - If you are replacing a power supply in a node that has only one DC power supply, shut down and remove power from the node as described in [Shutting Down and Removing Power From the Node](#), on page 69.
  - If you are replacing a power supply in a node that has two DC power supplies, you do not have to shut down the node.
- b) Remove the power cord from the power supply that you are replacing. Lift the connector securing clip slightly and then pull the connector from the socket on the power supply.
- c) Grasp the power supply handle while pinching the release lever toward the handle.
- d) Pull the power supply out of the bay.

**Step 2** Install a new DC power supply:



- a) Grasp the power supply handle and insert the new power supply into the empty bay.
- b) Push the power supply into the bay until the release lever locks.
- c) Connect the power cord to the new power supply. Press the connector into the socket until the securing clip clicks into place.
- d) Only if you shut down the node, press the Power button to boot the node to main power mode.

Figure 27: Replacing DC Power Supplies



1	Keyed cable connector (CAB-48DC-40A-8AWG)	3	PSU status LED
2	Keyed DC input socket	-	

## Installing DC Power Supplies (First Time Installation)



**Note** This procedure is for installing DC power supplies to the node for the first time. If you are replacing DC power supplies in a node that already has DC power supplies installed, see [Replacing DC Power Supplies, on page 126](#).



**Warning** A readily accessible two-poled disconnect device must be incorporated in the fixed wiring.  
Statement 1022



**Warning** This product requires short-circuit (overcurrent) protection, to be provided as part of the building installation. Install only in accordance with national and local wiring regulations.  
Statement 1045



**Warning** Installation of the equipment must comply with local and national electrical codes.

Statement 1074



**Note** Do not mix power supply types or wattages in the node. Both power supplies must be identical.



**Caution** As instructed in the first step of this wiring procedure, turn off the DC power source from your facility's circuit breaker to avoid electric shock hazard.

**Step 1** Turn off the DC power source from your facility's circuit breaker to avoid electric shock hazard.

**Note** The required DC input cable is Cisco part CAB-48DC-40A-8AWG. This 3-meter cable has a 3-pin connector on one end that is keyed to the DC input socket on the power supply. The other end of the cable has no connector so that you can wire it to your facility's DC power.


**Step 2** Wire the non-terminated end of the cable to your facility's DC power input source.

**Step 3** Connect the terminated end of the cable to the socket on the power supply. The connector is keyed so that the wires align for correct polarity and ground.

**Step 4** Return DC power from your facility's circuit breaker.

**Step 5** Press the Power button to boot the node to main power mode.

Figure 28: Replacing DC Power Supplies



1	Keyed cable connector (CAB-48DC-40A-8AWG)	3	PSU status LED
2	Keyed DC input socket	-	

**Step 6** See [Grounding for DC Power Supplies, on page 129](#) for information about additional chassis grounding.

---

## Grounding for DC Power Supplies

AC power supplies have internal grounding and so no additional grounding is required when the supported AC power cords are used.

When using a DC power supply, additional grounding of the node chassis to the earth ground of the rack is available. Two screw holes for use with your dual-hole grounding lug and grounding wire are supplied on the chassis rear panel.



---

**Note** The grounding points on the chassis are sized for 10-32 screws. You must provide your own screws, grounding lug, and grounding wire. The grounding lug must be dual-hole lug that fits 10-32 screws. The grounding cable that you provide must be 14 AWG (2 mm), minimum 60° C wire, or as permitted by the local code.

---

## Replacing a PCIe Riser

This node has two toolless PCIe risers for horizontal installation of PCIe cards. Each riser is available in multiple versions. See [PCIe Slot Specifications, on page 130](#) for detailed descriptions of the slots and features in each riser version.

---

- Step 1** Shut down and remove power from the node as described in [Shutting Down and Removing Power From the Node, on page 69](#).
- Step 2** Slide the node out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.
- Caution** If you cannot safely view and access the component, remove the node from the rack.
- Step 3** Remove the top cover from the node as described in [Removing the Node Top Cover, on page 70](#).
- Step 4** Remove the PCIe riser that you are replacing:
- Grasp the flip-up handle on the riser and the blue forward edge, and then lift up evenly to disengage its circuit board from the socket on the motherboard. Set the riser on an antistatic surface.
  - If the riser has a card installed, remove the card from the riser. See [Replacing a PCIe Card, on page 131](#).
- Step 5** Install a new PCIe riser:
- Note** The PCIe risers are not interchangeable. If you plug a PCIe riser into the wrong socket, the node will not boot. Riser 1 must plug into the motherboard socket labeled “RISER1.” Riser 2 must plug into the motherboard socket labeled “RISER2.”
- If you removed a card from the old PCIe riser, install the card to the new riser. See [Replacing a PCIe Card, on page 131](#).
  - Position the PCIe riser over its socket on the motherboard and over its alignment slots in the chassis.

- c) Carefully push down on both ends of the PCIe riser to fully engage its circuit board connector with the socket on the motherboard.

**Step 6** Replace the top cover to the node.

**Step 7** Replace the node in the rack, replace cables, and then fully power on the node by pressing the Power button.

## Replacing a PCIe Card



**Note** Cisco supports all PCIe cards qualified and sold by Cisco. PCIe cards not qualified or sold by Cisco are the responsibility of the customer. Although Cisco will always stand behind and support the C-Series rack-mount server nodes, customers using standard, off-the-shelf, third-party cards must go to the third-party card vendor for support if any issue with that particular card occurs.

## PCIe Slot Specifications

The node contains three toolless PCIe risers for horizontal installation of PCIe cards. Each riser is orderable in multiple versions.

- Riser 1 contains PCIe slots 1, 2, and 3 and is available in the following different options:
  - SFF node, I/O-Centric—Slots 1 (x8), 2 (x16), and 3 (x8). All slots are controlled by CPU 1 or CPU 2 depending on the node model.
  - SFF node, Storage Centric—
    - Slots 1 (Reserved), 2 (x4), and 3 (x4) for drive bays in 24-drive SAS/SATA model, 24-drive NVMe model, and 12 drive NVMe model. All slots are controlled by CPU 2.
    - Slots 1, 2, and 3 are not supported in 12 SAS/SATA model.
  - LFF node—Slots 1 (Reserved), 2 (x4), and 3 (x4). All slots are controlled by CPU 1.
- Riser 2 contains PCIe slots 4, 5 and 6 and is available in the following different options:
  - SFF node, I/O-Centric—Slots 4 (x8), 5 (x16), and 6 (x8). All slots are controlled by CPU 2.
  - SFF node, Storage Centric—Slots 4, 5, and 6 do not support storage devices in the SFF model of the node.
  - LFF node—Slots 4 (x8), 5 (x16), and 6 (x8). All slots are controlled by CPU 2.
- Riser 3 contains PCIe slots 7 and 8 and is available in the following different options:
  - SFF node, I/O-Centric—Slots 7 (x8) and 8 (x8) for SATA/SAS models. Slots 7 and 8 are controlled by CPU 2 for SATA/SAS nodes.
 


Slots 7 and 8 are not supported for NVMe-only models.
  - SFF node, Storage Centric—Slots 7 (x4) and 8 (x4) for drive bays in 24-drive and 12-drive SAS/SATA versions of the server. All slots are controlled by CPU 2 .

Slots 7 and 8 are not supported for NVMe-only models.

- LFF server—Slots 7 (x4) and 8 (x4) for drive bays. All slots are controlled by CPU 2.

The following illustration shows the PCIe slot numbering.

**Figure 29: Rear Panel, Showing PCIe Slot Numbering**



309418

## Replacing a PCIe Card



**Note** If you are installing a Cisco UCS Virtual Interface Card, there are prerequisite considerations. See [Cisco Virtual Interface Card \(VIC\) Considerations](#), on page 133.



**Note** RAID controller cards install into a dedicated motherboard socket. See [Replacing a SAS Storage Controller Card \(RAID or HBA\)](#), on page 135.



**Note** For instructions on installing or replacing double-wide GPU cards, see [GPU Card Installation](#), on page 189.

- Step 1** Shut down and remove power from the node as described in [Shutting Down and Removing Power From the Node](#), on page 69.
- Step 2** Slide the node out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.
- Caution** If you cannot safely view and access the component, remove the node from the rack.
- Step 3** Remove the top cover from the node as described in [Removing the Node Top Cover](#), on page 70.
- Step 4** Remove the PCIe card that you are replacing:
- Remove any cables from the ports of the PCIe card that you are replacing.
  - Use two hands to flip up and grasp the blue riser handle and the blue finger grip area on the front edge of the riser, and then lift straight up.
  - On the bottom of the riser, push the release latch that holds the securing plate, and then swing the hinged securing plate open.
  - Open the hinged card-tab retainer that secures the rear-panel tab of the card.

- e) Pull evenly on both ends of the PCIe card to remove it from the socket on the PCIe riser.  
If the riser has no card, remove the blanking panel from the rear opening of the riser.


**Step 5** Install a new PCIe card:

- a) With the hinged card-tab retainer open, align the new PCIe card with the empty socket on the PCIe riser.
- b) Push down evenly on both ends of the card until it is fully seated in the socket.
- c) Ensure that the card's rear panel tab sits flat against the riser rear-panel opening and then close the hinged card-tab retainer over the card's rear-panel tab.
- d) Swing the hinged securing plate closed on the bottom of the riser. Ensure that the clip on the plate clicks into the locked position.
- e) Position the PCIe riser over its socket on the motherboard and over the chassis alignment channels.
- f) Carefully push down on both ends of the PCIe riser to fully engage its connector with the sockets on the motherboard.

**Step 6** Replace the top cover to the node.

**Step 7** Replace the node in the rack, replace cables, and then fully power on the node by pressing the Power button.

**Figure 30: PCIe Riser Card Securing Mechanisms**



<b>1</b>	Release latch on hinged securing plate	<b>3</b>	Hinged card-tab retainer
----------	--	----------	--------------------------

306229

2	Hinged securing plate	-	
---	-----------------------	---	--

## Cisco Virtual Interface Card (VIC) Considerations

This section describes VIC card support and special considerations for this node.



**Note** If you use the *Cisco Card* NIC mode, you must also make a *VIC Slot* setting that matches where your VIC is installed. The options are Riser1, Riser2, and Flex-LOM. See [NIC Mode and NIC Redundancy Settings, on page 56](#) for more information about NIC modes.

**Table 17: VIC Support and Considerations in This Server**

VIC	How Many Supported in Server	Slots That Support VICs	Primary Slot For <i>Cisco Card</i> NIC Mode	Minimum Cisco IMC Firmware
Cisco HX VIC 1455 HX-PCIE-C25Q-04	2 PCIe	PCIe 2 PCIe 5	PCIe 2	4.0(1)
Cisco HX VIC 1495 HX-PCIE-C100-04	2 PCIe	PCIe 2 PCIe 5	PCIe 2	4.0(2)
Cisco HX VIC 1457 HX-MLOM-C25Q-04	1 mLOM	mLOM	mLOM	4.0(1)
Cisco HX VIC 1497 HX-MLOM-C100-04	1 mLOM	mLOM	mLOM	4.0(2)

- A total of 3 VICs are supported in the node: 2 PCIe style, and 1 mLOM style.



**Note** Single wire management is supported on only one VIC at a time. If multiple VICs are installed on a node, only one slot has NCSI enabled at a time. For single wire management, priority goes to the MLOM slot, then slot 2, then slot 5 for NCSI management traffic. When multiple cards are installed, connect the single-wire management cables in the priority order mentioned above.

- The primary slot for a VIC card in PCIe riser 1 is slot 2. The secondary slot for a VIC card in PCIe riser 1 is slot 1.




---

**Note** The NCSI protocol is supported in only one slot at a time in each riser. If a GPU card is present in slot 2, NCSI automatically shifts from slot 2 to slot 1.

---

- The primary slot for a VIC card in PCIe riser 2 is slot 5. The secondary slot for a VIC card in PCIe riser 2 is slot 4.




---

**Note** The NCSI protocol is supported in only one slot at a time in each riser. If a GPU card is present in slot 5, NCSI automatically shifts from slot 5 to slot 4.

---




---

**Note** PCIe riser 2 is not available in a single-CPU system.

---

## Replacing an mLOM Card

The node supports a modular LOM (mLOM) card to provide additional rear-panel connectivity. The mLOM socket is on the motherboard, under the storage controller card.

The mLOM socket provides a Gen-3 x16 PCIe lane. The socket remains powered when the node is in 12 V standby power mode, and it supports the network communications services interface (NCSI) protocol.




---

**Note** If your mLOM card is a Cisco UCS Virtual Interface Card (VIC), see [Cisco Virtual Interface Card \(VIC\) Considerations, on page 133](#) for more information and support details.

---

**Step 1** Remove any existing mLOM card (or a blanking panel):

- Shut down and remove power from the node as described in [Shutting Down and Removing Power From the Node, on page 69](#).
- Slide the node out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.

**Caution** If you cannot safely view and access the component, remove the node from the rack.

- Remove the top cover from the node as described in [Removing the Node Top Cover, on page 70](#).
- Remove any storage controller (RAID or HBA card) to provide clearance to the mLOM socket on the motherboard. See [Replacing a SAS Storage Controller Card \(RAID or HBA\), on page 135](#).
- Loosen the single captive thumbscrew that secures the mLOM card to the threaded standoff on the chassis floor.
- Slide the mLOM card horizontally to free it from the socket, then lift it out of the node.

**Step 2** Install a new mLOM card:

- Set the mLOM card on the chassis floor so that its connector is aligned with the motherboard socket.



- b) Push the card horizontally to fully engage the card's edge connector with the socket.
- c) Tighten the captive thumbscrew to secure the card to the chassis floor.
- d) Return the storage controller card to the node. See [Replacing a SAS Storage Controller Card \(RAID or HBA\)](#), on page 135.
- e) Replace the top cover to the node.
- f) Replace the node in the rack, replace cables, and then fully power on the node by pressing the Power button.

---

## Replacing a SAS Storage Controller Card (RAID or HBA)

For hardware-based storage control, the node can use a Cisco modular SAS RAID controller or SAS HBA that plugs into a dedicated, vertical socket on the motherboard.

### Storage Controller Card Firmware Compatibility

Firmware on the storage controller (RAID or HBA) must be verified for compatibility with the current Cisco IMC and BIOS versions that are installed on the node. If not compatible, upgrade or downgrade the storage controller firmware using the Host Upgrade Utility (HUU) for your firmware release to bring it to a compatible level.



---

**Note** **For nodes running in standalone mode only:** After you replace controller hardware (HX-RAID-M6T, HX-RAID-M6HD, HX-RAID-M6SD, HX-SAS-M6T, or HX-SAS-M6HD), you must run the Cisco UCS Host Upgrade Utility (HUU) to update the controller firmware, even if the firmware Current Version is the same as the Update Version. Running HUU is necessary to program any controller specific values to the storage controller for the specific node. If you do not run HUU, the storage controller may not be discovered..

---


See the HUU guide for your Cisco IMC release for instructions on downloading and using the utility to bring node components to compatible levels: [HUU Guides](#).

### Removing the Dual Storage Controller Cards

The front RAID assembly can contain either a single Storage controller card in a single tray, or two Storage controller cards each in its own tray. Use this procedure to remove each of the Storage controller cards. This procedure assumes that you have removed power from the node and removed the top cover.


- 
- Step 1** Locate the dual Storage controller cards.
- Each Storage controller card has its own tray, as shown.

## Removing the Dual Storage Controller Cards



**Step 2** Remove the fan tray.


For more information, see [Removing the Fan Tray](#), on page 100.



**Step 3** Disconnect the cables.

- a) For each Storage controller card, grasp the ribbon cable connector and disconnect it from the RAID card. You can leave the other end of the ribbon cable attached to the motherboard.

- b) For each Storage controller card, grasp the connector for the rear-drive cable, and disconnect it from the card. You can leave the other end of the rear-drive cable attached.



1	SAS cable connection on Storage controller card.	2	SAS cable that connects to rear drives in Riser 3B
3	SAS cable connection on rear Riser 3B.	4	SAS cable connection on Storage controller card.
5	Ribbon cables connecting Storage controller cards to motherboard.	6	SAS cable that connects to rear drives in Riser 1B
7	SAS cable connection on rear Riser 1B.		


**Step 4** Remove the Storage controller cards.

- a) Grasp the cable that leads to the rear drives and disconnect it from each card.
- b) Grasp the handle at the top of each card tray, and gently push it towards the rear of the node.


The handle should slide to the open position. This step disconnects the Storage controller card from a socket on an interior wall.

- c) Using a #2 Phillips screwdriver, loosen the captive screws at the edges of the trays.

## Removing the Dual Storage Controller Cards



- d) Grasp each card tray by the handle and lift the Storage controller cards out of the chassis.



---

**What to do next**

Reinsert the dual Storage controller cards. Go to [Installing the Dual Storage Controller Cards, on page 139](#).

## Installing the Dual Storage Controller Cards

Use this procedure to install the dual Storage controller cards into the node. The Storage controller cards are contained in tray, which is replaceable.


---

**Step 1** Grasp each card tray by the handle.

**Step 2** Install the Storage controller cards.


- a) Make sure that the handle of the tray is in the open position.
- b) Make sure that the cables do not obstruct installing the Storage controller cards.
- c) Orient the Storage controller card so that the thumbscrews align with their threaded standoffs on the motherboard.

- d) Holding the card tray by the handle, keep the tray level and lower it into the



node.

- e) Using a #2 Phillips head screwdriver, tighten the screws at the edges of each tray.




- f) Gently push the handle of the tray towards the front of the node.

This step seats each Storage controller card into its socket on the interior wall. You might feel some resistance as the card meets the socket. This resistance is normal.

**Step 3** Reconnect the cables.

**Step 4** Reinsert the fan tray.

For more information, see [Installing the Fan Tray, on page 101](#).



---

**What to do next**

Perform other maintenance tasks, if needed, or replace the top cover and restore facility power.


## Removing the Storage Controller Card

The Storage controller can contain either a single controller card in a single tray, or two controller cards each in its own tray. Use this procedure to remove the single Storage controller card. This procedure assumes that you have removed power from the node and removed the top cover.


---

**Step 1** Locate the Storage controller card.






**Step 2** Remove the fan tray.  
 For more information, see [Removing the Fan Tray](#), on page 100.



**Step 3** Disconnect the cables.


- a) Grasp the ribbon cable connectors and disconnect them from the Storage controller card.  
You can leave the other end of the ribbon cable attached to the motherboard.
- b) Grasp the connector for the rear-drive cables (1 and 4) and disconnect them from the Storage controller card.  
You can leave the other end of the rear-drive cable attached.




1	Storage controller card connector for rear drives (Riser 3B)	2	SAS/SATA cable for rear drives
3	Connector for PCI Riser 3	4	Storage controller card connector for rear drives (Riser 1B)
5	SAS/SATA cable for rear drives	6	Connector for PCI Riser 1

**Step 4** Remove the Storage controller card.

- a) Using both hands, grasp the handle at the top of the card tray, and gently push it towards the rear of the node.  
The handle should slide to the open position. This step disconnects the Storage controller card from a socket on an interior wall.
- b) Using a #2 Phillips screwdriver, loosen the captive screws at the edges of the tray.



- c) Using both hands, grasp the tray's handle, and keeping the Storage controller card tray level, lift it out of the chassis.



### What to do next

Reinsert the Storage controller card. Go to [Installing the Storage Controller Card](#), on page 146.


## Installing the Storage Controller Card

Use this procedure to install the single Storage controller card into the server. The Storage controller card is contained in tray, which is replaceable.


**Step 1** Grasp the card tray by the handle.

**Step 2** Install the Storage controller card.

- a) Make sure that the handle of the tray is in the open position.
- b) Make sure that the cables do not obstruct installing the Storage controller card.
- c) Orient the Storage controller card so that the thumbscrews align with their threaded standoffs on the motherboard.
- d) Using both hands, hold the card tray by the handle, keep the tray level, and lower it into the server.



- e) Using a #2 Phillips head screwdriver, tighten the screws at the edges of the tray.




- f) Using both hands, make sure to apply equal pressure to both sides of the handle, and gently push the handle of the tray towards the front of the server.

This step seats the Storage controller card into its sockets on the interior wall. You might feel some resistance as the card meets the socket. This resistance is normal.

**Step 3** Reconnect the cables.

**Step 4** Reinsert the fan tray.

For more information, see [Installing the Fan Tray, on page 101](#).




### What to do next

Perform other maintenance tasks, if needed, or replace the top cover and restore facility power.

## Verify Cabling

After installing a Storage controller card, the cabling between the card(s) and rear drives should be as follows.

- For a 24-drive node, verify the following:
  - the SAS/SATA cable is connected to the controller card and Riser 3B
  - the SAS/SATA cable is connected to the controller card and the Riser 1B
  - both ribbon cables are connected to the controller card and the motherboard

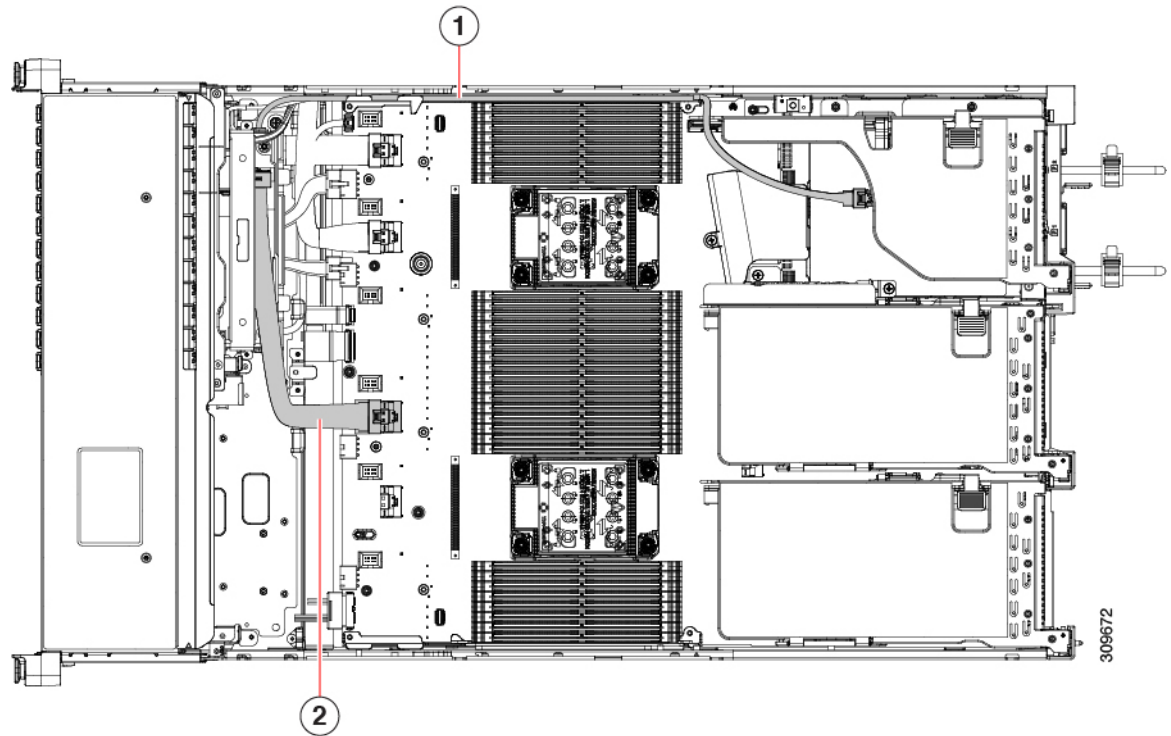


309673

1	SAS cable connection on Storage controller card.	2	SAS cable that connects to rear drives in Riser 3B
3	SAS cable connection on rear Riser 3B.	4	SAS cable connection on Storage controller card.
5	Ribbon cables connecting Storage controller cards to motherboard.	6	SAS cable that connects to rear drives in Riser 1B
7	SAS cable connection on rear Riser 1B.		

- For 12 drive node, verify the following:
  - the SAS/SATA cable is connected to the Storage controller card and the Riser 3B
  - the ribbon cable is connected to the Storage controller card and the motherboard.





1	SAS cable connects from rear drive to SAS card
2	Ribbon cable connects from motherboard to SAS card.

## Replacing a SATA Interposer Card (12-Drive SFF Server Only)



**Note** The only version of this node that supports the SATA Interposer card is the SFF, 12-drive version (HX-C240-M6S)

For software-based storage control for the front-loading drives, the node requires a SATA interposer card that plugs into a dedicated socket on the motherboard (the same socket used for SAS storage controllers). The interposer card sits between the front-loading drive bays and the fan tray and supports up to eight SATA drives in bays 1 through 8).



**Note** You cannot use a hardware RAID controller card and the embedded software RAID controller to control front drives at the same time.

**Step 1** Prepare the node for component installation:

- a) Shut down and remove power from the node as described in [Shutting Down and Removing Power From the Node, on page 69](#).
- b) Slide the node out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.

**Caution** If you cannot safely view and access the component, remove the node from the rack.

- c) Remove the top cover from the node as described in [Removing the Node Top Cover, on page 70](#).

## Step 2

Remove any existing SATA interposer card from the node:


**Note** A SATA interposer card for this node is preinstalled inside a carrier-frame that helps to secure the card to the inner chassis wall. You do not have to remove this carrier frame from the existing card.

- a) Disconnect the PCIe SAS cable from the existing card.


This is the only cable on the card, and it connects the card to the motherboard by a x8 slimSAS connector.

- b) Using a #2 Phillips screwdriver, completely loosen the captive screws.
- c) Gently push the handle toward the middle of the node.

This step moves the handle to the open position and unseats the interposer card from its socket on the interior wall.




- d) Making sure that the card is disconnected from its socket, grasp the handle and lift straight up to remove the interposer from the node.



3096696


**Step 3** Install a new SATA interposer card:

- a) Orient the interposer card so that the front edge connector is facing its socket on the interior wall.
- b) Align the captive screws on the card with their threaded standoffs.
- c) Making sure that the handle is completely in the open position and the PCIe SAS cable will not obstruct installation, lower the interposer card into the node.



3096694

- d) Keeping the card level, push the handle forward to seat the interposer card into its socket on the interior wall.
- e) Using a #2 Phillips screwdriver, tighten the captive screws.



f) Connect PCIe cables to the new card.

If this is a first-time installation, see [Storage Controller Cable Connectors and Backplanes](#), on page 184 for cabling instructions.

**Step 4** Replace the top cover to the node.

**Step 5** Replace the node in the rack, replace cables, and then fully power on the node by pressing the Power button.

## Replacing the Supercap (RAID Backup)

This node supports installation of one Supercap unit for 12 drive nodes, and two Supercap units for 24 drive nodes. The unit mounts to a bracket on the removable air baffle.


The Supercap provides approximately three years of backup for the disk write-back cache DRAM in the case of a sudden power loss by offloading the cache to the NAND flash.

**Step 1** Prepare the node for component installation:

- a) Shut down and remove power from the node as described in [Shutting Down and Removing Power From the Node, on page 69](#).
- b) Slide the node out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.


**Caution** If you cannot safely view and access the component, remove the node from the rack.

- c) Remove the top cover from the node as described in [Removing the Node Top Cover, on page 70](#).
- d) Locate the SuperCap unit(s) as shown below.



**Step 2** Remove an existing Supercap:

- a) Disconnect the Supercap cable from the RAID cable.
- b) Push aside the securing tab that holds the Supercap to its bracket.
- c) Lift the Supercap free of the bracket and set it aside.


**Step 3**

Install a new Supercap:

- Orient the SuperCap unit so that the SuperCap cable and the RAID cable connectors meet.
- Make sure that the RAID cable does not obstruct the SuperCap when you install it, then insert the new Supercap into the mounting bracket.

Make that the SuperCap unit is securely inserted into it bracket.

- Connect the Supercap cable from the RAID controller card to the connector on the Supercap cable.



**Step 4** Replace the top cover to the node.

**Step 5** Replace the node in the rack, replace cables, and then fully power on the node by pressing the Power button.

## Replacing a Boot-Optimized M.2 RAID Controller Module

The Cisco Boot-Optimized M.2 RAID Controller module connects to the mini-storage module socket on the motherboard. It includes slots for two SATA M.2 drives, plus an integrated 6-Gbps SATA RAID controller that can control the SATA M.2 drives in a RAID 1 array.

### Cisco Boot-Optimized M.2 RAID Controller Considerations

Review the following considerations:



**Note** The Cisco Boot-Optimized M.2 RAID Controller is not supported when the node is used as a compute-only node in Cisco HyperFlex configurations.

- The minimum version of Cisco IMC and Cisco UCS Manager that support this controller is 4.0(4) and later.
- This controller supports RAID 1 (single volume) and JBOD mode.





**Note** Do not use the node's embedded SW MegaRAID controller to configure RAID settings when using this controller module. Instead, you can use the following interfaces:

- Cisco IMC 4.2(1) and later
- BIOS HII utility, BIOS 4.2(1) and later
- Cisco UCS Manager 4.2(1) and later (UCS Manager-integrated nodes)

- 
- A SATA M.2 drive in slot 1 (the top) is the first SATA device; a SATA M.2 drive in slot 2 (the underside) is the second SATA device.
    - The name of the controller in the software is MSTOR-RAID.
    - A drive in Slot 1 is mapped as drive 253; a drive in slot 2 is mapped as drive 254.
  - When using RAID, we recommend that both SATA M.2 drives are the same capacity. If different capacities are used, the smaller capacity of the two drives is used to create a volume and the rest of the drive space is unusable.

JBOD mode supports mixed capacity SATA M.2 drives.
  - Hot-plug replacement is *not* supported. The node must be powered off.
  - Monitoring of the controller and installed SATA M.2 drives can be done using Cisco IMC and Cisco UCS Manager. They can also be monitored using other utilities such as UEFI HII, PMCLI, XMLAPI, and Redfish.
  - Updating firmware of the controller and the individual drives:
    - For standalone nodes, use the Cisco Host Upgrade Utility (HUU). Refer to the [HUU Documentation](#).
    - For nodes integrated with Cisco UCS Manager, refer to the [Cisco UCS Manager Firmware Management Guide](#).
  - The SATA M.2 drives can boot in UEFI mode only. Legacy boot mode is not supported.
  - If you replace a single SATA M.2 drive that was part of a RAID volume, rebuild of the volume is auto-initiated after the user accepts the prompt to import the configuration. If you replace both drives of a volume, you must create a RAID volume and manually reinstall any OS.
  - We recommend that you erase drive contents before creating volumes on used drives from another node. The configuration utility in the node BIOS includes a SATA secure-erase function.
  - The node BIOS includes a configuration utility specific to this controller that you can use to create and delete RAID volumes, view controller properties, and erase the physical drive contents. Access the utility by pressing **F2** when prompted during node boot. Then navigate to **Advanced > Cisco Boot Optimized M.2 RAID Controller**.

## Replacing a Cisco Boot-Optimized M.2 RAID Controller

This topic describes how to remove and replace a Cisco Boot-Optimized M.2 RAID Controller. The controller board has one M.2 socket on its top (Slot 1) and one M.2 socket on its underside (Slot 2).

- Step 1** Shut down and remove power from the node as described in [Shutting Down and Removing Power From the Node](#), on page 69.
- Step 2** Slide the node out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.


**Caution** If you cannot safely view and access the component, remove the node from the rack.

- Step 3** Remove the top cover from the node as described in [Removing the Node Top Cover](#), on page 70.


- Step 4** Remove a controller from its motherboard socket:

- a) Locate the controller in its socket between PCIe Riser 2 and 3.

**Figure 31: Cisco Boot-Optimized M.2 RAID Controller on Motherboard**



- b) Using a #2 Phillips screwdriver, loosen the captive screws and remove the M.2 module.
- c) At each end of the controller board, push outward on the clip that secures the carrier.
- d) Lift both ends of the controller to disengage it from the carrier.



- e) Set the carrier on an anti-static surface.


### Step 5

If you are transferring SATA M.2 drives from the old controller to the replacement controller, do that before installing the replacement controller:

**Note** Any previously configured volume and data on the drives are preserved when the M.2 drives are transferred to the new controller. The system will boot the existing OS that is installed on the drives.

- a) Use a #1 Phillips-head screwdriver to remove the single screw that secures the M.2 drive to the carrier.
- b) Lift the M.2 drive from its socket on the carrier.
- c) Position the replacement M.2 drive over the socket on the controller board.
- d) Angle the M.2 drive downward and insert the connector-end into the socket on the carrier. The M.2 drive's label must face up.
- e) Press the M.2 drive flat against the carrier.
- f) Install the single screw that secures the end of the M.2 SSD to the carrier.
- g) Turn the controller over and install the second M.2 drive.

Figure 32: Cisco Boot-Optimized M.2 RAID Controller, Showing M.2 Drive Installation



**Step 6** Install the controller to its socket on the motherboard:

- Position the controller over the socket, with the controller's connector facing down and at the same end as the motherboard socket. Two alignment pegs must match with two holes on the controller.
- Gently push down the socket end of the controller so that the two pegs go through the two holes on the controller.
- Push down on the controller so that the securing clips click over it at both ends.

**Step 7** Replace the top cover to the node.

**Step 8** Replace the node in the rack, replace cables, and then fully power on the node by pressing the Power button.

## Replacing a Chassis Intrusion Switch

The chassis intrusion switch is an optional security feature that logs an event in the system event log (SEL) whenever the cover is removed from the chassis.

**Step 1** Prepare the node for component installation:

- Shut down and remove power from the node as described in [Shutting Down and Removing Power From the Node, on page 69](#).
- Slide the node out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.

**Caution** If you cannot safely view and access the component, remove the node from the rack.

- Remove the top cover from the node as described in [Removing the Node Top Cover, on page 70](#).

**Step 2** Remove an existing intrusion switch:

- Disconnect the intrusion switch cable from the socket on the motherboard.

- b) Use a #1 Phillips-head screwdriver to loosen and remove the single screw that holds the switch mechanism to the chassis wall.
- c) Slide the switch mechanism straight up to disengage it from the clips on the chassis.

**Step 3** Install a new intrusion switch:

- a) Slide the switch mechanism down into the clips on the chassis wall so that the screwholes line up.
- b) Use a #1 Phillips-head screwdriver to install the single screw that secures the switch mechanism to the chassis wall.
- c) Connect the switch cable to the socket on the motherboard.

**Step 4** Replace the cover to the node.

**Step 5** Replace the node in the rack, replace cables, and then fully power on the node by pressing the Power button.

---

## Installing a Trusted Platform Module (TPM)

The trusted platform module (TPM) is a small circuit board that plugs into a motherboard socket and is then permanently secured with a one-way screw. The socket location is on the motherboard below PCIe riser 2.

### TPM Considerations

- This node supports either TPM version 1.2 or TPM version 2.0.
- Field replacement of a TPM is not supported; you can install a TPM after-factory only if the node does not already have a TPM installed.
- If there is an existing TPM 1.2 installed in the node, you cannot upgrade to TPM 2.0. If there is no existing TPM in the node, you can install TPM 2.0.
- If the TPM 2.0 becomes unresponsive, reboot the node.

### Installing and Enabling a TPM



---

**Note** Field replacement of a TPM is not supported; you can install a TPM after-factory only if the node does not already have a TPM installed.

---

This topic contains the following procedures, which must be followed in this order when installing and enabling a TPM:

1. Installing the TPM Hardware
2. Enabling the TPM in the BIOS
3. Enabling the Intel TXT Feature in the BIOS

## Installing TPM Hardware



---

**Note** For security purposes, the TPM is installed with a one-way screw. It cannot be removed with a standard screwdriver.

---

- Step 1** Prepare the node for component installation:
- Shut down and remove power from the node as described in [Shutting Down and Removing Power From the Node, on page 69](#).
  - Slide the node out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.
- Caution** If you cannot safely view and access the component, remove the node from the rack.
- Remove the top cover from the node as described in [Removing the Node Top Cover, on page 70](#).
- Step 2** Remove PCIe riser 2 from the node to provide clearance to the TPM socket on the motherboard.
- Step 3** Install a TPM:
- Locate the TPM socket on the motherboard.
  - Align the connector that is on the bottom of the TPM circuit board with the motherboard TPM socket. Align the screw hole on the TPM board with the screw hole that is adjacent to the TPM socket.
  - Push down evenly on the TPM to seat it in the motherboard socket.
  - Install the single one-way screw that secures the TPM to the motherboard.
- Step 4** Replace PCIe riser 2 to the node. See [Replacing a PCIe Riser, on page 129](#).
- Step 5** Replace the cover to the node.
- Step 6** Replace the node in the rack, replace cables, and then fully power on the node by pressing the Power button.
- Step 7** Continue with [Enabling the TPM in the BIOS, on page 164](#).
- 

## Enabling the TPM in the BIOS

After hardware installation, you must enable TPM support in the BIOS.



---

**Note** You must set a BIOS Administrator password before performing this procedure. To set this password, press the **F2** key when prompted during system boot to enter the BIOS Setup utility. Then navigate to **Security > Set Administrator Password** and enter the new password twice as prompted.

---

- Step 1** Enable TPM Support:
- Watch during bootup for the F2 prompt, and then press **F2** to enter BIOS setup.
  - Log in to the BIOS Setup Utility with your BIOS Administrator password.
  - On the BIOS Setup Utility window, choose the **Advanced** tab.
  - Choose **Trusted Computing** to open the TPM Security Device Configuration window.
  - Change TPM SUPPORT to **Enabled**.

- f) Press **F10** to save your settings and reboot the node.

**Step 2** Verify that TPM support is now enabled:

- a) Watch during bootup for the F2 prompt, and then press **F2** to enter BIOS setup.
- b) Log into the BIOS Setup utility with your BIOS Administrator password.
- c) Choose the **Advanced** tab.
- d) Choose **Trusted Computing** to open the TPM Security Device Configuration window.
- e) Verify that TPM SUPPORT and TPM State are Enabled.

**Step 3** Continue with [Enabling the Intel TXT Feature in the BIOS, on page 165](#).

---

## Enabling the Intel TXT Feature in the BIOS

Intel Trusted Execution Technology (TXT) provides greater protection for information that is used and stored on the business node. A key aspect of that protection is the provision of an isolated execution environment and associated sections of memory where operations can be conducted on sensitive data, invisibly to the rest of the system. Intel TXT provides for a sealed portion of storage where sensitive data such as encryption keys can be kept, helping to shield them from being compromised during an attack by malicious code.

---

**Step 1** Reboot the node and watch for the prompt to press F2.

**Step 2** When prompted, press **F2** to enter the BIOS Setup utility.

**Step 3** Verify that the prerequisite BIOS values are enabled:

- a) Choose the **Advanced** tab.
- b) Choose **Intel TXT(LT-SX) Configuration** to open the Intel TXT(LT-SX) Hardware Support window.
- c) Verify that the following items are listed as Enabled:
  - VT-d Support (default is Enabled)
  - VT Support (default is Enabled)
  - TPM Support
  - TPM State
- d) Do one of the following:
  - If VT-d Support and VT Support are already enabled, skip to step 4.
  - If VT-d Support and VT Support are not enabled, continue with the next steps to enable them.
- e) Press **Escape** to return to the BIOS Setup utility **Advanced** tab.
- f) On the Advanced tab, choose **Processor Configuration** to open the Processor Configuration window.
- g) Set Intel (R) VT and Intel (R) VT-d to **Enabled**.

**Step 4** Enable the Intel Trusted Execution Technology (TXT) feature:

- a) Return to the Intel TXT(LT-SX) Hardware Support window if you are not already there.
- b) Set TXT Support to **Enabled**.

**Step 5** Press **F10** to save your changes and exit the BIOS Setup utility.

---

# Removing the PCB Assembly (PCBA)

The PCBA is secured to the node's sheet metal. You must disconnect the PCBA from the tray before recycling the PCBA. The PCBA is secured by different types of fasteners.

## Before you begin



---

**Note** **For Recyclers Only!** This procedure is not a standard field-service option. This procedure is for recyclers who will be reclaiming the electronics for proper disposal to comply with local eco design and e-waste regulations.

---

To remove the printed circuit board assembly (PCBA), the following requirements must be met:

- The node must be disconnected from facility power.
- The node must be removed from the equipment rack.
- The node's top cover must be removed. See [Removing the Node Top Cover, on page 70](#).

Gather the following tools before beginning this procedure:

- Pliers
- T10 Torx screwdriver
- #2 Phillips screwdriver


---

**Step 1** Locate the PCBA's mounting screws.

The following figure shows the location of the mounting screws.



Figure 33: Screw Locations for Removing the HX C240 M6



Indicator	Fastener	Required Tool
Red circles (○)	M3.5x0.6 mm screws (18)	Torx T10 screwdriver
Green circles (○)	H15 M4x0.7 mm lock screws (2)	Pliers
Blue circles (○)	M3.5x0.6 mm Thumb screws (2)	Torx T10 screwdriver
Yellow circles (○)	H12 M4x0.7 mm locking screws (2)	Pliers
Purple circle (○)	M3.5 Thumb screw (1), on the M.2 riser cage	#2 Phillips screwdriver
Lavender circle (○)	M3.5 Thumb screw (1), on the air duct	#2 Phillips screwdriver

**Step 2** Using the appropriate tools, remove the screws.


**Step 3** Remove the PCBA from the sheet metal and dispose of each in compliance with your local e-waste and recycling regulations.

## Service Headers and Jumpers

This node includes blocks of headers and switches (SW12, CN3) that you can use for certain service and debug functions.

This section contains the following topics:

- Using the Clear CMOS Switch (SW12, Switch 9), on page 169
- Using the BIOS Recovery Switch (SW12, Switch 5), on page 170
- Using the Clear BIOS Password Switch (SW12, Switch 6), on page 172
- Using the Boot Alternate Cisco IMC Image Header (CN3, Pins 1-2), on page 172
- Using the System Firmware Secure Erase Header (CN3, Pins 3-4), on page 173



1	Location of header block CN3	5	Clear BIOS Password: SW12 switch 6 Default setting: Off Clear password: On.
2	Boot Cisco IMC from alternate image: CN pins 1-2 Default: Open. Place the jumper shunt over the pins to close the circuit.	6	Recover BIOS: SW12 switch 5 Default setting: Off Recovery mode: On.

3	System Firmware Secure Erase: CN3 pins 3-4 Default: Open. Place the jumper shunt over the pins to close the circuit.	7	Clear CMOS: SW12 switch 9 Default setting: Off Recovery mode: On. Gently push the switch to the right, which is the On position.
4	Location of switch block SW12	-	

## Using the Clear CMOS Switch (SW12, Switch 9)

You can use this switch to clear the node's CMOS settings in the case of a system hang. For example, if the node hangs because of incorrect settings and does not boot, use this switch to invalidate the settings and reboot with defaults.

You will find it helpful to refer to the location of the CN3 header. See [Service Headers and Jumpers](#), on page 167.



**Caution** Clearing the CMOS removes any customized settings and might result in data loss. Make a note of any necessary customized settings in the BIOS before you use this clear CMOS procedure.

- Step 1** Shut down and remove power from the node as described in [Shutting Down and Removing Power From the Node](#), on page 69.
- Step 2** Slide the node out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.
- Caution** If you cannot safely view and access the component, remove the node from the rack.
- Step 3** Remove the top cover from the node as described in [Removing the Node Top Cover](#), on page 70.
- Step 4** Using your finger, gently push the SW12 switch 9 to the side marked ON.
- Step 5** Reinstall the top cover and reconnect AC power cords to the node. The node powers up to standby power mode, indicated when the Power LED on the front panel is amber.
- Step 6** Return the node to main power mode by pressing the Power button on the front panel. The node is in main power mode when the Power LED is green.
- Note** You must allow the entire node to reboot to main power mode to complete the reset. The state of the switch cannot be determined without the host CPU running.
- Step 7** Press the Power button to shut down the node to standby power mode, and then remove AC power cords from the node to remove all power.
- Step 8** Remove the top cover from the node.
- Step 9** Using your finger, gently push switch 9 to its original position (OFF).
- Note** If you do not reset the switch to its original position (OFF), the CMOS settings are reset to the defaults every time you power-cycle the node.

- Step 10** Replace the top cover, replace the node in the rack, replace power cords and any other cables, and then power on the node by pressing the Power button.

## Using the BIOS Recovery Switch (SW12, Switch 5)

Depending on which stage the BIOS becomes corrupted, you might see different behavior.

- If the BIOS BootBlock is corrupted, you might see the system get stuck on the following message:

```
Initializing and configuring memory/hardware
```

- If it is a non-BootBlock corruption, a message similar to the following is displayed:

```
****BIOS FLASH IMAGE CORRUPTED****
Flash a valid BIOS capsule file using Cisco IMC WebGUI or CLI interface.
IF Cisco IMC INTERFACE IS NOT AVAILABLE, FOLLOW THE STEPS MENTIONED BELOW.
1. Connect the USB stick with bios.cap file in root folder.
2. Reset the host.
IF THESE STEPS DO NOT RECOVER THE BIOS
1. Power off the system.
2. Mount recovery jumper.
3. Connect the USB stick with bios.cap file in root folder.
4. Power on the system.
Wait for a few seconds if already plugged in the USB stick.
REFER TO SYSTEM MANUAL FOR ANY ISSUES.
```



**Note** As indicated by the message shown above, there are two procedures for recovering the BIOS. Try procedure 1 first. If that procedure does not recover the BIOS, use procedure 2.

### Procedure 1: Reboot With bios.cap Recovery File

- Step 1** Download the BIOS update package and extract it to a temporary location.
- Step 2** Copy the contents of the extracted recovery folder to the root directory of a USB drive. The recovery folder contains the bios.cap file that is required in this procedure.

**Note** The bios.cap file must be in the root directory of the USB drive. Do not rename this file. The USB drive must be formatted with either the FAT16 or FAT32 file system.

- Step 3** Insert the USB drive into a USB port on the node.
- Step 4** Reboot the node.
- Step 5** Return the node to main power mode by pressing the Power button on the front panel.

The node boots with the updated BIOS boot block. When the BIOS detects a valid bios.cap file on the USB drive, it displays this message:

```
Found a valid recovery file...Transferring to Cisco IMC
System would flash the BIOS image now...
System would restart with recovered image after a few seconds...
```

- Step 6** Wait for node to complete the BIOS update, and then remove the USB drive from the node.

**Note** During the BIOS update, Cisco IMC shuts down the node and the screen goes blank for about 10 minutes. Do not unplug the power cords during this update. Cisco IMC powers on the node after the update is complete.

---

## Using the BIOS Recovery Switch (SW12, Switch 5) and bios.cap File

You can use this switch to switch the server to use a recovery BIOS.

You will find it helpful to refer to the location of the CN3 header. See [Service Headers and Jumpers, on page 167](#).

---

**Step 1** Download the BIOS update package and extract it to a temporary location.

**Step 2** Copy the contents of the extracted recovery folder to the root directory of a USB drive. The recovery folder contains the bios.cap file that is required in this procedure.

**Note** The bios.cap file must be in the root directory of the USB drive. Do not rename this file. The USB drive must be formatted with either the FAT16 or FAT32 file system.

**Step 3** Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Node, on page 69](#). Disconnect power cords from all power supplies.

**Step 4** Slide the server out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.

**Caution** If you cannot safely view and access the component, remove the server from the rack.

**Step 5** Remove the top cover from the server as described in [Removing the Node Top Cover, on page 70](#).

**Step 6** Using your finger, gently slide the SW12 switch 5 to the ON position.

**Step 7** Reconnect AC power cords to the server. The server powers up to standby power mode.

**Step 8** Insert the USB thumb drive that you prepared in Step 2 into a USB port on the server.

**Step 9** Return the server to main power mode by pressing the Power button on the front panel.

The server boots with the updated BIOS boot block. When the BIOS detects a valid bios.cap file on the USB drive, it displays this message:

```
Found a valid recovery file...Transferring to Cisco IMC
System would flash the BIOS image now...
System would restart with recovered image after a few seconds...
```

**Step 10** Wait for server to complete the BIOS update, and then remove the USB drive from the server.

**Note** During the BIOS update, Cisco IMC shuts down the server and the screen goes blank for about 10 minutes. Do not unplug the power cords during this update. Cisco IMC powers on the server after the update is complete.

**Step 11** After the server has fully booted, power off the server again and disconnect all power cords.

**Step 12** Using your finger, gently slide the switch back to its original position (OFF).

**Note** If you do not reset the switch to its original position (OFF), after recovery completion you see the prompt, "Please remove the recovery jumper."

**Step 13** Replace the top cover, replace the server in the rack, replace power cords and any other cables, and then power on the server by pressing the Power button.

## Using the Clear BIOS Password Switch (SW12, Switch 6)

You can use this switch to clear the BIOS password.

You will find it helpful to refer to the location of the CN3 header. See [Service Headers and Jumpers](#), on page 167.

- 
- Step 1** Shut down and remove power from the node as described in [Shutting Down and Removing Power From the Node](#), on page 69. Disconnect power cords from all power supplies.
- Step 2** Slide the node out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.
- Caution** If you cannot safely view and access the component, remove the node from the rack.
- Step 3** Remove the top cover from the node as described in [Removing the Node Top Cover](#), on page 70.
- Step 4** Using your finger, gently slide the SW12 switch 6 to the ON position.
- Step 5** Reinstall the top cover and reconnect AC power cords to the node. The node powers up to standby power mode, indicated when the Power LED on the front panel is amber.
- Step 6** Return the node to main power mode by pressing the Power button on the front panel. The node is in main power mode when the Power LED is green.
- Note** You must allow the entire node to reboot to main power mode to complete the reset. The state of the switch cannot be determined without the host CPU running.
- Step 7** Press the Power button to shut down the node to standby power mode, and then remove AC power cords from the node to remove all power.
- Step 8** Remove the top cover from the node.
- Step 9** Reset the switch to its original position (OFF).
- Note** If you do not remove the switch to its original position (OFF), the BIOS password is cleared every time you power-cycle the node.
- Step 10** Replace the top cover, replace the node in the rack, replace power cords and any other cables, and then power on the node by pressing the Power button.
- 

## Using the Boot Alternate Cisco IMC Image Header (CN3, Pins 1-2)

You can use this Cisco IMC debug header to force the system to boot from an alternate Cisco IMC image.

You will find it helpful to refer to the location of the CN3 header. See [Service Headers and Jumpers](#), on page 167.

- 
- Step 1** Shut down and remove power from the node as described in [Shutting Down and Removing Power From the Node](#), on page 69. Disconnect power cords from all power supplies.
- Step 2** Slide the node out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.
- Caution** If you cannot safely view and access the component, remove the node from the rack.

- Step 3** Remove the top cover from the node as described in [Removing the Node Top Cover, on page 70](#).
- Step 4** Install a two-pin jumper across CN3 pins 1 and 2.
- Step 5** Reinstall the top cover and reconnect AC power cords to the node. The node powers up to standby power mode, indicated when the Power LED on the front panel is amber.
- Step 6** Return the node to main power mode by pressing the Power button on the front panel. The node is in main power mode when the Power LED is green.
- Note** When you next log in to Cisco IMC, you see a message similar to the following:
- ```
'Boot from alternate image' debug functionality is enabled.  
CIMC will boot from alternate image on next reboot or input power cycle.
```
- Note** If you do not remove the jumper, the node will boot from an alternate Cisco IMC image every time that you power cycle the node or reboot Cisco IMC.
- Step 7** To remove the jumper, press the Power button to shut down the node to standby power mode, and then remove AC power cords from the node to remove all power.
- Step 8** Remove the top cover from the node.
- Step 9** Remove the jumper that you installed.
- Step 10** Replace the top cover, replace the node in the rack, replace power cords and any other cables, and then power on the node by pressing the Power button.
- 

## Using the System Firmware Secure Erase Header (CN3, Pins 3-4)

You can use this header to securely erase system firmware from the node.

You will find it helpful to refer to the location of the CN3 header. See [Service Headers and Jumpers, on page 167](#).

---

- Step 1** Shut down and remove power from the node as described in [Shutting Down and Removing Power From the Node, on page 69](#). Disconnect power cords from all power supplies.
- Step 2** Slide the node out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.
- Caution** If you cannot safely view and access the component, remove the node from the rack.
- Step 3** Remove the top cover from the node as described in [Removing the Node Top Cover, on page 70](#).
- Step 4** Install a two-pin jumper across CN3 pins 3 and 4.
- Step 5** Reinstall the top cover and reconnect AC power cords to the node. The node powers up to standby power mode, indicated when the Power LED on the front panel is amber.
- Step 6** Return the node to main power mode by pressing the Power button on the front panel. The node is in main power mode when the Power LED is green.
- Note** You must allow the entire node to reboot to main power mode to complete the reset. The state of the jumper cannot be determined without the host CPU running.
- Step 7** Press the Power button to shut down the node to standby power mode, and then remove AC power cords from the node to remove all power.
- Step 8** Remove the top cover from the node.

**Step 9** Remove the jumper that you installed.

**Note** If you do not remove the jumper, the password is cleared every time you power-cycle the node.

**Step 10** Replace the top cover, replace the node in the rack, replace power cords and any other cables, and then power on the node by pressing the Power button.

---





# APPENDIX **A**

## Node Specifications

- [Physical Specifications, on page 175](#)
- [Environmental Specifications, on page 175](#)
- [Power Specifications, on page 176](#)
- [Power Cord Specifications, on page 179](#)

## Physical Specifications

The following table lists the physical specifications for the node.

**Table 18: Physical Specifications**

| Description                           | Specification                                                                  |
|---------------------------------------|--------------------------------------------------------------------------------|
| Height                                | 3.4 in. (86.4 mm)                                                              |
| Width                                 | 16.9 in. (429.0 mm)                                                            |
| Depth (length)                        | Server only: 29.5 in. (740.3 mm)<br>Server with slide rail: 31.0 in (787.4 mm) |
| Maximum weight (fully loaded chassis) | 57.5 lb. (26.1 Kg)                                                             |

## Environmental Specifications

The following table lists the environmental requirements and specifications for the node.

**Table 19: Physical Specifications**

| Description | Specification |
|-------------|---------------|
|-------------|---------------|

|                                                                                               |                                                                                                                                                                                                                                                                   |
|-----------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Temperature, Operating                                                                        | 41 to 95°F (5 to 35°C)<br><br>Derate the maximum temperature by 1°C per every 305 meters of altitude above sea level.<br><br>For general information, see the <a href="#">Cisco Unified Computing System Site Planning Guide: Data Center Power and Cooling</a> . |
| Temperature, non-operating<br>(when the node is stored or transported)                        | –40 to 149°F (–40 to 65°C)                                                                                                                                                                                                                                        |
| Humidity (RH), operating                                                                      | 10 to 90%                                                                                                                                                                                                                                                         |
| Humidity (RH), non-operating<br>(when the node is stored or transported)                      | 5 to 93%                                                                                                                                                                                                                                                          |
| Altitude, operating                                                                           | 0 to 10,000 feet                                                                                                                                                                                                                                                  |
| Altitude, non-operating<br>(when the node is stored or transported)                           | 0 to 40,000 feet                                                                                                                                                                                                                                                  |
| Sound power level<br>Measure A-weighted per ISO7779 LwAd (Bels)<br>Operation at 73°F (23°C)   | 5.5                                                                                                                                                                                                                                                               |
| Sound pressure level<br>Measure A-weighted per ISO7779 LpAm (dBA)<br>Operation at 73°F (23°C) | 40                                                                                                                                                                                                                                                                |

## Power Specifications



**Note** Do not mix power supply types or wattages in the node. Both power supplies must be identical.

You can get more specific power information for your exact node configuration by using the Cisco UCS Power Calculator:

<http://ucspowercalc.cisco.com>

The power specifications for the supported power supply options are listed in the following sections.

### 1050 W AC Power Supply

This section lists the specifications for each 1050 W AC power supply (Cisco part number HX-PSU1-1050W).

**Table 20: 1050 W AC Specifications**

| Description                  | Specification                                                               |
|------------------------------|-----------------------------------------------------------------------------|
| AC Input Voltage             | Nominal range: 100–120 VAC, 200–240 VAC<br>(Range: 90–132 VAC, 180–264 VAC) |
| AC Input Frequency           | Nominal range: 50 to 60Hz<br>(Range: 47–63 Hz)                              |
| Maximum AC Input current     | 12.5 A at 100 VAC<br>6.0 A at 208 VAC                                       |
| Maximum input volt-amperes   | 1250 VA at 100 VAC                                                          |
| Maximum inrush current       | 15 A (sub-cycle duration)                                                   |
| Maximum hold-up time         | 12 ms at 1050 W                                                             |
| Maximum output power per PSU | 800 W at 100–120 VAC<br>1050 W at 200–240 VAC                               |
| Power supply output voltage  | 12 VDC                                                                      |
| Power supply standby voltage | 12 VDC                                                                      |
| Efficiency rating            | Climate Savers Platinum Efficiency (80Plus Platinum certified)              |
| Form factor                  | RSP2                                                                        |
| Input connector              | IEC320 C14                                                                  |

## 1600 W AC Power Supply

This section lists the specifications for each 1600 W AC power supply (Cisco part number HX-PSU1-1600W).

**Table 21: 1600 W AC Specifications**

| Description                | Specification                                      |
|----------------------------|----------------------------------------------------|
| AC Input Voltage           | Nominal range: 200–240 VAC<br>(Range: 180–264 VAC) |
| AC Input Frequency         | Nominal range: 50 to 60Hz<br>(Range: 47–63 Hz)     |
| Maximum AC Input current   | 9.5 A at 200 VAC                                   |
| Maximum input volt-amperes | 1250 VA at 200 VAC                                 |

## 1050 W DC Power Supply

|                              |                                                                |
|------------------------------|----------------------------------------------------------------|
| Maximum inrush current       | 30 A at 35° C                                                  |
| Maximum hold-up time         | 80 ms at 1600 W                                                |
| Maximum output power per PSU | 1600 W at 200–240 VAC                                          |
| Power supply output voltage  | 12 VDC                                                         |
| Power supply standby voltage | 12 VDC                                                         |
| Efficiency rating            | Climate Savers Platinum Efficiency (80Plus Platinum certified) |
| Form factor                  | RSP2                                                           |
| Input connector              | IEC320 C14                                                     |

## 1050 W DC Power Supply

This section lists the specifications for each 1050 W DC power supply (Cisco part number HX-PSUV2-1050DC).

**Table 22: 1050 W DC Specifications**

| Description                  | Specification                                               |
|------------------------------|-------------------------------------------------------------|
| DC Input Voltage             | Nominal range: -48 to -60 VDC<br>(Range: -40 to -72 VDC)    |
| Maximum DC input current     | 32 A at -40 VDC                                             |
| Maximum input wattage        | 1234 W                                                      |
| Maximum inrush current       | 35 A (sub-cycle duration)                                   |
| Maximum hold-up time         | 5 ms at 100% load (1050 W main and 36 W standby)            |
| Maximum output power per PSU | 1050 W on 12 VDC main power<br>36 W on 12 VDC standby power |
| Power supply output voltage  | 12 VDC                                                      |
| Power supply standby voltage | 12 VDC                                                      |
| Efficiency rating            | ≥ 92% at 50% load                                           |
| Form factor                  | RSP2                                                        |
| Input connector              | Fixed 3-wire block                                          |

# Power Cord Specifications

Each power supply in the node has a power cord. Standard power cords or jumper power cords are available for connection to the node. The shorter jumper power cords, for use in racks, are available as an optional alternative to the standard power cords.



**Note** Only the approved power cords or jumper power cords listed below are supported.

**Table 23: Supported Power Cords**

| Description                                                                                               | Length (Feet) | Length (Meters) |
|-----------------------------------------------------------------------------------------------------------|---------------|-----------------|
| CAB-48DC-40A-8AWG<br>DC power cord, -48 VDC, 40 A, 8 AWG<br>Three-socket Mini-Fit connector to three-wire | 11.7          | 3.5             |
| CAB-C13-C14-AC<br>AC power cord, 10 A; C13 to C14, recessed receptacle                                    | 9.8           | 3.0             |
| CAB-250V-10A-AR<br>AC power cord, 250 V, 10 A<br>Argentina                                                | 8.2           | 2.5             |
| CAB-C13-C14-2M-JP<br>AC Power Cord, C13 to C14<br>Japan PSE Mark                                          | 6.6           | 2.0             |
| CAB-9K10A-EU<br>AC Power Cord, 250 V, 10 A; CEE 7/7 Plug<br>Europe                                        | 8.2           | 2.5             |
| CAB-250V-10A-IS<br>AC Power Cord, 250 V, 10 A<br>Israel                                                   | 8.2           | 2.5             |
| CAB-250V-10A-CN<br>AC power cord, 250 V, 10 A<br>PR China                                                 | 8.2           | 2.5             |
| CAB-ACTW<br>AC power cord, 250 V, 10 A<br>Taiwan                                                          | 7.5           | 2.3             |

|                                                                                       |     |      |
|---------------------------------------------------------------------------------------|-----|------|
| CAB-C13-CBN<br>AC cabinet jumper power cord, 250, 10 A,<br>C13 to C14                 | 2.2 | 0.68 |
| CAB-C13-C14-2M<br>AC cabinet jumper power cord, 250 V, 10 A,<br>C13 to C14            | 6.6 | 2.0  |
| CAB-9K10A-AU<br>AC power cord, 250 V, 10 A, 3112 plug,<br>Australia                   | 8.2 | 2.5  |
| CAB-N5K6A-NA<br>AC power cord, 200/240 V, 6 A,<br>North America                       | 8.2 | 2.5  |
| CAB-250V-10A-ID<br>AC power Cord, 250 V, 10 A,<br>India                               | 8.2 | 2.5  |
| CAB-9K10A-SW<br>AC power cord, 250 V, 10 A, MP232 plug<br>Switzerland                 | 8.2 | 2.5  |
| CAB-250V-10A-BR<br>AC power Cord, 250 V, 10 A<br>Brazil                               | 8.2 | 2.5  |
| CAB-9K10A-UK<br>AC power cord, 250 V, 10 A (13 A fuse), BS1363 plug<br>United Kingdom | 8.2 | 2.5  |
| CAB-9K12A-NA<br>AC power cord, 125 V, 13 A, NEMA 5-15 plug<br>North America           | 8.2 | 2.5  |
| CAB-AC-L620-C13<br>AC power cord, NEMA L6-20 to C13 connectors                        | 6.6 | 2.0  |
| CAB-9K10A-IT<br>AC power cord, 250 V, 10 A, CEI 23-16/VII plug<br>Italy               | 8.2 | 2.5  |

|                                                                                      |    |    |
|--------------------------------------------------------------------------------------|----|----|
| R2XX-DMYMPWRCORD<br>No power cord; PID option for ordering server with no power cord | NA | NA |
|--------------------------------------------------------------------------------------|----|----|







## APPENDIX **B**

# Storage Controller Considerations

---

This appendix provides storage controller (HBA) information.

- [Storage Controller Card Firmware Compatibility, on page 183](#)
- [RAID Backup \(Supercap\), on page 183](#)
- [Mixing Drive Types in RAID Groups, on page 184](#)
- [Storage Controller Cable Connectors and Backplanes, on page 184](#)
- [For More RAID Utility Information, on page 188](#)

## Storage Controller Card Firmware Compatibility

Firmware on the storage controller (RAID or HBA) must be verified for compatibility with the current Cisco IMC and BIOS versions that are installed on the node. If not compatible, upgrade or downgrade the storage controller firmware using the Host Upgrade Utility (HUU) for your firmware release to bring it to a compatible level.



**Note** **For nodes running in standalone mode only:** After you replace controller hardware (HX-RAID-M6T, HX-RAID-M6HD, HX-RAID-M6SD, HX-SAS-M6T, or HX-SAS-M6HD), you must run the Cisco UCS Host Upgrade Utility (HUU) to update the controller firmware, even if the firmware Current Version is the same as the Update Version. Running HUU is necessary to program any controller specific values to the storage controller for the specific node. If you do not run HUU, the storage controller may not be discovered..

See the HUU guide for your Cisco IMC release for instructions on downloading and using the utility to bring node components to compatible levels: [HUU Guides](#).

## RAID Backup (Supercap)

This server supports installation of one Supercap unit. The unit mounts to a bracket in-line with the fan modules.

The optional SCPM provides approximately three years of backup for the disk write-back cache DRAM in the case of a sudden power loss by offloading the cache to the NAND flash.

For Supercap unit replacement instructions, see [Replacing the Supercap \(RAID Backup\), on page 155](#).

## Mixing Drive Types in RAID Groups

For the best performance follow these guidelines:

- Use either all SAS or all SATA drives in a RAID group.
- Use the same capacity for each drive in the RAID group.
- Never mix HDDs and SSDs in the same RAID group.

## Storage Controller Cable Connectors and Backplanes

This section describes cabling for the storage controllers and backplanes. The SAS/SATA cables are factory-installed and are used for all supported internal controllers.

This section also contains diagrams that show the cable-to-drive mapping.

### **Cisco M6 12G Modular SAS RAID Controller or HBA For Up To 16 Drives (UCSC-RAID-M6T)**

The drive support differs by server version, as described in the following sections. These controllers are supported only in these server versions:


- SFF 12-Drives, SAS/SATA
- SFF 12-Drives NVMe

#### **SFF 12-Drives SAS/SATA**

This HW RAID or HBA option can control up to 12 front-loading SAS/SATA drives in this server version, plus two rear-loading SAS/SATA drives.

1. Connect a SAS/SATA cable from the small Slimline connector on the RAID card.
2. Connect a SAS/SATA cable to the Riser 3B connector on the PCE Riser 3 cage.

Figure 34: SFF, 12-Drive UCSC-C240-M6S Cabling With UCSC-RAID-M6T



#### Cisco 12G Modular SAS RAID Controller or HBA For Up To 28 Drives (UCSC-RAID-M6SD)

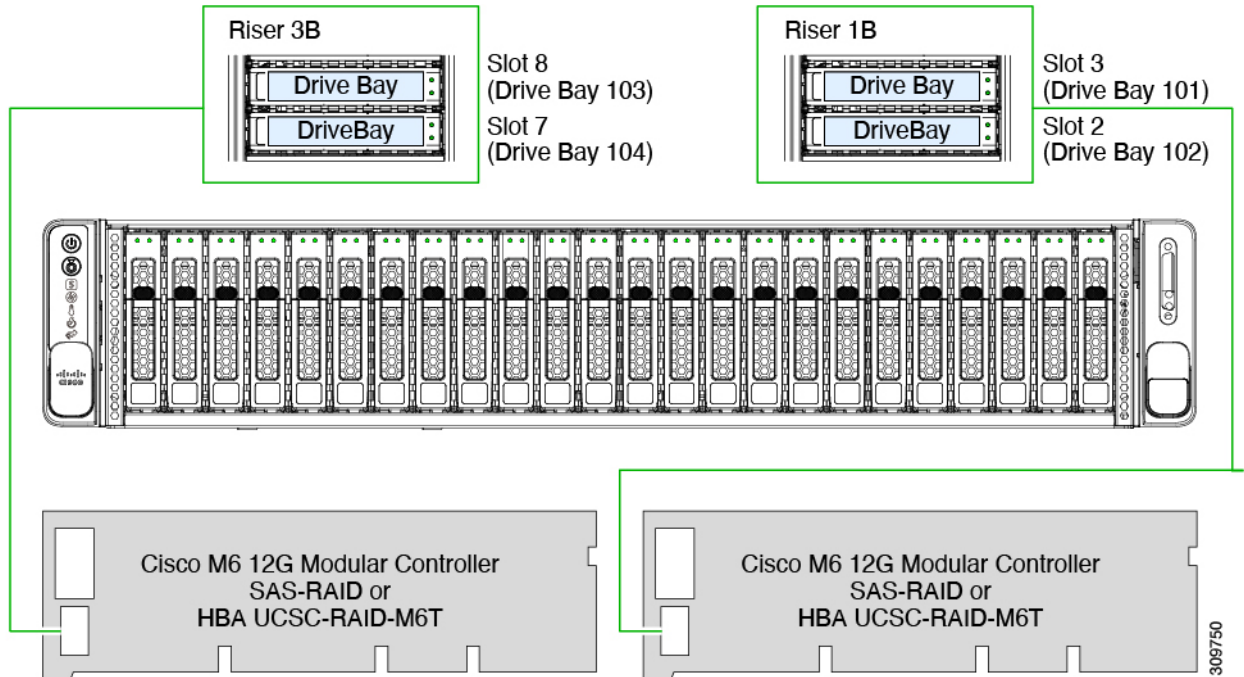
This controller is supported only in this server version:

- SFF 24-Drives SAS/SATA
- SFF 24 Drives NVMe

This HW RAID or HBA option can control up to 24 front-loading SAS/SATA drives in this server version, plus 2 rear-loading SAS/SATA drives.

1. Connect a SAS/SATA cable from the small Slimline connector on the RAID card.
2. Connect a SAS/SATA cable to the Riser 3B connector on the PCE Riser 3 cage.
3. Connect a SAS/SATA cable from the second small Slimline connector on the RAID card.
4. Connect a SAS/SATA cable to the Riser 1B connector on the PCE Riser 1 cage.

Figure 35: SFF, 24-Drive UCSC-C240-M6SX Cabling With UCSC-RAID-M6SD




### Cisco 12G Modular SAS RAID Controller or HBA For Up To 32 Drives (UCSC-RAID-M6HD)

#### LFF 12-Drives

This HW RAID or HBA option can control up to 12 front-loading SAS/SATA drives in this server version, plus 2 rear-loading SAS/SATA drives, and up to 4 optional mid mount drives. This option is only supported for the LFF drive version of the server.

To connect the RAID card to the front-loading drives, connect the split cable (Y cable) as follows:

1. Connect the single end (1) of the SAS/SATA cable to the RAID card.
2. Connect each of the dual-end connectors (2 and 3) to the two front backplane connectors.

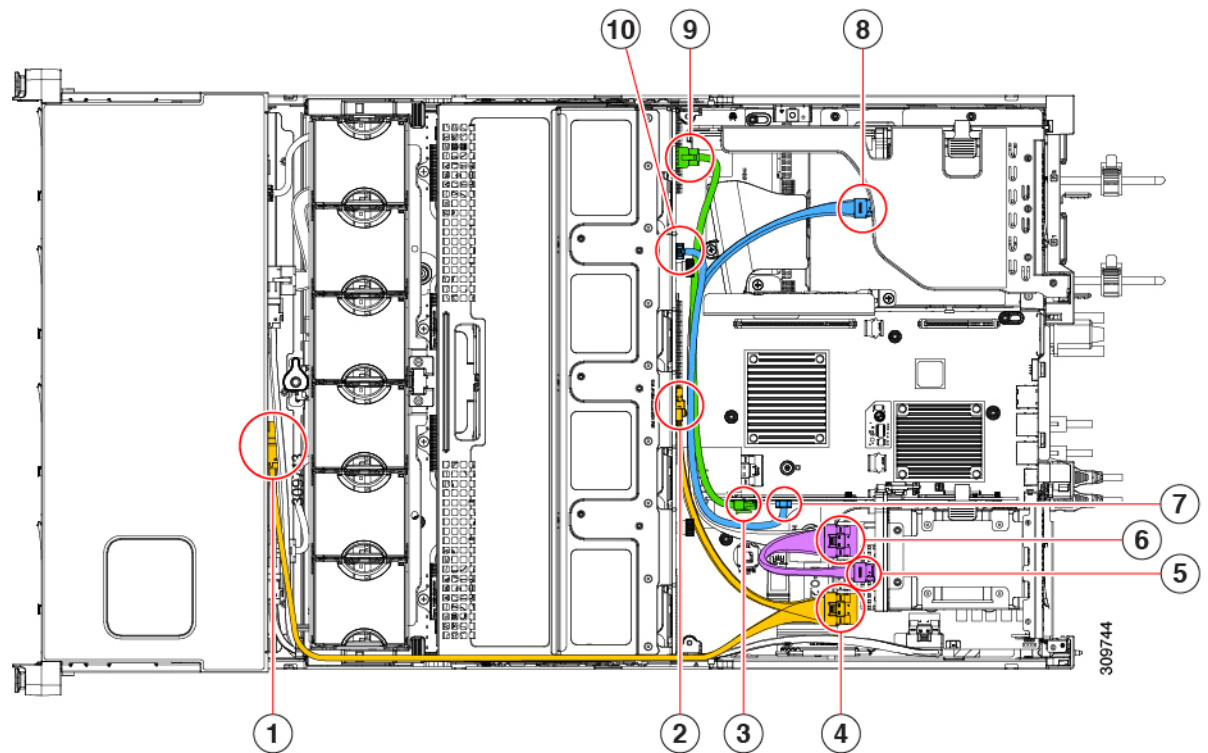


To connect the RAID card to the front drives and the mid-mount drives:

1. Connect one end of the SAS/SATA cable from the RAID card to the midplane connector.

To connect the rear drives:

1. Connect the SAS/SATA cable from the RAID card (6) to the Riser 1 connector (5).
2. Connect the SAS/SATA cable from the RAID card (7) to the Riser 3 connector (8) and the midplane connector (10).



## For More RAID Utility Information

### Cisco RAID Information

- For embedded SATA support of drives connected to SATA through the SATA Interposer card, the drives are managed directly as AHCI drives.
- The Cisco HX C Series M6 storage adapter supports only SAS/SATA drives. Tri-mode support is not supported.
- For basic information about RAID and for using the utilities for the RAID controller cards that are supported in Cisco servers, see the [Cisco UCS Servers RAID Guide](#).

### Broadcom RAID Information

- For hardware SAS MegaRAID configuration—[Broadcom 12Gb/s MegaRAID™ Tri-Mode Software User Guide](#).
- The Broadcom utilities have help documentation for more information about using the utilities. Go to [Broadcom StorCLI 12Gb/s MegaRAID™ Tri-Mode User Guide](#).



## APPENDIX **C**

# GPU Card Installation

This appendix contains configuration rules and installation procedures for the supported GPU cards.

- [Node Firmware Requirements, on page 189](#)
- [GPU Card Configuration Rules, on page 189](#)
- [Requirement For All GPUs: Memory-Mapped I/O Greater Than 4 GB, on page 190](#)
- [Installing a Single-Wide GPU Card, on page 191](#)
- [Installing a Double-Wide GPU Card, on page 194](#)
- [Replacing a Heatsink, on page 196](#)
- [Installing GRID Licenses From the NVIDIA Licensing Portal to the License Server, on page 202](#)

## Node Firmware Requirements

The following table lists the minimum node firmware versions for the supported GPU cards.

| GPU Card                                     | Cisco IMC/BIOS Minimum Version Required |
|----------------------------------------------|-----------------------------------------|
| NVIDIA Tesla A10 (HX-GPU-A10 or HX-GPU-A10=) | 4.2(1)                                  |
| NVIDIA Tesla A100 40GB (HX-GPU-A100=)        | 4.2(1)                                  |

## GPU Card Configuration Rules

Note the following rules when populating a node with GPU cards.

- The HX-C240-M6SX and HX-C240-M6N servers support a "GPU Ready" config which presets the node to accept a GPU at a later date.

This configuration sets the node with a with low-profile heatsink (HX-HSLP-M6) and GPU air duct (HX-ADGPU-240M6) so that only the GPU is required when you choose to expand the node's compute power. The GPU Ready config has the following considerations:

- The node must follow the same temperature limits as a node with a GPU installed, even if the node does not currently have a GPU installed. By following the temperature limits even when the GPU is not present, you ensure correct operation when the GPU is installed later.

- This configuration requires a low-profile heatsink and a GPU air duct. If you are ordering the GPU Ready configuration, you must select the GPU air duct PID to enable GPU ready configuration when ordering the node through the Cisco online ordering and configuration tool. Follow the additional rules displayed in the tool.
- All GPU cards must be procured from Cisco because of a unique SBIOS ID that is required by CIMC and UCSM.
- Do not mix different brands or models of GPU cards in the node.
- NVIDIA A10 Single Wide GPUs (HX-GPU-A10 or HX-GPU-A10) are supported:
  - This GPU can be populated in Riser 1A slots 2 (x16) and 3 (x8), Riser 2 slots 5 (x16) and 6 (x8), and Riser 3C slot 7 (x16).
  - Each node can support a maximum of five of these GPUs.
- NVIDIA A100 Double Wide GPUs (HX-GPU-A100) are supported:
  - This GPU requires x16 support.
  - This GPU can be populated in Riser 1A slot 2 (x16), Riser 2A slot 5 (x16), and Riser 3C slot 7 (x16).
  - Each node can support three of these GPUs.
- GPUs are not supported in Riser 1B or Riser 3B. Riser 3B cannot mechanically accept a GPU.
- The HX-C240M6-S and HX-C240M6-S servers support one full-height, full-length, double-wide GPU (PCIe slot 7 only) in Riser 3C.
- The HX-C240-M6N and HX-C240-M6N servers do not support any GPU in slot 7 of Riser 3C.
- HX managed servers are discoverable only if a PCIe VIC card is installed in slot 1 or slot 4 or an mLOM VIC card is installed in the mLOM slot. If you install double-width GPUs, they must be located in slots 2, 5, or 7. Therefore, if two GPUs are installed, UCSM managed servers are discoverable only if you install a VIC in slot 1, slot 4, or the mLOM slot. The node can support 2 PCIe VICs and 1 mLOM VIC along with 2 or 3 GPUs.
- Use the UCS power calculator at the following link to determine the power needed based on your node configuration: <http://ucspowercalc.cisco.com>

## Requirement For All GPUs: Memory-Mapped I/O Greater Than 4 GB

All supported GPU cards require enablement of the BIOS setting that allows greater than 4 GB of memory-mapped I/O (MMIO).

- Standalone Server: If the server is used in standalone mode, this BIOS setting is enabled by default:

```
Advanced > PCI Configuration > Memory Mapped I/O Above 4 GB [Enabled]
```

If you need to change this setting, enter the BIOS Setup Utility by pressing **F2** when prompted during bootup.



- If the server is integrated with Cisco UCS Manager and is controlled by a service profile, this setting is enabled by default in the service profile when a GPU is present.

To change this setting manually, use the following procedure.

**Step 1** Refer to the Cisco UCS Manager configuration guide (GUI or CLI) for your release for instructions on configuring service profiles:

[Cisco UCS Manager Configuration Guides](#)

**Step 2** Refer to the chapter on Configuring Server-Related Policies > Configuring BIOS Settings.

**Step 3** In the section of your profile for PCI Configuration BIOS Settings, set `Memory Mapped IO Above 4GB Config` to one of the following:

- **Disabled**—Does not map 64-bit PCI devices to 64 GB or greater address space.
- **Enabled**—Maps I/O of 64-bit PCI devices to 64 GB or greater address space.
- **Platform Default**—The policy uses the value for this attribute contained in the BIOS defaults for the server. Use this only if you know that the node BIOS is set to use the default enabled setting for this item.

**Step 4** Reboot the node.

**Note** Cisco UCS Manager pushes BIOS configuration changes through a BIOS policy or default BIOS settings to the Cisco Integrated Management Controller (CIMC) buffer. These changes remain in the buffer and do not take effect until the node is rebooted.

## Installing a Single-Wide GPU Card

Use the following procedure to install or replace an NVIDIA Tesla A10 Passive, 150W, 24GB, single-wide GPU or dual A10 GPUs.

With Cisco IMC version 4.2(1) and later, the node can support up to five NVIDIA Tesla A10 GPUs.

| Ordering PID   | Description |
|----------------|-------------|
| HX-RIS1A-240M6 | Riser 1A    |
| HX-RIS2A-240M6 | Riser 2A    |
| HX-RIS3C-240M6 | Riser 3C    |

**Step 1** Shut down and remove power from the node as described in [Shutting Down and Removing Power From the Node](#), on page 69.


**Step 2** Slide the node out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.

**Caution** If you cannot safely view and access the component, remove the node from the rack.

**Step 3** Remove the top cover from the node as described in [Removing the Node Top Cover, on page 70](#).

**Step 4** Remove the single-wide GPU card that you are replacing:


- a) Use two hands to flip up and grasp the blue riser handle and the blue finger grip area on the front edge of the riser, and then lift straight up.



309422

- b) On the bottom of the riser, push the release latch that holds the securing plate, and then swing the hinged securing plate open.
- c) Open the hinged card-tab retainer that secures the rear-panel tab of the card.

Figure 36: PCIe Riser Card Securing Mechanisms



|   |                                        |   |                          |
|---|----------------------------------------|---|--------------------------|
| 1 | Release latch on hinged securing plate | 3 | Hinged card-tab retainer |
| 2 | Hinged securing plate                  | - |                          |

- d) Disconnect the Y cable (UCSC-CBL-240M6) from the GPU ends (the black 8-pin connectors).

If you are inserting another GPU, you can leave the opposite end of the cable (the white connector) attached. If you are not inserting another GPU, you can leave the cable in place or completely remove it.

- e) Pull evenly on both ends of the single-wide GPU card to remove it from the socket on the PCIe riser.

If the riser has no card, remove the blanking panel from the rear opening of the riser.

### Step 5

Install a new single-wide GPU card:

- With the hinged card-tab retainer open, align the new single-wide GPU card with the empty socket on the PCIe riser.
- Push down evenly on both ends of the card until it is fully seated in the socket.
- Ensure that the card's rear panel tab sits flat against the riser rear-panel opening and then close the hinged card-tab retainer over the card's rear-panel tab.
- Swing the hinged securing plate closed on the bottom of the riser. Ensure that the clip on the plate clicks into the locked position.
- Position the PCIe riser over its socket on the motherboard and over the chassis alignment channels.
- Carefully push down on both ends of the PCIe riser to fully engage its connector with the sockets on the motherboard.
- Attach the Y cable to the new GPU by connecting the white end to the motherboard and one of the black ends to each GPU.

**Note** If the node has only one GPU in the riser, one of the black GPU cables will be unused. You can let it dangle in place.

- Step 6** Replace the top cover to the node.
- Step 7** Replace the node in the rack, replace power and network cables, and then fully power on the node by pressing the Power button.
- Step 8** Optional: Continue with [Installing Drivers to Support the GPU Cards, on page 206](#).

## Installing a Double-Wide GPU Card

Use the following procedure to install or replace the NVIDIA Tesla A100 250W 40GB Double-Wide GPU.

With Cisco IMC version 4.2(1) and later, the server can support up to three NVIDIA Tesla A100 GPUs.

| Server          | Component | Recommended T Ambient Spec, Normal |
|-----------------|-----------|------------------------------------|
| All C240 24 SFF | PCIe      | 35 C                               |
|                 | Rear HDD  | 30 C                               |
|                 | GPU A100  |                                    |
|                 | GPU A10   |                                    |
| C240 12 LFF     | Rear HDD  | 30 C                               |
|                 | GPU A100  | 28 C                               |
|                 | GPU A10   |                                    |


The NVIDIA GPU card might be shipped with two power cables: a straight cable and a Y-cable. The straight cable is used for connecting power to the GPU card in this server; do not use the Y-cable, which is used for connecting the GPU card in external devices only.

The supported NVIDIA GPU requires a C240 M5 NVIDIA Cable (HX-P100CBL-240M5).

- Step 1** Shut down and remove power from the node as described in [Shutting Down and Removing Power From the Node, on page 69](#).
- Step 2** Slide the node out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.
- Caution** If you cannot safely view and access the component, remove the node from the rack.
- Step 3** Remove the top cover from the node as described in [Removing the Node Top Cover, on page 70](#).
- Step 4** Remove an existing GPU card:
- Disconnect any existing cable from the GPU card.
  - Use two hands to grasp the metal bracket of the PCIe riser and lift straight up to disengage its connector from the socket on the motherboard. Set the riser on an antistatic surface.
  - On the bottom of the riser, press down on the clip that holds the securing plate.
  - Swing open the hinged securing plate to provide access.
  - Open the hinged plastic retainer that secures the rear-panel tab of the card.
  - Disconnect the GPU card's power cable from the power connector on the PCIe riser.

- g) Pull evenly on both ends of the GPU card to remove it from the socket on the PCIe riser.

**Figure 37: PCIe Riser Card Securing Mechanisms**



|   |                                        |   |                          |
|---|----------------------------------------|---|--------------------------|
| 1 | Release latch on hinged securing plate | 3 | Hinged card-tab retainer |
| 2 | Hinged securing plate                  | - |                          |

**Step 5** Install a new GPU card:

**Note** Observe the configuration rules for this node, as described in [GPU Card Configuration Rules, on page 189](#).

- Align the GPU card with the socket on the riser, and then gently push the card's edge connector into the socket. Press evenly on both corners of the card to avoid damaging the connector.
- Connect the GPU power cable. The straight power cable connectors are color-coded. Connect the cable's black connector into the black connector on the GPU card and the cable's white connector into the white GPU POWER connector on the PCIe riser.

**Caution** Do not reverse the straight power cable. Connect the *black* connector on the cable to the *black* connector on the GPU card. Connect the *white* connector on the cable to the *white* connector on the PCIe riser.


- Close the card-tab retainer over the end of the card.
- Swing the hinged securing plate closed on the bottom of the riser. Ensure that the clip on the plate clicks into the locked position.

- e) Position the PCIe riser over its socket on the motherboard and over the chassis alignment channels.
- f) Carefully push down on both ends of the PCIe riser to fully engage its connector with the sockets on the motherboard.

At the same time, align the GPU front support bracket (on the front end of the GPU card) with the securing latch that is on the node's air baffle.

**Step 6** Insert the GPU front support bracket into the latch that is on the air baffle:

- a) Pinch the latch release tab and hinge the latch toward the front of the node.
- b) Hinge the latch back down so that its lip closes over the edge of the GPU front support bracket.
- c) Ensure that the latch release tab clicks and locks the latch in place.



**Step 7** Replace the top cover to the node.

**Step 8** Replace the node in the rack, replace power and network cables, and then fully power on the node by pressing the Power button.

**Step 9** Optional: Continue with [Installing Drivers to Support the GPU Cards, on page 206](#).

## Replacing a Heatsink

For GPUs, the correct heatsink is the low-profile heatsink (HX-HSLP-M6), which has 4 T30 Torx screws on the main heatsink, and 2 Phillips-head screws on the extended heatsink. High profile heatsinks (HX-HSHP-240M6) cannot be used on a GPU.

Use the following procedures to replace the heatsink on a GPU.

- [Removing a Heat Sink, on page 197](#)
- [Installing a Heatsink, on page 200](#)

## Removing a Heat Sink


Use this procedure to remove a low-profile heatsink (HX-HSLP-M6) from a GPU.

**Step 1** [Removing the Node Top Cover.](#)

**Step 2** [Removing the Air Duct.](#)

**Step 3** Choose the appropriate GPU to remove.


- For a single-wide GPU
- a) Use two hands to flip up and grasp the blue riser handle and the blue finger grip area on the front edge of the riser, and then lift straight up.



309422

- b) On the bottom of the riser, push the release latch that holds the securing plate, and then swing the hinged securing plate open.
- c) Open the hinged card-tab retainer that secures the rear-panel tab of the card.

**Figure 38: PCIe Riser Card Securing Mechanisms**




|          |                                        |          |                          |
|----------|----------------------------------------|----------|--------------------------|
| <b>1</b> | Release latch on hinged securing plate | <b>3</b> | Hinged card-tab retainer |
| <b>2</b> | Hinged securing plate                  | -        |                          |

• For a double-wide GPU:

- a) Use two hands to grasp the metal bracket of the PCIe riser and lift straight up to disengage its connector from the socket on the motherboard. Set the riser on an antistatic surface.
- b) On the bottom of the riser, press down on the clip that holds the securing plate.
- c) Swing open the hinged securing plate to provide access.
- d) Open the hinged plastic retainer that secures the rear-panel tab of the card.
- e) Disconnect the GPU card's power cable from the power connector on the PCIe riser.
- f) Pull evenly on both ends of the GPU card to remove it from the socket on the PCIe riser.




**Figure 39: PCIe Riser Card Securing Mechanisms**



|          |                                        |          |                          |
|----------|----------------------------------------|----------|--------------------------|
| <b>1</b> | Release latch on hinged securing plate | <b>3</b> | Hinged card-tab retainer |
| <b>2</b> | Hinged securing plate                  | -        |                          |

**Step 4** Remove the CPU.

- a) Using a #2 Phillips screwdriver, loosen the two Phillips head screws for the extended heatsink.
- b) Using a T30 Torx driver, loosen the four Torx securing nuts.



- c) Push the rotating wires towards each other to move them to the unlocked position.

**Caution** Make sure that the rotating wires are as far inward as possible. When fully unlocked, the bottom of the rotating wire disengages and allows the removal of the CPU assembly. If the rotating wires are not fully in the unlocked position, you can feel resistance when attempting to remove the CPU assembly.

- d) Grasp the CPU and heatsink along the edge of the carrier and lift the CPU and heatsink off of the motherboard.

**Caution** While lifting the CPU assembly, make sure not to bend the heatsink fins. Also, if you feel any resistance when lifting the CPU assembly, verify that the rotating wires are completely in the unlocked position.

**Step 5** Remove the heatsink from the GPU.

**What to do next**

Install a low profile heatsink (HX-HSLP-M6) onto the GPU. See [Installing a Heatsink, on page 200](#).

## Installing a Heatsink

Use this procedure to install a low-profile heatsink (HX-HSLP-M6) on a GPU.

**Step 1** Apply new TIM, if needed.


**Note** The heatsink must have new TIM on the heatsink-to-CPU surface to ensure proper cooling and performance.

- If you are installing a new heatsink, it is shipped with a pre-applied pad of TIM. Go to step 2 below.
- If you are reusing a heatsink, you must remove the old TIM from the heatsink and then apply new TIM to the CPU surface from the supplied syringe. Continue with step a below.

- a) Apply the Bottle #1 cleaning solution that is included with the heatsink cleaning kit (UCSX-HSCK=), as well as the spare CPU package, to the old TIM on the heatsink and let it soak for a least 15 seconds.

- b) Wipe all of the TIM off the heatsink using the soft cloth that is included with the heatsink cleaning kit. Be careful to avoid scratching the heatsink surface.
- c) Completely clean the bottom surface of the heatsink using Bottle #2 to prepare the heatsink for installation.
- d) Using the syringe of TIM provided with the new CPU (HX-CPU-TIM=), apply 1.5 cubic centimeters (1.5 ml) of thermal interface material to the top of the CPU. Use the pattern shown in the following figure to ensure even coverage.

**Figure 40: Thermal Interface Material Application Pattern**




**Caution** Use only the correct heatsink for your CPU. CPU 1 uses heatsink HXB-HS-M6-R and CPU 2 uses heatsink HXB-HS-M6-F.

## Step 2 Install the heatsink.

- a) Push the rotating wires to the unlocked position so that they do not obstruct installation.
- b) Grasp the heatsink by the fins, align the pin 1 location on the heatsink with the pin 1 location on the GPU (2 in the following image), then seat the heatsink onto the CPU socket.
- c) Holding the CPU assembly level, orient it as shown and lower it onto the CPU socket.
- d) Push the rotating wires away from each other to lock the CPU assembly into the CPU socket.

**Caution** Make sure that you close the rotating wires completely before using the Torx driver to tighten the securing nuts.

- e) Set the T30 Torx driver to 12 in-lb of torque and tighten the 4 securing nuts to secure the CPU to the motherboard (3) first.
- f) Set the torque driver to 6 in-lb of torque and tighten the two Phillips head screws for the extended heatsink (4).



## Installing GRID Licenses From the NVIDIA Licensing Portal to the License Server

### Accessing the GRID License Server Management Interface

Open a web browser on the License Server host and access the URL <http://localhost:8080/licserver>.

If you configured the License Server host's firewall to permit remote access to the License Server, the management interface is accessible from remote machines at the URL <http://hostname:8080/licserver>

## Reading Your License Server's MAC Address

Your License Server's Ethernet MAC address is used as an identifier when registering the License Server with NVIDIA's Licensing Portal.

---

**Step 1** Access the GRID License Server Management Interface in a browser.

**Step 2** In the left-side License Server panel, select **Configuration**.

The License Server Configuration panel opens. Next to **Server host ID**, a pull-down menu lists the possible Ethernet MAC addresses.

**Step 3** Select your License Server's MAC address from the **Server host ID** pull-down.

**Note** It is important to use the same Ethernet ID consistently to identify the server when generating licenses on NVIDIA's Licensing Portal. NVIDIA recommends that you select one entry for a primary, non-removable Ethernet interface on the platform.

---

## Installing Licenses From the Licensing Portal

---

**Step 1** Access the GRID License Server Management Interface in a browser.

**Step 2** In the left-side License Server panel, select **Configuration**.

The License Server Configuration panel opens.

**Step 3** Use the License Server Configuration menu to install the .bin file that you generated earlier.

- a) Click **Choose File**.
- b) Browse to the license .bin file that you want to install and click **Open**.
- c) Click **Upload**.

The license file is installed on your License Server. When installation is complete, you see the confirmation message, "Successfully applied license file to license server."

---

## Viewing Available GRID Licenses

Use the following procedure to view which licenses are installed and available, along with their properties.

---

**Step 1** Access the GRID License Node Management Interface in a browser.

**Step 2** In the left-side License Server panel, select **Licensed Feature Usage**.

**Step 3** Click on a feature in the **Feature** column to see detailed information about the current usage of that feature.

---

## Viewing Current License Usage

Use the following procedure to view information about which licenses are currently in-use and borrowed from the server.

- 
- Step 1** Access the GRID License Node Management Interface in a browser.
  - Step 2** In the left-side License Node panel, select **Licensed Clients**.
  - Step 3** To view detailed information about a single licensed client, click on its **Client ID** in the list.
- 

## Managing GRID Licenses

Features that require GRID licensing run at reduced capability until a GRID license is acquired.

### Acquiring a GRID License on Windows

- 
- Step 1** Open the NVIDIA Control Panel using one of the following methods:
    - Right-click on the Windows desktop and select **NVIDIA Control Panel** from the menu.
    - Open Windows Control Panel and double-click the **NVIDIA Control Panel** icon.
  - Step 2** In the NVIDIA Control Panel left-pane under Licensing, select **Manage License**.

The Manage License task pane opens and shows the current license edition being used. The GRID software automatically selects the license edition based on the features that you are using. The default is Tesla (unlicensed).
  - Step 3** If you want to acquire a license for GRID Virtual Workstation, under License Edition, select **GRID Virtual Workstation**.
  - Step 4** In the **License Server** field, enter the address of your local GRID License Server. The address can be a domain name or an IP address.
  - Step 5** In the **Port Number** field, enter your port number or leave it set to the default used by the server, which is 7070.
  - Step 6** Select **Apply**.

The system requests the appropriate license edition from your configured License Server. After a license is successfully acquired, the features of that license edition are enabled.

**Note** After you configure licensing settings in the NVIDIA Control Panel, the settings persist across reboots.

---

### Acquiring a GRID License on Linux

- 
- Step 1** Edit the configuration file `/etc/nvidia/gridd.conf`:

```
sudo vi /etc/nvidia/gridd.conf
```
  - Step 2** Edit the `ServerUrl` line with the address of your local GRID License Server.

The address can be a domain name or an IP address. See the example file below.

**Step 3** Append the port number (default 7070) to the end of the address with a colon. See the example file below.

**Step 4** Edit the FeatureType line with the integer for the license type. See the example file below.

- GRID vGPU = 1
- GRID Virtual Workstation = 2

**Step 5** Restart the nvidia-gridd service.

```
sudo service nvidia-gridd restart
```

The service automatically acquires the license edition that you specified in the FeatureType line. You can confirm this in /var/log/messages.

**Note** After you configure licensing settings in the NVIDIA Control Panel, the settings persist across reboots.

Sample configuration file:

```
# /etc/nvidia/gridd.conf - Configuration file for NVIDIA Grid Daemon
# Description: Set License Server URL
# Data type: string
# Format: "<address>:<port>"
ServerUrl=10.31.20.45:7070

# Description: Set Feature to be enabled
# Data type: integer
# Possible values:
# 1 => for GRID vGPU
# 2 => for GRID Virtual Workstation
FeatureType=2
```

## Using gpumodeswitch

The command line utility `gpumodeswitch` can be run in the following environments:

- Windows 64-bit command prompt (requires administrator permissions)
- Linux 32/64-bit shell (including Citrix XenServer dom0) (requires root permissions)



**Note** Consult NVIDIA product release notes for the latest information on compatibility with compute and graphic modes.

The `gpumodeswitch` utility supports the following commands:

- `--listgpumodes`

Writes information to a log file named `listgpumodes.txt` in the current working directory.

- `--gpumode graphics`

Switches to graphics mode. Switches mode of all supported GPUs in the server unless you specify otherwise when prompted.

- `--gpumode compute`

Switches to compute mode. Switches mode of all supported GPUs in the server unless you specify otherwise when prompted.




---

**Note** After you switch GPU mode, reboot the server to ensure that the modified resources of the GPU are correctly accounted for by any OS or hypervisor running on the server.

---

## Installing Drivers to Support the GPU Cards

After you install the hardware, you must update to the correct level of server BIOS and then install GPU drivers and other software in this order:

1. Update the server BIOS.
2. Update the GPU drivers.

### 1. Updating the Server BIOS

Install the latest Cisco HX C240 M6 node BIOS by using the Host Upgrade Utility for the Cisco HX C240 M6 node.




---

**Note** You must do this procedure before you update the NVIDIA drivers.

---

- 
- Step 1** Navigate to the following URL: <http://www.cisco.com/cisco/software/navigator.html>.
  - Step 2** Click **Servers–Unified Computing** in the middle column.
  - Step 3** Click **Cisco HX C-Series Rack-Mount Standalone Node Software** in the right-hand column.
  - Step 4** Click the name of your model of node in the right-hand column.
  - Step 5** Click **Cisco HyperFlex Firmware**.
  - Step 6** Click the release number.
  - Step 7** Click **Download Now** to download the `hx-node platform-huu-version_number.iso` file.
  - Step 8** Verify the information on the next page, and then click **Proceed With Download**.
  - Step 9** Continue through the subsequent screens to accept the license agreement and browse to a location where you want to save the file.
  - Step 10** Use the Host Upgrade Utility to update the server BIOS.  
The user guides for the Host Upgrade Utility are at [Utility User Guides](#).
- 

### 2. Updating the GPU Card Drivers

After you update the node BIOS, you can install GPU drivers to your hypervisor virtual machine.



- 
- Step 1** Install your hypervisor software on a computer. Refer to your hypervisor documentation for the installation instructions.
- Step 2** Create a virtual machine in your hypervisor. Refer to your hypervisor documentation for instructions.
- Step 3** Install the GPU drivers to the virtual machine. Download the drivers from either:
- NVIDIA Enterprise Portal for GRID hypervisor downloads (requires NVIDIA login): <https://nvidia.flexnetoperations.com/>
  - NVIDIA public driver area: <http://www.nvidia.com/Download/index.aspx>
  - AMD: <http://support.amd.com/en-us/download>
- Step 4** Restart the node.
- Step 5** Check that the virtual machine is able to recognize the GPU card. In Windows, use the Device Manager and look under Display Adapters.
-

