



Cisco HX240c M5 HyperFlex Node (Hybrid and All-Flash Models) Installation Guide

First Published: 2017-09-15

Last Modified: 2019-11-13

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 NON-INFRINGEMENT 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.

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: www.cisco.com/go/trademarks. 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)

© 2017 Cisco Systems, Inc. All rights reserved.



CHAPTER 1

Overview

- [Overview, on page 1](#)
- [Cisco HyperFlex Systems Related Documentation, on page 2](#)
- [HX240c SFF Drives External Features, on page 2](#)
- [HX240c LFF Drives External Features, on page 3](#)
- [HX240c SFF Drives Serviceable Component Locations, on page 5](#)
- [HX240c LFF Drives Serviceable Component Locations, on page 7](#)
- [Summary of Node Features, on page 8](#)

Overview

This guide covers all Cisco HX240c Hybrid and All-Flash models.

The following table lists the minimum levels of Cisco HyperFlex software required for the small form-factor (SFF) 2.5-inch drives versions and the large form-factor (LFF) 3.5-inch drives version.



Note The software requirements given below are for the base chassis. Certain configurable components might require later software levels, as noted in this guide.

Table 1: HX240c M5 System Minimum Software Requirements

System Version	Cisco HyperFlex Software Minimum Level
HX240c M5 SFF Hybrid (HX240C-M5SX)	2.6.1a or later
HX240c M5 SFF All-Flash (HXAF240C-M5SX)	2.6.1a or later
HX240c M5 LFF Hybrid (HX240C-M5L)	3.0.1 or later

Cisco HyperFlex Systems Related Documentation

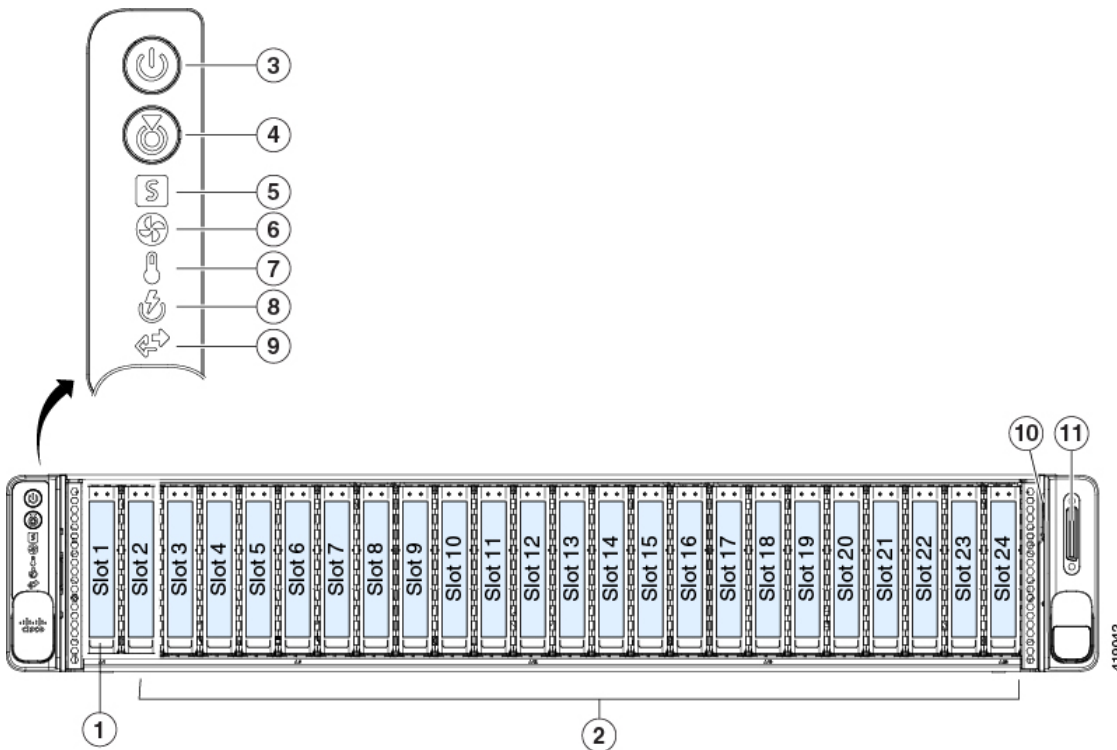
Links for related Cisco HyperFlex Systems documentation such as the Getting Started Guide, Administration Guide, and Release Notes are listed in the [Cisco HyperFlex Systems Documentation Roadmap](#).

HX240c SFF Drives External Features

HX240c SFF Drives Front Panel Features

For definitions of LED states, see [Front-Panel LEDs](#), on page 19.

Figure 1: SFF Front Panel




1	Drive bay 1: Housekeeping SSD for logs	7	Temperature status LED
2	2.5-inch drive bays 2 – 24: <ul style="list-style-type: none"> Hybrid HX240C-M5SX: 6 to 23 persistent data HDDs All-Flash HXAF240C-M5SX: 6 to 23 persistent data SSDs 	8	Power supply status LED
3	Power button/power status LED	9	Network link activity LED

4	Unit identification button/LED	10	Pull-out asset tag
5	System status LED	11	KVM connector (used with KVM cable that provides one DB-15 VGA, one DB-9 serial, and two USB connectors)
6	Fan status LED	-	

HX240c SFF Drives Rear Panel Features

For definitions of LED states, see [Rear-Panel LEDs](#), on page 22.

Figure 2: SFF Rear Panel




1	PCIe riser 1 (PCIe slot 1, 2, 3)	7	USB 3.0 ports (two)
2	PCIe riser 2B (PCIe slots 4, 5, 6)	8	Dual 1-Gb/10-Gb Ethernet ports (LAN1 and LAN2) The dual LAN ports can support 1 Gbps and 10 Gbps, depending on the link partner capability.
3	Rear drive bay 25: Caching SSD	9	VGA video port (DB-15 connector)
4	Power supplies (two, redundant as 1+1)	10	1-Gb Ethernet dedicated management port
5	Threaded holes for dual-hole grounding lug.	11	Serial port (RJ-45 connector)
6	Modular LAN-on-motherboard (mLOM) card slot (x16) for Cisco VIC card	12	Rear unit identification button/LED

HX240c LFF Drives External Features

HX240c LFF Drives Front Panel Features

For definitions of LED states, see [Front-Panel LEDs](#), on page 19.

Figure 3: LFF Front Panel




1	3.5-inch drive bays 1 – 12: • Hybrid HX240C-M5L: 6 to 12 persistent data HDDs	7	Power supply status LED
2	Power button/power status LED	8	Network link activity LED
3	Unit identification button/LED	9	Pull-out asset tag
4	System status LED	10	KVM connector (used with KVM cable that provides one DB-15 VGA, one DB-9 serial, and two USB connectors)
5	Fan status LED	-	
6	Temperature status LED	-	

HX240c LFF Drives Rear Panel Features

For definitions of LED states, see [Rear-Panel LEDs, on page 22](#).

Figure 4: LFF Rear Panel



1	PCIe riser 1 (PCIe slot 1, 2, 3)	8	USB 3.0 ports (two)
2	PCIe riser 2B (PCIe slots 4, 5, 6)	9	Dual 1-Gb/10-Gb Ethernet ports (LAN1 and LAN2) The dual LAN ports can support 1 Gbps and 10 Gbps, depending on the link partner capability.
3	Rear drive bay 13: Caching SSD	10	VGA video port (DB-15 connector)
4	Rear drive bay 14: Housekeeping SSD for logs	11	1-Gb Ethernet dedicated management port
5	Power supplies (two, redundant as 1+1)	12	Serial port (RJ-45 connector)
6	Threaded holes for dual-hole grounding lug.	13	Rear unit identification button/LED
7	Modular LAN-on-motherboard (mLOM) card slot (x16) for Cisco VIC card		

HX240c SFF Drives Serviceable Component Locations

This topic shows the locations of the field-replaceable components and service-related items.



Note The internal component *locations* are the same for the SFF and LFF versions of the node unless otherwise noted. The functions of the drives differ in SFF and LFF versions, as described in this document.

9	Power supplies (hot-swappable when redundant as 1+1)	-	
---	--	---	--

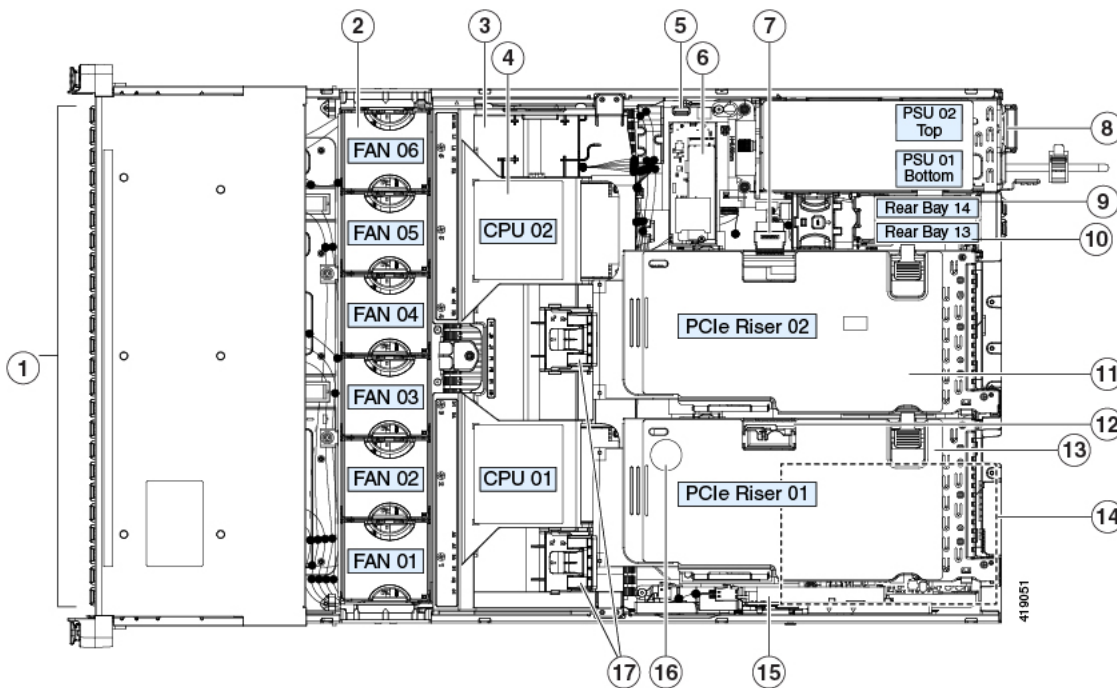
HX240c LFF Drives Serviceable Component Locations

This topic shows the locations of the field-replaceable components and service-related items for Hybrid HX240C-M5L.



Note The internal component *locations* are the same for the SFF and LFF versions of the node unless otherwise noted. The functions of the drives differ in SFF and LFF versions, as described in this document.

Figure 6: LFF Serviceable Component Locations



1	3.5-inch drive bays 1 – 12: • Hybrid HX240C-M5L: 6 to 12 persistent data HDDs	10	Rear drive bay 13: Caching SSD
2	Cooling fan modules (six, hot-swappable)	11	PCIe riser 2B (PCIe slots 4, 5, 6)
3	DIMM sockets on motherboard (up to 12 per CPU) Not visible under air baffle in this view.	12	Micro-SD card socket on PCIe riser 1

4	CPUs and heatsinks (two) Not visible under air baffle in this view.	13	PCIe riser 1 (PCIe slot 1, 2, 3)
5	Internal, vertical USB 3.0 port on motherboard	14	Modular LOM (mLOM) card bay on chassis floor (x16 PCIe lane), not visible in this view
6	Mini-storage module for SATA M.2 SSD boot drive	15	Cisco modular storage controller HBA slot (dedicated PCIe slot)
7	PCIe cable connector on PCIe riser 2B for rear-loading NVMe SSD	16	RTC battery, vertical socket on motherboard
8	Power supplies (hot-swappable when redundant as 1+1)	17	Securing clips for GPU cards on air baffle
9	Rear drive bay 14: Housekeeping SSD for logs	-	

Summary of Node Features

The following table lists a summary of node features.

Feature	Description
Chassis	Two rack-unit (2RU) chassis
Central Processor	Two identical CPUs from the Intel Xeon Processor Scalable Family.
Memory	24 DDR4 DIMM sockets on the motherboard (12 each CPU)
Multi-bit error protection	Multi-bit error protection is supported
Baseboard management	BMC, running Cisco Integrated Management Controller (Cisco IMC) firmware. Depending on your Cisco IMC settings, Cisco IMC can be accessed through the 1-Gb dedicated management port, the 1-Gb/10-Gb Ethernet LAN ports, or a Cisco virtual interface card.

Feature	Description
Network and management I/O	<p>Rear panel:</p> <ul style="list-style-type: none"> • One 1-Gb Ethernet dedicated management port (RJ-45 connector) • Two 1-Gb/10-Gb BASE-T Ethernet LAN ports (RJ-45 connectors) <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"> • One RS-232 serial port (RJ-45 connector) • One VGA video connector port (DB-15 connector) • Two USB 3.0 ports <p>Front panel:</p> <ul style="list-style-type: none"> • One front-panel keyboard/video/mouse (KVM) connector that is used with the KVM cable, which provides two USB 2.0, one VGA, and one DB-9 serial connector.
Modular LOM	One dedicated socket (x16 PCIe lane) that can be used to add a Cisco VIC mLOM card for additional rear-panel connectivity.
Power	<p>One or two power supplies, redundant as 1+1 when two power supplies are present. Several options are available:</p> <ul style="list-style-type: none"> • AC power supplies 1050 W AC each • AC power supplies 1600 W AC each • DC power supplies 1050 W DC each <p>Do not mix power supply types or wattages in the node.</p>
ACPI	The advanced configuration and power interface (ACPI) 4.0 standard is supported.
Cooling	Six hot-swappable fan modules for front-to-rear cooling.
PCIe I/O	Six horizontal PCIe expansion slots on two PCIe riser assemblies.
InfiniBand	The PCIe bus slots in this node support the InfiniBand architecture.

Feature	Description
Storage, external loading: Small form-factor (SFF) drive versions	Drives are installed into front- and rear-panel drive bays. <ul style="list-style-type: none"> • Front drive bay 1 is dedicated to one Housekeeping SSD for logs. • Front drive bays 2 - 24 are dedicated to persistent data drives: <ul style="list-style-type: none"> • Hybrid HX240C-M5SX: 6 to 23 HDDs. • All-Flash HXAF240C-M5SX: 6 to 23 SSDs. • Rear drive bay 25 is dedicated to one Caching SSD.
Storage, external loading: Large form-factor (LFF) drive versions	Drives are installed into front- and rear-panel drive bays. <ul style="list-style-type: none"> • Front drive bays 1 - 12 are dedicated to persistent data drives: <ul style="list-style-type: none"> • Hybrid HX240C-M5L: 6 to 12 HDDs. • Rear drive bay 13 is dedicated to one Caching SSD. • Rear drive bay 14 is dedicated to one Housekeeping SSD for logs.
Storage, internal (all versions)	The node has these internal storage options: <ul style="list-style-type: none"> • Mini-storage module that supports a SATA M.2 SSD Boot drive. • One Micro-SD card socket on PCIe riser 1. • One USB 3.0 port on the motherboard.
Storage management	The node has a dedicated internal PCIe slot that supports an HBA storage controller.
Integrated video	Integrated VGA video.



CHAPTER 2

Installing the Node

- [Preparing for Installation](#), on page 11
- [Installing the Node in a Rack](#), on page 13
- [Initial Node Setup](#), on page 17

Preparing for Installation

This chapter describes how to install the node in a rack and perform initial setup.

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 nodes 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 nodes, which could result in overheating, higher fan speeds, and higher power consumption. We recommend that you mount your nodes on rail kits when you are installing them into the rack because these rails provide the minimal spacing required between the nodes. No additional spacing between the nodes 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, which can have substantial current draw fluctuations from fluctuating data traffic patterns.

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

- 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 93](#).
- Ensure that the cabinet or rack meets the requirements listed in the [Rack Requirements, on page 13](#).
- Ensure that the site power meets the power requirements listed in the [Power Specifications, on page 94](#). 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-RAILB-M4= (ball-bearing slide rail kit)
- Cisco part UCSC-CMAF-M4= (cable management arm)

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

Figure 7: Attaching the Inner Rail to the Side of the node




1	Front of node	2	Inner rail
---	---------------	---	------------

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.

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 nodes 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 16 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 Node Setup

Refer to the following documentation:

- For instructions on physically installing the nodes in racks, see [Installing the Node in a Rack, on page 13](#).
- For installation instructions on configuring system fabric interconnects and connecting the HX Series HyperFlex nodes to the fabrics, refer to the Cisco HyperFlex Systems Getting Started Guide. See the [Cisco HyperFlex Systems Documentation Roadmap](#).
- For instructions on adding the HX220c node as a compute-only node, see the "Adding a Compute Node" section in the Cisco HyperFlex Systems Getting Started Guide. See the [Cisco HyperFlex Systems Documentation Roadmap](#).
- For instructions on starting the node in standalone mode *for troubleshooting purposes only*, see [Setting Up the Node in Standalone Mode, on page 88](#).



CHAPTER 3

Maintaining the Node

- [Status LEDs and Buttons, on page 19](#)
- [Decommissioning the Node Using Cisco UCS Manager, on page 25](#)
- [Shutting Down and Removing Power From the Node, on page 26](#)
- [Post-Maintenance Procedures, on page 29](#)
- [Removing the Node Top Cover, on page 30](#)
- [Removing and Replacing Components, on page 31](#)
- [Service Headers and Jumpers, on page 80](#)
- [Setting Up the Node in Standalone Mode, on page 88](#)

Status LEDs and Buttons

This section contains information for interpreting LED states.

Front-Panel LEDs

Figure 12: SFF Front Panel LEDs




Figure 13: LFF Front Panel LEDs

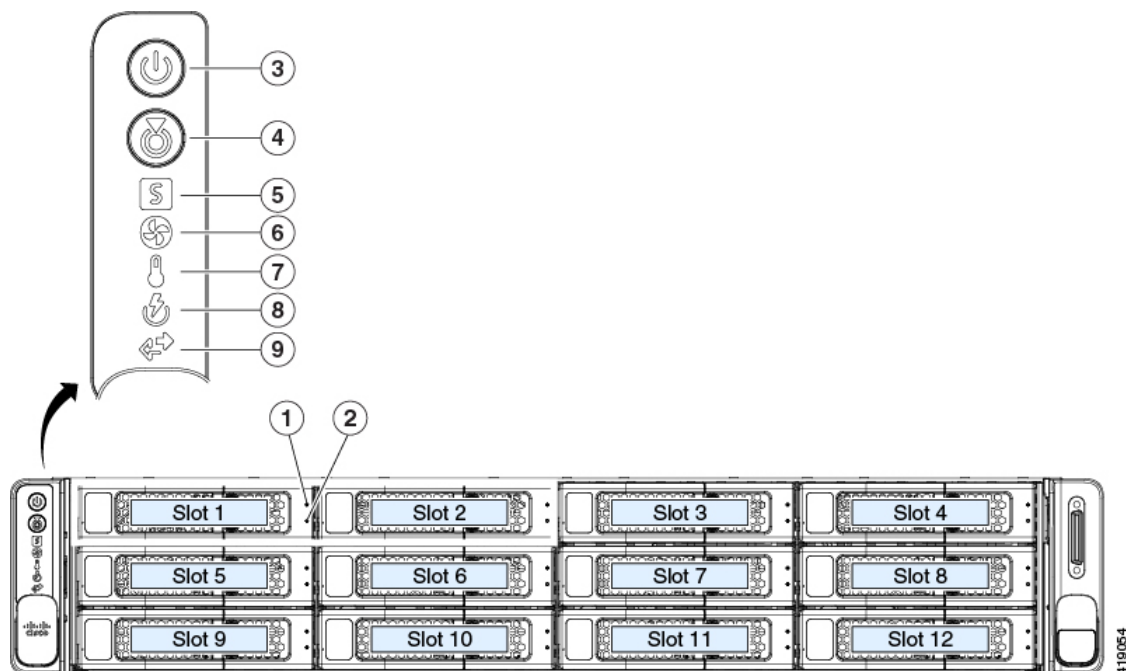


Table 2: Front Panel LEDs, Definition of States

	LED Name	States
1 SAS	SAS/SATA drive fault	<ul style="list-style-type: none"> • Off—The hard drive is operating properly. • Amber—Drive fault detected. • Amber, blinking—The device is rebuilding. • Amber, blinking with one-second interval—Drive locate function activated in the software.
2 SAS	SAS/SATA drive activity LED	<ul style="list-style-type: none"> • Off—There is no hard drive in the hard drive tray (no access, no fault). • Green—The hard drive is ready. • Green, blinking—The hard drive is reading or writing data.
3	Power button/LED	<ul style="list-style-type: none"> • Off—There is no AC power to the server. • Amber—The server is in standby power mode. Power is supplied only to the Cisco IMC and some motherboard functions. • Green—The server is in main power mode. Power is supplied to all server components.

4	Unit identification	<ul style="list-style-type: none"> • Off—The unit identification function is not in use. • Blue, blinking—The unit identification function is activated.
5	System health	<ul style="list-style-type: none"> • Green—The server is running in normal operating condition. • Green, blinking—The server is performing system initialization and memory check. • Amber, steady—The server is in a degraded operational state (minor fault). For example: <ul style="list-style-type: none"> • Power supply redundancy is lost. • CPUs are mismatched. • At least one CPU is faulty. • At least one DIMM is faulty. • At least one drive in a RAID configuration failed. • Amber, 2 blinks—There is a major fault with the system board. • Amber, 3 blinks—There is a major fault with the memory DIMMs. • Amber, 4 blinks—There is a major fault with the CPUs.
6	Power supply status	<ul style="list-style-type: none"> • Green—All power supplies are operating normally. • Amber, steady—One or more power supplies are in a degraded operational state. • Amber, blinking—One or more power supplies are in a critical fault state.
7	Fan status	<ul style="list-style-type: none"> • Green—All fan modules are operating properly. • Amber, blinking—One or more fan modules breached the non-recoverable threshold.
8	Network link activity	<ul style="list-style-type: none"> • Off—The Ethernet LOM port link is idle. • Green—One or more Ethernet LOM ports are link-active, but there is no activity. • Green, blinking—One or more Ethernet LOM ports are link-active, with activity.

9	Temperature status	<ul style="list-style-type: none"> • Green—The server is operating at normal temperature. • Amber, steady—One or more temperature sensors breached the critical threshold. • Amber, blinking—One or more temperature sensors breached the non-recoverable threshold.
---	--------------------	---

Rear-Panel LEDs

The rear panel LEDs are identical for SFF and LFF versions of the node. The rear drive bay numbering follows sequentially from the number of front drive bays.

Figure 14: Rear Panel LEDs




Table 3: Rear Panel LEDs, Definition of States

	LED Name	States
1	1-Gb/10-Gb Ethernet link speed (on both LAN1 and LAN2)	<ul style="list-style-type: none"> • Off—Link speed is 100 Mbps. • Amber—Link speed is 1 Gbps. • Green—Link speed is 10 Gbps.
2	1-Gb/10-Gb Ethernet link status (on both LAN1 and LAN2)	<ul style="list-style-type: none"> • Off—No link is present. • Green—Link is active. • Green, blinking—Traffic is present on the active link.
3	1-Gb Ethernet dedicated management link speed	<ul style="list-style-type: none"> • Off—Link speed is 10 Mbps. • Amber—Link speed is 100 Mbps. • Green—Link speed is 1 Gbps.

4	1-Gb Ethernet dedicated management link status	<ul style="list-style-type: none"> • Off—No link is present. • Green—Link is active. • Green, blinking—Traffic is present on the active link.
5	Rear unit identification	<ul style="list-style-type: none"> • Off—The unit identification function is not in use. • Blue, blinking—The unit identification function is activated.
6	Power supply status (one LED each power supply unit)	<p>AC power supplies:</p> <ul style="list-style-type: none"> • Off—No AC input (12 V main power off, 12 V standby power off). • Green, blinking—12 V main power off; 12 V standby power on. • Green, solid—12 V main power on; 12 V standby power on. • Amber, blinking—Warning threshold detected but 12 V main power on. • Amber, solid—Critical error detected; 12 V main power off (for example, over-current, over-voltage, or over-temperature failure). <p>DC power supplies:</p> <ul style="list-style-type: none"> • Off—No DC input (12 V main power off, 12 V standby power off). • Green, blinking—12 V main power off; 12 V standby power on. • Green, solid—12 V main power on; 12 V standby power on. • Amber, blinking—Warning threshold detected but 12 V main power on. • Amber, solid—Critical error detected; 12 V main power off (for example, over-current, over-voltage, or over-temperature failure).
7 SAS	<p>SAS/SATA drive fault</p> <p>Note NVMe solid state drive (SSD) drive tray LEDs have different behavior than SAS/SATA drive trays.</p>	<ul style="list-style-type: none"> • Off—The hard drive is operating properly. • Amber—Drive fault detected. • Amber, blinking—The device is rebuilding. • Amber, blinking with one-second interval—Drive locate function activated in the software.

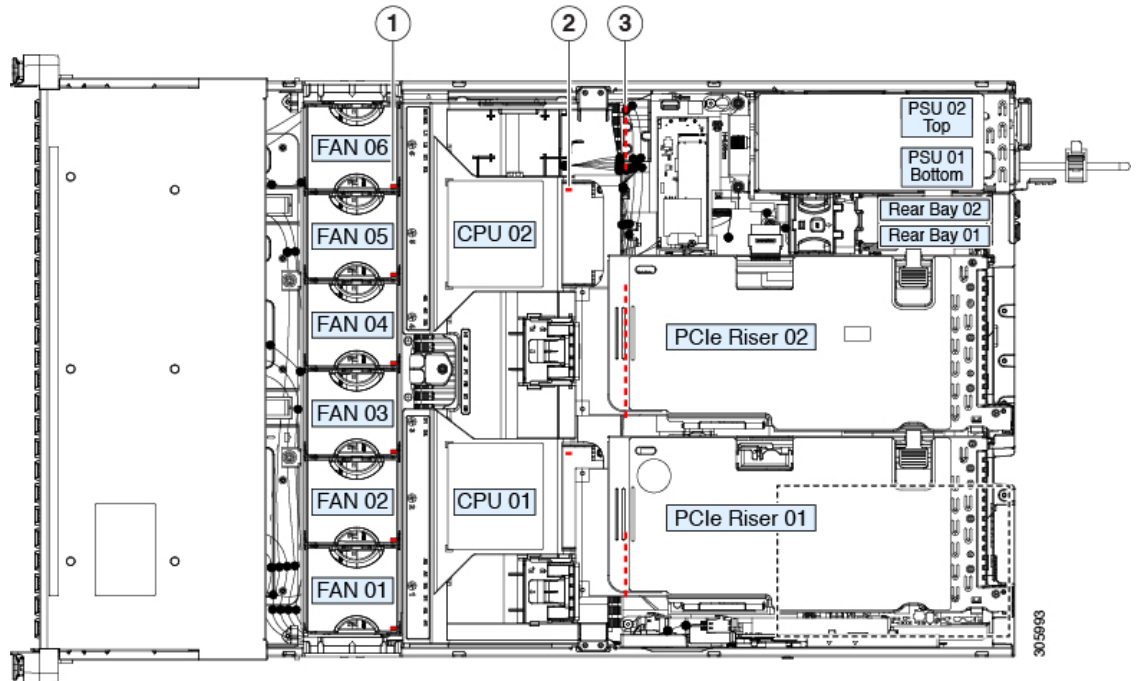
8 SAS	SAS/SATA drive activity LED	<ul style="list-style-type: none"> • Off—There is no hard drive in the hard drive tray (no access, no fault). • Green—The hard drive is ready. • Green, blinking—The hard drive is reading or writing data.
7 NVMe	NVMe SSD drive fault Note NVMe solid state drive (SSD) drive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

Internal Diagnostic LEDs

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

The internal LEDs are identical for the SFF and LFF versions of the node.

Figure 15: Internal Diagnostic LED Locations



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

Decommissioning the Node Using Cisco UCS Manager

Before replacing an internal component of a node, you must decommission the node to remove it from the Cisco UCS configuration. When you use this procedure to shut down an HX node, Cisco UCS Manager triggers the OS into a graceful shutdown sequence.

-
- Step 1** In the Navigation pane, click **Equipment**.
- Step 2** Expand **Equipment > Rack Mounts > Servers**.
- Step 3** Choose the node that you want to decommission.
- Step 4** In the **Work** pane, click the **General** tab.
- Step 5** In the **Actions** area, click **Server Maintenance**.
- Step 6** In the Maintenance dialog box, click **Decommission**, then click **OK**.
- The node is removed from the Cisco UCS configuration.
-

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 Cisco UCS Manager Equipment Tab

When you use this procedure to shut down an HX node, Cisco UCS Manager triggers the OS into a graceful shutdown sequence.



Note

If the **Shutdown Server** link is dimmed in the **Actions** area, the node is not running.

-
- Step 1** In the Navigation pane, click **Equipment**.
- Step 2** Expand **Equipment > Rack Mounts > Servers**.
- Step 3** Choose the node that you want to shut down.
- Step 4** In the Work pane, click the **General** tab.
- Step 5** In the Actions area, click **Shutdown Server**.
- Step 6** If a confirmation dialog displays, click **Yes**.

After the node has been successfully shut down, the **Overall Status** field on the **General** tab displays a power-off status.

- Step 7** 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 UCS Manager Service Profile

When you use this procedure to shut down an HX node, Cisco UCS Manager triggers the OS into a graceful shutdown sequence.



Note If the **Shutdown Server** link is dimmed in the **Actions** area, the node is not running.

Step 1 In the Navigation pane, click **Servers**.

Step 2 Expand **Servers > Service Profiles**.

Step 3 Expand the node for the organization that contains the service profile of the node that you are shutting down.

Step 4 Choose the service profile of the node that you are shutting down.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 In the **Actions** area, click **Shutdown Server**.

Step 7 If a confirmation dialog displays, click **Yes**.

After the node has been successfully shut down, the **Overall Status** field on the **General** tab displays a power-off status.

- Step 8** 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 vSphere With HX Maintenance Mode

Some procedures directly place the node into Cisco HX Maintenance mode. This procedure migrates all VMs to other nodes before the node is shut down and decommissioned from Cisco UCS Manager.

Step 1 Put the node in Cisco HX Maintenance mode by using the vSphere interface:

- Using the vSphere web client:
 - a. Log in to the vSphere web client.
 - b. Go to **Home > Hosts and Clusters**.
 - c. Expand the **Datacenter** that contains the **HX Cluster**.
 - d. Expand the **HX Cluster** and select the node.
 - e. Right-click the node and select **Cisco HX Maintenance Mode > Enter HX Maintenance Mode**.

- Using the command-line interface:
 - a. Log in to the storage controller cluster command line as a user with root privileges.
 - b. Identify the node ID and IP address:


```
# stcli node list --summary
```
 - c. Enter the node into HX Maintenance Mode:


```
# stcli node maintenanceMode (--id ID | --ip IP Address ) --mode enter
```

 (See also `stcli node maintenanceMode --help`).
 - d. Log into the ESXi command line of this node as a user with root privileges.
 - e. Verify that the node has entered HX Maintenance Mode:


```
# esxcli system maintenanceMode get
```

Step 2 Shut down the node using UCS Manager as described in [Shutting Down and Removing Power From the Node](#), on page 26.

Shutting Down Using the Power Button



Note This method is not recommended for a HyperFlex node, but the operation of the physical power button is explained here in case an emergency shutdown is required.

Step 1 Check the color of the Power Status LED:

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

Post-Maintenance Procedures

This section contains procedures that are referenced at the end of some maintenance procedures.

Recommissioning the Node Using Cisco UCS Manager

After replacing an internal component of a node, you must recommission the node to add it back into the Cisco UCS configuration.

-
- Step 1** In the Navigation pane, click **Equipment**.
 - Step 2** Expand **Equipment > Rack Mounts**.
 - Step 3** In the **Work** pane, click the **Decommissioned** tab.
 - Step 4** On the row for each node that you want to recommission, do the following:
 - a) In the **Recommission** column, check the check box.
 - b) Click **Save Changes**.
 - Step 5** If a confirmation dialog box displays, click **Yes**.
 - Step 6** (Optional) Monitor the progress of the server recommission and discovery on the **FSM** tab for the node.
-

Associating a Service Profile With an HX Node

Use this procedure to associate an HX node to its service profile after recommissioning.

-
- Step 1** In the Navigation pane, click **Servers**.
 - Step 2** Expand **Servers > Service Profiles**.
 - Step 3** Expand the node for the organization that contains the service profile that you want to associate with the HX node.
 - Step 4** Right-click the service profile that you want to associate with the HX node and then select **Associate Service Profile**.
 - Step 5** In the Associate Service Profile dialog box, select the **Server** option.
 - Step 6** Navigate through the navigation tree and select the HX node to which you are assigning the service profile.
 - Step 7** Click **OK**.
-

Exiting HX Maintenance Mode

Use this procedure to exit HX Maintenance Mode after performing a service procedure

Exit the node from Cisco HX Maintenance mode by using the vSphere interface:

- **Using the vSphere web client:**
 - a. Log in to the vSphere web client.

- b. Go to **Home > Hosts and Clusters**.
 - c. Expand the **Datacenter** that contains the **HX Cluster**.
 - d. Expand the **HX Cluster** and select the node.
 - e. Right-click the node and select **Cisco HX Maintenance Mode > Exit HX Maintenance Mode**.
- **Using the command-line interface:**
 - a. Log in to the storage controller cluster command line as a user with root privileges.
 - b. Identify the node ID and IP address:


```
# stcli node list --summary
```
 - c. Exit the node out of HX Maintenance Mode:


```
# stcli node maintenanceMode (--id ID | --ip IP Address) --mode exit
```


 (See also `stcli node maintenanceMode --help`).
 - d. Log into the ESXi command line of this node as a user with root privileges.
 - e. Verify that the node has exited HX Maintenance Mode:


```
# esxcli system maintenanceMode get
```

Removing the Node Top Cover

- Step 1** Remove the top cover:
- a) If the cover latch is locked, use a screwdriver to turn the lock 90-degrees counterclockwise to unlock it.
 - b) Lift on the end of the latch that has the green finger grip. The cover is pushed back to the open position as you lift the latch.
 - c) 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. The opening in the latch should fit over the peg that sticks up from the fan tray.
 - b) Press the cover latch down to the closed position. The cover is pushed forward to the closed position as you push down the latch.
 - c) If desired, lock the latch by using a screwdriver to turn the lock 90-degrees clockwise.

Figure 16: Removing the Top Cover



1	Cover latch	3	Serial number label location
2	Cover lock		

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.

This section describes how to remove and replace node components.


HX240c SFF Drives Serviceable Component Locations

This topic shows the locations of the field-replaceable components and service-related items.



Note The internal component *locations* are the same for the SFF and LFF versions of the node unless otherwise noted. The functions of the drives differ in SFF and LFF versions, as described in this document.

Figure 17: SFF Serviceable Component Locations



1	2.5-inch drive bays 2 – 24: <ul style="list-style-type: none"> Hybrid HX240C-M5SX: 6 to 23 persistent data HDDs All-Flash HXAF240C-M5SX: 6 to 23 persistent data SSDs 	10	Rear drive bay 25: Caching SSD
2	Drive bay 1: Housekeeping SSD for logs	11	PCIe riser 2B (PCIe slots 4, 5, 6)
3	Cooling fan modules (six, hot-swappable)	12	Micro-SD card socket on PCIe riser 1
4	DIMM sockets on motherboard (up to 12 per CPU) Not visible under air baffle in this view.	13	PCIe riser 1 (PCIe slot 1, 2, 3)
5	CPUs and heatsinks (two) Not visible under air baffle in this view.	14	Modular LOM (mLOM) card bay on chassis floor (x16 PCIe lane), not visible in this view
6	Internal, vertical USB 3.0 port on motherboard	15	Cisco modular storage controller HBA slot (dedicated PCIe slot)
7	Mini-storage module for SATA M.2 SSD boot drive	16	RTC battery, vertical socket on motherboard
8	PCIe cable connector on PCIe riser 2B for rear-loading NVMe SSD	17	Securing clips for GPU cards on air baffle
9	Power supplies (hot-swappable when redundant as 1+1)	-	


HX240c LFF Drives Serviceable Component Locations

This topic shows the locations of the field-replaceable components and service-related items for Hybrid HX240C-M5L.



Note The internal component *locations* are the same for the SFF and LFF versions of the node unless otherwise noted. The functions of the drives differ in SFF and LFF versions, as described in this document.

Figure 18: LFF Serviceable Component Locations



1	3.5-inch drive bays 1 – 12: • Hybrid HX240C-M5L: 6 to 12 persistent data HDDs	10	Rear drive bay 13: Caching SSD
2	Cooling fan modules (six, hot-swappable)	11	PCIe riser 2B (PCIe slots 4, 5, 6)
3	DIMM sockets on motherboard (up to 12 per CPU) Not visible under air baffle in this view.	12	Micro-SD card socket on PCIe riser 1
4	CPUs and heatsinks (two) Not visible under air baffle in this view.	13	PCIe riser 1 (PCIe slot 1, 2, 3)
5	Internal, vertical USB 3.0 port on motherboard	14	Modular LOM (mLOM) card bay on chassis floor (x16 PCIe lane), not visible in this view
6	Mini-storage module for SATA M.2 SSD boot drive	15	Cisco modular storage controller HBA slot (dedicated PCIe slot)
7	PCIe cable connector on PCIe riser 2B for rear-loading NVMe SSD	16	RTC battery, vertical socket on motherboard
8	Power supplies (hot-swappable when redundant as 1+1)	17	Securing clips for GPU cards on air baffle
9	Rear drive bay 14: Housekeeping SSD for logs	-	

Considerations For Upgrading Hardware in Multiple Nodes of a Cluster

This chapter contains removal and replacement procedures for components that are supported as field-replaceable. This topic describes additional considerations when multiple nodes in an existing cluster are upgraded with the addition or replacement of components.

The following procedure describes the general steps and considerations for upgrading hardware in the nodes of a cluster.



Note Hot-swappable components can be replaced or added without shutting down the system as described below. Those include certain drives, the internal fan modules, and the power supplies. Check the procedure for the component in this chapter to verify whether the shutdown steps are required.

- Step 1** Verify that the existing cluster is healthy.
- Step 2** Put the node in Cisco HX Maintenance Mode as described in [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#).
- Step 3** Shut down the node as described in [Shutting Down and Removing Power From the Node, on page 26](#).
- Step 4** Decommission the node from UCS Manager as described in [Decommissioning the Node Using Cisco UCS Manager, on page 25](#).
- 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.
- Step 5** Disconnect all power cables from all power supplies.
- Step 6** Remove and replace the existing component or add a new component following the supported population rules. Use the specific procedure in this chapter for the component.
- Step 7** Recommission the node in UCS Manager as described in [Recommissioning the Node Using Cisco UCS Manager, on page 29](#).
- Ensure that ESXi boots. The node is auto-discovered by Cisco UCS Manager and the ESXi operating system recognizes the new components.
- Step 8** Associate the node with its UCS Manager service profile as described in [Associating a Service Profile With an HX Node, on page 29](#).
- Step 9** Verify that ESXi is reconnected to HyperFlex vCenter.
- Step 10** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, on page 29](#).
- Step 11** Verify within 30 minutes that the HX cluster is again in a healthy state.
- Step 12** Move a test VM to the host. Ensure that it is working by performing tests.
- Note** You must perform the hardware upgrade on the next node in the cluster within two hours after upgrading the prior node and verifying that the cluster is healthy. The HyperFlex Data Platform (HXDP) runs a data rebalance operation every two hours and this timer starts after the cluster is again in a healthy state.
- Step 13** Repeat the steps above to upgrade the hardware component in each node until all node hardware is updated.
-

Replacing Drives

This section describes the supported drives by function and how to replace them.

Drive Population Rules

Drive bay numbering is shown in the following figure.

Figure 19: Front Drive Bay Numbering, SFF Versions





Figure 20: Front Drive Bay Numbering, LFF Versions



The numbering of the rear drive bays is sequential, following the number of front drive bays.

Figure 21: Rear Drive Bay Numbering, All Versions



Observe these drive population rules:

SFF Versions Drive Population Rules

- Front bay 1: Housekeeping SSD for SDS logs only
- Front bays 2 - 24:
 - HX240c Hybrid: 6 to 23 persistent data HDDs
 - HX240c All-Flash: 6 to 23 persistent data SSDs
- Rear bay 25: Caching SSD only
- When populating persistent data drives, add drives to the lowest-numbered bays first.
- Keep an empty drive blanking tray in any unused bays to ensure proper airflow.

- See [HX240c M5 Drive Configuration Comparison, on page 37](#) for supported drive configurations.

LFF Versions Drive Population Rules

- Front bays 1 - 12:
 - Hybrid HX240C-M5L: 6 to 12 persistent data HDDs
- Rear bay 13: Caching SSD only
- Rear bay 14: Housekeeping SSD for SDS logs only
- When populating persistent data drives, add drives to the lowest-numbered bays first.
- Keep an empty drive blanking tray in any unused bays to ensure proper airflow.
- See [HX240c M5 Drive Configuration Comparison, on page 37](#) for supported drive configurations.



Note Regarding drive capacity stated on drive labels vs reported capacity in software:

The capacity stated on the drive physical label and the capacity reported by the HyperFlex software differ because of the following reasons:

1. Drive label capacities are stated in decimal (base 10) notation, while the software-reported capacities are stated in binary (base 2) notation. For example, 1 TB in decimal notation would be reported as 909 GB in binary notation—these are actually the *same capacity*, similar to distance reported in miles vs kilometers being the same distance, but in different units. These examples show capacities stated as decimal notation vs binary notation:
 - 500 GB (decimal) = 465.5 GB (binary)
 - 1 TB (decimal) = 909 GB (binary)
 - 2 TB (decimal) = 1.82 TB (binary)
 - 3 TB (decimal) = 2.72 TB (binary)
2. Preinstalled software and partitions also reduce storage capacity.

HX240c M5 Drive Configuration Comparison

The following tables compare the drives supported by each function in the node. Also note the considerations listed after the tables.

Table 4: SFF Versions Drive Configuration Comparison

Component	HX240c M5 Hybrid SFF	HX240c M5 All-Flash SFF	HX240c M5 SED Hybrid SFF	HX240c M5 SED All-Flash SFF
Housekeeping SSD Slot 1	SSD: • HX-SD240G61X-EV	SSD: • HX-SD240G61X-EV	SSD: • HX-SD240G61X-EV	SSD: • HX-SD240G61X-EV

Component	HX240c M5 Hybrid SFF	HX240c M5 All-Flash SFF	HX240c M5 SED Hybrid SFF	HX240c M5 SED All-Flash SFF
Persistent data drives Slots 2 - 24	HDD: <ul style="list-style-type: none"> HX-HD12TB10K12N 	SSD: <ul style="list-style-type: none"> HX-SD960G61X-EV HX-SD38T61X-EV 	SED HDD: <ul style="list-style-type: none"> HX-HD12T10NK9 	SED SSD: <ul style="list-style-type: none"> HX-SD800GBENK9 HX-SD960GBE1NK9 HX-SD38TBE1NK9
Caching SSD Rear slot 25	SSD: <ul style="list-style-type: none"> HX-SD16T63X-EP HX-SD16T123X-EP 	SSD: <ul style="list-style-type: none"> HX-NVMEHW-H1600 HX-SD400G12TX-EP HX-NVMEXP-B375 	SED SSD: <ul style="list-style-type: none"> HX-SD16TBENK9 	SED SSD: <ul style="list-style-type: none"> HX-SD800GBENK9

Table 5: LFF Versions Drive Configuration Comparison

Component	HX240c M5 Hybrid LFF
Persistent data drives Slots 1 - 12	HDD (3.5-inch): <ul style="list-style-type: none"> HX-HD6T7KL4KN HX-HD8T7KL4KN
Caching SSD Rear slot 13	SSD: <ul style="list-style-type: none"> HX-SD32T123X-EP
Housekeeping SSD Rear bay 14	SSD: <ul style="list-style-type: none"> HX-SD240G61X-EV

Note the following considerations and restrictions for All-Flash HyperFlex nodes:

- The minimum Cisco HyperFlex software required for using Intel Optane NVMe SSD HX-NVMEXP-I375 is Release 3.0(1a) or later. If you use HX-NVMEXP-I375 as your caching drive in HX240c All-Flash nodes, all nodes in the cluster must use this same drive as the caching drive.
- HX240c All-Flash HyperFlex nodes are ordered as specific All-Flash PIDs; All-Flash configurations are supported only on those PIDs.
- Conversion from Hybrid HX240c configuration to All-Flash HX240c configuration is not supported.
- Mixing Hybrid nodes with All-Flash nodes within the same HyperFlex cluster is not supported.

Note the following considerations and restrictions for SED HyperFlex nodes:

- The minimum Cisco HyperFlex software required for SED configurations is Release 3.5(1a) or later.
- Mixing HX240c *Hybrid* SED HyperFlex nodes with HX240c *All-Flash* SED HyperFlex nodes within the same HyperFlex cluster is not supported.

Drive Replacement Overview

The three types of drives in the node require different replacement procedures.

Housekeeping SSD SFF: Front bay 1 LFF: Rear bay 14	The node must be put into Cisco HX Maintenance Mode before replacing the housekeeping SSD. See Replacing the Housekeeping SSD, on page 39 . Note Replacement requires additional technical assistance and cannot be completed by the customer.
Persistent data drives SFF: Front bays 2 - 24 LFF: Front bays 1 - 12	Hot-swap replacement is supported. See Replacing Persistent Data Drives, on page 42 . Note Hot-swap replacement includes hot-removal, so you can remove the drive while it is still operating.
Caching SSD SFF: Rear bay 25 LFF: Rear bay 13	Hot-swap replacement is supported. See #unique_45 . Note Hot-swap replacement for SAS/SATA drives includes hot-removal, so you can remove the drive while it is still operating. Note If an NVMe SSD is used as the Caching SSD, additional steps are required as described in the procedure.

Replacing the Housekeeping SSD

- SFF versions: The Housekeeping SSD must be installed in front bay 1.
- LFF versions: The Housekeeping SSD must be installed in rear bay 14.



Note This procedure requires assistance from technical support for additional software update steps after the hardware is replaced. It cannot be completed without technical support assistance.




Note Always replace the drive with the same type and size as the original drive.



Caution Put the node in Cisco HX Maintenance mode before replacing the Housekeeping SSD, as described in the procedure. Hot-swapping the Housekeeping SSD while the node is running causes the node to fail.

-
- Step 1** Put the node in Cisco HX Maintenance Mode as described in [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#).
- Step 2** Shut down the node as described in [Shutting Down and Removing Power From the Node, on page 26](#).
- Step 3** Decommission the node from UCS Manager as described in [Decommissioning the Node Using Cisco UCS Manager, on page 25](#).
- 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.
- Step 4** Disconnect all power cables from all power supplies.
- Step 5** Remove the Housekeeping SSD:
- SFF versions: The Housekeeping SSD must be installed in front bay 1.
 - LFF versions: The Housekeeping SSD must be installed in rear bay 14.
- 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) Remove the four drive-tray screws that secure the drive to the tray and then lift the drive out of the tray.
- Step 6** Install a new drive:
- a) Place a new drive 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.

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

- Step 7** Replace power cables and then power on the node by pressing the Power button.
- Step 8** Recommission the node in UCS Manager as described in [Recommissioning the Node Using Cisco UCS Manager, on page 29](#).
- Step 9** Associate the node with its UCS Manager service profile as described in [Associating a Service Profile With an HX Node, on page 29](#).
- Step 10** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, on page 29](#).

Note After you replace the Housekeeping SSD, you must contact technical support for additional software update steps.

Replacing Persistent Data Drives

- SFF versions: The persistent data drives must be installed in front bays 2 - 24.
- LFF versions: The persistent data drives must be installed in front bays 1 - 12.



Note Always replace the drive with the same type and size as the original drive.



Note Hot-swap replacement includes hot-removal, so you can remove the drive while it is still operating.


Step 1 Remove the drive that you are replacing or remove a blank drive tray from the bay:

- SFF versions: The persistent data drives must be installed in front bays 2 - 24.
 - LFF versions: The persistent data drives must be installed in front bays 1 - 12.
- 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.
- 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.

Figure 23: 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 the Caching SSD

- SFF versions: The Caching SSD must be installed in rear bay 25.
- LFF versions: The Caching SSD must be installed in rear bay 13.

Note the following considerations for NVMe SSDs, when used as the Caching SSD:

- NVMe SSDs are supported only in All-Flash nodes. NVMe SSDs are not supported in Hybrid nodes.
- NVMe SSDs are supported only in the Caching SSD position, in drive bay 13 for LFF versions or bay 25 for SFF versions. NVMe SSDs are not supported for persistent storage or as the Housekeeping drive.
- The locator (beacon) LED cannot be turned on or off on NVMe SSDs.



Note Always replace the drive with the same type and size as the original drive.



Note Upgrading or downgrading the Caching drive in an existing HyperFlex cluster is not supported. If the Caching drive must be upgraded or downgraded, then a full redeployment of the HyperFlex cluster is required.



Note When using a SAS/SATA drive, hot-swap replacement includes hot-removal, so you can remove the drive while it is still operating. NVMe drives cannot be hot-swapped.

Step 1 **Only if the caching drive is an NVMe SSD**, enter the ESXi host into HX Maintenance Mode (see [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#)). Otherwise, skip to step 2.

Step 2 Remove the Caching SSD:

- SFF versions: The Caching SSD must be installed in rear bay 25.
- LFF versions: The Caching SSD must be installed in rear bay 13.

- 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) Remove the four drive-tray screws that secure the drive to the tray and then lift the drive out of the tray.


Step 3 Install a new drive:

- a) Place a new drive 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 4 **Only if the Caching SSD is an NVMe SSD:**

- a) Reboot the ESXi host. This enables ESXi to discover the NVMe SSD.
- b) Exit the ESXi host from HX Maintenance Mode (see [Exiting HX Maintenance Mode, on page 29](#)).

Figure 24: 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 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 server to replace fan modules because they are hot-swappable. However, to maintain proper cooling, do not operate the server for more than one minute with any fan module removed.

Step 1 Remove an existing fan module:

- a) 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.

- b) Remove the top cover from the server as described in [Removing the Node Top Cover, on page 30](#).
- 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 server.
- 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 server.
- d) Replace the server in the rack, replace cables, and then fully power on the server by pressing the Power button.

Figure 25: Top View of Fan Module



1	Fan module release latches	2	Fan module fault LED
----------	----------------------------	----------	----------------------

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 server might result in system problems or damage to the motherboard.



Note To ensure the best server 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

This topic describes the rules and guidelines for maximum memory performance.

DIMM Slot Numbering

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

Figure 26: DIMM Slot Numbering



DIMM Population Rules

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

- Each CPU supports six memory channels.
 - CPU 1 supports channels A, B, C, D, E, F.
 - CPU 2 supports channels G, H, J, K, L, M.
- Each channel has two DIMM sockets (for example, channel A = slots A1, A2).
- In a single-CPU configuration, populate the channels for CPU1 only (A, B, C, D, E, F).

- 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 server has two CPUs, balance DIMMs evenly across the two CPUs as shown in the table.



Note The table below lists recommended configurations. Using 5, 7, 9, 10, or 11 DIMMs per CPU is not recommended.

Table 6: DIMM Population Order

Number of DIMMs per CPU (Recommended Configurations)	Populate CPU 1 Slot		Populate CPU2 Slots	
	Blue #1 Slots	Black #2 Slots	Blue #1 Slots	Black #2 Slots
1	(A1)	-	(G1)	-
2	(A1, B1)	-	(G1, H1)	-
3	(A1, B1, C1)	-	(G1, H1, J1)	-
4	(A1, B1); (D1, E1)	-	(G1, H1); (K1, L1)	-
6	(A1, B1); (C1, D1); (E1, F1)	-	(G1, H1); (J1, K1); (L1, M1)	-
8	(A1, B1); (D1, E1)	(A2, B2); (D2, E2)	(G1, H1); (K1, L1)	(G2, H2); (K2, L2)
12	(A1, B1); (C1, D1); (E1, F1)	(A2, B2); (C2, D2); (E2, F2)	(G1, H1); (J1, K1); (L1, M1)	(G2, H2); (J2, K2); (L2, M2)

- The maximum combined memory allowed in the 12 DIMM slots controlled by any one CPU is 768 GB. To populate the 12 DIMM slots with more than 768 GB of combined memory, you must use a high-memory CPU that has a PID that ends with an "M", for example, UCS-CPU-6134M.
- The NVIDIA Tesla P-Series GPU can support more than 1 TB of memory in the server. All other NVIDIA GPUs can support only 1 TB or less of memory in the server. Therefore, do not install DIMMs totalling more than 1 TB capacity when using an NVIDIA GPU card other than P-Series in this server.
- Observe the DIMM mixing rules shown in the following table.

Table 7: DIMM Mixing Rules

DIMM Parameter	DIMMs in the Same Channel	DIMMs in the Same Bank
DIMM Capacity RDIMM = 16GB or 32GB LRDIMM = 64GB TSV-RDIMM = 128GB	You can mix different capacity DIMMs in the same channel (for example, A1, A2).	You cannot mix DIMM capacities in a bank (for example A1, B1). Pairs of DIMMs must be identical (same PID and revision).

DIMM speed For example, 2666 GHz	You can mix speeds, but DIMMs will run at the speed of the slowest DIMMs/CPUs installed in the channel.	You cannot mix DIMM speeds in a bank (for example A1, B1). Pairs of DIMMs must be identical (same PID and revision).
DIMM type RDIMMs or LRDIMMs	You cannot mix DIMM types in a channel.	You cannot mix DIMM types in a bank.

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 24](#) for the locations of these LEDs. When the server is in standby power mode, these LEDs light amber to indicate a faulty DIMM.

-
- Step 1** Put the node in Cisco HX Maintenance Mode as described in [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#).
- Step 2** Shut down the node as described in [Shutting Down and Removing Power From the Node, on page 26](#).
- Step 3** Decommission the node from UCS Manager as described in [Decommissioning the Node Using Cisco UCS Manager, on page 25](#).
- 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.
- Step 4** Disconnect all power cables from all power supplies.
- Step 5** Remove an existing DIMM:
- 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.
 - Remove the top cover from the server as described in [Removing the Node Top Cover, on page 30](#).
 - Remove the air baffle that covers the front ends of the DIMM slots to provide clearance.
 - Locate the DIMM that you are removing, and then open the ejector levers at each end of its DIMM slot.
- Step 6** Install a new DIMM:
- Note** Before installing DIMMs, see the memory population rules for this server: [DIMM Population Rules and Memory Performance Guidelines , on page 47](#).
- Align the new DIMM with the empty slot on the motherboard. Use the alignment feature in the DIMM slot to correctly orient the DIMM.
 - 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.
 - Replace the top cover to the server.
 - Replace the server in the rack, replace cables, and then fully power on the server by pressing the Power button.
- Step 7** Recommission the node in UCS Manager as described in [Recommissioning the Node Using Cisco UCS Manager, on page 29](#).

- Step 8** Associate the node with its UCS Manager service profile as described in [Associating a Service Profile With an HX Node, on page 29](#).
- Step 9** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, on page 29](#).

Replacing CPUs and Heatsinks

This section contains CPU configuration rules and the procedure for replacing CPUs and heatsinks:

Special Information For *Upgrades to Second Generation Intel Xeon Scalable Processors*



Caution You must upgrade your node firmware and software to the required minimum levels before you upgrade to the Second Generation Intel Xeon Scalable Processors that are supported in this node. Older firmware versions cannot recognize the new CPUs and this would result in a non-bootable node.



Note You can use First Generation and Second Generation Intel Xeon Scalable processors in the same cluster. Do not mix First Generation and Second Generation processors within the same node.

The minimum software and firmware versions required for this node to support Second Generation Intel Xeon Scalable Processors are as follows:

Table 8: Minimum Requirements For Second Generation Intel Xeon Scalable Processors

Software or Firmware	Minimum Version
Node Cisco IMC/BIOS	4.0(4d)
Cisco UCS Manager	4.0(4d)
Cisco HyperFlex Data Platform	4.0(1b)

Do one of the following actions:

- If your server's firmware and Cisco UCS Manager software are already at the required minimums shown above (or later), you can replace the CPU hardware by using the procedure in this section.
- If your server's firmware and Cisco UCS Manager software are earlier than the required levels, upgrade your software. After you upgrade the software, return to this section as directed to replace the CPU hardware.

CPU Configuration Rules

This node has two CPU sockets on the motherboard. Each CPU supports six DIMM channels (12 DIMM slots).

- The minimum configuration is that the node must have two identical CPUs installed.

- The maximum combined memory allowed in the 12 DIMM slots controlled by any one CPU is 768 GB. To populate the 12 DIMM slots with more than 768 GB of combined memory, you must use a high-memory CPU that has a PID that ends with an "M", for example, UCS-CPU-6134M.

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 UCS-CPUAT=.
- Heatsink cleaning kit—Supplied with replacement CPU. Orderable separately as Cisco PID UCSX-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 UCS-CPU-TIM=.

New heatsinks have a pre-applied pad of TIM.

Replacing a CPU and Heatsink



Caution CPUs and their sockets are fragile and must be handled with extreme care to avoid damaging pins. The CPUs must be installed with heatsinks and thermal interface material to ensure cooling. Failure to install a CPU correctly might result in damage to the server.

Step 1 Put the node in Cisco HX Maintenance Mode as described in [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#).

Step 2 Shut down the node as described in [Shutting Down and Removing Power From the Node, on page 26](#).

Step 3 Decommission the node from UCS Manager as described in [Decommissioning the Node Using Cisco UCS Manager, on page 25](#).

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.

Step 4 Disconnect all power cables from all power supplies.

Step 5 Remove the existing CPU/heatsink assembly from the node:

- 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 30](#).

- Use the T-30 Torx driver that is supplied with the replacement CPU to loosen the four captive nuts that secure the assembly to the motherboard standoffs.

Note Alternate loosening the heatsink nuts evenly so that the heatsink remains level as it is raised. Loosen the heatsink nuts in the order shown on the heatsink label: 4, 3, 2, 1.

- d) Lift straight up on the CPU/heatsink assembly and set it heatsink-down on an antistatic surface.

Figure 27: Removing the CPU/Heatsink Assembly



1	Heatsink	4	CPU socket on motherboard
2	Heatsink captive nuts (two on each side)	5	T-30 Torx driver
3	CPU carrier (below heatsink in this view)	-	


Step 6

Separate the heatsink from the CPU assembly (the CPU assembly includes the CPU and the CPU carrier):

- a) Place the heatsink with CPU assembly so that it is oriented upside-down as shown below.

Note the thermal-interface material (TIM) breaker location. TIM BREAKER is stamped on the CPU carrier next to a small slot.

Figure 28: Separating the CPU Assembly From the Heatsink



1	CPU carrier	4	CPU-carrier inner-latch nearest to the TIM breaker slot
2	CPU	5	#1 flat-head screwdriver inserted into TIM breaker slot
3	TIM BREAKER slot in CPU carrier	-	

- b) Pinch inward on the CPU-carrier inner-latch that is nearest the TIM breaker slot and then push up to disengage the clip from its slot in the heatsink corner.
- c) Insert the blade of a #1 flat-head screwdriver into the slot marked TIM BREAKER.

Caution In the following step, do not pry on the CPU surface. Use gentle rotation to lift on the plastic surface of the CPU carrier at the TIM breaker slot. Use caution to avoid damaging the heatsink surface.

- d) Gently rotate the screwdriver to lift up on the CPU until the TIM on the heatsink separates from the CPU.

Note Do not allow the screwdriver tip to touch or damage the green CPU substrate.

- e) Pinch the CPU-carrier inner-latch at the corner opposite the TIM breaker and push up to disengage the clip from its slot in the heatsink corner.
- f) On the remaining two corners of the CPU carrier, gently pry outward on the outer-latches and then lift the CPU-assembly from the heatsink.


Note Handle the CPU-assembly by the plastic carrier only. Do not touch the CPU surface. Do not separate the CPU from the carrier.

Step 7 The new CPU assembly is shipped on a CPU assembly tool. Take the new CPU assembly and CPU assembly tool out of the carton.

If the CPU assembly and CPU assembly tool become separated, note the alignment features shown below for correct orientation. The pin 1 triangle on the CPU carrier must be aligned with the angled corner on the CPU assembly tool.

Caution CPUs and their sockets are fragile and must be handled with extreme care to avoid damaging pins.

Figure 29: CPU Assembly Tool, CPU Assembly, and Heatsink Alignment Features




1	CPU assembly tool	4	Angled corner on heatsink (pin 1 alignment feature)
2	CPU assembly (CPU in plastic carrier)	5	Triangle cut into carrier (pin 1 alignment feature)
3	Heatsink	6	Angled corner on CPU assembly tool (pin 1 alignment feature)

Step 8 Apply new TIM to the heatsink:

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 5.
- 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 cleaning solution that is included with the heatsink cleaning kit (UCSX-HSCK=) 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) Using the syringe of TIM provided with the new CPU (UCS-CPU-TIM=), apply 1.5 cubic centimeters (1.5ml) of thermal interface material to the top of the CPU. Use the pattern shown below to ensure even coverage.

Figure 30: Thermal Interface Material Application Pattern



Caution Use only the correct heatsink for your CPUs to ensure proper cooling. There are two different heatsinks: UCSC-HS-C240M5= for standard-performance CPUs 150 W and less; UCSC-HS2-C240M5= for high-performance CPUs above 150 W. Note the wattage described on the heatsink label.

Step 9

With the CPU assembly on the CPU assembly tool, set the heatsink onto the CPU assembly. Note the pin 1 alignment features for correct orientation. Push down gently until you hear the corner clips of the CPU carrier click onto the heatsink corners.

Caution In the following step, use extreme care to avoid touching or damaging the CPU contacts or the CPU socket pins.


Step 10

Install the CPU/heatsink assembly to the server:

- a) Lift the heatsink with attached CPU assembly from the CPU assembly tool.
- b) Align the CPU with heatsink over the CPU socket on the motherboard, as shown below.

Note the alignment features. The pin 1 angled corner on the heatsink must align with the pin 1 angled corner on the CPU socket. The CPU-socket posts must align with the guide-holes in the assembly.

Figure 31: Installing the Heatsink/CPU Assembly to the CPU Socket



1	Guide hole in assembly (two)	4	Angled corner on heatsink (pin 1 alignment feature)
2	CPU socket alignment post (two)	5	Angled corner on socket (pin 1 alignment feature)
3	CPU socket leaf spring	-	

- c) Set the heatsink with CPU assembly down onto the CPU socket.
- d) Use the T-30 Torx driver that is supplied with the replacement CPU to tighten the four captive nuts that secure the heatsink to the motherboard standoffs.

Caution Alternate tightening the heatsink nuts evenly so that the heatsink remains level while it is lowered. Tighten the heatsink nuts in the order shown on the heatsink label: 1, 2, 3, 4. The captive nuts must be fully tightened so that the leaf springs on the CPU socket lie flat.

- 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.

Step 11

Recommission the node in UCS Manager as described in [Recommissioning the Node Using Cisco UCS Manager, on page 29](#).

- Step 12** Associate the node with its UCS Manager service profile as described in [Associating a Service Profile With an HX Node, on page 29](#).
- Step 13** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, on page 29](#).

Additional CPU-Related Parts to Order with RMA Replacement CPUs

When a return material authorization (RMA) of the CPU is done on a Cisco UCS C-Series server, 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 58](#).

- Scenario 1—You are reusing the existing heatsinks:
 - Heatsink cleaning kit (UCSX-HSCK=)
 - One cleaning kit can clean up to four CPUs.
 - Thermal interface material (TIM) kit for M5 servers (UCS-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: UCSC-HS-C240M5= for CPUs 150 W and less; UCSC-HS2-C240M5= for CPUs above 150 W.

- Heatsink: UCSC-HS-C240M5= for CPUs 150 W and less; UCSC-HS2-C240M5= for CPUs above 150 W
 - New heatsinks have a pre-applied pad of TIM.
- Heatsink cleaning kit (UCSX-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 (UCS-M5-CPU-CAR=)
 - #1 flat-head screwdriver (for separating the CPU from the heatsink)
 - Heatsink cleaning kit (UCSX-HSCK=)
 - One cleaning kit can clean up to four CPUs.
 - Thermal interface material (TIM) kit for M5 servers (UCS-CPU-TIM=)
 - One TIM kit covers one CPU.

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

New heatsink 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 heatsinks. Therefore, even when you are ordering new heatsinks, you must order the heatsink 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 UCS C-Series server, you move existing CPUs to the new chassis.



Note Unlike previous generation CPUs, the M5 server 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.

Moving an M5 Generation CPU

Tool required for this procedure: T-30 Torx driver



Caution When you receive a replacement server for an RMA, it includes dust covers on all CPU sockets. These covers protect the socket pins from damage during shipping. You must transfer these covers to the system that you are returning, as described in this procedure.


Step 1 When moving an M5 CPU to a new server, you do not have to separate the heatsink from the CPU. Perform the following steps:

- Use a T-30 Torx driver to loosen the four captive nuts that secure the assembly to the board standoffs.

Note Alternate loosening the heatsink nuts evenly so that the heatsink remains level as it is raised. Loosen the heatsink nuts in the order shown on the heatsink label: 4, 3, 2, 1.

- Lift straight up on the CPU/heatsink assembly to remove it from the board.
- Set the CPUs with heatsinks aside on an anti-static surface.

Figure 32: Removing the CPU/Heatsink Assembly



1	Heatsink	4	CPU socket on motherboard
2	Heatsink captive nuts (two on each side)	5	T-30 Torx driver
3	CPU carrier (below heatsink in this view)	-	


Step 2

Transfer the CPU socket covers from the new system to the system that you are returning:

- a) Remove the socket covers from the replacement system. Grasp the two recessed finger-grip areas marked "REMOVE" and lift straight up.

Note Keep a firm grasp on the finger-grip areas at both ends of the cover. Do not make contact with the CPU socket pins.

Figure 33: Removing a CPU Socket Dust Cover



1	Finger-grip areas marked "REMOVE"	-	
---	-----------------------------------	---	--


- b) With the wording on the dust cover facing up, set it in place over the CPU socket. Make sure that all alignment posts on the socket plate align with the cutouts on the cover.

Caution In the next step, do not press down anywhere on the cover except the two points described. Pressing elsewhere might damage the socket pins.

- c) Press down on the two circular markings next to the word "INSTALL" that are closest to the two threaded posts (see the following figure). Press until you feel and hear a click.

Note You must press until you feel and hear a click to ensure that the dust covers do not come loose during shipping.

Figure 34: Installing a CPU Socket Dust Cover



-	Press down on the two circular marks next to the word INSTALL.	-	
---	--	---	--


Step 3

Install the CPUs to the new system:

- a) On the new board, align the assembly over the CPU socket, as shown below.

Note the alignment features. The pin 1 angled corner on the heatsink must align with the pin 1 angled corner on the CPU socket. The CPU-socket posts must align with the guide-holes in the assembly.

Figure 35: Installing the Heatsink/CPU Assembly to the CPU Socket



1	Guide hole in assembly (two)	4	Angled corner on heatsink (pin 1 alignment feature)
2	CPU socket alignment post (two)	5	Angled corner on socket (pin 1 alignment feature)
3	CPU socket leaf spring	-	

- b) On the new board, set the heatsink with CPU assembly down onto the CPU socket.
- c) Use a T-30 Torx driver to tighten the four captive nuts that secure the heatsink to the board standoffs.

Note Alternate tightening the heatsink nuts evenly so that the heatsink remains level while it is lowered. Tighten the heatsink nuts in the order shown on the heatsink label: 1, 2, 3, 4. The captive nuts must be fully tightened so that the leaf springs on the CPU socket lie flat.

Replacing a Mini-Storage Module or M.2 Boot Drive


The mini-storage module plugs into a motherboard socket to provide additional M.2 SSD internal storage. This node includes a SATA M.2 SSD that can be used as a boot drive.

Replacing a Mini-Storage Module Carrier

This topic describes how to remove and replace a mini-storage module carrier.

-
- Step 1** Put the node in Cisco HX Maintenance Mode as described in [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#).
- Step 2** Shut down the node as described in [Shutting Down and Removing Power From the Node, on page 26](#).
- Step 3** Decommission the node from UCS Manager as described in [Decommissioning the Node Using Cisco UCS Manager, on page 25](#).
- 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.
- Step 4** Disconnect all power cables from all power supplies.
- Step 5** 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 6** Remove the top cover from the server as described in [Removing the Node Top Cover, on page 30](#).
- Step 7** Remove a carrier from its socket:
- Locate the mini-storage module carrier in its socket just in front of power supply 1.
 - At each end of the carrier, push outward on the clip that secures 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 8** Install a carrier to its socket:
- Position the carrier over socket, with the carrier's connector facing down and at the same end as the motherboard socket. 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 9** Replace the top cover to the server.
- Step 10** Replace the server in the rack, replace cables, and then fully power on the server by pressing the Power button.

Figure 36: Mini-Storage Module Carrier



1	Location of socket on motherboard	3	Securing clips
2	Alignment pegs	-	

- Step 11** Recommission the node in UCS Manager as described in [Recommissioning the Node Using Cisco UCS Manager, on page 29](#).
- Step 12** Associate the node with its UCS Manager service profile as described in [Associating a Service Profile With an HX Node, on page 29](#).
- Step 13** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, on page 29](#).

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

This topic describes how to remove and replace an M.2 SSD in a mini-storage carrier for M.2 (UCS-MSTOR-M2).

- Step 1** Put the node in HX Maintenance Mode, shut down the node, decommission it, and then remove the mini-storage module carrier from the server as described in [Replacing a Mini-Storage Module Carrier, on page 63](#).
- 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) Angle the M.2 SSD downward and insert the connector-end into the socket on the carrier. The M.2 SSD's label must face 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 recommission the node, reassociate its profile, and exit HX Maintenance Mode as described in [Replacing a Mini-Storage Module Carrier, on page 63](#).

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]

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 ordered from Cisco (PID N20-MBLIBATT) or purchased from most electronic stores.

Step 1 Put the node in Cisco HX Maintenance Mode as described in [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#).

Step 2 Shut down the node as described in [Shutting Down and Removing Power From the Node, on page 26](#).

Step 3 Decommission the node from UCS Manager as described in [Decommissioning the Node Using Cisco UCS Manager, on page 25](#).

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.

Step 4 Disconnect all power cables from all power supplies.

Step 5 Remove the RTC battery:

- 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 30](#).
- c) Remove PCIe riser 1 from the node to provide clearance to the RTC battery socket that is on the motherboard.
- d) Locate the horizontal RTC battery socket.
- e) 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 6 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.
- c) Replace the top cover to the node.

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

Figure 37: RTC Battery Location on Motherboard



1	RTC battery in horizontal socket on motherboard	-	
----------	---	---	--

- Step 7** Reconfigure the node in UCS Manager as described in [Reconfiguring the Node Using Cisco UCS Manager](#), on page 29.
- Step 8** Associate the node with its UCS Manager service profile as described in [Associating a Service Profile With an HX Node](#), on page 29.
- Step 9** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode](#), on page 29.

Replacing Power Supplies

The node can have one or two power supplies. When two power supplies are installed they are redundant as 1+1.

Replacing AC Power Supplies



Note If two power supplies are present, you do not have to power off the node to replace one 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 power supply that you are replacing:

a) Perform one of the following actions:

- If your node has two power supplies, you do not have to shut down the server. Continue with step 2.
- If your node has only one power supply:
 1. Put the node in Cisco HX Maintenance Mode as described in [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#).
 2. Shut down the node as described in [Shutting Down and Removing Power From the Node, on page 26](#).
 3. Decommission the node from UCS Manager as described in [Decommissioning the Node Using Cisco UCS Manager, on page 25](#).


Step 2 Remove the power supply that you are replacing or a blank panel from an empty bay:

- a) Remove the power cord from the power supply that you are replacing.
- b) Grasp the power supply handle while pinching the release lever toward the handle.
- c) Pull the power supply out of the bay.

Step 3 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.

Figure 38: Replacing AC Power Supplies



1	Power supply release lever	2	Power supply handle
----------	----------------------------	----------	---------------------

Step 4 Only if you shut down the node, perform these steps:

- a) Recommission the node in UCS Manager as described in [Recommissioning the Node Using Cisco UCS Manager, on page 29](#).
- b) Associate the node with its UCS Manager service profile as described in [Associating a Service Profile With an HX Node, on page 29](#).

- c) After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, on page 29](#).

Replacing DC Power Supplies



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 two DC power supplies, you do not have to shut down the node. Continue with step 2.
 - If your node has only one power supply:
 1. Put the node in Cisco HX Maintenance Mode as described in [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#).
 2. Shut down the node as described in [Shutting Down and Removing Power From the Node, on page 26](#).
 3. Decommission the node from UCS Manager as described in [Decommissioning the Node Using Cisco UCS Manager, on page 25](#).
- 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 39: Replacing DC Power Supplies



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

Step 3

Only if you shut down the node, perform these steps:

- a) Recommission the node in UCS Manager as described in [Recommissioning the Node Using Cisco UCS Manager, on page 29](#).
- b) Associate the node with its UCS Manager service profile as described in [Associating a Service Profile With an HX Node, on page 29](#).
- c) After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, on page 29](#).

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 threaded holes for use with your dual-hole grounding lug and grounding wire are supplied on the chassis rear panel.



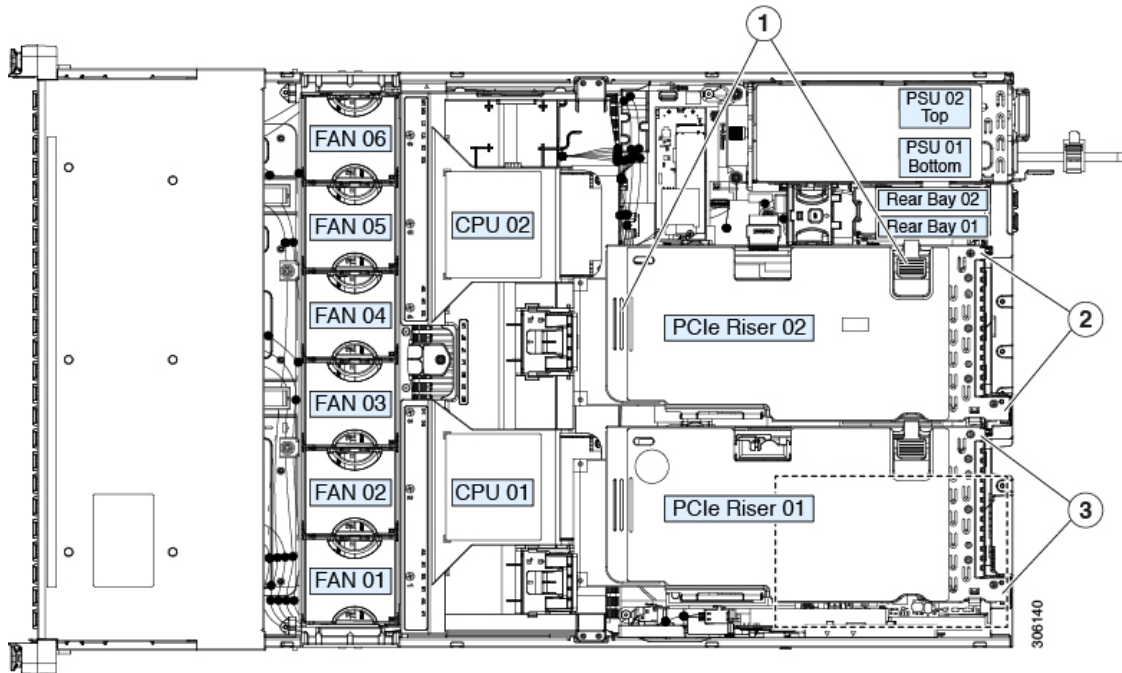
Note The grounding points on the chassis are sized for M5 screws. You must provide your own screws, grounding lug, and grounding wire. The grounding lug must be dual-hole lug that fits M5 screws. The grounding cable that you provide must be 14 AWG (2 mm), minimum 60° C wire, or as permitted by the local 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 71](#) for detailed descriptions of the slots and features in each riser version.

-
- Step 1** Put the node in Cisco HX Maintenance Mode as described in [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#).
- Step 2** Shut down the node as described in [Shutting Down and Removing Power From the Node, on page 26](#).
- Step 3** Decommission the node from UCS Manager as described in [Decommissioning the Node Using Cisco UCS Manager, on page 25](#).
- 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.
- Step 4** Disconnect all power cables from all power supplies.
- Step 5** 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 6** Remove the top cover from the node as described in [Removing the Node Top Cover, on page 30](#).
- Step 7** 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 73](#).
- Step 8** 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 73](#).
 - Position the PCIe riser over its socket on the motherboard and over its alignment slots in the chassis.
 - Carefully push down on both ends of the PCIe riser to fully engage its circuit board connector with the socket on the motherboard.
- Step 9** Replace the top cover to the node.
- Step 10** Replace the node in the rack, replace cables, and then fully power on the node by pressing the Power button.

Figure 40: PCIe Riser Alignment Features



1	Riser handling points (flip-up handle and blue forward edge)	3	Riser 1 alignment features in chassis
2	Riser 2 alignment features in chassis		

- Step 11** Recommission the node in UCS Manager as described in [Recommissioning the Node Using Cisco UCS Manager, on page 29](#).
- Step 12** Associate the node with its UCS Manager service profile as described in [Associating a Service Profile With an HX Node, on page 29](#).
- Step 13** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, on page 29](#).

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 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 server contains two toolless PCIe risers for horizontal installation of PCIe cards.

- Riser 1 contains PCIe slots 1 (x8), 2 (x16), and 3 (x8).

- Riser 2B contains PCIe slots 4 (x8), 5 (x16), and 6 (x8). It also includes one PCIe cable connector for rear-loading NVMe SSDs.

Figure 41: Rear Panel, Showing PCIe Slot Numbering



The following tables describe the specifications for the slots.

Table 9: PCIe Riser 1 (UCSC-PCI-1-C240M5) PCIe Expansion Slots

Slot Number	Electrical Lane Width	Connector Length	Maximum Card Length	Card Height (Rear Panel Opening)	NCSI Support	GPU Card Support
1	Gen-3 x8	x24 connector	¾ length	Full height	Yes ¹	No
2	Gen-3 x16	x24 connector	Full length	Full height	Yes	Yes
3	Gen-3 x8	x16 connector	Full length	Full height	No	No
Micro SD card slot	One socket for Micro SD card on the top of the riser.					

¹ NCSI is supported in only one slot at a time. If a GPU card is present in slot 2, NCSI support automatically moves to slot 1.

Table 10: PCIe Riser 2B (UCSC-PCI-2B-C240M5) PCIe Expansion Slots

Slot Number	Electrical Lane Width	Connector Length	Maximum Card Length	Card Height (Rear Panel Opening)	NCSI Support	GPU Card Support
4	Gen-3 x8	x24 connector	¾ length	Full height	Yes	No
5	Gen-3 x16	x24 connector	Full length	Full height	Yes ²	Yes
6	Gen-3 x8	x16 connector	Full length	Full height	No	No
Cable connector for rear NVMe SSD	Gen-3 x8	To rear drive backplane; supports rear-loading NVMe SSDs.				

² NCSI is supported in only one slot at a time.

Replacing a PCIe Card



Note A storage controller card installs into a dedicated motherboard socket. See [Replacing a SAS Storage Controller Card \(HBA\), on page 76](#).



Note For instructions on installing or replacing double-wide GPU cards, see [Installing a Double-Wide GPU Card, on page 104](#).



Note If the card you are replacing is a Cisco VIC 1455 (HX-PCIE-C25Q-04), note that this card requires Cisco HX 4.0(1a) or later.


- Step 1** Put the node in Cisco HX Maintenance Mode as described in [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#).
- Step 2** Shut down the node as described in [Shutting Down and Removing Power From the Node, on page 26](#).
- Step 3** Decommission the node from UCS Manager as described in [Decommissioning the Node Using Cisco UCS Manager, on page 25](#).
- 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.
- Step 4** Disconnect all power cables from all power supplies.
- Step 5** 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 6** Remove the top cover from the node as described in [Removing the Node Top Cover, on page 30](#).
- Step 7** 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 fingergrip 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.
 - 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 8** Install a new PCIe card:
- With the hinged card-tab retainer open, align the new PCIe 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.

- 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 9 Replace the top cover to the node.

Step 10 Replace the node in the rack, replace cables, and then fully power on the node by pressing the Power button.

Figure 42: PCIe Riser Card Securing Mechanisms



306229

1	Release latch on hinged securing plate	3	Hinged card-tab retainer
2	Hinged securing plate	-	

Step 11 Recommission the node in UCS Manager as described in [Recommissioning the Node Using Cisco UCS Manager](#), on page 29.

Step 12 Associate the node with its UCS Manager service profile as described in [Associating a Service Profile With an HX Node](#), on page 29.

Step 13 After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode](#), on page 29.

Replacing an mLOM Card

The node supports a modular LOM (mLOM) card to provide additional rear-panel connectivity, such as a Cisco VIC adapter. 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 the card you are replacing is a Cisco VIC 1457 (HX-MLOM-C25Q-04), note that this card requires Cisco HX 4.0(1a) or later.

- Step 1** Put the node in Cisco HX Maintenance Mode as described in [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#).
- Step 2** Shut down the node as described in [Shutting Down and Removing Power From the Node, on page 26](#).
- Step 3** Decommission the node from UCS Manager as described in [Decommissioning the Node Using Cisco UCS Manager, on page 25](#).
- 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.
- Step 4** Disconnect all power cables from all power supplies.
- Step 5** Remove any existing mLOM card (or a blanking panel):
- 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 30](#).
 - Remove any storage controller HBA card to provide clearance to the mLOM socket on the motherboard.
 - 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 6** Install a new mLOM card:
- Set the mLOM card on the chassis floor so that its connector is aligned with the motherboard socket.
 - Push the card horizontally to fully engage the card's edge connector with the socket.
 - Tighten the captive thumbscrew to secure the card to the chassis floor.
 - Return the storage controller card to the node.
 - Replace the top cover to the node.
 - Replace the node in the rack, replace cables, and then fully power on the node by pressing the Power button.

Figure 43: Location of the mLOM Card Socket Below the Storage Controller Card



1	Position of horizontal mLOM card socket	2	Position of mLOM card thumbscrew
---	---	---	----------------------------------

- Step 7** Recommission the node in UCS Manager as described in [Recommissioning the Node Using Cisco UCS Manager](#), on page 29.
- Step 8** Associate the node with its UCS Manager service profile as described in [Associating a Service Profile With an HX Node](#), on page 29.
- Step 9** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode](#), on page 29.

Replacing a SAS Storage Controller Card (HBA)

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

Storage Controller Card Firmware Compatibility

Firmware on the storage controller 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.

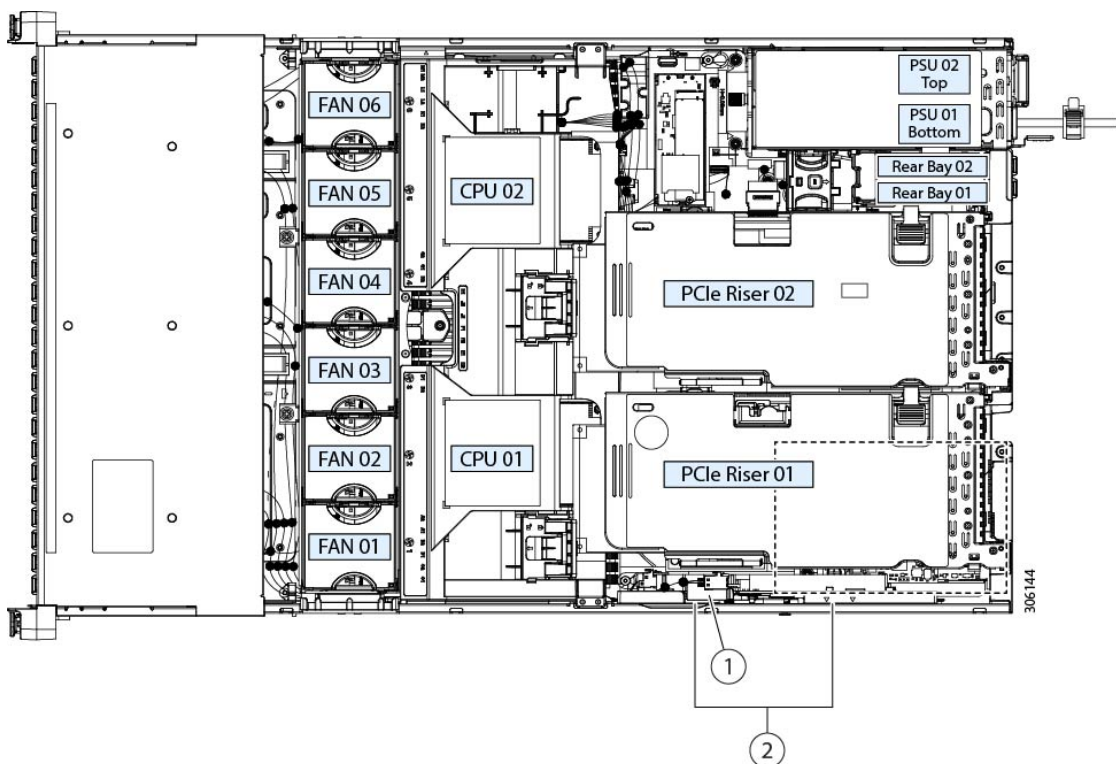
See the HUU guide for your Cisco IMC release for instructions on downloading and using the utility to bring node components to compatible levels: [HUU Guides](#).

Replacing a SAS Storage Controller Card (HBA)

A storage controller card is shipped in a plastic mounting frame. You do not have to remove this frame from an existing card because the frame is preinstalled on a replacement card.

-
- Step 1** Put the node in Cisco HX Maintenance Mode as described in [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#).
- Step 2** Shut down the node as described in [Shutting Down and Removing Power From the Node, on page 26](#).
- Step 3** Decommission the node from UCS Manager as described in [Decommissioning the Node Using Cisco UCS Manager, on page 25](#).
- 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.
- Step 4** Disconnect all power cables from all power supplies.
- Step 5** Prepare the node for component installation:
- 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 30](#).
- Step 6** Remove any existing storage controller card from the node:
- Note** A controller card for this node is mounted inside a plastic carrier-frame that helps to secure the card to the inner chassis wall. You do not have to remove this plastic carrier frame from the existing card.
- Disconnect cables from the existing card.
 - Lift up on the card's blue ejector lever to unseat it from the motherboard socket.
 - Lift straight up on the card's carrier frame to disengage the card from the motherboard socket and to disengage the frame from pegs on the chassis wall.
- Step 7** Install a new storage controller card:
- Carefully align the card edge with the motherboard socket. At the same time, align the two slots on the back of the carrier frame with the pegs on the chassis inner wall.
 - Push on both corners of the card to seat its connector in the riser socket. At the same time, ensure that the slots on the carrier frame engage with the pegs on the inner chassis wall.
 - Fully close the blue ejector lever on the card to lock the card into the socket.
 - Connect SAS/SATA cables to the new card.
- 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.

Figure 44: Replacing a Storage Controller Card



1	Blue ejector lever on card top edge	2	Pegs on inner chassis wall (four)
---	-------------------------------------	---	-----------------------------------

- Step 10** Recommission the node in UCS Manager as described in [Recommissioning the Node Using Cisco UCS Manager, on page 29](#).
- Step 11** Associate the node with its UCS Manager service profile as described in [Associating a Service Profile With an HX Node, on page 29](#).
- Step 12** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, on page 29](#).

Replacing a Micro SD Card

There is one socket for a Micro SD card on the top of PCIe riser 1.



Caution To avoid data loss, we do not recommend that you hot-swap the Micro SD card while it is operating, as indicated by its activity LED turning amber. The activity LED turns amber when the Micro SD card is updating or deleting.

- Step 1** Put the node in Cisco HX Maintenance Mode as described in [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#).
- Step 2** Shut down the node as described in [Shutting Down and Removing Power From the Node, on page 26](#).

Step 3 Decommission the node from UCS Manager as described in [Decommissioning the Node Using Cisco UCS Manager, on page 25](#).

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.

Step 4 Disconnect all power cables from all power supplies.

Step 5 Remove an existing Micro SD card:

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 30](#).

c) Locate the Micro SD card. The socket is on the top of PCIe riser 1, under a plastic cover.

d) Use your fingertip to push the retainer on the plastic socket cover open far enough to provide access to the Micro SD card, then push down and release the Micro SD card to make it spring up.

e) Grasp the Micro SD card and lift it from the socket.

Step 6 Install a new Micro SD card:


a) While holding the retainer on the plastic cover open with your fingertip, align the new Micro SD card with the socket.

b) Gently push down on the card until it clicks and locks in place in the socket.

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 45: Location of Internal Micro SD Card Socket



1	Location of Micro SD card socket on the top of PCIe riser 1	3	Plastic retainer (push aside to access socket)
2	Micro SD card socket under plastic retainer	4	Micro SD activity LED

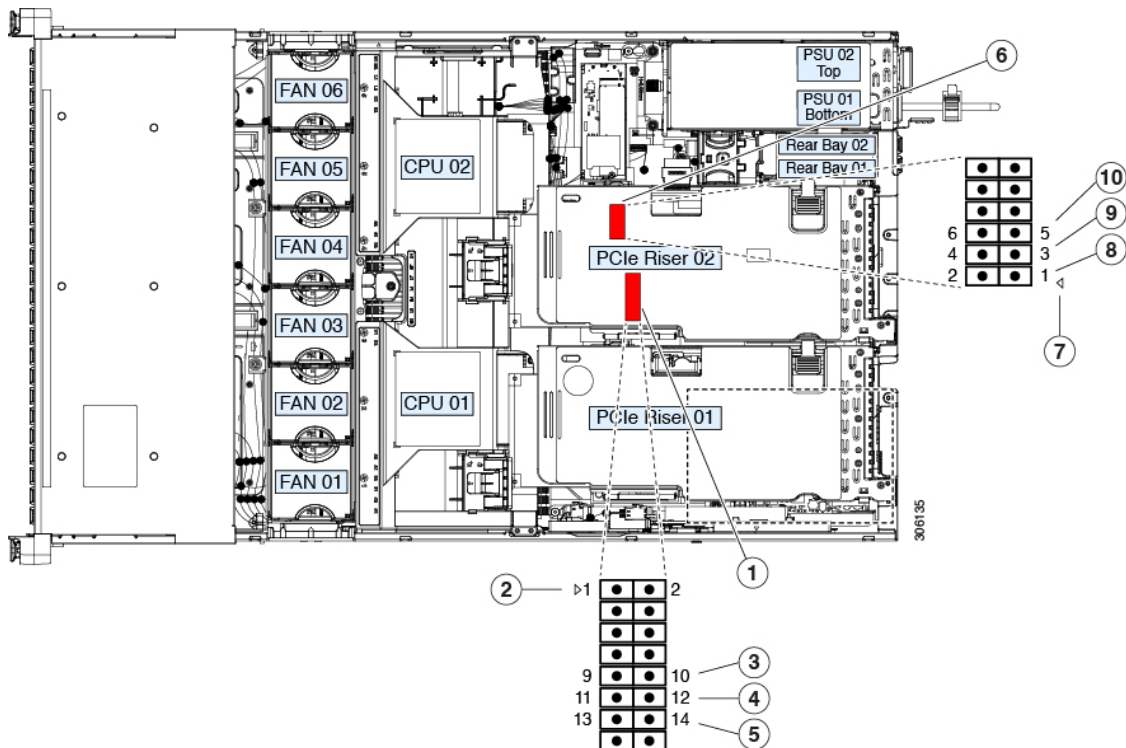
Step 7 Recommission the node in UCS Manager as described in [Recommissioning the Node Using Cisco UCS Manager, on page 29](#).

- Step 8** Associate the node with its UCS Manager service profile as described in [Associating a Service Profile With an HX Node, on page 29](#).
- Step 9** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, on page 29](#).

Service Headers and Jumpers

This node includes two blocks of headers (J38, J39) that you can jumper for certain service and debug functions.

Figure 46: Location of Service Header Blocks J38 and J39



1	Location of header block J38	6	Location of header block J39
2	J38 pin 1 arrow printed on motherboard	7	J39 pin 1 arrow printed on motherboard
3	Clear CMOS: J38 pins 9 - 10	8	Boot Cisco IMC from alternate image: J39 pins 1 - 2
4	Recover BIOS: J38 pins 11 - 12	9	Reset Cisco IMC password to default: J39 pins 3 - 4
5	Clear password: J38 pins 13 - 14	10	Reset Cisco IMC to defaults: J39 pins 5 - 6

Using the Clear CMOS Header (J38, Pins 9 - 10)

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 jumper to invalidate the settings and reboot with defaults.



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** Put the node in Cisco HX Maintenance Mode as described in [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#).
- Step 2** Shut down the node as described in [Shutting Down and Removing Power From the Node, on page 26](#).
- Step 3** Decommission the node from UCS Manager as described in [Decommissioning the Node Using Cisco UCS Manager, on page 25](#).
- 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.
- Step 4** Disconnect all power cables from all power supplies.
- Step 5** 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 6** Remove the top cover from the node as described in [Removing the Node Top Cover, on page 30](#).
- Step 7** Install a two-pin jumper across J38 pins 9 and 10.
- Step 8** 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 9** 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 10** 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 11** Remove the top cover from the node.
- Step 12** Remove the jumper that you installed.
- Note** If you do not remove the jumper, the CMOS settings are reset to the defaults every time you power-cycle the node.
- Step 13** 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.
- Step 14** Recommission the node in UCS Manager as described in [Recommissioning the Node Using Cisco UCS Manager, on page 29](#).
- Step 15** Associate the node with its UCS Manager service profile as described in [Associating a Service Profile With an HX Node, on page 29](#).

Step 16 After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, on page 29](#).

Using the BIOS Recovery Header (J38, Pins 11 - 12)

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 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 recovery.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.

Procedure 2: Use BIOS Recovery Header and bios.cap 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 Put the node in Cisco HX Maintenance Mode as described in [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#).

Step 4 Shut down the node as described in [Shutting Down and Removing Power From the Node, on page 26](#).

Step 5 Decommission the node from UCS Manager as described in [Decommissioning the Node Using Cisco UCS Manager, on page 25](#).

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.

Step 6 Disconnect all power cables from all power supplies.

Step 7 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 8 Remove the top cover from the node as described in [Removing the Node Top Cover, on page 30](#).

Step 9 Install a two-pin jumper across J38 pins 11 and 12.

Step 10 Reconnect AC power cords to the node. The node powers up to standby power mode.

Step 11 Insert the USB thumb drive that you prepared in Step 2 into a USB port on the node.

Step 12 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 13 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.

Step 14 After the node has fully booted, power off the node again and disconnect all power cords.

Step 15 Remove the jumper that you installed.

Note If you do not remove the jumper, after recovery completion you see the prompt, "Please remove the recovery jumper."

- Step 16** 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.
- Step 17** Recommission the node in UCS Manager as described in [Recommissioning the Node Using Cisco UCS Manager, on page 29](#).
- Step 18** Associate the node with its UCS Manager service profile as described in [Associating a Service Profile With an HX Node, on page 29](#).
- Step 19** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, on page 29](#).
-

Using the Clear Password Header (J38, Pins 13 - 14)

You can use this switch to clear the administrator password.

- Step 1** Put the node in Cisco HX Maintenance Mode as described in [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#).
- Step 2** Shut down the node as described in [Shutting Down and Removing Power From the Node, on page 26](#).
- Step 3** Decommission the node from UCS Manager as described in [Decommissioning the Node Using Cisco UCS Manager, on page 25](#).
- 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.
- Step 4** Disconnect all power cables from all power supplies.
- Step 5** 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 6** Remove the top cover from the node as described in [Removing the Node Top Cover, on page 30](#).
- Step 7** Install a two-pin jumper across J38 pins 13 and 14.
- Step 8** 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 9** 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 10** 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 11** Remove the top cover from the node.
- Step 12** 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 13** 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.
- Step 14** Recommission the node in UCS Manager as described in [Recommissioning the Node Using Cisco UCS Manager, on page 29](#).

- Step 15** Associate the node with its UCS Manager service profile as described in [Associating a Service Profile With an HX Node, on page 29](#).
- Step 16** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, on page 29](#).

Using the Boot Alternate Cisco IMC Image Header (J39, Pins 1 - 2)

You can use this Cisco IMC debug header to force the system to boot from an alternate Cisco IMC image.

-
- Step 1** Put the node in Cisco HX Maintenance Mode as described in [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#).
- Step 2** Shut down the node as described in [Shutting Down and Removing Power From the Node, on page 26](#).
- Step 3** Decommission the node from UCS Manager as described in [Decommissioning the Node Using Cisco UCS Manager, on page 25](#).
- 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.
- Step 4** Disconnect all power cables from all power supplies.
- Step 5** 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 6** Remove the top cover from the node as described in [Removing the Node Top Cover, on page 30](#).
- Step 7** Install a two-pin jumper across J39 pins 1 and 2.
- Step 8** 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 9** 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 10** 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 11** Remove the top cover from the node.
- Step 12** Remove the jumper that you installed.
- Step 13** 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.
- Step 14** Recommission the node in UCS Manager as described in [Recommissioning the Node Using Cisco UCS Manager, on page 29](#).
- Step 15** Associate the node with its UCS Manager service profile as described in [Associating a Service Profile With an HX Node, on page 29](#).

**Step 16** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, on page 29](#).

---

## Using the Reset Cisco IMC Password to Default Header (J39, Pins 3 - 4)

You can use this Cisco IMC debug header to force the Cisco IMC password back to the default.

---

- Step 1** Put the node in Cisco HX Maintenance Mode as described in [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#).
- Step 2** Shut down the node as described in [Shutting Down and Removing Power From the Node, on page 26](#).
- Step 3** Decommission the node from UCS Manager as described in [Decommissioning the Node Using Cisco UCS Manager, on page 25](#).

**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.

- Step 4** Disconnect all power cables from all power supplies.
- Step 5** 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 6** Remove the top cover from the node as described in [Removing the Node Top Cover, on page 30](#).
- Step 7** Install a two-pin jumper across J39 pins 3 and 4.
- Step 8** 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 9** 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:

```
'Reset to default CIMC password' debug functionality is enabled.
On input power cycle, CIMC password will be reset to defaults.
```

**Note** If you do not remove the jumper, the node will reset the Cisco IMC password to the default every time that you power cycle the node. The jumper has no effect if you reboot Cisco IMC.

- Step 10** 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 11** Remove the top cover from the node.
- Step 12** Remove the jumper that you installed.
- Step 13** 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.
- Step 14** Recommission the node in UCS Manager as described in [Recommissioning the Node Using Cisco UCS Manager, on page 29](#).
- Step 15** Associate the node with its UCS Manager service profile as described in [Associating a Service Profile With an HX Node, on page 29](#).



- Step 16** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, on page 29](#).
- 

## Using the Reset Cisco IMC to Defaults Header (J39, Pins 5 - 6)

You can use this Cisco IMC debug header to force the Cisco IMC settings back to the defaults.

---

- Step 1** Put the node in Cisco HX Maintenance Mode as described in [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#).
- Step 2** Shut down the node as described in [Shutting Down and Removing Power From the Node, on page 26](#).
- Step 3** Decommission the node from UCS Manager as described in [Decommissioning the Node Using Cisco UCS Manager, on page 25](#).
- 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.
- Step 4** Disconnect all power cables from all power supplies.
- Step 5** 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 6** Remove the top cover from the node as described in [Removing the Node Top Cover, on page 30](#).
- Step 7** Install a two-pin jumper across J39 pins 5 and 6.
- Step 8** 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 9** 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:
- ```
'CIMC reset to factory defaults' debug functionality is enabled.  
On input power cycle, CIMC will be reset to factory defaults.
```
- Note** If you do not remove the jumper, the node will reset the Cisco IMC to the default settings every time that you power cycle the node. The jumper has no effect if you reboot Cisco IMC.
- Step 10** 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 11** Remove the top cover from the node.
- Step 12** Remove the jumper that you installed.
- Step 13** 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.
- Step 14** Recommission the node in UCS Manager as described in [Recommissioning the Node Using Cisco UCS Manager, on page 29](#).
- Step 15** Associate the node with its UCS Manager service profile as described in [Associating a Service Profile With an HX Node, on page 29](#).

Step 16 After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, on page 29](#).

Setting Up the Node in Standalone Mode



Note The HX Series node is always managed in UCS Manager-controlled mode. This section is included only for cases in which a node might need to be put into standalone mode for troubleshooting purposes. Do not use this setup for normal operation of the HX Series node.

Initial Node Setup (Standalone)



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

Node 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 Node With the Cisco IMC Configuration Utility, on page 90](#).
- 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 server on the same network as the system. Your DHCP server 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.

Connecting to the Node Locally For Standalone 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 power outlet.
- 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:
- a) Press and hold the front panel power button for four seconds to boot the node.
 - b) During bootup, press **F8** when prompted to open the Cisco IMC Configuration Utility.
- Step 4** Continue with [Setting Up the Node With the Cisco IMC Configuration Utility, on page 90](#).
-

Connecting to the Node Remotely For Standalone 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 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.

- Step 1** Attach a power cord to each power supply in your node, and then attach each power cord to a grounded power outlet. 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 server to assign an IP address to the node.
- Step 4** Use the assigned IP address to access and log in to the Cisco IMC for the node. Consult with your DHCP node administrator to determine the IP address.
- Note** The default user name for the node is *admin*. The default password is *password*.
- Step 5** From the Cisco IMC node Summary page, click **Launch KVM Console**. A separate KVM console window opens.
- Step 6** From the Cisco IMC Summary page, click **Power Cycle node**. 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.
- Step 9** Continue with [Setting Up the Node With the Cisco IMC Configuration Utility](#), on page 90.
-

Setting Up the Node With the Cisco IMC Configuration Utility**Before you begin**

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

- Step 1** Set the NIC mode to choose which ports to use to access Cisco IMC for server 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 server 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 1.
 - If you select Riser2, you must install the VIC in slot 2.
 - 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 server 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 server with the range of MAC addresses for this server. The MAC address is printed on a label on the rear of the server. This server 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 server 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 server.

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 server reverts to the factory defaults.

Step 9 (Optional) Set a default user password.

Note The factory default username for the server 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 server 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 server in the next step.

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

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 server).



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



APPENDIX **A**

Node Specifications

- [Physical Specifications](#), on page 93
- [Environmental Specifications](#), on page 93
- [Power Specifications](#), on page 94
- [Power Cord Specifications](#), on page 97

Physical Specifications

The following table lists the physical specifications for the node.

Table 11: 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) 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 12: Physical Specifications

Description	Specification
-------------	---------------

Temperature, Operating	41 to 95°F (5 to 35°C) Derate the maximum temperature by 1°C per every 305 meters of altitude above sea level. For general information, see the Cisco Unified Computing System Site Planning Guide: Data Center Power and Cooling .
Temperature, non-operating (when the server is stored or transported)	–40 to 149°F (–40 to 65°C)
Humidity (RH), operating	10 to 90%
Humidity (RH), non-operating (when the server is stored or transported)	5 to 93%
Altitude, operating	0 to 10,000 feet
Altitude, non-operating (when the server is stored or transported)	0 to 40,000 feet
Sound power level Measure A-weighted per ISO7779 LwAd (Bels) Operation at 73°F (23°C)	5.5
Sound pressure level Measure A-weighted per ISO7779 LpAm (dBA) 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 13: 1050 W AC Specifications

Description	Specification
AC Input Voltage	Nominal range: 100–120 VAC, 200–240 VAC (Range: 90–132 VAC, 180–264 VAC)
AC Input Frequency	Nominal range: 50 to 60Hz (Range: 47–63 Hz)
Maximum AC Input current	12.5 A at 100 VAC 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 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 14: 1600 W AC Specifications

Description	Specification
AC Input Voltage	Nominal range: 200–240 VAC (Range: 180–264 VAC)
AC Input Frequency	Nominal range: 50 to 60Hz (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 15: 1050 W DC Specifications

Description	Specification
DC Input Voltage	Nominal range: -48 to -60 VDC (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 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 16: Supported Power Cords

Description	Length (Feet)	Length (Meters)
CAB-48DC-40A-8AWG DC power cord, -48 VDC, 40 A, 8 AWG Three-socket Mini-Fit connector to three-wire	11.7	3.5
CAB-C13-C14-AC AC power cord, 10 A; C13 to C14, recessed receptacle	9.8	3.0
CAB-250V-10A-AR AC power cord, 250 V, 10 A Argentina	8.2	2.5
CAB-C13-C14-2M-JP AC Power Cord, C13 to C14 Japan PSE Mark	6.6	2.0
CAB-9K10A-EU AC Power Cord, 250 V, 10 A; CEE 7/7 Plug Europe	8.2	2.5
CAB-250V-10A-IS AC Power Cord, 250 V, 10 A Israel	8.2	2.5
CAB-250V-10A-CN AC power cord, 250 V, 10 A PR China	8.2	2.5
CAB-ACTW AC power cord, 250 V, 10 A Taiwan	7.5	2.3

CAB-C13-CBN AC cabinet jumper power cord, 250, 10 A, C13 to C14	2.2	0.68
CAB-C13-C14-2M AC cabinet jumper power cord, 250 V, 10 A, C13 to C14	6.6	2.0
CAB-9K10A-AU AC power cord, 250 V, 10 A, 3112 plug, Australia	8.2	2.5
CAB-N5K6A-NA AC power cord, 200/240 V, 6 A, North America	8.2	2.5
CAB-250V-10A-ID AC power Cord, 250 V, 10 A, India	8.2	2.5
CAB-9K10A-SW AC power cord, 250 V, 10 A, MP232 plug Switzerland	8.2	2.5
CAB-250V-10A-BR AC power Cord, 250 V, 10 A Brazil	8.2	2.5
CAB-9K10A-UK AC power cord, 250 V, 10 A (13 A fuse), BS1363 plug United Kingdom	8.2	2.5
CAB-9K12A-NA AC power cord, 125 V, 13 A, NEMA 5-15 plug North America	8.2	2.5
CAB-AC-L620-C13 AC power cord, NEMA L6-20 to C13 connectors	6.6	2.0
CAB-9K10A-IT AC power cord, 250 V, 10 A, CEI 23-16/VII plug Italy	8.2	2.5

R2XX-DMYMPWRCORD 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.

- [Supported Storage Controller and Cables, on page 101](#)
- [Storage Controller Card Firmware Compatibility, on page 101](#)

Supported Storage Controller and Cables

This node supports a single PCIe-style, SAS HBA controller that plugs into a dedicated internal riser.

This node supports the storage controller options and cable requirements shown in the following table.

Controller	Maximum Drives Controlled	RAID Levels	Optional Supercap Backup?	Required Cables
HX240c SFF Versions: Cisco 12G Modular SAS HBA (for up to 26 drives) HX-SAS-M5HD	24 front-loading SAS/SATA drives, plus up to 2 rear-loading SAS/SATA drives.	Non-RAID	No	Use SAS/SATA cable included with chassis to connect controller to drive backplane.
HX240c LFF Versions: Cisco 12G Modular SAS HBA (for up to 14 drives) HX-SAS-M5	12 front-loading SAS/SATA drives, plus up to 2 rear-loading SAS/SATA drives.	Non-RAID	No	Use SAS/SATA cable included with chassis to connect controller to drive backplane.

Storage Controller Card Firmware Compatibility

Firmware on the storage controller 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.

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](#).



APPENDIX **C**

GPU Card Installation

This appendix contains configuration rules and installation procedures for the supported GPU cards.

- [GPU Card Configuration Rules, on page 103](#)
- [Requirement For All GPUs: Memory-Mapped I/O Greater Than 4 GB, on page 104](#)
- [Installing a Double-Wide GPU Card, on page 104](#)
- [Using NVIDIA GRID License Server For M-Series and P-Series GPUs, on page 107](#)
- [Installing Drivers to Support the NVIDIA GPU Cards, on page 114](#)

GPU Card Configuration Rules

Note the following rules when populating a node with GPU cards.



Caution

When using NVIDIA Tesla P40 GPU cards in this node, the maximum operating temperature (air inlet temperature) for the node is 32° C (89.6° F).

- Double-wide GPU cards are supported in PCIe riser 1--slot 2 and in PCIe riser 2--slot 5.
- A double-wide GPU card installed in slot 2 also covers slot 4; a double wide GPU card installed in slot 5 also covers slot 6.
- Do not mix different brands or models of GPU cards in the node.
- You can install a GPU card and a Cisco UCS VIC in the same riser. When you install a GPU card in slot 2, NCSI support in riser 1 automatically moves to slot 1. When you install a GPU card in slot 5, NCSI support in riser 2 automatically moves to slot 4.
- AMD FirePro S7150 X2 GPUs can support only less-than 1 TB memory in the server.
- NVIDIA M-Series GPUs can support only less-than 1 TB memory in the server.
- NVIDIA P-Series GPUs can support 1 TB or more memory in the server.

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 node: If the node 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 node 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 node-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 node. 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 Double-Wide GPU Card

Use the following procedure to install or replace the following supported GPU cards:

- NVIDIA Tesla M10
- NVIDIA Tesla P40
- AMD FirePro S7150 X2



Caution When using NVIDIA Tesla P40 GPU cards in this node, the maximum operating temperature (air inlet temperature) for the node is 32° C (89.6° F).

Table 17: HX240c M5 Operating Temperature Requirements For GPU Cards

GPU Card	Maximum node Operating Temperature (Air Inlet Temperature)
NVIDIA Tesla M10	35° C (95.0° F)
NVIDIA Tesla P40	32° C (89.6° F)
AMD FirePro S7150 X2	35° C (95.0° F)



Note **For NVIDIA GPUs:** 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 (such as the Magma chassis).

For AMD GPUs: The correct power cable is a Y-cable.

- Step 1** Put the node in Cisco HX Maintenance Mode as described in [Shutting Down Using vSphere With HX Maintenance Mode, on page 27](#).
- Step 2** Shut down the node as described in [Shutting Down and Removing Power From the Node, on page 26](#).
- Step 3** Decommission the node from UCS Manager as described in [Decommissioning the Node Using Cisco UCS Manager, on page 25](#).
- 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.
- Step 4** Disconnect all power cables from all power supplies.
- Step 5** 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 6** Remove the top cover from the node as described in [Removing the Node Top Cover, on page 30](#).
- Step 7** Remove an existing 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.
 - Pull evenly on both ends of the GPU card to remove it from the socket on the PCIe riser.
- Step 8** Install a new GPU card:

Note Observe the configuration rules for this node, as described in [GPU Card Configuration Rules, on page 103](#).

- a) 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.
- b) 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.

- c) Close the card-tab retainer over the end of the card.
- 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.


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 9

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.

Figure 47: GPU Front Support Bracket Inserted to Securing latch on Air Baffle



1	Front end of GPU card	3	Lip on securing latch
---	-----------------------	---	-----------------------

2	GPU front support bracket	4	Securing latch release tab
---	---------------------------	---	----------------------------

- Step 10** Replace the top cover to the node.
- Step 11** Replace the node in the rack, replace cables, and then fully power on the node by pressing the Power button.
- Step 12** Recommission the node in UCS Manager as described in [Recommissioning the Node Using Cisco UCS Manager, on page 29](#).
- Step 13** Associate the node with its UCS Manager service profile as described in [Associating a Service Profile With an HX Node, on page 29](#).
- Step 14** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, on page 29](#).
- Note** If you installed an NVIDIA Tesla M-series or P-Series GPU, you must install GRID licenses to use the optional GRID features. See [Using NVIDIA GRID License Server For M-Series and P-Series GPUs, on page 107](#).

Using NVIDIA GRID License Server For M-Series and P-Series GPUs

This section applies to NVIDIA Tesla M-Series and P-Series GPUs.

Use the topics in this section in the following order when obtaining and using NVIDIA GRID licenses.

1. Familiarize yourself with the NVIDIA GRID License Server.
[NVIDIA GRID License Server Overview, on page 107](#)
2. Register your product activation keys with NVIDIA.
[Registering Your Product Activation Keys With NVIDIA, on page 108](#)
3. Download the GRID software suite.
[Downloading the GRID Software Suite, on page 109](#)
4. Install the GRID License Server software to a host.
[Installing NVIDIA GRID License Server Software, on page 109](#)
5. Generate licenses on the NVIDIA Licensing Portal and download them.
[Installing Licenses From the Licensing Portal, on page 111](#)
6. Manage your GRID licenses.
[Managing GRID Licenses , on page 112](#)

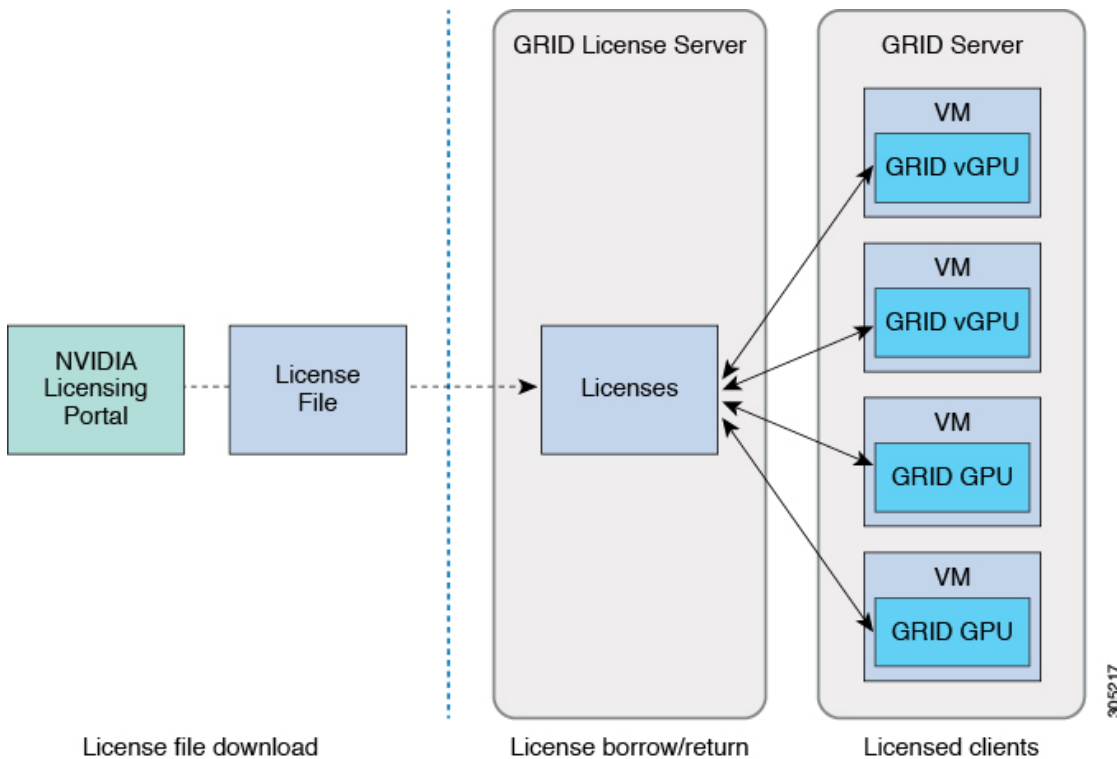
NVIDIA GRID License Server Overview

The NVIDIA M-Series GPUs combine Tesla and GRID functionality when the licensed GRID features such as GRID vGPU and GRID Virtual Workstation are enabled. These features are enabled during OS boot by

borrowing a software license that is served over the network from the NVIDIA GRID License Server virtual appliance. The license is returned to the license server when the OS shuts down.

You obtain the licenses that are served by the GRID License Server from NVIDIA's Licensing Portal as downloadable license files, which you install into the GRID License Server via its management interface.

Figure 48: NVIDIA GRID Licensing Architecture



There are three editions of GRID licenses, which enable three different classes of GRID features. The GRID software automatically selects the license edition based on the features that you are using.

GRID License Edition	GRID Feature
GRID Virtual GPU (vGPU)	Virtual GPUs for business desktop computing
GRID Virtual Workstation	Virtual GPUs for midrange workstation computing
GRID Virtual Workstation – Extended	Virtual GPUs for high-end workstation computing Workstation graphics on GPU pass-through

Registering Your Product Activation Keys With NVIDIA

After your order is processed, NVIDIA sends you a Welcome email that contains your product activation keys (PAKs) and a list of the types and quantities of licenses that you purchased.

Step 1 Select the **Log In** link, or the **Register** link if you do not already have an account.

The NVIDIA Software Licensing Center > License Key Registration dialog opens.

Step 2 Complete the License Key Registration form and then click **Submit My Registration Information**.

The NVIDIA Software Licensing Center > Product Information Software dialog opens.

Step 3 If you have additional PAKs, click **Register Additional Keys**. For each additional key, complete the form on the License Key Registration dialog and then click **Submit My Registration Information**.

Step 4 Agree to the terms and conditions and set a password when prompted.

Downloading the GRID Software Suite

Step 1 Return to the NVIDIA Software Licensing Center > Product Information Software dialog.

Step 2 Click the **Current Releases** tab.

Step 3 Click the **NVIDIA GRID** link to access the Product Download dialog. This dialog includes download links for:

- NVIDIA License Manager software
- The gpumodeswitch utility
- The host driver software

Step 4 Use the links to download the software.

Installing NVIDIA GRID License Server Software

For full installation instructions and troubleshooting, refer to the *NVIDIA GRID License Server User Guide*. Also refer to the *NVIDIA GRID License Server Release Notes* for the latest information about your release.

<http://www.nvidia.com>

Platform Requirements for NVIDIA GRID License Server

- The hosting platform can be a physical or a virtual machine. NVIDIA recommends using a host that is dedicated only to running the License Server.
- The hosting platform must run a supported Windows OS.
- The hosting platform must have a constant IP address.
- The hosting platform must have at least one constant Ethernet MAC address.
- The hosting platform's date and time must be set accurately.

Installing GRID License Server on Windows

The License Server requires a Java runtime environment and an Apache Tomcat installation. Apache Tomcat is installed when you use the NVIDIA installation wizard for Windows.

-
- Step 1** Download and install the latest Java 32-bit runtime environment from <https://www.oracle.com/downloads/index.html>.
- Note** Install the 32-bit Java Runtime Environment, regardless of whether your platform is Windows 32-bit or 64-bit.
- Step 2** Create a server interface:
- On the NVIDIA Software Licensing Center dialog, click **Grid Licensing > Create License Server**.
 - On the Create Server dialog, fill in your desired server details.
 - Save the .bin file that is generated onto your license server for installation.
- Step 3** Unzip the NVIDIA License Server installer Zip file that you downloaded previously and run setup.exe.
- Step 4** Accept the EULA for the NVIDIA License Server software and the Apache Tomcat software. Tomcat is installed automatically during the License Server installation.
- Step 5** Use the installer wizard to step through the installation.
- Note** On the Choose Firewall Options dialog, select the ports to be opened in the firewall. NVIDIA recommends that you use the default setting, which opens port 7070 but leaves port 8080 closed.
- Step 6** Verify the installation. Open a web browser on the License Server host and connect to the URL <http://localhost:8080/licserver>. If the installation was successful, you see the NVIDIA License Client Manager interface.
-

Installing GRID License Server on Linux

The License Server requires a Java runtime environment and an Apache Tomcat installation. You must install both separately before installing the License Server on Linux.

- Step 1** Verify that Java was installed with your Linux installation. Use the following command:
- ```
java -version
```
- If no Java version is displayed, use your Linux package manager to install with the following command:
- ```
sudo yum install java
```
- Step 2** Use your Linux package manager to install the tomcat and tomcat-webapps packages:
- Use the following command to install Tomcat:

```
sudo yum install tomcat
```
 - Enable the Tomcat service for automatic startup on boot:

```
sudo systemctl enable tomcat.service
```
 - Start the Tomcat service:

```
sudo systemctl start tomcat.service
```
 - Verify that the Tomcat service is operational. Open a web browser on the License Server host and connect to the URL <http://localhost:8080>. If the installation was successful, you see the Tomcat webapp.
- Step 3** Install the License Server:
- Unpack the License Server tar file using the following command:


```
tar xzf NVIDIA-linux-2015.09-0001.tgz
```

- b) Run the unpacked setup binary as root:

```
sudo ./setup.bin
```

- c) Accept the EULA and then continue with the installation wizard to finish the installation.

Note On the Choose Firewall Options dialog, select the ports to be opened in the firewall. NVIDIA recommends that you use the default setting, which opens port 7070 but leaves port 8080 closed.

- Step 4** Verify the installation. Open a web browser on the License Server host and connect to the URL <http://localhost:8080/licserver>. If the installation was successful, you see the NVIDIA License Client Manager interface.
-

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 Server 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 Server Management Interface in a browser.

Step 2 In the left-side License Server 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 NVIDIA GPU Cards

After you install the hardware, you must update to the correct level of server BIOS, activate the BIOS firmware, and then install NVIDIA drivers and other software in this order:

1. Updating the Node BIOS Firmware

Install the latest Cisco BIOS for your node by using Cisco UCS Manager.



Note You must do this procedure before you update the NVIDIA drivers.



Caution Do not remove the hardware that contains the endpoint or perform any maintenance on it until the update process completes. If the hardware is removed or otherwise unavailable due to maintenance, the firmware update fails. This failure might corrupt the backup partition. You cannot update the firmware on an endpoint with a corrupted backup partition.

-
- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** On the **Equipment** tab, expand **Equipment > Chassis > Chassis Number > Servers**.
- Step 3** Click the **Name** of the node for which you want to update the BIOS firmware.
- Step 4** On the **Properties** page in the **Inventory** tab, click **Motherboard**.
- Step 5** In the **Actions** area, click **Update BIOS Firmware**.
- Step 6** In the **Update Firmware** dialog box, do the following:
- From the **Firmware Version** drop-down list, select the firmware version to which you want to update the endpoint.
 - Click **OK**.
- Cisco UCS Manager copies the selected firmware package to the backup memory slot, where it remains until you activate it.
- Step 7** (Optional) Monitor the status of the update in the **Update Status** field.
- The update process can take several minutes. Do not activate the firmware until the firmware package you selected displays in the **Backup Version** field in the **BIOS** area of the **Inventory** tab.
-

What to do next

Activate the server BIOS firmware.

2. Activating the Node BIOS Firmware

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** On the **Equipment** tab, expand **Equipment > Chassis > Chassis Number > Servers**.
- Step 3** Click the **Name** of the server for which you want to activate the BIOS firmware.
- Step 4** On the **Properties** page in the **Inventory** tab, click **Motherboard**.
- Step 5** In the **Actions** area, click **Activate BIOS Firmware**.
- Step 6** In the **Activate Firmware** dialog box, do the following:
- Select the appropriate server BIOS version from the **Version To Be Activated** drop-down list.
 - If you want to set only the start-up version and not change the version running on the server, check **Set Startup Version Only**.
- If you configure **Set Startup Version Only**, the activated firmware moves into the pending-next-reboot state and the server is not immediately rebooted. The activated firmware does not become the running version of firmware until the server is rebooted.

c) Click **OK**.

What to do next

Update the NVIDIA drivers.

3. Updating the GPU Card Drivers

After you update the server 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 server.
- 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.
-