



Cisco HX240c M4 HyperFlex Node (Hybrid and All-Flash Models) Installation Guide

Updated April 11, 2019

Includes information for Hybrid, All-Flash, and SED versions

Cisco Systems, Inc.
www.cisco.com

Cisco has more than 200 offices worldwide.
Addresses, phone numbers, and fax numbers
are listed on the Cisco website at
www.cisco.com/go/offices.

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

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

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

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

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

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

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

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

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

CCDE, CCENT, CCSI, Cisco Eos, Cisco Explorer, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Nurse Connect, Cisco Pulse, Cisco SensorBase, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco TrustSec, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flipshare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), Cisco:Financed (Stylized), Cisco Store, Flip Gift Card, and One Million Acts of Green are service marks; and Access Registrar, Aironet, AllTouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Lumin, Cisco Nexus, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, EtherFast, EtherSwitch, Event Center, Explorer, Follow Me Browsing, GainMaker, iLNX, IOS, iPhone, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, PCNow, PIX, PowerKEY, PowerPanels, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROSA, SenderBase, SMARTnet, Spectrum Expert, StackWise, WebEx, and the WebEx logo are registered trademarks of Cisco and/or its affiliates in the United States and certain other countries.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at 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. (1005R)

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.

Cisco HX240c M4 HyperFlex Node (Hybrid and All-Flash Models) Installation Guide
© 2017 Cisco Systems, Inc. All rights reserved.



CHAPTER 1**Overview 1-1**

- Cisco HyperFlex Systems Related Documentation 1-1
- External Features Overview 1-1
- Replaceable Component Locations 1-4
- Summary of Node Features 1-5
- Cisco HX240c All-Flash HyperFlex Nodes Overview 1-6

CHAPTER 2**Installing the Node 2-1**

- Unpacking and Inspecting the Node 2-2
- Preparing for Node Installation 2-3
 - Installation Guidelines 2-3
 - Rack Requirements 2-4
 - Equipment Requirements 2-4
 - Supported Slide Rail Kits 2-4
 - Slide Rail Adjustment Range and Cable Management Arm Dimensions 2-4
- Installing the Node In a Rack 2-5
 - Installing the Slide Rails 2-5
 - Installing the Cable Management Arm (Optional) 2-8
 - Reversing the Cable Management Arm (Optional) 2-9
- Initial Node Setup 2-9

CHAPTER 3**Maintaining the Node 3-1**

- Status LEDs and Buttons 3-1
 - Front Panel LEDs 3-2
 - Rear Panel LEDs and Buttons 3-4
 - Internal Diagnostic LEDs 3-7
- Preparing for Component Installation 3-8
 - Required Equipment 3-8
 - Shutting Down the Node 3-8
 - Shutting Down the Node From the Equipment Tab in Cisco UCS Manager 3-9
 - Shutting Down the Node From the Service Profile in Cisco UCS Manager 3-9
 - Shutting Down the Node Through vSphere With Cisco HX Maintenance Mode 3-10

- Shutting Down the Node with the Node Power Button 3-11
- Decommissioning the Node Using Cisco UCS Manager 3-11
- Post-Maintenance Procedures 3-12
 - Recommissioning the Node Using Cisco UCS Manager 3-12
 - Associating a Service Profile With an HX Node 3-12
 - Exiting HX Maintenance Mode 3-13
- Removing and Replacing the Node Top Cover 3-14
- Serial Number Location 3-15
- Installing or Replacing Node Components 3-16
 - Replaceable Component Locations 3-17
 - Replacing Drives 3-19
 - Drive Population Guidelines 3-19
 - HX240c Drive Configuration Comparison 3-20
 - Drive Replacement Overview 3-21
 - Replacing Persistent Data Drives 3-22
 - Replacing the SSD Caching Drive (Bay 1) 3-23
 - Replacing the Internal Housekeeping SSDs for SDS Logs 3-24
 - Replacing Fan Modules 3-26
 - Replacing DIMMs 3-28
 - Memory Performance Guidelines and Population Rules 3-28
 - DIMM Replacement Procedure 3-31
 - Replacing CPUs and Heatsinks 3-33
 - Special Information For Upgrades to Intel Xeon v4 CPUs 3-33
 - CPU Configuration Rules 3-33
 - Replacing a CPU and Heatsink 3-34
 - Special Instructions for *Upgrades* to Intel Xeon v4 Series 3-37
 - Additional CPU-Related Parts to Order with RMA Replacement Motherboards 3-38
 - Replacing a Cisco Modular HBA Card 3-39
 - Replacing the Motherboard RTC Battery 3-41
 - Replacing an Internal SD Card 3-43
 - Enabling or Disabling the Internal USB Port 3-44
 - Replacing a PCIe Riser 3-45
 - Replacing a PCIe Card 3-47
 - PCIe Slots 3-47
 - Replacing a PCIe Card 3-48
 - Installing Multiple PCIe Cards and Resolving Limited Resources 3-50
 - Installing a Trusted Platform Module 3-52
 - Installing the TPM Hardware 3-52
 - Enabling TPM Support in the BIOS 3-53
 - Enabling the Intel TXT Feature in the BIOS 3-54

Replacing Power Supplies	3-56
Installing a DC Power Supply	3-57
Installing a 930W DC Power Supply, UCSC-PSU-930WDC	3-58
Replacing an mLOM Card (Cisco VIC 1227)	3-59
Replacing an mLOM Card	3-59
Special Considerations for Cisco UCS Virtual Interface Cards	3-61
Service DIP Switches	3-62
DIP Switch Location on the Motherboard	3-62
Using the BIOS Recovery DIP Switch	3-63
Procedure 1: Reboot with recovery.cap (or bios.cap) File	3-63
Procedure 2: Use BIOS Recovery DIP switch and recovery.cap (or bios.cap) File	3-64
Using the Clear Password DIP Switch	3-65
Using the Clear CMOS DIP Switch	3-67
Setting Up the Node in Standalone Mode	3-69
Connecting and Powering On the Node (Standalone Mode)	3-69
Local Connection Procedure	3-69
Remote Connection Procedure	3-70
Cisco IMC Configuration Utility Setup	3-71

APPENDIX A
Node Specifications A-1

Physical Specifications	A-1
Power Specifications	A-2
650 W AC Power Supply	A-2
1200 W AC Power Supply	A-3
1400 W AC Power Supply	A-4
930 W DC (Version 1) Power Supply, UCSC-PSU-930WDC	A-4
Environmental Specifications	A-5

APPENDIX B
Power Cord Specifications B-1

Supported Power Cords and Plugs	B-1
---------------------------------	-----

APPENDIX C
RAID Controller Considerations C-1

Supported HBA and Required Cables	C-1
HBA Firmware Compatibility	C-1
HBA Cabling	C-2
Cable Routing	C-2
Cisco HX240c Server HBA Cabling Instructions	C-3

GPU Card Installation D-1

GPU Card Configuration Rules D-1

Requirement For All Supported GPUs: Memory-Mapped I/O Greater than 4 GB D-2

Installing a GRID K1 or K2 GPU Card D-3

Installing a Tesla M60 GPU Card and 300 W GPU Conversion Kit D-5

Installation Overview D-5

Installing the NVIDIA M60 Hardware D-5

NVIDIA GRID License Server Overview D-13

Registering Your Product Activation Keys With NVIDIA D-14

Downloading the GRID Software Suite D-14

Installing NVIDIA GRID License Server Software D-15

Installing GRID Licenses From the NVIDIA Licensing Portal to the License Server D-17

Managing GRID Licenses D-19

Switching Between Compute Mode and Graphics Mode D-21

Installing Drivers to Support the NVIDIA GPU Cards D-23

1. Updating the Server BIOS Firmware D-23

2. Activating the Server BIOS Firmware D-24

3. Updating the GPU Drivers D-24

Overview

This chapter provides an overview of the Cisco HX240c HyperFlex Node features:

- [Cisco HyperFlex Systems Related Documentation, page 1-1](#)
- [External Features Overview, page 1-1](#)
- [Replaceable Component Locations, page 1-4](#)
- [Summary of Node Features, page 1-5](#)
- [Cisco HX240c All-Flash HyperFlex Nodes Overview, page 1-6](#)

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 [Documentation Roadmap](#).

External Features Overview

The figures in this chapter show an overview of external node features.


- The front-panel features are shown in [Figure 1-1](#).
- The rear panel features are shown in [Figure 1-2](#).

Figure 1-1 Front Panel Features



<p>1 Persistent drive bays:</p> <ul style="list-style-type: none"> • HX240c: bays 2–24 support HDD persistent data drives • HX240c All-Flash: bays 2–11 support SSD persistent data drives. (With Cisco HX Release 2.0, 10 SSDs are supported.) <p>See Replacing Drives, page 3-19 for information about supported drives.</p>	<p>7 Fan status LED</p>
<p>2 Drive bay 1: SSD caching drive</p> <p>The supported SSD differs between the HX240c and HX240c All-Flash nodes. See Replacing Drives, page 3-19.</p>	<p>8 Temperature status LED</p>
<p>3 Operations panel buttons and LEDs</p>	<p>9 Power supply status LED</p>
<p>4 Power button/LED</p>	<p>10 Network link activity LED</p>
<p>5 Unit Identification button/LED</p>	<p>11 Pull-out asset tag</p>
<p>6 Node status LED</p>	<p>12 KVM connector (used with KVM cable that provides two USB 2.0, one VGA, and one serial connector)</p>

Figure 1-2 Rear Panel Features




1	PCIe riser 1 (slots 1, 2, 3*) *Slot 3 is taken by two internal SDS log drives (SATA SSDs) in this node.	7	Serial port (RJ-45 connector)
2	PCIe riser 2 (slots 4, 5, 6)	8	Dual 1-Gb Ethernet ports (LAN1, LAN2)
3	Power supplies (DC power supplies shown)	9	VGA video port (DB-15 connector)
4	Modular LAN-on-motherboard (mLOM) card slot	10	Rear Unit Identification button/LED
5	USB 3.0 ports (two)	11	Grounding-lug holes (for DC power supplies)
6	1-Gb dedicated management port		

Replaceable Component Locations

Figure 1-3 shows the locations of the field-replaceable components. The view shown is from the top down with the top covers and air baffle removed.

Figure 1-3 Replaceable Component Locations



305245

<p>1 Persistent drive bays:</p> <ul style="list-style-type: none"> • HX240c: bays 2–24 support HDD persistent data drives • HX240c All-Flash: bays 2–11 support SSD persistent data drives. (With Cisco HX Release 2.0, 10 SSDs are supported.) <p>See Replacing Drives, page 3-19 for information about supported drives.</p>	<p>9 Trusted platform module (TPM) socket on motherboard, under PCIe riser 2</p>
<p>2 Drive bay 1: SSD caching drive</p> <p>The supported SSD differs between the HX240c and HX240c All-Flash nodes. See Replacing Drives, page 3-19.</p>	<p>10 PCIe riser 2 (PCIe slots 4, 5, 6)</p>
<p>3 Fan modules (six, hot-swappable)</p>	<p>11 PCIe riser 1 (PCIe slots 1, 2, 3*)</p> <p>*Slot 3 is taken by two internal SATA SSD sockets.</p>
<p>4 DIMM sockets on motherboard (24)</p>	<p>12 120 GB internal housekeeping SSDs for SDS logs (two SATA SSDs in PCIe riser 1 sockets)</p>
<p>5 CPUs and heatsinks (two)</p>	<p>13 mLOM card socket on motherboard under PCIe riser 1 for Cisco VIC 1227</p>
<p>6 Cisco SD card slots on motherboard (two)</p>	<p>14 Cisco UCS 12G SAS Modular HBA (dedicated slot and bracket)</p>
<p>7 USB socket on motherboard</p>	<p>15 RTC battery on motherboard</p>
<p>8 Power supplies (hot-swappable)</p>	

Summary of Node Features

Table 1-1 lists a summary of node features.

Table 1-1 Cisco HX240c HyperFlex Node Features

Chassis	Two rack-unit (2RU) chassis.
Processors	Two Intel Xeon E5-2600 v3 or v4 Series processors.
Memory	24 DDR4 DIMM ¹ 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 Ethernet LOM ports, or a Cisco virtual interface card.
Network and management I/O	The node provides these native connectors: <ul style="list-style-type: none"> • One 1-Gb Ethernet dedicated management port • Two 1-Gb BASE-T Ethernet LAN ports • One RS-232 serial port (RJ-45 connector) • One 15-pin VGA² connector • Two USB³ 3.0 connectors • One front-panel KVM connector that is used with the KVM cable, which provides two USB 2.0, one VGA, and one serial (DB-9) connector.
Modular I/O	A dedicated socket can be used to add an mLOM card for additional rear-panel connectivity.
WoL	1-Gb BASE-T Ethernet LAN ports support the wake-on-LAN (WoL) standard.
Power	Two power supplies: <ul style="list-style-type: none"> • AC power supplies optionally 650 W AC, 1200 W AC, or 1400 W AC each. • DC power supplies 930 W DC each. Do not mix power supply types or wattages in the node. Redundant as 1+1.
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	Five horizontal PCIe ⁴ expansion slots on two risers. Riser 1 contains PCIe slots 1 and 2, plus one internal SATA SSD. Riser 2 contains slots 4, 5, and 6.
InfiniBand	The InfiniBand architecture is supported by the bus slots.

Table 1-1 Cisco HX240c HyperFlex Node Features (continued)

Storage	The following storage disks: <ul style="list-style-type: none"> • One SSD cache drive in front bay 1. • Persistent data drives: <ul style="list-style-type: none"> – HX240c: Up to 23 HDD persistent data drives in bays 2–24. – HX240c All-Flash: Up to 10 SSD persistent data drives in bays 2–11. (With Cisco HX Release 2.0, 10 SSDs are supported.) • Two internal SATA SSDs for SDS logs in PCIe riser 1.
Internal USB	One internal USB 3.0 port on the motherboard that you can use with a USB thumb drive for additional storage.
SD cards	Two internal bays on the motherboard for up to two SD cards.
Disk Management	One Cisco UCS 12G SAS Modular HBA.
Native Video	VGA video resolution up to 1920 x 1200, 16 bpp at 60 Hz, and up to 256 MB of video memory.

1. DIMM = dual inline memory module
2. VGA = video graphics array
3. USB = universal serial bus
4. PCIe = peripheral component interconnect express

Cisco HX240c All-Flash HyperFlex Nodes Overview

The HX240c All-Flash HyperFlex node contains all SSDs, rather than the hybrid mix of SSDs and HDDs that is used in the HX240c HyperFlex node. Enterprise value SSDs are used for the persistent data drives. Enterprise high-endurance SSDs are used for the caching drives.

Note the following considerations and restrictions:

- The minimum Cisco HyperFlex software required is Release 2.0 or later.
- With Cisco HX Release 2.0, only 10 SSD persistent data drives are supported.)
- 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 HX240c All-Flash configuration is not supported.
- Mixing hybrid HX240c HyperFlex nodes with HX240c All-Flash HyperFlex nodes within the same HyperFlex cluster is not supported.

See [HX240c Drive Configuration Comparison, page 3-20](#) for specifics about drive PIDs supported in the node types.

Installing the Node

This chapter describes how to install the node, and it includes the following sections:

- [Unpacking and Inspecting the Node, page 2-2](#)
- [Preparing for Node Installation, page 2-3](#)
- [Installing the Node In a Rack, page 2-5](#)
- [Initial Node Setup, page 2-9](#)



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

Unpacking and Inspecting the Node

**Caution**

When handling internal node components, wear an ESD strap and handle modules by the carrier edges only.

**Tip**

Keep the shipping container in case the node requires shipping in the future.

**Note**

The chassis is thoroughly inspected before shipment. If any damage occurred during transportation or any items are missing, contact your customer service representative immediately.

-
- Step 1** Remove the node from its cardboard container and save all packaging material.
- Step 2** Compare the shipment to the equipment list provided by your customer service representative. Verify that you have all items.
- Step 3** Check for damage and report any discrepancies or damage to your customer service representative. Have the following information ready:
- Invoice number of shipper (see the packing slip)
 - Model and serial number of the damaged unit
 - Description of damage
 - Effect of damage on the installation
-

Preparing for Node Installation

This section provides information about preparing for node installation, and it includes the following topics:

- [Installation Guidelines, page 2-3](#)
- [Rack Requirements, page 2-4](#)
- [Equipment Requirements, page 2-4](#)
- [Slide Rail Adjustment Range and Cable Management Arm Dimensions, page 2-4](#)

Installation Guidelines



To prevent the node from overheating, do not operate it in an area that exceeds the maximum recommended ambient temperature of: 35° C (95° F).
Statement 1047



The plug-socket combination must be accessible at all times, because it serves as the main disconnecting device.
Statement 1019



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



Installation of the equipment must comply with local and national electrical codes.
Statement 1074



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.

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

- Ensure that there is adequate space around the node to allow for servicing 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 [Node Specifications, page A-1](#).
- Ensure that the cabinet or rack meets the requirements listed in the “[Rack Requirements](#)” section on [page 2-4](#).
- Ensure that the site power meets the power requirements listed in the [Node Specifications, page A-1](#). If available, you can use an uninterruptible power supply (UPS) to protect against power failures.

Rack Requirements

This section provides the requirements for the standard open racks.

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 supplied slide rails.
- The minimum vertical rack space per node must be two RUs, equal to 3.5 in. (88.9 mm).

Equipment Requirements

The slide rails supplied by Cisco for this node do not require tools for installation if you install them in a rack that has square 0.38-inch (9.6 mm), round 0.28-inch (7.1 mm), or #12-24 UNC threaded holes.

Supported Slide Rail Kits

This node supports one rail kit option: Cisco part UCSC-RAILB-M4= (ball-bearing rail kit).

Slide Rail Adjustment Range and Cable Management Arm Dimensions

The slide rails for this node have an adjustment range of 26 to 36 inches (660 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 34.4 inches (874 mm).

Installing the Node In a Rack

This section contains the following sections:

- [Installing the Slide Rails, page 2-5](#)
- [Installing the Cable Management Arm \(Optional\), page 2-8](#)
- [Reversing the Cable Management Arm \(Optional\), page 2-9](#)

Installing the Slide Rails

This section describes how to install the node in a rack using the rack kits that are sold by Cisco.



Warning

To prevent bodily injury when mounting or servicing this unit in a rack, you must take special precautions to ensure that the node 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:

- 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 (see [Figure 2-1](#)).
- 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.
- Install the second inner rail to the opposite side of the node.


Figure 2-1 Attaching Inner Rail to Side of Node



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

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

Figure 2-2 Front Securing Mechanism, Inside of Front End



1	Front mounting pegs	3	Securing plate shown pulled back to open position
2	Rack post		

Step 3 Install the slide rails into the rack:

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

The slide rail front-end wraps around the outside of the rack post and the mounting pegs enter the rack-post holes from the outside-front (see [Figure 2-2](#)).



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 with each other 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 67 pounds (59 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 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 release clip toward the rear on both inner rails ([Figure 2-3](#)), and then continue pushing the node into the rack until its front slam latches engage with the rack posts.

Figure 2-3 Inner Rail Release Clip



1	Inner rail release clip	3	Outer rail attached to rack post
2	Inner rail attached to node		

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 the 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 CMA is reversible left to right. To reverse the CMA, see [Reversing the Cable Management Arm \(Optional\)](#), page 2-9 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 (see [Figure 2-4](#)). Slide the tab over the end of the rail until it clicks and locks.
- 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 (see [Figure 2-4](#)). 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 (see [Figure 2-4](#)).
- 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 (see [Figure 2-4](#)). 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.

Figure 2-4 Attaching the Cable Management Arm to the Rear of the Slide Rails




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

Reversing the Cable Management Arm (Optional)

- Step 1** Rotate the entire CMA assembly 180 degrees. The plastic cable guides must remain pointing upward.
- Step 2** Flip the tabs at the end of each CMA arm 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 2-5 Reversing the CMA



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

Initial Node Setup

Refer to the following documentation:

- For instructions on physically installing the nodes in racks, see [Installing the Node In a Rack, page 2-5](#)
- For installation instructions on configuring system fabric interconnects and connecting the HX Series nodes to the fabrics, refer to the [Cisco HyperFlex Systems Getting Started Guide](#).
- For instructions on adding the HX240c node as a compute-only node, see the “Adding a Compute Node” section in the [Cisco HyperFlex Systems Getting Started Guide](#).
- For instructions on starting the node in standalone mode for troubleshooting purposes, see [Setting Up the Node in Standalone Mode, page 3-68](#).

Maintaining the Node

This chapter describes how to diagnose node problems using LEDs. It also provides information about how to install or replace hardware components, and it includes the following sections:

- [Status LEDs and Buttons, page 3-1](#)
- [Preparing for Component Installation, page 3-8](#)
- [Installing or Replacing Node Components, page 3-16](#)
- [Service DIP Switches, page 3-61](#)

Status LEDs and Buttons


This section describes the location and meaning of LEDs and buttons and includes the following topics

- [Front Panel LEDs, page 3-2](#)
- [Rear Panel LEDs and Buttons, page 3-4](#)
- [Internal Diagnostic LEDs, page 3-7](#)

Front Panel LEDs

Figure 3-1 shows the front panel LEDs. Table 3-1 defines the LED states.

Figure 3-1 Front Panel LEDs



1	Drive fault LED (on each drive tray)	6	Fan status LED
2	Drive activity LED (on each drive tray)	7	Temperature status LED
3	Power button/power status LED	8	Power supply status LED
4	Unit Identification button/LED	9	Network link activity LED
5	Node status LED		

Table 3-1 Front Panel LEDs, Definitions of States

	LED Name	State
1	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.
2	Drive activity	<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 node. Amber—The node is in standby power mode. Power is supplied only to the Cisco IMC and some motherboard functions. Green—The node is in main power mode. Power is supplied to all node components.
4	Unit Identification	<ul style="list-style-type: none"> Off—The unit identification function is not in use. Blue—The unit identification function is activated.


Table 3-1 Front Panel LEDs, Definitions of States (continued)

	LED Name	State
5	Node status	<ul style="list-style-type: none"> • Green—The node is running in a normal operating condition. • Green, blinking—The node is performing node initialization and memory check. • Amber, steady—The node is in a degraded operational state. 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, blinking—The node is in a critical fault state. For example: <ul style="list-style-type: none"> – Boot failed. – Fatal CPU and/or bus error is detected. – Node is in an over-temperature condition.
6	Fan status	<ul style="list-style-type: none"> • Green—All fan modules are operating properly. • Amber, steady—One or more fan modules breached the critical threshold. • Amber, blinking—One or more fan modules breached the non-recoverable threshold.
7	Temperature status	<ul style="list-style-type: none"> • Green—The node 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.
8	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.
9	Network link activity	<ul style="list-style-type: none"> • Off—The Ethernet 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.

Rear Panel LEDs and Buttons

Figure 3-2 shows the rear panel LEDs and buttons. Table 3-2 defines the LED states.

Figure 3-2 Rear Panel LEDs and Buttons



1	Power supply fault LED	5	1-Gb Ethernet dedicated management link status LED
2	Power supply status LED	6	1-Gb Ethernet link speed LED
3	mLOM card LEDs (Cisco VIC 1227) (not shown, see Table 3-2)	7	1-Gb Ethernet link status LED
4	1-Gb Ethernet dedicated management link speed LED	8	Unit Identification button/LED

Table 3-2 Rear Panel LEDs, Definitions of States

LED Name	State
1 Power supply fault This is a summary; for advanced power supply LED information, see Table 3-3.	<ul style="list-style-type: none"> Off—The power supply is operating normally. Amber, blinking—An event warning threshold has been reached, but the power supply continues to operate. Amber, solid—A critical fault threshold has been reached, causing the power supply to shut down (for example, a fan failure or an over-temperature condition).
2 Power supply status This is a summary; for advanced power supply LED information, see Table 3-3.	<p>AC power supplies:</p> <ul style="list-style-type: none"> Off—There is no AC power to the power supply. Green, blinking—AC power OK; DC output not enabled. Green, solid—AC power OK; DC outputs OK. <p>DC power supplies:</p> <ul style="list-style-type: none"> Off—There is no DC power to the power supply. Green, blinking—DC power OK; DC output not enabled. Green, solid—DC power OK; DC outputs OK.

Table 3-2 Rear Panel LEDs, Definitions of States (continued)

	LED Name	State
3	mLOM 10-Gb SFP+ (Cisco VIC 1227) (single status LED)	<ul style="list-style-type: none"> • Off—No link is present. • Green, steady—Link is active. • Green, blinking—Traffic is present on the active link.
4	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.
5	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.
6	1-Gb Ethernet 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.
7	1-Gb Ethernet 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.
8	Unit Identification	<ul style="list-style-type: none"> • Off—The unit identification function is not in use. • Blue—The unit identification function is activated.

In [Table 3-3](#), read the status and fault LED states together in each row to determine the event that cause this combination.

Table 3-3 Rear Power Supply LED States

Green PSU Status LED State	Amber PSU Fault LED State	Event
• Solid on	• Off	12V main on (main power mode)
• Blinking	• Off	12V main off (standby power mode)
• Off	• Off	No AC power input (all PSUs present)
• Off	• On	No AC power input (redundant supply active)
• Blinking	• Solid on	12V over-voltage protection (OVP)
• Blinking	• Solid on	12V under-voltage protection (UVP)
• Blinking	• Solid on	12V over-current protection (OCP)
• Blinking	• Solid on	12V short-circuit protection (SCP)
• Solid on	• Solid on	PSU fan fault/Lock (before OTP)
• Blinking	• Solid on	PSU fan fault/Lock (after OTP)
• Blinking	• Solid on	Over-temperature protection (OTP)
• Solid on	• Blinking	OTP warning
• Solid on	• Blinking	OCP warning
• Blinking	• Off	12V main off (CR slave PSU is in sleep mode)


Internal Diagnostic LEDs

The node is equipped with a supercap voltage source that can activate internal component fault LEDs up to 30 minutes after AC power is removed. The node has internal fault LEDs for CPUs, DIMMs, fan modules, SD cards, the RTC battery, and the mLOM card.

To use these LEDs to identify a failed component, press the front or rear Unit Identification button (see [Figure 3-1](#) or [Figure 3-2](#)) with AC power removed. An LED lights amber to indicate a faulty component.

See [Figure 3-3](#) for the locations of these internal LEDs.

Figure 3-3 Internal Diagnostic LED Locations



1	Fan module fault LEDs (one on each fan module)	4	SD card fault LEDs
2	DIMM fault LEDs (one directly in front of each DIMM socket on the motherboard)	5	RTC battery fault LED (under PCIe riser 1)
3	CPU fault LEDs	6	mLOM card fault LED (under PCIe riser 1)

Table 3-4 Internal Diagnostic LEDs, Definition of States

LED Name	State
Internal diagnostic LEDs (all)	<ul style="list-style-type: none"> Off—Component is functioning normally. Amber—Component has a fault.

Preparing for Component Installation

This section describes how to prepare for component installation, and it includes the following topics:

- [Required Equipment, page 3-8](#)
- [Shutting Down the Node, page 3-8](#)
- [Decommissioning the Node Using Cisco UCS Manager, page 3-11](#)
- [Post-Maintenance Procedures, page 3-12](#)
- [Removing and Replacing the Node Top Cover, page 3-14](#)
- [Serial Number Location, page 3-15](#)

Required Equipment

The following equipment is used to perform the procedures in this chapter:

- Number 2 Phillips-head screwdriver
- Electrostatic discharge (ESD) strap or other grounding equipment such as a grounded mat

Shutting Down the Node

The node can run in 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 the cooling fans and it is safe to power off the node from 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.

This section contains the following procedures, which are referenced from component replacement procedures. Alternate shutdown procedures are included.

- [Shutting Down the Node From the Equipment Tab in Cisco UCS Manager, page 3-9](#)
- [Shutting Down the Node From the Service Profile in Cisco UCS Manager, page 3-9](#)
- [Shutting Down the Node Through vSphere With Cisco HX Maintenance Mode, page 3-10](#)
- [Shutting Down the Node with the Node Power Button, page 3-11](#)

Shutting Down the Node From the Equipment Tab in Cisco UCS Manager

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.

Shutting Down the Node From the Service Profile in Cisco UCS Manager

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 server node you are shutting down.
 - Step 4** Choose the service profile of the server 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.

Shutting Down the Node Through vSphere With Cisco 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. Move the node into HX Maintenance Mode.

1. Identify the node ID and IP address:

```
# stcli node list --summary
```

2. Enter the node into HX Maintenance Mode.

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

(see also `stcli node maintenanceMode --help`)

- c. Log into the ESXi command line of this node as a user with root privileges.
- d. 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 the Node, page 3-8](#).

Shutting Down the Node with the Node 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 (see the “[Front Panel LEDs](#)” section on page 3-2).

- Green—The node is in main power mode and must be shut down before it can be safely powered off. Go to [Step 2](#).
- Amber—The node is already in standby mode and can be safely powered off. Go to .

Step 2

Invoke either a graceful shutdown or a hard shutdown:

**Caution**

To avoid data loss or damage to your operating node, 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 Status LED.
- Emergency shutdown—Press and hold the **Power** button for 4 seconds to force the main power off and immediately enter standby mode.

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.

Post-Maintenance Procedures

This section contains the following procedures, which are referenced from component replacement procedures:

- [Recommissioning the Node Using Cisco UCS Manager, page 3-12](#)
- [Associating a Service Profile With an HX Node, page 3-12](#)
- [Exiting HX Maintenance Mode, page 3-13](#)

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** Under **Equipment**, click the **Rack Mounts** node.
 - Step 3** In the **Work** pane, click the **Decommissioned** tab.
 - Step 4** On the row for each rack-mount server 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 server.
-

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 associating the service profile.
 - Step 7** Click **OK**.
-

Exiting HX Maintenance Mode

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

-
- Step 1** 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:**
 - a. Log in to the storage controller cluster command line as a user with root privileges.
 - b. Exit the node out of HX Maintenance Mode.
 1. Identify the node ID and IP address:

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

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

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

```
# esxcli system maintenanceMode get
```
-

Removing and Replacing the Node Top Cover

- Step 1** Remove the top cover (see [Figure 3-4](#)).
- If the cover latch is locked, use a screwdriver to turn the lock 90-degrees counterclockwise to unlock it. See [Figure 3-4](#).
 - 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.
 - Lift the top cover straight up from the node and set it aside.


- Step 2** Replace the top cover:



Note The latch must be in the fully open position when you set the cover back in place, which allows the opening in the latch to sit over a peg that is on the fan tray.

- 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.
- Press the cover latch down to the closed position. The cover is pushed forward to the closed position as you push down the latch.
- If desired, lock the latch by using a screwdriver to turn the lock 90-degrees clockwise.

Figure 3-4 Removing the Top Cover



1	Front cover panel	3	Locking cover latch
2	Top cover		

Serial Number Location

The serial number (SN) for the node is printed on a label on the top of the node, near the front.

Installing or Replacing Node 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 node unless all cards, faceplates, front covers, and rear covers are in place.

Statement 1029

**Caution**

When handling node components, wear an ESD strap to avoid damage.

**Tip**

You can press the Unit Identification button on the front panel or rear panel to turn on a flashing Unit Identification LED on the front and rear panels of the node. This button allows you to locate the specific node that you are servicing when you go to the opposite side of the rack. You can also activate these LEDs remotely by using the Cisco IMC interface. See the “[Status LEDs and Buttons](#)” section on [page 3-1](#) for locations of these LEDs.


This section describes how to install and replace node components, and it includes the following topics:

- [Replaceable Component Locations](#), page 3-17
- [Replacing Drives](#), page 3-18
- [Replacing the Internal Housekeeping SSDs for SDS Logs](#), page 3-23
- [Replacing Fan Modules](#), page 3-25
- [Replacing DIMMs](#), page 3-27
- [Replacing CPUs and Heatsinks](#), page 3-32
- [Replacing a Cisco Modular HBA Card](#), page 3-38
- [Replacing the Motherboard RTC Battery](#), page 3-40
- [Replacing an Internal SD Card](#), page 3-42
- [Enabling or Disabling the Internal USB Port](#), page 3-43
- [Replacing a PCIe Riser](#), page 3-44
- [Replacing a PCIe Card](#), page 3-46
- [Installing a Trusted Platform Module](#), page 3-51
- [Replacing Power Supplies](#), page 3-55
- [Replacing an mLOM Card \(Cisco VIC 1227\)](#), page 3-58

Replaceable Component Locations

Figure 3-5 shows the locations of the field-replaceable components. The view shown is from the top down with the top covers and air baffle removed.

Figure 3-5 Replaceable Component Locations



305245

<p>1 Persistent drive bays:</p> <ul style="list-style-type: none"> • HX240c: bays 2–24 support HDD persistent data drives • HX240c All-Flash: bays 2–11 support SSD persistent data drives. (With Cisco HX Release 2.0, 10 SSDs are supported.) <p>See Replacing Drives, page 3-18 for information about supported drives.</p>	<p>9 Trusted platform module (TPM) socket on motherboard, under PCIe riser 2</p>
<p>2 Drive bay 1: SSD caching drive</p> <p>The supported SSD differs between the HX240c and HX240c All-Flash nodes. See Replacing Drives, page 3-18.</p>	<p>10 PCIe riser 2 (PCIe slots 4, 5, 6)</p>
<p>3 Fan modules (six, hot-swappable)</p>	<p>11 PCIe riser 1 (PCIe slots 1, 2, 3*)</p> <p>*Slot 3 is taken by two internal SATA SSD sockets.</p>
<p>4 DIMM sockets on motherboard (24)</p>	<p>12 120 GB internal housekeeping SSDs for SDS logs (two SATA SSDs in PCIe riser 1 sockets)</p>
<p>5 CPUs and heatsinks (two)</p>	<p>13 mLOM card socket on motherboard under PCIe riser 1 for Cisco VIC 1227</p>
<p>6 Cisco SD card slots on motherboard (two)</p>	<p>14 Cisco UCS 12G SAS Modular HBA (dedicated slot and bracket)</p>
<p>7 USB socket on motherboard</p>	<p>15 RTC battery on motherboard</p>
<p>8 Power supplies (hot-swappable, accessed through rear panel)</p>	

Replacing Drives


This section includes the following information:

- [Drive Population Guidelines](#), page 3-18
- [HX240c Drive Configuration Comparison](#), page 3-19
- [Drive Replacement Overview](#), page 3-20
- [Replacing Persistent Data Drives](#), page 3-21
- [Replacing the SSD Caching Drive \(Bay 1\)](#), page 3-22

Drive Population Guidelines

The drive-bay numbering is shown in [Figure 3-6](#).

Figure 3-6 Drive Bay Numbering



1	Bay 1: solid state drive (SSD) caching drive	2	Bays 2 - 24: persistent data drives
---	--	---	-------------------------------------

Observe these drive population guidelines:

- Populate the SSD caching drive only in bay 1. See [Table 3-5](#) for the supported caching SSDs, which differ between supported drive configurations.
- Populate persistent data drives as follows:
 - HX240c: HDD persistent data drives—populate in bays 2–24.
 - HX240c All-Flash: SSD persistent data drives—populate in bays 2–11. With Cisco HyperFlex Release 2.0, only 10 SSD persistent data drives are supported.

See [Table 3-5](#) for the supported persistent drives, which differ between supported drive configurations.

- When populating persistent data drives, add drives in the lowest numbered bays first.
- Keep an empty drive blanking tray in any unused bays to ensure optimal airflow and cooling.
- See [HX240c Drive Configuration Comparison](#), page 3-19 for a comparison of supported drive configurations (hybrid, All-Flash, SED).

HX240c Drive Configuration Comparison

Table 3-5 Supported Drive Configurations

Component	HX240c Hybrid Drives	HX240c All-Flash Drives	HX240c SED Hybrid Drives	HX240c SED All-Flash Drives
Persistent data drives HX240c: Front bays 2–24 HX240c All-Flash: Front bays 2–11	HDD: • UCS-HD12TB10K12G	SSD: • UCS-SD960GBKS4-EV	HDD (SED): • UCS-HD12G10K9	HDD (SED): • UCS-SD800GBEK9
SSD caching SSD Front bay 1	SSD: • UCS-SD16TB12S3-EP	SSD: • UCS-SD16TB12S3-EP	SSD (SED): • UCS-SD16TBK9	SSD (SED): • UCS-SD800GBEK9
Housekeeping SSD for SDS logs Internal SATA Boot SSD on PCIe riser 1	SSD: • UCS-SD120GBKS4-EV	SSD: • UCS-SD120GBKS4-EV	SSD: • UCS-SD120GBKS4-EV	SSD: • UCS-SD120GBKS4-EV

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

- The minimum Cisco HyperFlex software required is Release 2.0 or later.
- With Cisco HX Release 2.0, only 10 SSD persistent data drives are supported.
- 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 HX240c All-Flash configuration is not supported.
- Mixing hybrid HX240c HyperFlex nodes with HX240c All-Flash HyperFlex nodes within the same HyperFlex cluster is not supported.

Note the following considerations and restrictions for self-encrypting drive (SED) HyperFlex nodes:

- The minimum Cisco HyperFlex software required for SED configurations is Release 2.1(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.

Table 3-6 Drive Replacement Overview

Persistent data drives	Hot-swap replacement. See Replacing Persistent Data Drives, page 3-21 .
HX240c: Front bays 2–24	NOTE: Hot-swap replacement includes hot-removal, so you can remove the drive while it is still operating.
HX240c All-Flash: Front bays 2–11	
SSD caching SSD front bay 1	Hot-swap replacement. See Replacing the SSD Caching Drive (Bay 1), page 3-22 . 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 drive, additional steps are required, as described in the procedure.
Housekeeping SSD for SDS logs (internal SATA SSD on PCIe riser 1)	Node must be put into Cisco HX Maintenance Mode before shutting down and powering off the node. Replacement requires additional technical assistance and cannot be completed by the customer. See Replacing the Internal Housekeeping SSDs for SDS Logs, page 3-23 .

Replacing Persistent Data Drives

The persistent data drives must be installed only as follows:


- HX240c: HDDs in bays 2–24
- HX240c All-Flash: SSDs in bays 2–11. With Cisco HyperFlex Release 2.0, only 10 SSD persistent data drives are supported.



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 an empty bay:
- Press the release button on the face of the drive tray. See [Figure 3-7](#).
 - Grasp and open the ejector lever and then pull the drive tray out of the slot.
 - If you are replacing an existing drive, remove the four drive-tray screws that secure the drive to the tray and then lift the drive out of the tray.+
- Step 2** Install a new drive:
- Place a new drive in the empty drive tray and replace the four drive-tray screws.
 - With the ejector lever on the drive tray open, insert the drive tray into the empty drive bay.
 - Push the tray into the slot until it touches the backplane, and then close the ejector lever to lock the drive in place.

Figure 3-7 Replacing Drives



1	Release button	3	Drive tray securing screws (four)
2	Ejector lever		–

Replacing the SSD Caching Drive (Bay 1)

The SSD caching drive must be installed in drive bay 1 (see [Figure 3-6](#)).

Note the following considerations and restrictions for NVMe SSDs When Used As the Caching SSD:

- NVMe SSDs are supported in HX240c and All-Flash nodes.
- NVMe SSDs are not supported in Hybrid nodes.
- NVMe SSDs are supported only in the Caching SSD position, in drive bay 1.
- 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 drive, hot-swap replacement includes hot-removal, so you can remove a SAS 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. Otherwise, skip to step 2.
- Step 2** Remove the SSD caching drive:
- a. Press the release button on the face of the drive tray (see [Figure 3-7](#)).
 - 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 SSD to the tray and then lift the SSD out of the tray.
- Step 3** Install a new SSD caching drive:
- a. Place a new SSD in the empty drive tray and replace 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 drive 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.
-

Replacing the Internal Housekeeping SSDs for SDS Logs

PCIe riser 1 includes sockets for two 120 GB SATA SSDs.

**Note**

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

**Caution**

Put the HX node in Cisco HX Maintenance mode before replacing the housekeeping SSD drive, as described in the procedure. Hot swapping the internal housekeeping/logs SSD while the HX node is running causes the HX node to fail.

- Step 1** Put the node in Cisco HX Maintenance mode as described in [Shutting Down the Node Through vSphere With Cisco HX Maintenance Mode](#), page 3-10.
- Step 2** Shut down the node as described in [Shutting Down the Node](#), page 3-8.
- Step 3** Decommission the node as described in [Decommissioning the Node Using Cisco UCS Manager](#), page 3-11.

**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 the power supplies.
- Step 5** Do one of the following:
- If the drive is a SED drive in front drive bay 1, remove it from front drive bay 1. Continue with [Step 6](#).
 - If the SSD is in PCIe riser 1, remove it from PCIe riser 1. Continue with [Step 8](#).
- Step 6** Remove the drive that you are replacing or remove a blank drive tray from an empty bay:
- a. Press the release button on the face of the drive tray. See [Figure 3-7](#).
 - 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 7** Install a new drive:
- a. Place a new drive in the empty drive tray and replace 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.
 - d. Go to [Step 16](#).
- Step 8** 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.
- Step 9** Remove the top cover as described in [Removing and Replacing the Node Top Cover](#), page 3-14.
- Step 10** Remove PCIe riser 1 from the node:
- a. Lift straight up on both ends of the riser to disengage its circuit board from the socket on the motherboard. Set the riser on an antistatic mat.

- b. On the bottom of the riser, loosen the single thumbscrew that holds the securing plate. See [Figure 3-18](#).
- c. Swing open the securing plate and remove it from the riser to provide access.

Step 11 Remove an existing SSD from PCIe riser 1.

Grasp the carrier-tabs on each side of the SSD and pinch them together as you pull the SSD from its cage and the socket on the PCIe riser.

Step 12 Install a new SSD to PCIe riser 1:

- a. Grasp the two carrier-tabs on either side of the SSD and pinch them together as you insert the SSD into the cage on the riser.
- b. Push the SSD straight into the cage to engage it with the socket on the riser. Stop when the carrier-tabs click and lock into place on the cage.

Step 13 Return PCIe riser 1 to the node:

- a. Return the securing plate to the riser. Insert the two hinge-tabs into the two slots on the riser, and then swing the securing plate closed.
- b. Tighten the single thumbscrew that holds the securing plate.
- c. Position the PCIe riser over its socket on the motherboard and over its alignment features in the chassis (see [Figure 3-16](#)).
- d. Carefully push down on both ends of the PCIe riser to fully engage its circuit board connector with the socket on the motherboard.

Step 14 Replace the top cover.

Step 15 Replace the node in the rack, replace power cables, and then power on the node by pressing the **Power** button.

Step 16 Recommission the node as described in [Recommissioning the Node Using Cisco UCS Manager](#), page 3-12.

Step 17 Associate the node to its service profile as described in [Associating a Service Profile With an HX Node](#), page 3-12.

Step 18 After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode](#), page 3-13.



Note

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

Replacing Fan Modules

The six hot-swappable fan modules in the node are numbered as follows when you are facing the front of the node.

Figure 3-8 Fan Module Numbering

FAN 6	FAN 5	FAN 4	FAN 3	FAN 2	FAN 1
-------	-------	-------	-------	-------	-------



Tip

A fault LED is on the top of each fan module that lights amber if the fan module fails. To operate these LEDs from the supercap power source, remove AC power cords and then press the Unit Identification button. See also [Internal Diagnostic LEDs, page 3-7](#).



Caution

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

Step 1

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 2

Remove the top cover as described in [Removing and Replacing the Node Top Cover, page 3-14](#).

Step 3

Identify a faulty fan module by looking for a fan fault LED that is lit amber (see [Figure 3-9](#)).

Step 4

Remove a fan module that you are replacing (see [Figure 3-9](#)):

- a. Grasp the top of the fan and pinch the green plastic latch toward the center.
- b. Lift straight up to remove the fan module from the node.

Step 5

Install a new fan module:

- a. Set the new fan module in place, aligning the connector on the bottom of the fan module with the connector on the motherboard.



Note The arrow label on the top of the fan module, which indicates the direction of airflow, should point toward the rear of the node.

- b. Press down gently on the fan module until the latch clicks and locks in place.


Step 6

Replace the top cover.

Step 7

Replace the node in the rack.

Figure 3-9 Fan Modules Latch and Fault LED



652959

1	Finger latch (on each fan module)	2	Fan module fault LED (on each fan module)
---	-----------------------------------	---	---

Replacing DIMMs

This section includes the following topics:

- [Memory Performance Guidelines and Population Rules, page 3-27](#)
- [DIMM Replacement Procedure, page 3-30](#)



Caution

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



Caution

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



Note

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

Memory Performance Guidelines and Population Rules


This section describes the type of memory that the node requires and its effect on performance. The section includes the following topics:

- [DIMM Socket Numbering, page 3-28](#)
- [DIMM Population Rules, page 3-28](#)
- [Memory Mirroring and RAS, page 3-29](#)
- [Lockstep Channel Mode, page 3-29](#)

DIMM Socket Numbering

Figure 3-10 shows the numbering of the DIMM sockets and CPUs.

Figure 3-10 CPUs and DIMM Socket Numbering on Motherboard



352815

DIMM Population Rules

Observe the following guidelines when installing or replacing DIMMs:

- Each CPU supports four memory channels.
 - CPU1 supports channels A, B, C, and D.
 - CPU2 supports channels E, F, G, and H.
- Each channel has three DIMM sockets (for example, channel A = slots A1, A2, and A3).
 - A channel can operate with one, two, or three DIMMs installed.
 - If a channel has only one DIMM, populate slot 1 first (the blue slot).
- When both CPUs are installed, populate the DIMM sockets of each CPU identically.
 - Fill blue #1 slots in the channels first: A1, E1, B1, F1, C1, G1, D1, H1
 - Fill black #2 slots in the channels second: A2, E2, B2, F2, C2, G2, D2, H2
 - Fill white #3 slots in the channels third: A3, E3, B3, F3, C3, G3, D3, H3
- Any DIMM installed in a DIMM socket for which the CPU is absent is not recognized. In a single-CPU configuration, populate the channels for CPU1 only (A, B, C, D).
- Memory mirroring reduces the amount of memory available by 50 percent because only one of the two populated channels provides data. When memory mirroring is enabled, you must install DIMMs in sets of 4, 6, 8, or 12 as described in [Memory Mirroring and RAS, page 3-29](#).
- Observe the DIMM mixing rules shown in [Table 3-7](#).

Table 3-7 DIMM Mixing Rules for HX240c Nodes

DIMM Parameter	DIMMs in the Same Channel	DIMMs in the Same Bank
DIMM Capacity: RDIMM = 8 or 16 GB LRDIMM = 64 GB	<ul style="list-style-type: none"> You can mix different capacity DIMMs in the same channel (for example, A1, A2, A3). 	<ul style="list-style-type: none"> You can mix different capacity DIMMs in the same bank. However, for optimal performance DIMMs in the same bank (for example, A1, B1, C1, D1) should have the same capacity.

Table 3-7 DIMM Mixing Rules for HX240c Nodes

DIMM Parameter	DIMMs in the Same Channel	DIMMs in the Same Bank
DIMM Speed: 2133 or 2400 MHz	You can mix speeds, but DIMMs will run at the speed of the slowest DIMMs/CPUs installed in the channel.	You can mix speeds, but DIMMs will run at the speed of the slowest DIMMs/CPUs installed in the bank.
DIMM Type: RDIMMs or LRDIMMs	You cannot mix DIMM types in a channel.	You cannot mix DIMM types in a bank.

Memory Mirroring and RAS

The Intel E5-2600 CPUs within the node support memory mirroring only when an even number of channels are populated with DIMMs. If one or three channels are populated with DIMMs, memory mirroring is automatically disabled. Furthermore, if memory mirroring is used, DRAM size is reduced by 50 percent for reasons of reliability.

Lockstep Channel Mode

When you enable lockstep channel mode, each memory access is a 128-bit data access that spans four channels.

Lockstep channel mode requires that all four memory channels on a CPU must be populated identically with regard to size and organization. DIMM socket populations within a channel (for example, A1, A2, A3) do not have to be identical but the same DIMM slot location across all four channels must be populated the same.

For example, DIMMs in sockets A1, B1, C1, and D1 must be identical. DIMMs in sockets A2, B2, C2, and D2 must be identical. However, the A1-B1-C1-D1 DIMMs do not have to be identical with the A2-B2-C2-D2 DIMMs.

DIMM Replacement Procedure

This section includes the following topics:

- [Identifying a Faulty DIMM, page 3-30](#)
- [Replacing DIMMs, page 3-30](#)

Identifying a Faulty DIMM

Each DIMM socket has a corresponding DIMM fault LED, directly in front of the DIMM socket. See [Figure 3-3](#) for the locations of these LEDs. The LEDs light amber to indicate a faulty DIMM. To operate these LEDs from the SuperCap power source, remove AC power cords and then press the Unit Identification button.

Replacing DIMMs

-
- Step 1** Put the node in Cisco HX Maintenance mode as described in [Shutting Down the Node Through vSphere With Cisco HX Maintenance Mode, page 3-10](#).
- Step 2** Shut down the node as described in [Shutting Down the Node, page 3-8](#).
- Step 3** Decommission the node as described in [Decommissioning the Node Using Cisco UCS Manager, page 3-11](#).



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 the 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 as described in [Removing and Replacing the Node Top Cover, page 3-14](#).
- Step 7** Remove the air baffle that sits over the DIMM sockets and set it aside.
- Step 8** Identify the faulty DIMM by observing the DIMM socket fault LEDs on the motherboard (see [Figure 3-3](#)).
- Step 9** Remove the DIMMs that you are replacing. Open the ejector levers at both ends of the DIMM socket, and then lift the DIMM out of the socket.
- Step 10** Install a new DIMM:



Note

Before installing DIMMs, see the population guidelines. See [Memory Performance Guidelines and Population Rules, page 3-27](#).

- Align the new DIMM with the empty socket on the motherboard. Use the alignment key in the DIMM socket 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.

- Step 11 Replace the air baffle.
 - Step 12 Replace the top cover.
 - Step 13 Replace the node in the rack, replace power cables, and then power on the node by pressing the **Power** button.
 - Step 14 Recommission the node as described in [Recommissioning the Node Using Cisco UCS Manager](#), page 3-12.
 - Step 15 Associate the node to its service profile as described in [Associating a Service Profile With an HX Node](#), page 3-12.
 - Step 16 After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode](#), page 3-13.
-

Replacing CPUs and Heatsinks

This section contains the following topics:

- [CPU Configuration Rules, page 3-32](#)
- [Replacing a CPU and Heatsink, page 3-33](#)
- [Additional CPU-Related Parts to Order with RMA Replacement Motherboards, page 3-37](#)



Note

You can use Xeon v3- and v4-based nodes in the same cluster. Do not mix Xeon v3 and v4 CPUs within the same node.

Special Information For *Upgrades* to Intel Xeon v4 CPUs



Caution

You must upgrade your server firmware to the required minimum level *before* you upgrade to Intel v4 CPUs. Older firmware versions cannot recognize the new CPUs and this results in a non-bootable server.

The minimum software and firmware versions required for the node to support Intel v4 CPUs are as follows:

Table 3-8 Minimum Requirements For Intel Xeon v4 CPUs

Software or Firmware	Minimum Version
Node Cisco IMC	2.0(10)
Node BIOS	2.0(10)
Cisco UCS Manager (UCSM-managed system only)	3.1(1)

Do one of the following actions:

- If your node's firmware and/or Cisco UCS Manager software are already at the required levels shown in [Table 3-8](#), you can replace the CPU hardware by using the procedure in this section.
- If your node's firmware and/or Cisco UCS Manager software is earlier than the required levels, use the instructions in [Special Instructions for Upgrades to Intel Xeon v4 Series, page 3-36](#) to upgrade your software. After you upgrade the software, return to the procedure in this section as directed to replace the CPU hardware.

CPU Configuration Rules

This node has two CPU sockets. Each CPU supports four DIMM channels (12 DIMM sockets). See [Figure 3-10](#).

- The minimum configuration is that the node must have at two CPUs installed.
- Do not mix Xeon v3 and v4 CPUs within the same node.

Replacing a CPU and Heatsink



Caution

CPUs and their motherboard sockets are fragile and must be handled with care to avoid damaging pins during installation. The CPUs must be installed with heatsinks and their thermal grease to ensure proper cooling. Failure to install a CPU correctly might result in damage to the node.



Note

This node uses the new independent loading mechanism (ILM) CPU sockets, so no Pick-and-Place tools are required for CPU handling or installation. Always grasp the plastic frame on the CPU when handling.

- Step 1** Put the node in Cisco HX Maintenance mode as described in [Shutting Down the Node Through vSphere With Cisco HX Maintenance Mode](#), page 3-10.
- Step 2** Shut down the node as described in [Shutting Down the Node](#), page 3-8.
- Step 3** Decommission the node as described in [Decommissioning the Node Using Cisco UCS Manager](#), page 3-11.



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 the 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 as described in [Removing and Replacing the Node Top Cover](#), page 3-14.
- Step 7** Remove the plastic air baffle that sits over the CPUs.
- Step 8** Remove the heatsink that you are replacing:
 - a. Use a Number 2 Phillips-head screwdriver to loosen the four captive screws that secure the heatsink.



Note Alternate loosening each screw evenly to avoid damaging the heatsink or CPU.






- b. Lift the heatsink off of the CPU.
- Step 9** Open the CPU retaining mechanism:
 - a. Unclip the first retaining latch labeled with the  icon, and then unclip the second retaining latch labeled with the  icon. See [Figure 3-11](#).
 - b. Open the hinged CPU cover plate.

Figure 3-11 CPU Socket



1	CPU retaining latch 	4	Hinged CPU seat
2	CPU retaining latch 	5	Finger-grips on plastic CPU frame
3	Hinged CPU cover plate		

Step 10 Remove any existing CPU:

- a. With the latches and hinged CPU cover plate open, swing the CPU in its hinged seat up to the open position, as shown in [Figure 3-11](#).
- b. Grasp the CPU by the finger-grips on its plastic frame and lift it up and out of the hinged CPU seat.
- c. Set the CPU aside on an antistatic surface.

Step 11 Install a new CPU:




- a. Grasp the new CPU by the finger-grips on its plastic frame and align the tab on the frame that is labeled “ALIGN” with the hinged seat, as shown in [Figure 3-12](#).
- b. Insert the tab on the CPU frame into the seat until it stops and is held firmly.
The line below the word “ALIGN” should be level with the edge of the seat, as shown in [Figure 3-12](#).
- c. Swing the hinged seat with the CPU down until the CPU frame clicks in place and holds flat in the socket.
- d. Close the hinged CPU cover plate.
- e. Clip down the CPU retaining latch with the  icon, and then clip down the CPU retaining latch with the  icon. See [Figure 3-11](#).

Figure 3-12 CPU and Socket Alignment Features



1	SLS mechanism on socket	2	Tab on CPU frame (labeled ALIGN)
---	-------------------------	---	----------------------------------

Step 12 Install a heat sink:



Caution

The heat sink must have new thermal grease on the heat sink-to-CPU surface to ensure proper cooling. If you are reusing a heat sink, you must remove the old thermal grease from the heatsink and the CPU surface. If you are installing a new heat sink, skip to Step c.

- a. Apply the cleaning solution, which is included with the heatsink cleaning kit (UCSX-HSCK=, shipped with spare CPUs), to the old thermal grease on the heatsink and CPU and let it soak for a least 15 seconds.
- b. Wipe all of the old thermal grease off the old heat sink and CPU using the soft cloth that is included with the heatsink cleaning kit. Be careful to not scratch the heat sink surface.



Note

New heatsinks come with a pre-applied pad of thermal grease. If you are reusing a heatsink, you must apply thermal grease from a syringe (UCS-CPU-GREASE3=).

- c. Align the four heatsink captive screws with the motherboard standoffs, and then use a Number 2 Phillips-head screwdriver to tighten the captive screws evenly.



Note

Alternate tightening each screw evenly to avoid damaging the heatsink or CPU.

Step 13 Replace the air baffle.

Step 14 Replace the top cover.

Step 15 Replace the node in the rack, replace power cables, and then power on the node by pressing the **Power** button.

Step 16 Recommission the node as described in [Recommissioning the Node Using Cisco UCS Manager](#), page 3-12.

- Step 17** Associate the node to its service profile as described in [Associating a Service Profile With an HX Node](#), page 3-12.
- Step 18** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode](#), page 3-13.
-

Special Instructions for *Upgrades* to Intel Xeon v4 Series



Caution

You must upgrade your node firmware to the required minimum level *before* you upgrade to Intel v4 CPUs. Older firmware versions cannot recognize the new CPUs and this results in a non-bootable node.

Use the following procedure to upgrade the node and CPUs.

- Step 1** Upgrade the Cisco UCS Manager software to the minimum version for your node (or later). See [Table 3-8](#).
- Use the procedures in the appropriate Cisco UCS Manager upgrade guide (depending on your current software version): [Cisco UCS Manager Upgrade Guides](#).
- Step 2** Use Cisco UCS Manager to upgrade and activate the node Cisco IMC to the minimum version for your node (or later). See [Table 3-8](#).
- Use the procedures in the GUI or CLI [Cisco UCS Manager Firmware Management Guide](#) for your release.
- Step 3** Use Cisco UCS Manager to upgrade and activate the node BIOS to the minimum version for your node (or later). See [Table 3-8](#).
- Use the procedures in the Cisco UCS Manager GUI or CLI [Cisco UCS Manager Firmware Management Guide](#) for your release.
- Step 4** Replace the CPUs with the Intel Xeon v4 Series CPUs.
- Use the CPU replacement procedures in [Replacing a CPU and Heatsink](#), page 3-33.
-

Additional CPU-Related Parts to Order with RMA Replacement Motherboards

When a return material authorization (RMA) of the motherboard or CPU is done, additional parts might not be included with the CPU or motherboard spare bill of materials (BOM). The TAC engineer might need to add the additional parts to the RMA to help ensure a successful replacement.

**Note**

This node uses the new independent loading mechanism (ILM) CPU sockets, so no Pick-and-Place tools are required for CPU handling or installation. Always grasp the plastic frame on the CPU when handling.

- Scenario 1—You are reusing the existing heatsinks:
 - Heat sink cleaning kit (UCSX-HSCK=)
 - Thermal grease kit for HX240c (UCS-CPU-GREASE3=)
- Scenario 2—You are replacing the existing heatsinks:
 - Heat sink (UCSC-HS-C240M4=)
 - Heat sink cleaning kit (UCSX-HSCK=)

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 thermal interface material and the other to prepare the surface of the heatsink.

New heatsink spares come with a pre-applied pad of thermal grease. It is important to clean the old thermal grease off of the CPU prior to installing the heatsinks. Therefore, when you are ordering new heatsinks, you must order the heatsink cleaning kit.

Replacing a Cisco Modular HBA Card

The node has an internal, dedicated PCIe slot on the motherboard for an HBA card (see [Figure 3-13](#)).

-
- Step 1** Put the node in Cisco HX Maintenance mode as described in [Shutting Down the Node Through vSphere With Cisco HX Maintenance Mode](#), page 3-10.
 - Step 2** Shut down the node as described in [Shutting Down the Node](#), page 3-8.
 - Step 3** Decommission the node as described in [Decommissioning the Node Using Cisco UCS Manager](#), page 3-11.


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 the 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 as described in [Removing and Replacing the Node Top Cover](#), page 3-14.
- Step 7** Remove an existing HBA controller card:
 - a. Disconnect the data cable from the card. Depress the tab on the cable connector and pull.
 - b. Disconnect the supercap power module cable from the transportable memory module (TMM), if present.
 - c. Lift straight up on the metal bracket that holds the card. The bracket lifts off of two pegs on the chassis wall.
 - d. Loosen the two thumbscrews that hold the card to the metal bracket and then lift the card from the bracket.
- Step 8** Install a new HBA controller card:


Caution


When installing the card to the bracket, be careful so that you do not scrape and damage electronic components on the underside of the card on features of the bracket. Also avoid scraping the card when you install the bracket to the pegs on the chassis wall.

- a. Set the new card on the metal bracket, aligned so that the thumbscrews on the card enter the threaded standoffs on the bracket. Tighten the thumbscrews to secure the card to the bracket.
 - b. Align the two slots on the back of the bracket with the two pegs on the chassis wall.

The two slots on the bracket must slide down over the pegs at the same time that you push the card into the motherboard socket.
 - c. Gently press down on both top corners of the metal bracket to seat the card into the socket on the motherboard.
 - d. Connect the supercap power module cable to its connector on the TMM, if present.
 - e. Connect the single data cable to the card.
- Step 9** Replace the top cover.
 - Step 10** Replace the node in the rack, replace power cables, and then power on the node by pressing the **Power** button.

- Step 11 [Recommission the node as described in *Recommissioning the Node Using Cisco UCS Manager*, page 3-12.](#)
- Step 12 [Associate the node to its service profile as described in *Associating a Service Profile With an HX Node*, page 3-12.](#)
- Step 13 [After ESXi reboot, exit HX Maintenance mode as described in *Exiting HX Maintenance Mode*, page 3-13.](#)

Figure 3-13 HBA Card Location



1	Thumbscrews on card	2	Cisco HBA bracket
---	---------------------	---	-------------------

Replacing the Motherboard 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 node settings when the node is disconnected from power. The battery type is CR2032. Cisco supports the industry-standard CR2032 battery, which can be purchased from most electronic stores.

- Step 1 Put the node in Cisco HX Maintenance mode as described in [Shutting Down the Node Through vSphere With Cisco HX Maintenance Mode](#), page 3-10.
- Step 2 Shut down the node as described in [Shutting Down the Node](#), page 3-8.
- Step 3 Decommission the node as described in [Decommissioning the Node Using Cisco UCS Manager](#), page 3-11.



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 the 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 as described in [Removing and Replacing the Node Top Cover](#), page 3-14.
- Step 7 Remove the battery from its holder on the motherboard (see [Figure 3-14](#)):
 - a. Use a small screwdriver or pointed object to press inward on the battery at the prying point (see [Figure 3-14](#)).
 - b. Lift up on the battery and remove it from the holder.
- Step 8 Install an RTC battery. Insert the battery into its holder and press down until it clicks in place.



Note The positive side of the battery marked "3V+" should face upward.

- Step 9 Replace the top cover.
- Step 10 Replace the node in the rack, replace power cables, and then power on the node by pressing the **Power** button.
- Step 11 Recommission the node as described in [Recommissioning the Node Using Cisco UCS Manager](#), page 3-12.
- Step 12 Associate the node to its service profile as described in [Associating a Service Profile With an HX Node](#), page 3-12.
- Step 13 After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode](#), page 3-13.

Figure 3-14 RTC Battery Location and Prying Point



1	RTC battery holder on motherboard	2	Prying point
---	-----------------------------------	---	--------------

Replacing an Internal SD Card

The node has two internal SD card bays on the motherboard. Dual SD cards are supported.

-
- Step 1 Put the node in Cisco HX Maintenance mode as described in [Shutting Down the Node Through vSphere With Cisco HX Maintenance Mode](#), page 3-10.
 - Step 2 Shut down the node as described in [Shutting Down the Node](#), page 3-8.
 - Step 3 Decommission the node as described in [Decommissioning the Node Using Cisco UCS Manager](#), page 3-11.



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 the 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 as described in [Removing and Replacing the Node Top Cover](#), page 3-14.
- Step 7 Remove an SD card (see [Figure 3-15](#)).
 - a. Push on the top of the SD card, and then release it to allow it to spring out from the slot.
 - b. Remove the SD card from the slot.
- Step 8 Install an SD card:
 - a. Insert the SD card into the slot with the label side facing up.
 - b. Press on the top of the card until it clicks in the slot and stays in place.
- Step 9 Replace the top cover.
- Step 10 Replace the node in the rack, replace power cables, and then power on the node by pressing the **Power** button.
- Step 11 Recommission the node as described in [Recommissioning the Node Using Cisco UCS Manager](#), page 3-12.
- Step 12 Associate the node to its service profile as described in [Associating a Service Profile With an HX Node](#), page 3-12.
- Step 13 After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode](#), page 3-13.

Figure 3-15 SD Card Bay Location and Numbering on the Motherboard



Enabling or Disabling the Internal USB Port



Caution

We do not recommend that you hot-swap the internal USB drive while the node is powered on.

The factory default is for all USB ports on the node to be enabled. However, the internal USB port can be enabled or disabled in the node BIOS. See [Figure 3-5](#) for the location of the internal USB 3.0 slot on the motherboard.

- Step 1 Enter the BIOS Setup Utility by pressing the **F2** key when prompted during bootup.
- Step 2 Navigate to the **Advanced** tab.
- Step 3 On the Advanced tab, select **USB Configuration**.
- Step 4 On the USB Configuration page, choose **USB Ports Configuration**.
- Step 5 Scroll to **USB Port: Internal**, press **Enter**, and then choose either **Enabled** or **Disabled** from the dialog box.
- Step 6 Press **F10** to save and exit the utility.

Replacing a PCIe Riser

The node contains two toolless PCIe risers for horizontal installation of PCIe cards. See [Replacing a PCIe Card, page 3-46](#) for the specifications of the PCIe slots on the risers.

- Step 1 Put the node in Cisco HX Maintenance mode as described in [Shutting Down the Node Through vSphere With Cisco HX Maintenance Mode, page 3-10](#).
- Step 2 Shut down the node as described in [Shutting Down the Node, page 3-8](#).
- Step 3 Decommission the node as described in [Decommissioning the Node Using Cisco UCS Manager, page 3-11](#).



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 the 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 as described in [Removing and Replacing the Node Top Cover, page 3-14](#).
- Step 7 Remove the PCIe riser that you are replacing (see [Figure 3-16](#)):
 - a. Grasp the top of the riser and lift straight up on both ends to disengage its circuit board from the socket on the motherboard. Set the riser on an antistatic mat.
 - b. If the riser has a card installed, remove the card from the riser. See [Replacing a PCIe Card, page 3-46](#).
- Step 8 Install a new PCIe riser:
 - a. If you removed a card from the old PCIe riser, install the card to the new riser (see [Replacing a PCIe Card, page 3-46](#)).
 - b. Position the PCIe riser over its socket on the motherboard and over its alignment slots in the chassis (see [Figure 3-16](#)). There are also two alignment pegs on the motherboard for each 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.”

- c. Carefully push down on both ends of the PCIe riser to fully engage its circuit board connector with the socket on the motherboard.
- Step 9 Replace the top cover.
- Step 10 Replace the node in the rack, replace power cables, and then power on the node by pressing the **Power** button.
- Step 11 Recommission the node as described in [Recommissioning the Node Using Cisco UCS Manager, page 3-12](#).
- Step 12 Associate the node to its service profile as described in [Associating a Service Profile With an HX Node, page 3-12](#).

Step 13 After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode](#), page 3-13.

Figure 3-16 PCIe Riser Alignment Features



1	Alignment peg locations on motherboard (two for each riser)	2	Alignment channel locations on chassis (two for each riser)
---	---	---	---

Replacing a PCIe Card



Caution

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 third-party card occurs.

This section includes the following topics:

- [PCIe Slots, page 3-46](#)
- [Replacing a PCIe Card, page 3-47](#)
- [Special Considerations for Cisco UCS Virtual Interface Cards, page 3-60](#)
- [Installing Multiple PCIe Cards and Resolving Limited Resources, page 3-49](#)

PCIe Slots

The system contains two toolless PCIe risers for horizontal installation of PCIe cards (see [Figure 3-17](#)).

- Riser 1: PCIe slots 1 and 2; slot 3 taken by two internal SATA SSD boot-drive sockets. See [Table 3-9](#).
- Riser 2 contains slots PCIe 4, 5, and 6. See [Table 3-10](#).

Figure 3-17 Rear Panel, Showing PCIe Slots




Table 3-9 Riser 1C (UCSC-PCI-1C-240M4) PCIe Expansion Slots

Slot Number	Electrical Lane Width	Connector Length	Card Length	Card Height	NCSI Support
1	Gen-3 x8	x16 connector	3/4 length	Full height	Yes
2	Gen-3 x16	x24 connector	Full length	Full height	Yes
SATA SSD sockets (two)	NA	NA	NA	NA	NA

Table 3-10 Riser 2 (UCSC-PCI-2-240M4) PCIe Expansion Slots

Slot Number	Electrical Lane Width	Connector Length	Card Length	Card Height	NCSI Support
4	Gen-3 x8	x24 connector	3/4 length	Full height	Yes
5	Gen-3 x16	x24 connector	Full length	Full height	Yes ¹
6	Gen-3 x8	x16 connector	Full length	Full height	No

1. NCSI is supported in only one slot at a time in this riser version.

Replacing a PCIe Card

**Note**

If you are installing a Cisco UCS Virtual Interface Card, there are prerequisite considerations. See [Special Considerations for Cisco UCS Virtual Interface Cards, page 3-60](#).

- Step 1** Put the node in Cisco HX Maintenance mode as described in [Shutting Down the Node Through vSphere With Cisco HX Maintenance Mode, page 3-10](#).
- Step 2** Shut down the node as described in [Shutting Down the Node, page 3-8](#).
- Step 3** Decommission the node as described in [Decommissioning the Node Using Cisco UCS Manager, page 3-11](#).

**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 the 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 as described in [Removing and Replacing the Node Top Cover, page 3-14](#).
- Step 7** Remove a PCIe card (or a blanking panel) from the PCIe riser:
- Lift straight up on both ends of the riser to disengage its circuit board from the socket on the motherboard. Set the riser on an antistatic mat.
 - On the bottom of the riser, loosen the single thumbscrew that holds the securing plate (see [Figure 3-18](#)).
 - Swing open the securing plate and remove it from the riser to provide access.
 - Swing open the card-tab retainer that secures the back-panel tab of the card (see [Figure 3-18](#)).
 - Pull evenly on both ends of the PCIe card to disengage it from the socket on the PCIe riser (or remove a blanking panel) and then set the card aside.
- Step 8** Install a PCIe card:
- 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.
Ensure that the card rear panel tab sits flat against the PCIe riser rear panel opening.
 - Close the card-tab retainer (see [Figure 3-18](#)).
 - Return the securing plate to the riser. Insert the two hinge-tabs into the two slots on the riser, and then swing the securing plate closed.
 - Tighten the single thumbscrew on the bottom of the securing plate.
 - Position the PCIe riser over its socket on the motherboard and over its alignment features in the chassis (see [Figure 3-16](#)).

- g. 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.
- Step 10** Replace the node in the rack, replace power cables, and then power on the node by pressing the **Power** button.
- Step 11** Recommission the node as described in [Recommissioning the Node Using Cisco UCS Manager](#), page 3-12.
- Step 12** Associate the node to its service profile as described in [Associating a Service Profile With an HX Node](#), page 3-12.
- Step 13** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode](#), page 3-13.

Figure 3-18 PCIe Riser Securing Features (Three-Slot Riser Shown)



353239

1	Securing plate hinge-tabs	3	GPU card power connector
2	Securing plate thumbscrew (knob not visible on underside of plate)	4	Card-tab retainer in open position

Installing Multiple PCIe Cards and Resolving Limited Resources

When a large number of PCIe add-on cards are installed in the node, the node might run out of the following resources required for PCIe devices:

- Option ROM memory space
- 16-bit I/O space

The topics in this section provide guidelines for resolving the issues related to these limited resources:

- [Resolving Insufficient Memory Space to Execute Option ROMs, page 3-49](#)
- [Resolving Insufficient 16-Bit I/O Space, page 3-50](#)

Resolving Insufficient Memory Space to Execute Option ROMs

The node has very limited memory to execute PCIe legacy option ROMs, so when a large number of PCIe add-on cards are installed in the node, the node BIOS might not be able to execute all of the option ROMs. The node BIOS loads and executes the option ROMs in the order that the PCIe cards are enumerated (slot 1, slot 2, slot 3, and so on).

If the node BIOS does not have sufficient memory space to load any PCIe option ROM, it skips loading that option ROM, reports a system event log (SEL) event to the Cisco IMC controller and reports the following error in the Error Manager page of the BIOS Setup utility:

ERROR CODE	SEVERITY	INSTANCE	DESCRIPTION
146	Major	N/A	PCI out of resources error. Major severity requires user intervention but does not prevent system boot.

To resolve this issue, disable the Option ROMs that are not needed for node booting. The BIOS Setup Utility provides the setup options to enable or disable the Option ROMs at the PCIe slot level for the PCIe expansion slots and at the port level for the onboard NICs. These options can be found in the BIOS Setup Utility **Advanced > PCI Configuration** page.

- **Guidelines for RAID controller booting**

If the node is configured to boot primarily from RAID storage, make sure that the option ROMs for the slots where your RAID controllers are installed are enabled in the BIOS, depending on your RAID controller configuration.

If the RAID controller does not appear in the node boot order even with the option ROMs for those slots are enabled, the RAID controller option ROM might not have sufficient memory space to execute. In that case, disable other option ROMs that are not needed for the node configuration to free up some memory space for the RAID controller option ROM.

- **Guidelines for onboard NIC PXE booting**

If the node is configured to primarily perform PXE boot from onboard NICs, make sure that the option ROMs for the onboard NICs to be booted from are enabled in the BIOS Setup Utility. Disable other option ROMs that are not needed to create sufficient memory space for the onboard NICs.

Resolving Insufficient 16-Bit I/O Space

The node has only 64 KB of legacy 16-bit I/O resources available. This 64 KB of I/O space is divided between the CPUs in the node because the PCIe controller is integrated into the CPUs. This node BIOS has the capability to dynamically detect the 16-bit I/O resource requirement for each CPU and then balance the 16-bit I/O resource allocation between the CPUs during the PCI bus enumeration phase of the BIOS POST.

When a large number of PCIe cards are installed in the node, the node BIOS might not have sufficient I/O space for some PCIe devices. If the node BIOS is not able to allocate the required I/O resources for any PCIe devices, the following symptoms have been observed:

- The node might get stuck in an infinite reset loop.
- The BIOS might appear to hang while initializing PCIe devices.
- The PCIe option ROMs might take excessive time to complete, which appears to lock up the node.
- PCIe boot devices might not be accessible from the BIOS.
- PCIe option ROMs might report initialization errors. These errors are seen before the BIOS passes control to the operating system.
- The keyboard might not work.

To work around this problem, rebalance the 16-bit I/O load using the following methods:

1. Physically remove any unused PCIe cards.
2. If the node has one or more Cisco virtual interface cards (VICs) installed, disable the PXE boot on the VICs that are not required for the node boot configuration by using the Network Adapters page in Cisco IMC Web UI to free up some 16-bit I/O resources. Each VIC uses a minimum 16 KB of 16-bit I/O resource, so disabling PXE boot on Cisco VICs would free up some 16-bit I/O resources that can be used for other PCIe cards that are installed in the node.

Installing a Trusted Platform Module

The trusted platform module (TPM) is a small circuit board that connects to a motherboard socket and is secured by a one-way screw. The socket location is on the motherboard under PCIe riser 2.

This section contains the following procedures, which must be followed in this order when installing and enabling a TPM:

1. [Installing the TPM Hardware, page 3-51](#)
2. [Enabling TPM Support in the BIOS, page 3-52](#)
3. [Enabling the Intel TXT Feature in the BIOS, page 3-53](#)

Installing the TPM Hardware

-
- Step 1** Put the node in Cisco HX Maintenance mode as described in [Shutting Down the Node Through vSphere With Cisco HX Maintenance Mode, page 3-10](#).
- Step 2** Shut down the node as described in [Shutting Down the Node, page 3-8](#).
- Step 3** Decommission the node as described in [Decommissioning the Node Using Cisco UCS Manager, page 3-11](#).

**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 the 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 as described in [Removing and Replacing the Node Top Cover, page 3-14](#).
- Step 7** Remove PCIe riser 2 to provide clearance. See [Replacing a PCIe Riser, page 3-44](#) for instructions.
- Step 8** Install a TPM:
- a. Locate the TPM socket on the motherboard, as shown in [Figure 3-19](#).
 - b. Align the connector that is on the bottom of the TPM circuit board with the motherboard TPM socket. Align the screw hole and standoff on the TPM board with the screw hole that is adjacent to the TPM socket.
 - c. Push down evenly on the TPM to seat it in the motherboard socket.
 - d. Install the single one-way screw that secures the TPM to the motherboard.
- Step 9** Replace PCIe riser 2 to the node. See [Replacing a PCIe Riser, page 3-44](#) for instructions.
- Step 10** Replace the top cover.
- Step 11** Replace the node in the rack, replace power cables, and then power on the node by pressing the **Power** button.
- Step 12** Recommission the node as described in [Recommissioning the Node Using Cisco UCS Manager, page 3-12](#).

- Step 13** Associate the node to its service profile as described in [Associating a Service Profile With an HX Node](#), page 3-12.
- Step 14** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode](#), page 3-13.
- Step 15** Continue with [Enabling TPM Support in the BIOS](#), page 3-52.

Figure 3-19 TPM Socket Location on Motherboard



- | | |
|---|---|
| 1 | TPM socket location on the motherboard (under PCIe riser 2) |
|---|---|

Enabling TPM Support in the BIOS



Note After hardware installation, you must enable TPM support in the BIOS.



Note You must set a BIOS Administrator password before performing this procedure. To set this password, press the **F2** key when prompted during system boot to enter the BIOS Setup utility. Then navigate to **Security > Set Administrator Password** and enter the new password twice as prompted.

- Step 1** Enable TPM support:
- a. Watch during bootup for the F2 prompt, and then press **F2** to enter BIOS setup.
 - b. Log in to the BIOS Setup Utility with your BIOS Administrator password.
 - c. On the BIOS Setup Utility window, choose the **Advanced** tab.
 - d. Choose **Trusted Computing** to open the TPM Security Device Configuration window.
 - e. Change TPM SUPPORT to **Enabled**.
 - f. Press **F10** to save your settings and reboot the node.

- Step 2** Verify that TPM support is now enabled:
- Watch during bootup for the F2 prompt, and then press **F2** to enter BIOS setup.
 - Log into the BIOS Setup utility with your BIOS Administrator password.
 - Choose the **Advanced** tab.
 - Choose **Trusted Computing** to open the TPM Security Device Configuration window.
 - Verify that TPM SUPPORT and TPM State are Enabled.
- Step 3** Continue with [Enabling the Intel TXT Feature in the BIOS, page 3-53](#).
-

Enabling the Intel TXT Feature in the BIOS

Intel Trusted Execution Technology (TXT) provides greater protection for information that is used and stored on the node. A key aspect of that protection is the provision of an isolated execution environment and associated sections of memory where operations can be conducted on sensitive data, invisibly to the rest of the node. Intel TXT provides for a sealed portion of storage where sensitive data such as encryption keys can be kept, helping to shield them from being compromised during an attack by malicious code.

- Step 1** Reboot the node and watch for the prompt to press F2.
- Step 2** When prompted, press **F2** to enter the BIOS Setup utility.
- Step 3** Verify that the prerequisite BIOS values are enabled:
- Choose the **Advanced** tab.
 - Choose **Intel TXT(LT-SX) Configuration** to open the Intel TXT(LT-SX) Hardware Support window.
 - Verify that the following items are listed as Enabled:
 - VT-d Support (default is Enabled)
 - VT Support (default is Enabled)
 - TPM Support
 - TPM State
 - If VT-d Support and VT Support are already enabled, skip to [Step 4](#).
 - If VT-d Support and VT Support are not enabled, continue with the next steps to enable them.
 - Press **Escape** to return to the BIOS Setup utility **Advanced** tab.
 - On the Advanced tab, choose **Processor Configuration** to open the Processor Configuration window.
 - Set Intel (R) VT and Intel (R) VT-d to **Enabled**.
- Step 4** Enable the Intel Trusted Execution Technology (TXT) feature:
- Return to the Intel TXT(LT-SX) Hardware Support window if you are not already there.
 - Set TXT Support to **Enabled**.
- Step 5** Press **F10** to save your changes and exit the BIOS Setup utility.
-

Replacing Power Supplies

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

- See [Power Specifications, page A-2](#) for more information about the supported power supplies.
- See [Rear Panel LEDs and Buttons, page 3-4](#) for information about the power supply LEDs.
- See [Installing a DC Power Supply, page 3-56](#) for information about wiring a DC power supply.

**Note**


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

**Note**

Do not mix power supply types in the node. Both power supplies must be the same wattage and Cisco product ID (PID).

- Step 1** Perform one of the following actions:
- If your node has two power supplies, you do not have to shut down the node. Continue with step 2.
 - If your node has only one power supply:
 - a. Put the node in Cisco HX Maintenance mode as described in [Shutting Down the Node Through vSphere With Cisco HX Maintenance Mode, page 3-10](#).
 - b. Shut down the node as described in [Shutting Down the Node, page 3-8](#).
 - c. Decommission the node as described in [Decommissioning the Node Using Cisco UCS Manager, page 3-11](#).
- Step 2** Remove the power cord from the power supply that you are replacing.
- Step 3** Grasp the power supply handle while pinching the green release lever towards the handle (see [Figure 3-20](#)).
- Step 4** Pull the power supply out of the bay.
- Step 5** 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.
- Step 6** Only if you shut down the node, perform these steps:
- a. Press the **Power** button to return the node to main power mode.
 - b. Recommission the node as described in [Recommissioning the Node Using Cisco UCS Manager, page 3-12](#).
 - c. Associate the node to its service profile as described in [Associating a Service Profile With an HX Node, page 3-12](#).
 - d. After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, page 3-13](#).

Figure 3-20 Power Supplies



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

Installing a DC Power Supply



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

When installing or replacing the unit, the ground connection must always be made first and disconnected last. Statement 1046



Warning

Installation of the equipment must comply with local and national electrical codes. Statement 1074



Warning

Hazardous voltage or energy may be present on DC power terminals. Always replace cover when terminals are not in service. Be sure uninsulated conductors are not accessible when cover is in place. Statement 1075

Installing a 930W DC Power Supply, UCSC-PSU-930WDC

If you are using a Version 1, 930W DC power supply, stripped wires connect power to the removable connector block.



Caution

Before beginning this wiring procedure, turn off the DC power source from your facility's circuit breaker to avoid electric shock hazard.

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

Step 2 Remove the DC power connector block from the power supply. (The spare PID for this connector is UCSC-CONN-930WDC=.)

To release the connector block from the power supply, push the orange plastic button on the top of the connector inward toward the power supply and pull the connector block out.

Step 3 Strip 15mm (.59 inches) of insulation off the DC wires that you will use.



Note

The recommended wire gauge is 8 AWG. The minimum wire gauge is 10 AWG.

Step 4 Orient the connector as shown in [Figure 3-21](#), with the orange plastic button toward the top.


Step 5 Use a small screwdriver to depress the spring-loaded wire retainer lever on the lower spring-cage wire connector. Insert your green (ground) wire into the aperture and then release the lever.

Step 6 Use a small screwdriver to depress the wire retainer lever on the middle spring-cage wire connector. Insert your black (DC negative) wire into the aperture and then release the lever.

Step 7 Use a small screwdriver to depress the wire retainer lever on the upper spring-cage wire connector. Insert your red (DC positive) wire into the aperture and then release the lever.

Step 8 Insert the connector block back into the power supply. Make sure that your red (DC positive) wire aligns with the power supply label, "+ DC".

Figure 3-21 Version 1 930 W, -48 VDC Power Supply Connector Block



1	Wire retainer lever	2	Orange plastic button on top of the connector
---	---------------------	---	---

Replacing an mLOM Card (Cisco VIC 1227)

The node uses a Cisco VIC 1227 mLOM adapter. The mLOM card socket remains powered when the node is in 12 V standby power mode and it supports the network communications services (NCSI) protocol.

- [Replacing an mLOM Card, page 3-58](#)
- [Special Considerations for Cisco UCS Virtual Interface Cards, page 3-60](#)

Replacing an mLOM Card

-
- Step 1** Put the node in Cisco HX Maintenance mode as described in [Shutting Down the Node Through vSphere With Cisco HX Maintenance Mode, page 3-10](#).
- Step 2** Shut down the node as described in [Shutting Down the Node, page 3-8](#).
- Step 3** Decommission the node as described in [Decommissioning the Node Using Cisco UCS Manager, page 3-11](#).



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 the 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 as described in [Removing and Replacing the Node Top Cover, page 3-14](#).
- Step 7** Remove PCIe riser 1 to provide clearance. See [Replacing a PCIe Riser, page 3-44](#) for instructions.
- Step 8** Remove any existing mLOM card or a blanking panel (see [Figure 3-22](#)):
- Loosen the single thumbscrew that secures the mLOM card to the chassis floor.
 - Slide the mLOM card horizontally to disengage its connector from the motherboard socket.
- Step 9** Install a new mLOM card:
- Set the mLOM card on the chassis floor so that its connector is aligned with the motherboard socket and its thumbscrew is aligned with the standoff on the chassis floor.
 - Push the card's connector into the motherboard socket horizontally.
 - Tighten the thumbscrew to secure the card to the chassis floor.
- Step 10** Return PCIe riser 1 to the node. See [Replacing a PCIe Riser, page 3-44](#) for instructions.
- Step 11** Replace the top cover.
- Step 12** Replace the node in the rack, replace power cables, and then power on the node by pressing the **Power** button.
- Step 13** Recommission the node as described in [Recommissioning the Node Using Cisco UCS Manager, page 3-12](#).

- Step 14 Associate the node to its service profile as described in [Associating a Service Profile With an HX Node](#), page 3-12.
- Step 15 After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode](#), page 3-13.

Figure 3-22 mLOM Card Location



1	mLOM card (VIC 1227) socket location on motherboard (under PCIe riser 1)	
---	--	--

Special Considerations for Cisco UCS Virtual Interface Cards

Table 3-11 describes the requirements for the supported Cisco UCS virtual interface cards (VICs).

Table 3-11 Cisco HX240c Requirements for Virtual Interface Cards

Virtual Interface Card (VIC)	Number of this VIC Supported in Node	Slots That Support VICs	Primary Slot for Cisco UCS Manager Integration	Primary Slot for Cisco Card NIC Mode	Minimum Cisco IMC Firmware	Minimum VIC Firmware
Cisco UCS VIC 1227 UCSC-MLOM-CSC-02	1 mLOM	mLOM	mLOM	mLOM	2.0(3)	4.0(0)



Note

The Cisco UCS VIC 1227 (UCSC-MLOM-CSC-02) is not compatible to use in *Cisco Card* NIC mode with a certain Cisco SFP+ module. Do not use a Cisco SFP+ module part number 37-0961-01 that has a serial number in the range MOC1238xxxx to MOC1309xxxx. If you use the Cisco UCS VIC 1227 in *Cisco Card* NIC mode, use a different part number Cisco SFP+ module, or you can use this part number 37-0961-01 if the serial number is *not* included in the range above. See the data sheet for this adapter for other supported SFP+ modules: [Cisco UCS VIC 1227 Data Sheet](#)

Service DIP Switches

This section includes the following topics:


- [DIP Switch Location on the Motherboard, page 3-61](#)
- [Using the BIOS Recovery DIP Switch, page 3-62](#)
- [Using the Clear Password DIP Switch, page 3-64](#)
- [Using the Clear CMOS DIP Switch, page 3-66](#)

DIP Switch Location on the Motherboard

See [Figure 3-23](#). The position of the block of DIP switches (SW8) is shown in red. In the magnified view, all switches are shown in the default position.

- BIOS recovery—switch 1.
- Clear password—switch 2.
- Not used—switch 3.
- Clear CMOS—switch 4.

Figure 3-23 Service DIP Switches



1	DIP switch block SW8	3	Clear password switch 2
2	Clear CMOS switch 4	4	BIOS recovery switch 1

Using the BIOS Recovery DIP Switch



Note

The following procedures use a recovery.cap recovery file. In Cisco IMC releases 3.0(1) and later, this recovery file has been renamed bios.cap

Depending on which stage the BIOS becomes corrupted, you might see different behavior.

- If the BIOS BootBlock is corrupted, you might see the node get stuck on the following message:

```
Initializing and configuring memory/hardware
```

- If it is a non-BootBlock corruption, the following message 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 recovery.cap (or 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 recovery.cap (or 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 recovery.cap (or 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 thumb drive. The recovery folder contains the recovery.cap (or bios.cap) file that is required in this procedure.



Note The recovery.cap (or 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 FAT16 or FAT32 file systems.

- 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 file on the USB thumb 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 thumb 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 DIP switch and recovery.cap (or bios.cap) File

See [Figure 3-23](#) for the location of the SW8 block of DIP switches.

- 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 thumb drive. The recovery folder contains the recovery.cap (or bios.cap) file that is required in this procedure.
-  **Note** The recovery.cap (or 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 FAT16 or FAT32 file systems.
- Step 3** Put the node in Cisco HX Maintenance mode as described in [Shutting Down the Node Through vSphere With Cisco HX Maintenance Mode, page 3-10](#).
- Step 4** Shut down the node as described in [Shutting Down the Node, page 3-8](#).
- Step 5** Decommission the node as described in [Decommissioning the Node Using Cisco UCS Manager, page 3-11](#).



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 the 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 as described in [Removing and Replacing the Node Top Cover, page 3-14](#).
- Step 9** Slide the BIOS recovery DIP switch from position 1 to the closed position (see [Figure 3-23](#)).
- 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 recovery file on the USB thumb 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 thumb 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 Slide the BIOS recovery DIP switch from the closed position back to the default position 1.



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

Step 16 Replace the top cover to the node.

Step 17 Replace the node in the rack, replace power cables, and then power on the node by pressing the **Power** button.

Step 18 Recommission the node as described in [Recommissioning the Node Using Cisco UCS Manager, page 3-12](#).

Step 19 Associate the node to its service profile as described in [Associating a Service Profile With an HX Node, page 3-12](#).

Step 20 After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, page 3-13](#).

Using the Clear Password DIP Switch

See [Figure 3-23](#) for the location of this DIP switch. 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 the Node Through vSphere With Cisco HX Maintenance Mode, page 3-10](#).

Step 2 Shut down the node as described in [Shutting Down the Node, page 3-8](#).

Step 3 Decommission the node as described in [Decommissioning the Node Using Cisco UCS Manager, page 3-11](#).



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 the 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 as described in [Removing and Replacing the Node Top Cover, page 3-14](#).

- Step 7** Slide the clear password DIP switch from position 2 to the closed position (see [Figure 3-23](#)).
- 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, not just the service processor, 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** Slide the clear CMOS DIP switch from the closed position back to default position 2 (see [Figure 3-23](#)).



Note If you do not move the jumper, the CMOS settings are reset to the default every time that you power-cycle the node.

- Step 13** Replace the top cover to the node.
- Step 14** Replace the node in the rack, replace power cables, and then power on the node by pressing the **Power** button.
- Step 15** Recommission the node as described in [Recommissioning the Node Using Cisco UCS Manager, page 3-12](#).
- Step 16** Associate the node to its service profile as described in [Associating a Service Profile With an HX Node, page 3-12](#).
- Step 17** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, page 3-13](#).
-

Using the Clear CMOS DIP Switch

See [Figure 3-23](#) for the location of this DIP switch. 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 the Node Through vSphere With Cisco HX Maintenance Mode](#), page 3-10.
- Step 2** Shut down the node as described in [Shutting Down the Node](#), page 3-8.
- Step 3** Decommission the node as described in [Decommissioning the Node Using Cisco UCS Manager](#), page 3-11.



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 the 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 as described in [Removing and Replacing the Node Top Cover](#), page 3-14.
- Step 7** Slide the clear CMOS DIP switch from position 4 to the closed position (see [Figure 3-23](#)).
- 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, not just the service processor, 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** Move the clear CMOS DIP switch from the closed position back to default position 4 (see [Figure 3-23](#)).



Note

If you do not move the jumper, the CMOS settings are reset to the default every time that you power-cycle the node.

- Step 13** Replace the top cover to the node.

- Step 14** Replace the node in the rack, replace power cables, and then power on the node by pressing the **Power** button.
 - Step 15** Recommission the node as described in [Recommissioning the Node Using Cisco UCS Manager, page 3-12](#).
 - Step 16** Associate the node to its service profile as described in [Associating a Service Profile With an HX Node, page 3-12](#).
 - Step 17** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, page 3-13](#).
-

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.

Connecting and Powering On the Node (Standalone Mode)

The node is shipped with these default settings:

- The NIC mode is Shared LOM EXT.

Shared LOM EXT mode enables the 1-Gb Ethernet ports and the ports on any installed Cisco virtual interface card (VIC) to access 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 [Step 1](#) of the following procedure.

- The NIC redundancy is active-active. All Ethernet ports are utilized simultaneously.
- DHCP is enabled.
- IPv4 is enabled. You can change this to IPv6.

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

- Local setup—Use this procedure if you want to connect a keyboard and monitor to the node for setup. This procedure requires a KVM cable (Cisco PID N20-BKVM). See [Local Connection Procedure, page 3-68](#).
- Remote setup—Use this procedure if you want to perform setup through your dedicated management LAN. See [Remote Connection Procedure, page 3-69](#).



Note

To configure the node remotely, you must have a DHCP server on the same network as the node. 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 (see [Figure 1-1](#)). 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.

Local Connection Procedure

- Step 1** Attach a power cord to each power supply in your node, and then attach each power cord to a grounded AC power outlet. See [Power Specifications, page A-2](#) for power specifications.

Wait for approximately two minutes to let the node boot in standby power during the first bootup.

You can verify node power status by looking at the node Power Status LED on the front panel (see [External Features Overview, page 1-1](#)). The node 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 a USB keyboard and VGA monitor to the corresponding connectors on the rear panel (see [External Features Overview, page 1-1](#)).

- Connect an optional KVM cable (Cisco PID N20-BKVM) to the KVM connector on the front panel (see [External Features Overview, page 1-1](#) for the connector location). Connect your USB keyboard and VGA monitor to the KVM cable.
- 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.
This utility has two windows that you can switch between by pressing F1 or F2.
- Step 4** Continue with [Cisco IMC Configuration Utility Setup, page 3-70](#).
-

Remote Connection Procedure

- Step 1** Attach a power cord to each power supply in your node, and then attach each power cord to a grounded AC power outlet.
- Wait for approximately two minutes to let the node boot in standby power during the first bootup. You can verify node power status by looking at the node Power Status LED on the front panel (see [External Features Overview, page 1-1](#)). The node 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 (see [External Features Overview, page 1-1](#)).
- 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 server 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 Summary page, click **Launch KVM Console**. A separate KVM console window opens.
- Step 6** From the Cisco IMC Summary page, click **Power Cycle System**. The node 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.
- This utility has two windows that you can switch between by pressing F1 or F2.
- Step 9** Continue with [Cisco IMC Configuration Utility Setup, page 3-70](#).
-

Cisco IMC Configuration Utility Setup

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

- Step 1** Set NIC mode and NIC redundancy:
- a. Set the NIC mode to choose which ports to use to access Cisco IMC for node management (see [Figure 1-2](#) for identification of the ports):
 - 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.

In this mode, DHCP replies are returned to both the shared LOM ports and the Cisco card ports. If the node determines that the Cisco card connection is not getting its IP address from a Cisco UCS Manager node because the node is in standalone mode, further DHCP requests from the Cisco card are disabled. Use the Cisco Card NIC mode if you want to connect to Cisco IMC through a Cisco card in standalone mode.
 - Dedicated—The dedicated management port is used to access Cisco IMC. You must select a NIC redundancy and IP setting.
 - Shared LOM—The 1-Gb Ethernet ports are used to access Cisco IMC. You must select a NIC redundancy and IP setting.
 - Cisco Card—The ports on an installed Cisco UCS virtual interface card (VIC) are used to access Cisco IMC. You must select a NIC redundancy and IP setting.

See also the required VIC Slot setting below.
 - VIC Slot—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, slot 2 is the primary slot, but you can use slot 1.
 - If you select Riser2, slot 5 is the primary slot, but you can use slot 4.
 - If you select Flex-LOM, you must use an mLOM-style VIC in the mLOM slot.
 - b. Use this utility to change the NIC redundancy to your preference. This node has three possible NIC redundancy settings:
 - None—The Ethernet ports operate independently and do not fail over if there is a problem. This setting can be used only with the Dedicated NIC mode.
 - Active-standby—If an active Ethernet port fails, traffic fails over to a standby port.
 - Active-active—All Ethernet ports are utilized simultaneously. Shared LOM EXT mode can have only this NIC redundancy setting. Shared LOM and Cisco Card modes can have both Active-standby and Active-active settings.
- Step 2** 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 node. The MAC address is printed on a label on the rear of the node. This node has a range of six MAC addresses assigned to Cisco IMC. The MAC address printed on the label is the beginning of the range of six contiguous MAC addresses.

The static IPv4 and IPv6 settings include the following:

- The Cisco IMC IP address.

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

- Step 3** (Optional) Use this utility to make VLAN settings.
- Step 4** 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 5** (Optional) Set a hostname for the node.
- Step 6** (Optional) Enable dynamic DNS and set a dynamic DNS (DDNS) domain.
- Step 7** (Optional) If you check the Factory Default check box, the node reverts to the factory defaults.
- Step 8** (Optional) Set a default user password.
- Step 9** (Optional) Enable auto-negotiation of port settings or set the port speed and duplex mode manually.



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

- Step 10** (Optional) Reset port profiles and the port name.
- Step 11** Press **F5** to refresh the settings that you made. You might have to wait about 45 seconds until the new settings appear and the message, “Network settings configured” is displayed before you reboot the node in the next step.
- Step 12** Press **F10** to save your settings and reboot the node.



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

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

Node Specifications

This appendix lists the technical specifications for the node and includes the following sections:

- [Physical Specifications, page A-1](#)
- [Power Specifications, page A-2](#)
- [Environmental Specifications, page A-5](#)

Physical Specifications

[Table A-1](#) lists the physical specifications for the node.

Table A-1 *Physical Specifications*

Description	Specification
Height	3.4 in. (8.70 cm)
Width (including slam latches)	19.0 in. (48.26 cm)
Depth (length)	29.0 in. (73.70 cm)
Depth, including slam latches and power supply handles	31.5 in. (80.00 cm)
Maximum Weight (fully loaded)	62.7 lb. (28.4 Kg)

Power Specifications

The power specifications for the four power supply options are listed in the following sections:

- [650 W AC Power Supply, page A-2](#)
- [1200 W AC Power Supply, page A-3](#)
- [1400 W AC Power Supply, page A-4](#)
- [930 W DC \(Version 1\) Power Supply, UCSC-PSU-930WDC, page A-4](#)

You can get more specific power information for your exact node configuration by using the Cisco UCS Power Calculator:

<http://ucspowercalc.cisco.com>

650 W AC Power Supply

[Table A-2](#) lists the specifications for each 650 W AC power supply (Cisco part number UCSC-PSU2V2-650W=).

Table A-2 Power Supply 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	7.6 A at 100 VAC 3.65 A at 208 VAC
Maximum input volt-amperes	760 VA at 100 VAC
Maximum output power per PSU	650 W
Maximum inrush current	35 A (sub-cycle duration)
Maximum hold-up time	12 ms at 650 W
Power supply output voltage	12 VDC
Power supply standby voltage	3.3 VDC
Efficiency rating	Climate Savers Platinum Efficiency (80Plus Platinum certified)
Form factor	RSP1
Input connector	IEC320 C14

1200 W AC Power Supply

Table A-3 lists the specifications for each 1200 W AC power supply (Cisco part number UCSC-PSU2V2-1200W=).

Table A-3 Power Supply 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	11 A at 100 VAC 7 A at 200 VAC
Maximum input volt-amperes	1456 VA
Maximum output power per PSU	800 W at 100–120 VAC 1200 W at 200–240 VAC
Maximum inrush current	35 A (sub-cycle duration)
Maximum hold-up time	12 ms at 1200 W
Power supply output voltage	12 VDC
Power supply standby voltage	12 VDC
Efficiency rating	Climate Savers Platinum Efficiency (80Plus Platinum certified)
Form factor	RSP1
Input connector	IEC320 C14

1400 W AC Power Supply

Table A-4 lists the specifications for each 1400 W AC power supply (Cisco part number UCSC-PSU2-1400=).

Table A-4 1400 W Power Supply 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	8.5 A at 200 VAC
Maximum input volt-amperes	1700 VA
Maximum output power per PSU	1400 W at 200–240 VAC
Maximum inrush current	35 A (sub-cycle duration)
Maximum hold-up time	12 ms at 1200 W
Power supply output voltage	12 VDC
Power supply standby voltage	12 VDC
Efficiency rating	Climate Savers Platinum Efficiency (80Plus Platinum certified)
Form factor	RSP1
Input connector	IEC320 C14

930 W DC (Version 1) Power Supply, UCSC-PSU-930WDC

Table A-5 lists the specifications for each 930 W DC power supply (Cisco part number UCSC-PSU-930WDC=).

Table A-5 930 W DC Power Supply Specifications

Description	Specification
DC input voltage range	Nominal range: –48 to –60 VDC nominal (Range: –40 to –60 VDC)
Maximum DC input current	23 A at –48 VDC
Maximum input W	1104 W
Maximum output power per PSU	930 W
Maximum inrush current	35 A (sub-cycle duration)
Maximum hold-up time	8 ms at 930 W
Power supply output voltage	12 VDC
Power supply standby voltage	12 VDC
Efficiency rating	> 92% at 50% load

Table A-5 930 W DC Power Supply Specifications

Description	Specification
Form factor	RSP1 (2U and 4U systems)
Input connector	Removable connector block UCSC-CONN-930WDC=

Environmental Specifications

Table A-6 lists the environmental specifications for the node.

Table A-6 Environmental 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.
Temperature, nonoperating (when the node is in storage or is transported)	-40 to 149°F (-40 to 65°C)
Humidity (RH), operating	10 to 90%
Humidity, non-operating	5 to 93%
Altitude, operating	0 to 10,000 feet
Altitude, nonoperating (when the node is in storage or is transported)	0 to 40,000 feet
Sound power level Measure A-weighted per ISO7779 LwAd (Bels) Operation at 73°F (23°C)	5.8
Sound pressure level Measure A-weighted per ISO7779 LpAm (dBA) Operation at 73°F (23°C)	43

Power Cord Specifications

This appendix provides supported power cable specifications.

Supported Power Cords and Plugs

Each power supply has a separate power cord. Standard power cords or jumper power cords are available for connection to the node. The 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 provided with the node are supported.

[Table B-1](#) lists the power cords for the node power supplies.

Table B-1 Supported Power Cords for the Node

Description	Length	
	Feet	Meters
CAB-250V-10A-AR Power Cord, 250 VAC 10 A IRAM 2073 Plug Argentina	8.2	2.5
CAB-9K10A-AU 250 VAC 10 A 3112 Plug, Australia	8.2	2.5
CAB-250V-10A-CN Power Cord, 250 VAC 10 A GB 2009 Plug China	8.2	2.5
CAB-ACTW AC power cord, 250 V, 10 A, C13 plug Taiwan	7.5	2.3
CAB-9K10A-EU Power Cord, 250 VAC 10 A M 2511 Plug Europe	8.2	2.5

Table B-1 Supported Power Cords for the Node (continued)

Description	Length	
	Feet	Meters
CAB-250V-10A-ID Power Cord, 250 VAC 16A EL-208 Plug South Africa, United Arab Emirates, India	8.2	2.5
CAB-250V-10A-IS Power Cord, 250 VAC 10 A SI32 Plug Israel	8.2	2.5
CAB-9K10A-IT Power Cord, 250 VAC 10 A CEI 23-16 Plug Italy	8.2	2.5
CAB-9K10A-SW Power Cord, 250 VAC 10 A MP232 Plug Switzerland	8.2	2.5
CAB-9K10A-UK Power Cord, 250 VAC 10 A BS1363 Plug (13 A fuse) United Kingdom	8.2	2.5
CAB-AC-250V/13A Power Cord, 250 VAC 13 A IEC60320 Plug North America	6.6	2.0
CAB-N5K6A-NA Power Cord, 250 VAC 13 A NEMA 6-15 Plug, North America	8.2	2.5
CAB-9K12A-NA Power cord, 125 VAC, 13 A, NEMA 5-15 Plug North America	8.2	2.5
CAB-C13-CBN Cabinet Jumper Power Cord, 250 VAC 10 A, C13-C14 Connectors	2.2	0.68
CAB-C13-C14-2M Cabinet Jumper Power Cord, 250 VAC 10 A, C13-C14 Connectors	6.6	2.0
CAB-C13-C14-AC Cabinet Jumper Power Cord, 250 VAC 10 A, C13-C14 Connectors	9.8	3.0
CAB-48DC-40A-8AWG Power cord, -48 VDC 40 A, Three-socket Mini-Fit connector to three-wire pigtail 8 AWG	11.7	3.5

RAID Controller Considerations

This appendix contains the following sections:

- [Supported HBA and Required Cables, page C-1](#)
- [HBA Firmware Compatibility, page C-1](#)
- [HBA Cabling, page C-2](#)

Supported HBA and Required Cables

This node supports the HBA and cable requirements shown in [Table C-1](#).

Table C-1 *Cisco Hx240c RAID Options*

Controller	Style	Supported Server Version/ Maximum Front-Facing Drives Controlled	SCPM	RAID Levels	Required Cables
Cisco UCS 12G SAS Modular HBA	PCIe	24 internal drives	No	Non-RAID	(UCS-240CBLMR24=)

HBA Firmware Compatibility

Firmware on the HBA must be verified for compatibility with the current Cisco IMC and BIOS versions that are installed on the server. If not compatible, upgrade or downgrade the HBA firmware accordingly.

Use the procedures in the GUI or CLI [Cisco UCS Manager Firmware Management Guide](#) for your release.

HBA Cabling


This section includes the following topics:

- [Cable Routing, page C-2](#)
- [Cisco HX240c Server HBA Cabling Instructions, page C-3](#)

Cable Routing

The red line in [Figure C-1](#) shows the recommended cable routing path from an HBA in the dedicated slot.

Figure C-1 HBA Cable Paths



1	Drive backplane	3	Cisco HBA slot (dedicated PCIe slot on motherboard)
2	Expander (separate from the backplane)		

Cisco HX240c Server HBA Cabling Instructions

This section contains cabling instructions for controlling front-facing drives using the supported cables:
The node has a backplane and a separate SAS expander.

Cisco UCS 12G Modular HBA

This option can control up to 24 drives.

The required UCS-240CBLMR24= cable kit has one cable with a mini-SAS HD connector on each end.

Make the following connections:

1. Connect one end of the cable to the HBA.
2. Connect the other end of the cable to the single mini-SAS HD connector on the expander.

GPU Card Installation

This appendix contains the following topics:

- [GPU Card Configuration Rules, page D-1](#)
- [Requirement For All Supported GPUs: Memory-Mapped I/O Greater than 4 GB, page D-2](#)
- [Installing a GRID K1 or K2 GPU Card, page D-3](#)
- [Installing a Tesla M60 GPU Card and 300 W GPU Conversion Kit, page D-5](#)
- [Installing Drivers to Support the NVIDIA GPU Cards, page D-23](#)

GPU Card Configuration Rules

Observe the following rules when installing a GPU in the node:



Caution

When using NVIDIA GPU cards, you must preserve at least 10 mm of space between nodes to ensure adequate air flow. When using GPU cards, the operating temperature range is 32° to 95° F (0° to 35° C).

- Only one GPU card per node is supported. The GPU card must be installed in PCIe riser 2, slot 5.
- All GPU cards require two CPUs in the node.
- NVIDIA GRID K1 and GRID K2 GPUs require two power supplies in the node (minimum 1200 W, recommended 1400 W).
- NVIDIA Tesla M60 GPUs require two 1400 W power supplies in the node.
- NVIDIA GPUs can support only less-than 1 TB of memory in the node. Therefore, do not install more than fourteen 64-GB DIMMs when using an NVIDIA GPU card in this node.

Requirement For All Supported 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 from a different setting, enter the BIOS Setup Utility by pressing F2 when prompted during bootup.

Cisco UCS Manager Controlled Node

If the node is integrated with Cisco UCS Manager and controlled by a service profile, this setting is not enabled by default in the service profile. You must enable it with a BIOS policy in your service profile.

-
- Step 1** Refer to the Cisco UCS Manager configuration guide (GUI or CLI) for your release for instructions on configuring service profiles:
- [Cisco UCS Manager Configuration Guides](#)
- Step 2** Refer to the chapter on [Configuring Server-Related Policies > Configuring BIOS Settings](#).
- Step 3** In the section of your profile for [PCI Configuration BIOS Settings](#), set `Memory Mapped IO Above 4GB Config` to one of the following:
- **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 GRID K1 or K2 GPU Card

Use this section to install or replace a GPU card that draws less than 300 W power. Use this section to install the following GPUs:

- GRID K1
- GRID K2

-
- Step 1** Put the node in Cisco HX Maintenance mode as described in [Shutting Down the Node Through vSphere With Cisco HX Maintenance Mode](#), page 3-10.
- Step 2** Shut down the node as described in [Shutting Down the Node](#), page 3-8.
- Step 3** Decommission the node as described in [Decommissioning the Node Using Cisco UCS Manager](#), page 3-11.



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 the 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 as described in [Removing and Replacing the Node Top Cover](#), page 3-14.
- Step 7** Remove PCIe riser 2 and any existing GPU card in slot 5:
- a. Lift straight up on both ends of the riser to disengage its circuit board from the socket on the motherboard. Set the riser on an antistatic mat.
 - b. On the bottom of the riser, loosen the single thumbscrew that holds the securing plate. See [Figure D-1](#).
 - c. Swing open the securing plate and remove it from the riser to provide access.
 - d. Swing open the card-tab retainer that secures the back-panel tab of the card (see [Figure D-1](#)).
 - e. Pull evenly on both ends of the GPU card to disengage it from the socket on the PCIe riser (or remove a blanking panel) and then set the card aside.
- Step 8** Install your GPU card into riser 2, PCIe slot 5. See [Figure 3-17](#) for the riser and slot locations.
- 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 card power cable (UCSC-GPUCBL-240M4) into the GPU card and into the GPU POWER connector on the PCIe riser (see [Figure D-1](#)).
 - c. Return the securing plate to the riser. Insert the two hinge-tabs into the two slots on the riser, and then swing the securing plate closed.
 - d. Tighten the single thumbscrew that holds the securing plate.
 - e. Close the card-tab retainer (see [Figure D-1](#)).
 - f. Position the PCIe riser over its socket on the motherboard and over its alignment features in the chassis (see [Figure 3-16](#)).
 - g. 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.
- Step 10** Replace the node in the rack, replace power cables, and then power on the node by pressing the **Power** button.
- Step 11** Recommission the node as described in [Recommissioning the Node Using Cisco UCS Manager, page 3-12](#).
- Step 12** Associate the node to its service profile as described in [Associating a Service Profile With an HX Node, page 3-12](#).
- Step 13** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, page 3-13](#).

Figure D-1 PCIe Riser Securing Features



1	Securing plate hinge-tabs	3	GPU card power connector
2	Securing plate thumbscrew (knob not visible on underside of plate)	4	Card-tab retainer in open position

Installing a Tesla M60 GPU Card and 300 W GPU Conversion Kit

Installation Overview

When installing an NVIDIA M60 GPU and the GRID software, use the topics in this section in the following order:

1. Install the hardware.
 - [Installing the NVIDIA M60 Hardware, page D-5](#)
2. Register your product activation keys with NVIDIA.
 - [NVIDIA GRID License Server Overview, page D-13](#)
 - [Registering Your Product Activation Keys With NVIDIA, page D-14](#)
3. Download the GRID software suite.
 - [Downloading the GRID Software Suite, page D-14](#)
4. Install the GRID License Server software to a host.
 - [Installing NVIDIA GRID License Server Software, page D-15](#)
5. Generate licenses on the NVIDIA Licensing Portal and download them.
 - [Installing GRID Licenses From the NVIDIA Licensing Portal to the License Server, page D-17](#)
6. Manage your GRID licenses.
 - [Managing GRID Licenses, page D-19](#)
7. Decide whether to use the GPU in compute mode or graphics mode.
 - [Switching Between Compute Mode and Graphics Mode, page D-21](#)

Installing the NVIDIA M60 Hardware

The NVIDIA Tesla M60 GPU requires a hardware conversion kit.

300 W GPU Card Conversion Kit

**Caution**

Do not operate the node with the 300W GPU kit installed, but no GPU card installed. The kit has been designed to provide adequate airflow for cooling only when at least one GPU card is installed.

The contents of the conversion kit are as follows:

- Replacement fan-module fan cage
- CPU cleaning kit
- Low-profile CPU heatsinks (two)
- Replacement air baffle (includes base, bridge, and filler panel)
- 300 W GPU card front support bracket
- 300 W GPU card straight power cable

**Note**

Your 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 node; do not use the Y-cable, which is used for connecting the GPU card to external devices only.

Installing the NVIDIA M60 GPU Card

- Step 1** Put the node in Cisco HX Maintenance mode as described in [Shutting Down the Node Through vSphere With Cisco HX Maintenance Mode](#), page 3-10.
- Step 2** Shut down the node as described in [Shutting Down the Node](#), page 3-8.
- Step 3** Decommission the node as described in [Decommissioning the Node Using Cisco UCS Manager](#), page 3-11.

**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 the 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 as described in [Removing and Replacing the Node Top Cover](#), page 3-14.
- Step 7** Remove the plastic air-baffle that covers the CPUs and DIMMs.
- Step 8** Remove the existing node fan cage (see [Figure D-2](#)):
 - a. Open the plastic locking-lever at each end of the existing fan cage to the upright 90-degree position.
 - b. Lift the existing fan cage with fan modules from the node. Set the fan cage with fan modules aside.
- Step 9** Install the new empty node fan cage from the conversion kit:
 - a. Open the plastic locking-lever at each end of the new fan cage to the upright 90-degree position.
 - b. Set the new fan cage into the guides on the chassis walls and then lower the cage evenly.
 - c. Close the plastic locking-lever at each end of the fan cage to the flat, locked position.
- Step 10** Move the six fan modules from the old fan cage to the new fan cage that you just installed:
 - a. Pinch the two finger latches on each fan module together, then lift up on the module to remove it from the cage (see [Figure D-2](#)).
 - b. Set the fan module in an open slot in the new fan cage, aligning the connector on the bottom of the fan module with the connector on the motherboard.

**Note**

The arrow label on the top of the fan module, which indicates the direction of airflow, should point toward the rear of the node.

- c. Press down gently on the fan module until the latch clicks and locks in place.
- d. Repeat until you have moved all fan modules into the new fan cage.

Figure D-2 Fan Cage and Fan Modules



1	Finger latches (on each fan module)	3	SuperCap power module position on removable air baffle (air baffle not shown)
2	Fan cage plastic locking-levers		

Step 11 Remove the existing heatsink from each CPU.

- a. Use a Number 2 Phillips-head screwdriver to loosen the four captive screws that secure the heatsink.



Note Alternate loosening each screw evenly to avoid damaging the heatsink or CPU.

- b. Lift the heatsink off of the CPU and set it aside.

Step 12 Use the heatsink cleaning kit that comes with the conversion kit to clean the existing thermal grease from the top surface of each CPU.

Step 13 Install the low-profile replacement heatsinks that come with the conversion kit:


- a. Remove the protective tape from the pre-applied pad of thermal grease that is on the underside of the new heatsink.
- b. Align the four heatsink captive screws with the motherboard standoffs, and then use a Number 2 Phillips-head screwdriver to tighten the captive screws evenly.



Note Alternate tightening each screw evenly to avoid damaging the heatsink or CPU.

Step 14 Set the base of the replacement air baffle into the node (see [Figure D-3](#)).


Figure D-3 Air Baffle Base



1	Fan cage	4	GPU front support bracket latches
2	Air baffle base	5	PCIe risers
3	Holder for RAID battery on air baffle		

- Step 15** Install the NVIDIA M60 GPU card front support bracket to the GPU card (see [Figure D-4](#)):
- Remove the three thumbscrews from the front end of the GPU card.
 - Set the three holes in the GPU card front support bracket over the three screw holes.
 - Insert and tighten the three thumbscrews into the three screw holes to secure the front bracket to the GPU card.

Figure D-4 Front Support Bracket



1	Front support bracket	2	Thumbscrews (three)
---	-----------------------	---	---------------------

- Step 16** Remove PCIe riser 2 from the node.
- Grasp the top of the riser and lift straight up on both ends to disengage its circuit board from the socket on the motherboard.
 - Set the riser on an antistatic surface.
 - On the bottom of the riser, loosen the single thumbscrew that holds the securing plate (see [Figure D-1](#)).
 - Swing open the securing plate and remove it from the riser to provide access.
 - Swing open the card-tab retainer (see [Figure D-1](#)).
- Step 17** Install the M60 GPU card into PCIe riser 2, slot 5. See [Figure 3-17](#) for the riser and slot locations.
- 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. The straight-cable connectors are color-coded. Connect the GPU card power cable *black* connector into *black* connector on the GPU card and the *white* connector into the *white* GPU POWER connector on the PCIe riser (see [Figure D-1](#)).

**Caution**

Do not reverse the GPU 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.

**Note**

Your 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 node; do not use the Y-cable, which is used for connecting the GPU card in external devices only.

- c. Close the card-tab retainer (see [Figure D-1](#)).
- d. Return the securing plate to the riser. Insert the two hinge-tabs into the two slots on the riser, and then swing the securing plate closed.
- e. Tighten the single thumbscrew that holds the securing plate.


Step 18 Install the PCIe riser with GPU card back into the node:

- a. Position the PCIe riser over its socket on the motherboard and over its alignment features in the chassis (see [Figure 3-16](#)).
- b. Carefully push down on both ends of the PCIe riser to fully engage its circuit board connector with the socket on the motherboard.

Step 19 Insert the front support bracket that you installed to the card in [Step 15](#) into the latch that is on the air baffle base that you installed in [Step 14](#).

- a. Pinch the latch release tab (see [Figure D-5](#)) and hinge the latch toward the front of the node.
- b. Hinge the latch back down so that its lip closes over the front edge of the support bracket that is attached to the GPU card (see [Figure D-5](#)).
- c. Ensure that the latch release tab clicks and locks the latch in place.


Figure D-5 GPU Front Support Bracket Inserted to Air Baffle Latch



1	Front end of GPU	3	Lip on support bracket latch
2	Front support bracket attached to GPU	4	Latch release tab

- Step 20** Install the filler panel to the air baffle base as shown in [Figure D-6](#). Slide the filler panel edges into the slots as shown. The filler panel shunts more airflow to PCIe riser 2 where the M60 GPU is installed.
- Step 21** Install the bridge to the air baffle. Ridges on the bridge legs slide into slots on the air baffle base (see [Figure D-6](#)).

Figure D-6 Air Baffle Filler Panel and Bridge



1	Air baffle base	3	Bridge
2	Filler panel		


- Step 22** Replace the top cover to the node.
- Step 23** Replace the node in the rack, replace power cables, and then power on the node by pressing the **Power** button.
- Step 24** Recommission the node as described in [Recommissioning the Node Using Cisco UCS Manager, page 3-12](#).
- Step 25** Associate the node to its service profile as described in [Associating a Service Profile With an HX Node, page 3-12](#).
- Step 26** After ESXi reboot, exit HX Maintenance mode as described in [Exiting HX Maintenance Mode, page 3-13](#).
- Step 27** To set up licenses for the M60 GPU, continue with [NVIDIA GRID License Server Overview, page D-13](#).

NVIDIA GRID License Server Overview

The NVIDIA Tesla M60 GPU combines 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 (see [Figure D-7](#)).

Figure D-7 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 (see [Table D-1](#)).

Table D-1 GRID Licensing Editions

GRID License Edition	GRID Features
GRID Virtual GPU (vGPU)	<ul style="list-style-type: none"> Virtual GPUs for business desktop computing
GRID Virtual Workstation	<ul style="list-style-type: none"> Virtual GPUs for midrange workstation computing
GRID Virtual Workstation – Extended	<ul style="list-style-type: none"> 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 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 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.

a. Use the following command to install Tomcat:

```
sudo yum install java
```

b. Enable the Tomcat service for automatic startup on boot:

```
sudo systemctl enable tomcat.service
```

c. Start the Tomcat service:

```
sudo systemctl start tomcat.service
```

d. 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:

a. Unpack the License Server tar file using the following command:

```
tar xfz 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. See [Accessing the GRID License Server Management Interface, page D-17](#).
- 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. See [Accessing the GRID License Server Management Interface, page D-17](#).
- 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 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. See [Accessing the GRID License Server Management Interface, page D-17](#).
 - 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. See [Accessing the GRID License Server Management Interface, page D-17](#).
 - 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

To acquire a GRID license on Windows, use the following procedure.

-
- 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

To acquire a GRID license on Linux, use the following procedure.

-
- 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 `gridd.conf`, 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=1
```

Switching Between Compute Mode and Graphics Mode

Overview of GPU Modes

The NVIDIA Tesla M60 GPU is shipped in compute mode, which is optimized for high-performance compute (HPC) applications. However, while compute mode is best for HPC usage, it can cause compatibility issues with OS and hypervisors if you use the GPU primarily as a graphics device.

The mode is determined at power-on, from settings stored in the GPU's non-volatile memory. You can use the command-line tool `gpumodeswitch` to toggle the GPU between compute mode and graphics mode. [Table D-2](#) and [Table D-3](#) compare the compute mode and graphic mode default settings.

Table D-2 *Compute Mode Default Settings*

Setting	Value	Notes
Classcode	3D Controller	This classcode tells the OS that the GPU is not intended for use as the primary display device.
Memory base address register (BAR)	8 GB	A large BAR is exposed for direct access to the frame buffer from the CPU and PCIe devices.
I/O base BAR	Disabled	The GPU does not consume any legacy I/O resources when used as a non-display device.
Error-correcting code (ECC) protection	Enabled	ECC is enabled on the GPU frame buffer to protect against single- and multi-bit memory errors.

Table D-3 *Graphic Mode Default Settings*

Setting	Value	Notes
Classcode	VGA Controller	This classcode tells the OS that the GPU can function as the primary display device.
Memory base address register (BAR)	256 MB	A smaller BAR is exposed for direct access to the frame buffer from the CPU and PCIe devices.
I/O base BAR	Enabled	The GPU exposes an I/O BAR to claim the resources required to operate as a VGA controller.
Error-correcting code (ECC) protection	Disabled	ECC protection is disabled.

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`

This command 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 Server BIOS Firmware, page D-23](#)
- 2. [Activating the Server BIOS Firmware, page D-24](#)
- 3. [Updating the GPU Drivers, page D-24](#)

1. Updating the Server BIOS Firmware

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

**Note**

You must perform 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 UCS Manager 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 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:
- a. From the Firmware Version drop-down list, select the firmware version to which you want to update the endpoint.
 - b. 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.
-

2. Activating the Server 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.
 - Click OK.
-

3. Updating the GPU 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:
- 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.
-