



Cisco UCS C240 M7 Server Installation and Service Guide

First Published: 2023-05-04

Last Modified: 2024-02-23

Americas Headquarters

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



CONTENTS

PREFACE

Preface	ix
Bias-Free Documentation	ix
Full Cisco Trademarks with Hardware License	ix
Communications, Services, and Additional Information	xi

CHAPTER 1

System Overview	1
Overview	1
External Features	3
PCIe Risers	12
Summary of Server Features	21
Serviceable Component Locations	24

CHAPTER 2

Installing the Server	27
Preparing for Installation	27
Installation Warnings and Guidelines	27
Rack Requirements	29
Installing the Server in a Rack	30
Installing the Cable Management Arm (Optional)	33
Reversing the Cable Management Arm (Optional)	34
Initial Server Setup	35
Connecting to the Server Locally For Setup	36
Connecting to the Server Remotely For Setup	37
Setting Up the System With the Cisco IMC Configuration Utility	38
NIC Mode and NIC Redundancy Settings	40
Updating the BIOS and Cisco IMC Firmware	41
Accessing the System BIOS	42

Smart Access Serial 42

Smart Access USB 43

CHAPTER 3**Maintaining the Server 45**

Status LEDs and Buttons 45

Front-Panel LEDs 46

Rear-Panel LEDs 49

Internal Diagnostic LEDs 51

Preparing For Component Installation 51

Required Equipment For Service Procedures 51

Shutting Down and Removing Power From the Server 52

Shutting Down Using the Power Button 52

Shutting Down Using The Cisco IMC CLI 52

Shutting Down Using The Cisco IMC GUI 53

Removing the Server Top Cover 53

Serial Number Location 55

Hot Swap vs Hot Plug 55

Replacing the Air Duct 55

Removing the Air Duct 58

Installing the Air Duct 59

Removing and Replacing Components 61

Serviceable Component Locations 61

Replacing Front-Loading SAS/SATA Drives 64

Front-Loading SAS/SATA Drive Population Guidelines 64

4K Sector Format SAS/SATA Drives Considerations 65

Replacing a Front-Loading SAS/SATA Drive 66

Replacing Rear-Loading SAS/SATA Drives 67

Rear-Loading SAS/SATA Drive Population Guidelines 67

Replacing a Rear-Loading SAS/SATA Drive 68

Basic Troubleshooting: Reseating a SAS/SATA Drive 69

Reseating a SAS/SATA Drive 70

Replacing Front-Loading NVMe SSDs 71

Front-Loading NVMe SSD Population Guidelines 71

Front-Loading NVMe SSD Requirements and Restrictions 71

Enabling Hot-Plug Support in the System BIOS	72
Replacing a Front-Loading NVMe SSD	72
Replacing Rear-Loading NVMe SSDs	74
Rear-Loading NVMe SSD Population Guidelines	74
Rear-Loading NVMe SSD Requirements and Restrictions	75
Replacing a Rear-Loading NVMe SSD	75
Replacing Fan Modules	77
Replacing the Fan Tray	79
Removing the Fan Tray	80
Installing the Fan Tray	81
Replacing CPUs and Heatsinks	81
CPU Configuration Rules	82
Tools Required For CPU Replacement	83
Removing CPUs and Heat Sinks	84
Installing the CPUs and Heatsinks	89
Upgrading to Intel Xeon Fifth Generation CPUs	93
Additional CPU-Related Parts to Order with RMA Replacement CPUs	94
Additional CPU-Related Parts to Order with RMA Replacement System Chassis	95
Replacing Memory DIMMs	95
DIMM Population Rules and Memory Performance Guidelines	96
Replacing DIMMs	98
Replacing a Mini-Storage Module	99
Replacing a Mini-Storage Module Carrier	99
Replacing an M.2 SSD in a Mini-Storage Carrier For M.2	100
Replacing the RTC Battery	101
Replacing Power Supplies	101
Supported Power Supplies	102
Replacing a PCIe Riser	106
Replacing a PCIe Card	107
PCIe Slot Specifications	107
Replacing a PCIe Card	108
Cisco Virtual Interface Card (VIC) Considerations	110
Replacing an mLOM Card	111
Replacing an OCP Card	112

Cisco VIC mLOM and OCP Card Replacement Considerations	112
Removing an OCP Card	113
Installing an OCP Card	116
Replacing a SAS Storage Controller Card (RAID or HBA)	118
Storage Controller Card Firmware Compatibility	118
Removing the Dual Storage Controller Cards	119
Installing the Dual Storage Controller Cards	123
Removing the Storage Controller Card	126
Installing the Storage Controller Card	130
Verify Cabling	133
Replacing the Supercap (RAID Backup)	134
Replacing a Boot-Optimized M.2 RAID Controller Module	137
Cisco Boot-Optimized M.2 RAID Controller Considerations	137
Replacing a Cisco Boot-Optimized M.2 RAID Controller	139
Replacing a Chassis Intrusion Switch	141
Installing a Trusted Platform Module (TPM)	142
TPM Considerations	142
Installing TPM Hardware	142
Enabling the TPM in the BIOS	143
Enabling the Intel TXT Feature in the BIOS	143
Service Headers and Jumpers	144
Using the Clear CMOS Switch (SW4, Switch 9)	145
Using the Clear BIOS Password Switch (SW4, Switch 6)	146
Using the Boot Alternate Cisco IMC Image Header (CN5, Pins 1-2)	147
CHAPTER 4	Recycling Server Components 149
Server Recycling and E-Waste	149
Battery Warning	149
Recycling the Main Motherboard PCB Assembly (PCBA)	150
Recycling the Front Mezzanine Module PCBA (NVMe Backplane)	152
Recycling the Front Mezzanine Module PCBA (SAS Backplane)	154
Recycling the Server KVM PCBA	156
Recycling the Server Front Panel PCBA	160

APPENDIX A**Server Specifications 165**

- Server Specifications 165
 - Physical Specifications 165
 - Environmental Specifications 166
- Power Specifications 167
 - 1050 W DC Power Supply 167
 - 1200 W AC Power Supply 168
 - 1600 W AC Power Supply 169
 - 2300 W AC Power Supply 169
- Power Cord Specifications 170

APPENDIX B**Storage Controller Considerations 173**

- Supported Storage Controllers and Cables 173
- Storage Controller Card Firmware Compatibility 174
- RAID Backup (Supercap) 175
- Mixing Drive Types in RAID Groups 175
- Storage Controller Cable Connectors and Backplanes 175
- For More RAID Utility Information 176

APPENDIX C**GPU Card Installation 179**

- Server Firmware Requirements 179
- GPU Card Configuration Rules 180
- Requirement For All GPUs: Memory-Mapped I/O Greater Than 4 GB 182
- Installing a Double-Wide GPU Card 183
- Replacing a Heatsink 186
 - Removing a Heat Sink 186
 - Installing a Heatsink 188
- Replacing the GPU Air Blocker 190
 - Removing the Riser 2 GPU Air Blocker 191
 - Installing the Riser 2 GPU Air Blocker 192
 - Removing the Riser 3 GPU Air Blocker 194
 - Installing the Riser 3 GPU Air Blocker 195
- Installing GRID Licenses From the NVIDIA Licensing Portal to the License Server 196

Reading Your License Server's MAC Address	196
Installing Licenses From the Licensing Portal	197
Viewing Available GRID Licenses	197
Viewing Current License Usage	197
Managing GRID Licenses	198
Acquiring a GRID License on Windows	198
Acquiring a GRID License on Linux	198
Using gpumodeswitch	199
Installing Drivers to Support the GPU Cards	200
1. Updating the Server BIOS	200
2. Updating the GPU Card Drivers	200

APPENDIX D	Installation for Cisco Software Management	203
	Installation For Cisco UCS Manager Integration	203



Preface

This preface contains the following topics:

- [Bias-Free Documentation](#), on page ix
- [Full Cisco Trademarks with Hardware License](#), on page ix
- [Communications, Services, and Additional Information](#), on page xi

Bias-Free Documentation



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

Full Cisco Trademarks with Hardware License

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

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

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated

in a commercial environment. This equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

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

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

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

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

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

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

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

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

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

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

Communications, Services, and Additional Information

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

Cisco Bug Search Tool

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



CHAPTER 1

System Overview

This chapter contains the following topics:

- [Overview, on page 1](#)
- [External Features, on page 3](#)
- [Summary of Server Features, on page 21](#)
- [Serviceable Component Locations, on page 24](#)

Overview

The Cisco UCS C240 M7 is a stand-alone 2U rack server chassis that can operate in both standalone environments and as part of the Cisco Unified Computing System (Cisco UCS).

Each Cisco UCS C240 M7 has two CPU sockets that can support the following Intel® Xeon® Scalable Processors, in either one or two CPU configurations.

- Fourth Generation Intel Xeon Scalable Server processors
- Fifth Generation Intel Xeon Scalable Server processors

Additionally, the server supports the following features with one CPU or two identical CPUs:

- 32 DDR5 DIMMs (RDIMM), up to 5600 MHz (1DPC) and 4400 MHz (2 DPC) support for RDIMMs. 16 DIMMs are supported per CPU for a total system memory of 8 TB (up to 256 GB DDR5 DIMMs) on servers with both Intel Fourth and Fifth Generation CPUs.
- DDR5 DIMM capacities vary based on the CPU type for the compute node. For more information, see [DIMM Population Rules and Memory Performance Guidelines, on page 96](#).
 - Intel Fourth Generation Xeon Scalable Server Processors support 16, 32, 64, 128, and 256 GB DDR5 DIMMs
 - Intel Fifth Generation Xeon Scalable Server Processors support 16, 32, 64, 96, and 128 GB DDR5 DIMMs.
- The server's DIMM configuration differs depending on which generation of CPU is populated on the server:
 - With Fourth Generation Intel Xeon Scalable Server processors, the compute node supports DDR5 DIMMs up to 4800 MT/s with 1DPC, and up to 4400 MT/s with 2DPC

- With Fifth Generation Intel Xeon Scalable Server processors, the compute node supports DDR5 DIMMs up to 5600 MT/s with 1 DPC, and up to 4400 MT/s with 2DPC
- The servers have different supported configurations of small form factor (SFF) front-loading drives.
- Up to 2 M.2 SATA RAID cards for server boot.
- Rear Storage risers (2 slots each)
- Rear PCIe risers
- Internal slot for a 24 G Tri-Mode RAID controller with SuperCap for write-cache backup, or for a SAS HBA.
- One mLOM/VIC card provides 10/25/40/50/100/200 Gbps. The following mLOMs are supported:
 - Cisco UCS VIC 15427 Quad Port CNA MLOM (UCSC-M-V5Q50GV2) supports:
 - a x16 PCIe Gen4 Host Interface to the rack server
 - four 10G/25G/50G SFP+/SFP28/SFP56 ports
 - 4GB DDR4 Memory, 3200 MHz
 - Integrated blower for optimal ventilation
 - Secure boot support
 - Cisco UCS VIC 15425 Quad Port 10G/25G/50G SFP56 CNA PCIe (UCSC-P-V5Q50G-D)
 - a x16 PCIe Gen4 Host Interface to the rack server
 - Four 10G/25G/50G QSFP56 ports
 - 4GB DDR4 Memory, 3200MHz
 - Integrated blower for optimal ventilation
 - Cisco UCS VIC 15237 Dual Port 40G/100G/200G QSFP56 mLOM (UCSC-M-V5D200GV2) supports:
 - a x16 PCIe Gen4 Host Interface to the rack server
 - two 40G/100G/200G QSFP/QSFP28/QSFP56 ports
 - 4GB DDR4 Memory, 3200 MHz
 - Integrated blower for optimal ventilation
 - Secure boot support
 - Cisco UCS VIC 15235 Dual Port 40G/100G/200G QSFP56 CNA PCIe (UCSC-P-V5D200G-D)
 - a x16 PCIe Gen4 Host Interface to the rack server
 - two 40G/100G/200G QSFP56 ports
 - 4GB DDR4 Memory, 3200MHz
 - Integrated blower for optimal ventilation

- Two power AC supplies (PSUs) that support N+1 power configuration and cold redundancy.
- Six modular, hot swappable fans.

Server Configurations, 24 SFF SAS/SATA

The SFF 24 SAS/SATA/U.3 NVMe configuration (UCSC-C240-M7SX) can be ordered as either an I/O-centric configuration or a storage centric configuration. This server supports the following:

- A maximum of 24 small form-factor (SFF) drives, with a 24-drive backplane.
 - Front-loading drive bays 1—24 support 2.5-inch SAS/SATA/U.3 drives as SSDs or HDDs.
 - Optionally, drive bays 1—4 can support 2.5-inch NVMe SSDs. In this configuration, any number of NVMe drives can be installed up to the maximum of 4.



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

- Drive bays 5 —24 support SAS/SATA/U.3 SSDs or HDDs only; no U.2 NVMe.
- Optionally, the rear-loading drive bays support four 4 2.5-inch SAS/SATA or NVMe drives.

Server Configurations, 24 NVMe

The SFF 24 NVMe configuration (UCSC-C240-M7SN) can be ordered as an NVMe-only server. The NVMe-optimized server requires two CPUs. This server supports the following:

- A maximum of 24 SFF NVMe drives as SSDs with a 24-drive backplane, NVMe-optimized.
 - Front-loading drive bays 1—24 support 2.5-inch NVMe PCIe SSDs *only*.
 - Optionally, the rear-loading drive bays support four 2.5-inch NVMe SSDs *only*. These drive bays are in the left and right (Riser 1 and Riser 3) of the rear panel.

External Features

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

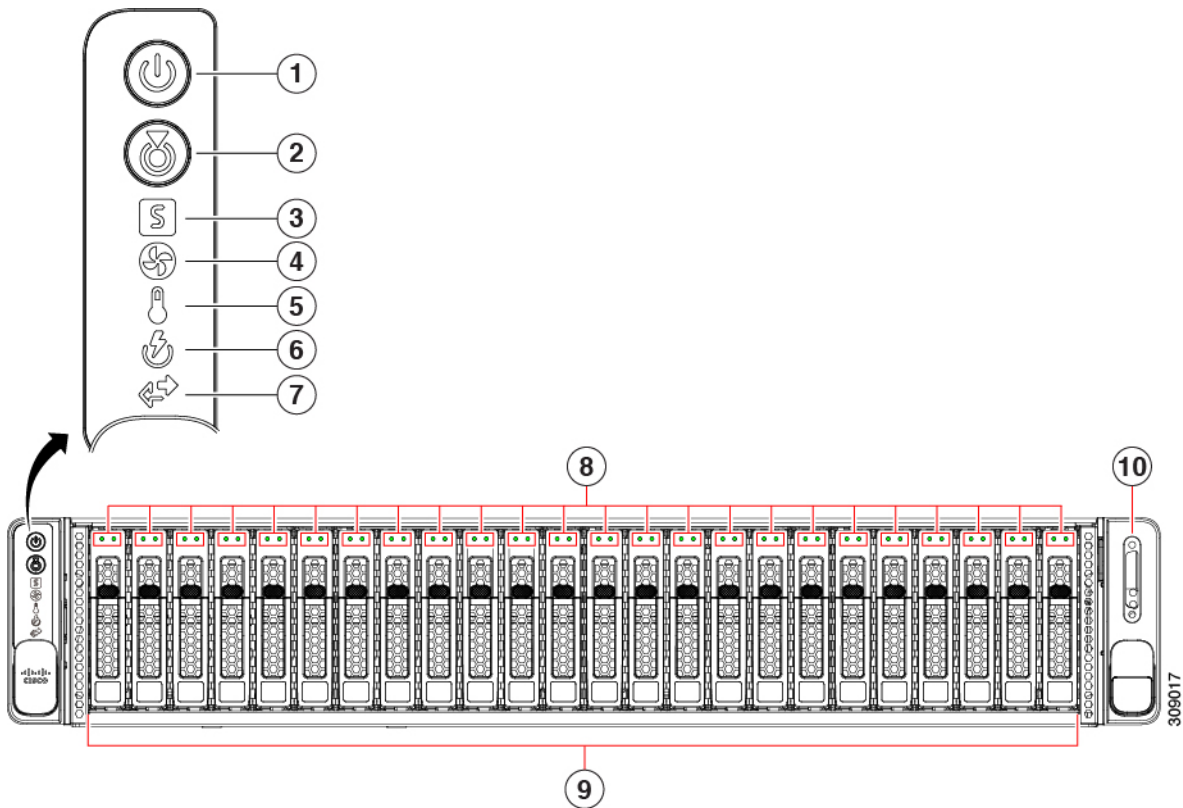
For definitions of LED states, see [Front-Panel LEDs, on page 46](#).

Cisco UCS C240 M7 Server 24 SAS/SATA Front Panel Features

The following figure shows the front panel features of Cisco UCS C240 M7SX, which is the small form-factor (SFF), 24 SAS/SATA/U.3 drive version of the server. Front-loading drives can be mix and match in slots 1 through 4 to support up to four SFF NVMe or SFF SAS/SATA drives. UCS C240 M7 servers with any number of NVMe drives must be dual CPU systems.

This configuration can support up to 4 optional universal HDD drives in the rear PCIe slots (riser 1 and Riser 3).

Figure 1: Cisco UCS C240 M7 Server 24 SAS/SATA Front Panel



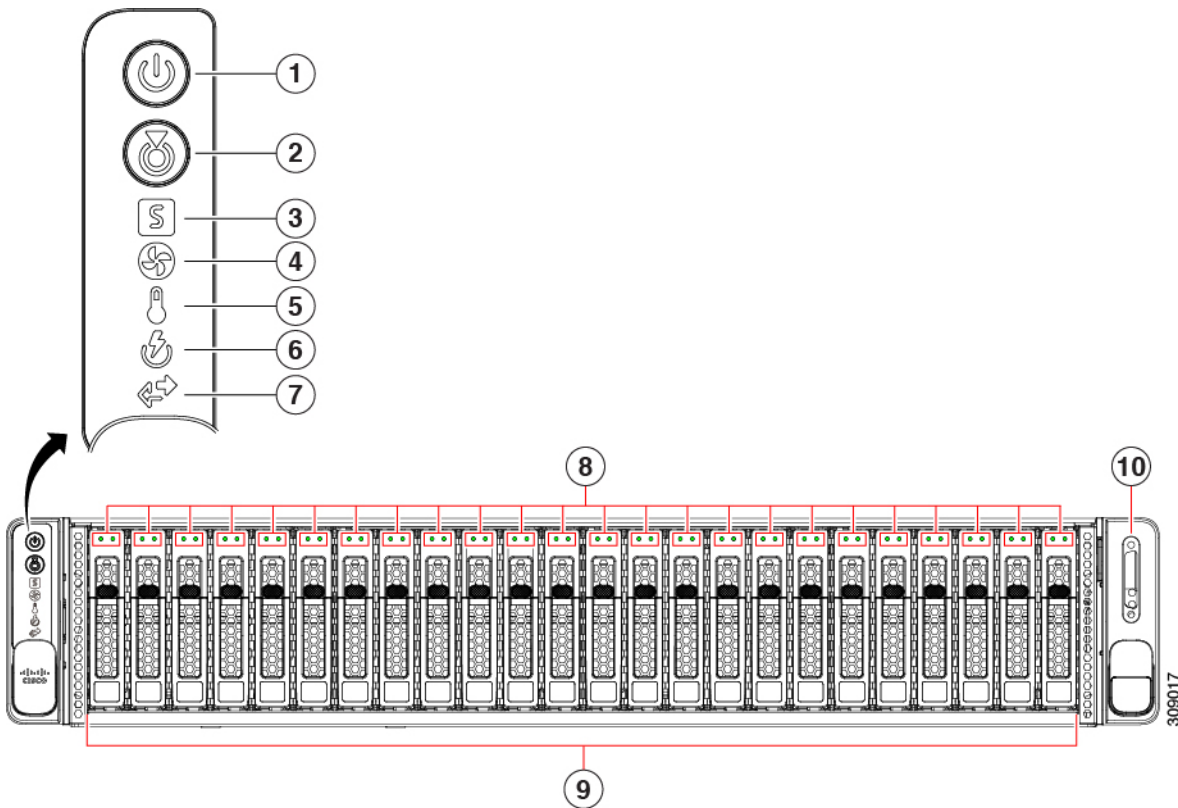
1	Power Button/Power Status LED	2	Unit Identification Button/Unit Identification LED
3	System Status LEDs	4	Fan Status LED
5	Temperature Status LED	6	Power Supply Status LED
7	Network Link Activity LED	8	Drive Status LEDs
9	<p>NVMe Drive Bays, front loading</p> <p>Drive bays 1—24 support front-loading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

Cisco UCS C240 M7 Server 24 NVMe Drives Front Panel Features

The following figure shows the front panel features of Cisco UCS C240 M7SN, which is the small form-factor (SFF) drive, 24 NVMe drive version of the server. Front-loading drives are all NVMe; SAS/SATA drives are not supported. UCS C240 M7 servers with any number of NVMe drives must be dual CPU systems.

This configuration can support up to 4 optional NVMe 2.5-inch drives in the rear PCIe slots (riser 1 and riser 3).


Figure 2: Cisco UCS C240 M7 Server 24 NVMe Front Panel



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

Common Rear Panel Features

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



473195


1	Rear hardware configuration options: <ul style="list-style-type: none"> • For I/O-Centric, these are PCIe slots. • For Storage-Centric, these are storage drives bays. This illustration shows the slots unpopulated	2	Power supplies (two, redundant as 1+1) See Power Specifications, on page 167 for specifications and supported options.
3	VGA video port (DB-15 connector)	4	Serial port (RJ-45 connector)
5	One dedicated 1Gbps management port	6	USB 3.0 ports, 2
7	Rear unit identification button/LED	8	Modular LAN-on-motherboard (mLOM) or OCP card slot (x16). This slot can contain either a Cisco mLOM or an Intel X710 OCP 3.0 card.

Cisco UCS C240 M7 Server 24 Drive Rear Panel, I/O Centric

The Cisco UCS C240 M7 24 SAS/SATA SFF version has a rear configuration option, for either I/O (I/O Centric) or Storage (Storage Centric) with the I/O Centric version of the server offering PCIe slots and the Storage Centric version of the server offering drive bays.

The following illustration shows the rear panel features for the I/O Centric version of the Cisco UCS C240 M7SX.

- For features common to all versions of the server, see [Common Rear Panel Features](#).
- For definitions of LED states, see [Rear-Panel LEDs, on page 49](#).



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

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

Table 1: Cisco UCS C240 M7 24 SFF SAS/SATA/NVMe (UCSC-C240-M7SX)

Riser	Options
<p>Riser 1</p> <p>This riser is I/O-centric and controlled by CPU 1.</p>	<p>Riser 1A supports three PCIe slots numbered bottom to top:</p> <ul style="list-style-type: none"> • Slot 1 is full-height, $\frac{3}{4}$ length, x8, NCSI • Slot 2 is full-height, full-length, x16, NCSI • Slot 3 is full-height, full-length, x8, no NCSI <p>Riser 1C supports two PCIe slots numbered bottom to top:</p> <ul style="list-style-type: none"> • Slot 1 is full-height, $\frac{3}{4}$-length, x16 Gen5, NCSI • Slot 2 is full-height, full-length, x16 Gen5, no NCSI
<p>Riser 2</p> <p>This riser is I/O-centric and controlled by CPU 2.</p>	<p>Riser 2A supports three PCIe slots:</p> <ul style="list-style-type: none"> • Slot 4 is full-height, $\frac{3}{4}$ length, x8, NCSI • Slot 5 is full-height, full-length, x16, NCSI • Slot 6 is full-height, full-length, x8, no NCSI <p>Riser 2C supports two PCIe slots, numbered bottom to top:</p> <ul style="list-style-type: none"> • Slot 4 is full-height, $\frac{3}{4}$-length, x16 Gen5, NCSI • Slot 5 is full-height, full-length, x16 Gen5, no NCSI


Riser	Options
<p>Riser 3</p> <p>This riser is I/O-centric and controlled by CPU 2.</p>	<p>Riser 3A supports two PCIe slots:</p> <ul style="list-style-type: none"> • Slot 7 is full-height, full-length, x8 • Slot 8 is full-height, full-length, x8 <p>Riser 3C supports a GPU only.</p> <ul style="list-style-type: none"> • Supports one full-height, full-length, double-wide GPU (PCIe slot 7 only), x16 • Slot 8 is blocked by double-wide GPU

Cisco UCS C240 M7 Server 24 NVMe Drive Rear Panel, I/O Centric

The Cisco UCS C240 M7 24 NVMe version has a rear configuration option, for either I/O (I/O Centric) or Storage (Storage Centric) with the I/O Centric version of the server offering PCIe slots and the Storage Centric version of the server offering drive bays.

The following illustration shows the rear panel features for the I/O Centric version of the Cisco UCS C240 M7SN.

- For features common to all versions of the server, see Common Rear Panel Features.
- For definitions of LED states, see [Rear-Panel LEDs, on page 49](#).



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

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

Table 2: Cisco UCS C240 M7 24 SFF NVMe (UCSC-C240M7-SN)


Riser	Options
<p>Riser 1</p> <p>This riser is I/O-centric and controlled by CPU 1.</p>	<p>Riser 1A supports three PCIe slots:</p> <ul style="list-style-type: none"> • Slot 1 is full-height, $\frac{3}{4}$ length, x8, NCSI • Slot 2 is full-height, full-length, x16, NCSI • Slot 3 is full-height, full-length, x8, no NCSI <p>Riser 1C supports two PCIe slots numbered bottom to top:</p> <ul style="list-style-type: none"> • Slot 1 is full-height, $\frac{3}{4}$-length, x16 Gen5, NCSI • Slot 2 is full-height, full-length, x16 Gen5, no NCSI
<p>Riser 2</p> <p>This riser is I/O-centric and controlled by CPU 2.</p>	<p>Riser 2A supports three PCIe slots:</p> <ul style="list-style-type: none"> • Slot 4 is full-height, $\frac{3}{4}$ length, x8 • Slot 5 is full-height, full-length, x16 • Slot 6 is full-height, full-length, x8 <p>Riser 2C supports two PCIe slots numbered bottom to top:</p> <ul style="list-style-type: none"> • Slot 4 is full-height, $\frac{3}{4}$-length, x16 Gen5, NCSI • Slot 5 is full-height, full-length, x16 Gen5, no NCSI
<p>Riser 3</p>	<p>Riser 3A supports two PCIe slots numbered bottom to top:</p> <ul style="list-style-type: none"> • Slot 7 is full-height, full-length, x8 • Slot 8 is full-height, full-length, x8 <p>Riser 3C supports a GPU only:</p> <ul style="list-style-type: none"> • Supports one full-height, full-length, double-wide GPU (PCIe slot 7 only), x16 • Slot 8 is blocked by double-wide GPU

Cisco UCS C240 M7 Server 24 Drive Rear Panel, Storage Centric

The Cisco UCS C240 M7 24 SAS/SATA SFF version has a rear configuration option, for either I/O (I/O Centric) or Storage (Storage Centric) with the I/O Centric version of the server offering PCIe slots and the Storage Centric version of the server offering drive bays.

The following illustration shows the rear panel features for the Storage Centric version of the Cisco UCS C240 M7SX.

- For features common to all versions of the server, see Common Rear Panel Features.
- For definitions of LED states, see [Rear-Panel LEDs](#), on page 49.



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

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

Table 3: Cisco UCS C240 M7 24 SFF SAS/SATA/NVMe (UCSC-C240-M7SX)

Riser	Options
<p>Riser 1</p> <p>This riser is Storage-centric and controlled by CPU 1.</p>	<p>Riser 1B supports two SFF SAS/SATA/NVMe drives</p> <ul style="list-style-type: none"> • Slot 1 is reserved • Slot 2 (drive bay 102), x16 • Slot 3 (drive bay 101), x8 <p>When the server uses a hardware RAID controller card, SAS/SATA HDDs or SSDs, or NVMe PCIe SSDs are supported in the rear bays.</p>


Riser	Options
<p>Riser 2</p> <p>This riser is I/O-centric and controlled by CPU 2.</p>	<p>Riser 2A and 2C are supported for the Storage-centric version of the server.</p> <p>Riser 2A supports three slots:</p> <ul style="list-style-type: none"> • Slot 4 is full-height, $\frac{3}{4}$ length, x8, NCSI • Slot 5 is full-height, full-length, x16, NCSI • Slot 6 is full-height, full-length, x8 <p>Riser 2C supports two slots:</p> <ul style="list-style-type: none"> • Slot 4 is full-height, $\frac{3}{4}$-length, x16, Gen 5, NCSI • Slot 5 is full-height, full-length, x16, Gen 5 <p>NCSI support is limited to one slot at a time.</p>
<p>Riser 3</p> <p>This riser is controlled by CPU 2.</p>	<p>Riser 3B has two drive slots that can support two universal HDD or NVMe SFF drives.</p> <ul style="list-style-type: none"> • Slot 7 (drive bay 107), x4 • Slot 8 (drive bay 106), x4 <p>When the server uses a hardware RAID controller card, SAS/SATA HDDs or SSDs, or NVME PCIe SSDs, are supported in the rear bays.</p>

Cisco UCS C240 M7 Server 24 NVMe Drive Rear Panel, Storage Centric

The Cisco UCS C240 M7 24 NVMe SFF version has a rear configuration option, for either I/O (I/O Centric) or Storage (Storage Centric) with the I/O Centric version of the server using PCIe slots and the Storage Centric version of the server offering drive bays.

The following illustration shows the rear panel features for the Storage Centric version of the Cisco UCS C240 M7SN.

- For features common to all versions of the server, see Common Rear Panel Features.
- For definitions of LED states, see [Rear-Panel LEDs, on page 49](#).



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

Table 4: Cisco UCS C240 M7 24 SFF NVMe (UCSC-C240M7-SN)

Riser	Options
<p>Riser 1B</p> <p>This riser is Storage centric and controlled by CPU 1.</p>	<p>Riser 1B supports two universal HDD/NVMe SFF drive slots:</p> <ul style="list-style-type: none"> • Slot 1 is reserved, and does not support HDD or NVMe drives. This slot does support one M.2 NVMe RAID card. • Slot 2 (drive bay 102), x4 • Slot 3 (drive bay 101), x4 <p>When the server uses a hardware RAID controller card, NVMe PCIe SSDs are supported in the rear bays.</p>
<p>Riser 2</p>	<p>Riser 2A and 2C are supported for the Storage-centric version of the server.</p> <p>Riser 2A supports three slots:</p> <ul style="list-style-type: none"> • Slot 4 is full-height, $\frac{3}{4}$ length, x8, NCSI • Slot 5 is full-height, full-length, x16, NCSI • Slot 6 is full-height, full-length, x8 <p>Riser 2C supports two slots:</p> <ul style="list-style-type: none"> • Slot 4 is full-height, $\frac{3}{4}$-length, x16, Gen 5, NCSI • Slot 5 is full-height, full-length, x16, Gen 5 <p>NCSI support is limited to one slot at a time.</p>
<p>Riser 3</p>	<p>In the Storage-Centric configuration, Riser 3B has two slots that can support two universal HDD/NVMe SFF drives.</p> <ul style="list-style-type: none"> • Slot 7 (drive bay 107), x4 • Slot 8 (drive bay 106), x4


PCIe Risers

The following different PCIe riser options are available.


Riser 1 Options

This riser supports the following options, Riser 1A, 1B (two HDDs only), and 1C.


PCIe Riser 1A (outside)



PCIe Riser 1A (inside)



1	PCIe slot 1, supports full-height, ¾ length, x8, Gen 4, NCSI support for one slot at a time	2	PCIe slot 2, full height, full length, x16, Gen 4, GPU capable, NCSI support for one slot at a time
3	PCIe slot 3, full height, full length, x8, Gen 4, no NCSI	4	Edge Connectors




473202


1	PCIe slot 1, Reserved for drive controller (NVMe M.2 RAID controller)	2	Drive Bay 102, x4, Gen 4, 2.5-inch Universal HDD, SSD, or NVMe
3	Drive 101, x4, Gen 4, 2.5-inch Universal HDD, SSD, or NVMe	4	Edge Connectors

Riser 1C supports two PCIe Gen5 x16 Slots.

The following illustration shows Riser 1C (inside)



The following illustration shows Riser 1C (outside).




1	PCIe slot 1, supports full-height, $\frac{3}{4}$ length, x16, Gen 5, NCSI support on one slot at a time	2	PCIe slot 2, supports full-height, full length, x16, Gen 5, no NCSI
---	---	---	---


Riser 2

This riser supports options Riser 2A and 2C, which have the same electrical and mechanical properties as Riser 1A and Riser 1C, but Riser 2 has a different mechanical holder.

PCIe Riser 2A (outside)




PCIe Riser 2A (inside)



309595

1	PCIe slot 4, supports full-height, ¾ length, x8, Gen 4, NCSI support on one slot at a time	2	PCIe slot 5, full height, full length, x16, Gen 4, GPU capable, NCSI support on one slot at a time
3	PCIe slot 6, full height, full length, x8, Gen 4, no NCSI	4	Edge Connectors

Riser 2C supports two PCIE Gen5 x16 Slots.




1	PCIe slot 4, supports full-height, ³ / ₄ length, x16, Gen 5, NCSI support on one slot at a time	2	PCIe slot 5, supports full-height, full length, x16, Gen 5, no NCSI
---	---	---	---


Riser 3

This riser supports three options, 3A, 3B (supports HDD, SSD, and NVMe), and 3C.

PCIe Riser 3A (outside)




PCIe Riser 3A (inside)




309556

1	PCIe slot 7, full height, full length x8, Gen 4, no NCSI	2	PCIe slot 8, full height, full length, x8, Gen 4, no NCSI
3	Edge Connectors		

PCIe Riser 3B (outside)




PCIe Riser 3B (inside)




1	PCIe Slot 7, Drive Bay 104, x4, Gen 4, no NCSI	2	PCIe Slot 8, Drive Bay 103, x4, Gen 4, no NCSI
3	Edge Connectors		

PCIe Riser 3C (outside)



PCIe Riser 3C (inside)



309598

<p>1</p>	<p>PCIe Slot 7, supports one full height, full length, double-wide GPU (slot 7 only), x16, Gen 4, no NCSI</p> <p>Note The other slot is non configurable with hardware.</p>	<p>2</p>	<p>Edge Connectors</p>
-----------------	--	-----------------	------------------------

Summary of Server Features

The following tables list a summary of the server features.

Table 5: Server Features, SFF

Feature	Description
Chassis	Two rack-unit (2RU) chassis
Central Processor	One or two 4th Generation Intel Xeon processors.
Chipset	Intel® C741 chipset
Memory	32 slots for registered DIMMs (RDIMMs)
Multi-bit error protection	Multi-bit error protection is supported
Video	The Cisco Integrated Management Controller (CIMC) provides video using the Aspeed AST2600 video/graphics controller.
Network and management I/O	<p>Rear panel:</p> <ul style="list-style-type: none"> • One 10-Gb/100-G/1000-Gb Ethernet dedicated management port (RJ-45 connector) • One RS-232 serial port (RJ-45 connector) • One VGA video connector port (DB-15 connector) • Two USB 3.0 ports • Identification Button/Identification LED • (Optional) Four mLOM ports, 1-Gb/10-GB Ethernet <p>Front panel:</p> <ul style="list-style-type: none"> • One front-panel keyboard/video/mouse (KVM) connector that is used with the KVM breakout cable. The breakout cable provides two type A USB 2.0 connectors, one VGA (DB-15) connector, and one DB-9 serial connector.
Power	<p>Two of the following Platinum-efficiency hot-swappable power supplies:</p> <ul style="list-style-type: none"> • 1050 W (DC) • 1200 W (AC) • 1600 W (AC) • 2300 W (AC) <p>Two power supplies are mandatory, and power supplies must be the same. Cold redundancy and 1 + 1 redundancy are supported as long power supplies are the same.</p> <p>For additional information, see Supported Power Supplies, on page 102</p>

Feature	Description
ACPI	The advanced configuration and power interface (ACPI) 4.0 standard is supported.
Front Panel	The front panel controller provides status indications and control buttons KVM connector.
Cooling	Six hot-swappable fan modules for front-to-rear cooling.
InfiniBand	The PCIe bus slots in this server support the InfiniBand architecture.
Expansion Slots	<p>For the SFF versions of the server, the following expansion slots are supported:</p> <ul style="list-style-type: none"> • Riser 1A (Three PCIe slots) • Riser 1B (Two drive bays) • Riser 1C (Two PCIe Gen5 slots) • Riser 2A (Three PCIe slots) • Riser 2C (Two PCIe Gen5 slots) • Riser 3A (Two PCIe slots) • Riser 3B (Two drive bays) • Riser 3C (One PCIe slot) <p>Note Not all risers are available in every server configuration option.</p>
Interfaces	<p>Rear panel:</p> <ul style="list-style-type: none"> • One 10/100/1000 Base-T RJ-45 management port • One RS-232 serial COM port (RJ45 connector) • One DB15 VGA connector • Two Type A USB 3.0 port connectors • One flexible modular LAN on motherboard (mLOM) slot that can accommodate various interface cards • Identification Button/Identification LED <p>Front panel supports one KVM console connector that supplies:</p> <ul style="list-style-type: none"> • two Type A USB 2.0 connectors • one DB-15 VGA video connector • one DB-9 serial port connector


Feature	Description
Internal Storage Devices	<ul style="list-style-type: none"> • UCSC-C240-M7SX: <ul style="list-style-type: none"> • Up to 24 front SFF SAS/SATA hard drives (HDDs) or SAS/SATA/U.3 SFF solid state drives (SSDs). • Optionally, up to four front SFF NVMe PCIe SSDs. These drives must be placed in front drive bays 1, 2, 3, and 4 only. The rest of the bays (5 - 24) can be populated with SAS/SATA/U.3 SSDs or HDDs. Two CPUs are required in a server that has any number of NVMe drives. • Optionally, up to four 2.5-inch rear-facing universal drives (SAS/SATA HDDs or SSDs, or NVMe SSDs) • UCSC-C240-M7SN: <ul style="list-style-type: none"> • Up to 24 front NVMe drives (only). • Optionally, up to 4 rear NVMe drives (only) • Two CPUs are required when choosing NVMe SSDs • Other Storage: <ul style="list-style-type: none"> • An optional mini-storage module connector on the motherboard supports a boot-optimized RAID controller. The controller can support up to two Dual M.2 2280 SATA SSDs which can be used as boot volumes with corresponding interposer/controller cards. <p>Mixing different capacity SATA M.2 SSDs is not supported.</p>
Integrated Management Processor	<p>Baseboard Management Controller (BMC) running Cisco Integrated Management Controller (CIMC) firmware.</p> <p>Depending on your CIMC settings, the CIMC can be accessed through the 1GE dedicated management port, the 1GE/10GE LOM ports or a Cisco virtual interface card (VIC).</p> <p>CIMC manages certain components within the server, such as the Cisco 12G SAS HBA.</p>

Feature	Description
Storage Controllers	<ul style="list-style-type: none"> • One Cisco M7 12G SAS RAID controller with 4GB FBWC (for UCSC-240-M7SX server) <ul style="list-style-type: none"> • RAID support (RAID 0, 1, 5, 6, 10, 50, and 60) and SRAID0 • Supports up to 28 internal drives • Two Cisco M7 12G SAS HBA (for UCSC-240-M7SX servers) <ul style="list-style-type: none"> • JBOD/Pass-through Mode support • Each HBA supports up to 14 SAS/SATA internal drives • Two Cisco 24G Tri-Mode RAID Controller with 4GB cache (UCSC-RAID-HP): <ul style="list-style-type: none"> • Supports up to 24 front-loading SFF SAS/SATA or U.3 NVMe drives plus four rear-loading SFF SAS/SATA or U.3 NVMe drives. • Provides RAID 0/1/5/6/10/50/60 • Supports RAID for U.3 NVMe drives only • Drives behind this controller are hot swappable regardless of media type
Modular LAN over Motherboard (mLOM) slot	The dedicated mLOM slot on the motherboard can flexibly accommodate the Cisco Virtual Interface Cards (VICs), Series 15xxx
Server Management	Cisco Intersight provides server management.
CIMC	Cisco Integrated Management Controller (CIMC) 4.3(1) or later is required for the server.

Serviceable Component Locations

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

Figure 3: Cisco UCS C240 M7 Server, Serviceable Component Locations



1	Front-loading drive bays.	2	Cooling fan modules (six, hot-swappable)
3	DIMM sockets on motherboard (16 per CPU) See DIMM Population Rules and Memory Performance Guidelines , on page 96 for DIMM slot numbering. Note An air baffle rests on top of the DIMM and CPUs when the server is operating. The air baffle is not displayed in this illustration.	4	CPU socket 2
5	CPU socket 1	6	M.2 RAID Controller

7	<p>PCIe riser 3 (PCIe slots 7 and 8 numbered from bottom to top), with the following options:</p> <ul style="list-style-type: none"> • 3A (Default Option)—Slot 7 (x24 mechanical, x8 electrical, Gen 4) Slot 8 (x16 mechanical, x8 electrical, Gen 4) • 3B (Storage Option)—Slots 7 and 8, both support x4 electrical, Gen 4. Both slots can accept universal SFF HDDs or NVMe SSDs. • 3C (GPU Option)—Slot 7 (x24 mechanical, x16 electrical). Slot 7 can support a full height, full length GPU card. 	8	<p>PCIe riser 2 (PCIe slots 4, 5, 6 numbered from bottom to top), with the following options:</p> <ul style="list-style-type: none"> • 2A (Default Option)—Slot 4 (x24 mechanical, x8 electrical, Gen 4). NCSI is supported on one slot at a time. Supports a full height, ¾ length card. Slot 5 (x24 mechanical, x16 electrical, Gen 4). NCSI is supported on one slot at a time. Supports one full height, full length card. Slot 6 (x16 mechanical, x8 electrical, Gen 4). Supports a full height, full length card. • 2C— Slots 4 (x24 mechanical, x16 electrical, Gen 5) NCSI supported on one slot at a time. Supports a full-height, full-length card. Slot 5 (x24 mechanical, x16 electrical Gen 5) Supports full-height, full-length drive.
9	<p>PCIe riser 1 (PCIe slot 1, 2, 3 numbered bottom to top), with the following options:</p> <ul style="list-style-type: none"> • 1A (Default Option)—Slot 1 (x24 mechanical, x8 electrical, Gen 4) NCSI is supported on one slot at a time. Supports full height, ¾ length card. Slot 2 (x24 mechanical, x16 electrical, Gen 4). NCSI is supported on one slot at a time. Supports full height, full length GPU card. Slot 3 (x16 mechanical, x8 electrical, Gen 4) Supports full height, full length card. • 1B (Storage Option)—Slot 1 supports an M.2 NVMe RAID card Slot 2 (x4 electrical), supports universal 2.5-inch HDD NVMe drive Slot 3 (x4 electrical), supports universal 2.5-inch HDD NVMe drive • 1C — (x24 mechanical, x16 electrical, Gen 5) NCSI supported on one slot at a time. Supports a full-height, ¾-length card. Slot 2 (x24 mechanical, x16 electrical, Gen 5) Supports a full-height, full-length card. 	-	

The Technical Specifications Sheets for all versions of this server, which include supported component part numbers, are at [Cisco UCS Servers Technical Specifications Sheets](#) (scroll down to *Technical Specifications*).



CHAPTER 2

Installing the Server

This chapter contains the following topics:

- [Preparing for Installation, on page 27](#)
- [Installing the Server in a Rack, on page 30](#)
- [Initial Server Setup, on page 35](#)
- [NIC Mode and NIC Redundancy Settings, on page 40](#)
- [Updating the BIOS and Cisco IMC Firmware, on page 41](#)
- [Accessing the System BIOS, on page 42](#)
- [Smart Access Serial, on page 42](#)
- [Smart Access USB, on page 43](#)

Preparing for Installation

This section contains the following topics:

Installation Warnings and Guidelines



Note

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



Warning**IMPORTANT SAFETY INSTRUCTIONS**

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

Statement 1071



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

Statement 1047



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

Statement 1019



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

Statement 1005



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

Statement 1074



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

Statement 1017



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



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

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

- Plan your site configuration and prepare the site before installing the server. See the [Cisco UCS Site Preparation Guide](#) for the recommended site planning tasks.

- Ensure that there is adequate space around the server to allow for accessing the server and for adequate airflow. The airflow in this server is from front to back.
- Ensure that the air-conditioning meets the thermal requirements listed in the [Environmental Specifications, on page 166](#).
- Ensure that the cabinet or rack meets the requirements listed in the [Rack Requirements, on page 29](#).
- Ensure that the site power meets the power requirements listed in the [Power Specifications, on page 167](#). If available, you can use an uninterruptible power supply (UPS) to protect against power failures.

Rack Requirements

The rack must be of the following type:

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

Supported Cisco Slide Rail Kits

The server supports the following rail kit options:

- Cisco part UCSC-RAIL-D (Ball bearing rail kit for C220 & C240 M7 rack servers)
- Cisco part UCSC-CMA-C220-D (Reversible CMA for ball bearing rail kit)

Rack Installation Tools Required

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

Slide Rail and Cable Management Arm Dimensions

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

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

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

Installing the Server in a Rack



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

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


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

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

Statement 1006

Step 1 Attach the inner rails to the sides of the server:

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




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

- Install the second inner rail to the opposite side of the server.

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

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

Figure 4: Front Securing Mechanism, Inside of Front End



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

Step 3

Install the outer slide rails into the rack:

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

The slide rail front-end wraps around the outside of the rack post and the mounting pegs enter the rack-post holes from the outside-front. See [Figure 4: Front Securing Mechanism, Inside of Front End, on page 31](#).

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

- b) Push the mounting pegs into the rack-post holes from the outside-front.
 c) Press the securing plate release button, marked PUSH. The spring-loaded securing plate closes to lock the pegs in place.
 d) Adjust the slide-rail length, and then push the rear mounting pegs into the corresponding rear rack-post holes. The slide rail must be level front-to-rear.

The rear mounting pegs enter the rear rack-post holes from the *inside* of the rack post.

- e) Attach the second slide-rail assembly to the opposite side of the rack. Ensure that the two slide-rail assemblies are at the same height and are level front-to-back.
 f) Pull the inner slide rails on each assembly out toward the rack front until they hit the internal stops and lock in place.


Step 4

Insert the server into the slide rails:

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


- a) Align the rear ends of the inner rails that are attached to the server 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



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

Figure 5: Inner-Rail Release Clip




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

Step 5

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

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




Installing the Cable Management Arm (Optional)



Note The cable management arm (CMA) is reversible left-to-right. To reverse the CMA, see [Reversing the Cable Management Arm \(Optional\)](#), on page 34 before installation.

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

Figure 6: Attaching the CMA to the Rear Ends of the Slide Rails



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


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

Reversing the Cable Management Arm (Optional)

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

- 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 server.

Figure 7: Reversing the CMA



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

Initial Server Setup



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

Server Default Settings

The server is shipped with these default settings:

- The NIC mode is *Shared OCP Extended*.

In this mode, DHCP replies are returned to the OCP adapter card (Intel X710 OCP 3.0 card) and the Cisco virtual interface card (VIC) ports. If the system determines that the Cisco VIC connection is not getting its IP address from a Cisco UCS Manager system because the server is in standalone mode, further DHCP requests from the Cisco VIC are disabled. If you want to use the 10/100/1000 dedicated management ports to access Cisco IMC, you can connect to the server and change the NIC mode as described in [Setting Up the System With the Cisco IMC Configuration Utility, on page 38](#).

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

Connection Methods

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

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



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

This section contains the following topics:

Connecting to the Server Locally For Setup

This procedure requires the following equipment:

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

Step 1 Attach a power cord to each power supply in your server, and then attach each power cord to a grounded AC power outlet. If you are using DC power supplies, see [Installing DC Power Supplies \(First Time Installation\)](#), on page 104.

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

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

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

Step 3 Open the Cisco IMC Configuration Utility:

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

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

The following are the requirements for Strong Password:

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

Step 4 Continue with [Setting Up the System With the Cisco IMC Configuration Utility, on page 38](#).

Connecting to the Server Remotely For Setup

This procedure requires the following equipment:

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

Before you begin



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

Step 1 Attach a power cord to each power supply in your server, and then attach each power cord to a grounded AC power outlet. If you are using DC power supplies, see [Installing DC Power Supplies \(First Time Installation\), on page 104](#).

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

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

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

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

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

- Step 5** From the Cisco IMC Server Summary page, click **Launch KVM Console**. A separate KVM console window opens.
- Step 6** From the Cisco IMC Summary page, click **Power Cycle Server**. The system reboots.
- Step 7** Select the KVM console window.
- Note** The KVM console window must be the active window for the following keyboard actions to work.
- Step 8** When prompted, press **F8** to enter the Cisco IMC Configuration Utility. This utility opens in the KVM console window.
- Note** The first time that you enter the Cisco IMC Configuration Utility, you are prompted to change the default password. The default password is *password*. The Strong Password feature is enabled.
- The following are the requirements for Strong Password:
- The password can have a minimum of 8 characters and a maximum of 14 characters.
 - The password must not contain the user's name.
 - The password must contain characters from three of the following four categories:
 - English uppercase letters (A through Z)
 - English lowercase letters (a through z)
 - Base 10 digits (0 through 9)
 - Non-alphabetic characters !, @, #, \$, %, ^, &, *, -, _, =, “
- Step 9** Continue with [Setting Up the System With the Cisco IMC Configuration Utility, on page 38](#).

Setting Up the System With the Cisco IMC Configuration Utility

Before you begin

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

-
- Step 1** Set the NIC mode to choose which ports to use to access Cisco IMC for server management:
- *Shared OCP Extended*—In this NIC mode, DHCP replies are returned to the OCP adapter card (Intel X710 OCP 3.0 card) and the Cisco virtual interface card (VIC) ports. If the system determines that the Cisco VIC connection is not getting its IP address from a Cisco UCS Manager system because the server is in standalone mode, further DHCP requests from the Cisco VIC are disabled.
 - *Shared OCP*—The OCP adapter card is used to access Cisco IMC. You must select either the Active-Active or Active-standby NIC redundancy setting in the following step.
 - *Dedicated*—The dedicated management port is used to access Cisco IMC. You must select the *None* NIC redundancy setting in the following step.
 - *Cisco Card*—The ports on an installed Cisco UCS Virtual Interface Card (VIC) are used to access the Cisco IMC. You must select either the *Active-Active* or *Active-standby* NIC redundancy setting in the following step.
- See also the required VIC Slot setting below.

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

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

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

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

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

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

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

Step 4 (Optional) Make VLAN settings.

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

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

Step 6 (Optional) Set a hostname for the server.

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

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

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

Step 9 (Optional) Set a default user password.

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

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

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

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

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

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

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

What to do next

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



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

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

NIC Mode and NIC Redundancy Settings

Table 6: Valid NIC Redundancy Settings For Each NIC Mode

NIC Mode	Valid NIC Redundancy Settings
Shared OCP Extended	Active-active
Dedicated	None
Shared OCP	Active-active Active-standby
Cisco Card	Active-active Active-standby

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

- *Shared OCP Extended* (default)—In this mode, DHCP replies are returned to the OCP adapter card (Intel X710 OCP 3.0 card) and the Cisco virtual interface card (VIC) ports. If the system determines that the Cisco VIC connection is not getting its IP address from a Cisco UCS Manager system because the server is in standalone mode, further DHCP requests from the Cisco VIC are disabled. If you want to use the 10/100/1000 dedicated management ports to access Cisco IMC, you can connect to the server and change the NIC mode as described in [Setting Up the System With the Cisco IMC Configuration Utility, on page 38](#).
- *Shared OCP*—The OCP adapter card is used to access Cisco IMC. You must select either the Active-Active or Active-standby NIC redundancy setting in the following step.
- *Dedicated*—The dedicated management port is used to access Cisco IMC. You must select the *None* NIC redundancy setting in the following step.
- *Cisco Card*—The ports on an installed Cisco UCS Virtual Interface Card (VIC) are used to access the Cisco IMC. You must select either the *Active-Active* or *Active-standby* NIC redundancy setting in the following step.

See also the required VIC Slot setting below.

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

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

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

Updating the BIOS and Cisco IMC Firmware



Caution

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

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

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

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

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

Accessing the System BIOS

- Step 1** Enter the BIOS Setup Utility by pressing the **F2** key when prompted during bootup.
- Note** The version and build of the current BIOS are displayed on the Main page of the utility.
- Step 2** Use the arrow keys to select the BIOS menu page.
- Step 3** Highlight the field to be modified by using the arrow keys.
- Step 4** Press **Enter** to select the field that you want to change, and then modify the value in the field.
- Step 5** Press the right arrow key until the Exit menu screen is displayed.
- Step 6** Follow the instructions on the Exit menu screen to save your changes and exit the setup utility (or press **F10**). You can exit without saving changes by pressing **Esc**.
-

Smart Access Serial

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

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

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



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

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

Smart Access USB

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

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



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

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



CHAPTER 3

Maintaining the Server

This chapter contains the following topics:

- [Status LEDs and Buttons, on page 45](#)
- [Preparing For Component Installation, on page 51](#)
- [Removing the Server Top Cover, on page 53](#)
- [Serial Number Location, on page 55](#)
- [Hot Swap vs Hot Plug, on page 55](#)
- [Replacing the Air Duct, on page 55](#)
- [Removing and Replacing Components, on page 61](#)
- [Service Headers and Jumpers, on page 144](#)

Status LEDs and Buttons

This section contains information for interpreting LED states.

Front-Panel LEDs

Figure 8: Front Panel LEDs




Table 7: Front Panel LEDs, Definition of States

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

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

5	System health	<ul style="list-style-type: none"> • Green—The server is running in normal operating condition. • Green, blinking—The server is performing system initialization and memory check. • Amber, steady—The server is in a degraded operational state (minor fault). For example: <ul style="list-style-type: none"> • Power supply redundancy is lost. • CPUs are mismatched. • At least one CPU is faulty. • At least one DIMM is faulty. • At least one drive in a RAID configuration failed. • Amber, Blinking—The server is experiencing a critical fault. For example: <ul style="list-style-type: none"> • Boot Failure • Fatal Processor and/or bus error detected • Loss of I/O • Over Temperature Condition
6	Fan status	<ul style="list-style-type: none"> • Green—All fan modules are operating properly. • Amber—Fans are operating in a degraded state. For example, one of the fans has a fault. • Amber, blinking—Two or more fan modules have a fault.
7	Temperature status	<ul style="list-style-type: none"> • Green—The server is operating at normal temperature, or the temperature sensor detects no error conditions. • Amber, steady—One or more temperature sensors breached a warning threshold. • Amber, blinking—One or more temperature sensors breached a critical threshold.
8	Power supply status	<ul style="list-style-type: none"> • Green—All power supplies are operating normally, and no error condition is detected. • 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 LOM port link is idle. • Green—One or more Ethernet LOM ports are link-active, but there is no activity on any of the links. • Green, blinking—One or more Ethernet LOM ports are link-active, with activity.
---	-----------------------	--

Rear-Panel LEDs

Figure 9: Rear Panel LEDs




Table 8: Rear Panel LEDs, Definition of States

	LED Name	States
1	Rear unit identification	<ul style="list-style-type: none"> • Off—The unit identification function is not in use. • Blue, blinking—The unit identification function is activated.
2	USB 3.0	
3	USB 3.0	
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 up. • Green, blinking—Link is up, and traffic is present on the link.

6	Power supply status (one LED each power supply unit)	<p>AC power supplies:</p> <ul style="list-style-type: none"> • Off—No AC input (12 V main power off, 12 V standby power off). • Green, blinking—12 V main power off; 12 V standby power on. • Green, solid—12 V main power on; 12 V standby power on. • Amber, blinking—Warning threshold detected but 12 V main power on. • Amber, solid—Critical error detected; 12 V main power off (for example, over-current, over-voltage, or over-temperature failure). <p>DC power supply (UCSC-PSUV2-1050DC):</p> <ul style="list-style-type: none"> • Off—No DC input (12 V main power off, 12 V standby power off). • Green, blinking—12 V main power off; 12 V standby power on. • Green, solid—12 V main power on; 12 V standby power on. • Amber, blinking—Warning threshold detected but 12 V main power on. • Amber, solid—Critical error detected; 12 V main power off (for example, over-current, over-voltage, or over-temperature failure).
7 SAS	<p>SAS/SATA drive fault</p> <p>Note NVMe solid state drive (SSD) drive tray LEDs have different behavior than SAS/SATA drive trays.</p>	<ul style="list-style-type: none"> • Off—The hard drive is operating properly. • Amber—Drive fault detected. • Amber, blinking—The device is rebuilding. • Amber, blinking with one-second interval—Drive locate function activated in the software.
8 SAS	SAS/SATA drive activity LED	<ul style="list-style-type: none"> • Off—There is no hard drive in the hard drive tray (no access, no fault). • Green—The hard drive is ready. • Green, blinking—The hard drive is reading or writing data.

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

Internal Diagnostic LEDs

The server has internal fault LEDs for CPUs, DIMMs, and fan modules at the base of the CPUs, DIMMs, and fan modules.

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

Preparing For Component Installation

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

Required Equipment For Service Procedures

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

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

Shutting Down and Removing Power From the Server

The server can run in either of two power modes:

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



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

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

Shutting Down Using the Power Button

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

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

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

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

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

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

Shutting Down Using The Cisco IMC CLI

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

Step 1 At the server prompt, enter:

Example:

```
server# scope chassis
```

Step 2 At the chassis prompt, enter:

Example:

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

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

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

Shutting Down Using The Cisco IMC GUI

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

Step 1 In the Navigation pane, click the **Server** tab.

Step 2 On the Server tab, click **Summary**.

Step 3 In the Actions area, click **Power Off Server**.

Step 4 Click **OK**.

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

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

Removing the Server Top Cover

Step 1 Remove the top cover:

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

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

b) Lift on the end of the latch so that it pivots vertically to 90 degrees.

c) Simultaneously, slide the cover back and lift the top cover straight up from the server and set it aside.

Step 2 Replace the top cover:



a) With the latch in the fully open position, place the cover on top of the server about one-half inch (1.27 cm) behind the lip of the front cover panel.

b) Slide the cover forward until the latch makes contact.

- c) Press the latch down to the closed position. The cover is pushed forward to the closed position as you push down the latch.
- d) Lock the latch by sliding the lock button to sideways to the left.

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

Figure 10: Removing the Top Cover



309157

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

Serial Number Location

The serial number for the server is printed on a label on the top of the server, near the front. See [Removing the Server Top Cover, on page 53](#).

Hot Swap vs Hot Plug

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

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

Replacing the Air Duct

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

The server has two different revisions of airduct (A0 and B0) with subtle differences.



Note You might be able to find the revision on packaging, such as an ESD bag, or on a label on the air duct (if present).


To help you identify each revision of air duct, compare the illustrations that follow.

- One air duct (UCSC-GPUAD-C240M7= Rev A0) is used for servers populated with Intel Fourth Generation Xeon Scalable Processors. In the following illustrations, notice the lower mesh on the front of the air duct and the legs at the rear of the air duct, which are different from the Rev B0 air duct. Also, the front wall (not pictured) is different from the Rev B0 air duct.




Note The A0 air duct will be phased out so that only the B0 air duct will be available.

Figure 11: Air Duct for Intel Fourth Generation Xeon Scalable Processors (Rev A0), Front View



481367

Figure 12: Air Duct for Intel Fourth Generation Xeon Scalable Processors (Rev A0), Rear View




481366

- One air duct UCSC-GPUAD-C240M7= Rev B0 is required for servers that have Intel Fifth Generation Xeon Scalable Processors.
 - By default, this air duct is pre-installed at the factory for new servers with Intel Fifth Generation Xeon Scalable Processors.
 - This air duct is required if you will be upgrading your server from Intel Fourth Generation Xeon Scalable Processors to Intel Fifth Generation Xeon Scalable Processors.
 - When upgrading to Intel Fifth Generation Xeon Scalable Processors, you must order the Rev B0 air duct from Cisco.


In the following illustrations, notice the lower mesh on the front of the air duct and the legs at the rear of the air duct, which are different from the Rev A0 air duct. Also, the front wall (not pictured) is different from the Rev A0 air duct.

Figure 13: Air Duct for Intel Fifth Generation Xeon Scalable Processors, Front View



481369

Figure 14: Air Duct for Intel Fifth Generation Xeon Scalable Processors, Rear View



481368

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

- [Removing the Air Duct, on page 58](#)
- [Installing the Air Duct, on page 59](#)

Removing the Air Duct

Use this procedure to remove the air duct when needed.



Note Your air duct might be somewhat different than shown in this topic based on its revision level (A0 or B0), but the overall procedure is applicable.

Before you begin


The air duct has triangular alignment features that match with similar features on the sidewall of the server. Notice their locations. You will use them to aid reinstalling the air duct.

-
- Step 1** Remove the server top cover.
- Step 2** Locate the detents for the air duct.

The following illustration shows the locations of the detents. You can use them to grasp the air duct. When removing the air duct, always grasp the detents closest to the chassis sidewalls (left and right).

Step 3 Grasp the left and right detent then lift the air duct out of the chassis.

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



481376

What to do next

When you are done servicing the server, install the air duct. See [Installing the Air Duct, on page 59](#).

Installing the Air Duct

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


Step 1 Orient the air duct as shown.

Step 2 Match the alignment features on each side of the air duct with their corresponding feature on the chassis sidewall.

Note Also, notice that the sheetmetal of the chassis wall has notches that accept the tabs on the air duct.

Step 3 Lower the air duct into place and gently press down to ensure that all of its edges sit flush.

If the air duct is not seated correctly, it can obstruct installing the server's top cover.



481377

Step 4 When the air duct is correctly seated, attach the server's top cover.

The server top cover should sit flush so that the metal tabs on the top cover match the indents in the top edges of the air duct.

Removing and Replacing Components



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

Statement 1029



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




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

This section describes how to install and replace server components.

Serviceable Component Locations

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

Figure 15: Cisco UCS C240 M7 Server, Serviceable Component Locations



1	Front-loading drive bays.	2	Cooling fan modules (six, hot-swappable)
3	DIMM sockets on motherboard (16 per CPU) See DIMM Population Rules and Memory Performance Guidelines , on page 96 for DIMM slot numbering. Note An air baffle rests on top of the DIMM and CPUs when the server is operating. The air baffle is not displayed in this illustration.	4	CPU socket 2
5	CPU socket 1	6	M.2 RAID Controller

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

The Technical Specifications Sheets for all versions of this server, which include supported component part numbers, are at [Cisco UCS Servers Technical Specifications Sheets](#) (scroll down to *Technical Specifications*).

Replacing Front-Loading SAS/SATA Drives



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

To replace rear-loading SAS/SATA drives, see [Replacing Rear-Loading SAS/SATA Drives, on page 67](#).


Front-Loading SAS/SATA Drive Population Guidelines

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

- Cisco UCS C240 M7 24 SAS/SATA—Small form-factor (SFF) drives, with 24-drive backplane.
 - Front-loading drive bays 1—24 support 2.5-inch SAS/SATA/U.3 drives.
 - Optionally, front-loading drive bays 1 through 4 can support 2.5-inch NVMe SSDs.
- Cisco UCS C240 M7 24 NVMe—SFF drives, with 24-drive backplane.
 - Front-loading drive bays 1—24 support 2.5-inch NVMe PCIe SSDs *only*.

Drive bay numbering is shown in the following figures.

Figure 16: Small Form-Factor Drive (24-Drive) Versions, Drive Bay Numbering



Observe these drive population guidelines for optimum performance:

- When populating drives, add drives to the lowest-numbered bays first.



Note For diagrams of which drive bays are controlled by particular controller cables on the backplane, see [Storage Controller Cable Connectors and Backplanes, on page 175](#).

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

4K Sector Format SAS/SATA Drives Considerations

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

Setting Up UEFI Mode Booting in the BIOS Setup Utility

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

Setting Up UEFI Mode Booting in the Cisco IMC GUI

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

Step 12 Save changes and reboot the server. The changes you made will be visible after the system reboots.

Replacing a Front-Loading SAS/SATA Drive



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


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

- a) Press the release button on the face of the drive tray.
- b) Grasp and open the ejector lever and then pull the drive tray out of the slot.
- c) If you are replacing an existing drive, remove the four drive-tray screws that secure the drive to the tray and then lift the drive out of the tray.

Step 2 Install a new drive:

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

Figure 17: Replacing a Drive in a Drive Tray



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

Replacing Rear-Loading SAS/SATA Drives



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

Rear-Loading SAS/SATA Drive Population Guidelines

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

- UCS C240 M7 24 SAS/SATA—Small form-factor (SFF) drives, with 24-drive backplane.
 - Hardware RAID—Rear drive bays support SAS or NVMe drives
 - Intel® Virtual RAID on CPU—Rear drive bays support NVMe drives only.
- UCSC240 M7 24 NVMe—SFF drives, with 24-drive backplane.
 - Rear drive bays support only NVMe SSDs.
- Rear bays are numbered bays 25 through 28. Drives are numbered 101 to 104 with drive 101 at the left bottom, 102 at the left top, drive 103 at the right bottom, and 104 at the right top positions.
- When populating drives, add drives to the lowest-numbered bays first.
- Keep an empty drive blanking tray in any unused bays to ensure proper airflow.
- You can mix SAS/SATA hard drives and SAS/SATA SSDs in the same server. However, you cannot configure a logical volume (virtual drive) that contains a mix of hard drives and SSDs. That is, when you create a logical volume, it must contain all SAS/SATA hard drives or all SAS/SATA SSDs.


Replacing a Rear-Loading SAS/SATA Drive



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

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

Figure 18: Replacing a Drive in a Drive Tray



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

Basic Troubleshooting: Reseating a SAS/SATA Drive

Sometimes it is possible for a false positive UBAD error to occur on SAS/SATA HDDs installed in the server.

- Only drives that are managed by the UCS MegaRAID controller are affected.
- Drives can be affected regardless of where they are installed in the server (front-loaded, rear-loaded, and so on).
- Both SFF and LFF form factor drives can be affected.
- Drives installed in all Cisco UCS C-Series servers with M3 processors and later can be affected.

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

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



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

To reseat the drive, see [Reseating a SAS/SATA Drive, on page 70](#).

Reseating a SAS/SATA Drive

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

Use the following procedure to reseat the drive.



Caution This procedure might require powering down the server. Powering down the server will cause a service interruption.

Before you begin

Before attempting this procedure, be aware of the following:

- Before reseating the drive, it is a best practice to back up any data on it.
- When reseating the drive, make sure to reuse the same drive bay.
 - Do not move the drive to a different slot.
 - Do not move the drive to a different server.
 - If you do not reuse the same slot, the Cisco management software (for example, Cisco IMM) might require a rescan/rediscovery of the server.
- When reseating the drive, allow 20 seconds between removal and reinsertion.

-
- Step 1** Attempt a hot reseat of the affected drive(s). Choose the appropriate option:
- For a front-loading drive, see [Replacing a Front-Loading SAS/SATA Drive, on page 66](#)
 - For a rear-loading drive, see [Replacing a Rear-Loading SAS/SATA Drive, on page 68](#)
- Step 2** During boot up, watch the drive's LEDs to verify correct operation.
See [Status LEDs and Buttons, on page 45](#).
- Step 3** If the error persists, cold reseat the drive, which requires a server power down. Choose the appropriate option:

- a) Use your server management software to gracefully power down the server.
See the appropriate Cisco management software documentation.
- b) If server power down through software is not available, you can power down the server by pressing the power button.
See [Status LEDs and Buttons, on page 45](#).
- c) Reseat the drive as documented in Step 1.
- d) When the drive is correctly reseated, restart the server, and check the drive LEDs for correct operation as documented in Step 2.

Step 4

If hot and cold reseating the drive (if necessary) does not clear the UBAD error, choose the appropriate option:

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

Replacing Front-Loading NVMe SSDs

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

Front-Loading NVMe SSD Population Guidelines

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

- UCS C240 M7 SFF 24 SAS/SATA—Small form-factor (SFF) drives, with 24-drive backplane. Drive bays 1 and 4 support 2.5-inch U.3 NVMe SSDs.
- UCS C240 M7 24 NVMe—SFF drives, with 24-drive backplane. Drive bay 1 - 24 support *only* 2.5-inch U.3 NVMe SSDs.

Front-Loading NVMe SSD Requirements and Restrictions

Observe these requirements:

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

Observe these restrictions:

- NVMe 2.5 SSDs support booting only in UEFI mode. Legacy boot is not supported. For instructions on setting up UEFI boot, see [Setting Up UEFI Mode Booting in the BIOS Setup Utility, on page 65](#) or [Setting Up UEFI Mode Booting in the Cisco IMC GUI, on page 65](#).
- You cannot control NVMe PCIe SSDs with a SAS RAID controller because NVMe SSDs interface with the server via the PCIe bus.
- You can combine NVMe SSDs in the same system, but the same partner brand must be used. For example, two *Intel* NVMe SFF 2.5-inch SSDs and two *HGST* SSDs is an invalid configuration.

- UEFI boot is supported in all supported operating systems. Hot-insertion and hot-removal are supported in all supported operating systems except VMWare ESXi.

Enabling Hot-Plug Support in the System BIOS

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

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

Enabling Hot-Plug Support Using the BIOS Setup Utility

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

Enabling Hot-Plug Support Using the Cisco IMC GUI

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

Replacing a Front-Loading NVMe SSD

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



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



Note OS-informed hot-insertion and hot-removal must be enabled in the system BIOS. See [Enabling Hot-Plug Support in the System BIOS, on page 72](#).

-
- Step 1** Remove an existing front-loading NVMe SSD:
 - a) Shut down the NVMe SSD to initiate an OS-informed removal. Use your operating system interface to shut down the drive, and then observe the drive-tray LED:

- Green—The drive is in use and functioning properly. Do not remove.
- Green, blinking—the driver is unloading following a shutdown command. Do not remove.
- Off—The drive is not in use and can be safely removed.

- b) Press the release button on the face of the drive tray.
- c) Grasp and open the ejector lever and then pull the drive tray out of the slot.
- d) Remove the four drive tray screws that secure the SSD to the tray and then lift the SSD out of the tray.

Note If this is the first time that front-loading NVMe SSDs are being installed in the server, you must install a PCIe cable with PCIe riser 2C. See [Serviceable Component Locations, on page 24](#).


Step 2 Install a new front-loading NVMe SSD:

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

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

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

Figure 19: Replacing a Drive in a Drive Tray



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

Replacing Rear-Loading NVMe SSDs

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

Rear-Loading NVMe SSD Population Guidelines

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

- UCS C240 M7 24 SAS/SATA—Small form-factor (SFF) drives, with 24-drive backplane.
 - Hardware RAID—Rear drive bays support SAS or NVMe drives

- UCS C240 M7 24 NVMe—SFF drives, with 24-drive backplane.
 - Hardware RAID—Rear drive bays support \NVMe drives only
- For 24-drive servers, rear bays are numbered bays 101 through 104
- When populating drives, add drives to the lowest-numbered bays first.
- Drives are hot pluggable, but each drive requires a 10-second delay between hot removal and hot insertion.
- Keep an empty drive blanking tray in any unused bays to ensure proper airflow.

Rear-Loading NVMe SSD Requirements and Restrictions

Observe these requirements:

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

Observe these restrictions:

- NVMe SSDs support booting only in UEFI mode. Legacy boot is not supported. For instructions on setting up UEFI boot, see [Setting Up UEFI Mode Booting in the BIOS Setup Utility, on page 65](#) or [Setting Up UEFI Mode Booting in the Cisco IMC GUI, on page 65](#).
- You cannot control NVMe PCIe SSDs with a SAS RAID controller because NVMe SSDs interface with the server via the PCIe bus.
- You can combine NVMe 2.5-inch SSDs in the same system, but the same partner brand must be used. For example, two *Intel* NVMe SFF 2.5-inch SSDs and two *HGST* SSDs is an invalid configuration.
- UEFI boot is supported in all supported operating systems. Hot-insertion and hot-removal are supported in all supported operating systems except VMWare ESXi.

Replacing a Rear-Loading NVMe SSD

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



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



Note OS-informed hot-insertion and hot-removal must be enabled in the system BIOS. See [Enabling Hot-Plug Support in the System BIOS, on page 72](#).

Step 1 Remove an existing rear-loading NVMe SSD:

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

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


Step 2 Install a new front-loading NVMe SSD:

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

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

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

Figure 20: Replacing a Drive in a Drive Tray



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

Replacing Fan Modules

The six fan modules in the server are numbered as shown in [Serviceable Component Locations, on page 24](#).



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



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

Step 1 Remove an existing fan module:

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


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

- b) Remove the top cover from the server as described in [Removing the Server Top Cover, on page 53](#).
- c) Grasp and squeeze the fan module release latches on its top. Lift straight up to disengage its connector from the motherboard.

Step 2 Install a new fan module:

- a) Set the new fan module in place. The arrow printed on the top of the fan module should point toward the rear of the server.
- b) Press down gently on the fan module to fully engage it with the connector on the motherboard.
- c) Replace the top cover to the server.
- d) Replace the server in the rack, replace cables, and then fully power on the server by pressing the Power button.

Figure 21: Top View of Fan Module



1	Fan module release latches	2	Fan module fault LED <ul style="list-style-type: none"> • Green—All fan modules are operating properly. • Amber—Fans are operating in a degraded state. For example, one or more of the fans has a fault.
----------	----------------------------	----------	---

Replacing the Fan Tray

The server has a fan tray that contains the 5 individual fans modules. Individual fan modules can be replaced, and the fan tray can also be completely removed if needed.

To remove individual fans, see [Replacing Fan Modules, on page 77](#)

Use the following procedure to replace the fan tray.


- [Removing the Fan Tray, on page 80](#)

- [Installing the Fan Tray, on page 81](#)

Removing the Fan Tray

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

-
- Step 1** Rotate the tool-less middle lockdown screws that secure the fan tray to the chassis.
- Locate the lockdown screws that secure the fan tray to the server.
 - Grasp the screws and rotate them a quarter of a turn (90 degrees) to loosen the screws.
- Step 2** Unhinge the handle on both sides of the fan tray.
- Step 3** Remove the fan tray from the server.
- Grasp the handles at the top of the fan tray.
 - Holding the fan tray level, lift the fan tray up until it is removed from the chassis.



What to do next


Reinsert the fan tray into the chassis. See [Installing the Fan Tray, on page 81](#).

Installing the Fan Tray

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

Step 1 Install the fan tray.

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



- d) Push down and rotate the middle lockdown screw clockwise to lock the fan tray into the chassis receiving bracket.

Step 2 Close the top cover, or perform additional procedures, if needed.

Replacing CPUs and Heatsinks

This section contains the following topics:

- [CPU Configuration Rules, on page 82](#)
- [Tools Required For CPU Replacement, on page 83](#)
- [Removing CPUs and Heat Sinks, on page 84](#)
- [Installing the CPUs and Heatsinks, on page 89](#)
- [Additional CPU-Related Parts to Order with RMA Replacement CPUs, on page 94](#)

CPU Configuration Rules

This server has two CPU sockets on the motherboard. Each CPU supports 8 DIMM channels (16 DIMM slots). See [DIMM Population Rules and Memory Performance Guidelines](#), on page 96.

- Depending on the GPU installed in your server, a new GPU airduct (UCSC-GPUAD-C240M7=) might be required when you upgrade from Intel Xeon Fourth Generation CPUs (Sapphire Rapids) to Intel Xeon Fifth Generation CPUs (Emerald Rapids).
 - If the server has no GPU(s) that are >75W installed in the Sapphire Rapids system, you can upgrade to Emerald Rapids CPUs and still use the existing GPU airduct.


If the server has at least one GPU that is >75W installed in the Sapphire Rapids system, you must install the new GPU airduct (UCSC-GPUAD-C240M7=) to upgrade to Sapphire Rapids CPUs.

- Fourth and Fifth Generation Intel Xeon Scalable Processors (Sapphire Rapids and Emerald Rapids CPUs, respectively) have the same physical dimensions, CPU alignment features, and use the same heatsinks, so field-replacement procedures are the same regardless of which generation of CPU is installed.
- The server can operate with one CPU, or two identical CPUs installed.
- The minimum configuration is that the server must have at least CPU 1 installed. Install CPU 1 first, and then CPU 2.
- The following restrictions apply when using a single-CPU configuration:
 - Any unused CPU socket must have the socket dust cover from the factory in place.
 - The maximum number of DIMMs is 16 (only CPU 1 channels A through H).
- Two different form factors exist for heatsinks, a low profile and a high profile. The server can be ordered with either, but you cannot mix high- and low-profile CPUs and heatsinks in the same server. A single server must have all of one type.

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


- Low profile (UCSC-HSLP-C220M7), which has 4 T30 Torx screws on the main heatsink, and 2 Phillips-head screws on the extended heatsink.

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



309698

- High profile (UCSC-HSHP-240M7), which has 4 T30 Torx screws.



309697

Tools Required For CPU Replacement

You need the following tools and equipment for this procedure:

- T-30 Torx driver—Supplied with replacement CPU.
- #1 flat-head screwdriver—Supplied with replacement CPU.
- CPU assembly tool—Supplied with replacement CPU. Orderable separately as Cisco PID UCS-CPUAT=.
- Heatsink cleaning kit—Supplied with replacement CPU. Orderable separately as Cisco PID UCSX-HSCK=.

One cleaning kit can clean up to four CPUs.

- Thermal interface material (TIM)—Syringe supplied with replacement CPU. Use only if you are reusing your existing heatsink (new heatsinks have a pre-applied pad of TIM). Orderable separately as Cisco PID UCS-CPU-TIM=.

One TIM kit covers one CPU.

See also [Additional CPU-Related Parts to Order with RMA Replacement CPUs, on page 94](#).

Removing CPUs and Heat Sinks

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

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


- For a CPU with a high-profile heatsink, proceed to step a.
- For a CPU with a low-profile heatsink, skip to step 2.

- a) Using a T30 Torx driver, loosen all the securing nuts.
- b) Push the rotating wires towards each other to move them to the unlocked position.

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

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

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



d) Go to step 3.

Step 2

Remove the CPU.

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

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

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

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

e) Go to step 3.

Step 3


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

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

Step 4 Attach a CPU dust cover to the CPU socket.


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

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



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

- a) Turn the CPU assembly upside down, so that the heatsink is pointing down.
This step enables access to the CPU securing clips.
- b) Gently lift the TIM breaker in a 90-degree upward arc to partially disengage the CPU clips on this end of the CPU carrier.




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

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

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

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

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




Step 6 Transfer the CPU and carrier to the fixture.

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

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

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



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

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

What to do next

Choose the appropriate option:


- If you will be installing a CPU, go to [Installing the CPUs and Heatsinks, on page 89](#).
- If you will not be installing a CPU, verify that a CPU socket cover is installed. This option is valid only for CPU socket 2 because CPU socket 1 must always be populated in a runtime deployment.

Installing the CPUs and Heatsinks

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

Step 1 Remove the CPU socket dust cover on the server motherboard.

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



- c) Store the dust cover for future use.

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


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

Step 3 Apply new TIM.

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

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

Figure 22: Thermal Interface Material Application Pattern



Caution Use only the correct heatsink for your CPU. CPUs use the same heatsink based on your server configuration. For non-GPU servers, use UCSC-HSHP-C240M7. For a GPU, or GPU-ready configuration, use UCSC-HSLP-C220M7.

Step 4 Attach the heatsink to the CPU fixture.


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

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


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

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

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



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



Upgrading to Intel Xeon Fifth Generation CPUs

Use this task to upgrade to Emerald Rapids CPUs.

Before you begin

To upgrade CPUs, you will need to power off the server, so be aware that this task will cause a service interruption.

Also, you will need some tools to perform this task. See [Tools Required For CPU Replacement, on page 83](#).

Step 1 If you have not done so already, review the [CPU Configuration Rules, on page 82](#).

Step 2 Remove the server's top cover.

See [Removing the Server Top Cover, on page 53](#)

- Step 3** Remove the air duct.
See [Removing the Air Duct, on page 58](#).
- Step 4** Remove the existing Intel Xeon Fourth Generation Scalable Server Processor(s).
See [Removing CPUs and Heat Sinks, on page 84](#).
- Step 5** Store the removed CPUs in an ESD-safe bag.
- Step 6** Install the new Intel Xeon Fifth Generation Scalable Server Processor(s).
See [Installing the CPUs and Heatsinks, on page 89](#).
- Step 7** Replace the air duct.
See [Installing the Air Duct, on page 59](#).
- Step 8** Replace the server's top cover.
- Step 9** Reapply power and return the server to operation.

Additional CPU-Related Parts to Order with RMA Replacement CPUs

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



Note The following items apply to CPU *replacement* scenarios. If you are replacing a system chassis and *moving* existing CPUs to the new chassis, you do not have to separate the heatsink from the CPU. See [Additional CPU-Related Parts to Order with RMA Replacement System Chassis, on page 95](#).

- Scenario 1—You are reusing the existing heatsinks:
 - Heat sink cleaning kit (UCSX-HSCK=)
 - One cleaning kit can clean up to four CPUs.
 - Thermal interface material (TIM) kit for M7 servers (UCS-CPU-TIM=)
 - One TIM kit covers one CPU.
- Scenario 2—You are replacing the existing heatsinks:



Caution Use only the correct heatsink for your CPUs to ensure proper cooling. There are two different heatsinks, a low profile (UCSC-HSLP-C220M7) which is used with GPUs, or a GPU-ready server, and a high-profile (UCSC-HSHP-C240M7) which is used with a non-GPU configuration.

- New heatsinks have a pre-applied pad of TIM.

- Heat sink cleaning kit (UCSX-HSCK=)
 - One cleaning kit can clean up to four CPUs.
- Scenario 3—You have a damaged CPU carrier (the plastic frame around the CPU):
 - CPU Carrier
 - #1 flat-head screwdriver (for separating the CPU from the heatsink)
 - Heatsink cleaning kit (UCSX-HSCK=)
 - One cleaning kit can clean up to four CPUs.
 - Thermal interface material (TIM) kit for M7 servers (UCS-CPU-TIM=)
 - One TIM kit covers one CPU.

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

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

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

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



Note Unlike previous generation CPUs, the M7 server CPUs do not require you to separate the heatsink from the CPU when you *move* the CPU-heatsink assembly. Therefore, no additional heatsink cleaning kit or thermal-interface material items are required.

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

Replacing Memory DIMMs

The server supports R-DIMM, DDR5 288 pin DIMM modules. Eight memory channels per socket are supported with 2 DIMMs per channel, for a total of 32 memory slots on the motherboard.

When installing or replacing DIMMs, be aware of the following:



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



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



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

DIMM Population Rules and Memory Performance Guidelines

The following sections provide partial information for memory usage, mixing, and population guidelines. For detailed information about memory usage and population, download the PDF of the [Cisco UCS/UCSX M7 Memory Guide](#).

DIMMs Supported By Intel Fourth Generation Xeon Scalable Server Processors

The following table shows the DIMMs supported by Intel Fourth Generation Xeon Scalable Server Processors.

Table 9: Available DDR5 4800 MT/s DIMMs

PID (Product ID)	PID Description
UCS-MRX16G1RE1	16GB DDR5-4800 RDIMM 1Rx8 (16 Gb)
UCS-MRX32G1RE1	32GB DDR5-4800 RDIMM 1Rx4 (16 Gb)
UCS-MRX64G2RE1	64GB DDR5-4800 RDIMM 2Rx4 (16 Gb)
UCS-MR128G4RE1	128GB DDR5-4800 RDIMM 4Rx4 (16 Gb)
UCS-MR256G8RE1	256GB DDR5-4800 RDIMM 8Rx4 (16 Gb)

DIMMs Supported By Intel Fifth Generation Xeon Scalable Server Processors

The following table shows the DIMMs supported by Intel Fifth Generation Xeon Scalable Server Processors.


Table 10: Available DDR5 5600 MT/s DIMMs

PID (Product ID)	PID Description
UCS-MRX16G1RE3	16GB DDR5-5600 RDIMM 1Rx8 (16 Gb)
UCS-MRX32G1RE3	32GB DDR5-5600 RDIMM 1Rx4 (16 Gb)
UCS-MRX64G2RE3	64GB DDR5-5600 RDIMM 2Rx4 (16 Gb)
UCS-MRX96G2RF3	96GB DDR5-5600 RDIMM 2Rx4 (24 Gb)
UCS-MR128G4RE3	128GB DDR5-5600 RDIMM 4Rx4 (16 Gb)

DIMM Slot Numbering

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

Figure 23: DIMM Slot Numbering



DIMM Population Rules

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

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



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

Memory Population Order

The Cisco UCS C240 M7 server supports RDIMM memory modules.

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

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

Table 11: DIMMs Population Order

Number of DDR5 DIMMs per CPU (Recommended Configurations)	Populate CPU 1 Slot		Populate CPU2 Slots	
	P1 Blue #1 Slots	P1 Black #2 Slots	P2 Blue #1 Slots	P2 Black #2 Slots
1	A1	-	A1	
2	A1, G1	-	A1, G1	
4	A1, C1, E1, G1	-	A1, C1, E1, G1	
6	A1, C1, D1, E1, F1, G1	-	A1, C1, D1, E1, F1, G1	
8	A1, C1, D1, E1, G1, H1, B1, F1	-	A1, C1, D1, E1, G1, H1, B1, F1	
12	A1, B1, C1, D1, E1, F1, G1, H1	A2, C2, E2, G2	A1, B1, C1, D1, E1, F1, G1, H1	A2, C2, E2, G2
16	All populated (A1 through H1)	All populated (A2 through H2)	All populated (A1 through H1)	All populated (A2 through H2)

Memory Mirroring

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

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

Replacing DIMMs

Identifying a Faulty DIMM

Each DIMM socket has a corresponding DIMM fault LED, directly in front of the DIMM socket. See [Internal Diagnostic LEDs, on page 51](#) for the locations of these LEDs. When the server is in standby power mode, these LEDs light amber to indicate a faulty DIMM.

Step 1

Remove an existing DIMM:

- Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 52](#).
- Slide the server out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.

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

- Remove the top cover from the server as described in [Removing the Server Top Cover, on page 53](#).
- Remove the air baffle that covers the front ends of the DIMM slots to provide clearance.

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

Step 2 Install a new DIMM:

Note Before installing DIMMs, see the memory population rules for this server: [DIMM Population Rules and Memory Performance Guidelines](#), on page 96.

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

Replacing a Mini-Storage Module

The mini-storage module plugs into a vertical riser card that attaches to motherboard by two captive screws to provide additional internal storage. The module is an M.2 SSD Carrier which provides two M.2 form-factor SSD sockets. See also [Replacing a Boot-Optimized M.2 RAID Controller Module](#), on page 137.



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

Replacing a Mini-Storage Module Carrier

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

The carrier has one vertical media socket on its front side (the side that faces PCI riser 3) and one socket on its rear side (which faces PCI riser 2). Use the following procedure for an M.2 SSD mini-storage module carrier.

Step 1 Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server](#), on page 52.

Step 2 Slide the server out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.

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

Step 3 Remove the top cover from the server as described in [Removing the Server Top Cover](#), on page 53.

Step 4 Locate the mini-storage module carrier in its socket between PCIe riser 2 and 3.

Step 5 Using a Phillips screwdriver, loosen each of the captive screws and lift the M.2 riser out of the server.

Step 6 Remove a carrier from its socket:

- a) Using a Phillips screwdriver, loosen the screw that holds the module to the carrier.
- b) Push outward on the securing clips that holds each end of the carrier.

- c) Lift both ends of the carrier to disengage it from the socket on the motherboard.
- d) Set the carrier on an anti-static surface.

Step 7 Install a carrier to its socket:

- a) Position carrier over socket, with the carrier's connector facing down. Two alignment pegs must match with two holes on the carrier.
- b) Gently push down the socket end of the carrier so that the two pegs go through the two holes on the carrier.
- c) Push down on the carrier so that the securing clips click over it at both ends.

Step 8 Replace the top cover to the server.

Step 9 Replace the server in the rack, replace cables, and then fully power on the server by pressing the Power button.

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

This topic describes how to remove and replace an M.2 SATA or NVMe SSD in a mini-storage carrier for M.2 (PID UCS-HWRAID-M2-D). The carrier is mounted vertically on the riser, and has one M.2 SSD socket on either side (front or back) of the vertical riser.

Population Rules For Mini-Storage M.2 SSDs

- Both M.2 SSDs must be either SATA or NVMe; do not mix types in the carrier.
 - You can use one or two M.2 SSDs in the carrier. It is a best practice to use two SSDs.
 - M.2 socket 1 is on the front side of the carrier, which faces PCI riser 3; M.2 socket 2 is on the back of the carrier, which faces PCI riser 2.
-

Step 1 Power off the server and then remove the mini-storage module carrier from the server as described in [Replacing a Mini-Storage Module Carrier, on page 99](#).

Step 2 Remove an M.2 SSD:

- a) Use a #1 Phillips-head screwdriver to remove the single screw that secures the M.2 SSD to the carrier.
- b) Remove the M.2 SSD from its socket on the carrier.

Step 3 Install a new M.2 SSD:

- a) Insert the new M.2 SSD connector-end into the socket on the carrier with its label side facing up.
- b) Press the M.2 SSD flat against the carrier.
- c) Install the single screw that secures the end of the M.2 SSD to the carrier.

Step 4 Install the mini-storage module carrier back into the server and then power it on as described in [Replacing a Mini-Storage Module Carrier, on page 99](#).

Replacing the RTC Battery



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

[Statement 1015]



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

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

Step 1

Remove the RTC battery:

- a) Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 52](#).
- b) Slide the server out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.

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

- c) Remove the top cover from the server as described in [Removing the Server Top Cover, on page 53](#).
- d) Remove PCIe riser 1 from the server to provide clearance to the RTC battery socket that is on the motherboard. See [Replacing a PCIe Riser, on page 106](#).
- e) Locate the horizontal RTC battery socket.
- f) Remove the battery from the socket on the motherboard. Gently pry the securing clip to the side to provide clearance, then lift up on the battery.

Step 2

Install a new RTC battery:

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

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

- b) Replace PCIe riser 1 to the server. See [Replacing a PCIe Riser, on page 106](#).
- c) Replace the top cover to the server.
- d) Replace the server in the rack, replace cables, and then fully power on the server by pressing the Power button.

Replacing Power Supplies

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

- See also [Power Specifications, on page 167](#)
- See also [Rear-Panel LEDs, on page 49](#) for information about the power supply LEDs.

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

Supported Power Supplies

The Cisco UCS C240 M7 supports the following power supplies.



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

For detailed information, see [Power Specifications, on page 167](#).

PSU Type	Supported In	Notes
1050 W DC	All UCS C240 M7 models	One power supply is mandatory; one more can be added for 1 + 1 redundancy as long power supplies are the same.
1200 W AC	All UCS C240 M7 models	One power supply is mandatory; one more can be added for 1 + 1 redundancy as long power supplies are the same.
1600 W AC	All UCS C240 M7 models	One power supply is mandatory; one more can be added for 1 + 1 redundancy as long power supplies are the same.
2300 W AC	All UCS C240 M7 models	One power supply is mandatory; one more can be added for 1 + 1 redundancy as long power supplies are the same.

Replacing AC Power Supplies



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



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



Caution DO NOT interchange power supplies of any earlier Cisco UCS servers (for example, any Cisco UCS C240 M6 server power supplies) with the Cisco UCS C240 M7 server.

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

a) Perform one of the following actions:

- If your server has only one power supply, shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 52](#).

- If your server has two power supplies, you do not have to shut down the server.
- b) Remove the power cord from the power supply that you are replacing.
- c) Grasp the power supply handle while pinching the release lever toward the handle.
- d) Pull the power supply out of the bay.

Step 2

Install a new power supply:

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

Replacing DC Power Supplies



Note This procedure is for replacing DC power supplies in a server that already has DC power supplies installed. If you are installing DC power supplies to the server for the first time, see [Installing DC Power Supplies \(First Time Installation\)](#), on page 104.



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

Statement 1022



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

Statement 1045



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

Statement 1074




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



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

- Step 1** Remove the DC power supply that you are replacing or a blank panel from an empty bay:
- a) Perform one of the following actions:
 - If you are replacing a power supply in a server that has only one DC power supply, shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 52](#).
 - If you are replacing a power supply in a server that has two DC power supplies, you do not have to shut down the server.
 - b) Remove the power cord from the power supply that you are replacing. Lift the connector securing clip slightly and then pull the connector from the socket on the power supply.
 - c) Grasp the power supply handle while pinching the release lever toward the handle.
 - d) Pull the power supply out of the bay.
- Step 2** Install a new DC power supply:
- a) Grasp the power supply handle and insert the new power supply into the empty bay.
 - b) Push the power supply into the bay until the release lever locks.
 - c) Connect the power cord to the new power supply. Press the connector into the socket until the securing clip clicks into place.
 - d) Only if you shut down the server, press the Power button to boot the server to main power mode.

Figure 24: Replacing DC Power Supplies



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

Installing DC Power Supplies (First Time Installation)



Note This procedure is for installing DC power supplies to the server for the first time. If you are replacing DC power supplies in a server that already has DC power supplies installed, see [Replacing DC Power Supplies, on page 103](#).



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

Statement 1022



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

Statement 1045



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

Statement 1074



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



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

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

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


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

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

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

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

Figure 25: Replacing DC Power Supplies



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

Step 6 See [Grounding for DC Power Supplies, on page 106](#) for information about additional chassis grounding.

Grounding for DC Power Supplies

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

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



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

Replacing a PCIe Riser

This server has two toolless PCIe risers for horizontal installation of PCIe cards. Each riser is available in multiple versions. See [PCIe Slot Specifications, on page 107](#) for detailed descriptions of the slots and features in each riser version.

Step 1 Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 52](#).

Step 2 Slide the server out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.

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

Step 3 Remove the top cover from the server as described in [Removing the Server Top Cover, on page 53](#).

Step 4 Remove the PCIe riser that you are replacing:

- a) Grasp the flip-up handle on the riser and the blue forward edge, and then lift up evenly to disengage its circuit board from the socket on the motherboard. Set the riser on an antistatic surface.
- b) If the riser has a card installed, remove the card from the riser. See [Replacing a PCIe Card, on page 108](#).

Step 5 Install a new PCIe riser:

Note The PCIe risers are not interchangeable. If you plug a PCIe riser into the wrong socket, the server will not boot. Riser 1 must plug into the motherboard socket labeled “RISER1.” Riser 2 must plug into the motherboard socket labeled “RISER2.”

- a) If you removed a card from the old PCIe riser, install the card to the new riser. See [Replacing a PCIe Card, on page 108](#).
- b) Position the PCIe riser over its socket on the motherboard and over its alignment slots in the chassis.
- c) Carefully push down on both ends of the PCIe riser to fully engage its circuit board connector with the socket on the motherboard.

Step 6 Replace the top cover to the server.

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

Replacing a PCIe Card



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

PCIe Slot Specifications


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

- Riser 1 contains PCIe slots 1, 2, and 3. Riser 1 is available with the following different options:
 - SFF and NVMe server, IO-Centric, PCIe Gen4 Riser— Slot 1 (x8 Gen 4), 2 (x16 Gen 4), 3 (x8 Gen 4). All slots are controlled by CPU 1
 - SFF and NVMe server, IO-Centric, PCIe Gen5 Riser— Slot 1 (x16 Gen 5), 2 (x16 Gen 5). All slots are controlled by CPU 1
 - SFF server and NVMe server, Storage Centric—Slots 1 (Reserved), 2 (x4 Gen4), and 3 (x4 Gen 4) for drive bays of both SFF SAS/SATA models and NVMe models. All slots are controlled by CPU 1.
- Riser 2 contains PCIe slots 4, 5, and 6 and is available with the following different options:
 - SFF and NVMe server, I/O-Centric, PCIe Gen 4 Riser—Slots 4 (x8 Gen 4), 5 (x16 Gen 4), and 6 (x8 Gen 4). All slots are controlled by CPU 2.
 - SFF and NVMe server, I/O-Centric, PCIe Gen5 Riser— Slot 4 (x16 Gen5), 5 (x16 Gen 5). All slots are controlled by CPU 2

- Riser 3 contains PCIe slots 7 and 8 and is available in the following different options:
 - SFF and NVMe server, I/O-Centric, PCIe Gen 4 Riser—Slots 7 (x8 Gen 4) and 8 (x8 Gen 4). Slots 7 and 8 are controlled by CPU 2.
 - SFF and NVMe server, Storage Centric, PCIe Gen 4 Riser—Slots 7 (x4 Gen 4) and 8 (x4 Gen 4). All slots are controlled by CPU 2.
 - SFF and NVMe server, GPU-capable or GPU-ready PCIe Gen 4 Riser—Slot 7 (x16 Gen 4), full-height, full length supports double-wide GPU.

The following illustration shows the PCIe slot numbering.

Figure 26: Rear Panel, Showing PCIe Slot Numbering



473199

Replacing a PCIe Card



Note If you are installing a Cisco UCS Virtual Interface Card, there are prerequisite considerations. See [Cisco Virtual Interface Card \(VIC\) Considerations, on page 110](#).



Note RAID controller cards install into a dedicated motherboard socket. See [Replacing a SAS Storage Controller Card \(RAID or HBA\), on page 118](#).

- Step 1** Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 52](#).
- Step 2** Slide the server out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.
- Caution** If you cannot safely view and access the component, remove the server from the rack.
- Step 3** Remove the top cover from the server as described in [Removing the Server Top Cover, on page 53](#).
- Step 4** Remove the PCIe card that you are replacing:
- Remove any cables from the ports of the PCIe card that you are replacing.
 - Use two hands to flip up and grasp the blue riser handle and the blue finger grip area on the front edge of the riser, and then lift straight up.
 - On the bottom of the riser, push the release latch that holds the securing plate, and then swing the hinged securing plate open.

- d) Open the hinged card-tab retainer that secures the rear-panel tab of the card.
- e) Pull evenly on both ends of the PCIe card to remove it from the socket on the PCIe riser.

If the riser has no card, remove the blanking panel from the rear opening of the riser.

Step 5

Install a new PCIe card:

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


Step 6

Replace the top cover to the server.

Step 7

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

Figure 27: PCIe Riser Card Securing Mechanisms



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

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

Cisco Virtual Interface Card (VIC) Considerations

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



Note If you use the *Cisco Card* NIC mode, you must also make a *VIC Slot* setting that matches where your VIC is installed. The options are Riser1, Riser2, and Flex-LOM. See [NIC Mode and NIC Redundancy Settings, on page 40](#) for more information about NIC modes.

If you want to use the Cisco UCS VIC card for Cisco UCS Manager integration, see also the [Cisco UCS C-Series Server Integration with Cisco UCS Manager Guides](#) for details about supported configurations, cabling, and other requirements.

Table 12: VIC Support and Considerations in This Server

VIC	How Many Supported in Server	Slots That Support VICs	Primary Slot For Cisco UCS Manager Integration	Primary Slot For <i>Cisco Card</i> NIC Mode	Minimum Cisco IMC Firmware
Cisco UCS VIC 15425 UCSC-P-V5Q50G-D	2 PCIe	PCIe 2 PCIe 5	PCIe 2	PCIe 2	4.0(1)
Cisco UCS VIC 15235 UCSC-P-V5D200G-D	2 PCIe	PCIe 2 PCIe 5	PCIe 2	PCIe 2	4.0(2)
Cisco UCS VIC 15428 UCSC-M-V5Q50G	1 mLOM	mLOM	mLOM	mLOM	4.0(1)
Cisco UCS VIC 15238 UCSC-M-V5D200G	1 mLOM	mLOM	mLOM	mLOM	4.0(2)

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



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

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



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

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



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



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

Replacing an mLOM Card

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

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



Note If your mLOM card is a Cisco UCS Virtual Interface Card (VIC), see [Cisco Virtual Interface Card \(VIC\) Considerations, on page 110](#) for more information and support details.

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

- Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 52](#).
- Slide the server out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.

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

- Remove the top cover from the server as described in [Removing the Server Top Cover, on page 53](#).
- Remove any storage controller (RAID or HBA card) to provide clearance to the mLOM socket on the motherboard. See [Replacing a SAS Storage Controller Card \(RAID or HBA\), on page 118](#).
- Loosen the single captive thumbscrew that secures the mLOM card to the threaded standoff on the chassis floor.
- Slide the mLOM card horizontally to free it from the socket, then lift it out of the server.

Step 2 Install a new mLOM card:

- Set the mLOM card on the chassis floor so that its connector is aligned with the motherboard socket.
- Push the card horizontally to fully engage the card's edge connector with the socket.
- Tighten the captive thumbscrew to secure the card to the chassis floor.

- d) Return the storage controller card to the server. See [Replacing a SAS Storage Controller Card \(RAID or HBA\)](#), on page 118.
- e) Replace the top cover to the server.
- f) Replace the server in the rack, replace cables, and then fully power on the server by pressing the Power button.

Replacing an OCP Card

As a hardware option, the server can be configured with an Open Compute Project (OCP) 3.0 NIC in the rear mezzanine mLOM slot. To support this option, the server requires the Intel Ethernet Network Adapter X710 OCP 3.0 card (UCSC-O-ID10GC).



Note In addition to an OCP card, the server can support a Cisco mLOM in the rear mezzanine mLOM slot. The server can support either an OCP card or an mLOM, but not both. For information about replacing an mLOM, see [Replacing an mLOM Card](#), on page 111.

See the following topics.

- [Cisco VIC mLOM and OCP Card Replacement Considerations](#), on page 112
- [Removing an OCP Card](#), on page 113
- [Installing an OCP Card](#), on page 116

Cisco VIC mLOM and OCP Card Replacement Considerations

In Cisco UCS C240 M7 servers, Cisco IMC network connection may be lost in the following situations, while replacing Cisco VIC mLOM and OCP cards:

- If an OCP card is replaced by a Cisco VIC card in the mLOM Slot and the NIC mode is set to **Shared OCP** or **Shared OCP Extended**.
- If a Cisco VIC Card in the mLOM Slot is replaced by an OCP Card and the NIC mode is set to **Cisco-card MLOM**.

Follow these recommendations while replacing Cisco VIC mLOM or OCP cards in Cisco UCS C240 M7 servers to avoid loss of connectivity:

- Before replacing the card, configure any of the NIC modes that has network connected, other than **Cisco card MLOM**, **Shared OCP**, or **Shared OCP Extended**. After replacing the card, configure the appropriate NIC mode.

To set the NIC mode, refer *Server NIC Configuration* section in [Configuration Guides](#) for your Cisco IMC release.

- Or, after replacing the card, configure the appropriate NIC mode using Cisco IMC Configuration Utility/F8.

See [Connecting to the Server Remotely For Setup](#), on page 37.

- Or, after replacing the card, perform factory default settings using Cisco IMC Configuration Utility/F8 then perform the following steps:
 1. Once the server is rebooted, boot the system to Cisco IMC Configuration Utility/F8 then change the default password.
 2. Configure the appropriate NIC mode settings.

Table 13: Factory Default Settings

VIC in mLOM slot	Intel OCP 3.0 NIC (Intel X710) in mLOM Slot	VIC in Riser Slot	Dedicated Management Port	NIC Mode for CIMC Access
Yes	No	No	Yes	Cisco Card mode with the card in mLOM Slot
No	Yes	No	Yes	Shared OCP Extended
No	Yes	Yes	Yes	Shared OCP Extended
No	No	Yes	Yes	Cisco Card with VIC SLOT based on precedence: For C220 M7: 1. Riser 1 - Slot 1 2. Riser 3 - Slot 3 For C240 M7: 1. Riser 1 - Slot 2 2. Riser 2 - Slot 5 3. Riser 1 - Slot 1 4. Riser 2 - Slot 4
No	No	No	Yes	Dedicated

Removing an OCP Card

The OCP card (UCSC-O-ID10GC) mounts into the rear mezzanine mLOM slot. You will need to open the server top cover to remove or install the OCP card.

Use the following procedure to remove the OCP card from a server with full-height risers.

Before you begin


Gather a #2 Phillips screwdriver.

Step 1 If you have not removed the server's top cover, do so now.

See [Removing the Server Top Cover, on page 53](#).


Step 2 Remove the OCP bracket.

- a) Locate the two screws that secure the bracket to the server sheetmetal.
- b) Using a #2 Phillips screwdriver, loosen the screws.




473988

- c) Remove the screws and lift the bracket off of the server.



d) Holding the OCP card level, slide it out of the server.



Step 3 Choose the appropriate option:

- Reinstall an OCP Card. See [Installing an OCP Card, on page 116](#).
- If you want to install an mLOM instead, see [Replacing an mLOM Card, on page 111](#).
- Replace the top cover and return the server to operation.

Installing an OCP Card

The OCP 3.0 card (UCSC-O-ID10GC) installs into the rear mezzanine mLOM slot and connects to an adapter, not directly to the motherboard. To install the OCP card, the server's top cover must be opened to gain access to screws that secure the OCP card in place.

Use the following task to install an OCP 3.0 card.

Before you begin


Gather a #2 Phillips screwdriver.

Step 1 If you have not removed the server's top cover, do so now.


See [Removing the Server Top Cover, on page 53](#).

Step 2 Install the OCP card.


a) Holding the OCP card level, slide it into the slot on the rear of the server.



- b) Install the OCP bracket onto the server, making sure to align the screwholes at each end of the bracket with the screwholes on the OCP/mLOM slot.



Step 3 Using a #2 Phillips screwdriver, tighten the screws to secure the OCP bracket and OCP card to the server.



473987

What to do next

Replace the server top cover.

Replacing a SAS Storage Controller Card (RAID or HBA)

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

Storage Controller Card Firmware Compatibility

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




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

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


Removing the Dual Storage Controller Cards

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

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




- Step 2** Remove the fan tray.
For more information, see [Removing the Fan Tray, on page 80](#).

**Step 3**

Disconnect the cables.

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



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


Step 4 Remove the Storage controller cards.

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


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

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

Removing the Dual Storage Controller Cards



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



What to do next

Reinsert the dual Storage controller cards. Go to [Installing the Dual Storage Controller Cards](#), on page 123.


Installing the Dual Storage Controller Cards

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


Step 1 Grasp each card tray by the handle.

Step 2 Install the Storage controller cards.

- a) Make sure that the handle of the tray is in the open position.
- b) Make sure that the cables do not obstruct installing the Storage controller cards.
- c) Orient the Storage controller card so that the thumbscrews align with their threaded standoffs on the motherboard.
- d) Holding the card tray by the handle, keep the tray level and lower it into the server.



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




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

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

Step 3 Reconnect the cables.

Step 4 Reinsert the fan tray.

For more information, see [Installing the Fan Tray, on page 81](#).




What to do next

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


Removing the Storage Controller Card

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

Step 1 Locate the Storage controller card.




Step 2 Remove the fan tray.
 For more information, see [Removing the Fan Tray](#), on page 80.



Step 3 Disconnect the cables.

Removing the Storage Controller Card


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




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

Step 4 Remove the Storage controller card.

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



- c) Using both hands, grasp the tray's handle, and keeping the Storage controller card tray level, lift it out of the chassis.



What to do next

Reinsert the Storage controller card. Go to [Installing the Storage Controller Card, on page 130](#).


Installing the Storage Controller Card

Use this procedure to install the single Storage controller card into the server. The Storage controller card is contained in tray, which is replaceable.


Step 1 Grasp the card tray by the handle.

Step 2 Install the Storage controller card.

- a) Make sure that the handle of the tray is in the open position.
- b) Make sure that the cables do not obstruct installing the Storage controller card.
- c) Orient the Storage controller card so that the thumbscrews align with their threaded standoffs on the motherboard.
- d) Using both hands, hold the card tray by the handle, keep the tray level, and lower it into the server.



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




- f) Using both hands, make sure to apply equal pressure to both sides of the handle, and gently push the handle of the tray towards the front of the server.

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

Step 3 Reconnect the cables.

Step 4 Reinsert the fan tray.

For more information, see [Installing the Fan Tray, on page 81](#).




What to do next

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

Verify Cabling

After installing a Storage controller card, the cabling between the card(s) and rear drives should be as follows.

- For a 24-drive server, verify the following:
 - the SAS/SATA cable is connected to the controller card and Riser 3B
 - the SAS/SATA cable is connected to the controller card and the Riser 1B
 - both ribbon cables are connected to the controller card and the motherboard



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

Replacing the Supercap (RAID Backup)

This server supports installation of two Supercap units for 24 drive servers. The unit mounts to a bracket on the removable air baffle.


The Supercap provides approximately three years of backup for the disk write-back cache DRAM in the case of a sudden power loss by offloading the cache to the NAND flash.

Step 1 Prepare the server for component installation:

- a) Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 52](#).
- b) Slide the server out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.


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

- c) Remove the top cover from the server as described in [Removing the Server Top Cover, on page 53](#).
- d) Locate the SuperCap unit(s) as shown below.



Step 2 Remove an existing Supercap:

- a) Disconnect the Supercap cable from the RAID cable.
- b) Push aside the securing tab that holds the Supercap to its bracket.
- c) Lift the Supercap free of the bracket and set it aside.



309684


Step 3

Install a new Supercap:

- a) Orient the SuperCap unit so that the SuperCap cable and the RAID cable connectors meet.
- b) Make sure that the RAID cable does not obstruct the SuperCap when you install it, then insert the new Supercap into the mounting bracket.

Make that the SuperCap unit is securely inserted into it bracket.

- c) Connect the Supercap cable from the RAID controller card to the connector on the Supercap cable.



309665

Step 4 Replace the top cover to the server.

Step 5 Replace the server in the rack, replace cables, and then fully power on the server by pressing the Power button.

Replacing a Boot-Optimized M.2 RAID Controller Module

The Cisco Boot-Optimized M.2 RAID Controller module connects to the mini-storage module socket on the motherboard. It includes slots for two SATA M.2 drives, plus an integrated 6-Gbps SATA RAID controller that can control the SATA M.2 drives in a RAID 1 array.

Cisco Boot-Optimized M.2 RAID Controller Considerations

Review the following considerations:



Note The Cisco Boot-Optimized M.2 RAID Controller is not supported when the server is used as a compute-only node in Cisco HyperFlex configurations.

- The minimum version of Cisco IMC and Cisco UCS Manager that support this controller is 4.0(4) and later.
- This controller supports RAID 1 (single volume) and JBOD mode.



Note Do not use the server's embedded SW MegaRAID controller to configure RAID settings when using this controller module. Instead, you can use the following interfaces:

- Cisco IMC 4.2(1) and later
- BIOS HII utility, BIOS 4.2(1) and later
- Cisco UCS Manager 4.2(1) and later (UCS Manager-integrated servers)

-
- A SATA M.2 drive in slot 1 (the top) is the first SATA device; a SATA M.2 drive in slot 2 (the underside) is the second SATA device.
 - The name of the controller in the software is MSTOR-RAID.
 - A drive in Slot 1 is mapped as drive 253; a drive in slot 2 is mapped as drive 254.
 - When using RAID, we recommend that both SATA M.2 drives are the same capacity. If different capacities are used, the smaller capacity of the two drives is used to create a volume and the rest of the drive space is unusable.

JBOD mode supports mixed capacity SATA M.2 drives.
 - Hot-plug replacement is *not* supported. The server must be powered off.
 - Monitoring of the controller and installed SATA M.2 drives can be done using Cisco IMC and Cisco UCS Manager. They can also be monitored using other utilities such as UEFI HII, PMCLI, XMLAPI, and Redfish.
 - Updating firmware of the controller and the individual drives:
 - For standalone servers, use the Cisco Host Upgrade Utility (HUU). Refer to the [HUU Documentation](#).
 - For servers integrated with Cisco UCS Manager, refer to the [Cisco UCS Manager Firmware Management Guide](#).
 - The SATA M.2 drives can boot in UEFI mode only. Legacy boot mode is not supported.
 - If you replace a single SATA M.2 drive that was part of a RAID volume, rebuild of the volume is auto-initiated after the user accepts the prompt to import the configuration. If you replace both drives of a volume, you must create a RAID volume and manually reinstall any OS.
 - We recommend that you erase drive contents before creating volumes on used drives from another server. The configuration utility in the server BIOS includes a SATA secure-erase function.
 - The server BIOS includes a configuration utility specific to this controller that you can use to create and delete RAID volumes, view controller properties, and erase the physical drive contents. Access the utility by pressing **F2** when prompted during server boot. Then navigate to **Advanced > Cisco Boot Optimized M.2 RAID Controller**.

Replacing a Cisco Boot-Optimized M.2 RAID Controller

This topic describes how to remove and replace a Cisco Boot-Optimized M.2 RAID Controller. The controller board has one M.2 socket on its top (Slot 1) and one M.2 socket on its underside (Slot 2).

Step 1 Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server](#), on page 52.

Step 2 Slide the server out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.


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

Step 3 Remove the top cover from the server as described in [Removing the Server Top Cover](#), on page 53.


Step 4 Remove a controller from its motherboard socket:

- a) Locate the controller in its socket between PCIe Riser 2 and 3.

Figure 28: Cisco Boot-Optimized M.2 RAID Controller on Motherboard



- b) Using a #2 Phillips screwdriver, loosen the captive screws and remove the M.2 module.
- c) At each end of the controller board, push outward on the clip that secures the carrier.
- d) Lift both ends of the controller to disengage it from the carrier.



e) Set the carrier on an anti-static surface.


Step 5

If you are transferring SATA M.2 drives from the old controller to the replacement controller, do that before installing the replacement controller:

Note Any previously configured volume and data on the drives are preserved when the M.2 drives are transferred to the new controller. The system will boot the existing OS that is installed on the drives.

- a) Use a #1 Phillips-head screwdriver to remove the single screw that secures the M.2 drive to the carrier.
- b) Lift the M.2 drive from its socket on the carrier.
- c) Position the replacement M.2 drive over the socket on the controller board.
- d) Angle the M.2 drive downward and insert the connector-end into the socket on the carrier. The M.2 drive's label must face up.
- e) Press the M.2 drive flat against the carrier.
- f) Install the single screw that secures the end of the M.2 SSD to the carrier.
- g) Turn the controller over and install the second M.2 drive.

Figure 29: Cisco Boot-Optimized M.2 RAID Controller, Showing M.2 Drive Installation



806244

- Step 6** Install the controller to its socket on the motherboard:
- Position the controller over the socket, with the controller's connector facing down and at the same end as the motherboard socket. Two alignment pegs must match with two holes on the controller.
 - Gently push down the socket end of the controller so that the two pegs go through the two holes on the controller.
 - Push down on the controller so that the securing clips click over it at both ends.
- Step 7** Replace the top cover to the server.
- Step 8** Replace the server in the rack, replace cables, and then fully power on the server by pressing the Power button.

Replacing a Chassis Intrusion Switch

The chassis intrusion switch is an optional security feature that logs an event in the system event log (SEL) whenever the cover is removed from the chassis.

- Step 1** Prepare the server for component installation:
- Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 52](#).
 - Slide the server out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.

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

- c) Remove the top cover from the server as described in [Removing the Server Top Cover, on page 53](#).

Step 2 Remove an existing intrusion switch:

- a) Disconnect the intrusion switch cable from the socket on the motherboard.
- b) Use a #1 Phillips-head screwdriver to loosen and remove the single screw that holds the switch mechanism to the chassis wall.
- c) Slide the switch mechanism straight up to disengage it from the clips on the chassis.

Step 3 Install a new intrusion switch:

- a) Slide the switch mechanism down into the clips on the chassis wall so that the screwholes line up.
- b) Use a #1 Phillips-head screwdriver to install the single screw that secures the switch mechanism to the chassis wall.
- c) Connect the switch cable to the socket on the motherboard.

Step 4 Replace the cover to the server.

Step 5 Replace the server in the rack, replace cables, and then fully power on the server by pressing the Power button.

Installing a Trusted Platform Module (TPM)

The trusted platform module (TPM) is a small circuit board that plugs into a motherboard socket and is then permanently secured with a one-way screw. The socket location is on the motherboard below PCIe riser 2.

TPM Considerations

- This server supports TPM version 2.0 (UCSX-TPM-002C) as defined by the Trusted Computing Group (TCG). The TPM is also SPI-based.
- Field replacement of a TPM is not supported; you can install a TPM after-factory only if the server does not already have a TPM installed.
- If there is an existing TPM 1.2 installed in the server, you cannot upgrade to TPM 2.0. If there is no existing TPM in the server, you can install TPM 2.0.
- If the TPM 2.0 becomes unresponsive, reboot the server.

Installing TPM Hardware



Note For security purposes, the TPM is installed with a one-way screw. It cannot be removed with a standard screwdriver.

Step 1 Prepare the server for component installation:

- a) Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 52](#).
- b) Slide the server out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.

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

- c) Remove the top cover from the server as described in [Removing the Server Top Cover, on page 53](#).

Step 2 Remove PCIe riser 2 from the server to provide clearance to the TPM socket on the motherboard.

Step 3 Install a TPM:

- a) Locate the TPM socket on the motherboard.
- b) Align the connector that is on the bottom of the TPM circuit board with the motherboard TPM socket. Align the screw hole on the TPM board with the screw hole that is adjacent to the TPM socket.
- 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 4 Replace PCIe riser 2 to the server. See [Replacing a PCIe Riser, on page 106](#).

Step 5 Replace the cover to the server.

Step 6 Replace the server in the rack, replace cables, and then fully power on the server by pressing the Power button.

Step 7 Continue with [Enabling the TPM in the BIOS, on page 143](#).

Enabling the TPM in the BIOS

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



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

Step 1 Enable TPM Support:

- 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 server.

Step 2 Verify that TPM support is now enabled:

- a) Watch during bootup for the F2 prompt, and then press **F2** to enter BIOS setup.
- b) Log into the BIOS Setup utility with your BIOS Administrator password.
- c) Choose the **Advanced** tab.
- d) Choose **Trusted Computing** to open the TPM Security Device Configuration window.
- e) Verify that TPM SUPPORT and TPM State are Enabled.

Step 3 Continue with [Enabling the Intel TXT Feature in the BIOS, on page 143](#).

Enabling the Intel TXT Feature in the BIOS

Intel Trusted Execution Technology (TXT) provides greater protection for information that is used and stored on the business server. A key aspect of that protection is the provision of an isolated execution environment

and associated sections of memory where operations can be conducted on sensitive data, invisibly to the rest of the system. Intel TXT provides for a sealed portion of storage where sensitive data such as encryption keys can be kept, helping to shield them from being compromised during an attack by malicious code.


- Step 1** Reboot the server 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
 - Do one of the following:
 - If VT-d Support and VT Support are already enabled, skip to step 4.
 - If VT-d Support and VT Support are not enabled, continue with the next steps to enable them.
 - 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.
-

Service Headers and Jumpers

This server includes blocks of headers and switches (SW12, CN3) that you can use for certain service and debug functions.

This section contains the following topics:

- [Using the Clear CMOS Switch \(SW4, Switch 9\) , on page 145](#)
- [Using the Clear BIOS Password Switch \(SW4, Switch 6\), on page 146](#)
- [Using the Boot Alternate Cisco IMC Image Header \(CN5, Pins 1-2\), on page 147](#)



1	Location of header block CN5	4	Clear BIOS Password Switch (SW4 Switch 6) Clear CMOS Switch (SW4 Switch 9)
2	Boot Alternate Cisco IMC Header: CN5 pins 1 - 2	5	Clear CMOS Switch (SW4 Switch 9)
3	Location of SW4 DIP switches	-	

Using the Clear CMOS Switch (SW4, Switch 9)

You can use this switch to clear the server’s CMOS settings in the case of a system hang. For example, if the server hangs because of incorrect settings and does not boot, use this switch to invalidate the settings and reboot with defaults.

You will find it helpful to refer to the location of the SW4 switch block. See [Service Headers and Jumpers, on page 144](#).



Caution

Clearing the CMOS removes any customized settings and might result in data loss. Make a note of any necessary customized settings in the BIOS before you use this clear CMOS procedure.

-
- Step 1** Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 52](#).
- Step 2** Slide the server out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.
- Caution** If you cannot safely view and access the component, remove the server from the rack.
- Step 3** Remove the top cover from the server as described in [Removing the Server Top Cover, on page 53](#).
- Step 4** Using your finger, gently push the SW4 switch 9 to the side marked ON.
- Step 5** Reinstall the top cover and reconnect AC power cords to the server. The server powers up to standby power mode, indicated when the Power LED on the front panel is amber.
- Step 6** Return the server to main power mode by pressing the Power button on the front panel. The server is in main power mode when the Power LED is green.
- Note** You must allow the entire server to reboot to main power mode to complete the reset. The state of the switch cannot be determined without the host CPU running.
- Step 7** Press the Power button to shut down the server to standby power mode, and then remove AC power cords from the server to remove all power.
- Step 8** Remove the top cover from the server.
- Step 9** Using your finger, gently push switch 9 to its original position (OFF).
- Note** If you do not reset the switch to its original position (OFF), the CMOS settings are reset to the defaults every time you power-cycle the server.
- Step 10** Replace the top cover, replace the server in the rack, replace power cords and any other cables, and then power on the server by pressing the Power button.
-

Using the Clear BIOS Password Switch (SW4, Switch 6)

You can use this switch to clear the BIOS password.

You will find it helpful to refer to the location of the SW4 switch block. See [Service Headers and Jumpers, on page 144](#).

-
- Step 1** Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 52](#). Disconnect power cords from all power supplies.
- Step 2** Slide the server out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.
- Caution** If you cannot safely view and access the component, remove the server from the rack.
- Step 3** Remove the top cover from the server as described in [Removing the Server Top Cover, on page 53](#).
- Step 4** Using your finger, gently slide the SW4 switch 6 to the ON position.
- Step 5** Reinstall the top cover and reconnect AC power cords to the server. The server powers up to standby power mode, indicated when the Power LED on the front panel is amber.
- Step 6** Return the server to main power mode by pressing the Power button on the front panel. The server is in main power mode when the Power LED is green.

Note You must allow the entire server to reboot to main power mode to complete the reset. The state of the switch cannot be determined without the host CPU running.

Step 7 Press the Power button to shut down the server to standby power mode, and then remove AC power cords from the server to remove all power.

Step 8 Remove the top cover from the server.

Step 9 Reset the switch to its original position (OFF).

Note If you do not remove the switch to its original position (OFF), the BIOS password is cleared every time you power-cycle the server.

Step 10 Replace the top cover, replace the server in the rack, replace power cords and any other cables, and then power on the server by pressing the Power button.

Using the Boot Alternate Cisco IMC Image Header (CN5, Pins 1-2)

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

You will find it helpful to refer to the location of the CN5 header. See [Service Headers and Jumpers](#), on page 144.

Step 1 Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server](#), on page 52. Disconnect power cords from all power supplies.

Step 2 Slide the server out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.

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

Step 3 Remove the top cover from the server as described in [Removing the Server Top Cover](#), on page 53.

Step 4 Install a two-pin jumper across CN5 pins 1 and 2.

Step 5 Reinstall the top cover and reconnect AC power cords to the server. The server powers up to standby power mode, indicated when the Power LED on the front panel is amber.

Step 6 Return the server to main power mode by pressing the Power button on the front panel. The server is in main power mode when the Power LED is green.

Note When you next log in to Cisco IMC, you see a message similar to the following:

```
'Boot from alternate image' debug functionality is enabled.  
CIMC will boot from alternate image on next reboot or input power cycle.
```

Note If you do not remove the jumper, the server will boot from an alternate Cisco IMC image every time that you power cycle the server or reboot Cisco IMC.

Step 7 To remove the jumper, press the Power button to shut down the server to standby power mode, and then remove AC power cords from the server to remove all power.

Step 8 Remove the top cover from the server.

Step 9 Remove the jumper that you installed.

- Step 10** Replace the top cover, replace the server in the rack, replace power cords and any other cables, and then power on the server by pressing the Power button.
-



CHAPTER 4

Recycling Server Components

This chapter contains the following topics:

- [Server Recycling and E-Waste](#), on page 149
- [Battery Warning](#), on page 149
- [Recycling the Main Motherboard PCB Assembly \(PCBA\)](#), on page 150
- [Recycling the Front Mezzanine Module PCBA \(NVMe Backplane\)](#), on page 152
- [Recycling the Front Mezzanine Module PCBA \(SAS Backplane\)](#), on page 154
- [Recycling the Server KVM PCBA](#), on page 156
- [Recycling the Server Front Panel PCBA](#), on page 160

Server Recycling and E-Waste

The server has various components that can be recycled. When recycling the server and its components, always comply with your local laws governing recycling and e-waste.



Warning

The procedures in this chapter are destructive and can render the server unusable, so this content is not for standard use or FRU procedures! These procedures are for recyclers only.

Battery Warning

The server has a round button-style battery which is used for the real-time clock.



Warning

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

For information about removing the battery, see [Replacing the RTC Battery](#), on page 101.

Recycling the Main Motherboard PCB Assembly (PCBA)

The PCBA is secured to the server's sheet metal. You must disconnect the PCBA from the tray before recycling the PCBA. The PCBA is secured by different types of fasteners.

Before you begin



Note **For Recyclers Only!** This procedure is not a standard field-service option. This procedure is for recyclers who will be reclaiming the electronics for proper disposal to comply with local eco design and e-waste regulations.

To remove the printed circuit board assembly (PCBA), the following requirements must be met:

- The server must be disconnected from facility power.
- The server must be removed from the equipment rack.
- The server's top cover must be removed. See [Removing the Server Top Cover, on page 53](#).


Gather the following tools before beginning this procedure:

- Pliers
- T10 Torx screwdriver
- #2 Phillips screwdriver

Step 1 Locate the PCBA's mounting screws.

The following figure shows the location of the mounting screws.

Figure 30: Screw Locations for Removing the UCS C240 M7 PCBA



Indicator	Fastener	Required Tool
Red circles (○)	M3.5x0.6 mm screws (18)	Torx T10 screwdriver
Green circles (○)	H15 M4x0.7 mm lock screws (2)	Pliers
Blue circles (○)	M3.5x0.6 mm Thumb screws (2)	Torx T10 screwdriver
Yellow circles (○)	H12 M4x0.7 mm locking screws (2)	Pliers
Purple circle (○)	M3.5 Thumb screw (1), on the M.2 riser cage	#2 Phillips screwdriver
Lavender circle (○)	M3.5 Thumb screw (1), on the air duct	#2 Phillips screwdriver

Step 2 Using the appropriate tools, remove the screws.

Step 3 Remove the PCBA from the sheet metal and dispose of each in compliance with your local e-waste and recycling regulations.

Recycling the Front Mezzanine Module PCBA (NVMe Backplane)

The server's front mezzanine module contains one PCBA, which is the vertical PCB for the NVMe drive backplane. The PCBA is attached to the server's sheetmetal by 13 T10 screws.

You must disconnect the PCBA from the sheetmetal before recycling the PCBA.

Before you begin



Note **For Recyclers Only!** This procedure is not a standard field-service option. This procedure is for recyclers who will be reclaiming the electronics for proper disposal to comply with local eco design and e-waste regulations.


To remove the printed circuit board assembly (PCBA), the following requirements must be met:

- The server must be disconnected from facility power.
- The server must be removed from the equipment rack.
- The server's top cover must be removed. See [Removing the Server Top Cover, on page 53](#).

Gather a T10 Torx screwdriver.

Step 1 Remove the front mezzanine top cover.


- a) Using a T10 Torx driver, remove the two screws on the top of the front mezzanine module.
- b) Grasp the top of the front mezzanine module, and slide it off of the server chassis.



472747


Step 2 Disconnect the vertical drive backplane.

- a) Using a T10 Torx driver, remove the two screws on the top of the vertical drive backplane near the server chassis sidewalls.
- b) Grasp the drive backplane and disconnect it from the sheetmetal frame.




Step 3 Disconnect the PCBA from the sheetmetal frame.

- a) Using a T10 Torx screwdriver, remove the nine screws that secure the PCBA to the sheetmetal frame.



- b) Grasp the PCBA and detach it from the frame.



Step 4 Remove the PCBA and dispose of it properly in accordance with your local recycling and e-waste laws.

Recycling the Front Mezzanine Module PCBA (SAS Backplane)

The server's front mezzanine module contains one PCBA, which is the vertical PCB for the SAS drive backplane. The PCBA is attached to the server's sheetmetal by 13 T10 screws.

You must disconnect the PCBA from the sheetmetal before recycling the PCBA.

Before you begin



Note **For Recyclers Only!** This procedure is not a standard field-service option. This procedure is for recyclers who will be reclaiming the electronics for proper disposal to comply with local eco design and e-waste regulations.


To remove the printed circuit board assembly (PCBA), the following requirements must be met:

- The server must be disconnected from facility power.
- The server must be removed from the equipment rack.
- The server's top cover must be removed. See [Removing the Server Top Cover](#), on page 53.

Gather a T10 Torx screwdriver.

Step 1 Remove the front mezzanine top cover.


- a) Using a T10 Torx driver, remove the two screws on the top of the front mezzanine module.
- b) Grasp the top of the front mezzanine module, and slide it off of the server chassis.



472743

Step 2 Disconnect the vertical drive backplane.


- a) Using a T10 Torx driver, remove the two screws on the top of the vertical drive backplane near the server chassis sidewalls.
- b) Grasp the drive backplane and disconnect it from the sheetmetal frame.




472744

Step 3 Disconnect the PCBA from the sheetmetal frame.

- a) Using a T10 Torx screwdriver, remove the nine screws that secure the PCBA to the sheetmetal frame.



b) Grasp the PCBA and detach it from the frame.



Step 4 Remove the PCBA and dispose of it properly in accordance with your local recycling and e-waste laws.

Recycling the Server KVM PCBA

The server contains a KVM port, which is vertically mounted in the front right mounting ear. The KVM has a small PCB attached to a flat ribbon cable.

- Two size 2.5 slotted screws on the face of the tray.
- One #1 Phillips-head screw.
- One T10 Torx screws that secure the PCBA to the interior of the tray.

You must disconnect the PCBA from the server before recycling the PCBA.

Before you begin

Note For Recyclers Only! This procedure is not a standard field-service option. This procedure is for recyclers who will be reclaiming the electronics for proper disposal to comply with local eco design and e-waste regulations.

To remove the printed circuit board assembly (PCBA), the following requirements must be met:

- The server must be disconnected from facility power.
- The server must be removed from the equipment rack.
- The server's top cover must be removed. See [Removing the Server Top Cover, on page 53](#).


Gather the following tools:

- A size 2.5 slotted screwdriver.
- A #1 Phillips screwdriver
- A T10 Torx screwdriver.


Step 1

Disconnect the KVM port from the front panel.


- a) Using a 2.5 Phillips screwdriver, remove the two screws on the face of the front panel.




- b) On the back of the front panel, use a #1 Phillips head screwdriver to remove the 1 screw that secures the KVM port to the front panel.
- c) On the side of the chassis, use a T10 Torx driver to remove the one T10 screw.
- d) Grasp the side trim piece and remove it.



- Step 2** Disconnect the KVM port from the server.
- Grasp the cable and pull it backwards.
 - Grasp the rear of the KVM port and slide it backwards through the opening in the front panel.



- c) Continue to peel the cable off of the server to remove the PCBA.



Step 3 Remove the PCBA and dispose of it properly in accordance with your local recycling and e-waste laws.

Recycling the Server Front Panel PCBA

The server contains a front panel PCBA, which is vertically mounted in the front left mounting ear. The front panel has a small PCBA that is attached to the server by two flat ribbon cables. To recycle the PCBA, you will need to remove the following

- One #1 Phillips-head screw.
- One T10 Torx screws that secure the PCBA to the interior of the tray.

You must disconnect the PCBA from the server before recycling the PCBA.

Before you begin



Note **For Recyclers Only!** This procedure is not a standard field-service option. This procedure is for recyclers who will be reclaiming the electronics for proper disposal to comply with local eco design and e-waste regulations.

To remove the printed circuit board assembly (PCBA), the following requirements must be met:

- The server must be disconnected from facility power.
- The server must be removed from the equipment rack.
- The server's top cover must be removed. See [Removing the Server Top Cover, on page 53](#).


Gather the following tools:

- A #1 Phillips screwdriver
- A T10 Torx screwdriver.

Step 1

Disconnect the front panel from the server.

- a) On the back of the front panel, use a #1 Phillips head screwdriver to remove the 1 screw that secures the KVM port to the front panel.
- b) On the side of the chassis, use a T10 Torx driver to remove one T10 screw.




- c) Grasp the side trim piece and remove it.


Step 2

Remove the front panel PCBA.

- a) Grasp the PCBA and slide it backwards until it is halfway out of the front panel.



- b) Grasp the top cable and simultaneously lift it up and pull it toward the front panel to disconnect the cable.
- c) Grasp the bottom cable and simultaneously lift it up and pull it away from the front panel to disconnect the cable.
- d) When both cables are disconnected, completely slide the PCBA out of the front panel.



472757

Step 3 Dispose of the PCBA properly in accordance with your local recycling and e-waste laws.



APPENDIX **A**

Server Specifications

This chapter contains the following topics:

- [Server Specifications, on page 165](#)
- [Power Specifications, on page 167](#)
- [Power Cord Specifications, on page 170](#)

Server Specifications

This appendix lists the physical, environmental, and power specifications for the server.

- [Physical Specifications, on page 165](#)
- [Environmental Specifications, on page 166](#)
- [Power Specifications, on page 167](#)

Physical Specifications

The following table lists the physical specifications for the server versions.

Table 14: Physical Specifications

Description	Specification
Height	3.42 in. (86.9 mm)
Width	16.9 in. (429.3 mm)
Depth (length)	Server only: 30.5 in. (775 mm) Server with slide rail: 30.5 in. (775 mm) Server with bezel: 31.5 in. (800 mm)

Server weight	<ul style="list-style-type: none"> • SFF 24-drive server: <ul style="list-style-type: none"> • Maximum, fully configured with rail kit: 61.7 lb (26.67 kg) • Minimum, empty chassis, no rail kit: 33.14 lb (15.03 kg)
---------------	---

Environmental Specifications

As a Class A2 product, the server has the following environmental specifications.

Table 15: Environmental Specifications

Description	Specification
Temperature, Operating	Dry bulb temperature of 10°C to 35°C (50°F to 95°F) Maximum temperature change of 20°C (36°F) per hour (a temperature change within a specified period of time and not a rate of change) Humidity condition: Uncontrolled, not to exceed 50% RH starting condition Derate the maximum temperature by 1°C (33.8°F) per every 305 meters of altitude above 900m
Temperature, Extended Operating	5°C to 40°C (41°F to 104°F) with no direct sunlight Humidity condition: Uncontrolled, not to exceed 50% RH starting condition Derate the maximum temperature by 1°C (33.8°F) per every 305 meters of altitude above 900m
Temperature, non-operating (when the server is stored or transported)	Dry bulb temperature of 40 °C to 65 °C (-40°F to 149 °F)
Humidity (RH), operating	10% to 90% and 28°C (82.4°F) maximum dew-point temperature, non-condensing environment Minimum to be higher (more moisture) of -12 °C (10.4 °F) dew point or 8% relative humidity Maximum to be 24 °C (75.2 °F) dew point or 90% relative humidity
Humidity (RH), non-operating (when the server is stored or transported)	5% to 93% relative humidity, non-condensing, with a maximum wet bulb temperature of 28 °C across the 20 °C to 40 °C dry bulb range.
Altitude, operating	A maximum elevation of 3050 meters (10,006 feet)
Altitude, non-operating (when the server is stored or transported)	An elevation of 0 to 12,000 meters (39,370 feet)
Maximum Operating Duration	Unlimited

Sound power level Measure A-weighted per ISO7779 LwAd (Bels) Operation at 73°F (23°C)	5.5
Sound pressure level Measure A-weighted per ISO7779 LpAm (dBA) Operation at 73°F (23°C)	40

Power Specifications



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

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

<http://ucspowercalc.cisco.com>

The power specifications for the supported power supply options are listed in the following sections.

1050 W DC Power Supply

This section lists the specifications for each 1050 W DC power supply (Cisco part number UCSC-PSUV21050D-D).

Table 16: 1050 W DC Specifications

Description	Specification
DC Input Voltage	Nominal range: -48 to -60 VDC (Range: -40 to -72 VDC)
Maximum DC input current	N32 A at -40 VDC
Maximum input wattage	1234 W
Maximum inrush current	35 A (sub-cycle duration)
Maximum hold-up time	5 ms at 100% load (1050 W main and 36 W standby)
Maximum output power per PSU	1050 W on 12 VDC main power 36 W on 12 VDC standby power
Power supply output voltage	12 VDC
Power supply standby voltage	12 VDC

Efficiency rating	≥ 92% at 50% load
Form factor	RSP2
Input connector	Fixed 3-wire block

1200 W AC Power Supply

This section lists the specifications for each 1200 W AC power supply (Cisco part number UCSC-PSU1-1200W-D). This power supply is Titanium rated.

Parameter	Specification			
Input Connector	IEC320 C14			
Input Voltage Range (V rms)	100 to 240			
Maximum Allowable Input Voltage Range (V rms)	90 to 264			
Frequency Range (Hz)	50 to 60			
Maximum Allowable Frequency Range (Hz)	47 to 63			
Maximum Rated Output (W) Limited to 800W when operating at low-line input voltage, 100-127 V	1100	1200		
Maximum Rated Standby Output (W)	48			
Nominal Input Voltage (V rms)	100	120	208	230
Nominal Input Current (A rms)	12.97	10.62	6.47	5.84
Maximum Input at Nominal Input Voltage (W)	1300	1264	1343	1340
Maximum Input at Nominal Input Voltage (VA)	1300	1266	1345	1342
Maximum Rated Efficiency (%) Minimum rating required to achieve 80PLUS platinum certification.	90	90	91	91
Maximum Rated Power Factor Minimum rating required to achieve 80PLUS platinum certification.	0.97	0.97	0.97	0.97
Maximum Inrush Current (peak A)	20			
Maximum Inrush Current (ms)	0.2			

Parameter	Specification
Maximum Ride-Through Time The time that the output voltage remains within regulation limits at 100% load, during input voltage dropout	12

1600 W AC Power Supply

This section lists the specifications for each 1600 W AC power supply (Cisco part number UCSC-PSU1-1600W-D).

Table 17: 1600 W AC Specifications

Description	Specification
AC Input Voltage	Nominal range: 200–240 VAC (Range: 180–264 VAC)
AC Input Frequency	Nominal range: 50 to 60Hz (Range: 47–63 Hz)
Maximum AC Input current	9.5 A at 200 VAC
Maximum input volt-amperes	1250 VA at 200 VAC
Maximum inrush current	30 A at 35° C
Maximum hold-up time	80 ms at 1600 W
Maximum output power per PSU	1600 W at 200–240 VAC
Power supply output voltage	12 VDC
Power supply standby voltage	12 VDC
Efficiency rating	Climate Savers Platinum Efficiency (80Plus Platinum certified)
Form factor	RSP2
Input connector	IEC320 C14

2300 W AC Power Supply



Note For the 80PLUS platinum certification documented in the following table, you can find test results at <https://www.clearesult.com/80plus/>.

This section lists the specifications for each 2300 W AC power supply (Cisco part number UCSC-PSU1-2300W-D).

Parameter	Specification			
Input Connector	IEC320 C20			
Input Voltage Range (V rms)	100 to 240			
Maximum Allowable Input Voltage Range (V rms)	90 to 264			
Frequency Range (Hz)	50 to 60			
Maximum Allowable Frequency Range (Hz)	47 to 63			
Maximum Rated Output (W) Limited to 800W when operating at low-line input voltage, 100-127 V	2300			
Maximum Rated Standby Output (W)	36			
Nominal Input Voltage (V rms)	100	120	208	230
Nominal Input Current (A rms)	13	11	12	10.8
Maximum Input at Nominal Input Voltage (W)	1338	1330	2490	2480
Maximum Input at Nominal Input Voltage (VA)	1351	1343	2515	2505
Maximum Rated Efficiency (%) Minimum rating required to achieve 80PLUS platinum certification.	92	92	93	93
Maximum Rated Power Factor Minimum rating required to achieve 80PLUS platinum certification.	0.99	0.99	0.97	0.97
Maximum Inrush Current (peak A)	30			
Maximum Inrush Current (ms)	0.2			
Maximum Ride-Through Time Time output voltage remains within regulation limits at 100% load, during input voltage dropout	12			

Power Cord Specifications

Each power supply in the server has a power cord. Standard power cords or jumper power cords are available for connection to the server. The shorter jumper power cords, for use in racks, are available as an optional alternative to the standard power cords.



Note Only the approved power cords or jumper power cords listed below are supported.

Table 18: Supported Power Cords

Description	Length (Feet)	Length (Meters)
CAB-48DC-40A-8AWG DC power cord, -48 VDC, 40 A, 8 AWG Three-socket Mini-Fit connector to three-wire	11.7	3.5
CAB-C13-C14-AC AC power cord, 10 A; C13 to C14, recessed receptacle	9.8	3.0
CAB-250V-10A-AR AC power cord, 250 V, 10 A Argentina	8.2	2.5
CAB-C13-C14-2M-JP AC Power Cord, C13 to C14 Japan PSE Mark	6.6	2.0
CAB-9K10A-EU AC Power Cord, 250 V, 10 A; CEE 7/7 Plug Europe	8.2	2.5
CAB-250V-10A-IS AC Power Cord, 250 V, 10 A Israel	8.2	2.5
CAB-250V-10A-CN AC power cord, 250 V, 10 A PR China	8.2	2.5
CAB-ACTW AC power cord, 250 V, 10 A Taiwan	7.5	2.3
CAB-C13-CBN AC cabinet jumper power cord, 250, 10 A, C13 to C14	2.2	0.68

CAB-C13-C14-2M AC cabinet jumper power cord, 250 V, 10 A, C13 to C14	6.6	2.0
CAB-9K10A-AU AC power cord, 250 V, 10 A, 3112 plug, Australia	8.2	2.5
CAB-N5K6A-NA AC power cord, 200/240 V, 6 A, North America	8.2	2.5
CAB-250V-10A-ID AC power Cord, 250 V, 10 A, India	8.2	2.5
CAB-9K10A-SW AC power cord, 250 V, 10 A, MP232 plug Switzerland	8.2	2.5
CAB-250V-10A-BR AC power Cord, 250 V, 10 A Brazil	8.2	2.5
CAB-9K10A-UK AC power cord, 250 V, 10 A (13 A fuse), BS1363 plug United Kingdom	8.2	2.5
CAB-9K12A-NA AC power cord, 125 V, 13 A, NEMA 5-15 plug North America	8.2	2.5
CAB-AC-L620-C13 AC power cord, NEMA L6-20 to C13 connectors	6.6	2.0
CAB-9K10A-IT AC power cord, 250 V, 10 A, CEI 23-16/VII plug Italy	8.2	2.5
R2XX-DMYMPWRCORD No power cord; PID option for ordering server with no power cord	NA	NA



APPENDIX **B**

Storage Controller Considerations

This appendix contains the following topics:

- [Supported Storage Controllers and Cables, on page 173](#)
- [Storage Controller Card Firmware Compatibility, on page 174](#)
- [RAID Backup \(Supercap\), on page 175](#)
- [Mixing Drive Types in RAID Groups, on page 175](#)
- [Storage Controller Cable Connectors and Backplanes, on page 175](#)
- [For More RAID Utility Information, on page 176](#)

Supported Storage Controllers and Cables

This server supports a single, PCIe-style, SAS RAID or HBA controller that plugs into a dedicated internal socket.



Note NVMe PCIe SSDs cannot be controlled by a SAS/SATA RAID controller.

This server supports the RAID and HBA controller options and cable requirements shown in the following table.

Storage Adapter (PID)	Product Name	Supported Server	Maximum Number of Drives Supported	Supported RAID Type	Cache Size (GB)
UCSC-RAID-SD-D	Cisco 12G SAS RAID Controller with 4GB FBWC (28 Drives)	UCSC-C240-M7SX	28 24 2.5-inch SFF front-loading + 4 2.5-inch rear-loading drives Slots 1-24 plus 101 - 104	RAID	4

Storage Adapter (PID)	Product Name	Supported Server	Maximum Number of Drives Supported	Supported RAID Type	Cache Size (GB)
UCSC-SAS-T-D	Cisco 12G SAS HBA (16 Drives)	UCSC-C240-M7SX	28 2 controllers, each manages 14 drives. Controller 1: Slots 1-12 plus 103 and 104 Controller 2: Slots 13-24 plus 101 and 102	SAS HBA	NA
UCSC-9500-8E-D	Cisco 9500-8e 12G SAS HBA for external JBOD attach	All models of UCS C220 M7 and UCS C240 M7 server	NA	SAS HBA	NA
UCSC-RAID-HP-D	Cisco UCSC 24G Trim-Mode RAID Controller with 4GB FBWC (16 drives each)	Only the UCSC-C240-M7SX server	28 Two controllers, each manages 14 drives Controller 1: Slots 1-12 plus 103 and 104 Controller 2: Slots 13-24 plus 101 and 102	RAID	4GB FBWC

Storage Controller Card Firmware Compatibility

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



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

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

RAID Backup (Supercap)

This server supports installation of one Supercap unit. The unit mounts to a bracket in-line with the fan modules.

The optional SCPM provides approximately three years of backup for the disk write-back cache DRAM in the case of a sudden power loss by offloading the cache to the NAND flash.

For Supercap unit replacement instructions, see [Replacing the Supercap \(RAID Backup\)](#), on page 134.

Mixing Drive Types in RAID Groups

For the best performance follow these guidelines:

- Use either all SAS or all SATA drives in a RAID group.
- Use the same capacity for each drive in the RAID group.
- Never mix HDDs and SSDs in the same RAID group.

Storage Controller Cable Connectors and Backplanes

This section describes cabling for the storage controllers and backplanes. The SAS/SATA/U.3 cables are factory-installed and are used for all supported internal controllers.

This section also contains diagrams that show the cable-to-drive mapping.

- SFF 24-Drives SAS/SATA/U.3 NVMe
- SFF 24 Drives NVMe

Cisco 24G Tri-Mode RAID Controller or HBA for Up To 16 Drives

This controller is supported only in this server version:

- SFF 16-Drives SAS/SATA
- SFF 16 Drives U.3 NVMe

This HW RAID or HBA option can control up to 24 front-loading SAS/SATA/U.3 drives in this server version, plus 4 rear-loading SAS/SATA/U.3 drives.

1. Connect a SAS/SATA/U.3 cable from the small Slimline connector on the RAID card.
2. Connect a SAS/SATA/U.3 cable to the Riser 3B connector on the PCE Riser 3 cage.
3. Connect a SAS/SATA/U.3 cable from the second small Slimline connector on the RAID card.
4. Connect a SAS/SATA/U.3 cable to the Riser 1B connector on the PCE Riser 1 cage.


Cisco 12G Modular SAS RAID Controller or HBA For Up To 28 Drives (UCSC-RAID-SD-D or UCSC-SAS-T-D)

This controller is supported only in the SFF 24-drive SAS/SATA version of the server.

This HW RAID or HBA option can control up to 24 front-loading SAS/SATA drives in this server version, plus 4 rear-loading SAS/SATA drives.

1. Connect a SAS/SATA cable from the small Slimline connector on the RAID card.
2. Connect a SAS/SATA cable to the Riser 3B connector on the PCE Riser 3 cage.
3. Connect a SAS/SATA cable from the second small Slimline connector on the RAID card.
4. Connect a SAS/SATA cable to the Riser 1B connector on the PCE Riser 1 cage.

Figure 31: SFF, 24-Drive UCSC-C240-M7SX Cabling With UCSC-RAID-SD-D or UCSC-SAS-T-D



For More RAID Utility Information

Cisco RAID Information

- The Cisco UCS C Series M7 storage adapter supports SAS/SATA/U.3 drives. Tri-mode RAID support is supported.
- For basic information about RAID and for using the utilities for the RAID controller cards that are supported in Cisco servers, see the [Cisco UCS Servers RAID Guide](#).

Broadcom RAID Information

- For hardware SAS MegaRAID configuration—[Broadcom 12Gb/s MegaRAID® Tri-Mode Software User Guide](#).

- The Broadcom utilities have help documentation for more information about using the utilities. Go to [Broadcom StorCLI 12Gb/s MegaRAID® Tri-Mode User Guide](#).



APPENDIX **C**

GPU Card Installation

This chapter contains the following topics:

- [Server Firmware Requirements, on page 179](#)
- [GPU Card Configuration Rules, on page 180](#)
- [Requirement For All GPUs: Memory-Mapped I/O Greater Than 4 GB, on page 182](#)
- [Installing a Double-Wide GPU Card, on page 183](#)
- [Replacing a Heatsink, on page 186](#)
- [Replacing the GPU Air Blocker, on page 190](#)
- [Installing GRID Licenses From the NVIDIA Licensing Portal to the License Server, on page 196](#)

Server Firmware Requirements

The following table lists the minimum server firmware versions for the supported GPU cards.

GPU Card	Cisco IMC/BIOS Minimum Version Required
Intel Flex 140 PCIe, 75W, Gen4 x8	4.1(3)
Intel Flex 170 PCIe, 150W, Gen4 x16	4.1(3)
NVIDIA Tesla A100, 80GB, 300W, Passive (UCSC-GPUA100-80-D or HX-GPU-A100-80-D)	4.1(3)
NVIDIA Tesla A40 RTX, 48GB, 300 W, Passive (UCSC-GPU-A40-D)	4.1(3)
NVIDIA Tesla A30, 24GB, 180 W, Passive (UCSC-GPU-A30-D)	4.1(3)
NVIDIA A16 PCIe, 250W, 64GB (4x16GB), (UCSC-GPU-A16-D)	4.1(3)
NVIDIA H100 PCIe, 350W, Gen 5 x16, (UCSC-GPU-H100-D)	4.1(3)
NVIDIA L4 PCIe, 72W, Gen 4 x16, (UCSC-GPU-L4-D)	4.1(3)

GPU Card	Cisco IMC/BIOS Minimum Version Required
NVIDIA L40 PCIe, 300W, Gen 4 x16, (UCSC-GPU-L40-D)	4.1(3)

GPU Card Configuration Rules

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

- The UCSC-C240-M7SX and UCSC-C240-M7SN servers support a GPU Ready configuration which presets the server to accept a GPU at a later date.

This configuration sets the server with a with low-profile heatsink (UCSC-HSLP-C220M7) and GPU air blocker that installs in the middle slot of some individual risers (Riser 2 slot 4, and Riser 3 slot 8) (UCSC-RISAB-24XM7). The GPU air blocker is a double width part that blocks the slot in which it is installed, plus the slot above, to properly control airflow and ventilation.

The GPU air blocker is pre-installed in the GPU Ready configuration so that only the GPU is required when you choose to expand the server's compute power. The GPU Ready config has the following considerations:

- The server must follow the same temperature limits as a server with a GPU installed, even if the server does not currently have a GPU installed. By following the temperature limits even when the GPU is not present, you ensure correct operation when the GPU is installed later.
- This configuration requires a low-profile heatsink and a GPU air blocker. If you are ordering the GPU Ready configuration, you must select the GPU air blocker PID to enable GPU ready configuration when ordering the server through the Cisco online ordering and configuration tool. Follow the additional rules displayed in the tool.
- Two versions of air blocker exist. One is for systems with Sapphire Rapids CPUs and servers that have GPUs that are rated less than 75W of power consumption. One is for servers that have Emerald Rapids CPUs and one or more GPUs that are rated greater than 75W power consumption. For information about these GPU air blockers, see [Replacing the GPU Air Blocker, on page 190](#).
- The GPU air blocker is required in any empty GPU slots in a GPU-configured server or a GPU-ready server.
 - In these servers, the GPU air blocker is installed at the factory where needed.
 - However, if you remove a NIC or GPU from the GPU slot, the air blocker must be installed to ensure proper airflow.
- All GPU cards must be procured from Cisco because of a unique SBIOS ID that is required by CIMC and UCSM.
- Do not mix different brands or models of GPU cards in the server.
- GPUs are not supported in Riser 1B or Riser 3B. Riser 3B cannot mechanically accept a GPU.
- The UCSC-C240M7SX and UCSC-C240M7SN servers support one full-height, full-length, double-wide GPU (PCIe slot 7 only) in Riser 3C up to 300W & PCIe Gen4 speeds.

Both Riser 1A and 2A can support full-height, full-length, double-wide GPUs of up to 300W & PCIe Gen4 speeds.

Risers 1C and 2C can support full-height, full-length, double-wide GPU of up to 350W & PCIe Gen5 speeds.

- The following table shows additional details for the different supported GPUs.

GPU	GPU Information	Riser and Installation Notes
Intel Flex 140	HHHL, 75W, PCIe Gen 4 x8	Both Gen 4 or Gen 5 Risers, maximum of 5 GPUs supported <ul style="list-style-type: none"> • In Gen 4 risers, GPU is supported in slots 2, 3, 5,6, and 7 (Riser 3C) • In Gen 5 risers, GPU is supported in slots 1, 2, 4, and 5 (Risiers 1C and 2C). Also, slot 7 (Gen 4 Riser 3C)
Intel Flex 170	FHFL, single wide GPU, 150W, PCIe Gen 4 x16	Both Gen 4 risers (maximum of 5 GPUs supported) or Gen 5 risers (maximum of 3 GPUs supported). <ul style="list-style-type: none"> • In Gen 4 risers, GPU is supported in slots 2, 5, and 7 (Riser 3C) • In Gen 5 risers, GPU is supported in slots 2 and 5 (Risiers 1C and 2C). Also, slot 7 (Gen 4 Riser 3C) Requires power cable (UCS-M10CBL-C240M5)
Nvidia H100	FHFL, double wide GPU, 350W, PCIe Gen 5 x16	Gen 5 risers only, maximum of 2 GPUs supported in slots 2 and 5. Requires power cable (UCS-G5GPU-C240M7)
Nvidia L4	HHHL, 72W, PCIe Gen 4 x16	Both Gen 4 risers (maximum of 8 GPUs supported) or Gen 5 risers (maximum of 5 GPUs supported). <ul style="list-style-type: none"> • In Gen 4 risers, GPU is supported in all slots (Riser 1A, 2A, and 3A) • In Gen 5 risers, GPU is supported in slots 1, 2, 4, and 5 (Risiers 1C and 2C). Also, slot 7 (Gen 4 Riser 3C)

GPU	GPU Information	Riser and Installation Notes
Nvidia L40	FHFL, double wide GPU, 300W, PCIe Gen 4 x16	<p>Both Gen 4 risers (maximum of 3 GPUs supported) or Gen 5 risers (maximum of 3 GPUs supported as 2 GPUs in Gen 5 risers plus 1 in Gen 4 Riser 3).</p> <ul style="list-style-type: none"> In Gen 4 risers, GPU is supported in slots 2, 5, and 7. In Gen 5 risers, GPU is supported in slots 2, 5, and 7. <p>Requires power cable (CBL-L40GPU-C240M7)</p>

- Use the UCS power calculator at the following link to determine the power needed based on your server configuration: <http://ucspowercalc.cisco.com>

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

All supported GPU cards require enablement of the BIOS setting that allows greater than 4 GB of memory-mapped I/O (MMIO).

- Standalone Server: If the server is used in standalone mode, this BIOS setting is enabled by default:

Advanced > PCI Configuration > Memory Mapped I/O Above 4 GB [**Enabled**]

If you need to change this setting, enter the BIOS Setup Utility by pressing **F2** when prompted during bootup.

- If the server is integrated with Cisco UCS Manager and is controlled by a service profile, this setting is enabled by default in the service profile when a GPU is present.

To change this setting manually, use the following procedure.

Step 1 Refer to the Cisco UCS Manager configuration guide (GUI or CLI) for your release for instructions on configuring service profiles:

[Cisco UCS Manager Configuration Guides](#)

Step 2 Refer to the chapter on Configuring Server-Related Policies > Configuring BIOS Settings.

Step 3 In the section of your profile for PCI Configuration BIOS Settings, set `Memory Mapped IO Above 4GB Config` to one of the following:

- Disabled**—Does not map 64-bit PCI devices to 64 GB or greater address space.
- Enabled**—Maps I/O of 64-bit PCI devices to 64 GB or greater address space.
- Platform Default**—The policy uses the value for this attribute contained in the BIOS defaults for the server. Use this only if you know that the server BIOS is set to use the default enabled setting for this item.

Step 4 Reboot the server.

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 server is rebooted.

Installing a Double-Wide GPU Card

Use the following procedure to install or replace an NVIDIA Double-Wide GPU.

With Cisco IMC version 4.3(1) and later, the server can support up to three NVIDIA GPUs. For a list of supported GPUs, see [Server Firmware Requirements, on page 179](#).

The following table shows the ambient temperature thresholds for servers with PCIe HDDs and SSDs.

Table 19: PCIe Server Ambient Temperature

SKU Details	Storage Hardware Options	
	All NVMe or All SAS Storage	All NVMe or All SAS Plus 4 Rear HDDs
XCC, 350 W, Gen 4 or Gen 5	35 C normal ambient T	30 C normal ambient T
MCC, 300 W Gen 4 or Gen 5		

The following table shows the ambient temperature threshold for a server with PCIe SSDs and GPUs.

Table 20: GPU Server Ambient Temperature

SKU Details	Storage Hardware Options	
	All NVMe Storage plus GPUs	Four NVMe SSDs Plus 20 SAS HDDs Plus GPUs
XCC, 350 W, Gen 4 or Gen 5	30 C normal ambient T	30 C normal ambient T
MCC, 300 W Gen 4 or Gen 5		

The NVIDIA GPU card might be shipped with two power cables: a straight cable and a Y-cable. The straight cable is used for connecting power to the GPU card in this server; do not use the Y-cable, which is used for connecting the GPU card in external devices only.



Caution A GPU air blocker must be installed in any empty GPU slot in a GPU-configured or GPU-ready server! Do not operate the server with an empty GPU slot!

The supported NVIDIA GPU requires a C240 M5 NVIDIA Cable (UCS-P100CBL-240M5).

Step 1 Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 52](#).

Step 2 Slide the server out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.


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

Step 3 Remove the top cover from the server as described in [Removing the Server Top Cover, on page 53](#).

Step 4 Remove an existing GPU card:

- a) Disconnect any existing cable from the GPU card.
- b) Use two hands to grasp the metal bracket of the PCIe riser and lift straight up to disengage its connector from the socket on the motherboard. Set the riser on an antistatic surface.
- c) On the bottom of the riser, press down on the clip that holds the securing plate.
- d) Swing open the hinged securing plate to provide access.
- e) Open the hinged plastic retainer that secures the rear-panel tab of the card.
- f) Disconnect the GPU card's power cable from the power connector on the PCIe riser.
- g) Pull evenly on both ends of the GPU card to remove it from the socket on the PCIe riser.

Figure 32: PCIe Riser Card Securing Mechanisms



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

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

Step 5 Install a new GPU card:

Note Observe the configuration rules for this server, as described in [GPU Card Configuration Rules, on page 180](#).

- a) Align the GPU card with the socket on the riser, and then gently push the card's edge connector into the socket. Press evenly on both corners of the card to avoid damaging the connector.
- b) Connect the GPU power cable. The straight power cable connectors are color-coded. Connect the cable's black connector into the black connector on the GPU card and the cable's white connector into the white GPU POWER connector on the PCIe riser.


Caution Do not reverse the straight power cable. Connect the *black* connector on the cable to the *black* connector on the GPU card. Connect the *white* connector on the cable to the *white* connector on the PCIe riser.

- c) Close the card-tab retainer over the end of the card.
- d) Swing the hinged securing plate closed on the bottom of the riser. Ensure that the clip on the plate clicks into the locked position.
- e) Position the PCIe riser over its socket on the motherboard and over the chassis alignment channels.
- f) Carefully push down on both ends of the PCIe riser to fully engage its connector with the sockets on the motherboard.

At the same time, align the GPU front support bracket (on the front end of the GPU card) with the securing latch that is on the server's air baffle.

Step 6 Insert the GPU front support bracket into the latch that is on the air baffle:

- a) Pinch the latch release tab and hinge the latch toward the front of the server.
- b) Hinge the latch back down so that its lip closes over the edge of the GPU front support bracket.
- c) Ensure that the latch release tab clicks and locks the latch in place.



Step 7 Replace the top cover to the server.

Step 8 Replace the server in the rack, replace power and network cables, and then fully power on the server by pressing the Power button.

Step 9 Optional: Continue with [Installing Drivers to Support the GPU Cards, on page 200](#).

Replacing a Heatsink

For GPUs, the correct heatsink is the low-profile heatsink (UCSC-HSLP-C220M7), which has 4 T30 Torx screws on the main heatsink, and 2 Phillips-head screws on the extended heatsink. High profile heatsinks (UCSC-HSHP-C240M7) cannot be used on a GPU.

Use the following procedures to replace the heatsink on a GPU.

- [Removing a Heat Sink, on page 186](#)
- [Installing a Heatsink, on page 188](#)

Removing a Heat Sink

Use this procedure to remove a low-profile heatsink (UCSC-HSLP-C220M7) from a GPU.


Step 1 [Removing the Server Top Cover.](#)

Step 2 [Removing the Air Duct.](#)

Step 3 Remove the double-wide GPU.

- Use two hands to grasp the metal bracket of the PCIe riser and lift straight up to disengage its connector from the socket on the motherboard. Set the riser on an antistatic surface.
- On the bottom of the riser, press down on the clip that holds the securing plate.
- Swing open the hinged securing plate to provide access.
- Open the hinged plastic retainer that secures the rear-panel tab of the card.
- Disconnect the GPU card's power cable from the power connector on the PCIe riser.
- Pull evenly on both ends of the GPU card to remove it from the socket on the PCIe riser.


Figure 33: PCIe Riser Card Securing Mechanisms



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

Step 4 Remove the CPU.

- a) Using a #2 Phillips screwdriver, loosen the two Phillips head screws for the extended heatsink.
- b) Using a T30 Torx driver, loosen the four Torx securing nuts.



- c) Push the rotating wires towards each other to move them to the unlocked position.

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

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

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

Step 5 Remove the heatsink from the GPU.

What to do next

Install a low profile heatsink (UCSC-HSLP-C220M7) onto the GPU. See [Installing a Heatsink, on page 188](#).

Installing a Heatsink

Use this procedure to install a low-profile heatsink (UCSC-HSLP-C220M7) on a GPU.

Step 1 Apply new TIM, if needed.


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

- If you are installing a new heatsink, it is shipped with a pre-applied pad of TIM. Go to step 2 below.
- If you are reusing a heatsink, you must remove the old TIM from the heatsink and then apply new TIM to the CPU surface from the supplied syringe. Continue with step a below.

- a) Apply the Bottle #1 cleaning solution that is included with the heatsink cleaning kit (UCSX-HSCK=), as well as the spare CPU package, to the old TIM on the heatsink and let it soak for a least 15 seconds.

- b) Wipe all of the TIM off the heatsink using the soft cloth that is included with the heatsink cleaning kit. Be careful to avoid scratching the heatsink surface.
- c) Completely clean the bottom surface of the heatsink using Bottle #2 to prepare the heatsink for installation.
- d) Using the syringe of TIM provided with the new CPU (UCS-CPU-TIM=), apply 1.5 cubic centimeters (1.5 ml) of thermal interface material to the top of the CPU. Use the pattern shown in the following figure to ensure even coverage.

Figure 34: Thermal Interface Material Application Pattern




Caution Use only the correct heatsink. For GPU servers and GPU-ready servers, use UCSC-HSLP-C220M7.

Step 2 Install the heatsink.

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

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

- e) Set the T30 Torx driver to 12 in-lb of torque and tighten the 4 securing nuts to secure the CPU to the motherboard (3) first.
- f) Set the torque driver to 6 in-lb of torque and tighten the two Phillips head screws for the extended heatsink (4).



Replacing the GPU Air Blocker

The GPU air blocker is a molded part that installs into the PCIe risers in slots 2, 5, or 7 and extends upward to cover the slot above it. The GPU air blocker provides proper airflow and reduces the levels of dust or other potential particulate contaminants.

For GPU-configured or GPU-ready servers, a GPU air blocker is installed where needed as part of the riser (UCSC-RISAB-24MX7). The air blocker is also available as a separately orderable part (UCSC-RISAB-24MX7=).

To replace the GPU air blockers, use the appropriate task:



Note Riser 1 does not accept a GPU air blocker. If you are performing service tasks, do not attempt to install an air blocker for Riser 1.

- [Removing the Riser 2 GPU Air Blocker, on page 191](#)
- [Installing the Riser 2 GPU Air Blocker, on page 192](#)
- [Removing the Riser 3 GPU Air Blocker, on page 194](#)
- [Installing the Riser 3 GPU Air Blocker, on page 195](#)

Removing the Riser 2 GPU Air Blocker


The GPU air blocker is a molded part that installs into slot 5 in Riser 2 and extends upward to cover the slot above it (slot 6).



Caution In a GPU-configured or GPU-ready server, the GPU air blocker is required in any slot that does not contain a GPU! Do not operate the server with any empty GPU slot!

Use this procedure to remove the GPU air blocker.

-
- Step 1** Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 52](#).
- Step 2** Slide the server out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.
- Caution** If you cannot safely view and access the component, remove the server from the rack.
- Step 3** Remove the top cover from the server as described in [Removing the Server Top Cover, on page 53](#).
- Step 4** Remove the air blocker from the riser cage:
- Use two hands to grasp the metal bracket of the PCIe riser and lift straight up to disengage its connector from the socket on the motherboard.
 - Set the riser on an antistatic surface.
 - On the bottom of the riser, press down on the clip that holds the securing plate.
 - Swing open the hinged securing plate to provide access.
 - Open the hinged plastic retainer that secures the rear-panel tab of the card.
 - Grasp the air blocker, and holding it level, pull it horizontally out of the riser cage.



1	Release latch on hinged securing plate	2	GPU Air Blocker
----------	--	----------	-----------------

What to do next

Choose the appropriate option.


- Install a GPU. See [Installing a Double-Wide GPU Card](#), on page 183.
- Install a GPU air blocker. See [Installing the Riser 2 GPU Air Blocker](#), on page 192.

Installing the Riser 2 GPU Air Blocker


After removing a GPU or a GPU air blocker, you must reinstall a GPU or air blocker. You cannot operate the GPU configured or GPU-enabled server with an empty GPU slot.

Use the following procedure to install the GPU air blocker in slot 5 of Riser 2.

- Step 1** Orient the part so that it will insert into the connector of slot 5.
- Step 2** Holding the air blocker level, slide it into the riser cage and insert it into the riser cage connector.
- When the air blocker is installed, the air blocker should fit snugly into the riser and cover slot 5 and slot 6.
- Step 3** Close and latch the hinged door on the riser cage.



When the GPU air blocker is correctly installed, the hinged door easily closes, the air blocker is seated level, and it fits securely in the riser cage.



Step 4 Install the riser cage into the server.

What to do next

If no other maintenance work is required, replace the top cover and return the server to operation. Otherwise, continue with the additional maintenance tasks.

Removing the Riser 3 GPU Air Blocker


The GPU air blocker is a molded part that installs into slot 7 in Riser 3 and extends upward to cover the slot above it (slot 8).



Caution In a GPU-configured or GPU-ready server, the GPU air blocker is required in any slot that does not contain a GPU. Do not operate the server with any empty GPU slot!

Use this procedure to remove the GPU air blocker.

- Step 1** Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 52](#).
- Step 2** Slide the server out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.
- Caution** If you cannot safely view and access the component, remove the server from the rack.
- Step 3** Remove the top cover from the server as described in [Removing the Server Top Cover, on page 53](#).
- Step 4** Remove the air blocker from the riser cage:
- Use two hands to grasp the metal bracket of the PCIe riser and lift straight up to disengage its connector from the socket on the motherboard.
 - Set the riser on an antistatic surface.
 - On the bottom of the riser, press down on the clip that holds the securing plate.
 - Swing open the hinged securing plate to provide access.
 - Open the hinged plastic retainer that secures the rear-panel tab of the card.
 - Grasp the air blocker, and holding it level, pull it horizontally out of the riser cage.



1	Release latch on hinged securing plate	2	GPU Air blocker
---	--	---	-----------------

What to do next

Choose the appropriate option.


- Install a GPU. See [Installing a Double-Wide GPU Card, on page 183](#).
- Install a GPU air blocker. See [Installing the Riser 3 GPU Air Blocker, on page 195](#).

Installing the Riser 3 GPU Air Blocker


After removing a GPU or a GPU air blocker, you must reinstall a GPU or air blocker. You cannot operate the GPU configured or GPU-enabled server with an empty GPU slot.

Use the following procedure to install the GPU air blocker in slot 7 of Riser 3.

-
- Step 1** Orient the part so that it will insert into the riser cage connector of slot 7.
- Step 2** Holding the air blocker level, slide it into the riser cage.
- When the air blocker is installed, the air blocker should fit snugly into the riser and cover slot 7 and slot 8.
- Step 3** Close the hinged door on the riser cage.



When the GPU air blocker is correctly installed, the hinged door easily closes, the air blocker is seated level, and it fits securely in the riser cage.



Step 4 Install the riser cage into the server.

What to do next

If no other maintenance work is required, replace the top cover and return the server to operation. Otherwise, continue with the additional maintenance tasks.

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

Accessing the GRID License Server Management Interface

Open a web browser on the License Server host and access the URL <http://localhost:8080/licserver>.

If you configured the License Server host's firewall to permit remote access to the License Server, the management interface is accessible from remote machines at the URL <http://hostname:8080/licserver>

Reading Your License Server's MAC Address

Your License Server's Ethernet MAC address is used as an identifier when registering the License Server with NVIDIA's Licensing Portal.

Step 1 Access the GRID License Server Management Interface in a browser.

Step 2 In the left-side License Server panel, select **Configuration**.

The License Server Configuration panel opens. Next to **Server host ID**, a pull-down menu lists the possible Ethernet MAC addresses.

Step 3 Select your License Server's MAC address from the **Server host ID** pull-down.

Note It is important to use the same Ethernet ID consistently to identify the server when generating licenses on NVIDIA's Licensing Portal. NVIDIA recommends that you select one entry for a primary, non-removable Ethernet interface on the platform.

Installing Licenses From the Licensing Portal

Step 1 Access the GRID License Server Management Interface in a browser.

Step 2 In the left-side License Server panel, select **Configuration**.

The License Server Configuration panel opens.

Step 3 Use the License Server Configuration menu to install the .bin file that you generated earlier.

- a) Click **Choose File**.
- b) Browse to the license .bin file that you want to install and click **Open**.
- c) Click **Upload**.

The license file is installed on your License Server. When installation is complete, you see the confirmation message, "Successfully applied license file to license server."

Viewing Available GRID Licenses

Use the following procedure to view which licenses are installed and available, along with their properties.

Step 1 Access the GRID License Server Management Interface in a browser.

Step 2 In the left-side License Server panel, select **Licensed Feature Usage**.

Step 3 Click on a feature in the **Feature** column to see detailed information about the current usage of that feature.

Viewing Current License Usage

Use the following procedure to view information about which licenses are currently in-use and borrowed from the server.

Step 1 Access the GRID License Server Management Interface in a browser.

Step 2 In the left-side License Server panel, select **Licensed Clients**.

Step 3 To view detailed information about a single licensed client, click on its **Client ID** in the list.

Managing GRID Licenses

Features that require GRID licensing run at reduced capability until a GRID license is acquired.

Acquiring a GRID License on Windows

Step 1 Open the NVIDIA Control Panel using one of the following methods:

- Right-click on the Windows desktop and select **NVIDIA Control Panel** from the menu.
- Open Windows Control Panel and double-click the **NVIDIA Control Panel** icon.

Step 2 In the NVIDIA Control Panel left-pane under Licensing, select **Manage License**.

The Manage License task pane opens and shows the current license edition being used. The GRID software automatically selects the license edition based on the features that you are using. The default is Tesla (unlicensed).

Step 3 If you want to acquire a license for GRID Virtual Workstation, under License Edition, select **GRID Virtual Workstation**.

Step 4 In the **License Server** field, enter the address of your local GRID License Server. The address can be a domain name or an IP address.

Step 5 In the **Port Number** field, enter your port number or leave it set to the default used by the server, which is 7070.

Step 6 Select **Apply**.

The system requests the appropriate license edition from your configured License Server. After a license is successfully acquired, the features of that license edition are enabled.

Note After you configure licensing settings in the NVIDIA Control Panel, the settings persist across reboots.

Acquiring a GRID License on Linux

Step 1 Edit the configuration file `/etc/nvidia/gridd.conf`:

```
sudo vi /etc/nvidia/gridd.conf
```

Step 2 Edit the `ServerUrl` line with the address of your local GRID License Server.

The address can be a domain name or an IP address. See the example file below.

Step 3 Append the port number (default 7070) to the end of the address with a colon. See the example file below.

Step 4 Edit the `FeatureType` line with the integer for the license type. See the example file below.

- GRID vGPU = 1
- GRID Virtual Workstation = 2

Step 5 Restart the `nvidia-gridd` service.

```
sudo service nvidia-gridd restart
```

The service automatically acquires the license edition that you specified in the FeatureType line. You can confirm this in /var/log/messages.

Note After you configure licensing settings in the NVIDIA Control Panel, the settings persist across reboots.

Sample configuration file:

```
# /etc/nvidia/gridd.conf - Configuration file for NVIDIA Grid Daemon
# Description: Set License Server URL
# Data type: string
# Format: "<address>:<port>"
ServerUrl=10.31.20.45:7070

# Description: Set Feature to be enabled
# Data type: integer
# Possible values:
# 1 => for GRID vGPU
# 2 => for GRID Virtual Workstation
FeatureType=2
```

Using gpumodeswitch

The command line utility `gpumodeswitch` can be run in the following environments:

- Windows 64-bit command prompt (requires administrator permissions)
- Linux 32/64-bit shell (including Citrix XenServer dom0) (requires root permissions)



Note Consult NVIDIA product release notes for the latest information on compatibility with compute and graphic modes.

The `gpumodeswitch` utility supports the following commands:

- `--listgpumodes`

Writes information to a log file named `listgpumodes.txt` in the current working directory.

- `--gpumode graphics`

Switches to graphics mode. Switches mode of all supported GPUs in the server unless you specify otherwise when prompted.

- `--gpumode compute`

Switches to compute mode. Switches mode of all supported GPUs in the server unless you specify otherwise when prompted.



Note After you switch GPU mode, reboot the server to ensure that the modified resources of the GPU are correctly accounted for by any OS or hypervisor running on the server.

Installing Drivers to Support the GPU Cards

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

1. Update the server BIOS.
2. Update the GPU drivers.

1. Updating the Server BIOS

Install the latest Cisco UCS C240 M4 server BIOS by using the Host Upgrade Utility for the Cisco UCS C240 M4 server.



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

-
- Step 1** Navigate to the following URL: <http://www.cisco.com/cisco/software/navigator.html>.
 - Step 2** Click **Servers–Unified Computing** in the middle column.
 - Step 3** Click **Cisco UCS C-Series Rack-Mount Standalone Server Software** in the right-hand column.
 - Step 4** Click the name of your model of server in the right-hand column.
 - Step 5** Click **Unified Computing System (UCS) Server Firmware**.
 - Step 6** Click the release number.
 - Step 7** Click **Download Now** to download the `ucs-server_platform-huu-version_number.iso` file.
 - Step 8** Verify the information on the next page, and then click **Proceed With Download**.
 - Step 9** Continue through the subsequent screens to accept the license agreement and browse to a location where you want to save the file.
 - Step 10** Use the Host Upgrade Utility to update the server BIOS.
The user guides for the Host Upgrade Utility are at [Utility User Guides](#).
-

2. Updating the GPU Card Drivers

After you update the server BIOS, you can install GPU drivers to your hypervisor virtual machine.

-
- Step 1** Install your hypervisor software on a computer. Refer to your hypervisor documentation for the installation instructions.
 - Step 2** Create a virtual machine in your hypervisor. Refer to your hypervisor documentation for instructions.
 - Step 3** Install the GPU drivers to the virtual machine. Download the drivers from either:
 - NVIDIA Enterprise Portal for GRID hypervisor downloads (requires NVIDIA login): <https://nvidia.flexnetoperations.com/>
 - NVIDIA public driver area: <http://www.nvidia.com/Download/index.aspx>
 - 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.



APPENDIX **D**

Installation for Cisco Software Management

This chapter contains the following topics:

- [Installation For Cisco UCS Manager Integration, on page 203](#)

Installation For Cisco UCS Manager Integration

The Cisco UCS Manager integration instructions are in the integration guides found here:

[Cisco UCS C-Series Server Integration with UCS Manager Configuration Guides](#)

Refer to the guide that is for the version of Cisco UCS Manager that you are using.

Also refer to the release notes for Cisco UCS Manager software and C-Series Cisco IMC software for any special considerations regarding integration in your release.

- [Cisco UCS Manager Release Notes](#)
- [Cisco C-Series Software Release Notes](#)



INDEX

A

- air duct, installing [59](#)
- air duct, removing [58](#)
- alternate boot image, CIMC [147](#)

B

- BIOS password, clearing [146](#)
- boot image, alternate [147](#)

C

- CIMC, alternate boot image [147](#)
- clear BIOS password [146](#)
- clearing CMOS [145](#)
- CMOS clearing [145](#)
- CPU, upgrading [93](#)
- CPUs, installing [89](#)
- CPUs, removing [84](#)

D

- drive (SAS/SATA), reseating [70](#)
- dual storage controllers, installing [123](#)
- dual storage controllers, removing [119](#)

F

- fan tray, installing [81](#)
- fan tray, removing [80](#)
- front mezzanine PCB assembly, recycling [152, 154](#)
- front panel PCB assembly, recycling [156, 160](#)

G

- gpu air blocker, riser 1 [194](#)
- gpu air blocker, riser 2 [191–192](#)
- gpu air blocker, riser 3 [195](#)

H

- heatsink (gpu), removing [186, 188](#)
- heatsinks, installing [89](#)

- heatsinks, removing [84](#)

I

- installing CPUs [89](#)
- installing heatsinks [89](#)
- installing OCP card [116](#)
- installing, air duct [59](#)
- installing, fan tray [81](#)
- installing, GPU air blocker [192, 195](#)
- installing, storage controller [130](#)
- installing, storage controllers [123](#)

M

- main motherboard PCBA, recycling [150](#)

O

- OCP card [116](#)
- OCP card, installing [116](#)
- OCP card, removing [113](#)

P

- password (BIOS), clearing [146](#)

R

- recycling, front mezzanine PCBA [152, 154](#)
- recycling, front panel PCBA [160](#)
- recycling, front panel PCBAs [156](#)
- recycling, main motherboard PCBA [150](#)
- removing CPUs [84](#)
- removing heatsinks [84](#)
- removing OCP card [113](#)
- removing, air duct [58](#)
- removing, dual storage controllers [119](#)
- removing, fan tray [80](#)
- removing, GPU air blocker [191, 194](#)
- removing, heatsink (gpu) [186, 188](#)
- removing, storage controller [126](#)
- reseating drives, SAS/SATA [70](#)
- riser 1, GPU air blocker [194](#)

riser 2, GPU air blocker [191–192](#)
riser 3, GPU air blocker [195](#)

S

SAS/SATA drive, reseating [70](#)

storage controller, installing [130](#)
storage controller, removing [126](#)

U

upgrading CPU, [93](#)