



Cisco UCS C480 M5 Server Installation and Service Guide

First Published: 2017-10-12

Last Modified: 2021-01-19

Americas Headquarters

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

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

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

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

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

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

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

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

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

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

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

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

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

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)

© 2017– 2021 Cisco Systems, Inc. All rights reserved.



CHAPTER 1

Overview

- [Overview, on page 1](#)
- [External Features, on page 1](#)
- [Serviceable Component Locations, on page 4](#)
- [Summary of Server Features, on page 9](#)

Overview

This chapter provides a summary overview of the Cisco UCS C480 M5 server.


External Features

This topic shows the external features of the server.

Cisco UCS C480 M5 Server Front Panel Features

For definitions of LED states, see [Front-Panel LEDs, on page 33](#).

Figure 1: Cisco UCS C480 M5 Server Front Panel




1	Power button/LED	8	Left bay module (drive bays 1 - 8) • Bays 1, 2, 7, 8 support SAS/SATA and NVMe drives. Front NVMe drives are not supported in a single-CPU module system. • Bays 3, 4, 5, 6 support SAS/SATA drives only. Note An NVMe-only front drive module is available that supports up to 8 NVMe SSDs. You cannot mix this NVMe-only module with SAS/SATA modules or change module types in the field.
2	Identification button/LED	9	Center bay module (drive bays 9 - 16) • Bays 9, 10, 15, 16 support SAS/SATA or NVMe drives. Front NVMe drives are not supported in a single-CPU module system. • Bays 11, 12, 13, 14 support SAS/SATA drives only.
3	System status LED	10	Right bay module, supports either: • Optional DVD drive module • Drive bays 17 - 24 (shown) • Bays 17, 18, 23, 24 support SAS/SATA or NVMe drives. Front NVMe drives are not supported in a single-CPU module system. • Bays 19, 20, 21, 22 support SAS/SATA drives only.
4	Fan status LED	11	KVM console connector (used with a KVM cable that provides two USB, one VGA, and one serial connector)
5	Temperature status LED	12	Pull-out asset tag

6	Power supply status LED	13	CPU module bay 1 The system must have at least one CPU module in bay 1 to boot. It must also have either a CPU module or a blank filler module in bay 2.
7	Network link activity LED	14	CPU module bay 2 If no CPU module is present in bay 2, there must be a blank filler module in bay 2 or the system will not boot.

Cisco UCS C480 M5 Server Rear Panel Features

For definitions of LED states, see [Rear-Panel LEDs](#), on page 36.

Figure 2: Cisco UCS C480 M5 Server Rear Panel



1	Serial port COM 1 (DB-9 connector)	7	Rear identification button/LED
2	VGA video port (DB-15 connector)	8	USB 3.0 ports (three)
3	Not used at this time	9	Power supplies 1 – 4 (hot-swappable, redundant as 2+2 (default) or 3+1) See Power Specifications , on page 152 for specifications and supported options.
4	1-Gb/10-Gb Ethernet ports (LAN1 upper, LAN2 lower) The dual LAN ports can support 1 Gbps and 10 Gbps, depending on the link-partner capability.	10	Threaded holes for dual-hole grounding lug

5	10/100/1000 Ethernet dedicated management port (Base-T)	11	PCIe slots 1 – 12 See PCIe Slot Specifications and Restrictions , on page 94 for slot specifications.
6	Not used at this time	-	

Serviceable Component Locations


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

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*).

- [Serviceable Components Inside the Main Chassis](#), on page 5
- [Serviceable Components Inside a CPU Module](#), on page 8
- [Serviceable Components Inside an I/O Module](#), on page 9

Serviceable Components Inside the Main Chassis

Figure 3: Serviceable Component Locations Inside the Main Chassis




1	RAID controller card for front-loading drives. (not visible in this view; position is near chassis floor under CPU modules)	11	PCIe slot 01: Primary slot for Cisco UCS VIC adapter card. (Secondary slot for Cisco UCS VIC is slot 02.)
2	Supercap (RAID backup) for front RAID controller (not visible in this view; mounting bracket position is on chassis wall under CPU modules)	12	Power connectors for high-power GPU cards (six)
3	Fan modules (four modules with two fans each; hot-swappable)	13	Trusted platform module socket (TPM) on motherboard
4	Air diffuser for auxiliary rear drive module This diffuser is required only when using SAS/SATA drives in the rear drive module.	14	CPU modules (up to two, front-loading)
5	Position of the supercap unit (RAID backup) for the rear RAID controller. The clip for the supercap is on the inside surface of the air diffuser.	15	Left bay module (drive bays 1 - 8) • Bays 1, 2, 7, 8 support SAS/SATA or NVMe drives. Front NVMe drives are not supported in a single-CPU module system. • Bays 3, 4, 5, 6 support SAS/SATA drives only. Note An NVMe-only front drive module is available that supports up to 8 NVMe SSDs. You cannot mix this NVMe-only module with SAS/SATA modules or change module types in the field.
6	Auxiliary rear drive module; holds either (no mixing): • Up to eight 2.5-inch SAS/SATA drives • Up to eight 2.5-inch NVMe SSDs	16	Center bay module (drive bays 9 - 16) • Bays 9, 10, 15, 16 support SAS/SATA or NVMe drives. Front NVMe drives are not supported in a single-CPU module system. • Bays 11, 12, 13, 14 support SAS/SATA drives only.

7	Internal USB 2.0 socket on motherboard	17	Right bay module, supports either: • Drive bays 17 - 24 (shown) • Bays 17, 18, 23, 24 support SAS/SATA or NVMe drives. Front NVMe drives are not supported in a single-CPU module system. • Bays 19, 20, 21, 22 support SAS/SATA drives only. • Optional DVD drive module
8	PCIe slots 1 – 12 For PCIe slot specifications, see PCIe Slot Specifications and Restrictions, on page 94. PCIe slot 12 is not available when the auxiliary internal drive cage is used because of internal clearance.	18	I/O module Note The I/O module is not field replaceable, nor can you move an I/O module from one chassis to another. This module contains a security chip that requires it to stay with the PCIe module in the same chassis, as shipped from the factory.
9	PCIe slot 11: Default slot for rear RAID controller when the rear drive module is used with SAS/SATA drives. Note In systems with only one CPU module, slot 11 is not supported. In this case, the rear RAID controller must be installed in slot 10 and a blanking panel must be installed in slot 11.	19	Power supplies 1 – 4 (hot-swappable, redundant as 2+2 (default) or 3+1) All power supplies in the system must be identical (no mixing).
10	PCIe slot 10: Required slot for NVMe switch card when the rear drive module is used with NVMe SSDs. This slot must also be used for the rear RAID controller in systems with only one CPU module.	-	

Serviceable Components Inside a CPU Module

Figure 4: Serviceable Component Locations Inside a CPU Module




1	CPU number differs depending on the CPU module location: • CPU 2 and heatsink (when module is in lower bay 1) • CPU 4 and heatsink (when module is in upper bay 2) Note The CPUs in CPU module 1 must be identical with the CPUs in CPU module 2 (no mixing).	4	DIMM sockets controlled by CPU 1 or 3 (channels A, B, C, D, E, F.)
2	DIMM sockets controlled by CPU 2 or 4 (channels G, H, J, K, L, M.) See DIMM Population Rules and Memory Performance Guidelines, on page 114 for DIMM slot numbering.	5	Release levers for module (two each module)

3	CPU number differs depending on the CPU module location: • CPU 1 and heatsink (when module is in lower bay 1) • CPU 3 and heatsink (when module is in upper bay 2) Note The CPUs in CPU module 1 must be identical with the CPUs in CPU module 2 (no mixing).	-	
---	---	---	--

Serviceable Components Inside an I/O Module

Figure 5: Serviceable Component Locations Inside an I/O Module



1	Micro SD card socket	3	RTC battery vertical socket
2	Mini-storage module socket. Options: • SD card module with two SD card slots • M.2 module with slots for either two SATA M.2 drives or two NVMe M.2 drives • Cisco Boot-Optimized M.2 RAID Controller (module with two slots for SATA M.2 drives, plus an integrated SATA RAID controller that can control the two M.2 drives in a RAID 1 array)	-	

Summary of Server Features

The following table lists a summary of server features.

Feature	Description
Chassis	Four rack-unit (4RU) chassis
Central Processor	The server supports one or two removable CPU modules, each with two CPUs. Up to four CPUs from the Intel Xeon Processor Scalable Family. This includes CPUs from the following series: • Intel Xeon Gold 5XXX Processors • Intel Xeon Gold 6XXX Processors • Intel Xeon Platinum 8XXX Processors
Memory	Each of the CPUs support up to 12 DIMMs for a total of 48 DIMMs. The server supports up to two removable CPU modules, each with 24 DIMM sockets.
Multi-bit error protection	Multi-bit error protection is supported
Baseboard management	BMC, running Cisco Integrated Management Controller (Cisco IMC) firmware. Depending on your Cisco IMC settings, Cisco IMC can be accessed through the 1-Gb dedicated management port, the 1-Gb/10-Gb Ethernet LAN ports, or a Cisco virtual interface card.
Network and management I/O	The network and management I/O ports for this server are on a removable I/O module: • One 10/100/1000 Ethernet dedicated management port (RJ-45 connector) • One 10/100/1000 Ethernet private inter-chassis port (RJ-45 connector) • Two 1-Gb/10-Gb BASE-T Ethernet LAN ports (RJ-45 connectors) • One RS-232 serial port (DB-9 connector) • One VGA video connector port (DB-15 connector) • Two USB 3.0 ports Front panel: • One front-panel keyboard/video/mouse (KVM) connector that is used with the KVM cable, which provides two USB 2.0, one VGA, and one DB-9 serial connector.

Feature	Description
Power	Four power supplies, redundant as 2+2 (default) or 3+1: • AC power supplies 1600 W AC each Do not mix power supply types or wattages in the server.
ACPI	The advanced configuration and power interface (ACPI) 4.0 standard is supported.
Cooling	Four hot-swappable fan modules with two fans in each for front-to-rear cooling.
PCIe I/O	Twelve vertical PCIe expansion slots on the chassis motherboard. See PCIe Slot Specifications and Restrictions, on page 94 for specifications of the slots.
InfiniBand	The PCIe bus slots in this server support the InfiniBand architecture.
Storage, front-panel	The server can hold up to 24 front-loading, 2.5-inch drives. Front drive bays are divided across 3 removable drive bay modules. Each drive bay module has 8 drive bays for a total of 24 front-loading drive bays. • All 24 front drive bays support SAS/SATA drives. • Each of the 3 drive bay modules have 4 bays that support NVMe SSDs, for a total of 12 bays that support NVMe SSDs. Note An NVMe-only front drive module is available that supports up to 8 NVMe SSDs. You cannot mix this NVMe-only module with SAS/SATA modules or change module types in the field.

Feature	Description
Storage, internal	The server has these internal storage options: • The optional, rear drive bay module has 8 drive bays. All 8 bays support either all SAS/SATA drives or all NVMe SSDs. Mixing is not supported. • One USB 2.0 port on the chassis motherboard. • One micro-SD card socket on the I/O module board. • Mini-storage module socket on the I/O module board, optionally with either: • SD card carrier. Supports up to two SD cards. • M.2 SSD carrier. Supports two SATA M.2 SSDs. • Cisco Boot-Optimized M.2 RAID Controller (module with two slots for SATA M.2 drives, plus an integrated SATA RAID controller that can control the two SATA M.2 drives in a RAID 1 array)
Other removable media	A DVD drive module option is available in place of the right drive bay module.
Storage management	• Front-loading storage: the server has a dedicated internal socket near the chassis front for a single storage controller card (RAID or HBA). This controller card can control up to 24 front-loading drives. • Internal, auxiliary storage. The server supports these options: • When SAS/SATA drives are populated in the auxiliary rear drive bays, they can be controlled by a storage controller card in PCIe slot 11 or 10. • When NVMe SSDs are populated in the auxiliary drive bays, you must have an NVMe switch card in PCIe slot 10. • Cisco Boot-Optimized M.2 RAID Controller (module with two slots for SATA M.2 drives, plus an integrated SATA RAID controller that can control the two SATA M.2 drives in a RAID 1 array) For a detailed list of storage controller options, see Supported Storage Controllers and Cables, on page 157.

Feature	Description
RAID supercap backup	The server supports the following options when RAID cards are installed: • There is a bracket on the chassis wall for a supercap unit that backs up a front RAID controller for front-loading drives. • There is a bracket on the rear drive module diffuser for a supercap unit that backs up a rear RAID controller.
Integrated video	Integrated VGA video.



CHAPTER 2

Installing the Server

- [Preparing for Installation, on page 15](#)
- [Installing the Server in a Rack, on page 18](#)
- [Initial Server Setup, on page 24](#)
- [NIC Mode and NIC Redundancy Settings, on page 29](#)
- [Updating the BIOS and Cisco IMC Firmware, on page 30](#)
- [Accessing the System BIOS, on page 31](#)

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

**Warning**

This equipment must be grounded. Never defeat the ground conductor or operate the equipment in the absence of a suitably installed ground conductor. Contact the appropriate electrical inspection authority or an electrician if you are uncertain that suitable grounding is available.

Statement 1024

**Warning**

For Nordic countries (Norway, Finland, Sweden and Denmark) this system must be installed in a Restricted Access Location, where the voltage of the main ground connection of all equipment is the same (equipotential earth) and the system is connected to a grounded electrical outlet.

Statement 328

**Warning**

High leakage current – earth connection essential before connection to system power supply.

Statement 342

**Warning**

This equipment must be externally grounded using a customer-supplied ground wire before power is applied. Contact the appropriate electrical inspection authority or an electrician if you are uncertain that suitable grounding is available.

Statement 366

**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 151](#).
- Ensure that the cabinet or rack meets the requirements listed in the [Rack Requirements, on page 17](#).
- Ensure that the site power meets the power requirements listed in the [Power Specifications, on page 152](#). 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 four rack units (RUs), equal to 7.0 in. (177.8 mm).

Supported Cisco Slide Rail Kits

The server supports the following rail kit options:

- Cisco part UCSC-RAIL-4U-M5= (ball-bearing slide rail kit)
- Cisco part UCSC-CMA-4U-M5= (cable management arm)

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

Remove the inner rail from the slide rail assembly:

- a) Slide out the intermediate and inner slide rails until they click and lock in the fully open position.
- b) Hold down the inner rail release clip and at the same time, pull the inner rail free from the assembly.
- c) Push down the rail release latch while you collapse the intermediate rail back into the rail assembly.

Figure 6: Removing the Inner Rail From the Assembly



1	Rail release latch	2	Inner rail release clip
---	--------------------	---	-------------------------


Step 2

Attach the inner rails to the sides of the server:

Note The inner rails are not identical: there is a left rail and a right rail (as viewed from the chassis front). The inner rails are marked with “L” for left or “R” for right.

- Align the left inner rail marked “L” with the left side of the chassis (as viewed from the front). Align the 10 keyed slots in the rail with the 10 pegs on the side of the chassis.
- Set the keyed slots over the pegs, and then slide the rail toward the front to lock it in place on the pegs.
- Install the right inner rail marked “R” to the right side of the chassis (as viewed from the front).

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



306753

1	Left side of chassis	3	Right-side inner rail marked "R"
2	Right side of chassis	-	

Step 3 Install the slide rail assemblies into the rack:

Note The slide rail assemblies are not identical; there is a left rail and a right rail (as viewed from the rack front). The assemblies are marked with “L” for left or “R” for right.

- a) Align the front end of the left-side slide-rail assembly (marked “L”) with the left-front rack-post (as you face the front of the rack).

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.

The bottom of the slide rail assembly lines up with the intended bottom of the rack unit.

- b) Push the front mounting pegs into the rack-post holes until you hear them click and lock.
c) Adjust the slide-rail length until it reaches the rear rack post perfectly level.

Note Ensure that the rail is perfectly level and that the same height rack-post holes are used in the front and rear posts.


- d) Hold open the rear-peg spring latch, then push the rear mounting pegs into the rear rack-post-holes

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

- e) Release the rear-peg spring latch to lock the rear pegs in place.
f) 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.

Caution Ensure that all pegs are fully inserted into the rack post holes before installing the server to the rack.

Figure 8: Attaching the Rail Assembly to the Rack Post



1	Front mounting pegs, entering rack-post holes from the outside front	3	Rear peg spring-latch
2	Rear mounting pegs, entering rack-post holes from inside rear	-	

Step 4 Pull the inner slide rails on each assembly out toward the rack front until they hit the internal stops and lock in place.

Step 5 Insert the server into the slide rails:

Caution This server can weigh up to 146 pounds (66.2 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.

Note The rack rail channels are fragile to side loads. Install the chassis gently to avoid damaging the rails.

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


Caution Ensure that the inner rails on the chassis are parallel with the intermediate rails on the rack. This could require adjusting the mechanical lift position up-down and left-to-right. Consider using shims if your lift does not support these motions. The ends of the inner rail must align up-down and side-to-side with the ends of the intermediate rails.

- b) Very slowly push the chassis toward the rack. Ensure that the rail ends mesh to each other and are fully engaged at the top and bottom of the rail channels.

- c) Slowly push the inner rails into the slide rails on the rack until they stop at the internal stops.
- d) Press the release clip on each inner rail inward, and then continue pushing the server into the rack until its front slam-latches engage with the rack posts.

Caution Ensure that both inner rail release clips are pushed in before pushing the server into the rack. Push the server into the rails slowly to avoid damaging the rails. Let go of the release clip buttons as the server begins to push in.

Figure 9: Inner-Rail Release Clip



1	Intermediate rail extended from outer rail	3	Inner rail release clip
2	Inner rail attached to server	-	

Step 6 (Optional) Secure the server in the rack more permanently by using the two screws that are provided with the slide rails. Always 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.

Caution Depressing the release clips on the inner rails allows the chassis to slide all the way out of the intermediate rails and could result in injury or equipment damage. When you pull the chassis outward from the rack, it stops at internal locking stops. Do not depress the inner-rail release clips unless you intend to slide the chassis back into the rack or to fully remove the chassis from the rack. It is recommended that when you test pulling the chassis out from the rack for the first time that you place the mechanical lift under the chassis to avoid an accidental drop.


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 24 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 10: 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 11: Reversing the CMA



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

Initial Server Setup



Note This section describes how to power on the server, assign an IP address, and connect to server management when using the server in standalone mode. To use the server in Cisco UCS Manager integration, specific cabling and settings are required. See [Installation For Cisco UCS Manager Integration, on page 189](#).

Server Default Settings

The server is shipped with these default settings:

- The NIC mode is *Shared LOM EXT*.

Shared LOM EXT mode enables the 1-Gb/10-Gb Ethernet ports *and* the ports on any installed Cisco virtual interface card (VIC) to access the Cisco Integrated Management Interface (Cisco IMC). If you want to use the 10/100/1000 dedicated management port 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 27](#).

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

Connection Methods

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

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

Connecting to the Server Locally For Setup

This procedure requires the following equipment:

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

Step 1

Attach a power cord to each power supply in your server, and then attach each power cord to a grounded AC power outlet.

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) 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 bootstrap, 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 minimum 8 characters; maximum 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 27](#).

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.

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 user name 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 minimum 8 characters; maximum 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 27](#).

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

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

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

See also the required VIC Slot setting below.

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

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

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

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

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

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

Step 4 (Optional) Make VLAN settings.

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

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

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

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

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

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 1: Valid NIC Redundancy Settings For Each NIC Mode

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

Shared LOM	Active-active
	Active-standby
Cisco Card	Active-active
	Active-standby

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

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

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

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

See also the required VIC Slot setting below.

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 LOM and Cisco Card modes can each use either *Active-standby* or *Active-active* settings.
- *Active-active* (default)—All Ethernet ports are utilized simultaneously. The Shared LOM EXT mode must use only this NIC redundancy setting. Shared LOM and Cisco Card modes can each use either *Active-standby* or *Active-active* settings.

Updating the BIOS and Cisco IMC Firmware



Caution

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

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

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

- **Recommended method for firmware update:** Use the Cisco Host Upgrade Utility to simultaneously upgrade the Cisco IMC, BIOS, and component firmware to compatible levels.

See the *Cisco Host Upgrade Utility Quick Reference Guide* for your firmware release at the documentation roadmap link below.

- You can upgrade the Cisco IMC and BIOS firmware by using the Cisco IMC GUI interface.

See the *Cisco UCS C-Series Rack-Mount Server Configuration Guide*.

- You can upgrade the Cisco IMC and BIOS firmware by using the Cisco IMC CLI interface.

See the *Cisco UCS C-Series Rack-Mount Server CLI Configuration Guide*.

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

Accessing the System BIOS

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

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

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

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

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

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

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



CHAPTER 3

Maintaining the Server

This chapter contains the following sections:

- [Status LEDs and Buttons, on page 33](#)
- [Preparing For Component Installation, on page 38](#)
- [Serviceable Component Locations, on page 43](#)
- [Replacing Components Inside the Main Chassis, on page 49](#)
- [Replacing Components Inside a CPU Module, on page 99](#)
- [Replacing Components Inside an I/O Module, on page 122](#)
- [Recycling the PCB Assembly \(PCBA\), on page 133](#)
- [Service DIP Switches, on page 143](#)

Status LEDs and Buttons

This section contains information for interpreting LED states.

Front-Panel LEDs

Figure 12: Front Panel LEDs




Table 2: Front Panel LEDs, Definition of States

	LED Name	States
1	Power button/LED	• 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.
2	Unit identification button/LED	• Off—The unit identification function is not in use. • Blue, blinking—The unit identification function is activated.
3	System status LED	• Green—The server is running in normal operating condition. • Amber, steady—The server is in a degraded operational state (minor fault). For example: • 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 in a critical fault state. For example: • Boot failure • Fatal processor and/or bus error detected • Over-temperature condition
4	Fan status LED	• Green—All power supplies are operating normally. • Amber, steady—One or more power supplies are in a degraded operational state. • Amber, blinking—One or more power supplies are in a critical fault state.

5	Temperature status LED	• Green—All fan modules are operating properly. • Amber, steady—Fan modules are in a degraded state. One fan module has a fault. • Amber, blinking—Two or more fan modules have faults.
6	Power supply status LED	• Green—All power supplies are operating normally. • Amber, steady—One or more power supplies are in a degraded operational state. • Amber, blinking—One or more power supplies are in a critical fault state.
7	Network link activity LED	• Green—The server is operating at normal temperature. No error conditions detected. • Amber, steady—One or more temperature sensors exceeded a warning threshold. • Amber, blinking—One or more temperature sensors exceeded a critical non-recoverable threshold.
8 SAS	SAS/SATA drive fault LED Note NVMe solid state drive (SSD) drive tray LEDs have different behavior than SAS/SATA drive trays.	• 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.
9 SAS	SAS/SATA drive activity LED	• 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.
8 NVMe	NVMe SSD drive fault LED Note NVMe solid state drive (SSD) drive tray LEDs have different behavior than SAS/SATA drive trays.	• 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.

9 NVMe	NVMe SSD activity LED	- Off—No drive activity. - Green, blinking—There is drive activity.
10	CPU module power status LED	- Green—The CPU module is correctly seated and receiving power. - Off—There is no power to the CPU module or it is incorrectly seated.
11	CPU module fault LED	- Off—There is no fault with the CPUs or DIMMs on the CPU module board. - Amber—There is a fault with a CPU or DIMM on the CPU module board, such as an over-temperature condition.
-	DVD drive activity LED (optional DVD module not shown)	- Off—The drive is idle. - Green, steady—The drive is spinning up a disk. - Green, blinking—The drive is accessing data.

Rear-Panel LEDs

Figure 13: Rear Panel LEDs




Table 3: Rear Panel LEDs, Definition of States

LED Name	States
----------	--------

1	1-Gb/10-Gb Ethernet link speed (on both LAN1 and LAN2) These ports auto-negotiate link speed based on the link-partner capability.	• Off—Link speed is 100 Mbps. • Amber—Link speed is 1 Gbps. • Green—Link speed is 10 Gbps.
2	1-Gb/10-Gb Ethernet link status (on both LAN1 and LAN2)	• Off—No link is present. • Green—Link is active. • Green, blinking—Traffic is present on the active link.
3	1-Gb Ethernet dedicated management link speed	• Off—Link speed is 10 Mbps. • Amber—Link speed is 100 Mbps. • Green—Link speed is 1 Gbps.
4	1-Gb Ethernet dedicated management link status	• Off—No link is present. • Green—Link is active. • Green, blinking—Traffic is present on the active link.
5	Rear unit identification	• Off—The unit identification function is not in use. • Blue, blinking—The unit identification function is activated.
6	Power supply status (one LED each power supply unit)	AC power supplies: • 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).

Internal Diagnostic LEDs


The system has the following internal fault LEDs to help with identifying a failing component:

- Each chassis fan module has a fault LED on top of the module. These fan LEDs operate only when the system is in standby power mode.

- The CPU module has internal fault LEDs for CPUs and DIMMs on the CPU module board. POST and runtime error detection routines are stored in on-board registers. The contents of the registers are preserved for a limited time by a supercap voltage source.

To operate the LEDs, press switch SW1 on the board after the CPU module is removed from the chassis.

Figure 14: Internal Diagnostic LED Locations



1	CPU fault LEDs (one behind each CPU socket on the board). • Amber—CPU has a fault. • Off—CPU is OK.	3	DIMM fault LEDs (one next to each DIMM socket on the board) • Amber—DIMM has a fault. • Off—DIMM is OK.
2	Switch SW1 SW1 is labeled, "PRESS HERE TO SEE FAULTS".	-	

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 (supplied with replacement CPUs for heatsink removal)
- #1 Phillips-head screwdriver (for M.2 SSD 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 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.
-

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 UCS Manager Equipment Tab

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

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

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

- Step 7** 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 UCS Manager Service Profile


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

-
- Step 1** In the Navigation pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that contains the service profile of the server that you are shutting down.
- Step 4** Choose the service profile of the server that you are shutting down.
- Step 5** In the **Work** pane, click the **General** tab.
- Step 6** In the **Actions** area, click **Shutdown Server**.
- Step 7** If a confirmation dialog displays, click **Yes**.
- The operating system performs a graceful shutdown and the server goes to standby mode, which is indicated by an amber Power button/LED.
- Step 8** 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:
- If the cover latch is locked, use a screwdriver to turn the lock 90-degrees counterclockwise to unlock it.
 - Lift on the end of the latch that has the green finger grip. The cover is pushed back to the open position as you lift the latch.
 - Lift the top cover straight up from the server and set it aside.
- Step 2** Replace the top cover:
- 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. The opening in the latch should fit over the peg that sticks up from the fan tray.
 - Press the cover latch down to the closed position. The cover is pushed forward to the closed position as you push down the latch.
 - If desired, lock the latch by using a screwdriver to turn the lock 90-degrees clockwise.

Figure 15: Removing the Top Cover



1	Top cover	3	Serial number location on label
2	Locking cover latch		

Serial Number Location

The serial number for the server is printed on a label on the top of the server, near the front.

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 2+2 or 1+1)

- Hot-plug replacement—You must take the component offline before removing it for the following component:
 - NVMe PCIe solid state drives

Serviceable Component Locations


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

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*).

- [Serviceable Components Inside the Main Chassis, on page 44](#)
- [Serviceable Components Inside a CPU Module, on page 47](#)
- [Serviceable Components Inside an I/O Module, on page 48](#)

Serviceable Components Inside the Main Chassis

Figure 16: Serviceable Component Locations Inside the Main Chassis




1	RAID controller card for front-loading drives. (not visible in this view; position is near chassis floor under CPU modules)	11	PCIe slot 01: Primary slot for Cisco UCS VIC adapter card. (Secondary slot for Cisco UCS VIC is slot 02.)
2	Supercap (RAID backup) for front RAID controller (not visible in this view; mounting bracket position is on chassis wall under CPU modules)	12	Power connectors for high-power GPU cards (six)
3	Fan modules (four modules with two fans each; hot-swappable)	13	Trusted platform module socket (TPM) on motherboard
4	Air diffuser for auxiliary rear drive module This diffuser is required only when using SAS/SATA drives in the rear drive module.	14	CPU modules (up to two, front-loading)
5	Position of the supercap unit (RAID backup) for the rear RAID controller. The clip for the supercap is on the inside surface of the air diffuser.	15	Left bay module (drive bays 1 - 8) • Bays 1, 2, 7, 8 support SAS/SATA or NVMe drives. Front NVMe drives are not supported in a single-CPU module system. • Bays 3, 4, 5, 6 support SAS/SATA drives only. Note An NVMe-only front drive module is available that supports up to 8 NVMe SSDs. You cannot mix this NVMe-only module with SAS/SATA modules or change module types in the field.
6	Auxiliary rear drive module; holds either (no mixing): • Up to eight 2.5-inch SAS/SATA drives • Up to eight 2.5-inch NVMe SSDs	16	Center bay module (drive bays 9 - 16) • Bays 9, 10, 15, 16 support SAS/SATA or NVMe drives. Front NVMe drives are not supported in a single-CPU module system. • Bays 11, 12, 13, 14 support SAS/SATA drives only.

7	Internal USB 2.0 socket on motherboard	17	Right bay module, supports either: • Drive bays 17 - 24 (shown) • Bays 17, 18, 23, 24 support SAS/SATA or NVMe drives. • Front NVMe drives are not supported in a single-CPU module system. • Bays 19, 20, 21, 22 support SAS/SATA drives only. • Optional DVD drive module
8	PCIe slots 1 – 12 For PCIe slot specifications, see PCIe Slot Specifications and Restrictions , on page 94. PCIe slot 12 is not available when the auxiliary internal drive cage is used because of internal clearance.	18	I/O module Note The I/O module is not field replaceable, nor can you move an I/O module from one chassis to another. This module contains a security chip that requires it to stay with the PCIe module in the same chassis, as shipped from the factory.
9	PCIe slot 11: Default slot for rear RAID controller when the rear drive module is used with SAS/SATA drives. Note In systems with only one CPU module, slot 11 is not supported. In this case, the rear RAID controller must be installed in slot 10 and a blanking panel must be installed in slot 11.	19	Power supplies 1 – 4 (hot-swappable, redundant as 2+2 (default) or 3+1) All power supplies in the system must be identical (no mixing).
10	PCIe slot 10: Required slot for NVMe switch card when the rear drive module is used with NVMe SSDs. This slot must also be used for the rear RAID controller in systems with only one CPU module.	-	

Serviceable Components Inside a CPU Module

Figure 17: Serviceable Component Locations Inside a CPU Module




1	CPU number differs depending on the CPU module location: • CPU 2 and heatsink (when module is in lower bay 1) • CPU 4 and heatsink (when module is in upper bay 2) Note The CPUs in CPU module 1 must be identical with the CPUs in CPU module 2 (no mixing).	4	DIMM sockets controlled by CPU 1 or 3 (channels A, B, C, D, E, F.)
2	DIMM sockets controlled by CPU 2 or 4 (channels G, H, J, K, L, M.) See DIMM Population Rules and Memory Performance Guidelines, on page 114 for DIMM slot numbering.	5	Release levers for module (two each module)

3	CPU number differs depending on the CPU module location: • CPU 1 and heatsink (when module is in lower bay 1) • CPU 3 and heatsink (when module is in upper bay 2) Note The CPUs in CPU module 1 must be identical with the CPUs in CPU module 2 (no mixing).	-	
---	---	---	--

Serviceable Components Inside an I/O Module

Figure 18: Serviceable Component Locations Inside an I/O Module



1	Micro SD card socket	3	RTC battery vertical socket
2	Mini-storage module socket. Options: • SD card module with two SD card slots • M.2 module with slots for either two SATA M.2 drives or two NVMe M.2 drives • Cisco Boot-Optimized M.2 RAID Controller (module with two slots for SATA M.2 drives, plus an integrated SATA RAID controller that can control the two M.2 drives in a RAID 1 array)	-	

Replacing Components Inside the Main Chassis



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 main chassis components. See also:

- [Replacing Components Inside a CPU Module, on page 99](#)
- [Replacing Components Inside an I/O Module, on page 122](#)

Replacing a CPU Module

CPU Module Population Rules:

- The server can operate with one or two CPU modules.
- If you have only one CPU module, populate lower bay 1 first.
- If no CPU module is present in upper bay 2, you must insert a blank filler module or the system will not boot.
- The following restrictions apply when using only a two-CPU configuration (CPU module 2 is not present):
 - The maximum number of DIMMs is 24 (only CPU 1 and CPU 2 memory channels).
 - Some PCIe slots are unavailable when CPU module 2 is not present:

PCIe Slots Controlled by CPU Module 1 (CPUs 1 and 2)	PCIe Slots Controlled by CPU Module 2 (CPUs 3 and 4)
1, 2, 5, 8, 9, 10	3, 4, 6, 7, 11, 12

- Only four double-wide GPUs are supported, in PCIe slots 1, 2, 8, and 10.
- No front NVMe drives are supported.
- The optional NVMe-only drive bay module UCSC-C480-8NVME is not supported.
- If a rear RAID controller is used, it must be installed in PCIe slot 10 rather than the default slot 11. A blank filler must be installed in slot 11.



Note Each CPU module has a fault LED on its front that turns amber to help to identify which CPU module has a fault.



Caution Never remove a CPU module without shutting down and removing power from the server.

Step 1 Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server](#), on page 39.


You do not have to pull the server out from the rack or remove the cover because the CPU modules are accessed from the front of the chassis.

Step 2 Remove an existing CPU module:

Note Verify that the power LED on the front of the CPU module is off before removing the module.

- Grasp the two ejector levers on the module and pinch their latches to release the levers.
- Rotate both levers to the outside at the same time to evenly disengage the module from the midplane connectors.
- Pull the module straight out from the chassis and then set it on an antistatic surface.

Figure 19: CPU Module Front



1	Ejector levers (two each CPU module)	3	CPU module fault LED
2	CPU module power status LED	-	

Step 3 If you are moving CPUs from the old CPU module to the new CPU module, see [Moving an M5 Generation CPU](#), on page 109.

Step 4 If you are moving DIMMs from the old CPU module to the new CPU module, perform the following steps:

- Open the ejector lever at each end of the DIMM slot and pick the DIMM straight up from the old CPU module board.
- On the new CPU module board, align the new DIMM with an empty slot. Use the alignment feature in the DIMM slot to correctly orient the DIMM.

- c) Push down evenly on the top corners of the DIMM until it is fully seated and the ejector levers on both ends lock into place.

Step 5 Install a new CPU module to the chassis:

- With the two ejector levers open, align the new CPU module with an empty bay.
- Push the module into the bay until it engages with the midplane connectors and is flush with the chassis front.
- Rotate both ejector levers toward the center until they lay flat and their latches lock into the front of the module.

Step 6 Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).

Step 7 Fully power on the server by pressing the Power button.

Note Verify that the power LED on the front of the CPU module returns to solid green.

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.

Front-Loading SAS/SATA Drive Population Guidelines

The front drives in the server are installed into three removable drive bay modules.

Figure 20: Drive Bay Numbering



- SAS/SATA/NVMe drive bay modules (UCSC-C480-8HDD):
 - Left drive bay module: Bays 1, 2, 7, 8 support SAS/SATA or NVMe drives; bays 3, 4, 5, 6 support SAS/SATA drives only.



Note Front NVMe drives are not supported in a system with only one CPU module.

- Center drive bay module: Bays 9, 10, 15, 16 support SAS/SATA or NVMe drives; bays 11, 12, 13, 14 support SAS/SATA drives only.
- Right bay module: Bays 17, 18, 23, 24 support SAS/SATA or NVMe drives; bays 19, 20, 21, 22 support SAS/SATA drives only.

Observe these drive population guidelines for optimum performance:

- 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.

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:
- Press the release button on the face of the drive tray.
 - Grasp and open the ejector lever and then pull the drive tray out of the slot.
 - If you are replacing an existing drive, remove the four drive-tray screws that secure the drive to the tray and then lift the drive out of the tray.
- Step 2** Install a new drive:
- Place a new drive in the empty drive tray and install the four drive-tray screws.
- Note** When you insert the drive tray in the slot, the LEDs on the drive tray must be on the upper side. The ejector lever closes upward.
- With the ejector lever on the drive tray open, insert the drive tray into the empty drive bay.
 - Push the tray into the slot until it touches the backplane, and then close the ejector lever to lock the drive in place.

Figure 21: Replacing a Drive in a Drive Tray



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

Replacing Rear (Internal) 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 SAS/SATA Drive Population Guidelines


The server supports an internal, rear drive module that holds up to eight 2.5-inch drives.

- When using SAS/SATA drives, the eight drives must be all SAS/SATA; no mixing with NVMe drives is allowed.
- 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 cage. 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.



Note See also [4K Sector Format SAS/SATA Drives Considerations](#), on page 52.

Figure 22: Internal Drive Module Bays (Top View)



Rear SAS/SATA Drive Requirements

Observe these requirements:

- The optional, rear drive module (UCSC-C480-8HDD).
- The rear drive module requires minimum Cisco IMC and BIOS 3.1(3) or later.
- The rear drive-bay module must have air-diffuser UCSC-DIFF-C480M5 installed when SAS/SATA drives are installed.

- For RAID support: RAID controller card (UCSC-SAS9460-8i) installed in PCIe slot 11.

**Note**

In a system with only one CPU module, this RAID controller must be installed in PCIe slot 10 and a blank filler is required in slot 11 to ensure adequate air flow.

- For RAID support: RAID cable (CBL-AUX-SAS-M5). This cable connects the rear RAID card to the drive-bay module.
- For RAID support: Supercap RAID backup unit (UCSC-SCAP-M5). This unit installs to a clip on the inside of the air diffuser. It cables to the rear RAID controller.

Replacing a Rear (Internal) 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 Prepare the server for component installation:

- a) Slide the server out the front of the rack far enough so that you can remove the top cover.

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 41](#).

Step 2 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 up out of the bay.
- 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 3 Install a new drive:


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

Step 4 Replace the top cover to the server.

Step 5 Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).

Step 6 Fully power on the server by pressing the Power button.

Figure 23: Replacing a Drive in a Drive Tray



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

Replacing Front-Loading NVMe SSDs



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 59](#).




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

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, depending on what type of drive bay module is installed (NVMe-only or SAS/SATA/NVMe), and the number of CPU modules in the system:

Figure 24: Drive Bay Numbering



Note Front NVMe drives are not supported in a single CPU-module system. Front NVMe support requires two CPU modules in the system.

There are two types of front drive bay modules that support NVMe drives:



Note You cannot mix front drive module types in the same system.

- **UCSC-C480-8HDD:** SAS/SATA/NVMe drive bay modules that support up to four NVMe drives each:
 - Left drive-bay module: Bays 1, 2, 7, 8 support SAS/SATA or NVMe drives; bays 3, 4, 5, 6 support SAS/SATA drives only.
Front NVMe drives are not supported in a single-CPU module system.
 - Center drive-bay module: Bays 9, 10, 15, 16 support SAS/SATA or NVMe drives; bays 11, 12, 13, 14 support SAS/SATA drives only.
Front NVMe drives are not supported in a single-CPU module system.
 - Right drive-bay module: Bays 17, 18, 23, 24 support SAS/SATA or NVMe drives; bays 19, 20, 21, 22 support SAS/SATA drives only.
Front NVMe drives are not supported in a single-CPU module system.
- **UCSC-C480-8NVME:** NVMe-only drive bay modules. All eight bays support only NVMe drives.

In a single CPU-module system, this NVMe-only module is not supported.

Observe these drive population guidelines for optimum performance:

- When populating drives, add drives to the lowest-numbered bays first.
- Keep an empty blanking tray in any unused bays to ensure proper airflow.

Front-Loading NVMe SSD Requirements and Restrictions

Observe these requirements:

- 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. See [Enabling Hot-Plug Support in the System BIOS, on page 59](#).

Observe these restrictions:

- NVMe 2.5-inch 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 52](#) or [Setting Up UEFI Mode Booting in the Cisco IMC GUI, on page 53](#).
- 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 and HHHL form-factor 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* HHHL form-factor SSDs is an invalid configuration. A valid configuration is two *HGST* NVMe SFF 2.5-inch SSDs and two *HGST* HHHL form-factor SSDs.
- 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.

Procedure

	Command or Action	Purpose
Step 1		

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-inch 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 59](#).


- Step 1** Remove an existing front-loading NVMe SSD:
- Shut down the NVMe SSD to initiate an OS-informed removal. Use your operating system interface to shut down the drive, and then observe the drive-tray LED:
 - Green—The drive is in use and functioning properly. Do not remove.
 - Green, blinking—the driver is unloading following a shutdown command. Do not remove.
 - Off—The drive is not in use and can be safely removed.
 - Press the release button on the face of the drive tray.
 - Grasp and open the ejector lever and then pull the drive tray out of the slot.
 - Remove the four drive tray screws that secure the SSD to the tray and then lift the SSD out of the tray.
- Step 2** Install a new front-loading NVMe SSD:
- Place a new SSD in the empty drive tray and install the four drive-tray screws.

Note When you insert the drive tray in the slot, the LEDs on the drive tray must be on the upper side. The ejector lever closes upward.
 - With the ejector lever on the drive tray open, insert the drive tray into the empty drive bay.
 - Push the tray into the slot until it touches the backplane, and then close the ejector lever to lock the drive in place.

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

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

Figure 25: 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 NVMe SSDs



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 59](#).


This section is for replacing 2.5-inch form-factor NVMe solid-state drives (SSDs) in the internal, rear drive-bay module.

Rear NVMe SSD Population Guidelines

The server supports an rear, internal drive-bay module that holds up to eight 2.5-inch drives.

- When using NVMe drives, the eight drives must be all NVMe; no mixing with SAS/SATA drives is allowed.
- 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.

Figure 26: Internal Drive Module Bays (Top View)



Rear NVMe SSD Requirements and Restrictions

Observe these requirements:

- The optional, rear drive-bay module. When using NVMe drives, the eight drives must be all NVMe; no mixing with SAS/SATA drives is allowed.
- NVMe switch card (UCSC-NVME-SC). This card must be installed in PCIe slot 10.
- NVMe cable (CBL-AUX-NVME-M5). This cable connects the NVMe switch card to the module 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 52](#) or [Setting Up UEFI Mode Booting in the Cisco IMC GUI, on page 53](#).
- 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 and HHHL form-factor 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* HHHL form-factor SSDs is an invalid configuration. A valid configuration is two HGST NVMe SFF 2.5-inch SSDs and two HGST HHHL form-factor SSDs.
- 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 (Internal) NVMe Drive

This topic describes how to replace 2.5-inch form-factor NVMe SSDs in the internal drive bays. You do not have to shut down the server, but you must shut down the NVMe drive before removal to avoid data loss.



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 59](#).

Step 1

Prepare the server for component installation:


- Slide the server out the front of the rack far enough so that you can remove the top cover.

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 41](#).

- Step 2** Remove the drive that you are replacing or remove a blank drive tray from the bay:
- Shut down the NVMe SSD to initiate an OS-informed removal. Use your operating system interface to shut down the drive, and then observe the drive-tray LED:
 - Green—The drive is in use and functioning properly. Do not remove.
 - Green, blinking—the driver is unloading following a shutdown command. Do not remove.
 - Off—The drive is not in use and can be safely removed.
 - Press the release button on the face of the drive tray.
 - Grasp and open the ejector lever and then pull the drive tray up out of the bay.
 - 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 3** Install a new drive:
- Place a new drive in the empty drive tray and install the four drive-tray screws.
 - With the ejector lever on the drive tray open, insert the drive tray into the empty drive bay.
 - Push the tray into the slot until it touches the backplane, and then close the ejector lever to lock the drive in place.
- Step 4** 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.
- Step 5** Replace the top cover to the server.
- Step 6** Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).
- Step 7** Fully power on the server by pressing the Power button.

Figure 27: 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 HHHL Form-Factor NVMe Solid State Drives

This section is for replacing half-height, half-length (HHHL) form-factor NVMe SSDs in the PCIe slots.

HHHL SSD Population Guidelines

Observe the following population guidelines when installing HHHL form-factor NVMe SSDs:

- Dual CPU-Module systems—You can populate up to 12 HHHL form-factor SSDs, using PCIe slots 1 – 12.



Note Other installed components affect how many PCIe slots are open to use. For example:

- If the auxiliary, internal drive module is installed, PCIe slot 12 is not available because of internal clearance.
 - If the server has a rear RAID controller card, it must be installed in PCIe slot 11 (or slot 10 in single CPU-module systems).
 - If the server has a rear NVMe switch card, it must be installed in PCIe slot 10.
-
- Single CPU-Module systems—In a single CPU-module system (CPU module 2 is not present), PCIe slots 3, 4, 6, 7, 11, and 12 are not available. Therefore, the maximum number of HHHL form-factor SSDs you can populate is 6, using only PCIe slots 1, 2, 5, 8, 9, and 10.

Number of CPU Modules	PCIe Slots Supported
Dual CPU-Module System (4 CPUs)	1 - 12 (all)
Single CPU-Module System (2 CPUs)	1, 2, 5, 8, 9, 10

HHHL Form-Factor NVMe SSD Restrictions

Observe these restrictions:

- You cannot boot from an HHHL form-factor NVMe SSD.
- You cannot control HHHL NVMe SSDs with a SAS RAID controller because NVMe SSDs interface with the server via the PCIe bus.
- You can combine NVMe SFF 2.5- or 3.5-inch SSDs and HHHL form-factor 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* HHHL form-factor SSDs is an invalid configuration. A valid configuration is two *HGST* NVMe SFF 2.5-inch SSDs and two *HGST* HHHL form-factor SSDs.

Replacing HHHL NVMe Drives

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 39](#).
- 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 41](#).

Step 2 Remove any existing HHHL drive from the slot (or a blanking panel):

- a) Open the hinged retainer bar that covers the top of the PCIe slot from which you are removing the HHHL drive.

Use your fingertips to pull back on the wire locking-latches at each end of the retainer bar, and then hinge the bar open to expose the tops of the PCIe slots.

- b) Pull both ends of the HHHL drive's card vertically to disengage the card from the socket, and then set it aside.

Step 3

Install a new HHHL drive:

- Carefully align the HHHL drive's card edge with the PCIe socket.
- Push on both corners of the card to seat its connector in the socket.
- Close the hinged retainer bar over the top of the PCIe slots.

Use your fingertips to pull back on the wire locking-latches at each end of the retainer bar, and then hinge it closed to lock in the tops of the PCIe slots. Push the wire locking-latching back to the forward, locked position.

Step 4

Replace the top cover to the server.


Step 5

Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).

Step 6

Fully power on the server by pressing the Power button.

Figure 28: PCIe Slot Hinged Retainer Bars



1	Wire locking latches for left PCIe retainer bar (slots 10 - 12)	2	Wire locking latches for right PCIe retainer bar (slots 1 - 9)
---	---	---	--

Replacing a Front Drive Bay Module

The front drive bays are divided across three removable drive bay modules that have eight bays each. There are two types of drive bay modules:


- SAS/SATA and NVMe (UCSC-C480-8HDD)

- NVMe only (UCSC-C480-8NVME)



Note Mixing these two module types in the same chassis is not supported.

- Step 1** Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 39](#).
- 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 41](#).
- Step 4** Remove all CPU modules from the chassis to provide clearance:
- Grasp the two ejector levers on the module and pinch their latches to release the levers.
 - Rotate both levers to the outside at the same time to evenly disengage the module from the midplane connectors.
 - Pull the module straight out from the chassis and then set it on an antistatic surface.
- Step 5** Remove an existing drive bay module:
- Remove any drives from the existing module and set them aside.
 - From the top of the chassis, loosen the single captive screw that secures the module to the chassis brace.
 - Disconnect any SAS cables from the rear of the module.
 - Push the module out the front of the chassis.
 - Pull the module and its attached interposer board out the front of the chassis and then set it aside.
- Step 6** Install a new drive module:
- Insert the new module with attached interposer into the opening in the chassis front.
 - Gently slide the module into the opening, ensuring that the connector on the end of the interposer board engages with the socket on the chassis midplane. Press until the front edges of the module align evenly with the chassis.
 - Tighten the single captive screw that secures the module to the chassis brace.
- Step 7** Connect any SAS cables that you disconnected earlier to the new drive module.
- Step 8** Install your drives to the bays in the module.
- Step 9** Reinstall the CPU modules to the chassis:
- With the two ejector levers open, align the new CPU module with an empty bay.
 - Push the module into the bay until it engages with the midplane connectors and is flush with the chassis front.
 - Rotate both ejector levers toward the center until they lay flat and their latches lock into the front of the module.
- Step 10** Replace the top cover to the server.
- Step 11** Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).
- Step 12** Fully power on the server by pressing the Power button.

Figure 29: Front Drive-Bay Module Securing Screws (CPU Modules Removed)

306267

1	Front of server (view of front compartment shown with both CPU modules removed)	3	Thumbscrews that secure drive bay modules (one each module)
2	Chassis brace		

Replacing a Front RAID Controller Card

For detailed information about storage controllers in this server, see [Supported Storage Controllers and Cables, on page 157](#).

The server supports one front RAID controller card for control of up to 24 front-loading SAS/SATA drives. The card installs to a dedicated, horizontal socket on the chassis midplane. The socket is below the CPU modules and can be accessed from the top of the server after the CPU modules are removed.

Firmware on the storage controller 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.

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](#).



Note **For servers running in standalone mode only:** After you replace front controller hardware (UCSC-RAID-M5HD), 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. This is necessary to program the controller's suboem-id to the correct value for the server SKU. If you do not do this, drive enumeration might not display correctly in the software. This issue does not affect servers controlled in UCSM mode.

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 39](#).
- 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 41](#).

Step 2 Remove all CPU modules from the chassis to provide clearance:

- Grasp the two ejector levers on the module and pinch their latches to release the levers.
- Rotate both levers to the outside at the same time to evenly disengage the module from the midplane connectors.
- Pull the module straight out from the chassis and then set it on an antistatic surface.

Step 3 Remove any existing front RAID controller card from the server:

- Disconnect any SAS and supercap cables from the existing card.
- Remove the metal retainer plate that secures the front edge of the RAID card. Loosen its two captive screws and then lift the plate out of the chassis and set it aside.
- Open the card's ejector lever to unseat it from the horizontal socket on the midplane.

- d) Pull both ends of the card horizontally to disengage the card from the socket, and then set it aside.

Step 4 Install a new front RAID controller card:

- a) Carefully align the card edge with the dedicated horizontal socket on the midplane.
- b) Push on both corners of the card to seat its connector in the socket.
- c) Fully close the ejector lever on the card to lock the card into the socket.
- d) Reinstall the metal retainer plate. Align it over the two threaded standoffs, and then tighten both captive screws.
- e) Reconnect any SAS and supercap cables to the new card.

Card connectors A1-A2 connect to SAS drive bay 1; card connectors B1-B2 connect to SAS drive bay 2; card connectors C1-C2 connect to SAS drive bay 3.

Step 5 Reinstall the CPU modules to the chassis:

- a) With the two ejector levers open, align the new CPU module with an empty bay.
- b) Push the module into the bay until it engages with the midplane connectors and is flush with the chassis front.
- c) Rotate both ejector levers toward the center until they lay flat and their latches lock into the front of the module.

Step 6 Replace the top cover to the server.

Step 7 Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).


Step 8 Fully power on the server by pressing the Power button.

Step 9 If your server is running in standalone mode, use the Cisco UCS Host Upgrade Utility to update the controller firmware and program the correct suboem-id for the controller.

Note **For servers running in standalone mode only:** After you replace front controller hardware (UCSC-RAID-M5HD), 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. This is necessary to program the controller's suboem-id to the correct value for the server SKU. If you do not do this, drive enumeration might not display correctly in the software. This issue does not affect servers controlled in UCSM mode.

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](#).

Figure 30: Front RAID Controller Card Location (CPU Modules removed)



1	Location of front RAID card in dedicated horizontal socket (view of the front compartment shown with the CPU modules removed)	3	Card ejector lever (magnified view)
2	Metal retainer plate securing screws		


Replacing the Front RAID Supercap Unit

This server supports installation of up to two supercap units, one for a front RAID controller and one for a rear RAID controller. The front supercap unit mounts to a bracket on the inner chassis wall, below the CPU modules.

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:
- Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 39](#).
 - 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 41](#).
- Step 2** Remove all CPU modules from the chassis to provide clearance:
- Grasp the two ejector levers on the module and pinch their latches to release the levers.
 - Rotate both levers to the outside at the same time to evenly disengage the module from the midplane connectors.
 - Pull the module straight out from the chassis and then set it on an antistatic surface.
- Step 3** Remove an existing supercap unit:
- Disconnect the supercap cable from the existing supercap.
 - Lift gently on the top securing tab that holds the supercap unit to its bracket.
 - Lift the supercap unit free of the bracket and set it aside.
- Step 4** Install a new supercap unit:
- Lift gently on the top securing tab on the bracket while you set the supercap unit into the bracket. Relax the tab so that it closes over the top of the supercap.
 - Connect the supercap cable from the RAID controller card to the connector on the new supercap cable.
- Step 5** Reinstall the CPU modules to the chassis:
- With the two ejector levers open, align the new CPU module with an empty bay.
 - Push the module into the bay until it engages with the midplane connectors and is flush with the chassis front.
 - Rotate both ejector levers toward the center until they lay flat and their latches lock into the front of the module.
- Step 6** Replace the top cover to the server.
- Step 7** Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).
- Step 8** Fully power on the server by pressing the Power button.

Figure 31: Front Supercap Bracket Location (Below CPU Modules)



306269

1	Supercap bracket location on inner chassis wall (view of the front compartment shown is with the CPU modules removed)	-	
---	---	---	--

Replacing a Rear (Internal) Drive-Bay Module

The optional, rear drive-bay module provides eight drive bays.



Note When the rear drive-bay module is used, PCIe slot 12 is not available because there is not enough clearance.



Note When the rear drive-bay module is populated with SAS/SATA drives, the air diffuser must be installed.

Step 1 Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 39](#).

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 41](#).

Step 4 Remove any existing rear drive-bay module:

- Remove any drives from the existing rear drive-bay module and set them aside.
- If the air diffuser is present on the module, remove the diffuser. Lift straight up on the diffuser and set it aside.

It is not necessary to remove the rear supercap unit from the diffuser.

- Disconnect any cable from a RAID controller or NVMe switch card from the module connectors.
- Loosen the two screws that secure the module to the chassis.
- Grasp the module at each end and lift up evenly to disengage its connector from the socket on the motherboard.

Step 5 Install a new rear drive-bay module:


- While holding the new module level, align it over the socket on the motherboard and the two screw-holes.
- Gently press the module connector to the motherboard socket. Stop when the module frame sits flat over the screw-holes.
- Install the two screws that secure the module to the chassis.
- Connect any cable from a RAID card or NVMe switch card to the new module backplane.
- Reinstall the air diffuser to the module if you removed one earlier (required only if the module is populated with SAS/SATA drives).

Note In a system with only one CPU module, an additional filler panel is required in PCIe slot 11 to ensure adequate air flow. See [Replacing a Rear RAID Controller Card, on page 81](#) for more information.

- Install your drives to the bays in the new module.

- Step 6** Replace the top cover to the server.
- Step 7** Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).
- Step 8** Fully power on the server by pressing the Power button.

Figure 32: Internal Rear Drive Module




1	Air diffuser top view This diffuser is required when SAS/SATA drives are installed in the rear drive module.	4	Alignment flange on chassis floor
2	Diffuser alignment points against the chassis mid-brace	5	Two drive module securing screws
3	Rear RAID supercap unit location on the inside surface of the diffuser	-	

Replacing an Air Diffuser on the Rear Drive Module

The air diffuser UCSC-BAFF-C480-M5 must be installed on the rear drive module when SAS/SATA hard drives or solid state drives are installed. The diffuser includes a clip for the rear supercap unit on its inside surface.

-
- 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 39](#).
 - 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 41](#).
- Step 2** Remove the air diffuser:
- Grasp the two top edges of the diffuser and lift straight up to free it from the grooves on the chassis mid-brace.
 - If there is a supercap unit present in the clip on the inside of the diffuser, gently pry the supercap from its clip and set it aside. Do not disconnect the supercap cable.
- Step 3** Install the new air diffuser:
- Set the supercap unit into the clip on the inside of the air diffuser and push gently until it clicks in place and is secured.
 - Position SAS and supercap cables so that they do not interfere with the diffuser installation. Cables must route out the rear of the diffuser.
 - Set the air diffuser in place and carefully lower it, using the grooves in the chassis mid-brace as guides. Make sure that the diffuser alignment flange sits flat on the chassis floor and against PCIe slot 11.
- Step 4** Replace the top cover to the server.
- Step 5** Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).
- Step 6** Fully power on the server by pressing the Power button.

Figure 33: Rear Drive Module Air Diffuser



1	Air diffuser top view This diffuser is required when SAS/SATA drives are installed in the rear drive module.	3	Rear RAID supercap unit location on the inside surface of the diffuser
2	Diffuser alignment points against the chassis mid-brace	4	Alignment flange on chassis floor

Replacing the Rear RAID Supercap Unit

This server supports installation of up to two supercap units, one for a front RAID controller and one for a rear RAID controller. The rear supercap unit mounts to a clip on the air diffuser that wraps around the internal drive module.

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:

- Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server](#), on page 39.

- 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 41](#).

Step 2 Remove an existing rear supercap unit:

- a) Grasp the two top edges of the air diffuser and lift straight up to free it from the grooves on the chassis mid-brace.
- b) Remove the supercap unit from the clip that is on the inside of the air diffuser.
- c) Disconnect the supercap cable from the rear RAID controller.

Step 3 Install a new supercap:


- a) Set the new supercap unit into the clip on the inside of the air diffuser and push gently until it clicks in place and is secured.
- b) Connect the supercap cable to the rear RAID controller card.
- c) Position SAS and supercap cables so that they do not interfere with the diffuser installation. Cables must route out the rear of the diffuser.
- d) Set the air diffuser in place and carefully lower it, using the grooves in the chassis mid-brace as guides. Make sure that the diffuser alignment flange sits flat on the chassis floor and against PCIe slot 11.

Step 4 Replace the top cover to the server.

Step 5 Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).

Step 6 Fully power on the server by pressing the Power button.

Figure 34: Rear Drive Module Air Diffuser and Supercap Unit Location



1	Air diffuser top view This diffuser is required when SAS/SATA drives are installed in the rear drive module.	3	Rear RAID supercap unit location on the inside surface of the diffuser
2	Diffuser alignment points against the chassis mid-brace	4	Alignment flange on chassis floor

Replacing a Rear RAID Controller Card

The server supports one rear RAID controller card for control of up to eight internal SAS/SATA drives in the optional auxiliary drive module.



Note The default slot for a rear RAID controller is PCIe slot 11. However, in a single CPU-module system, slot 11 is not supported. In this case, install the rear RAID controller in PCIe slot 10 and install the required blank filler to PCIe slot 11 to ensure adequate air flow.

For detailed information about storage controllers in this server, see [Supported Storage Controllers and Cables, on page 157](#).

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.

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](#).



Note **For servers running in standalone mode only:** After you replace rear controller hardware, 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. This is necessary to program the controller's suboem-id to the correct value for the server SKU. If you do not do this, drive enumeration might not display correctly in the software. This issue does not affect servers controlled in UCSM mode.

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 39](#).
- 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 41](#).

Step 2 Remove an existing rear RAID card:

- a) Disconnect the SAS and supercap cables from the existing card.
- b) Open the hinged retainer bar that covers the top of PCIe slot 11 or 10.

Use your fingertips to pull back on the wire locking-latches at each end of the retainer bar, and then hinge the bar open to expose the tops of the PCIe slots. See [Replacing a PCIe Card, on page 96](#).

- c) Open the card's blue ejector lever to unseat it from the slot.
- d) Pull both ends of the card vertically to disengage the card from the socket, and then set it aside.

Step 3 Install a new rear RAID controller card:

- a) Carefully align the card edge with the socket of PCIe slot 11 (or 10 in a single CPU module system).
- b) Push on both corners of the card to seat its connector in the socket.
- c) Fully close the blue ejector lever on the card to lock the card into the socket.
- d) Connect SAS cable (CBL-AUX-SAS-M5) and the supercap cable to the new card.
- e) Close the hinged retainer bar over the top of the PCIe slots.

Use your fingertips to pull back on the wire locking-latches at each end of the retainer bar, and then hinge it closed to lock in the tops of the PCIe slots. Push the wire locking-latching back to the forward, locked position.

Step 4 Replace the top cover to the server.

Step 5 Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).

Step 6 Fully power on the server by pressing the Power button.


Step 7

If your server is running in standalone mode, use the Cisco UCS Host Upgrade Utility to update the controller firmware and program the correct suboem-id for the controller.

Note **For servers running in standalone mode only:** After you replace rear controller hardware, 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. This is necessary to program the controller's suboem-id to the correct value for the server SKU. If you do not do this, drive enumeration might not display correctly in the software. This issue does not affect servers controlled in UCSM mode.

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](#).

Figure 35: Rear RAID Card and PCIe Slot 11 Filler (Single CPU Module System Shown)



1	PCIe slot 11 (primary position for a rear RAID controller card) In this view, a blank filler is installed in slot 11. This is required only for single CPU module systems to ensure air flow.	2	PCIe slot 10, secondary slot for rear RAID controller In a single CPU module system slot 11 is not supported, so the controller must install to slot 10.
----------	--	----------	---

Replacing a Rear NVMe Switch Card

When you install NVMe drives in the rear drive-bay module, you must have an NVMe switch card in PCIe slot 10. A PCIe cable connects the switch card to the drive-bay module backplane.



Note If a rear NVMe switch card is used, it must be installed in PCIe slot 10.

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 39](#).
- 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 41](#).

Step 2 Remove any existing rear NVMe switch card from PCIe slot 10:

- a) Disconnect the PCIe cable from the existing card.
- b) Open the hinged retainer bar that covers the top of PCIe slot 10. See [Replacing a PCIe Card, on page 96](#).

Use your fingertips to pull back on the wire locking-latches at each end of the retainer bar, and then hinge the bar open to expose the tops of the PCIe slots.

- c) Open the card's blue ejector lever to unseat it from PCIe slot 10.
- d) Pull both ends of the card vertically to disengage the card from the socket, and then set it aside.

Step 3 Install a new rear NVMe switch card:

- a) Carefully align the card edge with the socket of PCIe slot 10.
- b) Push on both corners of the card to seat its connector in the socket.
- c) Fully close the blue ejector lever on the card to lock the card into the socket.
- d) Connect the PCIe cable (CBL-AUX-NVME-M5) from the internal drive module backplane to the new switch card.
- e) Close the hinged retainer bar over the top of the PCIe slots.


Use your fingertips to pull back on the wire locking-latches at each end of the retainer bar, and then hinge it closed to lock in the tops of the PCIe slots. Push the wire locking-latching back to the forward, locked position.

Step 4 Replace the top cover to the server.

Step 5 Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).

Step 6 Fully power on the server by pressing the Power button.

Figure 36: Rear NVMe Switch Card Location (PCIe 10)



1	NVMe switch card in required location (PCIe 10)	-
---	---	---

Replacing Fan Modules

The four hot-swappable fan modules in the server are numbered as shown in [Figure 3: Serviceable Component Locations Inside the Main Chassis, on page 5](#). Each fan module contains two fans.



Tip

There is a fault LED on the top of each fan module. This LED lights green when the module is correctly seated and is operating OK. The LED lights amber when the module 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:

- 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 41](#).
- 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.

Figure 37: Top View of Fan Module



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

Replacing an Internal USB Drive

This section includes procedures for installing a USB drive and for enabling or disabling the internal USB port.

Replacing a USB Drive

The server has one horizontal USB 2.0 socket on the motherboard.



Caution We do not recommend that you hot-swap the internal USB drive while the server is powered on because of the potential for data loss.

Step 1 Remove an existing internal USB drive:

- Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 39](#).
- 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 41](#).
- Locate the USB socket on the motherboard as shown below, near PCIe slot 12.
- Grasp the USB drive and pull it horizontally to free it from the socket.


Step 2 Install a new internal USB drive:

- Align the USB drive with the socket.
- Push the USB drive horizontally to fully engage it with the socket.
- Replace the top cover to the server.

Step 3 Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).

Step 4 Fully power on the server by pressing the Power button.

Figure 38: Internal USB 2.0 Socket Location



1	Location of horizontal USB socket on motherboard	-	
---	--	---	--

Enabling or Disabling the Internal USB Port

The factory default is that all USB ports on the server are enabled. However, the internal USB port can be enabled or disabled in the server BIOS.

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

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.

TPM Considerations

- This server supports either TPM version 1.2 or TPM version 2.0. The TPM 2.0, UCSX-TPM2-002B(=), is compliant with Federal Information Processing (FIPS) Standard 140-2. FIPS support has existed, but FIPS 140-2 is now supported.
- 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 and Enabling a TPM



Note 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.

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

1. Installing the TPM Hardware

2. Enabling the TPM in the BIOS
3. Enabling the Intel TXT Feature in the BIOS


Installing TPM Hardware



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

- Step 1** Prepare the 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 39](#).
 - 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 41](#).
- Step 2** Install a TPM:
- a) Locate the TPM socket on the motherboard, as shown below.
 - 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 3** Replace the cover to the server.
- Step 4** Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).
- Step 5** Fully power on the server by pressing the Power button.
- Step 6** Continue with [Enabling the TPM in the BIOS, on page 90](#).

Figure 39: TPM Socket Location



1	TPM socket location on motherboard	-	
---	------------------------------------	---	--

Enabling the TPM in the BIOS

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



Note

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

Step 1 Enable TPM Support:

- Watch during bootup for the F2 prompt, and then press **F2** to enter BIOS setup.
- Log in to the BIOS Setup Utility with your BIOS Administrator password.
- On the BIOS Setup Utility window, choose the **Advanced** tab.
- Choose **Trusted Computing** to open the TPM Security Device Configuration window.
- Change TPM SUPPORT to **Enabled**.
- Press **F10** to save your settings and reboot the server.

Step 2 Verify that TPM support is now enabled:

- 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 91](#).

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:

- a) Choose the **Advanced** tab.
- b) Choose **Intel TXT(LT-SX) Configuration** to open the Intel TXT(LT-SX) Hardware Support window.
- c) Verify that the following items are listed as Enabled:
 - VT-d Support (default is Enabled)
 - VT Support (default is Enabled)
 - TPM Support
 - TPM State
- d) Do one of the following:
 - If VT-d Support and VT Support are already enabled, skip to step 4.
 - If VT-d Support and VT Support are not enabled, continue with the next steps to enable them.
- e) Press **Escape** to return to the BIOS Setup utility **Advanced** tab.
- f) On the Advanced tab, choose **Processor Configuration** to open the Processor Configuration window.
- g) Set Intel (R) VT and Intel (R) VT-d to **Enabled**.

Step 4 Enable the Intel Trusted Execution Technology (TXT) feature:

- a) Return to the Intel TXT(LT-SX) Hardware Support window if you are not already there.
- b) Set TXT Support to **Enabled**.

Step 5 Press **F10** to save your changes and exit the BIOS Setup utility.

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:

- a) Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 39](#).
- 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 41](#).

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 Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).

Step 6 Fully power on the server by pressing the Power button.

Figure 40: Chassis Intrusion Switch



Replacing Power Supplies

The server requires four power supplies. When four power supplies are installed they are redundant as 2+2 by default. You can change this to 3+1 redundancy in the system BIOS.



Note The power supplies are hot-swappable and are accessible from the external rear of the server, so you do not have to pull the server out from the rack or remove the server cover.

- See also [Power Specifications, on page 152](#) for more information about the supported power supplies.
- See also [Rear-Panel LEDs, on page 36](#) for information about the power supply LEDs.

Replacing AC Power Supplies



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

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

- Remove the power cord from the power supply that you are replacing.


- b) Grasp the power supply handle while pinching the release latch toward the handle.
- c) 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.

Figure 41: AC Power Supplies



1	Power supply status LED	3	Power supply handle
2	Power supply release latch	-	

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 and Restrictions

The server provides 12 PCIe slots for vertical installation of up to 12 PCIe expansion cards.

The following figure shows a top view of the PCIe sockets and the corresponding PCIe slot openings in the rear panel. Some rear-panel openings are not used at this time.

Figure 42: PCIe Slot Numbering



PCIe Slot Specifications

Table 4: PCIe Slot Specifications

Slot Number	Electrical Lane Width	Connector Length	Maximum Card Length	Card Height (Rear Panel Opening)	NCSI Support	GPU Card Support	Cisco VIC Card Support
1	Gen-3 x16	x24 connector	Full length	Full height	Yes	Yes	Yes (primary slot)

2	Gen-3 x16	x24 connector	Full length	Full height	Yes	Yes	Yes (secondary slot)
3	Gen-3 x8	x24 connector	Full length	Full height	Yes	No	Yes
4	Gen-3 x16	x24 connector	Full length	Full height	Yes	Yes	Yes
5	Gen-3 x8	x24 connector	Full length	Full height	Yes	No	Yes
6	Gen-3 x16	x24 connector	Full length	Full height	Yes	Yes	Yes
7	Gen-3 x8	x24 connector	Full length	Full height	Yes	No	Yes
8	Gen-3 x16	x24 connector	Full length	Full height	Yes	Yes	Yes
9	Gen-3 x8	x24 connector	Full length	Full height	No	No	No
10	Gen-3 x16	x24 connector	Full length	Full height	No	Yes	No
11	Gen-3 x8	x24 connector	Full length	Full height	No	No	No
12	Gen-3 x8	x8 connector	Full length	Full height	No	No	No

PCIe Population Guidelines and Restrictions

Note the following guidelines and restrictions:

- Control of the PCIe sockets is divided between the CPUs that are present in the system. Some PCIe slots are not available if your system does not have CPU module 2 installed:
 - If your system has four CPUs, all PCIe slots are supported.
 - If your system has only two CPUs (CPU module 2 is not present), see the following table for the PCIe slots that are supported.

PCIe Slots Controlled by CPU Module 1 (CPUs 1 and 2)	PCIe Slots Controlled by CPU Module 2 (CPUs 3 and 4)
1, 2, 5, 8, 9, 10	3, 4, 6, 7, 11, 12


- If the rear drive-bay module is installed, PCIe slot 12 is not available because of internal clearance.
- If the server has a rear RAID controller card, it must be installed in PCIe slot 11 or slot 10.
- If the server has a rear NVMe switch card, it must be installed in PCIe slot 10.

Replacing a PCIe Card

Before installing PCIe cards, see [PCIe Slot Specifications and Restrictions, on page 94](#).

-
- 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 39](#).
 - 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 41](#).
- Step 2** Remove any existing card or a blanking panel:
- Open the hinged retainer bar that covers the top of the PCIe slot.
- Use your fingertips to pull back on the wire locking-latches at each end of the retainer bar, and then hinge the bar open to expose the tops of the PCIe slots.
- Pull both ends of the card vertically to disengage the card from the socket, and then set it aside.
- Step 3** Install a new PCIe card:
- Carefully align the card edge with the socket while you align the card's rear tab with the rear panel opening.
 - Push down on both corners of the card to seat its edge connector in the socket.
 - Close the hinged retainer bar over the top of the PCIe slots.
- Use your fingertips to pull back on the wire locking-latches at each end of the retainer bar, and then hinge it closed to secure the tops of the PCIe slots. Push the wire locking-latches back to the forward, locked position.
- Step 4** Replace the top cover to the server.
- Step 5** Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).
- Step 6** Fully power on the server by pressing the Power button.

Figure 43: PCIe Slot Hinged Retainer Bars



1	Wire locking latches for left PCIe retainer bar (slots 10 - 12)	2	Wire locking latches for right PCIe retainer bar (slots 1 - 9)
---	---	---	--

Cisco Virtual Interface Card (VIC) Considerations

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

If you want to use the Cisco UCS VIC card for Cisco UCS Manager integration, see also the [Installation For Cisco UCS Manager Integration, on page 189](#) for details about supported configurations, cabling, and other requirements.

Table 5: VIC Support and Considerations in This Server

VIC	How Many Supported in Server	Slots That Support the VIC	Primary Slot For Cisco UCS Manager Integration	Primary Slot For Cisco Card NIC Mode	Minimum Cisco IMC Firmware
Cisco UCS VIC 1385 UCSC-PCIE-C40Q-03	8	PCIe 1 - 8	PCIe 1	PCIe 1	3.1(2)
Cisco UCS VIC 1455 UCSC-PCIE-C25Q-04	8	PCIe 1 - 8	PCIe 1	PCIe 1	4.0(1)
Cisco UCS VIC 1495 UCSC-PCIE-C100-04	8	PCIe 1 - 8	PCIe 1	PCIe 1	4.0(2)

- The primary slot for a VIC card is slot 1; the secondary slot for a VIC card is slot 2.
- The system can support up to two VIC cards total in UCSM mode. Only the VIC card installed in slot 1 can be used for both UCS Manager management and data traffic. A second VIC installed in slots 2 - 8 is used for data traffic only.
- The VICs are supported in slots 1 - 8. Of these 8 slots, CPU module 1 (CPU 1 and 2) supports slots 1, 2, 5, 8; CPU module 2 (CPU 3 and 4) supports slots 3, 4, 6, 7.

Replacing Components Inside a CPU Module



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.

This section describes how to install and replace CPUs and DIMMs inside a CPU module.



Caution Never remove a CPU module without shutting down and removing power from the server.

See also:

- [Replacing Components Inside the Main Chassis, on page 49](#)
- [Replacing Components Inside an I/O Module, on page 122](#)

Replacing CPUs and Heatsinks

This section contains information for replacing CPUs and heatsinks inside a CPU module.

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



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

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

Table 6: Minimum Requirements For Second Generation Intel Xeon Scalable processors


Software or Firmware	Minimum Version
Server Cisco IMC	4.0(4)
Server BIOS	4.0(4)

CPU Configuration Rules

The CPUs in this server install to sockets inside one or two removable CPU modules. Each CPU module has two CPU sockets.

- The system numbers the CPUs in CPU module 1 (the lower bay) as CPU 1 and CPU 2.
- The system numbers the CPUs in CPU module 2 (the upper bay) as CPU 3 and CPU 4.

Figure 44: CPU Numbering



- The server can operate with one or two CPU modules (two or four identical CPUs) installed.



Note The CPUs in CPU module 1 must be identical with the CPUs in CPU module 2 (no mixing).

- The minimum configuration is that the server must have at least CPU module 1 installed in the lower CPU module bay. Install CPU module 1 first, and then CPU module 2 in the upper bay.



Note If CPU module 2 is not present in the upper bay, you must have a blank filler module in the upper bay or the server will not boot.

- **For Intel Xeon Scalable processors (first generation):** The maximum combined memory allowed in the 12 DIMM slots controlled by any one CPU is 768 GB. To populate the 12 DIMM slots with more than 768 GB of combined memory, you must use a high-memory CPU that has a PID that ends with an "M", for example, UCS-CPU-6134M.
- **For Second Generation Intel Xeon Scalable processors:** These Second Generation CPUs have three memory tiers. These rules apply on a *per-socket* basis:
 - If the CPU socket has up to 1 TB of memory installed, a CPU with no suffix can be used (for example, Gold 6240).
 - If the CPU socket has 1 TB or more (up to 2 TB) of memory installed, you must use a CPU with an M suffix (for example, Platinum 8276M).
 - If the CPU socket has 2 TB or more (up to 4.5 TB) of memory installed, you must use a CPU with an L suffix (for example, Platinum 8270L).
- The following restrictions apply when using only a two-CPU configuration (CPU module 2 is not present):

- The maximum number of DIMMs is 24 (only CPU 1 and CPU 2 memory channels).
- Some PCIe slots are unavailable when CPU module 2 is not present:

PCIe Slots Controlled by CPU Module 1 (CPUs 1 and 2)	PCIe Slots Controlled by CPU Module 2 (CPUs 3 and 4)
1, 2, 5, 8, 9, 10	3, 4, 6, 7, 11, 12

- Only four double-wide GPUs are supported, in PCIe slots 1, 2, 8, and 10.
- No front NVMe drives are supported.
- The optional NVMe-only drive bay module UCSC-C480-8NVME is not supported.
- If a rear RAID controller is used, it must be installed in PCIe slot 10 rather than the default slot 11. A blank filler must be installed in slot 11.
- The following NVIDIA GPUs are not supported with Second Generation Intel Xeon Scalable processors:
 - NVIDIA Tesla P100 12G
 - NVIDIA Tesla P100 16G

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 108](#).

Replacing a CPU and Heatsink



Caution

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

Step 1

Caution Never remove a CPU module without shutting down and removing power from the server.

Prepare the server for component removal:

- a) Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 39](#).

Note You do not have to pull the server out of the rack or remove the server cover because the CPU modules are accessible from the front of the server.

Step 2

Remove an existing CPU module from the chassis:

Note Verify that the power LED on the front of the CPU module is off before removing the module.

- a) Grasp the two ejector levers on the front of the CPU module and pinch their latches to release the levers.
- b) Rotate both levers to the outside at the same time to evenly disengage the module from the midplane connectors.
- c) Pull the module straight out from the chassis and then set it on an antistatic surface.

Step 3


Remove the existing CPU/heatsink assembly from the CPU module:

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

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

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

Figure 45: Removing the CPU/Heatsink Assembly



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


Step 4

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

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

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

Figure 46: Separating the CPU Assembly From the Heatsink



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

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

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

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

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

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

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


Step 5

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

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

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

Figure 47: CPU Assembly Tool, CPU Assembly, and Heatsink Alignment Features



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


Step 6

Apply new TIM to the heatsink:

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

- If you are installing a new heatsink, it is shipped with a pre-applied pad of TIM. Go to step 5.
 - If you are reusing a heatsink, you must remove the old TIM from the heatsink and then apply new TIM to the CPU surface from the supplied syringe. Continue with step **a** below.
- a) Apply the cleaning solution that is included with the heatsink cleaning kit (UCSX-HSCK=) to the old TIM on the heatsink and let it soak for a least 15 seconds.
 - b) Wipe all of the TIM off the heatsink using the soft cloth that is included with the heatsink cleaning kit. Be careful to avoid scratching the heatsink surface.
 - c) Using the syringe of TIM provided with the new CPU (UCS-CPU-TIM=), apply 1.5 cubic centimeters (1.5ml) of thermal interface material to the top of the CPU. Use the pattern shown below to ensure even coverage.

Figure 48: Thermal Interface Material Application Pattern



Step 7

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

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


Step 8

Install the CPU/heatsink assembly to the server:

- a) Lift the heatsink with attached CPU assembly from the CPU assembly tool.
- b) Align the assembly over the CPU socket on the board, as shown in the following figure.

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

Figure 49: Installing the Heatsink/CPU Assembly to the CPU Socket



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

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

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

Step 9

Return the CPU module to the chassis:

- a) With the two ejector levers open, align the CPU module with an empty bay.
- b) Push the module into the bay until it engages with the midplane connectors and is flush with the chassis front.
- c) Rotate both ejector levers toward the center until they lay flat and their latches lock into the front of the module.

Step 10 Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).

Step 11 Fully power on the server by pressing the Power button.

Note Verify that the power LED on the front of the CPU module returns to solid green.

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

If you are moving existing CPUs to a new CPU module, it is not necessary to separate the CPU and heatsink. They can be moved as one assembly. See [Additional CPU-Related Parts to Order with RMA Replacement CPU Modules, on page 109](#).

- 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 M5 servers (UCS-CPU-TIM=)

One TIM kit covers one CPU.
- Scenario 2—You are replacing the existing heatsinks:
 - Heat sink (UCSC-HS-02-EX=)

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: UCS-M5-CPU-CAR=
 - #1 flat-head screwdriver (for separating the CPU from the heatsink)
 - Heatsink cleaning kit (UCSX-HSCK=)

One cleaning kit can clean up to four CPUs.
 - Thermal interface material (TIM) kit for M5 servers (UCS-CPU-TIM=)

One TIM kit covers one CPU.

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

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

Additional CPU-Related Parts to Order with RMA Replacement CPU Modules

When a return material authorization (RMA) of the CPU module is done on a C480 M5 CPU module, you move existing CPUs to the new CPU module.



Note Unlike previous generation CPUs, the M5 server CPUs do not require you to separate the heatsink from the CPU when you *move* the CPU-heatsink assembly. Therefore, no additional heatsink cleaning kit or thermal-interface material items are required.

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

To move a CPU to a new CPU module, use the procedure in [Moving an M5 Generation CPU](#), on page 109.

Moving an M5 Generation CPU

Tool required for this procedure: T-30 Torx driver



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

Step 1


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

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

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

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

Figure 50: Removing the CPU/Heatsink Assembly



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


Step 2

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

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

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

Figure 51: Removing a CPU Socket Dust Cover



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


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

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

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

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

Figure 52: Installing a CPU Socket Dust Cover



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


Step 3

Install the CPUs to the new system:

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

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

Figure 53: Installing the Heatsink/CPU Assembly to the CPU Socket



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

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

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

Replacing Memory DIMMs


Caution

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


Caution

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


Note

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

DIMM Population Rules and Memory Performance Guidelines

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



Note

You must use DIMM blanking panels in any DIMM slots that do not have DIMMs installed to ensure adequate air flow.

DIMM Slot Numbering

The following figure shows the numbering of the DIMM slots on the CPU module board. When a CPU module is in bay 1 (the lower bay), the system numbers the CPUs as CPU 1 and CPU 2. When a CPU module is in bay 2 (the upper bay), the system numbers the CPUs as CPU 3 and CPU 4.

Figure 54: DIMM Slot Numbering



DIMM Population Rules

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

- Each CPU supports six memory channels.
 - CPU 1/3 supports channels A, B, C, D, E, F.
 - CPU 2/4 supports channels G, H, J, K, L, M.

- Each channel has two DIMM sockets (for example, channel A = slots A1, A2).
- For optimal performance, populate DIMMs in the order shown in the following table, depending on the number of DIMMs per CPU. Balance DIMMs evenly across the two CPUs as shown in the table.



Note The table below lists recommended configurations. Using 5, 6, 7, 9, 10, or 11 DIMMs per CPU is not recommended.



Note The CPU numbering in the lower CPU module 1 is CPU 1 and CPU 2; in the upper CPU module 2, the system numbers the CPUs as CPU 3 and CPU 4. The channel lettering is the same in both CPU modules. Balance the DIMMs evenly across all four CPUs, if present.

Table 7: DIMM Population Order

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

- The maximum combined memory allowed in the 12 DIMM slots controlled by any one CPU is 768 GB. To populate the 12 DIMM slots with more than 768 GB of combined memory, you must use a high-memory CPU that has a PID that ends with an "M", for example, UCS-CPU-6134M.
- All DIMMs must be DDR4 DIMMs that support ECC. Non-buffered UDIMMs and non-ECC DIMMs are not supported.
- Memory mirroring reduces the amount of memory available by 50 percent because only one of the two populated channels provides data. When memory mirroring is enabled, you must install DIMMs in even numbers of channels.
- NVIDIA M-Series GPUs can support only less-than 1 TB memory in the server.
- NVIDIA P-Series GPUs can support 1 TB or more memory in the server.
- AMD FirePro S7150 X2 GPUs can support only less-than 1 TB memory in the server.
- Observe the DIMM mixing rules shown in the following table.

Table 8: DIMM Mixing Rules

DIMM Parameter	DIMMs in the Same Channel	DIMMs in the Same Bank
DIMM Capacity For example, 16GB, 32GB, 64GB, 128GB	You can mix different capacity DIMMs in the same channel (for example, A1, A2).	You cannot mix DIMMs with different capacities and Revisions in the same bank (for example A1, B1). The Revision value depends on the manufactures. Two DIMMs with the same PID can have different Revisions.
DIMM speed For example, 2666 GHz	You can mix speeds, but DIMMs will run at the speed of the slowest DIMMs/CPUs installed in the channel.	You cannot mix DIMMs with different speeds and Revisions in the same bank (for example A1, B1). The Revision value depends on the manufactures. Two DIMMs with the same PID can have different Revisions.
DIMM type RDIMMs or LRDIMMs	You cannot mix DIMM types in a channel.	You cannot mix DIMM types in a bank.

Memory Mirroring

The Intel CPUs within 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 37](#) for the locations of these LEDs.

Step 1 **Caution** Never remove a CPU module without shutting down and removing power from the server.

Prepare the server for component removal:

- a) Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 39](#).

Note You do not have to pull the server out of the rack or remove the server cover because the CPU modules are accessible from the front of the server.

Step 2 Remove an existing CPU module from the chassis:

Note Verify that the power LED on the front of the CPU module is off before removing the module.

- a) Grasp the two ejector levers on the front of the CPU module and pinch their latches to release the levers.

- b) Rotate both levers to the outside at the same time to evenly disengage the module from the midplane connectors.
- c) Pull the module straight out from the chassis and then set it on an antistatic surface.

Step 3 Remove an existing DIMM (or DIMM blank) from the CPU module:

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

Step 4 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 114](#).

Note You must use DIMM blanking panels in any DIMM slots that do not have DIMMs installed to ensure adequate air flow.

- a) Align the new DIMM with the empty slot on the CPU module board. 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.

Step 5 Return the CPU module to the chassis:

- a) With the two ejector levers open, align the CPU module with an empty bay.
- b) Push the module into the bay until it engages with the midplane connectors and is flush with the chassis front.
- c) Rotate both ejector levers toward the center until they lay flat and their latches lock into the front of the module.

Step 6 Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).

Step 7 Fully power on the server by pressing the Power button.

Note Verify that the power LED on the front of the CPU module returns to solid green.

Replacing Intel Optane DC Persistent Memory Modules

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



Caution

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



Note

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



Note

Intel Optane DC persistent memory modules require Second Generation Intel Xeon Scalable processors. You must upgrade the server firmware and BIOS to version 4.0(4) or later and install the supported Second Generation Intel Xeon Scalable processors before installing DCPMMs.

DCPMMs can be configured to operate in one of three modes:

- **Memory Mode (default):** The module operates as 100% memory module. Data is volatile and DRAM acts as a cache for DCPMMs. This is the factory default mode.
- **App Direct Mode:** The module operates as a solid-state disk storage device. Data is saved and is non-volatile.
- **Mixed Mode (25% Memory Mode + 75% App Direct):** The module operates with 25% capacity used as volatile memory and 75% capacity used as non-volatile storage.


Intel Optane DC Persistent Memory Module Population Rules and Performance Guidelines

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

DIMM Slot Numbering

The following figure shows the numbering of the DIMM slots on the CPU module board. When a CPU module is in bay 1 (the lower bay), the system numbers the CPUs as CPU 1 and CPU 2. When a CPU module is in bay 2 (the upper bay), the system numbers the CPUs as CPU 3 and CPU 4.

Figure 55: DIMM Slot Numbering



Configuration Rules

Observe the following rules and guidelines:

- To use DCPMMs in this server, four CPUs must be installed.
- Intel Optane DC persistent memory modules require Second Generation Intel Xeon Scalable processors. You must upgrade the server firmware and BIOS to version 4.0(4) or later and then install the supported Second Generation Intel Xeon Scalable processors before installing DCPMMs.
- When using DCPMMs in a server:
 - The DDR4 DIMMs installed in the server must all be the same size.
 - The DCPMMs installed in the server must all be the same size and must have the same SKU.
- The DCPMMs run at 2666 MHz. If you have 2933 MHz RDIMMs or LRDIMMs in the server and you add DCPMMs, the main memory speed clocks down to 2666 MHz to match the speed of the DCPMMs.
- Each DCPMM draws 18 W sustained, with a 20 W peak.

- The following table shows supported DCPMM configurations for this server. Fill the DIMM slots for CPU 1 and CPU 2 in CPU module 1 as shown, depending on which DCPMM:DRAM ratio you want to populate. If CPU module 2 is present, fill the DIMM slots for CPU 3 and CPU 4 as shown.

Figure 56: Supported DCPMM Configurations for Quad-CPU Configurations

DIMM to DCPMM Count	CPU 1 (lower server node)											
	IMC1						IMC0					
	Channel 2		Channel 1		Channel 0		Channel 2		Channel 1		Channel 0	
	F2	F1	E2	E1	D2	D1	C2	C1	B2	B1	A2	A1
6 to 2		DIMM		DIMM	DCPMM	DIMM		DIMM		DIMM	DCPMM	DIMM
6 to 4		DIMM	DCPMM	DIMM	DCPMM	DIMM		DIMM	DCPMM	DIMM	DCPMM	DIMM
6 to 6	DCPMM	DIMM	DCPMM	DIMM	DCPMM	DIMM	DCPMM	DIMM	DCPMM	DIMM	DCPMM	DIMM

DIMM to DCPMM Count	CPU 2 (lower server node)											
	IMC1						IMC0					
	Channel 2		Channel 1		Channel 0		Channel 2		Channel 1		Channel 0	
	M2	M1	L2	L1	K2	K1	J2	J1	H2	H1	G2	G1
6 to 2		DIMM		DIMM	DCPMM	DIMM		DIMM		DIMM	DCPMM	DIMM
6 to 4		DIMM	DCPMM	DIMM	DCPMM	DIMM		DIMM	DCPMM	DIMM	DCPMM	DIMM
6 to 6	DCPMM	DIMM	DCPMM	DIMM	DCPMM	DIMM	DCPMM	DIMM	DCPMM	DIMM	DCPMM	DIMM

DIMM to DCPMM Count	CPU 3 (upper server node)											
	IMC1						IMC0					
	Channel 2		Channel 1		Channel 0		Channel 2		Channel 1		Channel 0	
	F2	F1	E2	E1	D2	D1	C2	C1	B2	B1	A2	A1
6 to 2		DIMM		DIMM	DCPMM	DIMM		DIMM		DIMM	DCPMM	DIMM
6 to 4		DIMM	DCPMM	DIMM	DCPMM	DIMM		DIMM	DCPMM	DIMM	DCPMM	DIMM
6 to 6	DCPMM	DIMM	DCPMM	DIMM	DCPMM	DIMM	DCPMM	DIMM	DCPMM	DIMM	DCPMM	DIMM

DIMM to DCPMM Count	CPU 4 (upper server node)											
	IMC1						IMC0					
	Channel 2		Channel 1		Channel 0		Channel 2		Channel 1		Channel 0	
	M2	M1	L2	L1	K2	K1	J2	J1	H2	H1	G2	G1
6 to 2		DIMM		DIMM	DCPMM	DIMM		DIMM		DIMM	DCPMM	DIMM
6 to 4		DIMM	DCPMM	DIMM	DCPMM	DIMM		DIMM	DCPMM	DIMM	DCPMM	DIMM
6 to 6	DCPMM	DIMM	DCPMM	DIMM	DCPMM	DIMM	DCPMM	DIMM	DCPMM	DIMM	DCPMM	DIMM

453/90

Installing Intel Optane DC Persistent Memory Modules



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

Understand which mode your DCPMM is operating in. App Direct mode has some additional considerations in this procedure.



Caution Replacing a DCPMM in App-Direct mode requires all data to be wiped from the DCPMM. Make sure to backup or offload data before attempting this procedure.

Step 1 **Caution** Never remove a CPU module without shutting down and removing power from the server.

Prepare the server for component removal:

- a) Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 39](#).

Note You do not have to pull the server out of the rack or remove the server cover because the CPU modules are accessible from the front of the server.

Step 2 Remove an existing CPU module from the chassis:

Note Verify that the power LED on the front of the CPU module is off before removing the module.

- a) Grasp the two ejector levers on the front of the CPU module and pinch their latches to release the levers.
- b) Rotate both levers to the outside at the same time to evenly disengage the module from the midplane connectors.
- c) Pull the module straight out from the chassis and then set it on an antistatic surface.

Step 3 For App Direct mode, backup the existing data stored in all Optane DIMMs to some other storage.

Step 4 For App Direct mode, remove the Persistent Memory policy which will remove goals and namespaces automatically from all Optane DIMMs.

Step 5 Remove an existing DCPMM:

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

- a) Locate the DCPMM that you are removing, and then open the ejector levers at each end of its DIMM slot.
- b) Lift straight up on the DCPMM and set it aside.

Step 6 Install a new DCPMM:

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

- a) Align the new DCPMM with the empty slot on the motherboard. Use the alignment feature in the DIMM slot to correctly orient the DCPMM.
- b) Push down evenly on the top corners of the DCPMM until it is fully seated and the ejector levers on both ends lock into place.

Step 7 Return the CPU module to the chassis:

- a) With the two ejector levers open, align the CPU module with an empty bay.
- b) Push the module into the bay until it engages with the midplane connectors and is flush with the chassis front.
- c) Rotate both ejector levers toward the center until they lay flat and their latches lock into the front of the module.

Step 8 Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).

Step 9 Fully power on the server by pressing the Power button.

Note Verify that the power LED on the front of the CPU module returns to solid green.

Step 10 Perform post-installation actions:

Note If your Persistent Memory policy is Host Controlled, you must perform the following actions from the OS side.

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

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

- To use the server's BIOS Setup Utility, see [Server BIOS Setup Utility Menu for DCPMM, on page 121](#).
- To use Cisco IMC or Cisco UCS Manager, see the [Cisco UCS: Configuring and Managing Intel Optane DC Persistent Memory Modules](#) guide.

Server BIOS Setup Utility Menu for DCPMM



Caution

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

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

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

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

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

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

Advanced > Intel Optane DC Persistent Memory Configuration

From this tab, you can access other menu items:

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

- Monitor health
- Update firmware
- Configure security

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

- Configure data policy

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

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

- Create goal config

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

Existing namespace attributes such as the size cannot be modified. You can only add or delete namespaces.

- **Total capacity:** Displays the total resource allocation across the server.

Updating the DCPMM Firmware Using the BIOS Setup Utility

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

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

Replacing Components Inside an I/O Module

**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.

**Caution**

Never remove an I/O module without shutting down and removing power from the server.

This section describes how to install and replace I/O module components.

**Note**

The I/O module is not field replaceable, nor can you move an I/O module from one chassis to another. This module contains a security chip that requires it to stay with the PCIe module in the same chassis, as shipped from the factory.

See also:

- [Replacing Components Inside the Main Chassis, on page 49](#)
- [Replacing Components Inside a CPU Module, 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.

**Caution**

Removing the RTC battery impacts the following:

- Real clock time gets reset to default value.
- CMOS setting of the server is lost. You should reset the system setting after replacing the RTC battery.

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.

**Caution**

Never remove an I/O module without shutting down and removing power from the server.

Step 1

Prepare the server for component removal:

- a) Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 39](#).

Note You do not have to pull the server out of the rack or remove the server cover because the I/O module is accessible from the rear of the server.

Step 2 Remove an I/O module from the chassis:

- a) Disconnect any cables from the ports on the I/O module.
- b) Push down on the locking clip on the I/O module's ejector-handle, and then hinge the handle upward to disengage the module's connector from the chassis midplane.
- c) Pull the module straight out from the chassis and then set it on an antistatic surface.

Step 3 Remove the RTC battery:


- a) Locate the vertical RTC battery socket on the I/O module board.
- b) Remove the battery from the socket. Gently pry the securing clip to the side to provide clearance, then lift up on the battery.

Step 4 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 flat, positive side of the battery marked “3V+” should face the clip on the socket (toward the module rear).

Figure 57: RTC Battery Socket Location Inside an I/O Module



1	RTC battery in vertical socket	-	
---	--------------------------------	---	--

Step 5 Return the I/O module to the chassis:

- a) With the ejector-handle open, align the I/O module with the empty bay.
- b) Push the module into the bay until it engages with the midplane connector.
- c) Hinge the ejector-handle down until it sits flat and its locking clip clicks. The module face must be flush with the rear panel of the chassis.
- d) Reconnect cables to the ports on the I/O module.

Step 6 Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).

Step 7 Fully power on the server by pressing the Power button.

Replacing a Micro SD Card

There is one socket for a Micro SD card on the I/O module board.



Caution Never remove a CPU module without shutting down and removing power from the server.

Step 1 Prepare the server for component removal:

- a) Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 39](#).

Note You do not have to pull the server out of the rack or remove the server cover because the I/O module is accessible from the rear of the server.

Step 2 Remove an I/O module from the chassis:

- a) Disconnect any cables from the ports on the I/O module.
- b) Push down on the locking clip on the I/O module's ejector-handle, and then hinge the handle upward to disengage the module's connector from the chassis midplane.
- c) Pull the module straight out from the chassis and then set it on an antistatic surface.


Step 3 Remove an existing Micro SD card:

- a) Locate the Micro SD card.
- b) Push horizontally on the Micro SD card and release it to make it spring out from the socket.
- c) Grasp the Micro SD card and lift it from the socket.

Step 4 Install a new Micro SD card:

- a) Align the new Micro SD card with the socket.
- b) Gently push down on the card until it clicks and locks in place in the socket.

Figure 58: Micro SD Card Location Inside an I/O Module



1	Location of Micro SD card socket on the I/O module board	-	
---	--	---	--

Step 5 Return the I/O module to the chassis:

- a) With the ejector-handle open, align the I/O module with the empty bay.
- b) Push the module into the bay until it engages with the midplane connector.

- c) Hinge the ejector-handle down until it sits flat and its locking clip clicks. The module face must be flush with the rear panel of the chassis.
- d) Reconnect cables to the ports on the I/O module.

Step 6 Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).

Step 7 Fully power on the server by pressing the Power button.

Replacing a Mini-Storage Module

The mini-storage module plugs into an I/O module board socket to provide additional internal storage. The mini-storage module is available in two different versions:

- SD card carrier—provides two SD card sockets.
- M.2 SSD Carrier—provides two M.2 form-factor SSD sockets.



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-MSTOR-M2). The M.2 drives are not listed in Cisco IMC inventory, nor can they be managed by Cisco IMC. This is expected behavior.

Replacing a Mini-Storage Module Carrier

This topic describes how to remove and replace a mini-storage module carrier. The carrier has one media socket on its top and one socket on its underside. Use the following procedure for any type of mini-storage module carrier (SD card or M.2 SSD).



Caution

Never remove an I/O module without shutting down and removing power from the server.

Step 1 Prepare the server for component removal:

- a) Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 39](#).

Note

You do not have to pull the server out of the rack or remove the server cover because the I/O module is accessible from the rear of the server.

Step 2 Remove an I/O module from the chassis:

- a) Disconnect any cables from the ports on the I/O module.
- b) Push down on the locking clip on the I/O module's ejector-handle, and then hinge the handle upward to disengage the module's connector from the chassis midplane.
- c) Pull the module straight out from the chassis and then set it on an antistatic surface.

Step 3 Remove a carrier from its socket:

- a) Locate the mini-storage module 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 4

Install a new carrier to its socket:

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

Figure 59: Mini-Storage Module Location on I/O Module Board



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

Step 5

Return the I/O module to the chassis:

- a) With the ejector-handle open, align the I/O module with the empty bay.
- b) Push the module into the bay until it engages with the midplane connector.
- c) Hinge the ejector-handle down until it sits flat and its locking clip clicks. The module face must be flush with the rear panel of the chassis.
- d) Reconnect cables to the ports on the I/O module.

Step 6

Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).

Step 7

Fully power on the server by pressing the Power button.

Replacing an SD Card in a Mini-Storage Carrier For SD

This topic describes how to remove and replace an SD card in a mini-storage carrier for SD (UCS-MSTOR-SD). The carrier has one SD card socket on its top and one socket on its underside.

Population Rules For Mini-Storage SD Cards

- You can use one or two SD cards in the carrier.
- Dual SD cards can be configured in a RAID 1 array through the Cisco IMC interface.
- SD socket 1 is on the top side of the carrier; SD socket 2 is on the underside of the carrier (the same side as the carrier's motherboard connector).



Caution

Never remove an I/O module without shutting down and removing power from the server.

- Step 1** Power off the server and then remove the mini-storage module carrier from the I/O module as described in [Replacing a Mini-Storage Module Carrier, on page 126](#).
- Step 2** Remove an SD card:
- a) Push on the top of the SD card, and then release it to allow it to spring out from the socket.
 - b) Grasp and remove the SD card from the socket.
- Step 3** Install a new SD card:
- a) Insert the new SD card into the socket with its label side facing up (away from the carrier).
 - b) Press on the top of the SD card until it clicks in the socket and stays in place.
- Step 4** Install the mini-storage module carrier back into the I/O module as described in [Replacing a Mini-Storage Module Carrier, on page 126](#).

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 SSD in a mini-storage carrier for M.2 (UCS-MSTOR-M2). The carrier has one M.2 SSD socket on its top and one socket on its underside.

Population Rules For Mini-Storage M.2 SSDs

-
- You can use one or two M.2 SSDs in the carrier.
- M.2 slot 1 is on the top side of the carrier; M.2 slot 2 is on the underside of the carrier (the same side as the carrier's motherboard connector).



Note

If you use the server's embedded software RAID controller with M.2 SATA SSDs, note that the numbering of the slots in the software interfaces is different than the physical slot numbering. Physical slot 1 is seen as slot 0 in the software; physical slot 2 is seen as slot 2 in the software.

- Dual SATA M.2 SSDs can be configured in a RAID 1 array through the BIOS Setup Utility's embedded SATA RAID interface. See [Embedded SATA RAID Controller, on page 160](#).



Note You cannot control the M.2 SATA SSDs in the server with a HW RAID controller.



Note The embedded SATA RAID controller requires that the server is set to boot in UEFI mode rather than Legacy mode.



Caution Never remove an I/O module without shutting down and removing power from the server.

-
- 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 126](#).
- Step 2** Remove an M.2 SSD:
- Use a #1 Phillips-head screwdriver to remove the single screw that secures the M.2 SSD to the carrier.
 - Grasp the M.2 SSD and lift up on the end that is opposite its socket on the carrier.
 - Remove the M.2 SSD from its socket on the carrier.
- Step 3** Install a new M.2 SSD:
- Angle downward and insert the new M.2 SSD connector-end into the socket on the carrier with its label side facing up.
 - Press the M.2 SSD flat against the carrier.
 - 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 126](#).
-

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 I/O module board. 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 4.0(4a) and later
- BIOS HII utility, BIOS 4.0(4a) and later
- Cisco UCS Manager 4.0(4a) 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 Prepare the server for component removal:

- a) Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 39](#).

Note You do not have to pull the server out of the rack or remove the server cover because the I/O module is accessible from the rear of the server.


Step 2 Remove an I/O module from the chassis:

- a) Disconnect any cables from the ports on the I/O module.
- b) Push down on the locking clip on the I/O module's ejector-handle, and then hinge the handle upward to disengage the module's connector from the chassis midplane.
- c) Pull the module straight out from the chassis and then set it on an antistatic surface.

Step 3 Remove a controller from its socket:

- a) At each end of the controller board, push outward on the clip that secures the carrier.
- b) Lift both ends of the controller to disengage it from the socket on the motherboard.
- c) Set the carrier on an anti-static surface.

Figure 60: Cisco Boot-Optimized M.2 RAID Controller on Motherboard




1	Location of socket on I/O module board	3	Securing clips
2	Alignment pegs	-	

Step 4 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 61: Cisco Boot-Optimized M.2 RAID Controller, Showing M.2 Drive Installation



Step 5 Install the controller to its socket on the motherboard:

- a) 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.
- b) Gently push down the socket end of the controller so that the two pegs go through the two holes on the controller.
- c) Push down on the controller so that the securing clips click over it at both ends.

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.

Recycling the PCB Assembly (PCBA)

The PCBA is secured to the server's sheet metal. To recycle the PCBA, you will need to remove a large assembly of components from the server, then breakdown the large assembly into its smaller sub assemblies and components. The assemblies and sub assemblies are secured to the chassis and held together by various screws:

- M3x0.6mm
- M3.5x0.6mm
- M4x0.7mm

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 41](#).

Step 1


On the exterior right and left side of the server's chassis, use a screwdriver to remove the mounting screws. The following image shows the locations of the mounting screws on each side of the chassis.

Figure 62: Location of Exterior Mounting Screws (Horizontal View)**Step 2**

Remove the top-level screws and fobs.

- Using a screwdriver, rotate each of the screws counter clockwise until it disengages.
- When all screws are removed, grasp the plastic fobs and remove them by hand.

The following image shows the locations of the screws and components.

Figure 63: Locations of Mounting Screws and Components (Top Down View)**Step 3**

Continue disassembly.


- a) Grasp the ribbon cable connector and disconnect it by hand.
- b) Using a screwdriver, remove the interior mounting screws.

Note Six screws to the right of the fan cage can be partially covered by the top sheet metal flange of the Midplane assembly. These screws are hard to locate and access, but they are accessible with a small angled screwdriver or similar tool.

- c) Detach the fan cage from the Midplane assembly.
- d) Using a screwdriver, rotate each of the screws for the latch bracket counter-clockwise until it disengages.

The following image shows the location of these screws.

Figure 64: Location of Interior Mounting Screws (Top Down View)




Step 4

Remove the I/O module.

- a) Lift the I/O Module latch.
- b) Slide the I/O Module out of the chassis.

The following image shows the location of this part.

Figure 65: Location of the I/O Module




Step 5

Remove the RAID card (if present).

- Disconnect the supercap cable and remove battery pack (if present) that is connected to it.
- Using a screwdriver, rotate each of the screws for the cable management bracket counter-clockwise until it disengages.
- Remove the cable bracket.


The following image shows the location of the bracket and its screws.

Figure 66: Location of Cable Management Bracket

- d) Pull the blue RAID card lever towards you to unseat the RAID card from its socket.
- e) Keeping the RAID card level, slide it toward you, then lift it out of the RAID card bracket.
- f) Using a screwdriver, rotate each of the RAID card bracket screws counter-clockwise until it disengages.
- g) Grasp the ends of the RAID card bracket and lift it straight up to disengage it from the metal pins that hold it in place.
- h) Using a screwdriver, rotate each of the screws in the black plastic supercap bracket counter-clockwise until it disengages.
- i) Remove the black plastic supercap bracket.

The following image shows the location of these screws and brackets.

Figure 67: Location of RAID Card Bracket and Supercap Bracket




Step 6

Remove the KVM card.

- Using a screwdriver, rotate the KVM card's security screw counter-clockwise until it disengages.
- Placing your fingers on the metal card guide near the socket connector, pull to disconnect the KVM card and slide it out of the chassis.


The following image shows the location of this component.

Figure 68: Location of KVM Card and Security Screw**Step 7**

Remove the mounting screws and Bridge card from the Midplane assembly:


- Using a screwdriver, rotate each of the screws counter clockwise until it disengages.
- Grasp the Bridge card (the vertical card) and remove it by hand.
- Grasp the Midplane stiffener and remove it by hand.

The following image shows the location of these screws and components.


Figure 69: Location of Mounting Screws, Bridge Card, and Midplane Stiffener**Step 8**

Using a screwdriver, continue disassembling the Bridge card by rotating each of its screws counter clockwise until it disengages.

The following image shows the location of the screws.

Figure 70: Location of Bridge Card Screws**Step 9**

Grasp the Midplane assembly handle and the Midplane frame and lift the entire midplane assembly out of the chassis. The following illustration shows where to grasp the Midplane Assembly.

Figure 71: Location of Hand Holds for Removing the Midplane Assembly (Horizontal View)**Step 10**


Remove the rear sub assembly.

- Using a screwdriver, rotate each of the screws counter clockwise until it disengages from the midplane frame.
- Grasp the rear sub assembly and disconnect it from the Midplane frame.
- Grasp the PCIE module and separate it from the Midplane frame.

The following illustration shows the location of the screws.

Note The following image is straight on showing the rear of the midplane assembly.

Figure 72: Location of Mounting Screws for Rear Sub Assembly




Step 11

Remove the PCBA, which includes additional components.

- Using a screwdriver, rotate each of the screws counter clockwise until it disengages, then detach the motherboard from the sheet metal tray.
- Remove the vertical metal PCBA handle.
- Remove the plastic baffle.

The following image shows these components.

Figure 73: Location of Mounting Screws, Baffle, and PCBA Bracket (Top Down View)



Step 12

Disassemble the power distribution board.


- a) Flip the PCBA over so that the component-side is facing down.

This step exposes the Power Distribution board and its mounting screws.

- b) Using a screwdriver, rotate each of the screws counter clockwise until it disengages.
- c) Detach the Power Distribution Board from the PCBA.

The following image shows the location of these components.

Figure 74: Underside of PCBA Showing Location of Mounting Screws for Power Distribution Board




Step 13 Properly dispose of the PCBA and all the components you disassembled.

Service DIP Switches

This server includes a block of DIP switches (SW1) that you can use for certain service and Cisco IMC debug functions. The block is located on the chassis motherboard, as shown in the following figure.

The switches in the following figure are shown in the default, open position (off).

Figure 75: Location of DIP Switches on Chassis Motherboard



1	Location of DIP switch block SW1	-	
---	----------------------------------	---	--

DIP Switch Function	Pin Numbers (Open - Closed)
Boot from alternate Cisco IMC image	8 - 9
Reset Cisco IMC to factory defaults	7 - 10
Reset Cisco IMC password to default	6 - 11
Clear CMOS	3 - 14
Recover BIOS	2 - 15
Password clear	1 - 16

Using the Clear Password Switch (Positions 1 - 16)

You can use this switch to clear the administrator password.

-
- Step 1** Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 39](#). 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 41](#).
- Step 4** Locate DIP switch block SW1 and the switch for pins 1 - 16 (see [Figure 75: Location of DIP Switches on Chassis Motherboard, on page 144](#)).
- Step 5** Move the DIP switch from position 1 to the closed, on position.
- Step 6** 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 7** 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 jumper cannot be determined without the host CPU running.
- Step 8** Press the Power button to shut down the server to standby power mode.
- Step 9** Remove AC power cords from the server to remove all power.
- Step 10** Remove the top cover from the server.
- Step 11** Move the DIP switch back to its default, off position.
- Note** If you do return the switch back to the default, open position, the password is cleared every time you power-cycle the server.
- Step 12** Replace the top cover to the server.
- Step 13** Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).
- Step 14** Fully power on the server by pressing the Power button.
-

Using the BIOS Recovery Switch (Positions 2 - 15)

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

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

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

- If it is a non-BootBlock corruption, a message similar to the following is displayed:

```
****BIOS FLASH IMAGE CORRUPTED****
Flash a valid BIOS capsule file using Cisco IMC WebGUI or CLI interface.
IF Cisco IMC INTERFACE IS NOT AVAILABLE, FOLLOW THE STEPS MENTIONED BELOW.
1. Connect the USB stick with bios.cap file in root folder.
2. Reset the host.
IF THESE STEPS DO NOT RECOVER THE BIOS
1. Power off the system.
2. Mount recovery jumper.
3. Connect the USB stick with bios.cap file in root folder.
```

4. Power on the system.
 Wait for a few seconds if already plugged in the USB stick.
 REFER TO SYSTEM MANUAL FOR ANY ISSUES.



Note As indicated by the message shown above, there are two procedures for recovering the BIOS. Try procedure 1 first. If that procedure does not recover the BIOS, use procedure 2.

Procedure 1: Reboot With recovery.cap File

- Step 1** Download the BIOS update package and extract it to a temporary location.
- Step 2** Copy the contents of the extracted recovery folder to the root directory of a USB drive. The recovery folder contains the bios.cap file that is required in this procedure.
- Note** The bios.cap file must be in the root directory of the USB drive. Do not rename this file. The USB drive must be formatted with either the FAT16 or FAT32 file system.
- Step 3** Insert the USB drive into a USB port on the server.
- Step 4** Reboot the server to standby power.
- The server boots with the updated BIOS boot block. When the BIOS detects a valid bios.cap file on the USB drive, it displays this message:
- ```
Found a valid recovery file...Transferring to Cisco IMC
System would flash the BIOS image now...
System would restart with recovered image after a few seconds...
```
- Step 5** Wait for server to complete the BIOS update, and then remove the USB drive from the server.
- Note** During the BIOS update, Cisco IMC shuts down the server and the screen goes blank for about 10 minutes. Do not unplug the power cords during this update. Cisco IMC powers on the server after the update is complete.

## Procedure 2: Use BIOS Recovery Switch and bios.cap File

- Step 1** Download the BIOS update package and extract it to a temporary location.
- Step 2** Copy the contents of the extracted recovery folder to the root directory of a USB drive. The recovery folder contains the bios.cap file that is required in this procedure.
- Note** The bios.cap file must be in the root directory of the USB drive. Do not rename this file. The USB drive must be formatted with either the FAT16 or FAT32 file system.
- Step 3** Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 39](#). Disconnect power cords from all power supplies.
- Step 4** Slide the server out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rear panel to provide clearance.
- Caution** If you cannot safely view and access the component, remove the server from the rack.
- Step 5** Remove the top cover from the server as described in [Removing the Server Top Cover, on page 41](#).

- Step 6** Locate DIP switch block SW1 and the switch for pins 2 - 15 (see [Figure 75: Location of DIP Switches on Chassis Motherboard, on page 144](#)).
- Step 7** Move the DIP switch from position 2 to the closed, on position.
- Step 8** Insert the USB thumb drive that you prepared in Step 2 into a USB port on the server.
- Step 9** Reconnect power cords to all power supplies and allow the server to boot to standby power.
- You do not have to return the server to main power for the change to take effect. Only Cisco IMC (the BMC) must reboot. The change takes effect after Cisco IMC finishes booting.
- Cisco IMC boots with the updated BIOS boot block. When the BIOS detects a valid bios.cap file on the USB drive, it displays this message:
- ```
Found a valid recovery file...Transferring to Cisco IMC
System would flash the BIOS image now...
System would restart with recovered image after a few seconds...
```
- Step 10** Wait for the BIOS update to complete, and then remove the USB drive from the server.
- Note** During the BIOS update, Cisco IMC shuts down the server and the screen goes blank for about 10 minutes. Do not unplug the power cords during this update. Cisco IMC powers on the server to standby power after the update is complete.
- Step 11** Remove all power cords again to fully remove power from the server.
- Step 12** Move the DIP switch back to its default, off position.
- Note** If you do not return the switch to the default open position, after recovery completion you see the prompt, "Please remove the recovery jumper."
- Step 13** Replace the top cover to the server.
- Step 14** Reconnect power cords to all power supplies and then allow the server to boot to standby power mode.
- Step 15** Fully power on the server to main power by pressing the Power button.

Using the Clear CMOS Switch (Positions 3 - 14)

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 jumper to invalidate the settings and reboot with defaults.



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

- Step 1** Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 39](#). 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 41](#).

- Step 4** Locate DIP switch block SW1 and the switch for pins 3 - 14 (see [Figure 75: Location of DIP Switches on Chassis Motherboard, on page 144](#)).
- Step 5** Move the DIP switch from position 3 to the closed, on position.
- Step 6** 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 7** 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 jumper cannot be determined without the host CPU running.
- Step 8** Press the Power button to shut down the server to standby power mode.
- Step 9** Remove AC power cords from the server to remove all power.
- Step 10** Remove the top cover from the server.
- Step 11** Move the DIP switch back to its default, off position.
- Note** If you do not return the switch to the default, open position, the CMOS settings are reset to the defaults every time you power-cycle the server.
- Step 12** Replace the top cover to the server.
- Step 13** Reconnect power cords to all power supplies and then allow the server to boot to standby power mode.
- Step 14** Fully power on the server to main power by pressing the Power button.

Using the Reset Cisco IMC Password to Default Switch (Positions 6 - 11)

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

- Step 1** Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 39](#). 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 41](#).
- Step 4** Locate DIP switch block SW1 and the switch for pins 6 - 11 (see [Figure 75: Location of DIP Switches on Chassis Motherboard, on page 144](#)).
- Step 5** Move the DIP switch from position 6 to the closed, on position.
- Step 6** 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.
- You do not have to return the server to main power for the change to take effect. Only Cisco IMC (the BMC) must reboot. The change takes effect after Cisco IMC finishes booting.
- Note** When you next log in to Cisco IMC, you see a message similar to the following:
- ```
'Reset to default CIMC password' debug functionality is enabled.
On input power cycle, CIMC password will be reset to defaults.
```

**Note** If you do not move the switch back to the default, open position, the server will reset the Cisco IMC password to the default every time that you power-cycle the server. The switch has no effect if you reboot Cisco IMC.

- Step 7** Remove AC power cords from the server to remove all power.
- Step 8** Remove the top cover from the server.
- Step 9** Move the DIP switch back to its default, off position.
- Step 10** Replace the top cover to the server.
- Step 11** Reconnect power cords to all power supplies and then allow the server to boot to standby power mode.
- Step 12** Fully power on the server by pressing the Power button.

## Using the Reset Cisco IMC to Defaults Switch (Positions 7 - 10)

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

- Step 1** Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 39](#). 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 41](#).
- Step 4** Locate DIP switch block SW1 and the switch for pins 7 - 10 (see [Figure 75: Location of DIP Switches on Chassis Motherboard, on page 144](#)).
- Step 5** Move the DIP switch from position 7 to the closed, on position.
- Step 6** 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.

You do not have to return the server to main power for the change to take effect. Only Cisco IMC (the BMC) must reboot. The change takes effect after Cisco IMC finishes booting.

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

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

**Note** If you do not move the switch back to the default, open position, the server will reset the Cisco IMC to the default settings every time that you power-cycle the server. The switch has no effect if you reboot Cisco IMC.

- Step 7** Remove AC power cords from the server to remove all power.
- Step 8** Remove the top cover from the server.
- Step 9** Move the DIP switch back to its default, off position.
- Step 10** Replace the top cover to the server.
- Step 11** Reconnect power cords to all power supplies and then allow the server to boot to standby power mode.
- Step 12** Fully power on the server by pressing the Power button.

## Using the Boot Alternate Cisco IMC Image Switch (Positions 8 - 9)

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

- 
- Step 1** Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 39](#). 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 41](#).
- Step 4** Locate DIP switch block SW1 and the switch for pins 8 - 9 (see [Figure 75: Location of DIP Switches on Chassis Motherboard, on page 144](#)).
- Step 5** Move the DIP switch from position 8 to the closed, on position.
- Step 6** 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.
- You do not have to return the server to main power for the change to take effect. Only Cisco IMC (the BMC) must reboot. The change takes effect after Cisco IMC finishes booting.
- 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 move the switch back to the default, open position, the server will boot from an alternate Cisco IMC image every time that you power cycle the server or reboot Cisco IMC.
- Step 7** Remove AC power cords from the server to remove all power.
- Step 8** Remove the top cover from the server.
- Step 9** Move the DIP switch back to its default, off position.
- Step 10** Replace the top cover to the server.
- Step 11** Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).
- Step 12** Fully power on the server by pressing the Power button.
-



APPENDIX **A**

Server Specifications

- [Server Specifications, on page 151](#)
- [Power Cord Specifications, on page 153](#)

Server Specifications

This appendix lists the physical, environmental, and power specifications for the server.

- [Physical Specifications, on page 151](#)
- [Environmental Specifications, on page 151](#)
- [Power Specifications, on page 152](#)

Physical Specifications

The following table lists the physical specifications for the server versions.

Table 9: Physical Specifications

Description	Specification
Height	7.0 in. (177.8 mm) 4 rack-unit (4RU)
Width	19.0 in. (482.6 mm)
Depth (length including front handles and power supplies)	32.7 in. (830.6 mm)
Maximum weight (fully loaded chassis)	146 lb. (66.2 Kg)

Environmental Specifications

The following table lists the environmental requirements and specifications for the server.

Table 10: Physical Specifications

Description	Specification
Temperature, Operating	41 to 95°F (5 to 35°C) Derate the maximum temperature by 1°C per every 305 meters of altitude above sea level. For more information, see the Cisco Unified Computing System Site Planning Guide: Data Center Power and Cooling .
Temperature, non-operating (when the server is stored or transported)	–40 to 149°F (–40 to 65°C)
Humidity (RH), operating	10 to 90%
Humidity (RH), non-operating (when the server is stored or transported)	5 to 93%
Altitude, operating	0 to 10,000 feet (0 to 3048 m)
Altitude, non-operating (when the server is stored or transported)	0 to 39,370 feet (0 to 12,000 m)
Sound pressure level Measure A-weighted per ISO7779 LpAm (dBA)	• Minimum configuration: 57.6 dBA • Typical configuration: 63.5 dBA • Maximum configuration: 70.5 dBA
Sound power	• Minimum configuration: 7.08 bels • Typical configuration: 7.67 bels • Maximum configuration: 8.24 bels

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.

1600 W AC Power Supply

This section lists the specifications for each 1600 W AC power supply (Cisco part number UCSC-PSU1-1600W).

Table 11: 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

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 12: Supported Power Cords

Description	Length (Feet)	Length (Meters)
-------------	---------------	-----------------

CAB-250V-10A-AR AC power cord, 250 V, 10 A Argentina	8.2	2.5
CAB-C13-C14-2M AC cabinet jumper power cord, 250 V, 10 A, C13 to C14	6.6	2.0
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-9K10A-AU AC power cord, 250 V, 10 A, 3112 plug, Australia	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-AC-L620-C13 AC power cord, NEMA L6-20 to C13 connectors	6.6	2.0
CAB-9K10A-IT AC power cord, 250 V, 10 A, CEI 23-16/VII plug Italy	8.2	2.5
R2XX-DMYMPWRCORD No power cord; PID option for ordering server with no power cord	NA	NA



APPENDIX **B**

Storage Controller Considerations

This appendix provides storage controller (RAID) information.

- [Supported Storage Controllers and Cables, on page 157](#)
- [Storage Controller Card Firmware Compatibility, on page 158](#)
- [RAID Backup \(Supercap\), on page 158](#)
- [Write-Cache Policy for Cisco 12G SAS Modular RAID Controllers, on page 158](#)
- [Mixing Drive Types in RAID Groups, on page 159](#)
- [Storage Controller Cable Connectors and Backplanes, on page 159](#)
- [Embedded SATA RAID Controller, on page 160](#)
- [For More RAID Utility Information, on page 169](#)

Supported Storage Controllers and Cables

This server supports up to two PCIe-style, SAS RAID controllers (front and rear). Optionally, the server has a software-based SATA RAID controller embedded in the system that you can use to control two internal M.2 SATA SSDs.



Note NVMe PCIe SSDs cannot be controlled by a SAS/SATA RAID controller.

This server supports the RAID controller options and cable requirements shown in the following table.

Controller	Maximum Drives Controlled	RAID Levels	Optional Supercap Backup?	Required SAS Cables
Embedded RAID (PCH SATA)	Two internal M.2 SATA SSDs.	0, 1	No	No cables are required for control of internal SATA M.2 drives.
Cisco 12G Modular <i>Front</i> RAID Controller UCSC-RAID-M5HD Includes 4-GB cache; controls up to 24 drives	24 front-loading SAS/SATA drives	0, 1, 5, 6, 10, 50, 60	Yes	Use the SAS/SATA cables that come with the chassis (not orderable separately).

Cisco 12G Modular <i>Rear</i> RAID Controller UCSC-SAS9460-8i Includes 2-GB cache; controls up to 8 internal drives in the rear drive-bay module.	8 internal SAS/SATA drives in the rear drive-bay module	0, 1, 5, 6, 10, 50, 60	Yes	Rear RAID cable CBL-AUX-SAS-M5
Cisco 12G 9400-8e HBA for external JBOD attach UCSC-9400-8E	Supported in all server versions: 8 external SAS/SATA ports, controlling up to 1024 external drives.	Non-RAID	No	External drive cables not sold by Cisco. NOTE: This HBA does not support optical cables for connection to external storage (copper only).

Storage Controller Card Firmware Compatibility

Firmware on the storage controller 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.

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 up to two supercap units. The front supercap unit backs up the front RAID controller for front-loading drives. The rear supercap unit backs up the optional, rear RAID controller for internal auxiliary drives.

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.

For supercap unit replacement instructions, see [Replacing the Front RAID Supercap Unit, on page 72](#).

Write-Cache Policy for Cisco 12G SAS Modular RAID Controllers

For this server and other Cisco Generation M5 servers, the default write-cache policy for the Cisco Modular RAID controllers is *Write Through* (irrespective of the presence of a charged supercap or “good BBU”). This utilizes the optimal performance characteristics of the controller.

The write policy can be set to *Write Back*, if preferred. You can set the write policy using the following methods:

- For standalone servers, use the Cisco IMC interface to set Virtual Drive Properties > Write Policy. See the “Managing Storage Adapters” section in your Cisco IMC Configuration Guide.

[Cisco IMC GUI and CLI Configuration Guides](#)

- For Cisco UCS-integrated servers, use the Cisco UCS Manager interface to set the write-cache policy as part of virtual drive configuration in your storage profile.

[Cisco UCS Manager Configuration Guides](#)

- Use the LSI Option ROM Configuration Utility.

Mixing Drive Types in RAID Groups

The following table lists the technical capabilities for mixing hard disk drive (HDD) and solid state drive (SSD) types in a RAID group. However, see the recommendations that follow for the best performance.

Table 13: Mixing Drive Types

Mix of Drive Types in RAID Group	Allowed?
SAS HDD + SATA HDD	Yes
SAS SSD + SATA SSD	Yes
HDD + SSD	No

Drive Type Mixing Best Practices

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.

Embedded SATA RAID

This software RAID option controls only two internal M.2 SATA SSDs. No cabling or other hardware is required.

Cisco 12G Modular SAS RAID Controller With 4-GB Cache (UCSC-RAID-M5HD)

This hardware RAID option can control up to 24 front-loading SAS/SATA drives. The card plugs into a dedicated, horizontal socket on the drive midplane. SAS/SATA cables are used to connect the controller to the backplanes of the front drive modules.

1. Cable card connectors A1 and A2 to the two connectors on front drive module 1.
2. Cable card connectors B1 and B2 to the two connectors on front drive module 2.
3. Cable card connectors C1 and C2 to the two connectors on front drive module 3.

Cisco 12G Modular SAS RAID Controller With 2-GB Cache (UCSC-SAS9460-8i)

This hardware RAID option can control up to 8 internal SAS/SATA drives in the rear drive-bay module. The primary slot for this card is PCIe slot 11. In single CPU-module systems, slot 11 is not supported, and so this card installs to slot 10 in this case. It connects to the rear drive-bay module with two SAS/SATA cables.

1. Connect SAS/SATA cable A1 from the A1 card connector to the A1 auxiliary drive module connector.
2. Connect SAS/SATA cable A2 from the A2 card connector to the A2 auxiliary drive module connector.

Embedded SATA RAID Controller

The server includes an embedded SATA MegaRAID controller that can be used to control internal SATA M.2 drives. This controller supports RAID levels 0 and 1.



Note The VMware ESX/ESXi operating system is not supported with the embedded SATA MegaRAID controller in SW RAID mode. You can use VMWare in AHCI mode.



Note The Microsoft Windows Server 2016 Hyper-V hypervisor is supported for use with the embedded MegaRAID controller in SW RAID mode, but all other hypervisors are not supported. All Hypervisors are supported in AHCI mode.



Note You cannot control the M.2 SATA SSDs in the server with a HW RAID controller.

Embedded SATA RAID Requirements

The embedded SATA RAID controller requires the following items:

- The embedded SATA RAID controller must be enabled in the server BIOS. If you ordered the server with embedded SATA RAID, it is enabled at the factory.
- M.2 mini-storage module with two SATA M.2 SSDs.
- The software RAID controller requires UEFI boot mode; Legacy boot mode is not supported.
- (Optional) LSI MegaSR drivers for Windows or Linux.
- If you use an embedded RAID controller with Linux, both the pSATA and the sSATA controller must be set to LSI SW RAID mode.

Embedded SATA RAID Controller Considerations

Note the following considerations:

- The default setting for this embedded controller hub is SATA RAID 0 and 1 support for two M.2 SATA drives. The hub is divided into two SATA controllers that have different functions. See [Embedded SATA RAID: Two SATA Controllers, on page 161](#).
- When you order the server with this embedded controller, the controller is enabled in the BIOS. Instructions for enabling the controller are included for the case in which a server is reset to defaults. See [Enabling SATA Mode , on page 161](#).
- The required drivers for this controller are already installed and ready to use. However, if you will use this controller with Windows or Linux, you must download and install additional drivers for those operating systems. See [Installing LSI MegaSR Drivers For Windows and Linux, on page 162](#).

Embedded SATA RAID: Two SATA Controllers

The embedded RAID platform controller hub (PCH) is split into two controllers: primary SATA (pSATA) and secondary SATA (sSATA). These two controllers are seen as separate RAID controllers in the Cisco IMC interface and are configurable separately.

- The primary pSATA controller controls only the optional DVD drive; otherwise, it is disabled.
- The secondary sSATA controller controls two internal M.2 SATA drives, when they are present in the M.2 mini-storage module option.
- Each controller is listed separately in the BIOS. You can enable or disable the controllers in the BIOS. See [Enabling SATA Mode , on page 161](#).

Enabling SATA Mode

This procedure uses the server's BIOS Setup Utility

Step 1

Set the SATA mode:

- a) Boot the server and press **F2** when prompted to enter the BIOS Setup utility.
- b) Choose the **Advanced** tab, and then choose **LOM and PCIe Slots Configuration**.
- c) For the primary pSATA controller, select **pSATA** and then choose one of the options from the dialog:
 - **SWR**—Enable the embedded pSATA RAID controller.
 - **AHCI**—Enable control of a DVD drive by AHCI through your OS rather than the embedded RAID controller.
 - **Disabled**—Disable the embedded pSATA RAID controller.
- d) For the secondary sSATA controller, select **M.2** and then choose one of the options from the dialog:
 - **SWR**—Enable the embedded sSATA RAID controller for control of internal SATA M.2 drives.
 - **AHCI**—Enable control of the internal SATA M.2 drives by AHCI through your OS rather than the embedded RAID controller.
 - **Disabled**—Disable the embedded sSATA RAID controller.

Step 2 Press **F10** to save your changes and exit the utility.

Accessing the LSI Software RAID Configuration Utility

To configure RAID settings for the embedded SATA RAID controller, use the utility that is built into the BIOS. Each controller is controlled by its own instance of the utility.

Step 1 Boot the server and press **F2** when prompted to enter the BIOS Setup utility.

Step 2 Choose the **Advanced** tab.

Step 3 Select the instance of the utility that is for the controller that you want to manage (primary or secondary):

- For the pSATA controller, select **LSI Software RAID Configuration Utility (SATA)**.
 - For the sSATA controller, select **LSI Software RAID Configuration Utility (sSATA)**.
-

Installing LSI MegaSR Drivers For Windows and Linux



Note The required drivers for this controller are already installed and ready to use. However, if you will use this controller with Windows or Linux, you must download and install additional drivers for those operating systems.

This section explains how to install the LSI MegaSR drivers for the following supported operating systems:

- Microsoft Windows Server
- Red Hat Enterprise Linux (RHEL)
- SUSE Linux Enterprise Server (SLES)

For the specific supported OS versions, see the [Hardware and Software Compatibility Matrix](#) for your server release.

Downloading the MegaSR Drivers

The MegaSR drivers are included in the C-Series driver ISO for your server and OS.

Step 1 Find the drivers ISO file download for your server online and download it to a temporary location on your workstation:

- a) See the following URL: <http://www.cisco.com/cisco/software/navigator.html>.
- b) Type the name of your server in the **Select a Product** search field and then press **Enter**.
- c) Click **Unified Computing System (UCS) Drivers**.
- d) Click the release number that you are downloading.
- e) Click the Download icon to download the drivers ISO file.

- Step 2** Continue through the subsequent screens to accept the license agreement and then browse to a location where you want to save the driver ISO file.
-

Microsoft Windows Server Drivers

Installing Microsoft Windows Server Drivers

The Windows Server operating system automatically adds the driver to the registry and copies the driver to the appropriate directory.

Before you begin

Before you install this driver on the sSATA embedded controller, you must configure a RAID drive group.

To access the configuration utility, open the BIOS Setup Utility, go to the **Advanced** tab, and then choose the utility instance for the sSATA embedded controller: **LSI Software RAID Configuration Utility (sSATA)**.

- Step 1** Download the Cisco UCS C-Series drivers' ISO, as described in [Downloading the MegaSR Drivers, on page 162](#).
- Step 2** Prepare the drivers on a USB thumb drive:
- Burn the ISO image to a disk.
 - Browse the contents of the drivers folders to the location of the embedded MegaRAID drivers:
/<OS>/Storage/Intel/C600-M5/
 - Expand the Zip file, which contains the folder with the MegaSR driver files.
 - Copy the expanded folder to a USB thumb drive.
- Step 3** Start the Windows driver installation using one of the following methods:
- To install from local media, connect an external USB DVD drive to the server (if the server does not have a DVD drive installed), and then insert the first Windows installation disk into the DVD drive. Skip to Step 6.
 - To install from remote ISO, log in to the server's Cisco IMC interface and continue with the next step.
- Step 4** Launch a Virtual KVM console window and click the **Virtual Media** tab.
- Click **Add Image** and browse to select your remote Windows installation ISO file.
 - Check the check box in the **Mapped** column for the media that you just added, and then wait for mapping to complete.
- Step 5** Power cycle the server.
- Step 6** Press **F6** when you see the F6 prompt during bootup. The Boot Menu window opens.
- Step 7** On the Boot Manager window, choose the physical disk or virtual DVD and press **Enter**. The Windows installation begins when the image is booted.
- Step 8** Press **Enter** when you see the prompt, "Press any key to boot from CD."
- Step 9** Observe the Windows installation process and respond to prompts in the wizard as required for your preferences and company standards.
- Step 10** When Windows prompts you with "Where do you want to install Windows," install the drivers for embedded MegaRAID:
- Click **Load Driver**. You are prompted by a Load Driver dialog box to select the driver to be installed.
 - Connect the USB thumb drive that you prepared in Step 3 to the target server.
 - On the Windows Load Driver dialog, click **Browse**.

- d) Use the dialog box to browse to the location of the drivers folder on the USB thumb drive, and then click **OK**.
Windows loads the drivers from the folder and when finished, the driver is listed under the prompt, “Select the driver to be installed.”
- e) Click **Next** to install the drivers.

Updating Microsoft Windows Server Drivers

- Step 1** Click **Start**, point to **Settings**, and then click **Control Panel**.
 - Step 2** Double-click **System**, click the **Hardware** tab, and then click **Device Manager**. Device Manager starts.
 - Step 3** In Device Manager, double-click **SCSI and RAID Controllers**, right-click the device for which you are installing the driver, and then click **Properties**.
 - Step 4** On the Driver tab, click **Update Driver** to open the Update Device Driver wizard, and then follow the wizard instructions to update the driver.
-

Linux Drivers

Downloading the Driver Image File

See [Downloading the MegaSR Drivers, on page 162](#) for instructions on downloading the drivers. The Linux driver is included in the form of `dud-[driver version].img`, which is the boot image for the embedded MegaRAID stack.



Note The LSI MegaSR drivers that Cisco provides for Red Hat Linux and SUSE Linux are for the original GA versions of those distributions. The drivers do not support updates to those OS kernels.

Preparing Physical Thumb Drive for Linux

This topic describes how to prepare physical Linux thumb drive from the driver image files.

This procedure requires a CD or DVD drive that you can use to burn the ISO image to disk; and a USB thumb drive.

Alternatively, you can mount the `dud.img` file as a virtual floppy disk, as described in the installation procedures.

For RHEL and SLES, you can use a driver disk utility to create disk images from image files.

-
- Step 1** Download the Cisco UCS C-Series drivers ISO, as described in [Downloading the MegaSR Drivers, on page 162](#) and save it to your Linux system.
 - Step 2** Extract the `dud.img` or `dd.iso` driver file:

Note For RHEL 7.1 and later, there is no `dud.img` file--the driver is contained in a `dd.iso` file.

- a) Burn the Cisco UCS C-Series Drivers ISO image to a disc.
- b) Browse the contents of the drivers folders to the location of the embedded MegaRAID drivers:
/`<OS>`/Storage/Intel/C600-M5/

c) Expand the Zip file, which contains the folder with the driver files.

Step 3 Copy the driver update disk image `dud-[driver version].img` (or `dd.iso`) to your Linux system.

Step 4 Insert a blank USB thumb drive into a port on your Linux system.

Step 5 Create a directory and mount the `dud.img` or `dd.iso` image to that directory:

Example:

```
mkdir <destination_folder>
mount -o loop <driver_image> <destination_folder>
```

Step 6 Copy the contents in the directory to your USB thumb drive.

Installing the Red Hat Enterprise Linux Driver

For the specific supported OS versions, see the [Hardware and Software Compatibility Matrix](#) for your server release.

This topic describes the fresh installation of the RHEL device driver on systems that have the embedded MegaRAID stack.



Note If you use an embedded RAID controller with Linux and a DVD drive is present on the pSATA controller, both the pSATA and the sSATA controller must be set to `LSI SW RAID` mode.

Before you begin

Before you install this driver on the sSATA embedded controller, you must configure a RAID drive group.

To access the configuration utility, open the BIOS Setup Utility, go to the **Advanced** tab, and then choose the utility instance for the sSATA embedded controller: **LSI Software RAID Configuration Utility (sSATA)**.

Step 1 Prepare the `dud.img` file using one of the following methods:

Note For RHEL 7.1 and later, there is no `dud.img` file--the driver is contained in a `dd.iso` file.

- To install from physical disk, use the procedure in [Preparing Physical Thumb Drive for Linux, on page 164](#), then continue with step 3.
- To install from *virtual* disk, download the Cisco UCS C-Series drivers' ISO, as described in [Downloading the MegaSR Drivers, on page 162](#), then continue with the next step.

Step 2 Extract the `dud.img` (or `dd.iso`) file:

- Burn the Cisco UCS C-Series Drivers ISO image to a disk.
- Browse the contents of the drivers folders to the location of the embedded MegaRAID drivers:
`/<OS>/Storage/Intel/C600-M5/`
- Copy the `dud-<driver version>.img` (or `dd.iso`) file to a temporary location on your workstation.
- If you are using RHEL 7.x, rename the saved `dd.iso` to `dd.img`.

Note If you are using RHEL 7.x, renaming the dd.iso file to dd.img simplifies this procedure and saves time. The Cisco UCS virtual drive mapper can map only one .iso at a time, and only as a virtual CD/DVD. Renaming the file to dd.img allows you to mount the RHEL installation ISO as a virtual CD/DVD and the renamed dd.img as a virtual floppy disk or removable disk at the same time. This avoids the steps of unmounting and remounting the RHEL ISO when the dd.iso driver file is prompted for.

Step 3 Start the Linux driver installation using one of the following methods:

- To install from local media, connect an external USB CD/DVD drive to the server and then insert the first RHEL installation disk into the drive. Then continue with Step 5.
- To install from virtual disk, log in to the server's Cisco IMC interface. Then continue with the next step.

Step 4 Launch a Virtual KVM console window and click the **Virtual Media** tab.

- a) Click **Add Image** and browse to select your remote RHEL installation ISO image.

Note An .iso file can be mapped only as a virtual CD/DVD.

- b) Click **Add Image** again and browse to select your RHEL 6.x `dud.img` or the RHEL 7.x `dd.img` file that you renamed in step 2.

Note Map the .img file as a virtual floppy disk or virtual removable disk.

- c) Check the check boxes in the **Mapped** column for the media that you just added, then wait for mapping to complete.

Step 5 Power-cycle the target server.

Step 6 Press **F6** when you see the F6 prompt during bootup. The Boot Menu window opens.

Note Do not press Enter in the next step to start the installation. Instead, press **e** to edit installation parameters.

Step 7 On the Boot Menu window, use the arrow keys to select **Install Red Hat Enterprise Linux** and then press **e** to edit installation parameters.

Step 8 Append one of the following blacklist commands to the end of the line that begins with **linuxefi**:

- For RHEL 6.x (32- and 64-bit), type:

```
linux dd blacklist=isci blacklist=ahci nodmraid noprobe=<atadrive number>
```

Note The noprobe values depend on the number of drives. For example, to install RHEL 6.x on a RAID 5 configuration with three drives, type:

```
Linux dd blacklist=isci blacklist=ahci nodmraid noprobe=ata1 noprobe=ata2
```

- For RHEL 7.x (32- and 64-bit), type:

```
linux dd modprobe.blacklist=ahci nodmraid
```

Step 9 **Optional:** To see full, verbose installation status steps during installation, delete the **Quiet** parameter from the line.

Step 10 On the Boot Menu window, press **Ctrl+x** to start the interactive installation.

Step 11 Below **Driver disk device selection**, select the option to install your driver .img file. (Type **r** to refresh the list if it is not populated.)

Note The installer recognizes the driver file as an .iso file, even though you renamed it to dd.img for mapping.

Type the number of the driver device ISO in the list. Do *not* select the RHEL ISO image. In the following example, type **6** to select device sdb:

```
5) sr0 iso9660 RHEL-7.6\x20Server.x
6) sdb iso9660 CDROM
# to select, 'r' - refresh, or 'c' -continue: 6
```

The installer reads the driver file and lists the drivers.

Step 12 Under **Select drivers to install**, type the number of the line that lists the megasr driver. In the following example, type **1**:

```
1) [ ] /media/DD-1/rpms/x86_61/kmod-megasr-18.01.2010.1107_e17.6-1.x86_61.rpm
# to toggle selection, or 'c' -continue: 1
```

Your selection is displayed with an X in brackets.

```
1) [X] /media/DD-1/rpms/x86_61/kmod-megasr-18.01.2010.1107_e17.6-1.x86_61.rpm
```

Step 13 Type **c** to continue.

Step 14 Follow the RHEL installation wizard to complete the installation.

Step 15 When the wizard's Installation Destination screen is displayed, ensure that **LSI MegaSR** is listed as the selection. If it is not listed, the driver did not load successfully. In that case, select **Rescan Disc**.

Step 16 After the installation completes, reboot the target server.

Installing the SUSE Linux Enterprise Server Driver

For the specific supported OS versions, see the [Hardware and Software Compatibility Matrix](#) for your server release.

This topic describes the fresh installation of the SLES driver on systems that have the embedded MegaRAID stack.



Note If you use an embedded RAID controller with Linux and a DVD drive is present on the pSATA controller, both the pSATA and the sSATA controller must be set to `LSI SW RAID` mode.

Before you begin

Before you install this driver on the sSATA embedded controller, you must configure a RAID drive group.

To access the configuration utility, open the BIOS Setup Utility, go to the **Advanced** tab, and then choose the utility instance for the sSATA embedded controller: **LSI Software RAID Configuration Utility (sSATA)**.

Step 1 Prepare the `dud.img` (or `.iso`) file using one of the following methods:

- To install from physical disk, use the procedure in [Preparing Physical Thumb Drive for Linux, on page 164](#), then continue with step 4.
- To install from *virtual* disk, download the Cisco UCS C-Series drivers' ISO, as described in [Downloading the MegaSR Drivers, on page 162](#), then continue with the next step.

- Step 2** Extract the `dud.img` file that contains the driver:
- Burn the ISO image to a disk.
 - Browse the contents of the drivers folders to the location of the embedded MegaRAID drivers:
`/<OS>/Storage/Intel/C600-M5/...`
 - Within the SLES folder for your version, the `dud-<driver version>.img` file is packaged in a compressed `.gz` file. Extract the `.img` file from the `.gz` file.
 - Copy the `dud-<driver version>.img` file to a temporary location on your workstation.
- Step 3** Start the Linux driver installation using one of the following methods:
- To install from local media, connect an external USB DVD drive to the server and then insert the first SLES installation disk into the drive. Then continue with Step 5.
 - To install from remote ISO, log in to the server's Cisco IMC interface. Then continue with the next step.
- Step 4** Launch a Virtual KVM console window and click the **Virtual Media** tab.
- Click **Add Image** and browse to select your remote SLES installation ISO file.
 - Click **Add Image** again and browse to select your `dud-<driver version>.img` file.
 - Check the check boxes in the **Mapped** column for the media that you just added, then wait for mapping to complete.
- Step 5** Power-cycle the target server.
- Step 6** Press **F6** when you see the F6 prompt during bootup. The Boot Menu window opens.
- Step 7** On the Boot Manager window, select the physical or virtual SLES installation ISO and press **Enter**.
The SLES installation begins when the image is booted.
- Step 8** When the first SLES screen appears, select **Installation**.
- Step 9** Press **e** to edit installation parameters.
- Step 10** Append the following parameter to the end of the line that begins with **linuxefi**:
- ```
brokenmodules=ahci
```
- Step 11** **Optional:** To see detailed status information during the installation, add the following parameter to the line that begins with **linuxefi**:
- ```
splash=verbose
```
- Step 12** Press **Ctrl+x** to start the installation.
- The installation proceeds. The installer finds the LSI driver automatically in the `dud-<driver version>.img` file that you provided. With verbose status messages, you see the driver being installed when LSI MegaRAID SW RAID Module is listed.
- Step 13** Follow the SLES installation wizard to complete the installation. Verify installation of the driver when you reach the **Suggested Partitioning** screen:
- On the **Suggested Partitioning** screen, select **Expert Partitioner**.
 - Navigate to **Linux > Hard disks** and verify that there is a device listed for the LSI - LSI MegaSR driver. The device might be listed as a type other than `sda`. For example:
- ```
dev/sdd: LSI - LSI MegaSR
```
- If no device is listed, the driver did not install properly. In that case, repeat the steps above.

**Step 14** When installation is complete, reboot the target server.

---

## For More RAID Utility Information

The Broadcom utilities have help documentation for more information about using the utilities.

- 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](#).
- For hardware SAS MegaRAID configuration—[Broadcom 12Gb/s MegaRAID SAS Software User Guide, Version 2.8](#)
- For embedded software MegaRAID and the utility that is accessed via the server BIOS (refer to Chapter 4)—[Broadcom Embedded MegaRAID Software User Guide, March 2018](#).





## APPENDIX C

# GPU Card Installation

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

- [Server Firmware Requirements, on page 171](#)
- [GPU Card Configuration Rules, on page 172](#)
- [Requirement For All GPUs: Memory-Mapped I/O Greater Than 4 GB, on page 173](#)
- [Installing a Double-Wide GPU Card, on page 174](#)
- [Installing a Single-Wide GPU Card, on page 176](#)
- [Using NVIDIA GRID License Server For P-Series and T-Series GPUs, on page 178](#)
- [Installing Drivers to Support the GPU Cards, on page 186](#)

## 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 |
|---------------------------|-----------------------------------------|
| NVIDIA Tesla A100         | 4.1(3)                                  |
| NVIDIA Tesla M10          | 3.1(2)                                  |
| NVIDIA Tesla P40          | 3.1(2)                                  |
| NVIDIA Tesla P100 12GB    | 3.1(2)                                  |
| NVIDIA Tesla P100 16GB    | 3.1(2)                                  |
| AMD FirePro S7150 X2      | 3.1(2)                                  |
| NVIDIA Tesla M60          | 3.1(2c)                                 |
| NVIDIA Tesla V100<br>16GB | 3.1(3)                                  |
| NVIDIA Tesla V100<br>32GB | 4.0(1)                                  |

|           |                                                                                                                                                                                                                                          |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NVIDIA T4 | <p>4.0(4a): the server supports up to 6 T4 GPUs</p> <p>4.0(4e) and later: the server supports up to 10 T4 GPUs</p> <p>Note for UCS Manager-integrated servers: Cisco UCS Manager 4.0(4c) or later is required to support 10 T4 GPUs.</p> |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## GPU Card Configuration Rules

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



### Caution

When using NVIDIA Tesla V100 16GB or V100 32GB GPU cards in this server, there are special temperature requirements. See [Installing a Double-Wide GPU Card, on page 174](#).

- This server supports up to six high-power, double-wide GPU cards. The motherboard includes six GPU card power connectors. Double-wide GPU cards can be installed in PCIe slots 1, 2, 4, 6, 8, 10.

See the figure below for the double-wide GPU card spacing and the correlation to the six GPU-power connectors.

- The server supports up to 6 single-wide NVIDIA T4 GPUs with server firmware release 4.0(4a).

The server supports up to 10 single-wide NVIDIA T4 GPUs with server firmware release 4.0(4e) and later.

Note for UCS Manager-integrated servers: Cisco UCS Manager 4.0(4c) or later is required to support 10 T4 GPUs.

- In a single CPU-module system, only PCIe slots 1, 2, 8, and 10 can be used for double-wide GPUs.
- The number of slots available for GPUs might be reduced because of the other cards installed in the server. For example, if present, the NVMe switch card must be installed in slot 11.
- Use the UCS power calculator at the following link to determine the power needed based on your server configuration: <http://ucspowercalc.cisco.com>
- Do not mix different brands or models of GPU cards in the server.
- Some GPUs have a limitation on whether they can support 1 TB or more memory in the server:


**Table 14: Total Server Memory Support**

| GPU               | Support 1TB or More Memory in Server |
|-------------------|--------------------------------------|
| NVIDIA Tesla A100 | Yes                                  |
| NVIDIA M-Series   | No                                   |
| NVIDIA P-Series   | Yes                                  |
| NVIDIA T-Series   | Yes                                  |
| NVIDIA V-Series   | Yes                                  |

| GPU                  | Support 1TB or More Memory in Server |
|----------------------|--------------------------------------|
| AMD FirePro S7150 X2 | No                                   |

- The following NVIDIA GPUs are not supported with Second Generation Intel Xeon Scalable processors:
  - NVIDIA Tesla P100 12G
  - NVIDIA Tesla P100 16G

Figure 76: Double-Wide GPU Spacing in PCIe Slots



## 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



**Note** Make sure that you have the correct power cable for your GPU card, as described in the following table.

In the table below, the cable that is used with the GPU is listed. It is also indicated whether the cable is included in the GPU BOM or must be ordered separately.

- **Separate** = Cable must be ordered separately when the ordering tool prompts you.
- **Included** = Cable is included with the GPU; no additional action is needed.

**Table 15: Double-Wide GPU Required Power Cables**

| GPU                    | GPU Power Cable     | Cable Included When the GPU Card is Ordered With a System Order? | Cable Included When the GPU Card is Ordered as a Spare? |
|------------------------|---------------------|------------------------------------------------------------------|---------------------------------------------------------|
| Nvidia Tesla A100      | UCS-P100CBL-C480-M5 | Included                                                         | Separate                                                |
| Nvidia Tesla M10       | UCSC-P10CBL-C480M5  | Included                                                         | Separate                                                |
| Nvidia Tesla P40       | UCSC-P100CBL-C480M5 | Included                                                         | Separate                                                |
| Nvidia Tesla P100 12GB | UCSC-P100CBL-C480M5 | Included                                                         | Separate                                                |

| GPU                    | GPU Power Cable     | Cable Included When the GPU Card is Ordered With a System Order? | Cable Included When the GPU Card is Ordered as a Spare? |
|------------------------|---------------------|------------------------------------------------------------------|---------------------------------------------------------|
| Nvidia Tesla P100 16GB | UCSC-P100CBL-C480M5 | Included                                                         | Separate                                                |
| Nvidia Tesla M60       | UCSC-P100CBL-C480M5 | Included                                                         | Separate                                                |
| Nvidia Tesla V100 16GB | UCSC-P100CBL-C480M5 | Included                                                         | Included                                                |
| Nvidia Tesla V100 32GB | UCSC-P100CBL-C480M5 | Included                                                         | Included                                                |
| AMD FirePro S7150 X2   | UCSC-AMDCBL-C480M5  | Included                                                         | Separate                                                |

**Caution**

When using NVIDIA Tesla V100 16GB or V100 32GB GPU cards in this server, there are special temperature requirements, as described in the following table.

**Table 16: C480 M5 Operating Temperature Requirements For Double-Wide GPU Cards**

| GPU Card                | Maximum Server Operating Temperature (Air Inlet Temperature)                                                                                                                                    |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AMD FirePro S7150 X2    | 35° C (95.0° F)                                                                                                                                                                                 |
| NVIDIA Tesla A100       | 35° C (95.0° F)                                                                                                                                                                                 |
| NVIDIA Tesla M10        | 35° C (95.0° F)                                                                                                                                                                                 |
| NVIDIA Tesla M60        | 35° C (95.0° F)                                                                                                                                                                                 |
| NVIDIA Tesla P40        | 35° C (95.0° F)                                                                                                                                                                                 |
| NVIDIA Tesla P100 12GB  | 35° C (95.0° F)                                                                                                                                                                                 |
| NVIDIA Tesla P100 16GB  | 35° C (95.0° F)                                                                                                                                                                                 |
| NVIDIA Tesla V100 16 GB | Differs by number of GPUs installed: <ul style="list-style-type: none"> <li>• Four or less GPUs installed: 35° C (95.0° F)</li> <li>• More than four GPUs installed: 32° C (89.6° F)</li> </ul> |
| NVIDIA Tesla V100 32 GB | Differs by number of GPUs installed: <ul style="list-style-type: none"> <li>• Four or less GPUs installed: 35° C (95.0° F)</li> <li>• More than four GPUs installed: 32° C (89.6° F)</li> </ul> |





**Note** For **NVIDIA GPUs:** The NVIDIA GPU card might be shipped with two power cables: a straight cable and a Y-cable. The straight cable is used for connecting power to the GPU card in this server; do not use the Y-cable, which is used for connecting the GPU card in external devices only (such as the Magma chassis).

For **AMD GPUs:** The correct power cable is a Y-cable (UCS-AMDCBL-C480M5).

**Step 1** Shut down and remove power from the server as described in [Shutting Down and Removing Power From the Server, on page 39](#).

**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 41](#).

**Step 4** Remove any existing card or a blanking panel:

- a) Open the hinged retainer bar that covers the top of the PCIe slot.

Use your fingertips to pull back on the wire locking-latches at each end of the retainer bar, and then hinge the bar open to expose the tops of the PCIe slots.

- b) Pull both ends of the card vertically to disengage the card from the socket, and then set it aside.

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

**Note** See [GPU Card Configuration Rules, on page 172](#) for which PCIe slots support double-wide GPUs.

- a) Carefully align the card edge with the socket while you align the card's rear tab with the rear panel opening.
- b) Push down on both corners of the card to seat its edge connector in the socket.
- c) Close the hinged retainer bar over the top of the PCIe slots.

Use your fingertips to pull back on the wire locking-latches at each end of the retainer bar, and then hinge it closed to secure the tops of the PCIe slots. Push the wire locking-latching back to the forward, locked position.

- d) Connect the GPU card power cable to the front of the GPU card and to the motherboard power connector.

**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.

**Step 8** Optional: Continue with [Installing Drivers to Support the GPU Cards, on page 186](#).

**Note** If you installed an NVIDIA Tesla M-series or P-Series GPU, you must install GRID licenses to use the GRID features. See [Using NVIDIA GRID License Server For P-Series and T-Series GPUs, on page 178](#).

## Installing a Single-Wide GPU Card

Use the following procedure to install or replace the following supported single-wide GPU card:

- NVIDIA T4: These half-height, half-length (HHHL) GPU cards are supported in all PCIe slots.

With Cisco IMC 4.0(4a), this server can support up to six single-wide NVIDIA T4 GPU cards.

With Cisco IMC 4.0(4e) and later, this server can support up to 10 single-wide NVIDIA T4 GPU cards.

Note for UCS Manager-integrated servers: Cisco UCS Manager 4.0(4c) or later is required to support 10 T4 GPUs.

For information about PCIe slot specifications and restrictions, see [PCIe Slot Specifications and Restrictions, on page 94](#).

---

**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 39](#).
- 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 41](#).

**Step 2**

Remove any existing card or a blanking panel:

- a) Open the hinged retainer bar that covers the top of the PCIe slot.

Use your fingertips to pull back on the wire locking-latches at each end of the retainer bar, and then hinge the bar open to expose the tops of the PCIe slots.

- b) Pull both ends of the card vertically to disengage the card from the socket, and then set it aside.

**Step 3**

Install a new PCIe card:

- a) Carefully align the card edge with the socket while you align the card's rear tab with the rear panel opening.
- b) Push down on both corners of the card to seat its edge connector in the socket.
- c) Close the hinged retainer bar over the top of the PCIe slots.

Use your fingertips to pull back on the wire locking-latches at each end of the retainer bar, and then hinge it closed to secure the tops of the PCIe slots. Push the wire locking-latches back to the forward, locked position.

**Step 4**

Replace the top cover to the server.


**Step 5**

Reconnect power cords to all power supplies and then allow the server to boot to standby power mode (indicated when the front panel Power button LED lights amber).

**Step 6**

Fully power on the server by pressing the Power button.

Figure 77: PCIe Slot Hinged Retainer Bars



## Using NVIDIA GRID License Server For P-Series and T-Series GPUs

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

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

1. Familiarize yourself with the NVIDIA GRID License Server.

[NVIDIA GRID License Server Overview, on page 179](#)

2. Register your product activation keys with NVIDIA.

[Registering Your Product Activation Keys With NVIDIA, on page 180](#)

3. Download the GRID software suite.

[Downloading the GRID Software Suite, on page 180](#)

4. Install the GRID License Server software to a host.

[Installing NVIDIA GRID License Server Software, on page 180](#)

5. Generate licenses on the NVIDIA Licensing Portal and download them.

[Installing GRID Licenses From the NVIDIA Licensing Portal to the License Server, on page 182](#)

## 6. Manage your GRID licenses.


[Managing GRID Licenses](#) , on page 183

# NVIDIA GRID License Server Overview

The NVIDIA M-Series GPUs combine Tesla and GRID functionality when the licensed GRID features such as GRID vGPU and GRID Virtual Workstation are enabled. These features are enabled during OS boot by borrowing a software license that is served over the network from the NVIDIA GRID License Server virtual appliance. The license is returned to the license server when the OS shuts down.

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

Figure 78: NVIDIA GRID Licensing Architecture



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

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

## Registering Your Product Activation Keys With NVIDIA

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

- 
- Step 1** Select the **Log In** link, or the **Register** link if you do not already have an account.  
The NVIDIA Software Licensing Center > License Key Registration dialog opens.
- Step 2** Complete the License Key Registration form and then click **Submit My Registration Information**.  
The NVIDIA Software Licensing Center > Product Information Software dialog opens.
- Step 3** If you have additional PAKs, click **Register Additional Keys**. For each additional key, complete the form on the License Key Registration dialog and then click **Submit My Registration Information**.
- Step 4** Agree to the terms and conditions and set a password when prompted.
- 

## Downloading the GRID Software Suite

- 
- Step 1** Return to the NVIDIA Software Licensing Center > Product Information Software dialog.
- Step 2** Click the **Current Releases** tab.
- Step 3** Click the **NVIDIA GRID** link to access the Product Download dialog. This dialog includes download links for:
- NVIDIA License Manager software
  - The gpumodeswitch utility
  - The host driver software
- Step 4** Use the links to download the software.
- 

## Installing NVIDIA GRID License Server Software

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

<http://www.nvidia.com>

### Platform Requirements for NVIDIA GRID License Server

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

- The hosting platform's date and time must be set accurately.

## Installing GRID License Server on Windows

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

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

## Installing GRID License Server on Linux

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

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

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

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

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

- d) Verify that the Tomcat service is operational. Open a web browser on the License Server host and connect to the URL <http://localhost:8080>. If the installation was successful, you see the Tomcat webapp.

### Step 3 Install the License Server:

- a) Unpack the License Server tar file using the following command:

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

- b) Run the unpacked setup binary as root:

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

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

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

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

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

### Accessing the GRID License Server Management Interface

Open a web browser on the License Server host and access the URL <http://localhost:8080/licserver>.

If you configured the License Server host's firewall to permit remote access to the License Server, the management interface is accessible from remote machines at the URL <http://hostname:8080/licserver>

## Reading Your License Server's MAC Address

Your License Server's Ethernet MAC address is used as an identifier when registering the License Server with NVIDIA's Licensing Portal.

- Step 1** Access the GRID License Server Management Interface in a browser.

- Step 2** In the left-side License Server panel, select **Configuration**.

The License Server Configuration panel opens. Next to **Server host ID**, a pull-down menu lists the possible Ethernet MAC addresses.

- Step 3** Select your License Server's MAC address from the **Server host ID** pull-down.

**Note** It is important to use the same Ethernet ID consistently to identify the server when generating licenses on NVIDIA's Licensing Portal. NVIDIA recommends that you select one entry for a primary, non-removable Ethernet interface on the platform.

## Installing Licenses From the Licensing Portal

---

**Step 1** Access the GRID License Server Management Interface in a browser.

**Step 2** In the left-side License Server panel, select **Configuration**.

The License Server Configuration panel opens.

**Step 3** Use the License Server Configuration menu to install the .bin file that you generated earlier.

- a) Click **Choose File**.
- b) Browse to the license .bin file that you want to install and click **Open**.
- c) Click **Upload**.

The license file is installed on your License Server. When installation is complete, you see the confirmation message, “Successfully applied license file to license server.”

---

## Viewing Available GRID Licenses

Use the following procedure to view which licenses are installed and available, along with their properties.

---

**Step 1** Access the GRID License Server Management Interface in a browser.

**Step 2** In the left-side License Server panel, select **Licensed Feature Usage**.

**Step 3** Click on a feature in the **Feature** column to see detailed information about the current usage of that feature.

---

## Viewing Current License Usage

Use the following procedure to view information about which licenses are currently in-use and borrowed from the server.

---

**Step 1** Access the GRID License Server Management Interface in a browser.

**Step 2** In the left-side License Server panel, select **Licensed Clients**.

**Step 3** To view detailed information about a single licensed client, click on its **Client ID** in the list.

---

## Managing GRID Licenses

Features that require GRID licensing run at reduced capability until a GRID license is acquired.

## Acquiring a GRID License on Windows

---

**Step 1** Open the NVIDIA Control Panel using one of the following methods:

- Right-click on the Windows desktop and select **NVIDIA Control Panel** from the menu.



- Open Windows Control Panel and double-click the **NVIDIA Control Panel** icon.

**Step 2** In the NVIDIA Control Panel left-pane under Licensing, select **Manage License**.

The Manage License task pane opens and shows the current license edition being used. The GRID software automatically selects the license edition based on the features that you are using. The default is Tesla (unlicensed).

**Step 3** If you want to acquire a license for GRID Virtual Workstation, under License Edition, select **GRID Virtual Workstation**.

**Step 4** In the **License Server** field, enter the address of your local GRID License Server. The address can be a domain name or an IP address.

**Step 5** In the **Port Number** field, enter your port number or leave it set to the default used by the server, which is 7070.

**Step 6** Select **Apply**.

The system requests the appropriate license edition from your configured License Server. After a license is successfully acquired, the features of that license edition are enabled.

**Note** After you configure licensing settings in the NVIDIA Control Panel, the settings persist across reboots.

---

## Acquiring a GRID License on Linux

---

**Step 1** Edit the configuration file `/etc/nvidia/gridd.conf`:

```
sudo vi /etc/nvidia/gridd.conf
```

**Step 2** Edit the `ServerUrl` line with the address of your local GRID License Server.

The address can be a domain name or an IP address. See the example file below.

**Step 3** Append the port number (default 7070) to the end of the address with a colon. See the example file below.

**Step 4** Edit the `FeatureType` line with the integer for the license type. See the example file below.

- GRID vGPU = 1
- GRID Virtual Workstation = 2

**Step 5** Restart the `nvidia-gridd` service.

```
sudo service nvidia-gridd restart
```

The service automatically acquires the license edition that you specified in the `FeatureType` line. You can confirm this in `/var/log/messages`.

**Note** After you configure licensing settings in the NVIDIA Control Panel, the settings persist across reboots.

Sample configuration file:

```
/etc/nvidia/gridd.conf - Configuration file for NVIDIA Grid Daemon
Description: Set License Server URL
Data type: string
Format: "<address>:<port>"
ServerUrl=10.31.20.45:7070

Description: Set Feature to be enabled
Data type: integer
```

```
Possible values:
1 => for GRID vGPU
2 => for GRID Virtual Workstation
FeatureType=2
```

---

## Using gpumodeswitch

The command line utility `gpumodeswitch` can be run in the following environments:

- Windows 64-bit command prompt (requires administrator permissions)
- Linux 32/64-bit shell (including Citrix XenServer dom0) (requires root permissions)



---

**Note** Consult NVIDIA product release notes for the latest information on compatibility with compute and graphic modes.

---

The `gpumodeswitch` utility supports the following commands:

- `--listgpumodes`

Writes information to a log file named `listgpumodes.txt` in the current working directory.

- `--gpumode graphics`

Switches to graphics mode. Switches mode of all supported GPUs in the server unless you specify otherwise when prompted.

- `--gpumode compute`

Switches to compute mode. Switches mode of all supported GPUs in the server unless you specify otherwise when prompted.



---

**Note** After you switch GPU mode, reboot the server to ensure that the modified resources of the GPU are correctly accounted for by any OS or hypervisor running on the server.

---

## 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>
  - AMD: <http://support.amd.com/en-us/download>
- Step 4** Restart the server.
- Step 5** Check that the virtual machine is able to recognize the GPU card. In Windows, use the Device Manager and look under Display Adapters.
-





## APPENDIX **D**

# Installation For Cisco UCS Manager Integration

---

- [Installation For Cisco UCS Manager Integration, on page 189](#)

## 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](#)

