Cisco UCS C4200 Server Chassis Installation and Service Guide

First Published: 2018-08-09

Last Modified: 2020-08-10

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)

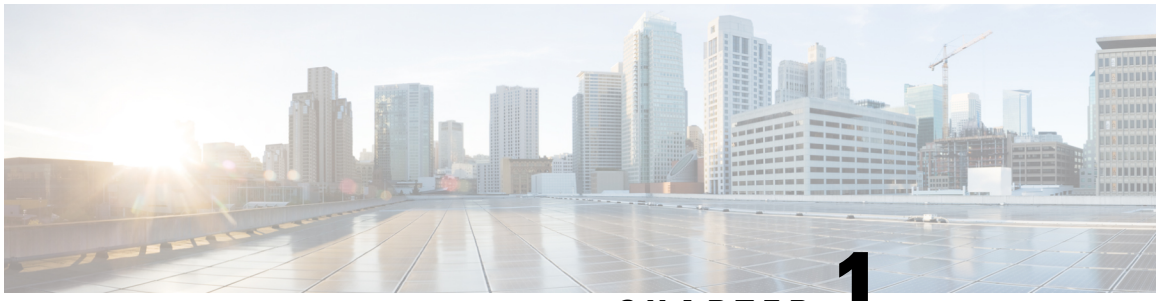

CHAPTER 1

Overview

Overview

The Cisco UCS C4200 Server Chassis is a 2RU, rack-mount chassis that provides shared storage, cooling, and power for up to four removeable compute nodes. Each of the four removable compute nodes can control 6 front-loading drives in the chassis, for a total of up to 24 small form-factor (SFF), 2.5-inch, SAS/SATA HDDs or SSDs.

With Cisco IMC 4.0(2) and later, each of the four compute nodes can control two front-loading NVMe SSDs, for a total of up to eight NVME SSDs in the chassis.

For information about compute nodes, see the service note for your compute node:

- Cisco UCS C125 Compute Node Service Note

External Features

This topic shows the external features of the server chassis.

Cisco UCS C4200 Chassis Front Panel Features (SFF, 24-Drive)

The following figure shows the front panel features.

For definitions of LED states, see Front-Panel LEDs, on page 25.

Figure 1: Front Panel

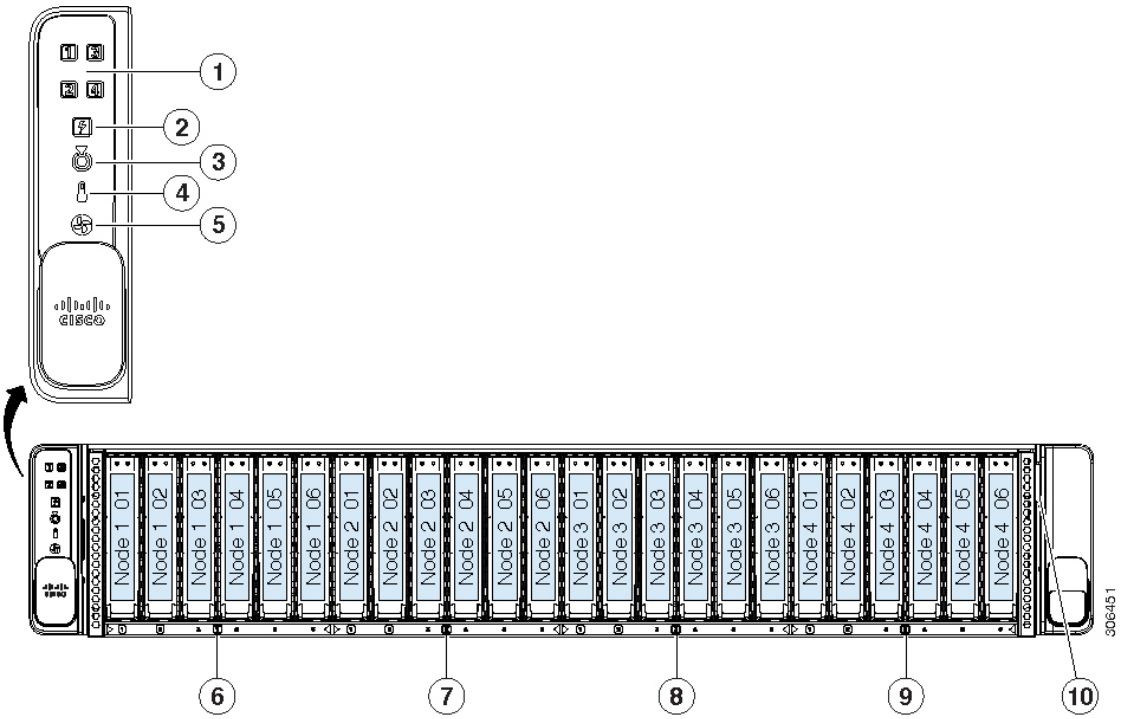

1	Node health LEDs	**6**	Node 1-controlled drive bays 1—6 All six bays support SAS/SATA drives; bays 1 and 2 also support NVME drives.
2	Power supply status LED	**7**	Node 2-controlled drive bays 1—6 All six bays support SAS/SATA drives; bays 1 and 2 also support NVME drives.
3	Locator beacon LED Activating the locator beacon of any installed compute node activates this chassis locator beacon.	**8**	Node 3-controlled drive bays 1—6 All six bays support SAS/SATA drives; bays 1 and 2 also support NVME drives.
4	Temperature status LED	**9**	Node 4-controlled drive bays 1—6 All six bays support SAS/SATA drives; bays 1 and 2 also support NVME drives.
5	Fan status LED	**10**	Pull-out asset tag

Cisco UCS C4200 Chassis Rear Panel Features

The exact features depend on how many compute nodes are installed in the node bays and which cards are installed in the nodes. The sample figure below shows a chassis with four Cisco UCS C125 M5 compute nodes installed.

Although the power supplies are the only components native to the chassis in the view below, features of a removeable compute node are defined to explain network connections for the system. For information about node components, see the Cisco UCS C125 Compute Node Service Note.

Note All node bays must have either a compute node or a node blank installed to ensure adequate air flow.

For definitions of LED states, see Rear-Panel LEDs, on page 27.

Figure 2: Cisco UCS C4200 Chassis Rear Panel (Shown With Four C125 M5 Compute Nodes Installed)

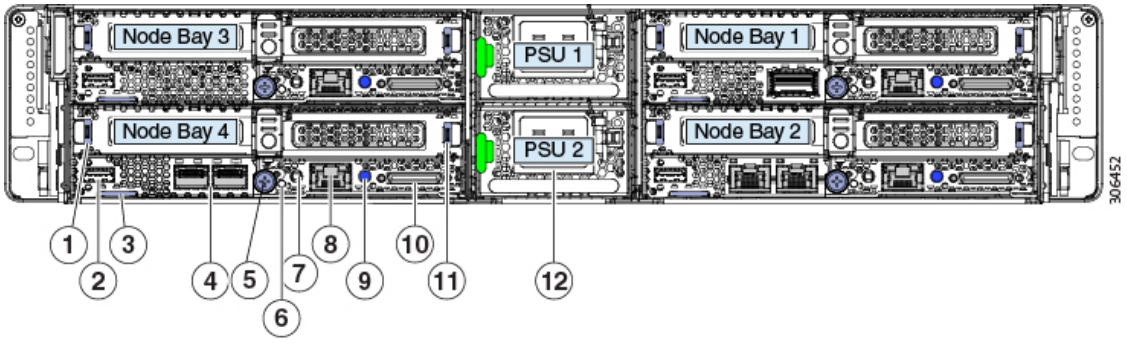

1	PCIe riser 1 handle (one each node) • Node PCIe riser 1/slot 1 (half-height, half length, x8 slot)	**7**	Node Power button/Power status LED (one each node)
2	Node USB 3.0 port (one each node)	**8**	Node 1 Gb Ethernet dedicated management port (one each node)
3	Node pull-out asset tag (one each node)	**9**	Node locator button/LED (one each node)
4	Node OCP adapter card Ethernet LAN ports (one each node, if this optional adapter card is installed) Depending on which adapter card is installed, these ports can be either: • Dual 10 Gb Base-T (RJ-45 connectors) • Dual 10/25 Gb (SFP 28 connectors) • Single 100 Gb (QSFP 28 connector)	**10**	Node KVM local debug console port (one each node) Used with KVM cable that provides one DB-15 VGA, one DB-9 serial, and two USB 2.0 connectors.
5	Node securing thumbscrew and release lever (one each node)	**11**	PCIe riser 2 handle (one each node) • Node PCIe riser 2/slot 2 (half-height, half length, x16 slot)
6	Node Health Status LED	**12**	Chassis power supplies (two, redundant 1+1)

Serviceable Components in the Chassis

The figure in this topic shows the locations of the serviceable components in the chassis.

For components inside a compute node, see the service note for your compute node:

- Cisco UCS C125 Compute Node Service Note

Figure 3: Cisco UCS C4200 Chassis Serviceable Component Locations

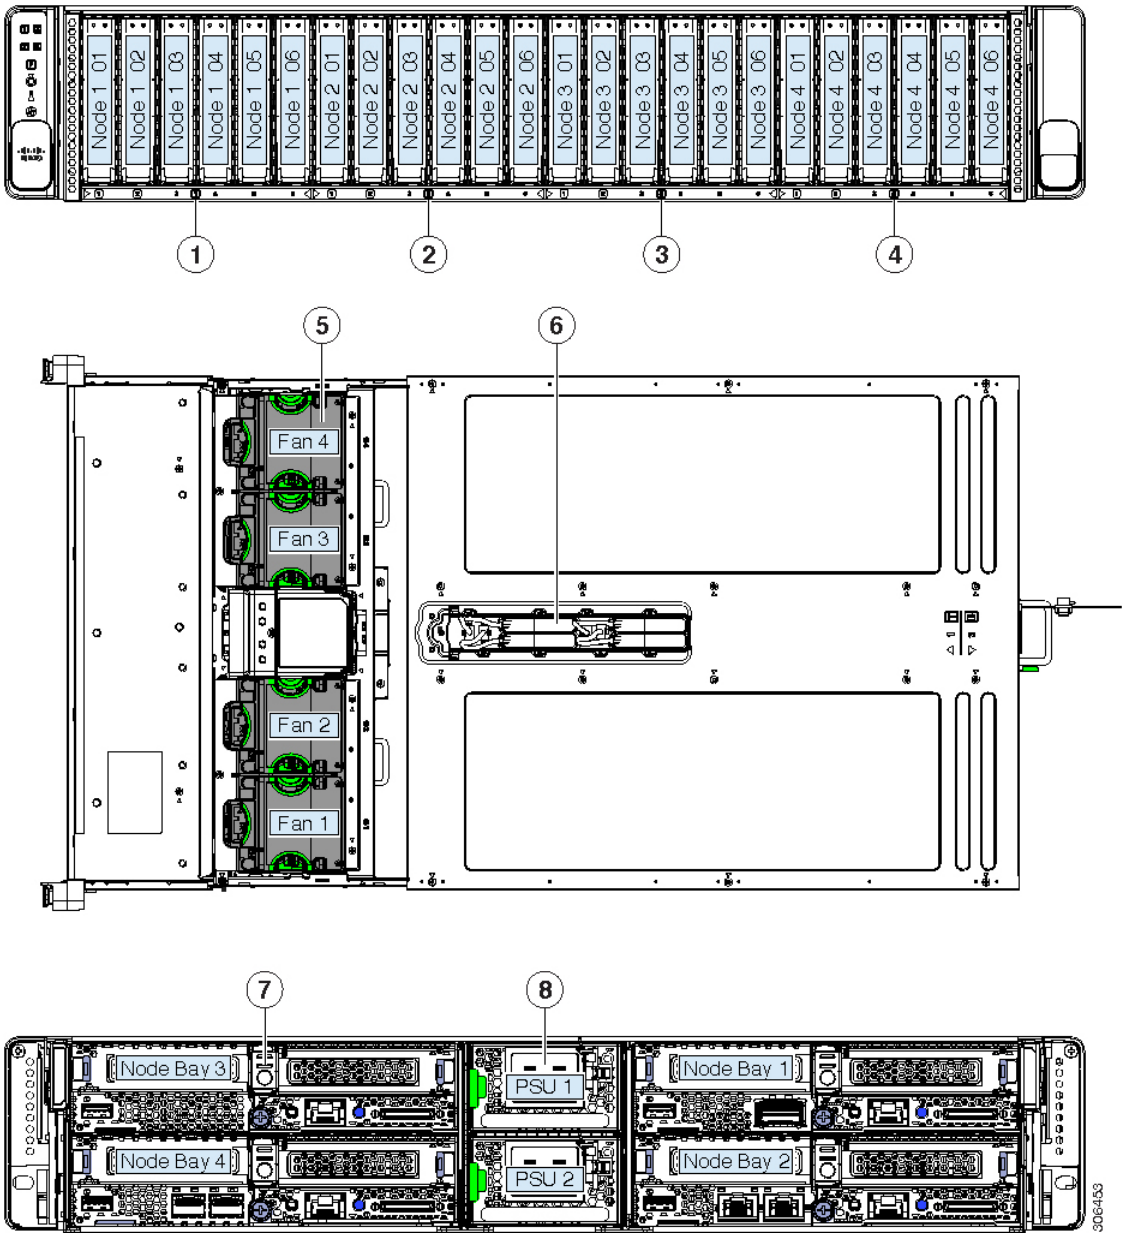

1	Front-loading drives Node 1-controlled drive bays 1—6 All six bays support SAS/SATA drives; bays 1 and 2 also support NVME drives.	**5**	Cooling fan modules (four) Each fan module contains two fans for redundancy.
2	Front-loading drives Node 2-controlled drive bays 1—6 All six bays support SAS/SATA drives; bays 1 and 2 also support NVME drives.	**6**	Supercap units (RAID backup) Each supercap unit backs up one RAID controller in the corresponding node (numbered 1—4).
3	Front-loading drives Node 3-controlled drive bays 1—6 All six bays support SAS/SATA drives; bays 1 and 2 also support NVME drives.	**7**	Compute node (up to four)
4	Front-loading drives Node 4-controlled drive bays 1—6 All six bays support SAS/SATA drives; bays 1 and 2 also support NVME drives.	**8**	Power supplies (two, redundant 1+1)

Summary of Server Features

The following table lists a summary of server features.

Feature	Description
Chassis	Two rack-unit (2RU) chassis
Central Processor	The chassis supports one to four removable compute nodes, each with two CPUs. With four nodes, the system can total up to eight CPUs from the AMD EPYC 7000 Series.
Memory	The chassis supports one to four removable compute nodes, each with two CPUs. Each CPU supports up to eight DIMMs. With four nodes, the system can total up to 64 DIMMs.
Multi-bit error protection	Multi-bit error protection is supported
Baseboard management	Each compute node has a BMC, running Cisco Integrated Management Controller (Cisco IMC) firmware. Depending on your settings, Cisco IMC can be accessed on each node through its 1-Gb dedicated management port or an adapter card.

Feature	Description

Network and management I/O	The network and management I/O ports for this chassis are on the removeable compute nodes. Each compute node has these connectors accessible from the rear of the chassis:
	• One 10/100/1000 Ethernet dedicated management port (RJ-45 connector)
	• One keyboard/video/mouse (KVM) console connector that is used with a KVM cable, which provides two USB 2.0, one DB-15 VGA, and one DB-9 serial connector.
	• One USB 3.0 port
	• Optional OCP adapter-card Ethernet LAN ports. Depending on which adapter is installed, these ports can be:
	  • Dual 10 Gb BASE-T (RJ-45 connectors)
	  • Dual 10/25 Gb (SFP 28 connectors)
	  • Single 100 Gb (QSFP 28 connector)

Power	Two power supplies, redundant as 1+1: • AC power supplies 2400 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 for front-to-rear cooling. Each fan module contains two fans for redundancy.
PCIe I/O	Each removeable compute node has two PCIe risers for horizontal installation of PCIe cards such as a RAID controller or Cisco Virtual Interface Card (VIC).
Storage, front-panel	The chassis can hold up to 24 front-loading, 2.5-inch drives. Each of the four removeable compute nodes can control six of the front drives. All six bays controlled by a compute node support SAS/SATA drives. Bays 1 and 2 of the six controlled by a compute node also support NVMe drives, for a total of up to eight NVMe drives supported in the chassis.

Feature	Description

Storage, internal	Each of the four compute nodes have these internal storage options:
	• Mini-storage module socket, optionally with either:
	  • SD card carrier. Supports up to two SD cards.
	  • M.2 SSD carrier. Supports two SATA M.2 SSDs.
	• One micro-SD card socket.

Storage management	The system has these options via the installed compute nodes (each node can control six of the front-panel drives): • SAS RAID control via one RAID controller card in each compute node. • SATA pass-through JBOD control via the on-board controller in each compute node.
RAID supercap backup	Up to four supercap units are supported, one for the RAID controller card in each node. The supercap units have numbered bays and numbered cable connectors in the top of the chassis, corresponding to each numbered compute node.
Integrated video	Integrated VGA video in each compute node. The DB-15 VGA connector is on the KVM cable that can be used with the KVM connector on each node.

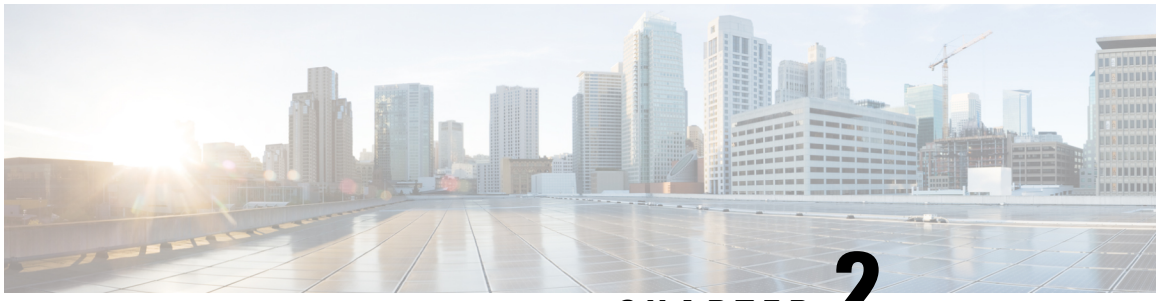

CHAPTER 2

Installation and Initial Setup

This chapter includes rack and installation requirements, installation instructions, and initial setup instructions for connecting to your network.

Installation Warnings and Guidelines

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

Warning

IMPORTANT SAFETY INSTRUCTIONS

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

Statement 1071

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

Statement 1047

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

Statement 1019

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

Statement 1005

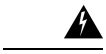

Warning **Installation of the equipment must comply with local and national electrical codes.**

Statement 1074

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

Statement 1017

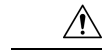

Caution To ensure proper airflow it is necessary to rack the servers using rail kits. Physically placing the units on top of one another or "stacking" without the use of the rail kits blocks the air vents on top of the servers, which could result in overheating, higher fan speeds, and higher power consumption.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

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

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

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

- Ensure that there is adequate space around the server to allow for accessing the server and for adequate airflow. The airflow in this server is from front to back.
- Ensure that the air-conditioning meets the thermal requirements listed in the Environmental Specifications, on page 51.
- Ensure that the cabinet or rack meets the requirements listed in the Rack Requirements, on page 11.
- Ensure that the site power meets the power requirements listed in the Power Specifications, on page 52. If available, you can use an uninterruptible power supply (UPS) to protect against power failures.

Rack Requirements

This section provides the requirements for installing the server chassis in a standard open rack, assuming an external ambient air temperature range of 50 to 95°F (10 to 35°C).

The Cisco R-Series Racks are an ideal choice. If you use other racks, the rack must be of the following type:

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

Supported Cisco Slide Rail Kits

The server supports the following rail kit options:

- Cisco part UCSC-RAILB-C4200= (ball-bearing slide rail kit)
- Cisco part UCSC-CMAF-C4200= (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 Chassis in a Rack

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

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

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

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

Statement 1006

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

a) Align an inner rail with one side of the server so that the four keyed slots in the rail align with the four pegs on the side of the server.

b) Set the keyed slots over the pegs, and then slide the rail toward the front to lock it in place on the pegs.

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

Figure 4: Attaching the Inner Rail to the Side of the Server

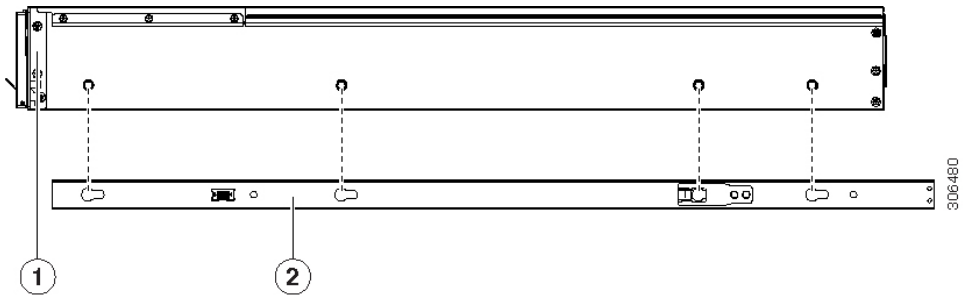

1	Front of server	**2**	Inner rail

Step 2 Install the outer slide rails into the rack:

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

The slide rail front-end wraps around the outside of the rack post and the mounting pegs enter the rack-post holes from the outside-front. See Figure 5: Attaching the Slide Rails to the Rack (View from Inside Rack), on page 13.

b) Push the mounting pegs into the rack-post holes from the outside-front.

c) Adjust the slide-rail length, and then push the rear mounting pegs into the corresponding rear rack-post holes from the *inside* of the rack post. You must push aside the spring-loaded securing clip that wraps around the rear rack post.

Note The slide rail must be level front-to-rear.

Figure 5: Attaching the Slide Rails to the Rack (View from Inside Rack)

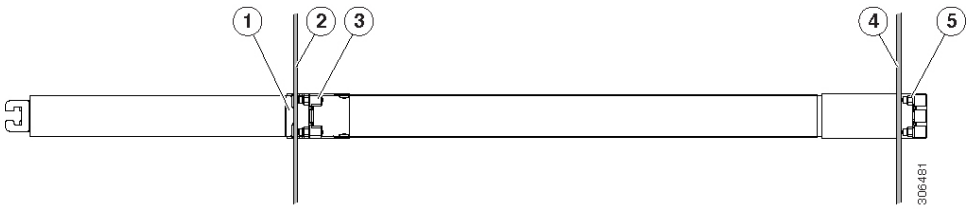

1	Rear peg securing clip, wrapping around rear rack post.	**4**	Front rack post
2	Rear rack post	**5**	Front pegs entering rack post from the outside
3	Rear pegs entering rack post from the inside	-	

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

Step 3 Insert the server into the slide rails:

Caution This system can weigh up to 95.8 pounds (43.5 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. Consider removing compute nodes from the chassis while you lift the chassis.

a) Align the rear ends of the inner rails that are attached to the server sides with the front ends of the empty slide rails on the rack.
b) Push the inner rails into the slide rails on the rack until they stop at the internal stops.
c) Slide the inner-rail release clip toward the rear on both inner rails, and then continue pushing the server into the rack until its front slam-latches engage with the rack posts.

Figure 6: Inner-Rail Release Clip

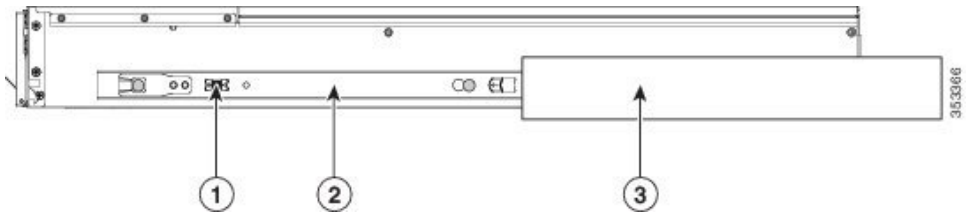

1	Inner-rail release clip	**3**	Outer slide rail attached to rack post
2	Inner rail attached to server and inserted into outer slide rail	-	

Step 4 (Optional) Secure the server in the rack more permanently by using the two screws that are provided with the slide rails. Perform this step if you plan to move the rack with servers installed.

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

Installing the Cable Management Arm (Optional)

Note The cable management arm (CMA) is reversible left-to-right. To reverse the CMA, see Reversing the Cable Management Arm (Optional), on page 15 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 7: Attaching the CMA to the Rear Ends of the Slide Rails

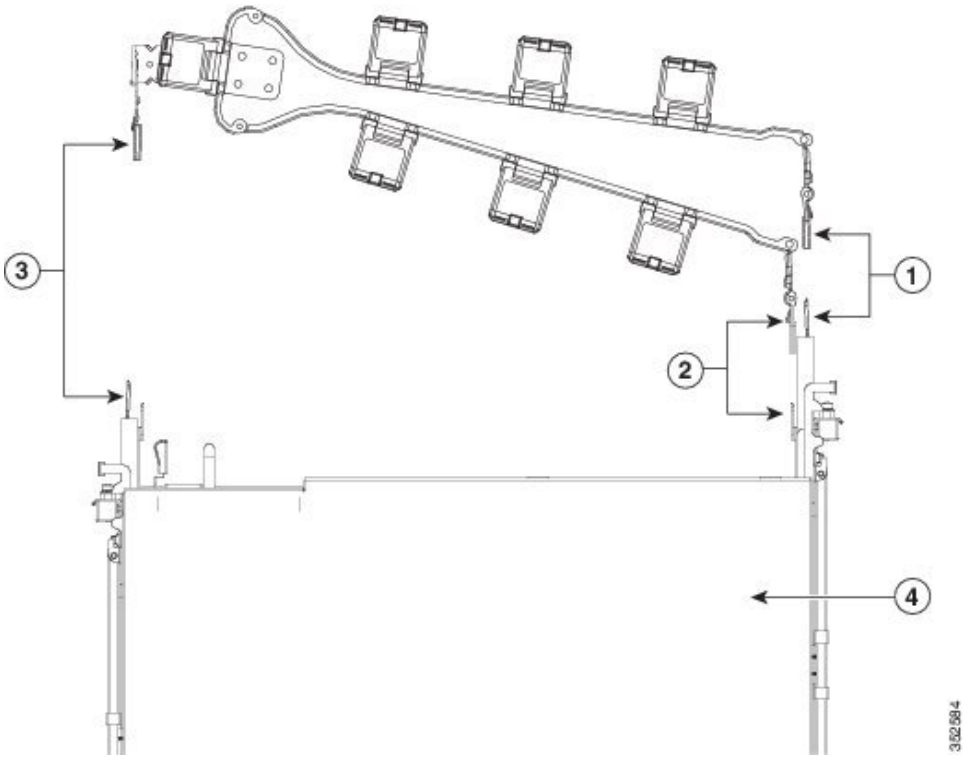

1	CMA tab on arm farthest from server attaches to end of stationary outer slide rail.	**3**	CMA tab on width-adjustment slider attaches to end of stationary outer slide rail.
2	CMA tab on arm closest to the server attaches to end of inner slide rail attached to server.	**4**	Rear of server

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 8: Reversing the CMA

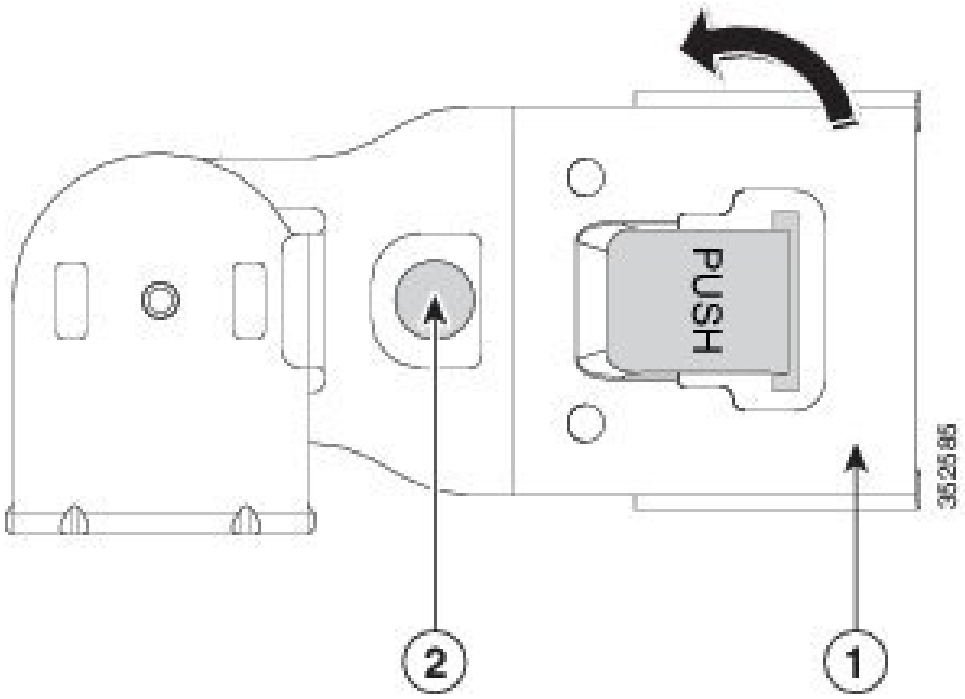

1	CMA tab on end of width-adjustment slider	**2**	Metal button on outside of tab

Initial Compute Node Setup

Note This section describes how to power on the system, assign an IP address to each compute node BMC, and connect to Cisco IMC management when using the system in *standalone* mode. To use the system in Cisco UCS Manager integration (*UCSM mode*), specific cabling and settings are required. See Installation For Cisco UCS Manager Integration, on page 57.

Note All nodes in a C4200 chassis must be either all managed in standalone mode or all managed in UCSM mode. Mixing in a chassis is not supported.

Note Each compute node in your chassis is managed independently and so the following initial setup procedures must be repeated for each node.

Node Default Settings

The node is shipped with these default settings:

- The NIC mode is *Shared LOM EXT*.

  Shared LOM Extended mode enables the Ethernet ports on an installed OCP adapter card *and* the ports on an installed Cisco virtual interface card (VIC) to access the Cisco Integrated Management Interface (Cisco IMC). If you want to use the 10/100/1000 dedicated management ports to access Cisco IMC, you can connect to the compute node and change the NIC mode.

- 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 compute node 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 compute node. The MAC address is printed on a label that is on the pull-out asset tag. This compute node has a range of six MAC addresses assigned to the Cisco IMC. The MAC address printed on the label is the beginning of the range of six contiguous MAC addresses.

Connecting to the Node Locally For Setup

This procedure requires the following equipment:

- VGA monitor
- USB keyboard
- The supported Cisco KVM cable (Cisco PID N20-BKVM).

Step 1 Connect a power cord to each power supply in your chassis, and then connect each power cord to a grounded AC power outlet.

Wait for approximately two minutes to let installed compute nodes boot to standby power during the first bootup. You can verify compute node power status by looking at the node power status LED on the face of the node. The node is in standby power mode when the LED is amber.

Step 2 Connect an optional KVM cable (Cisco PID N20-BKVM) to the KVM connector on the face of the node. Connect your USB keyboard and VGA monitor to the KVM cable.

Step 3 Open the Cisco IMC Configuration Utility:

a) Press and hold the node power button for four seconds to boot the node.

b) During bootup, press **F8** when prompted to open the Cisco IMC Configuration Utility.

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

The following are the requirements for Strong Password:

- The password can have 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 Node With the Cisco IMC Configuration Utility, on page 19.

Connecting to the Node 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 compute node. The MAC address is printed on a label that is on the pull-out asset tag. This compute 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 Connect a power cord to each power supply in your chassis, and then connect each power cord to a grounded AC power outlet.

Wait for approximately two minutes to let the compute node boot to standby power during the first bootup. You can verify node power status by looking at the power status LED on the face of the node. The node is in standby power mode when the LED is amber.

Step 2 Plug your management Ethernet cable into the dedicated management port on the face of the node.

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

Step 4 Use the assigned IP address to access and log in to the Cisco IMC for the 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 Node With the Cisco IMC Configuration Utility, on page 19.

Setting Up the Node With the Cisco IMC Configuration Utility

Before you begin

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

Note Each compute node in your chassis is managed independently and so the following initial setup procedure must be repeated for each node.

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 OCP adapter card LOM ports and the Cisco virtual interface card (VIC) ports. If the system determines that the Cisco VIC connection is not getting its IP address from a Cisco UCS Manager system because the server is in standalone mode, further DHCP requests from the Cisco VIC are disabled.

  If you want to connect to Cisco IMC through a Cisco VIC in standalone mode, it is recommended that you use the *Cisco Card* NIC mode.

- *Shared LOM*—The OCP adapter card LOM 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 on the node 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.

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

- *None*—The Ethernet ports operate independently and do not fail over if there is a problem. This setting can be used only with the Dedicated NIC mode.

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

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

Note 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 compute node. The MAC address is printed on a label that is on the pull-out asset tag. This compute 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.

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 node.

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

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

Step 9 (Optional) Set a default user password.

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

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

Note Auto-negotiation is applicable only when you use the Dedicated NIC mode. Auto-negotiation sets the port speed and duplex mode automatically based on the switch port to which the 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 node in the next step.

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

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

Step 14 Repeat the setup for all compute nodes in your chassis.

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 Rack-Mount Server Configuration Guide* or the *Cisco UCS C-Series Rack-Mount Server CLI Configuration Guide* for instructions on using those interfaces for your Cisco IMC release: Cisco IMC Configuration Guides

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
Shared LOM	Active-active Active-standby
Dedicated	None
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 OCP adapter card LOM ports and the Cisco virtual interface card (VIC) ports. If the system determines that the Cisco VIC connection is not getting its IP address from a Cisco UCS Manager system because the server is in standalone mode, further DHCP requests from the Cisco VIC are disabled.

  If you want to connect to Cisco IMC through a Cisco VIC in standalone mode, it is recommended that you use the *Cisco Card* NIC mode.

- *Shared LOM*—The OCP adapter card LOM 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 on the node 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.

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 tab.

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

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

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

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

Smart Access Serial

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

- This feature has the following requirements when used with compute nodes:
  - A DB-9 serial cable connection, using the KVM cable (Cisco PID N20-BKVM) on the KVM console connector on the face of the compute node.
  - Console redirection must be enabled in the server BIOS.
  - Terminal type must be set to VT100+ or VTUFT8.
  - Serial-over-LAN (SOL) must be disabled (SOL is disabled by default).
- To switch from host serial to Cisco IMC CLI, press **Esc+9**.

  You must enter your Cisco IMC credentials to authenticate the connection.
- To switch from Cisco IMC CLI to host serial, press **Esc+8**.

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

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

Smart Access USB

This server supports the Smart Access USB feature. The board management controller (BMC) in the compute node can accept a USB mass storage device and access the data on it. This feature allows you to use a USB

device as a medium to transfer data between the BMC and the user without need for network connectivity. This can be useful, for example, when remote BMC interfaces are not yet available, or are not accessible due to network misconfiguration.

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

    

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

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

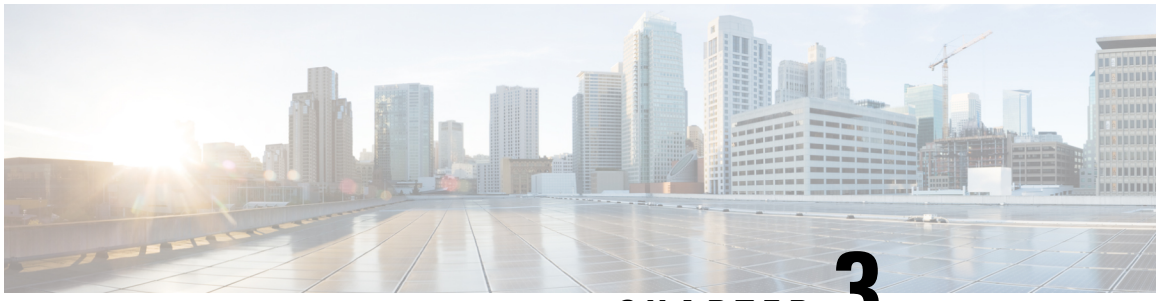

CHAPTER 3

Maintaining the Server Chassis

This chapter contains information about system LEDs and supported component installation or replacement.

Status LEDs and Buttons

This section contains information for interpreting LED states.

Front-Panel LEDs

Figure 9: Front Panel LEDs

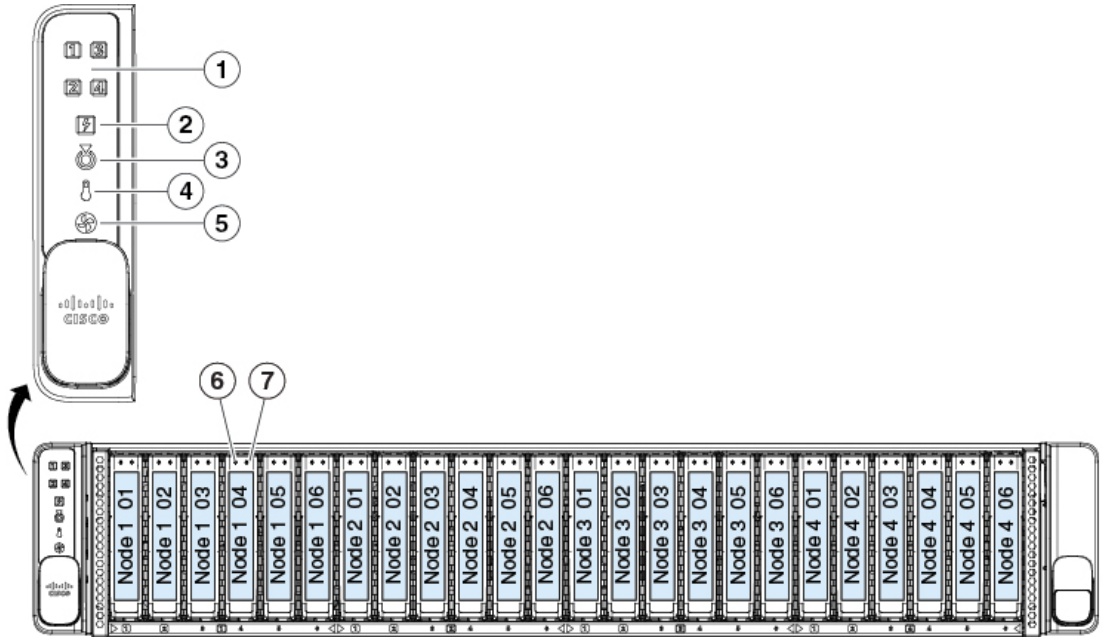

Table 2: Front Panel LEDs, Definition of States

	LED Name	States

1	Node health
	The numbers 1 - 4 correspond to the numbered node bays.
		• Off—No node is detected in the node bay.
		• Green—The node is operating normally.
		• Amber, steady—The node 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 node is in a critical fault state. For example:
		  • Boot failure
		  • Fatal processor and/or bus error detected
		  • Over-temperature condition detected

2	Power supply status	• 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.
3	Locator beacon	Activating the locator beacon on any installed compute node activates this chassis locator beacon. • Off—The unit identification function is not in use. • Blue, blinking—The unit identification function is activated.
4	Temperature status	• Green—The system is operating at normal temperature. • Amber, steady—One or more temperature sensors breached the critical threshold. • Amber, blinking—One or more temperature sensors breached the non-recoverable threshold.

5	Fan status	• Green—All fan modules are operating properly. • Amber, steady—One fan has a fault. • Amber, blinking—Two or more fan modules have a fault.
6 **SAS**	SAS/SATA drive fault **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.
7 **SAS**	SAS/SATA drive activity **Note** NVMe solid state drive (SSD) drive tray LEDs have different behavior than SAS/SATA drive trays.	• 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.
6 **NVMe**	NVMe drive fault	• 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.
7 **NVMe**	NVMe drive activity LED	• Off—No drive activity. • Green, blinking—There is drive activity.

Rear-Panel LEDs

The power supply LEDs are the only rear-panel LEDs native to the chassis. Compute node LEDs that repeat on each compute node are also described below. The rear ports and LEDs vary, depending on which adapter card and PCIe cards are installed.

Figure 10: Rear Panel LEDs

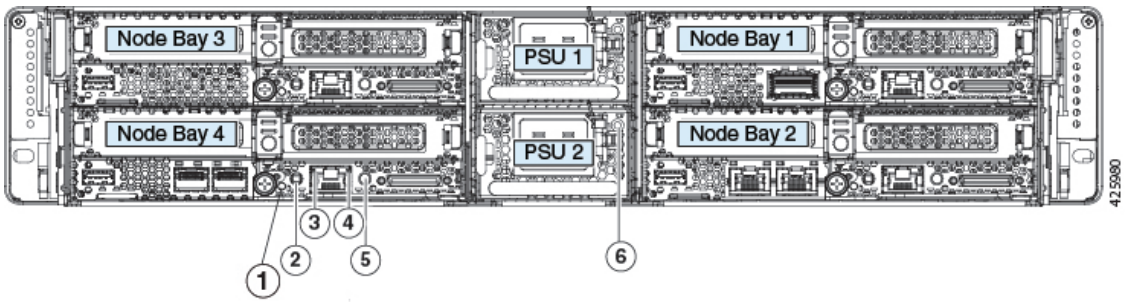

Table 3: Rear Panel LEDs, Definition of States

	LED Name	States

1	Node Health Status	• Green—The node is operating normally.
		• Green, Blinking—The node is in standby power mode.
		• Amber—The node is in a degraded condition (for example, one or more of the following conditions):
		  • Faulty or mismatched CPUs
		  • DIMM failure
		  • Failed drive in a RAID configuration
		• Amber, Blinking—The node is in a critical condition (for example, one or more of the following conditions):
		  • Boot failure
		  • Fatal CPU and/or bus errors detected
		  • Fatal uncorrectable memory errors
		  • Excessive thermal conditions

2	Node Power button/Node Power status (One each node)	• Off—There is no AC power to the node. • Amber—The node is in standby power mode. Power is supplied only to the Cisco IMC and some motherboard functions. • Green—The node is in main power mode. Power is supplied to all node components.
3	Node 1-Gb Ethernet dedicated management link speed (One each node)	• Off—Link speed is 10 Mbps. • Amber—Link speed is 100 Mbps. • Green—Link speed is 1 Gbps.

4	Node 1-Gb Ethernet dedicated management link status (One each node)	• Off—No link is present. • Green—Link is active. • Green, blinking—Traffic is present on the active link.
5	Node locator beacon (One each node)	• Off—The unit identification function is not in use. • Blue, blinking—The unit identification function is activated.
6	Power supply status (One bi-color LED each power supply unit)	**AC power supplies:** • Off—No AC input to any power supplies in the system (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- or under-current, over-voltage, or over-temperature failure). • Amber, solid—If no AC input is supplied to one power supply, the LED lights amber from the shared system standby bus.

Internal Diagnostic LEDs in the Chassis

The fan tray in the chassis includes a fault LED for each of the fan modules. The four LEDs are numbered to correspond to the four numbered fan modules.

Figure 11: Chassis Internal Diagnostic LED Location on Fan Tray

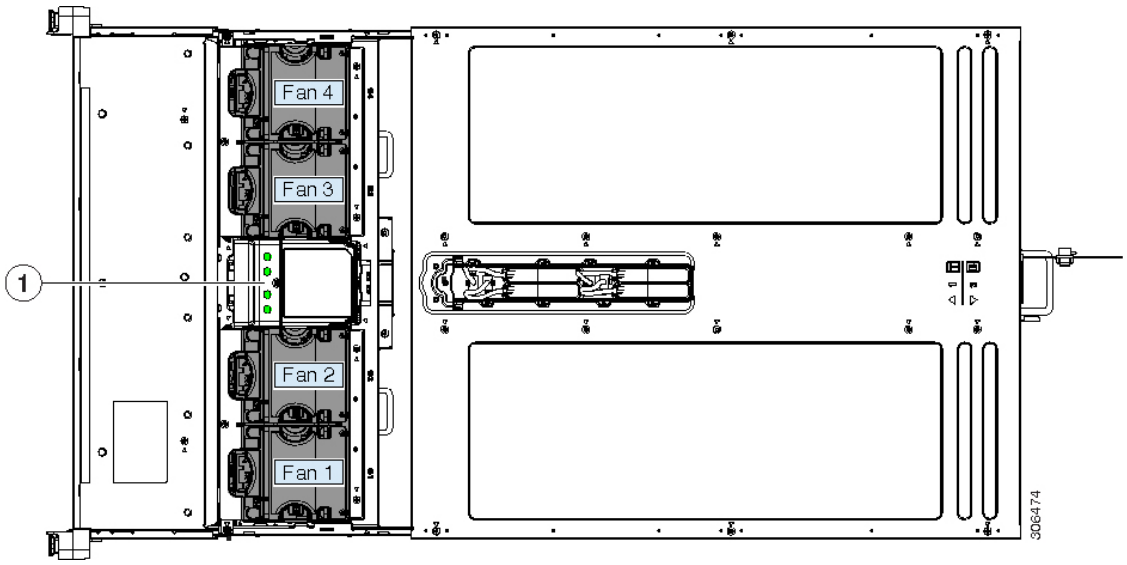

1	Fan module fault LEDs on fan tray (one LED for each fan module) • Green—Fan is OK. • Amber—Fan has a fault or is not fully seated.	-	

Preparing For Component Installation

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

Required Equipment For Service Procedures

The replaceable components in the system chassis require the following tools for removal or installation:

- #1 Phillips-head screwdriver (for opening the supercap compartment cover)
- An electrostatic discharge (ESD) strap or other grounding equipment such as a grounded mat is recommended to protect components.

Shutting Down and Removing Power From the System

Chassis Power

The C4200 system chassis does not include a physical power button. All of the component replacement in the chassis can be performed without removing chassis power (assuming two power supplies with 1+1 redundancy).

Caution The chassis top covers are designed to allow access to replaceable components without exposing the user to high voltages. However, to completely remove power when moving a chassis, you must disconnect all power cords from the power supplies in the chassis.

Compute Node Power

Compute nodes include a physical power button so that you can shut down the individual node and any installed operating system by using the power button on the node or the software interface.

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

The compute node 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 the node from the chassis in this mode.

Shutting Down a Node Using the Power Button

Step 1 Check the color of the Power button/LED on the face of the compute node:

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

Step 2 Invoke a graceful shutdown by pressing and releasing the **Power** button.

Caution To avoid data loss or damage to your operating system, you should always invoke a graceful shutdown of the operating system. Do not power off a node if any firmware or BIOS updates are in progress.

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

As a best practice, attempt a graceful shutdown first. As an option, you can also invoke an Emergency shutdown, by pressing and holding the **Power** button for 4 seconds to force the main power off and immediately enter standby mode.

Shutting Down a Node 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 **Chassis** tab.

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

Step 3 In the toolbar above the work pane, click the **Host Power** link.

The Server Power Management dialog opens. This dialog lists all servers that are present in the system.

Step 4 In the **Server Power Management** dialog, select Shut Down for the server that you want to shut down. Shut Down performs a graceful shutdown of the operating system.

Note To avoid data loss or damage to your operating system, you should always invoke a graceful shutdown of the operating system. Do not power off a node if any firmware or BIOS updates are in progress.

- It is safe to remove the node from the chassis when the Chassis Status pane shows the Power State as Off for the node that you are removing.
- The physical power button on the node also turns amber when it is safe to remove the node from the chassis.
- The **Server Power Management** dialog also has a Power Off option, but you should use Shut down as a best practice. The Power Off option is a forced shutdown that powers off the chosen node even if tasks are running on that server. Use Power Off only if the Shut Down option does not complete successfully.

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:

/chassis# **power shutdown**

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

Shutting Down a Node 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** > **Chassis** > ***Chassis Number*** > **Servers**.

Step 3 Choose the node that you want to shut down.

Step 4 In the Work pane, click the **General** tab.

Step 5 In the Actions area, click **Shutdown Server**.

Step 6 If a confirmation dialog displays, click **Yes**.

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

Shutting Down a Node 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 associated service profile.

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

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

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

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

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

Opening the Chassis Compartment Covers

The server chassis has been designed so that only small compartment covers are opened to allow access to replaceable components (cooling fans and supercap units for RAID backup).

Caution Never remove the overall chassis cover; open only the compartment covers. The overall cover protects the user from exposure to hazardous voltages that are present in the chasssis.

Figure 12: Fan and Supercap Compartment Covers

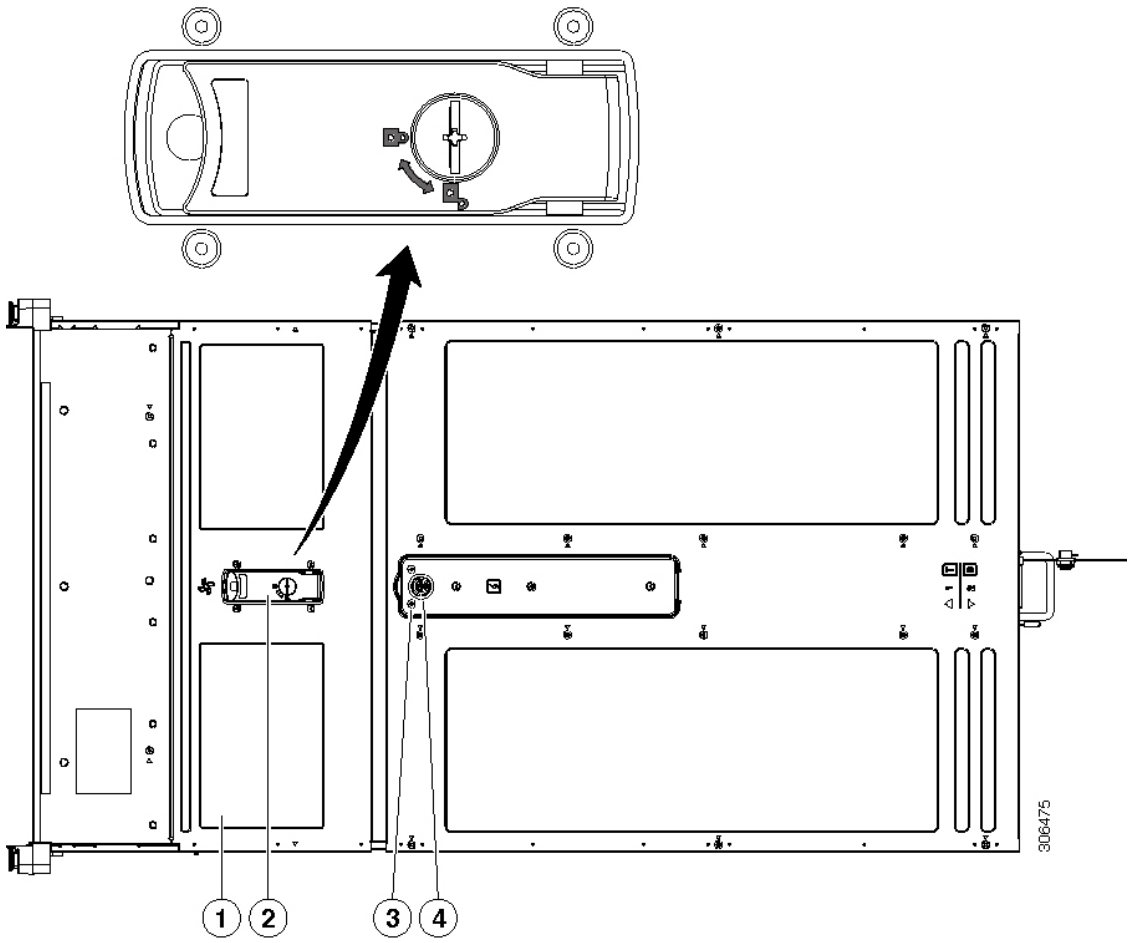

1	Fan compartment cover	**3**	Supercap compartment cover
2	Fan compartment cover latch and lock	**4**	Supercap compartment cover securing screw

Opening the Fan Compartment Cover

Step 1 Open the hinged cover:

a) If the cover latch is locked, use a screwdriver to turn the lock 90-degrees counterclockwise to unlock it.
b) Lift on the end of the latch that has the green finger grip. The cover is pushed back as you lift the latch.
c) Open the hinged cover.

Step 2 Close the hinged cover:

a) With the latch in the fully open position, close the hinged cover.
b) Press the cover latch down to the closed position. The cover is pushed forward.

c) If desired, lock the latch by using a screwdriver to turn the lock 90-degrees clockwise.

Opening the Supercap Compartment Cover

Step 1 Open the supercap compartment cover:

a) Use a #2 Phillips-head screwdriver to loosen the single captive screw on the cover.

b) Lift on the end of the cover next to the captive screw and then completely remove the cover from the chassis.

Step 2 Replace the cover:

a) Set the cover in place. The end with the captive screw should be toward the chassis front.

b) Tighten the single captive screw on the cover.

Removing and Replacing Components

This section describes how to install or replace components in the Cisco UCS C4200 Server Chassis. For information on replacing components inside an installed compute node, see the service note for your compute node:

- Cisco UCS C125 Compute Node Service Note

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.

Serviceable Components in the Chassis

The figure in this topic shows the locations of the serviceable components in the chassis.

For components inside a compute node, see the service note for your compute node:

- Cisco UCS C125 Compute Node Service Note

Figure 13: Cisco UCS C4200 Chassis Serviceable Component Locations

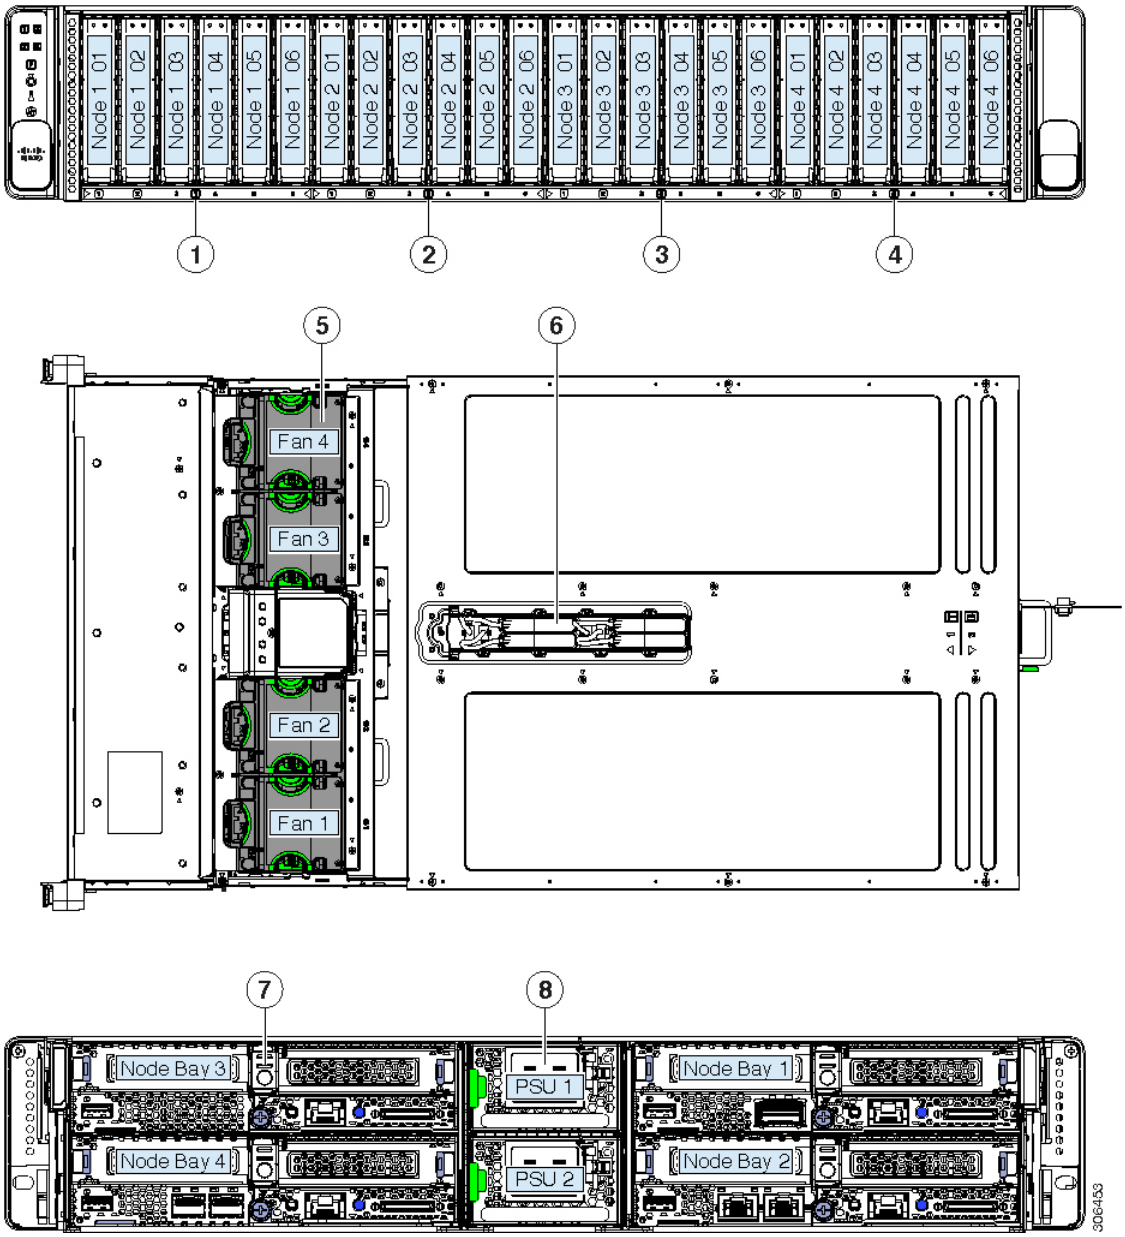

1	Front-loading drives Node 1-controlled drive bays 1—6 All six bays support SAS/SATA drives; bays 1 and 2 also support NVME drives.	**5**	Cooling fan modules (four) Each fan module contains two fans for redundancy.

2	Front-loading drives Node 2-controlled drive bays 1—6 All six bays support SAS/SATA drives; bays 1 and 2 also support NVME drives.	**6**	Supercap units (RAID backup) Each supercap unit backs up one RAID controller in the corresponding node (numbered 1—4).
3	Front-loading drives Node 3-controlled drive bays 1—6 All six bays support SAS/SATA drives; bays 1 and 2 also support NVME drives.	**7**	Compute node (up to four)
4	Front-loading drives Node 4-controlled drive bays 1—6 All six bays support SAS/SATA drives; bays 1 and 2 also support NVME drives.	**8**	Power supplies (two, redundant 1+1)

Replacing Front-Loading SAS/SATA Drives

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

SAS/SATA Drive Population Guidelines

The chassis can hold up to 24 front-loading, 2.5-inch drives. Each installed compute node can control the six drives that correspond to the node number in the chassis.

- The four compute node groups are marked on the bottom lip of the chassis (below the drives).
- In each of the four compute node groups, the drives are enumerated 1- 6.
- In each of the compute node groups, populate the lowest numbered bays first.
- Drives installed in front-panel bays that do not have a corresponding compute node are not seen by the system.
- 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.

Figure 14: Drive Bay Numbering

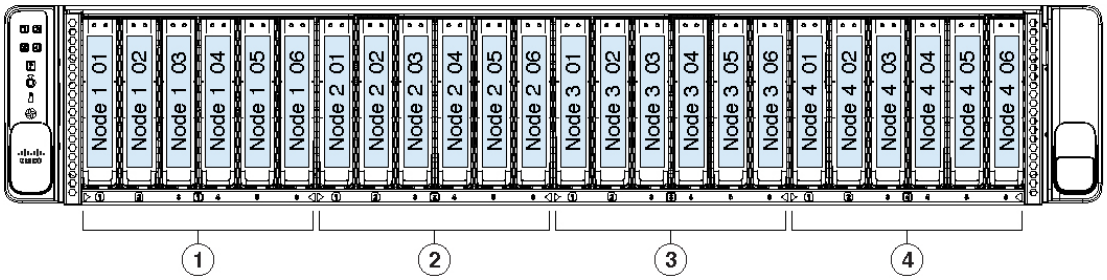

1	Drive bays controlled by compute node 1	**3**	Drive bays controlled by compute node 3
2	Drive bays controlled by compute node 2	**4**	Drive bays controlled by compute node 4

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 Boot the compute node and 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 compute node to log into the Cisco IMC GUI management interface.
Step 2 Navigate to **Server** > **BIOS**.
Step 3 Under Actions, click **Configure BIOS**.

Step 4 In the Configure BIOS Parameters dialog, select the **Advanced** tab.

Step 5 Go to the **LOM** and **PCIe Slot Configuration** section.

Step 6 Set the **PCIe Slot: HBA Option ROM** to **UEFI Only**.

Step 7 Click **Save Changes**. The dialog closes.

Step 8 Under BIOS Properties, set **Configured Boot Order** to **UEFI**.

Step 9 Under Actions, click **Configure Boot Order**.

Step 10 In the Configure Boot Order dialog, click **Add Local HDD**.

Step 11 In the Add Local HDD dialog, enter the information for the 4K sector format drive and make it first in the boot order.

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

Replacing a Front-Loading SAS/SATA Drive

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

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

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

Step 2 Install a new drive:

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

Figure 15: Replacing a Drive in a Drive Tray

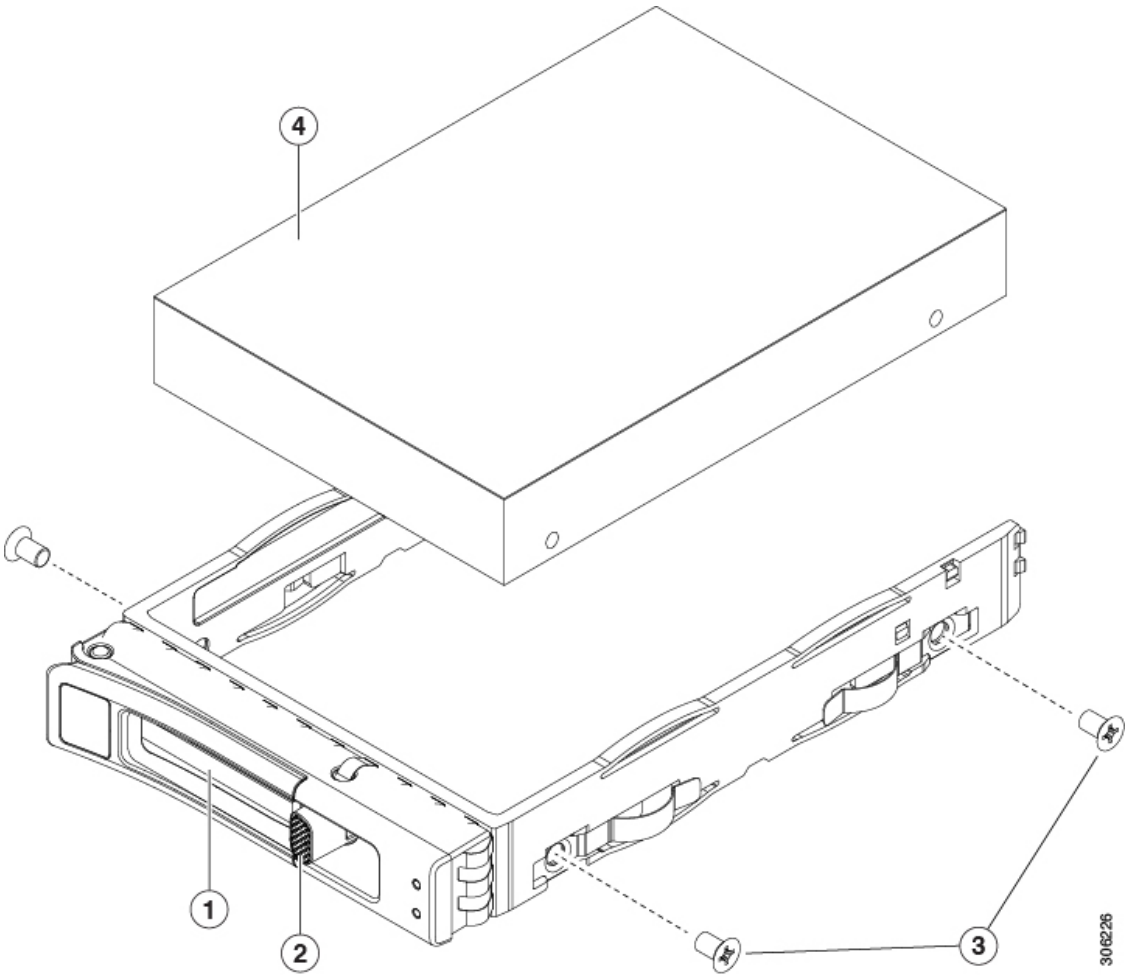

1	Ejector lever	**3**	Drive tray screws (two on each side)
2	Release button	**4**	Drive removed from drive tray

Replacing Front-Loading NVMe SSDs

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

Caution NVMe drives are not hot-swappable. You can replace them while the system is running, but you must shut down the drive in the software or OS before removal.

Note NVMe drives are supported only with Cisco IMC/BIOS 4.0(2) and later; they are supported with Cisco UCS Manager 4.0(2) and later.

Front-Loading NVMe SSD Population Guidelines

Each compute node in the chassis controls six drive bays. Drive bays 1 and 2 in each set of six bays support NVMe SSDs. In the following figure only the drive bays outlined in red support NVMe drives.

- The four compute node groups are marked on the bottom lip of the chassis (below the drives). Drive bays 1 and 2 in each group are marked with the drive bay numbers in squares to indicate that those bays support NVMe drives.
- In each of the compute node groups, populate the lowest numbered bays first.
- Drives installed in front-panel bays that do not have a corresponding compute node are not seen by the system.
- Keep an empty drive blanking tray in any unused bays to ensure proper airflow.
- In the following figure only the drive bays outlined in red support NVMe drives.

Figure 16: Drive Bay Numbering and NVMe Drive Support

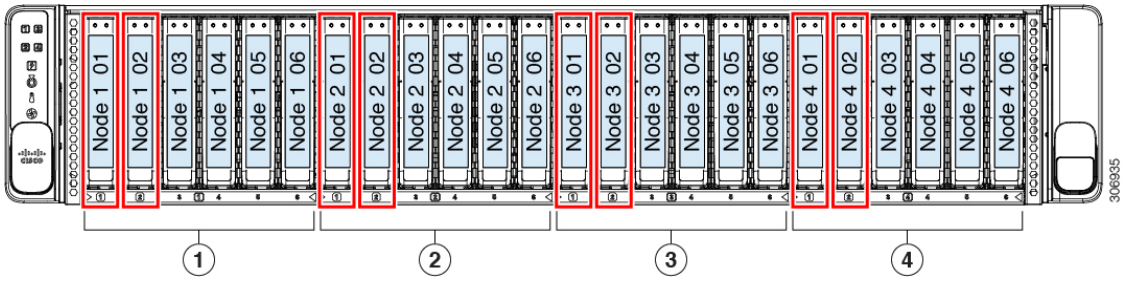

1	Drive bays controlled by compute node 1	**3**	Drive bays controlled by compute node 3
2	Drive bays controlled by compute node 2	**4**	Drive bays controlled by compute node 4

Front-Loading NVME SSD Requirements and Restrictions

Caution NVMe drives are not hot-swappable. You can replace them while the system is running, but you must shut down the drive in the software or OS before removal.

Observe these requirements and restrictions:

- Cisco IMC/BIOS 4.0(2) and later; if using Cisco UCS Manager, release 4.0(2) and later.
- Informed hot plug (safe hot plug) is supported. This functionality is enabled by default. However, uninformed hot plug (surprise hot removal) is not supported.
- NVMe SSDs support booting only in UEFI mode. Legacy boot is not supported.

- You cannot control NVMe SSDs with a SAS RAID controller because NVMe SSDs interface with the server via the PCIe bus.

Replacing a Front-Loading NVMe SSD

This topic describes how to replace 2.5-inch 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.

Step 1 Remove an existing front-loading NVMe SSD:

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

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

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

Step 2 Install a new front-loading NVMe SSD:

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

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

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

Figure 17: Replacing a Drive in a Drive Tray

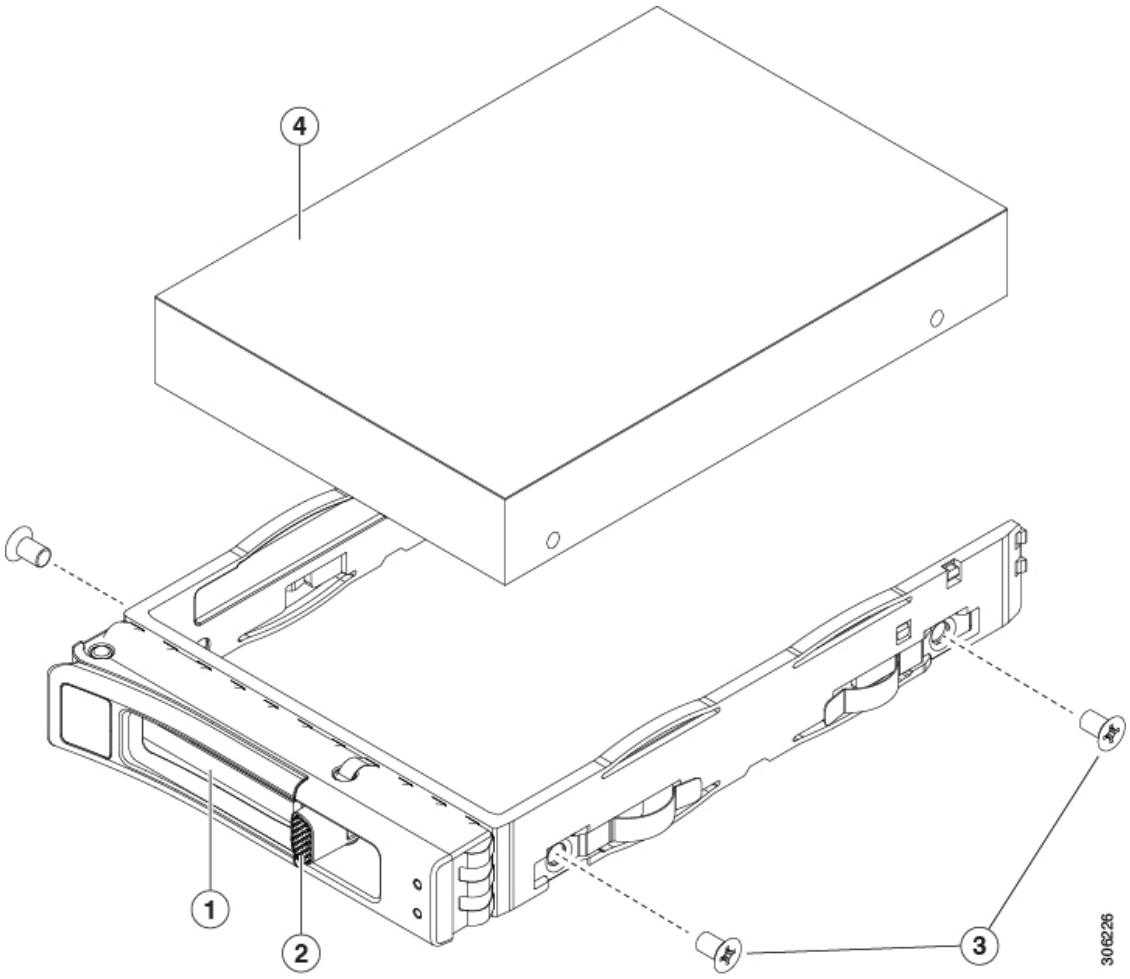

1	Ejector lever	**3**	Drive tray screws (two on each side)
2	Release button	**4**	Drive removed from drive tray

Replacing Fan Modules

Tip There are four fault LEDs on the fan tray, each numbered to a corresponding fan module. These LEDs light green when the fan is correctly seated and is operating OK. The LED lights amber when the fan has a fault or is not correctly seated.

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

Step 1 Slide the server out the front of the rack far enough so that you can open the fan compartment cover on top of the chassis. 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 2 Open the hinged cover:

a) If the cover latch is locked, use a screwdriver to turn the lock 90-degrees counterclockwise to unlock it.
b) Lift on the end of the latch that has the green finger grip. The cover is pushed back as you lift the latch.
c) Open the hinged cover.

Step 3 Grasp and squeeze the fan module release latches on its top. Lift straight up to disengage its connector from the motherboard.

Step 4 Set the new fan module in place.

Note The arrows printed on the top of the fan module must point toward the rear of the server.

Step 5 Press down gently on the fan module to fully engage it with the connector on the motherboard.

Step 6 Close the hinged cover:

a) With the latch in the fully open position, close the hinged cover.
b) Press the cover latch down to the closed position. The cover is pushed forward.
c) If desired, lock the latch by using a screwdriver to turn the lock 90-degrees clockwise.

Figure 18: Top View of Fan Modules

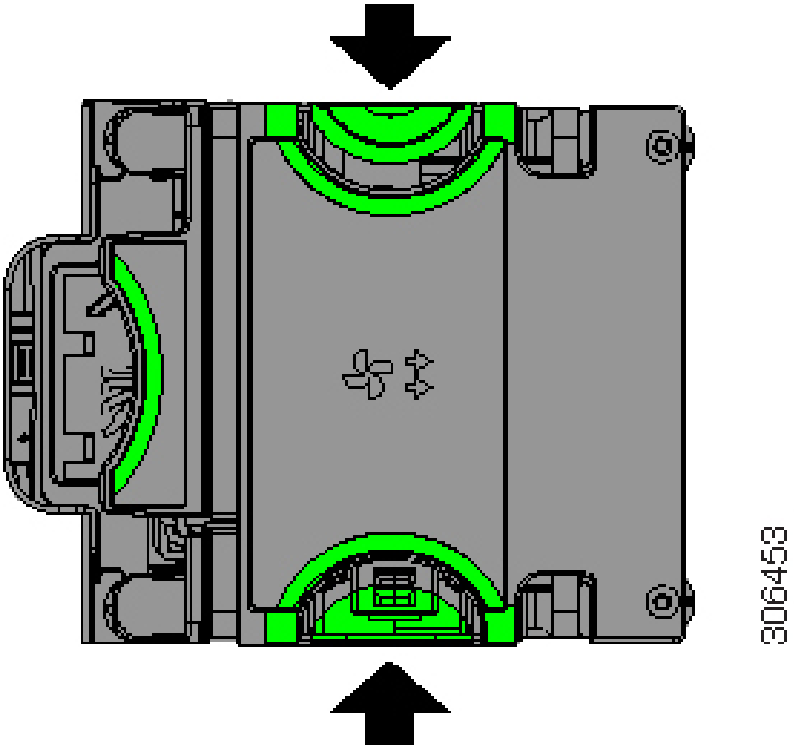

1	Fan module release latches	**2**	Fan module fault LEDs

Step 7 Return the chassis to the rack and replace any cables that you removed.

Replacing the Supercap (RAID Backup)

This chassis supports installation of up to four supercap units, one for each installed compute node. The units install to numbered bays and the supercap cables connect to numbered sockets in the supercap compartment.

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

Step 1 Prepare the server for component installation:

a) Shut down the compute node that corresponds to the supercap unit that you are replacing as described in Shutting Down and Removing Power From the System, on page 30.

b) Slide the server out the front of the rack far enough so that you can open the supercap compartment 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 2 Open the supercap compartment cover:

a) Use a #1 Phillips-head screwdriver to loosen the single captive screw on the cover.

b) Lift on the end of the cover next to the captive screw and then completely remove the cover from the chassis.

Step 3 Remove an existing supercap:

a) Disconnect the supercap cable from the existing supercap.

b) Pull straight up on the supercap unit and set it aside.

Step 4 Install a new supercap:

a) Set the new supercap into the empty bay.

b) Connect the supercap cable to the connector on the supercap cable.

Note Make sure that the number of the supercap unit bay, the number of the cable connector, and the number of the compute node match. A supercap cabled to a connector for an absent compute node is not seen by the system.

Step 5 Replace the supercap compartment cover:

a) Set the cover in place. The end with the captive screw should be toward the chassis front.

b) Tighten the single captive screw on the cover.

Step 6 Replace the chassis in the rack, replace cables, and then fully power on the compute node.

Figure 19: Supercap Bays and Cable Connectors (Compartment Cover Removed)

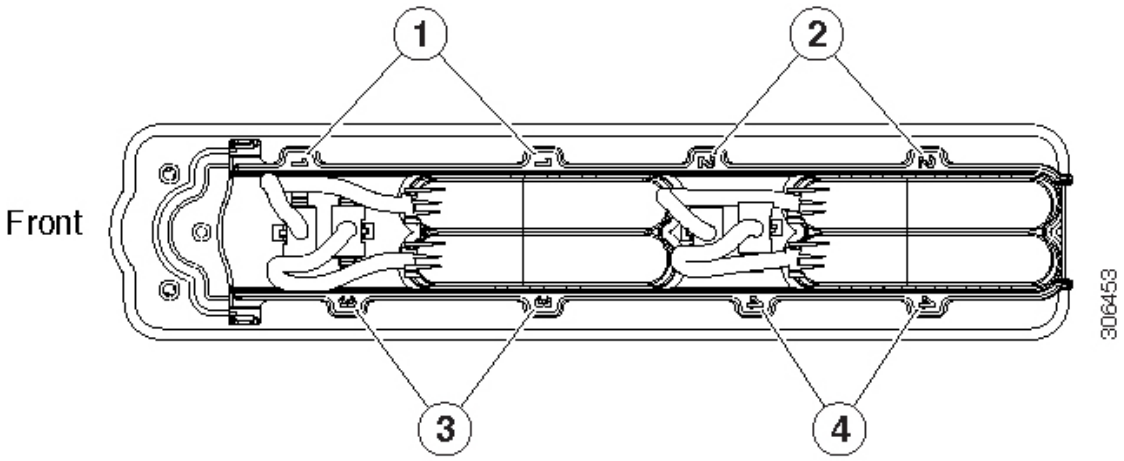

1	Supercap bay and supercap cable connector for node 1	**3**	Supercap bay and supercap cable connector for node 3
2	Supercap bay and supercap cable connector for node 2	**4**	Supercap bay and supercap cable connector for node 4

Replacing Power Supplies

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

- See also Power Specifications, on page 52 for more information about the supported power supplies.

- See also Rear-Panel LEDs, on page 27 for information about the power supply LEDs.

This section includes procedures for replacing AC power supply units.

Replacing AC Power Supplies

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

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

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

a) Perform one of the following actions:

- If your chassis has only one power supply, shut down all installed compute nodes as described in Shutting Down and Removing Power From the System, on page 30.
- If your chassis has two power supplies, you do not have to shut down the compute nodes.

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

Step 2 Install a new power supply:

a) Grasp the power supply handle and insert the new power supply into the empty bay.
b) Push the power supply into the bay until the release lever locks.
c) Connect the power cord to the new power supply.
d) Only if you shut down the compute nodes, reboot each to main power mode.

Figure 20: Replacing AC Power Supplies

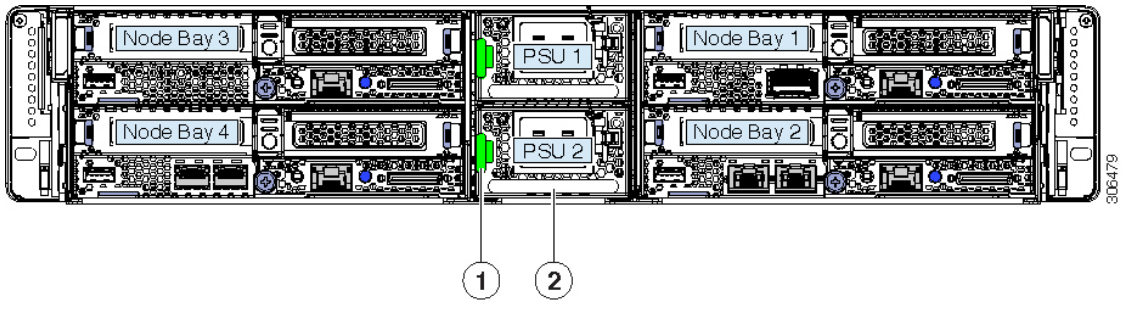

1	Power supply release lever	**2**	Power supply handle

Replacing Compute Nodes and Node Components

Information and procedures for replacing compute nodes and the components inside nodes are provided in separate node service notes, as supported.

Caution Always shut down the node before removing it from the chassis, as described in the procedures. Failure to shut down the node before removal results in the corresponding RAID supercap cache being discarded and other data might be lost.

Replacing a Compute Node

To replace a compute node, including exporting and importing its configuration, see the service note for your compute node:

- Cisco UCS C125 Compute Node Service Note

Replacing DIMMs Inside a Compute Node

For information about replacing memory DIMMs inside a compute node, including supported memory population, see the service note for your compute node:

- Cisco UCS C125 Compute Node Service Note

Replacing CPUs and Heatsinks Inside a Compute Node

For information about replacing AMD CPUs and their heatsinks inside a compute node, see the service note for your compute node:

- Cisco UCS C125 Compute Node Service Note

Installing a Trusted Platform Module (TPM) Inside a Compute Node

For information about installing a trusted platform module (TPM) inside a compute node, see the service note for your compute node:

- Cisco UCS C125 Compute Node Service Note

Replacing an RTC Battery Inside a Compute Node

For information about replacing a real-time clock (RTC) battery inside a compute node, see the service note for your compute node:

- Cisco UCS C125 Compute Node Service Note

Replacing Mini-Storage (SD or M.2) Inside a Compute Node

For information about replacing a mini-storage carrier with SD cards or M.2 SATA drives inside a compute node, see the service note for your compute node:

- Cisco UCS C125 Compute Node Service Note

Replacing a Micro-SD Card Inside a Compute Node

For information about replacing a Micro-SD card inside a compute node, see the service note for your compute node:

- Cisco UCS C125 Compute Node Service Note

Replacing an OCP Adapter Card Inside a Compute Node

For information about replacing an OCP adapter card inside a compute node, see the service note for your compute node:

- Cisco UCS C125 Compute Node Service Note

Replacing a PCIe Riser Inside a Compute Node

For information about PCIe slots and replacing a PCIe riser inside a compute node, see the service note for your compute node:

- Cisco UCS C125 Compute Node Service Note

Replacing a PCIe Card Inside a Compute Node

For information about replacing a PCIe card inside a compute node, see the service note for your compute node:

- Cisco UCS C125 Compute Node Service Note

Replacing a Storage Controller Inside a Compute Node

For information about supported storage controllers and replacing a controller card inside a compute node, see the service note for your compute node:

- Cisco UCS C125 Compute Node Service Note

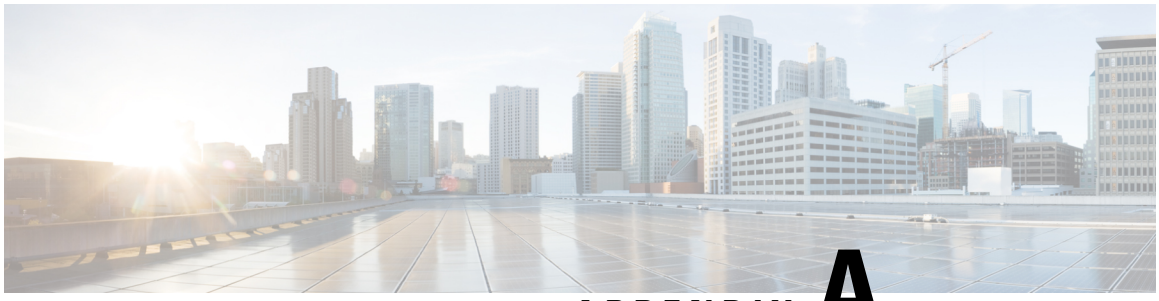

APPENDIX A

System Specifications

C4200 Server Chassis Specifications

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

Physical Specifications

The following table lists the physical specifications for the C4200 server chassis.

Table 4: Physical Specifications

Description	Specification
Height	3.4 in. (87.2 mm)
Width	16.9 in. (429.3 mm) (width of the chassis rail-receiving surface; not including front latch-handles)
Depth (length)	Server only: 32.6 in. (827.6 mm) Server with Cisco slide rails and CMA: 38.0 in (965.2 mm)
Maximum weight (fully loaded chassis with four compute nodes)	95.8 lb. (43.5 Kg)

Environmental Specifications

The following table lists the environmental requirements for the C4200 server chassis.

Table 5: 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
Altitude, non-operating (when the server is stored or transported)	0 to 40,000 feet
Sound power level Measure A-weighted per ISO7779 LwAd (Bels) Operation at 73°F (23°C)	5.5
Sound pressure level Measure A-weighted per ISO7779 LpAm (dBA) Operation at 73°F (23°C)	40

Power Specifications

Note Do not mix power supply types or wattages in the 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.

2400 W AC Power Supply

This section lists the specifications for each 2400 W AC power supply (Cisco part number UCSC-PSU3-2400W).

Table 6: 2400 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	16 A at 200 VAC
Maximum inrush current	35 A at 35° C
Maximum output power per PSU	2400 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	IEC60320 C20

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 7: Supported Power Cords

Description	Length (Feet)	Length (Meters)
Argentina CAB-IR2073-C19-AR AC power cord, 16 A, 250 VAC Plugs: IRSM 2073 to IEC 60320 C19	14	4.26

Australia and New Zealand CAB-AC-16A-AUS AC power cord, 16 A, 250 VAC Plugs: AU20S3 to IEC 60320 C19	14	4.26
Brazil CAB-9K16A-BRZ AC power cord, 16 A, 250 VAC Plugs: EL224 to IEC 60320 C19	14	4.26
Europe, Continental CAB-AC-2500W-EU AC power cord, 16 A, 250 VAC Plugs: CEE 7/7 to IEC 60320 C19	14	4.26
India CAB-SABS-C19-IND AC power cord, 16 A, 250 VAC Plugs: SABS 164-1 to IEC 60320 C19	14	4.26
International CAB-AC-2500W-INT AC power cord, 16 A, 250 VAC Plugs: IEC 309 to IEC 60320 C19	14	4.26
Israel CAB-AC-2500W-ISRL AC power cord, 16 A, 250 VAC Plugs: SI16S3 to IEC 60320 C19	14	4.26
Israel CAB-S132-C19 AC power cord, 16 A, 250 VAC Plugs: S132 to IEC 60320 C19	14	4.26
Italy CAB-C2316-C19-IT AC power cord, 16 A, 250 VAC Plugs: CEI 23-16 to IEC 60320 C19	14	4.26

Japan CAB-C19-C20-3M-JP AC power cord, 16 A, 250 VAC Plugs: C20 to IEC 60320 C19 Includes Japan PSE mark	10	3
Korea CAB-9K16A-KOR AC power cord, 16 A, 250 VAC Plugs: Src plug to IEC 60320 C19	14	4.26
North America and Japan **Non-locking** CAB-AC-2500W-US1 AC power cord, 16 A, 250 VAC Plugs: NEMA 6-20 to IEC 60320 C19	14	4.26
North America and Japan **Locking** CAB-AC-C6K-TWLK AC power cord, 16 A, 250 VAC Plugs: NEMA L6-20 to IEC 60320 C19	14	4.26
Peoples Republic of China CAB-AC-16A-CH AC power cord, 16 A, 250 VAC Plugs: NGB16C to IEC 60320 C19	14	4.26
Taiwan CAB-AC-C19-TW AC power cord, 16 A, 250 VAC Plugs: TW1-15P to IEC 60320 C19	7.5	2.3
Switzerland CAB-ACS-16 AC power cord, 16 A, 250 VAC Plugs: SEV 5934-2, Type 23 to IEC 60320 C19	8	2.5

United Kingdom CAB-BS1363-C19-UK AC power cord, 16 A, 250 VAC Plugs: BS-1363 to IEC 60320 C19	14	4.26
United States CAB-L520P-C19-US AC power cord, 16 A, 250 VAC Plugs: NEMA L5-20 to IEC 60320 C19	6	1.8
United States CAB-US515-C19-US AC power cord, 16 A, 250 VAC Plugs: NEMA 5-15 to IEC 60320 C19	13	4
United States CAB-US520-C19-US AC power cord, 16 A, 250 VAC Plugs: NEMA 5-20 to IEC 60320 C19	14	4.26
United States CAB-US620-C19-US AC power cord, 16 A, 250 VAC Plugs: NEMA 6-20 to IEC 60320 C19	13	4
Cabinet Jumper Power Cord CAB-C19-CBN AC power cord, 16 A, 250 VAC Plugs: IEC 60320 C20 to IEC 60320 C19	9	2.7
No power cord option R2XX-DMYMPWRCORD	NA	NA

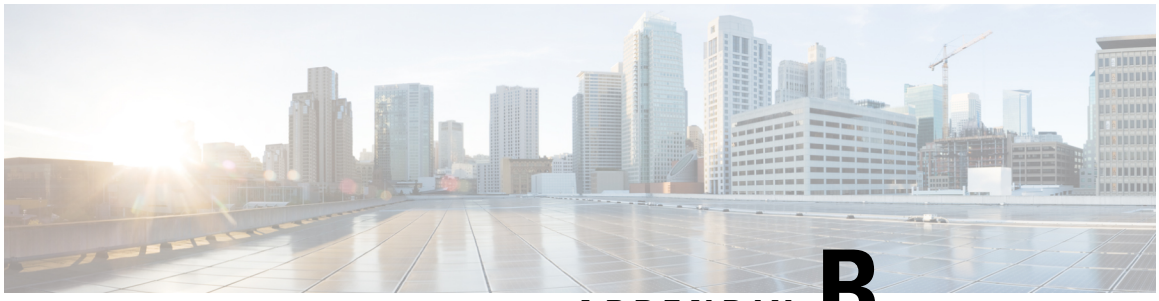

APPENDIX B

Installation For Cisco UCS Manager Integration

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.

Note All nodes in a C4200 chassis must be either all managed in standalone mode or all managed in UCSM mode. Mixing in a chassis is not supported.

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

Installation For Cisco UCS Manager Integration