



Maintaining the Server Chassis

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

- [Status LEDs and Buttons](#), on page 1
- [Preparing For Component Installation](#), on page 6
- [Opening the Chassis Compartment Covers](#), on page 9
- [Removing and Replacing Components](#), on page 11

Status LEDs and Buttons

This section contains information for interpreting LED states.

Front-Panel LEDs

Figure 1: Front Panel LEDs

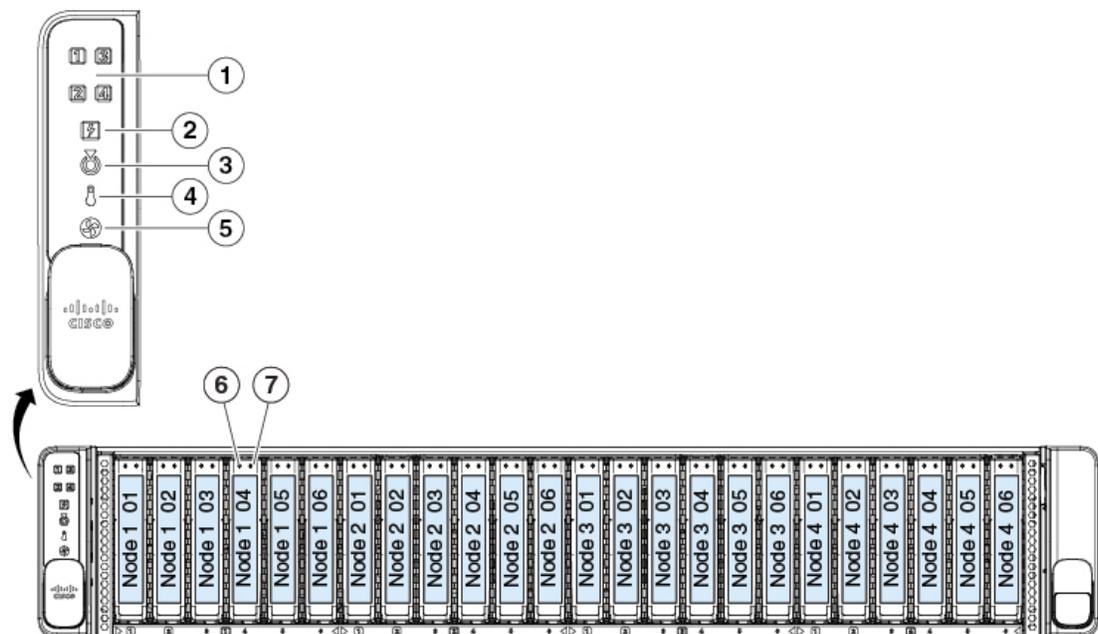


Table 1: Front Panel LEDs, Definition of States

	LED Name	States
1	Node health The numbers 1 - 4 correspond to the numbered node bays.	<ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • Power supply redundancy is lost. • CPUs are mismatched. • At least one CPU is faulty. • At least one DIMM is faulty. • At least one drive in a RAID configuration failed. • Amber, blinking—The node is in a critical fault state. For example: <ul style="list-style-type: none"> • Boot failure • Fatal processor and/or bus error detected • Over-temperature condition detected
2	Power supply status	<ul style="list-style-type: none"> • Green—All power supplies are operating normally. • Amber, steady—One or more power supplies are in a degraded operational state. • Amber, blinking—One or more power supplies are in a critical fault state.
3	Locator beacon	<p>Activating the locator beacon on any installed compute node activates this chassis locator beacon.</p> <ul style="list-style-type: none"> • Off—The unit identification function is not in use. • Blue, blinking—The unit identification function is activated.
4	Temperature status	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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	<p>SAS/SATA drive fault</p> <p>Note NVMe solid state drive (SSD) drive tray LEDs have different behavior than SAS/SATA drive trays.</p>	<ul style="list-style-type: none"> • Off—The hard drive is operating properly. • Amber—Drive fault detected. • Amber, blinking—The device is rebuilding. • Amber, blinking with one-second interval—Drive locate function activated in the software.
7 SAS	<p>SAS/SATA drive activity</p> <p>Note NVMe solid state drive (SSD) drive tray LEDs have different behavior than SAS/SATA drive trays.</p>	<ul style="list-style-type: none"> • Off—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	<ul style="list-style-type: none"> • Off—The drive is not in use and can be safely removed. • Green—The drive is in use and functioning properly. • Green, blinking—the driver is initializing following insertion or the driver is unloading following an eject command. • Amber—The drive has failed. • Amber, blinking—A drive Locate command has been issued in the software.
7 NVMe	NVMe drive activity LED	<ul style="list-style-type: none"> • 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 2: Rear Panel LEDs

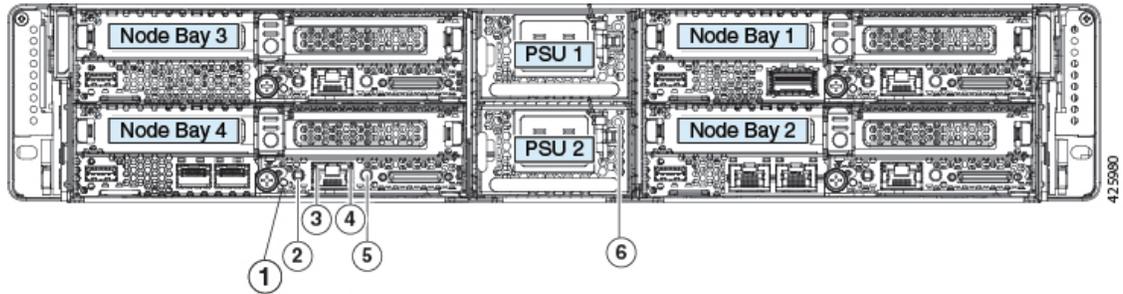


Table 2: Rear Panel LEDs, Definition of States

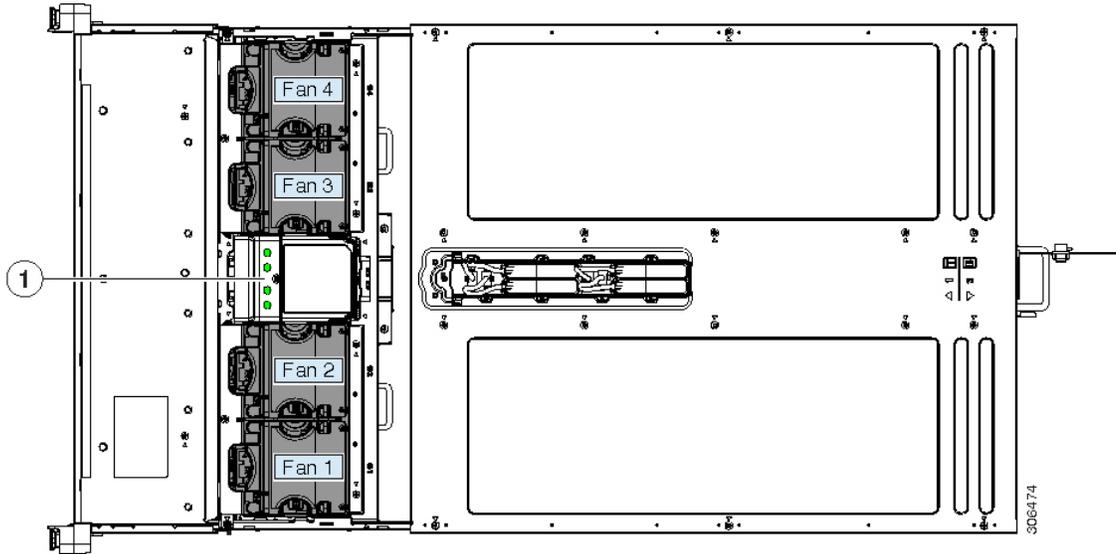
	LED Name	States
1	Node Health Status	<ul style="list-style-type: none"> • Green—The node is operating normally. • Green, Blinking—The node is in standby power mode. • Amber—The node is in a degraded condition (for example, one or more of the following conditions): <ul style="list-style-type: none"> • 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): <ul style="list-style-type: none"> • 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)	<ul style="list-style-type: none"> • Off—There is no AC power to the node. • Amber—The node is in standby power mode. Power is supplied only to the Cisco IMC and some motherboard functions. • Green—The node is in main power mode. Power is supplied to all node components.
3	Node 1-Gb Ethernet dedicated management link speed (One each node)	<ul style="list-style-type: none"> • Off—Link speed is 10 Mbps. • Amber—Link speed is 100 Mbps. • Green—Link speed is 1 Gbps.

4	Node 1-Gb Ethernet dedicated management link status (One each node)	<ul style="list-style-type: none"> • Off—No link is present. • Green—Link is active. • Green, blinking—Traffic is present on the active link.
5	Node locator beacon (One each node)	<ul style="list-style-type: none"> • Off—The unit identification function is not in use. • Blue, blinking—The unit identification function is activated.
6	Power supply status (One bi-color LED each power supply unit)	<p>AC power supplies:</p> <ul style="list-style-type: none"> • 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 3: Chassis Internal Diagnostic LED Location on Fan Tray



1	<p>Fan module fault LEDs on fan tray (one LED for each fan module)</p> <ul style="list-style-type: none"> • 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.

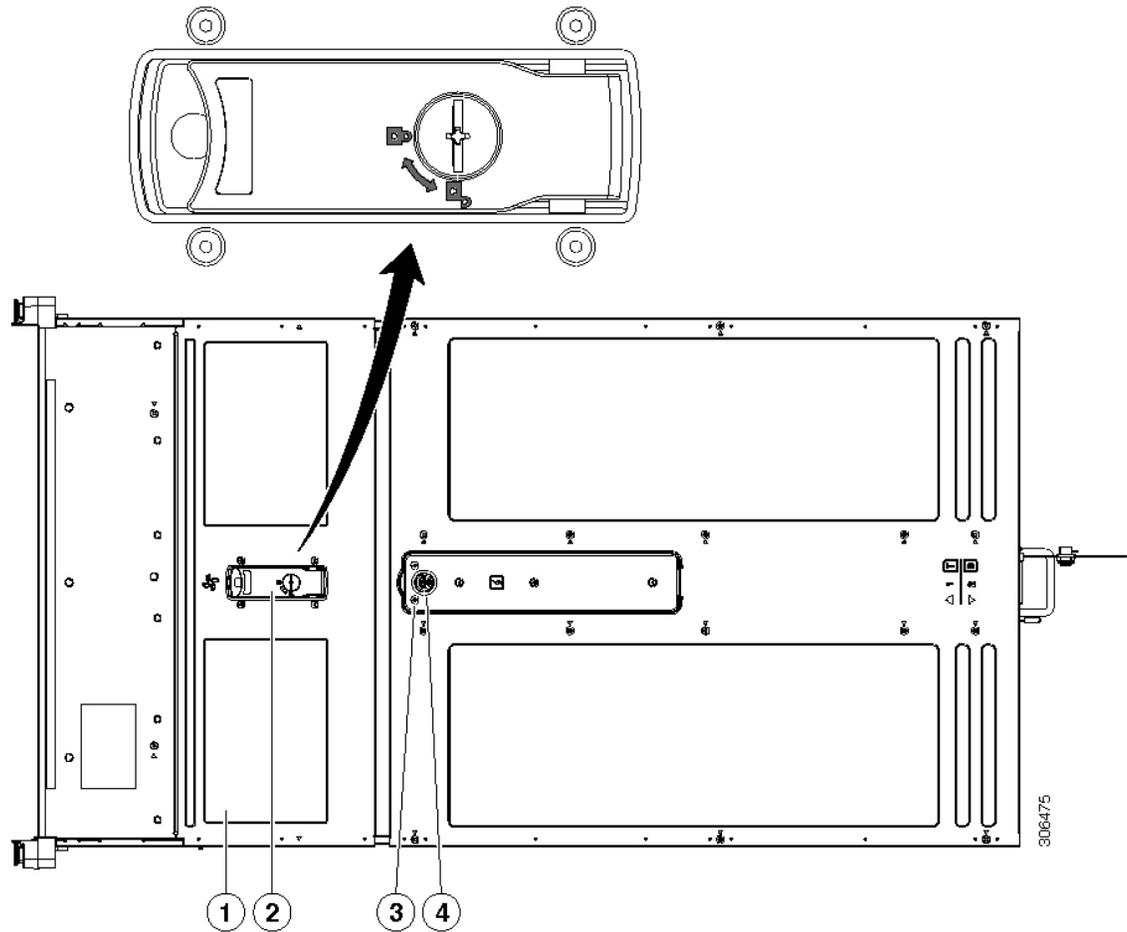
Opening the Chassis Compartment Covers

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 chassis.

Figure 4: 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:
- 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 as you lift the latch.
 - Open the hinged cover.
- Step 2** Close the hinged cover:
- With the latch in the fully open position, close the hinged cover.
 - 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.

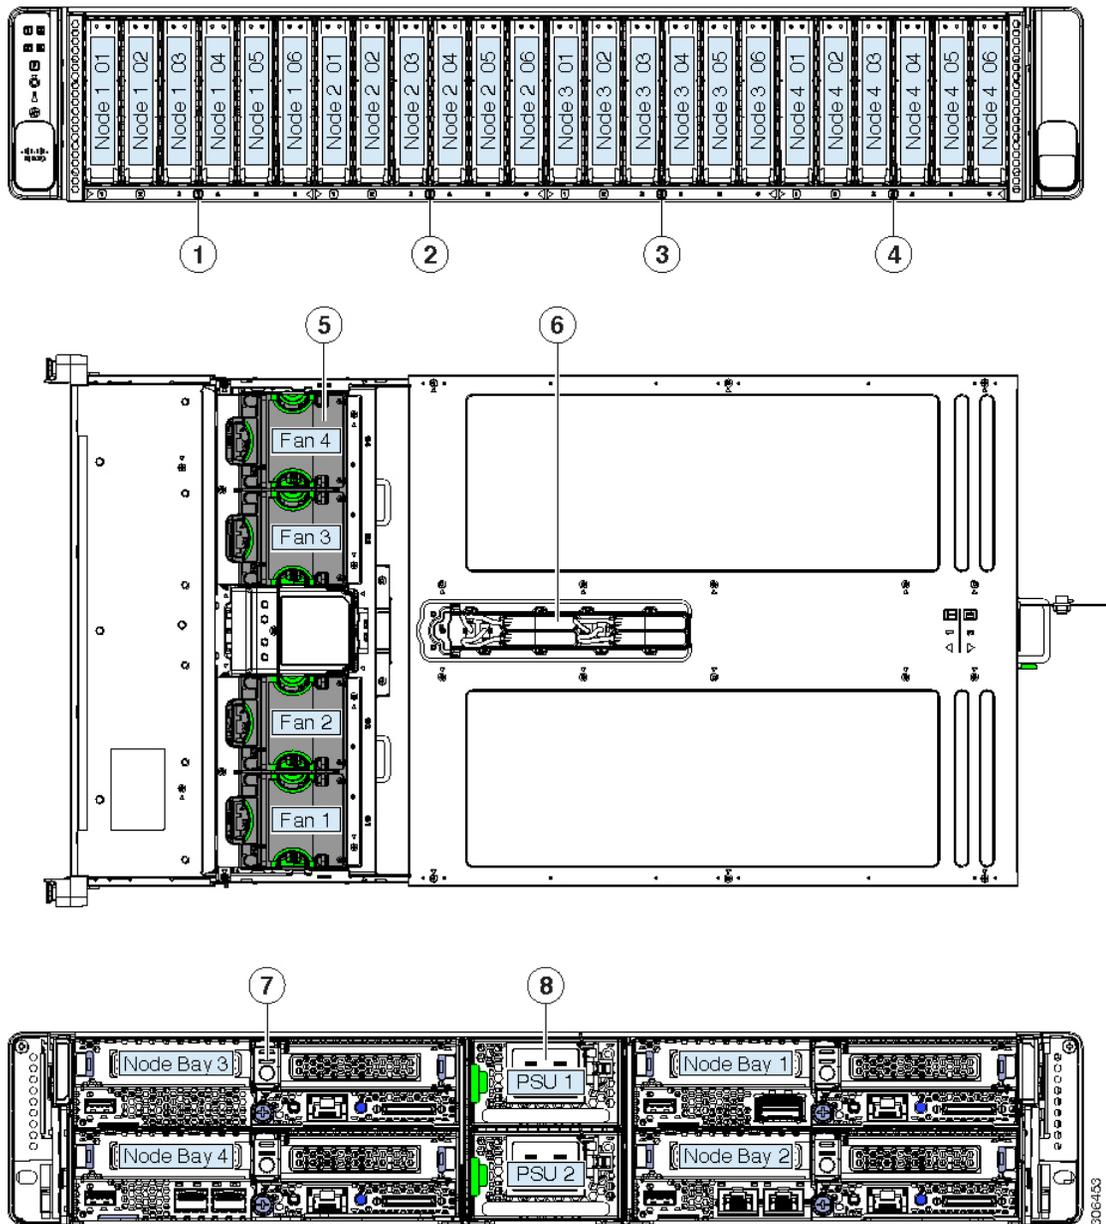
Serviceable Components in the Chassis

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 5: Cisco UCS C4200 Chassis Serviceable Component Locations



<p>1</p>	<p>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.</p>	<p>5</p>	<p>Cooling fan modules (four) Each fan module contains two fans for redundancy.</p>
-----------------	---	-----------------	---

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



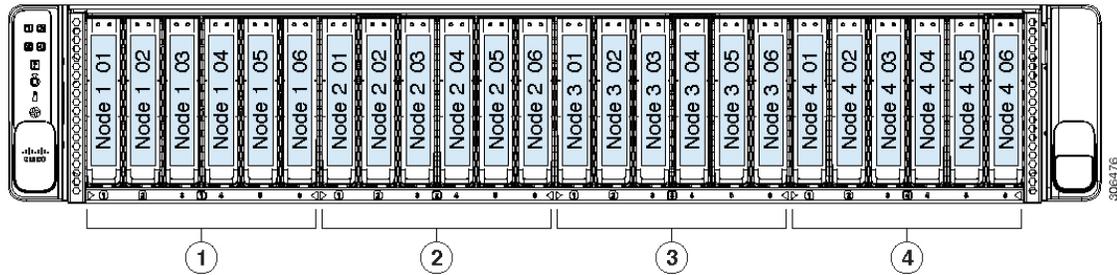
Note You do not have to shut down the drive or the corresponding compute node to replace SAS/SATA hard drives or SSDs because they are hot-swappable.

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 6: 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:
- Enter the BIOS setup utility by pressing the **F2** key when prompted during bootup.
 - Go to the **Boot Options** tab.
 - 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.

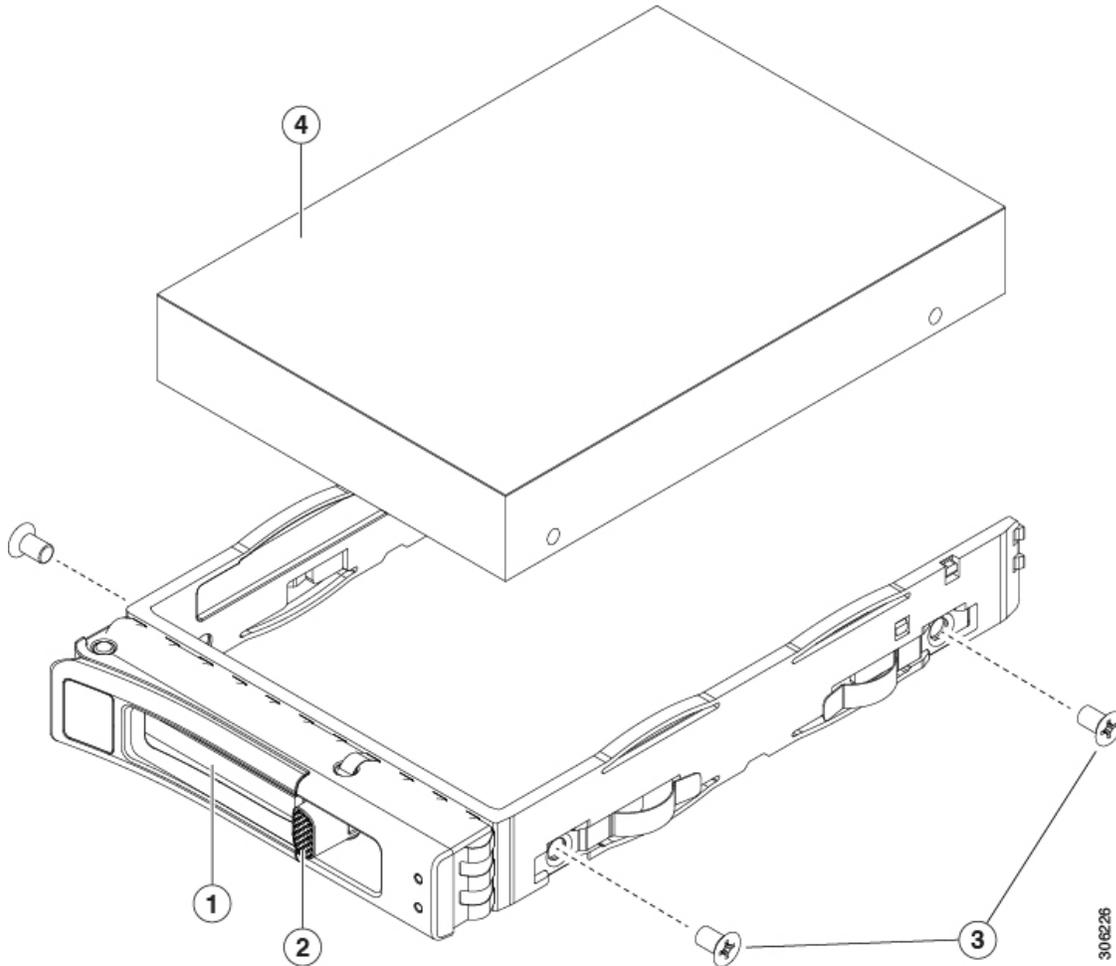
Replacing a Front-Loading SAS/SATA Drive



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:
- 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.
 - 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 7: 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.



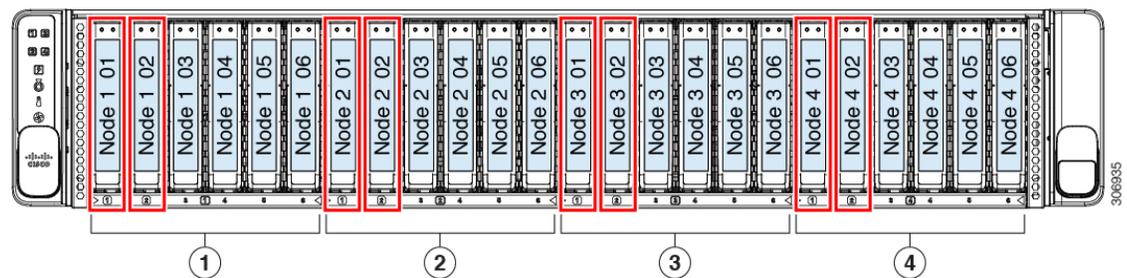
Note NVMe drives are supported only with Cisco IMC/BIOS 4.0(2) and later; they are supported with Cisco UCS Manager 4.0(2) and later.

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 8: 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.

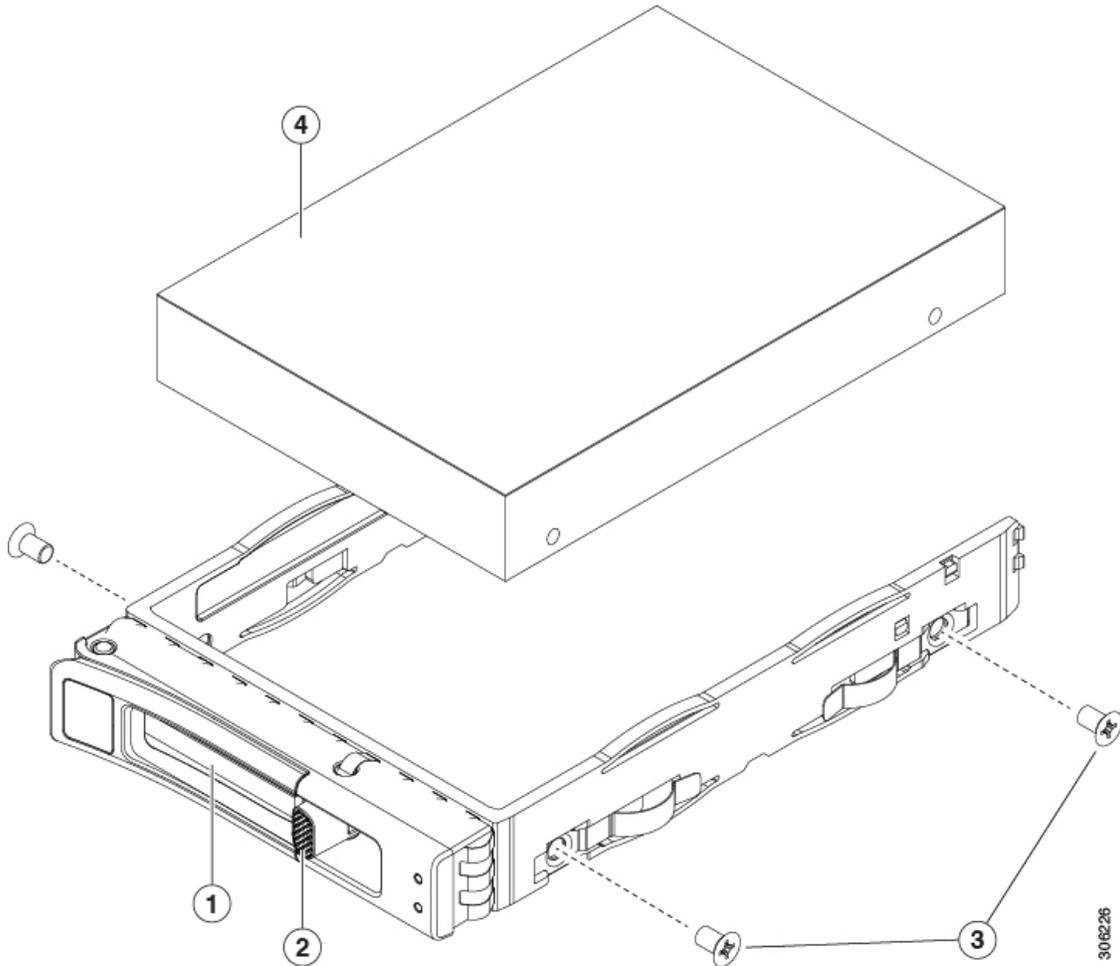
Step 2 Install a new front-loading NVMe SSD:

- 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 9: 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.

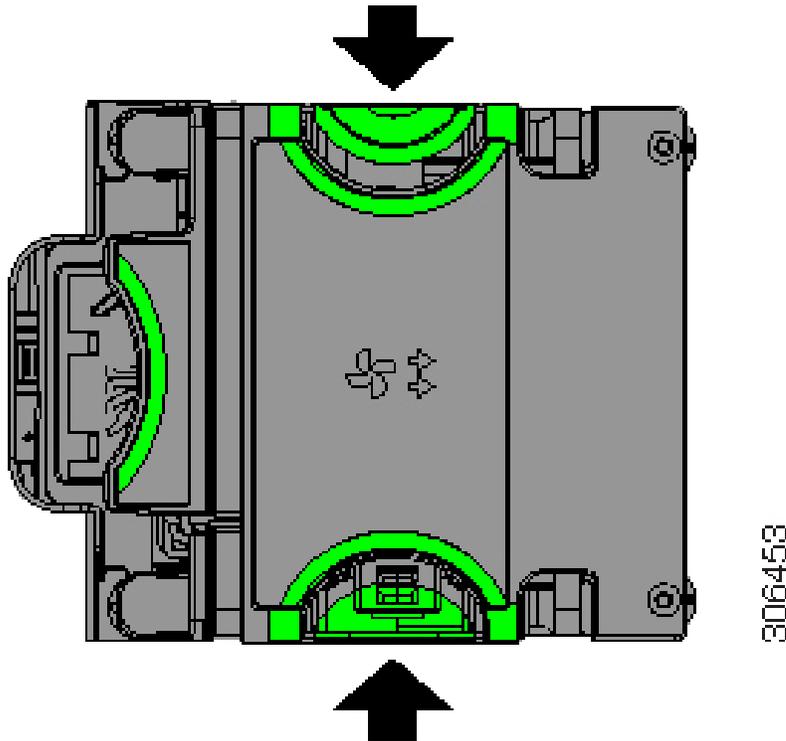
Note The arrows printed on the top of the fan module must point toward the rear of the server.

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 10: 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 6](#).
- 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.

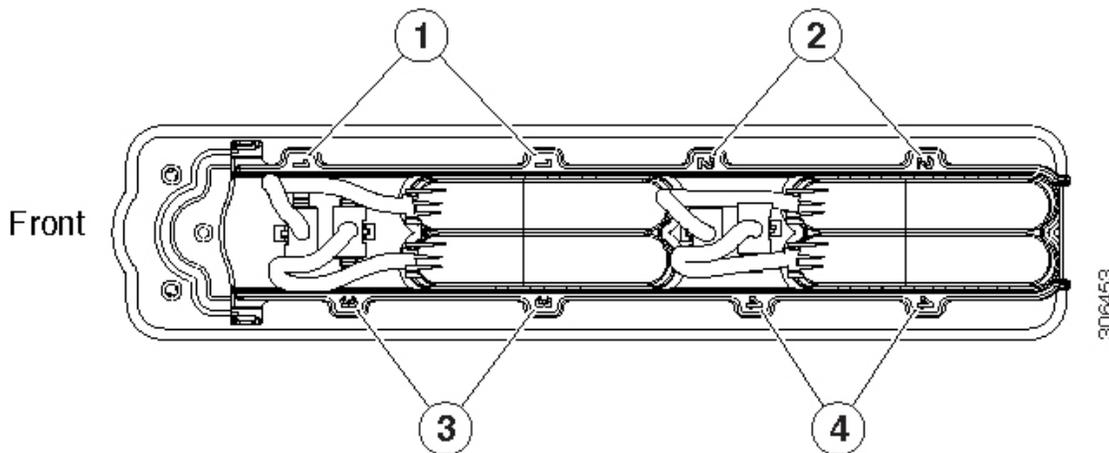
Note Make sure that the number of the supercap unit bay, the number of the cable connector, and the number of the compute node match. A supercap cabled to a connector for an absent compute node is not seen by the system.

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 11: 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](#) for more information about the supported power supplies.

- See also [Rear-Panel LEDs, on page 3](#) 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.

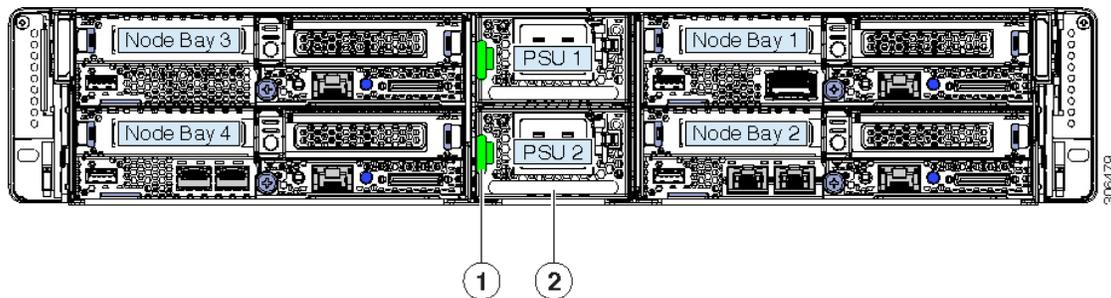
Step 1 Remove the power supply that you are replacing or a blank panel from an empty bay:

- 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 6](#).
 - 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 12: 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](#)

