CHAPTER 3

Maintaining the Server

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

Server Monitoring and Management Tools

Cisco Integrated Management Interface

You can monitor the server inventory, health, and system event logs by using the built-in Cisco Integrated Management Controller (CIMC) GUI or CLI interfaces. See the user documentation for your firmware release at the following URL:

http://www.cisco.com/en/US/products/ps10739/products_installation_and_configuration_guides_list.html

Server Configuration Utility

Cisco has also developed the Cisco Server Configuration Utility for C-Series servers, which can aid and simplify the following tasks:

- Monitoring server inventory and health
- Diagnosing common server problems with diagnostic tools and logs
- Setting the BIOS booting order
- Configuring some RAID configurations
- Installing operating systems

This utility is preinstalled on an internal Cisco Flexible Flash (Cisco FlexFlash) card inside the server (see Overview of the Preinstalled Cisco FlexFlash Drive, page 3-39). You can also download the ISO from Cisco.com. See the user documentation for this utility at the following URL:

http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/ucsscu/user/guide/20/SCUUG20.html

Status LEDs and Buttons

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

Front-Panel LEDs

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

Figure 3-1 Front-Panel LEDs

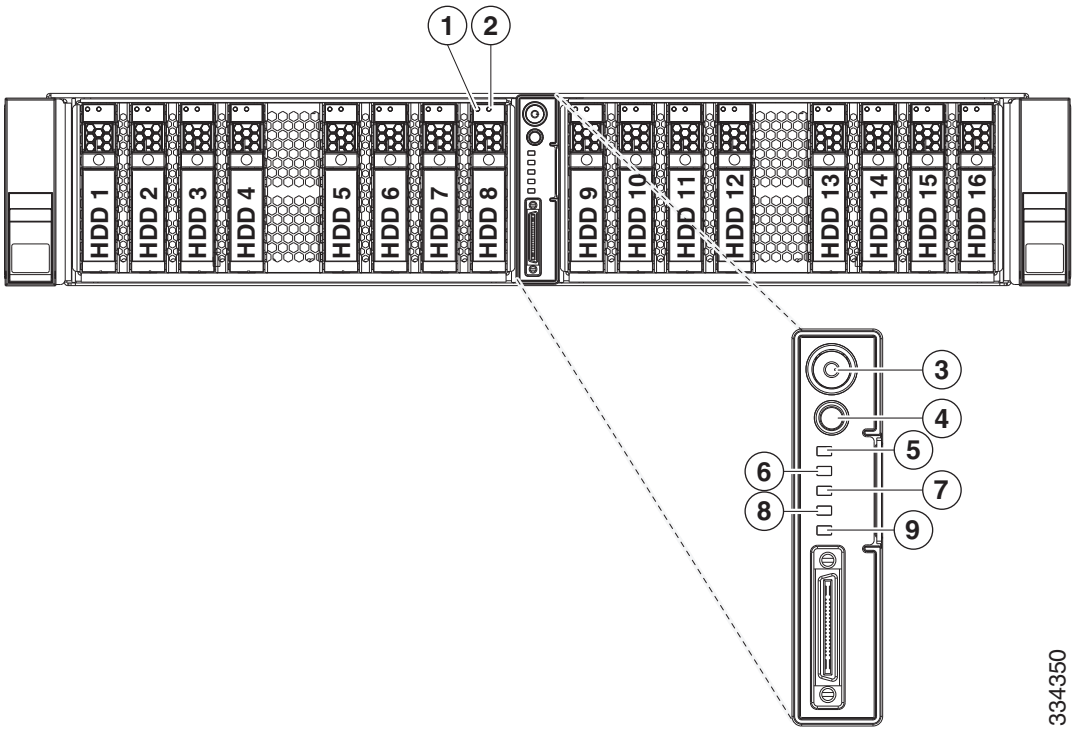

1	Hard drive fault LED	**6**	Fan status LED
2	Hard drive activity LED	**7**	Temperature status LED
3	Power button and Power status LED	**8**	Power supply status LED
4	Identification button and LED	**9**	Network link activity LED
5	System status LED		

Table 3-1 Front-Panel LEDs States

LED Name	State
Hard drive activity	• 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.
Hard drive fault	• Off—The hard drive is operating properly. • Amber—This hard drive has failed. • Amber, blinking—The device is rebuilding.
Power button/Power status LED	• Off—There is no AC power to the server. • Amber—The server is in standby power mode. Power is supplied only to the CIMC and some motherboard functions. • Green—The server is in main power mode. Power is supplied to all server components.
Identification	• Off—The Identification LED is not in use. • Blue—The Identification LED is activated.
System status	• Green—The server is running in normal operating condition. • Green, blinking—The server is performing system initialization and memory check. • Amber, steady—The server is in a degraded operational state. For example: – Power supply redundancy is lost. – CPUs are mismatched. – At least one CPU is faulty. – At least one DIMM is faulty. – At least one drive in a RAID configuration failed. • Amber, blinking—The server is in a critical fault state. For example: – Boot failed. – Fatal CPU and/or bus error is detected. – Server is in an over-temperature condition.
Fan status	• Green—All fan modules are operating properly. • Amber, steady—One fan module has failed. • Amber, blinking—Critical fault; two or more fan modules have failed.
Temperature status	• Green—The server is operating at normal temperature. • Amber, steady—One or more temperature sensors have exceeded a warning threshold. • Amber, blinking—One or more temperature sensors have exceeded a critical threshold.

Table 3-1 Front-Panel LEDs States (continued)

LED Name	State
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.
Network link activity	• Off—The Ethernet link is idle. • Green—One or more Ethernet LOM ports are link-active.

Rear-Panel LEDs and Buttons

Figure 3-2 shows the rear-panel LEDs and buttons. Table 3-2 on page 3-5 defines the rear-panel LED states.

Figure 3-2 Rear-Panel LEDs and Buttons

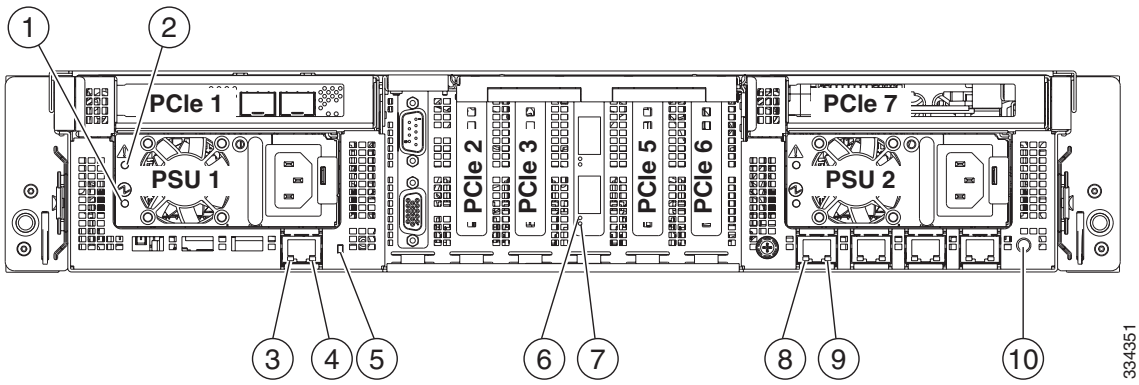

1	Power supply status	**6**	10-Gb Ethernet link status (present only if 10-Gb VIC is installed)
2	Power supply fault	**7**	10-Gb Ethernet link speed (present only if 10-Gb VIC is installed)
3	10/100/1000 Ethernet link speed	**8**	1-Gb Ethernet link speed
4	10/100/1000 Ethernet link status	**9**	1-Gb Ethernet link status
5	System status	**10**	Rear identification button and LED

Table 3-2 Rear-Panel LED States

LED Name	State
Power supply status	• Green, steady—The power supply is operating normally and supplying AC power to the server. • Green, blinking—The power supply is off and in cold-redundancy mode.
Power supply fault	• Off—The power supply is operating normally. • Amber, blinking—An event warning threshold has been reached, but the power supply continues to operate. • Amber, steady—A critical fault threshold has been reached, causing the power supply to shut down.
10/100/1000 Ethernet link speed	• Off—Link speed is 10 Mbps. • Green—Link speed is 100 Mbps.
10/100/1000 Ethernet link status	• Off—No link is present. • Green—Link is active. • Green, blinking—Traffic is present on the active link.
System status	• Green—The server is running in normal operating condition. • Green, blinking—The server is performing system initialization and memory check. • Amber, steady—The server is in a degraded operational state. For example: – Power supply redundancy is lost. – CPUs are mismatched. – At least one CPU is faulty. – At least one DIMM is faulty. – At least one drive in a RAID configuration failed. • Amber, blinking—The server is in a critical fault state. For example: – Boot failed. – Fatal CPU and/or bus error is detected. – Server is in an over-temperature condition.
10-Gb Ethernet link speed	• Off—Link speed is 10 Mbps. • Green—Link speed is 10 Gbps.
10-Gb Ethernet link status	• Off—No link is present. • Green—Link is active. • Green, blinking—Traffic is present on the active link.
1-Gb Ethernet link speed	• Off—Link speed is 10 Mbps. • Amber—Link speed is 100 Mbps. • Green—Link speed is 1 Gbps.

Table 3-2 Rear-Panel LED States (continued)

LED Name	State
1-Gb Ethernet link status	• Off—No link is present. • Green—Link is active. • Green, blinking—Traffic is present on the active link.
Identification	• Off—The identification LED is not in use. • Blue—The identification LED is activated.

Internal Diagnostic LEDs

The server is equipped with a supercap voltage source that can activate internal component fault LEDs up to 30 minutes after AC power is removed. The server has internal fault LEDs for fan modules, CPU sockets, DIMM sockets, the motherboard RTC battery, PCIe sockets, TPM socket, I/O riser socket, and Cisco FlexFlash cards.

To use these LEDs to identify a failed component, press the front or rear identification button with AC power removed (see Figure 3-1 or Figure 3-2 for the identification button location). See Figure 3-3 for the locations of these internal LEDs.

Figure 3-3 Internal Diagnostic LED Locations

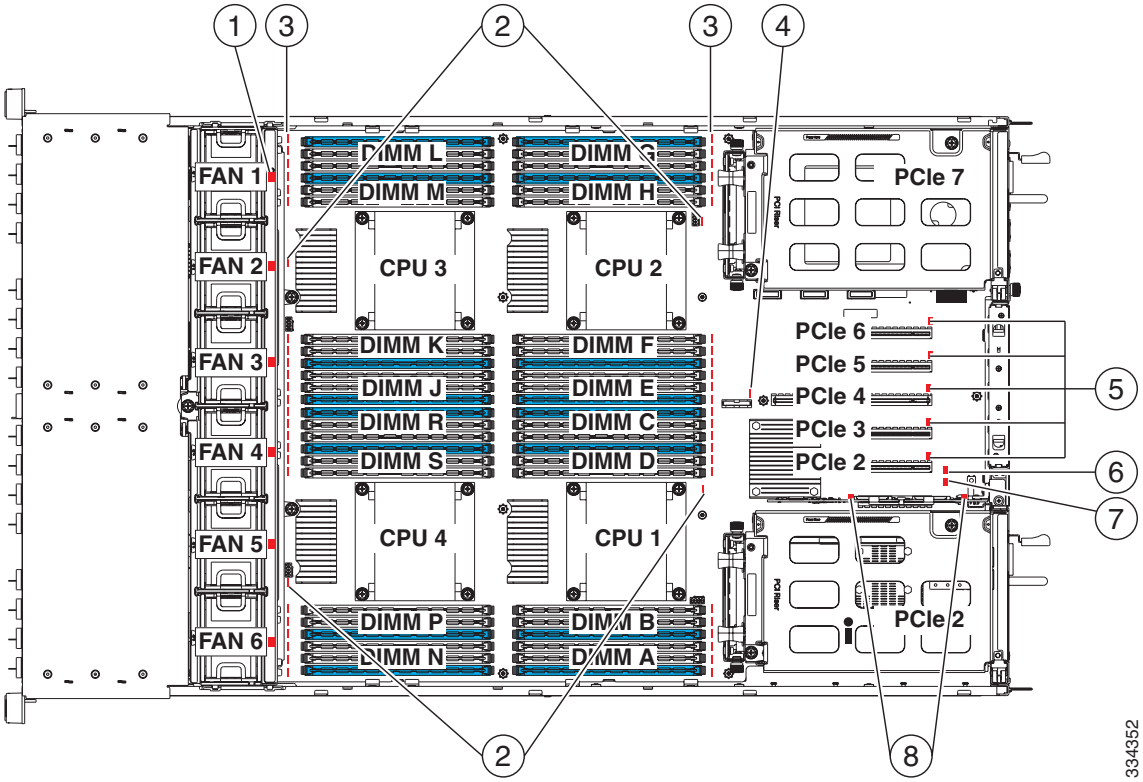

1	Fan module fault LEDs (one on each fan)	**5**	PCIe card fault LEDs (one for each PCIe socket on motherboard)
2	CPU fault LEDs on motherboard: • CPU1 LED = CR69 • CPU2 LED = CR68 • CPU3 LED = CR42 • CPU4 LED = CR43	**6**	TPM fault LED on motherboard (CR40)
3	DIMM fault LEDs on motherboard (one for each DIMM socket)	**7**	I/O riser fault LED on motherboard (CR39)
4	RTC battery fault LED on motherboard	**8**	Cisco FlexFlash card fault LEDs on I/O riser

Table 3-3 Internal Diagnostic LED States

LED Name	State
Internal diagnostic LEDs (all)	• Off—Component is functioning normally. • Amber—Component has failed.

Preparing for Server Component Installation

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

Required Equipment

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

- Number 1 Phillips-head screwdriver
- Number 2 Phillips-head screwdriver
- Number 10 Torx screwdriver
- Electrostatic discharge (ESD) strap or other grounding equipment such as a grounded mat

Shutting Down and Powering Off the Server

The server can run in two power modes:

- Main power mode—Power is supplied to all server components and any operating system on your hard drives can run.
- Standby power mode—Power is supplied only to the service processor and the cooling fans and it is safe to power off the server from this mode.

You can invoke a graceful shutdown or a hard shutdown by using either the CIMC management interface or the Power button on the front panel.

To use the Power button, follow these steps:

Step 1 Check the color of the Power Status LED (see the "Front-Panel LEDs" section on page 3-2).

- Green—The server is in main power mode and must be shut down before it can be safely powered off. Go to Step 2.
- Amber—The server is already in standby mode and can be safely powered off. Go to Step 3.

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

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

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

Step 3 Disconnect the power cords from the power supplies in your server to completely power off the server.

Removing and Replacing the Server Top Cover

To remove or replace the top cover of the server, follow these steps:

Tip You do not have to remove the cover to replace hard drives or power supplies.

Step 1 Remove the top cover:

a. Simultaneously press the two release buttons. See Figure 3-4.

b. Push the top cover toward the server rear about one-half inch (1.27 cm), until it stops.

c. Lift the top cover straight up from the server and set it aside.

Step 2 Replace the top cover:

a. Place the cover on top of the server about one-half inch (1.27 cm) behind the lip of the chassis front cover panel. The cover should sit flat when the edge flanges are sitting in the grooves in the chassis.

b. Slide the top cover toward the front cover panel until it stops and the release buttons lock.

Figure 3-4 Removing the Top Cover

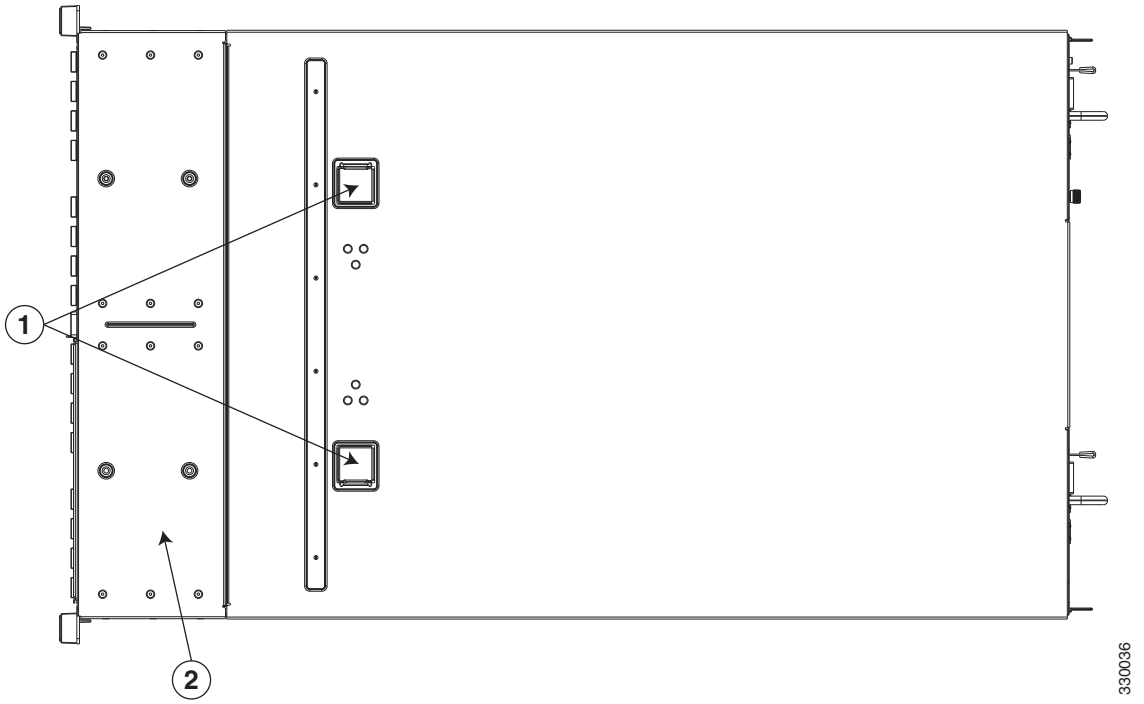

1	Release buttons	**2**	Front cover panel

Replaceable Component Locations

This section shows the locations of the components that are discussed in this chapter. The view in Figure 3-5 is from the top down with the top cover removed.

Figure 3-5 Replaceable Component Locations

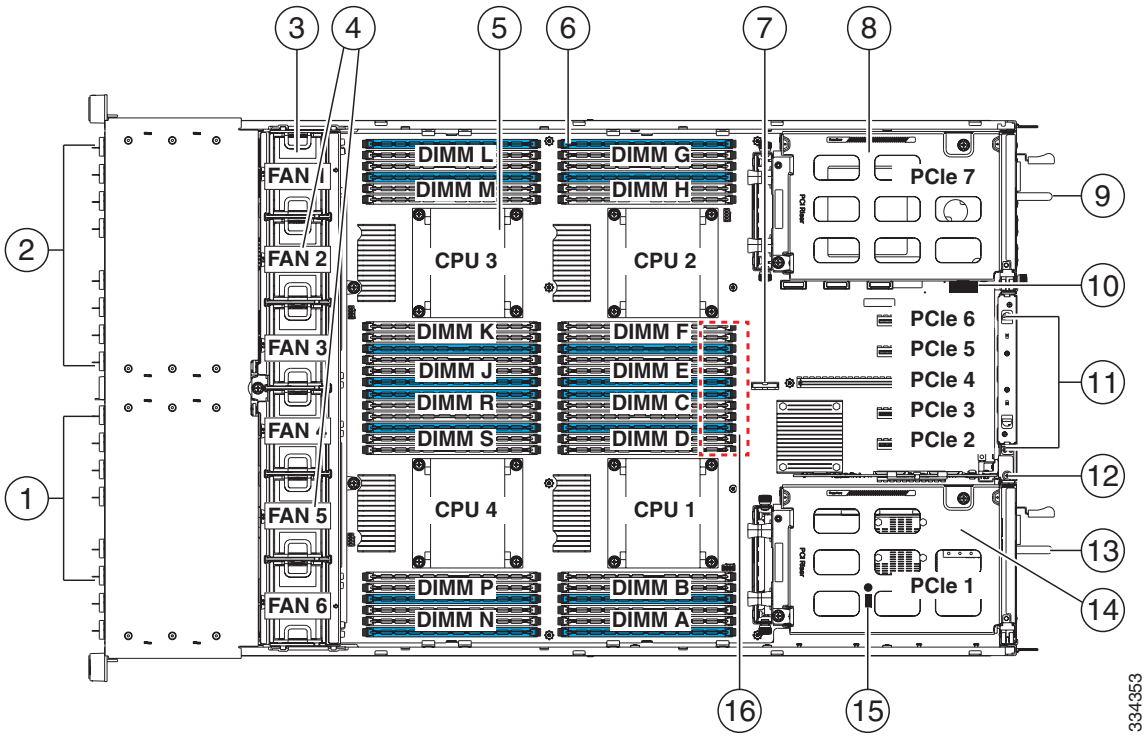

1	Drive bay module 2 (up to eight 2.5-inch drives, hot-pluggable)	**9**	Power supply 2 (hot-pluggable)
2	Drive bay module 1 (up to eight 2.5-inch drives, hot-pluggable)	**10**	USB 2.0 slot on motherboard
3	Fan tray, with six hot-pluggable fan modules	**11**	PCIe slots 2–6 on motherboard See also Figure 3-23 on page 3-43 for details.
4	Drive backplane transition cards (up to two on chassis floor, not visible under fan tray in this view)	**12**	I/O riser (includes two sockets for Cisco FlexFlash cards)
5	CPUs and heatsinks (two or four)	**13**	Power supply 1 (hot-pluggable)
6	DIMM sockets on motherboard (up to 48)	**14**	PCIe riser 2 (horizontal PCIe slot 1)
7	RTC battery on motherboard	**15**	TPM socket (on motherboard, not visible under power supply in this view)
8	PCIe riser 1 (horizontal PCIe slot 7)	**16**	RAID backup unit (supercap power module) mounting location (two, on air baffle not shown in this view)

Color-Coded Touch Points

This server has color-coded touch points that indicate thumbscrews and latches on replaceable and hot-swappable components.

- Hot-swappable components have green plastic touch points, which include the internal cooling fans and the power supplies. (Exceptions are the drive trays on the front panel, which are hot-swappable but not green).
- Some replaceable but non-hot-swappable components have light-blue plastic touch points, which include PCIe risers, PCIe slots, and the fan tray.

Installing or Replacing Server Components

Warning **Blank faceplates and cover panels serve three important functions: they prevent exposure to hazardous voltages and currents inside the chassis; they contain electromagnetic interference (EMI) that might disrupt other equipment; and they direct the flow of cooling air through the chassis. Do not operate the system unless all cards, faceplates, front covers, and rear covers are in place.**
Statement 1029

Warning **Class 1 laser product.**
Statement 1008

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

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

Replacing Hard Drives or Solid State Drives

This section includes the following topics:

Drive Population Guidelines

The server can operate with one or two drive bay modules. Each drive bay module can hold up to eight 2.5-inch drives for a total of 16 drives. Figure 3-6 shows the drive bay modules and the drive bay numbering.

Figure 3-6 Drive Bay Modules and Drive Bay Numbering

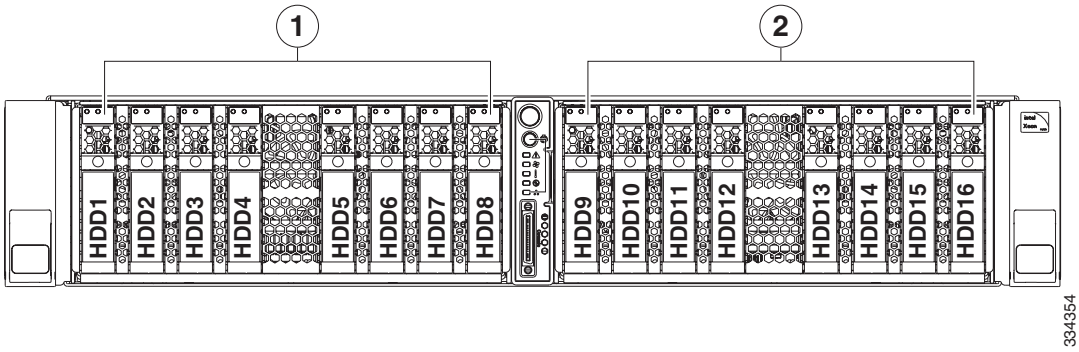

1	Drive bay module 1	**2**	Drive bay module 2

Observe these drive population guidelines for optimum performance:

- When populating drives in a drive bay module, add drives in empty slots from left to right.
- If your server has two drive bay modules, populate all eight bays in module 1 before you populate the module 2 bays.
- Keep an empty drive blanking tray in any unused slots to ensure proper airflow.
- You can mix hard drives and SSDs in the same server. However, you cannot configure a logical volume (virtual drive) that contains a mix of hard drives and SSDs. When you create a logical volume, it must contain all hard drives or all SSDs.

Drive Replacement Procedure

To replace or install a hot-pluggable hard drive, follow these steps:

Tip You do not have to shut down or power off the server to replace hard drives or solid state drives (SSDs) because they are hot-pluggable.

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

a. Press the release button on the face of the drive tray. See Figure 3-7.

b. Grasp and open the ejector lever and then pull the drive tray out of the slot.

c. 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 replace the four drive tray screws.

b. With the ejector lever on the drive tray open, insert the drive tray into the empty drive bay.

c. Push the tray into the slot until it touches the backplane, and then close the ejector lever to lock the drive in place.

Figure 3-7 Removing and Replacing Hard Drives

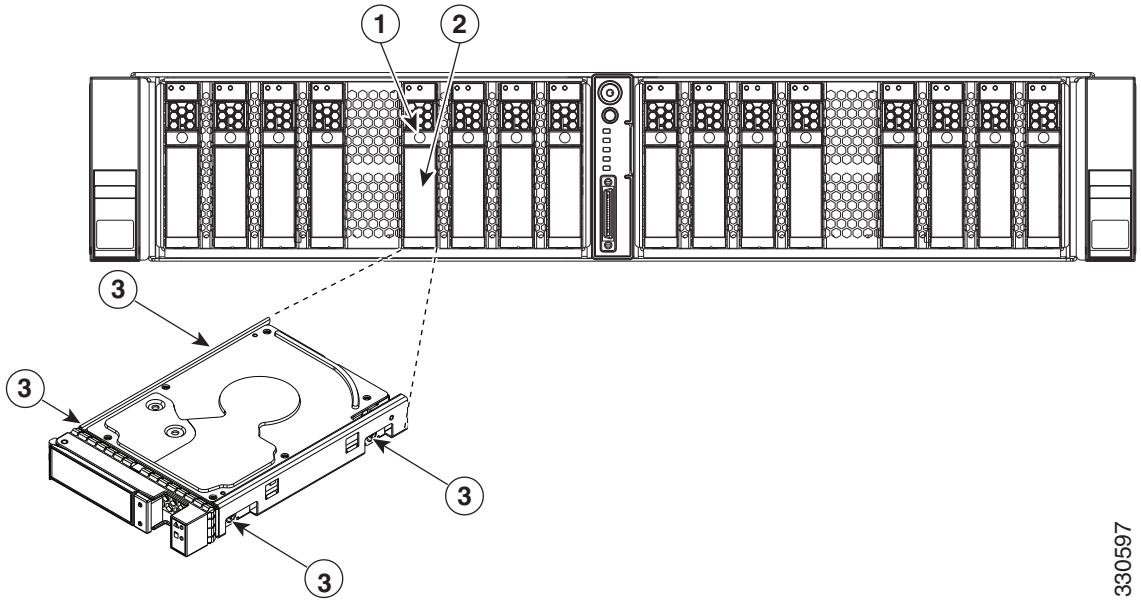

1	Release button	3	Drive tray securing screws (4)
2	Ejector lever		

Replacing a Modular Drive Bay Assembly

The server allows modular configuration of the drive bays, so that you can use one or two drive bays depending on your needs.

Each drive bay assembly includes the chassis-steel drive bay, the attached drive backplane, and a transition card (expander or nonexpander version) that provides a connection to your RAID controllers and the motherboard.

- Each of the two modular chassis drive bays can hold up to eight 2.5-inch drives.
- Each drive backplane connects to one transition card. The transition card connects the drive backplane to the motherboard and connects to cables from your RAID controller.
- The transition card is available in two versions:
  - Nonexpander—Two internal connections to the RAID controller card are required to control eight drives. This version includes two connectors for the cables from your RAID controller. One connector allows control of drives 1 through 4 on the backplane. The other connector allows control of drives 5 through 8 on the backplane.
  - Expander—One internal connection to the RAID controller card is required to control eight drives. This one connector allows control of drives 1 through 8 on the backplane.
- To use more than eight 2.5-inch drives in the server (up to 16), two drive bay modules are required, each with its own backplane and transition card.

Figure 3-8 Drive Bay Assembly and Transition Card (Two Versions)

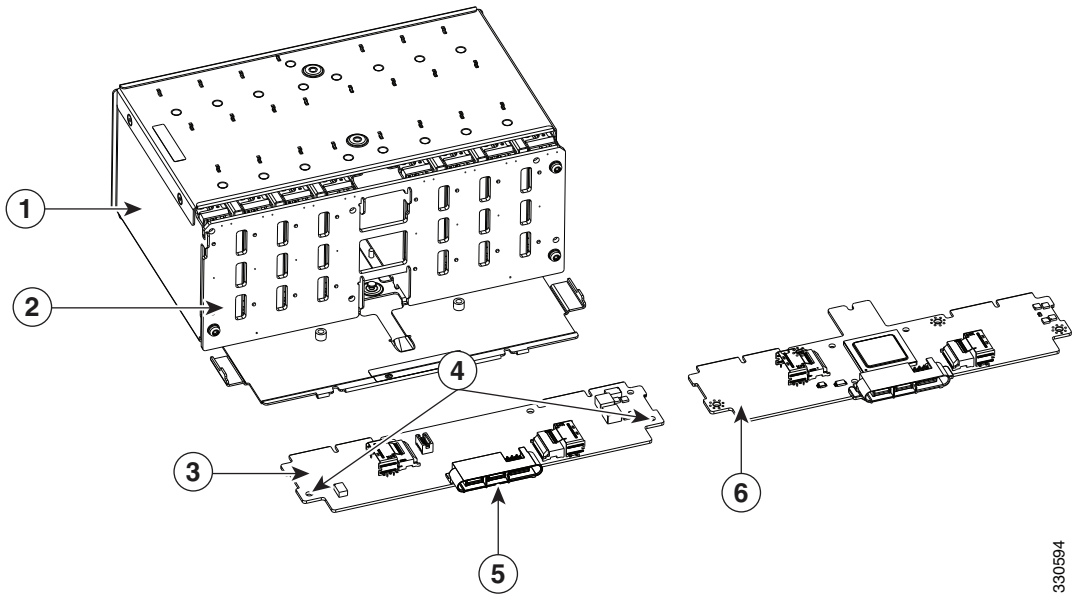

1	Chassis drive bay module (rear view)	**4**	Transition card securing screws (two)
2	Drive backplane	**5**	Transition card connector to motherboard
3	Transition card (nonexpander version)	**6**	Transition card (expander version)

To install or replace the modular drive bay assembly, follow these steps:

Step 1 Remove the assembly that you are replacing. See Figure 3-8 and Figure 3-9:

a. Power off the server as described in Shutting Down and Powering Off the Server, page 3-8.

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

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

c. Remove the top cover as described in Removing and Replacing the Server Top Cover, page 3-9.

d. Remove all drives from the drive-bay module that you are replacing.

e. Remove the fan tray from the chassis, as described in Replacing a Fan Tray, page 3-19. You do not have to remove the fan modules from the fan tray.

f. Disconnect RAID controller cables from the transition card that you are replacing.

Tip Label the cables before you disconnect them to aid replacement.

g. Use a Number 2 Phillips screwdriver to remove the two screws that secure the transition card to the chassis floor (see Figure 3-9).

h. Slide the drive bay with the attached backplane and transition card out the front of the chassis. The transition card disengages from its motherboard connector.

Note Tilt the module downward as you slide it out the chassis opening to provide clearance over an alignment peg on the chassis floor.

Step 2 Install a new drive bay assembly:

a. Carefully slide the drive bay with the attached backplane and transition card into the front chassis opening.

Stop when the connector on the rear edge of the transition card is fully engaged with the motherboard connector and the front of the drive bay is even with the chassis front panel.

b. Replace the two screws that secure the transition card to the chassis floor.

c. Replace the RAID controller cables to the connectors on the transition card.

d. Replace the fan tray as described in Replacing a Fan Tray, page 3-19.

e. Replace the top cover.

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

Figure 3-9 Removing and Replacing the Drive-Bay Module

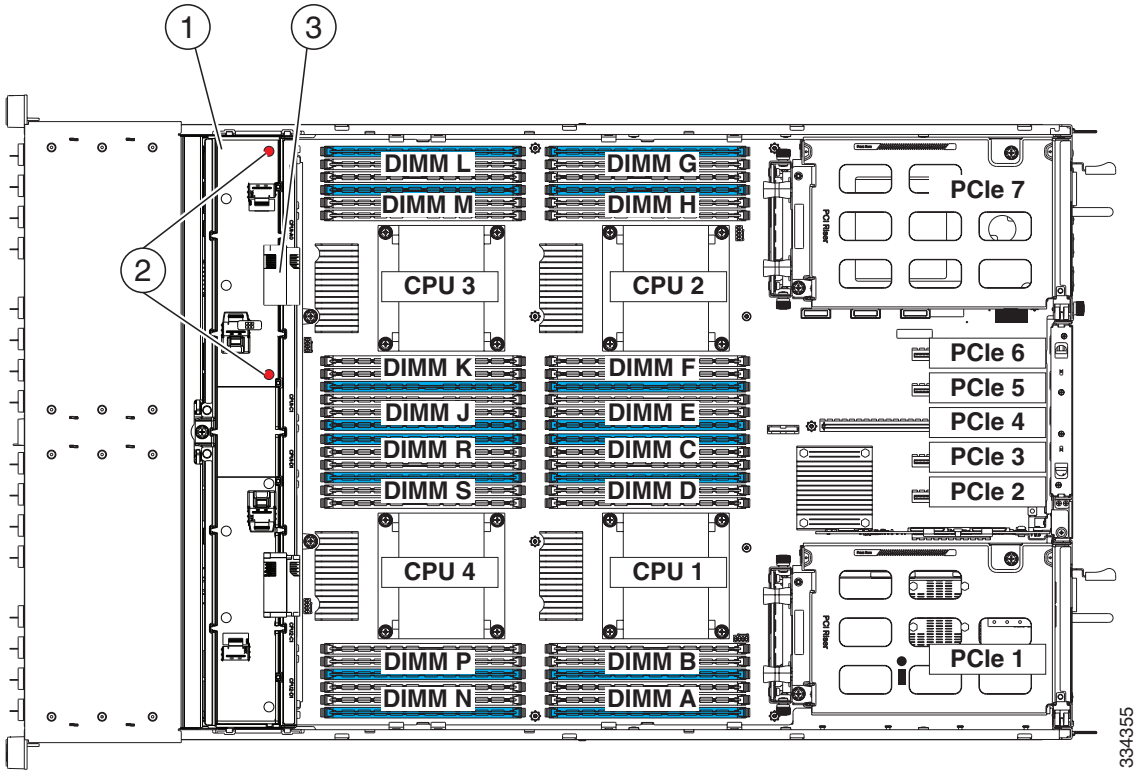

1	Transition card on chassis floor (shown with fan tray removed)	**3**	Transition card connector to motherboard
2	Transition card securing screws (2)		

Replacing Fan Modules

The six fan modules in the server are numbered as shown in Figure 3-10 when you are facing the front of the server.

Figure 3-10 Fan Module Numbering

FAN 1	FAN 2	FAN 3	FAN 4	FAN 5	FAN 6
Front of Server					

Tip Each fan module has a fault LED that lights amber if the fan module fails.

To replace or install a hot-pluggable fan module, follow these steps:

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

Step 1 Remove the fan module that you are replacing (see Figure 3-11):

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

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

b. Remove the top cover as described in Removing and Replacing the Server Top Cover, page 3-9.

c. Insert your thumb and forefinger into the two green release latches on the top of the fan module.

d. Squeeze the release latches together and lift out the fan module.

Step 2 Install a new fan module:

a. Grasp the fan module by the release latches and align it with the empty fan bay and the connector on the floor of the fan tray. See Figure 3-11.

b. Press down on the top corners of the fan module until the connector is fully seated and the release latches lock in place.

c. Replace the top cover.

d. Replace the server in the rack.

Figure 3-11 Removing and Replacing Fan Modules

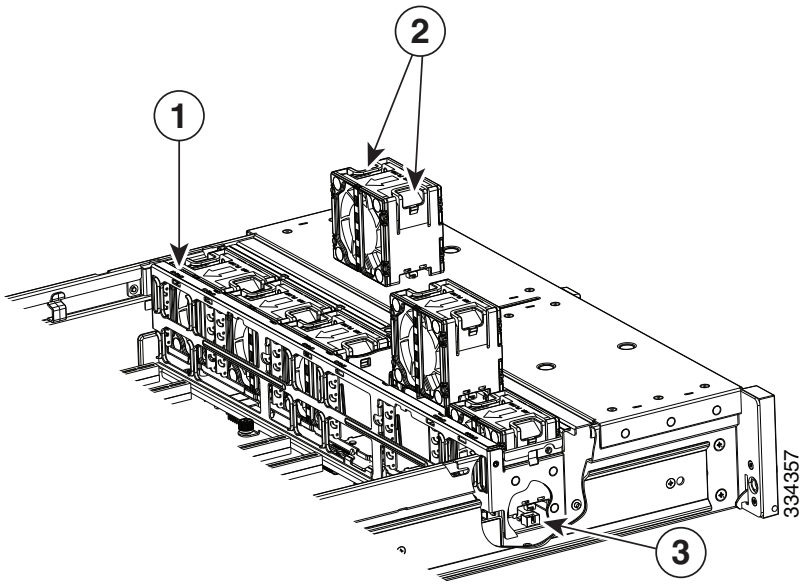

1	Fan tray	**3**	Connector location on underside of fan module
2	Fan module release latches		

Replacing a Fan Tray

To replace a fan tray, follow these steps:

Step 1 Remove the fan tray (see Figure 3-12):

a. Power off the server as described in Shutting Down and Powering Off the Server, page 3-8.

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

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

c. Remove the top cover as described in Removing and Replacing the Server Top Cover, page 3-9.

d. Use a Number 2 Phillips-head screwdriver to loosen the three captive thumbscrews that secure the fan tray to the chassis.

e. Lift the fan tray straight up and out of the chassis.

Step 2 Install a new fan tray (see Figure 3-12):

a. Carefully align the new tray with the chassis and lower it in place:

- Align the connector on the underside of the tray with its socket on the motherboard.
- Align the two guide pegs on each end of the tray with the slots in the chassis.

b. Tighten the three captive thumbscrews that secure the tray to the motherboard and chassis.

c. Replace the top cover.

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

Figure 3-12 Removing and Replacing a Fan Tray

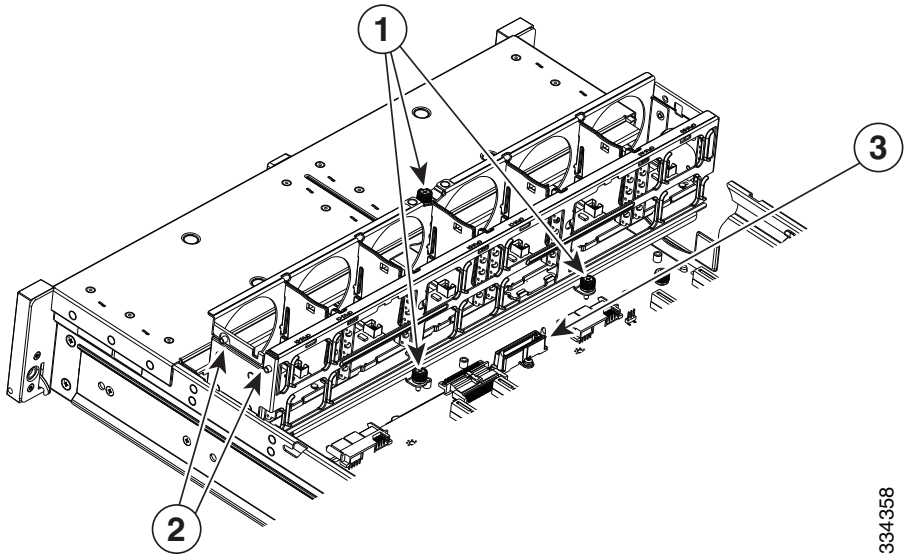

1	Captive thumbscrews (3)	**3**	Motherboard connector location
2	Guide pegs (2 on each end of fan tray)		

Replacing DIMMs

This section includes the following topics:

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

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

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

DIMM Performance Guidelines and Population Rules

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

DIMM Sockets

Figure 3-13 shows the DIMM sockets on the motherboard.

- Each CPU supports four DDR3 memory channels.
- Channels are labeled with letters as shown in Figure 3-13.
- Each channel has three DIMM sockets. The blue socket in a channel is always socket 1.

Figure 3-13 DIMM and CPU Sockets on the Motherboard

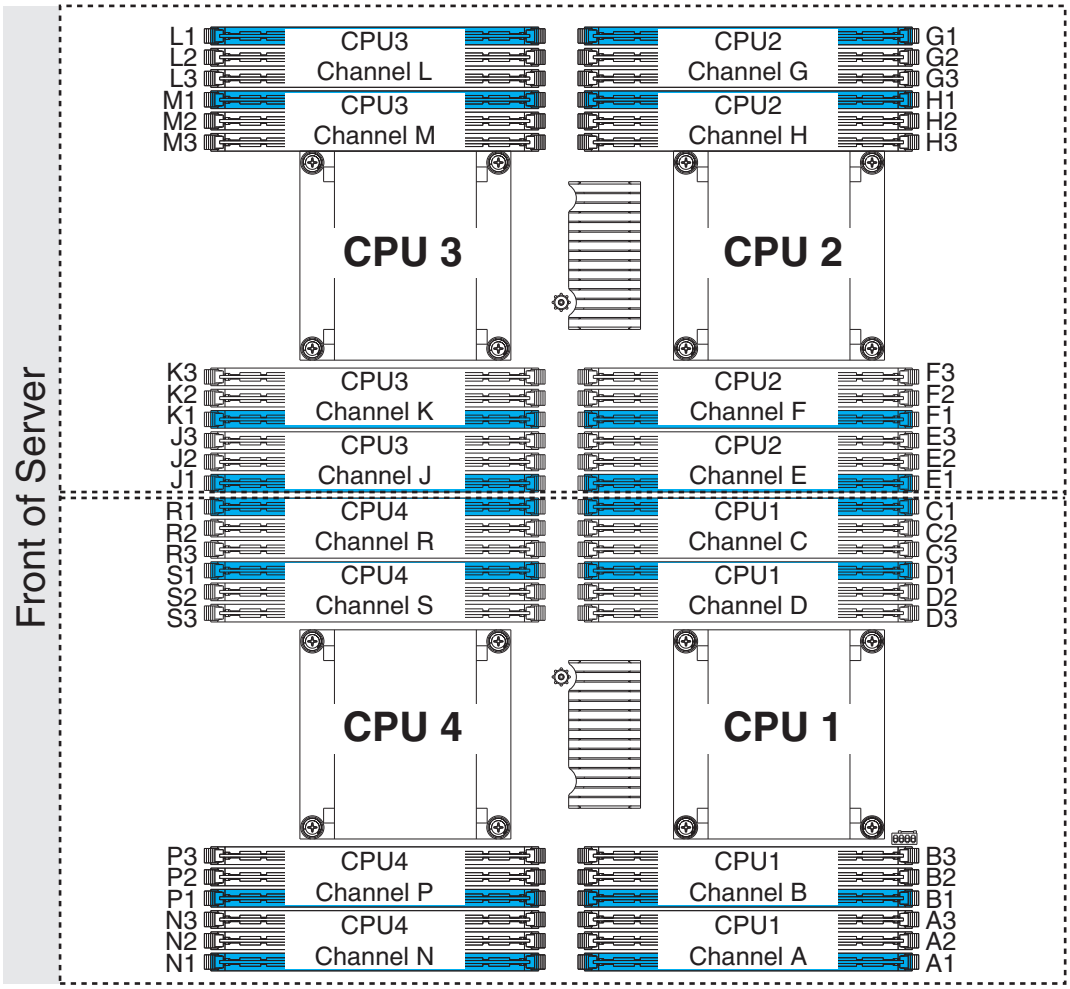

DIMM Population Rules

Observe the following guidelines when installing or replacing DIMMs (see Figure 3-13):

- The minimum configuration is one DIMM installed in any of CPU1's blue sockets (A1, B1, C1, D1).
- When populating a channel, fill the blue socket first, and then fill the black sockets (1, 2, 3).
- The server supports one, two, or three DIMMs per channel for single- or dual-rank DIMMs.
- The server supports only one or two quad-rank registered DIMMs (RDIMMs) per channel. The server supports one, two, or three quad-rank load-reduced DIMMs (LRDIMMs) per channel.
- The server supports RDIMMs or LRDIMMs. Do not mix RDIMMs and LRDIMMs in a server.
- Unregistered DIMMs (UDIMMs) and non-ECC DIMMs are not supported.
- For memory mirroring population rules, see Memory Mirroring Mode, page 3-25.
- There is no requirement to install DIMMs evenly across all CPUs. However, for optimal performance, you should populate sockets in the order that is shown in Table 3-4 for two-CPU systems or Table 3-5 for four-CPU systems.

Table 3-4 DIMM Population for Two-CPU System Optimum Performance

# of DIMMs Installed	CPU	DIMM Socket	Color
First: 1	CPU1	C1	Blue
2	CPU2	E1	Blue
3	CPU1	D1	Blue
4	CPU2	F1	Blue
5	CPU1	B1	Blue
6	CPU2	H1	Blue
7	CPU1	A1	Blue
8	CPU2	G1	Blue
9	CPU1	C2	Black
10	CPU2	E2	Black
11	CPU1	D2	Black
12	CPU2	F2	Black
13	CPU1	B2	Black
14	CPU2	H2	Black
15	CPU1	A2	Black
16	CPU2	G2	Black
17	CPU1	C3	Black
18	CPU2	E3	Black
19	CPU1	D3	Black
20	CPU2	F3	Black
21	CPU1	B3	Black
22	CPU2	H3	Black
23	CPU1	A3	Black
Last: 24	CPU2	G3	Black

Table 3-5 DIMM Population For 4-CPU System Optimum Performance

# of DIMMs Installed	CPU	DIMM Socket	Color	# of DIMMs Installed	CPU	DIMM Socket	Color
First: 1	CPU1	C1	Blue	25	CPU1	B2	Black
2	CPU2	E1	Blue	26	CPU2	H2	Black
3	CPU3	J1	Blue	27	CPU3	M2	Black
4	CPU4	R1	Blue	28	CPU4	P2	Black
5	CPU1	D1	Blue	29	CPU1	A2	Black
6	CPU2	F1	Blue	30	CPU2	G2	Black
7	CPU3	K1	Blue	31	CPU3	L2	Black
8	CPU4	S1	Blue	32	CPU4	N2	Black
9	CPU1	B1	Blue	33	CPU1	C3	Black
10	CPU2	H1	Blue	34	CPU2	E3	Black
11	CPU3	M1	Blue	35	CPU3	J3	Black
12	CPU4	P1	Blue	36	CPU4	R3	Black
13	CPU1	A1	Blue	37	CPU1	D3	Black
14	CPU2	G1	Blue	38	CPU2	F3	Black
15	CPU3	L1	Blue	39	CPU3	K3	Black
16	CPU4	N1	Blue	40	CPU4	S3	Black
17	CPU1	C2	Black	41	CPU1	B3	Black
18	CPU2	E2	Black	42	CPU2	H3	Black
19	CPU3	J2	Black	43	CPU3	M3	Black
20	CPU4	R2	Black	44	CPU4	P3	Black
21	CPU1	D2	Black	45	CPU1	A3	Black
22	CPU2	F2	Black	46	CPU2	G3	Black
23	CPU3	K2	Black	47	CPU3	L3	Black
24	CPU4	S2	Black	**Last**: 48	CPU4	N3	Black

Memory Mirroring Mode

Note You cannot use memory sparing and memory mirroring at the same time.

When you enable memory mirroring mode, the memory subsystem simultaneously writes identical data to two channels. If a memory read from one of the channels returns incorrect data due to an uncorrectable memory error, the system automatically retrieves the data from the other channel. A transient or soft error in one channel does not affect the mirrored data, and operation continues.

Memory mirroring reduces the amount of memory available to the operating system by 50 percent because only one of the two populated channels provides data.

If you choose to enable memory mirroring, you must install sets of identical DIMMs as follows:

- For a two-CPU system, you must install DIMMs in identical sets of four. Use each set of four identical DIMMs to populate the sockets in the order shown in Table 3-4.
- For a four-CPU system, you must install DIMMs in identical sets of eight. Use each set of eight identical DIMMs to populate the DIMM sockets in the order shown in Table 3-5.

Lockstep Channel Mode

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

Lockstep channel mode requires that all four memory channels on a CPU must be populated identically with regards to size and organization. DIMM socket populations within a channel do not have to be identical but the same DIMM slot location across all four channels must be populated the same.

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

If you choose to enable lockstep channel mode, you must install sets of identical DIMMs as follows:

- For a two-CPU system, you must install DIMMs in identical sets of four. Use each set of four identical DIMMs to populate the sockets in the order shown in Table 3-4.
- For a four-CPU system, you must install DIMMs in identical sets of eight. Use each set of eight identical DIMMs to populate the DIMM sockets in the order shown in Table 3-5.

DIMM Replacement Procedure

This section includes the following topics:

Identifying a Faulty DIMM

Each DIMM socket on the motherboard has a corresponding DIMM fault LED. See Internal Diagnostic LEDs, page 3-7 for the locations of the LEDs and how to activate them.

Replacing DIMMs

To install a DIMM assembly, follow these steps:

Step 1 Remove the DIMMs that you are replacing:

a. Power off the server as described in Shutting Down and Powering Off the Server, page 3-8.

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

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

c. Remove the top cover as described in "Removing and Replacing the Server Top Cover" section on page 3-9.

d. Remove the internal air baffle that covers the DIMM sockets.

e. Locate the faulty DIMM and remove it from the socket on the motherboard by opening the ejector levers at both ends of the DIMM socket.

Step 2 Install a new DIMM:

Note Before installing risers or DIMMs, refer to the population guidelines. See DIMM Performance Guidelines and Population Rules, page 3-21.

a. Align the new DIMM with the socket on the motherboard. Use the alignment key in the DIMM socket to correctly orient the DIMM.

b. Push the DIMM into the connector until it is fully seated and the ejector levers on either side of the connector lock into place.

c. Replace the internal air baffle, and then replace the top cover.

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

Replacing CPUs and Heatsinks

The server supports up to four EP-Series CPUs. See Figure 3-13 on page 3-22 for the CPU socket numbering.

- The minimum configuration is two CPUs (CPU1 and CPU2).
- The server must have either two or four CPUs to operate; you should populate CPU1 and CPU2 first, and then populate CPU3 and CPU4.

This section contains the following topics:

Additional CPU-Related Parts To Order With RMA Replacement Motherboards

When a return material authorization (RMA) of the motherboard or CPU is done on a Cisco UCS C-series server, there are additional parts that might not be included with the CPU or motherboard spare bill of materials (BOM). The TAC engineer might need to add the additional parts to the RMA to help ensure a successful replacement.

- Scenario 1—You are re-using the existing heatsinks:
  - Heat sink cleaning kit (UCSX-HSCK=)
  - Thermal grease kit for C420 (UCS-CPU-GREASE=)
  - Intel CPU Pick-n-Place tool for EP CPUs (UCS-CPU-EP-PNP=)
- Scenario 2—You are replacing the existing heatsinks:
  - Heat sink (UCSB-HS-01-EP=)
  - Heat sink cleaning kit (UCSX-HSCK=)
  - Intel CPU Pick-n-Place tool for EP CPUs (UCS-CPU-EP-PNP=)

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

New heatsink spares have preinstalled thermal interface material covered by a small sheet of plastic. It is important to clean the old thermal interface material off of the CPU prior to installing the heatsinks. Therefore, when ordering new heatsinks it is still necessary to order the heatsink cleaning kit at a minimum.

CPU Replacement Procedure

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

Caution The Pick-and-Place tools used in this procedure are required to prevent damage to the contact pins between the motherboard and the CPU. Do not attempt this procedure without the required tools, which are included with each CPU option kit. If you do not have the tool, you can order a spare (Cisco PID UCS-CPU-EP-PNP).

To install or replace a CPU heatsink and CPU, follow these steps:

Step 1 Prepare the server for replacement:

a. Power off the server as described in Shutting Down and Powering Off the Server, page 3-8.

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

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

c. Remove the top cover as described in Removing and Replacing the Server Top Cover, page 3-9.

d. Remove the internal air baffle that covers the CPU sockets.

Step 2 Use a Number 2 Phillips-head screwdriver to loosen the four captive screws that secure the heatsink and lift it off of the CPU.

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

Step 3 Unclip the first CPU retaining latch that is labeled with the icon, and then unclip the second retaining latch that is labeled with the icon. See Figure 3-14.

Step 4 Open the hinged CPU cover plate. See Figure 3-14.

Figure 3-14 CPU Retaining Latches

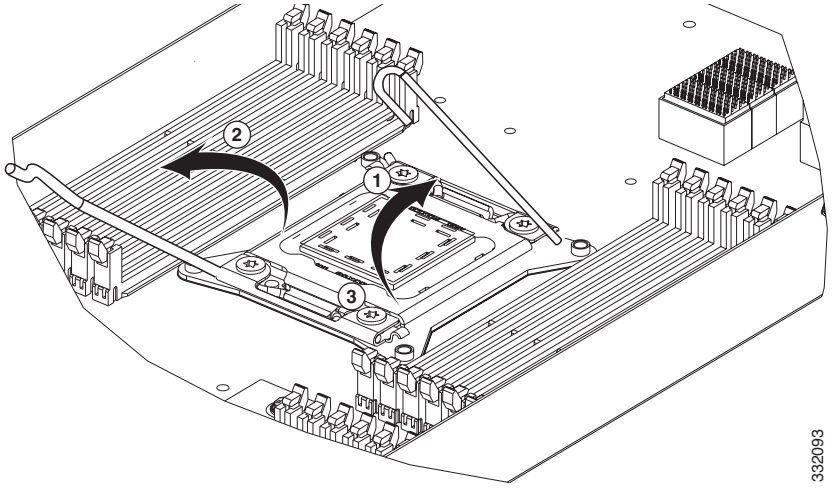

1	CPU retaining latch,	**3**	Hinged CPU cover plate
2	CPU retaining latch,		

Step 5 Remove a protective cap from the socket (if present).

If you are installing a new CPU to a socket that was shipped empty, the socket has a protective cap that is intended to prevent bent contact pins. If you are removing an old CPU instead, skip to Step 6.

Use the tool as shown in Figure 3-15 to grasp the protective cap and then pivot to remove the cap.

Figure 3-15 Protective Cap Removal Tool

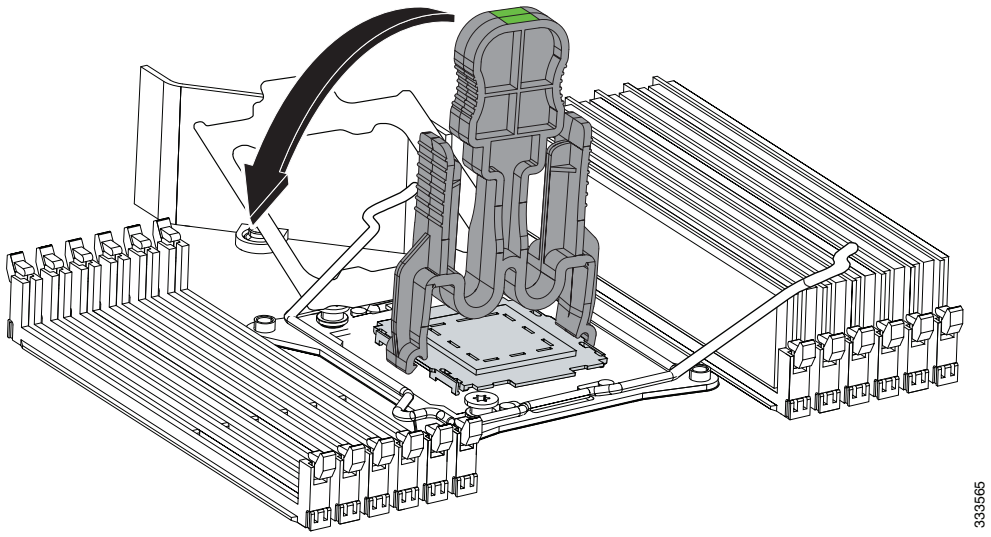

Step 6 Remove the old CPU:

a. Set the Pick-and-Place tool on the CPU in the socket, aligning the arrow on the tool with the registration mark on the socket (the small triangular mark). See Figure 3-16.

b. Press the top button on the tool to grasp the installed CPU.

c. Lift the tool and CPU straight up.

d. Press the top button on the tool to release the old CPU on an antistatic surface.

Figure 3-16 Removing or Inserting a CPU

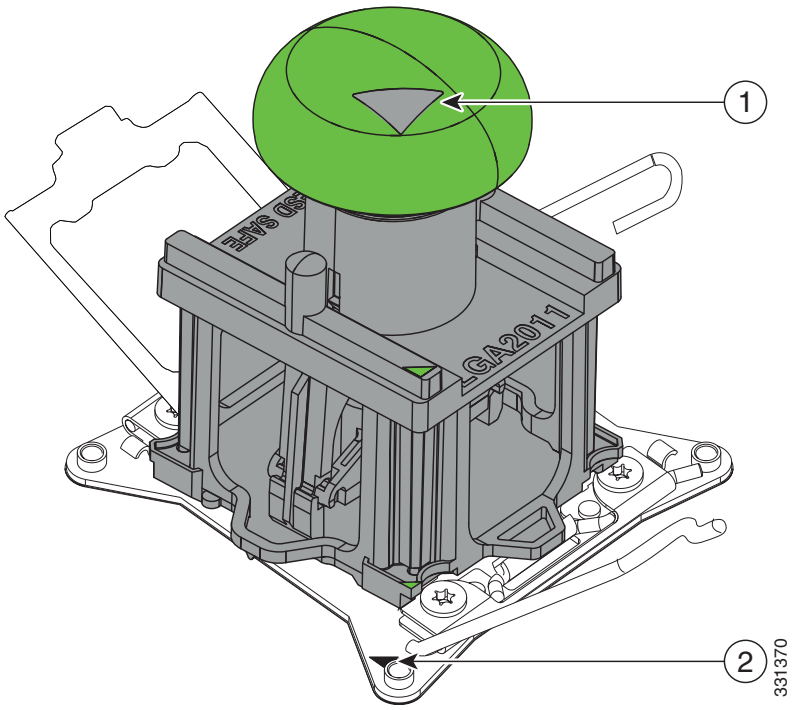

1	Arrow on tool	**2**	Registration mark on CPU socket

Step 7 Insert the new CPU into the Pick-and-Place tool:

a. Remove the new CPU from the packaging and place it on the pedestal that is included in the kit. Align the registration mark on the corner of the CPU with the arrow on the corner of the pedestal (see Figure 3-17).

b. Press down on the top button of the tool to lock it open.

c. Set the Pick-and-Place tool on the CPU pedestal, aligning the arrow on the tool with the arrow on the corner of the pedestal. Make sure that the tabs on the tool are fully seated in the slots on the pedestal.

d. Press the side lever on the tool to grasp and lock in the CPU.

e. Lift the tool and CPU straight up off the pedestal.

Figure 3-17 CPU and Pick-and-Place Tool on Pedestal

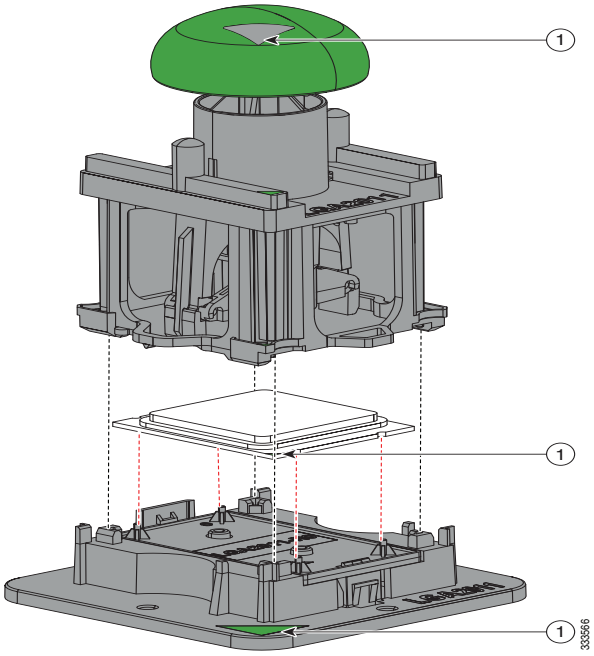

1	Arrow marks for alignment		

Step 8 Install a new CPU:

a. Set the Pick-and-Place tool that is holding the CPU over the empty CPU socket on the motherboard.

Note Align the arrow on the top of the tool with the registration mark (small triangle) that is stamped on the metal of the CPU socket, as shown in Figure 3-16 on page 3-30.

b. Press the top button on the tool to set the CPU into the socket. Remove the empty tool.

c. Close the hinged CPU cover plate.

d. Clip down the CPU retaining latch with the icon, then clip down the CPU retaining latch with the icon. See Figure 3-14 on page 3-29.

Step 9 Install a heatsink:

Caution The heatsink must have a new, undamaged thermal pad on the heatsink-to-CPU surface to ensure proper cooling. If you are reusing a heatsink, you must remove the old thermal pad. If you are installing a new heatsink, skip to Step d. below.

a. Apply an alcohol-based cleaning solution to the old thermal pad and let it soak for a least 15 seconds.

b. Wipe all of the old thermal pad off the old heatsink using a soft cloth that will not scratch the heatsink surface.

c. Apply thermal grease from an included syringe to the top of the CPU.

Apply about 2 cubic centimeters of grease (about half the syringe contents) to the top of the CPU in the pattern that is shown in Figure 3-18.

Note If you do not have a syringe of thermal grease, you can order a spare (Cisco PID UCS-CPU-GREASE).

Figure 3-18 Thermal Grease Application Pattern

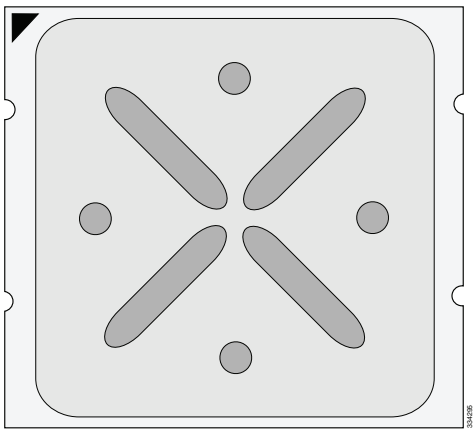

d. For a new heatsink, peel the protective film from the thermal pad that is on the bottom of the new heatsink.

Note Do not apply a syringe of thermal grease if you are installing a new heatsink that already has a preapplied thermal pad.

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

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

f. Replace the internal air baffle and then the top cover.

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

Replacing the Motherboard RTC Battery

Warning **There is danger of explosion if the battery is replaced incorrectly. Replace the battery only with the same or equivalent type recommended by the manufacturer. Dispose of used batteries according to the manufacturer's instructions.** [Statement 1015]

The CMOS real-time clock (RTC) battery retains system settings when the server is disconnected from power. The battery type is Panasonic CR2032 or equivalent.

To replace or install the motherboard CMOS battery, follow these steps:

Step 1 Remove the RTC battery (see Figure 3-19):

a. Power off the server as described in Shutting Down and Powering Off the Server, page 3-8.

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

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

c. Remove the top cover as described in Removing and Replacing the Server Top Cover, page 3-9.

d. Locate the RTC battery. See Figure 3-19.

e. Bend the battery retaining clip away from the battery and use a pair of needle-nose pliers to pull the battery from the socket.

Step 2 Install the RTC battery:

a. Bend the retaining clip away from the battery socket and insert the battery in the socket.

Note The flat, positive side of the battery marked "+" should face the retaining clip.

b. Push the battery into the socket until it is fully seated and the retaining clip clicks over the top of the battery.

c. Replace the top cover.

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

Figure 3-19 Removing and Replacing the Motherboard RTC Battery

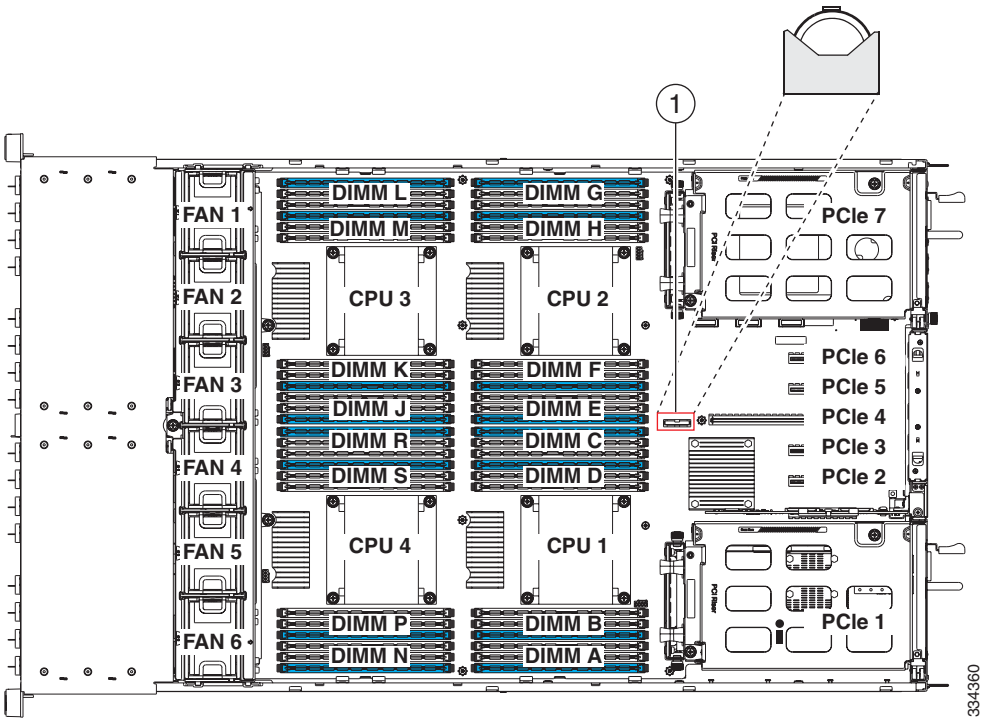

1	RTC battery socket on motherboard	**2**	Retaining clip (enlarged view) The flat, positive (+) side of the battery faces the clip.

Replacing a PCIe Riser Assembly

The server contains two PCIe risers for horizontal installation of PCIe cards. These risers wrap around the power supply bays and connect to dedicated PCIe slots 1 and 7 on the motherboard. These risers each provide one standard-profile, half-length, x16 horizontal slot.

To install or replace a PCIe riser assembly, follow these steps:

Step 1 Remove the PCIe riser assembly that you are replacing (see Figure 3-20):

a. Power off the server as described in Shutting Down and Powering Off the Server, page 3-8.

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

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

c. Remove the top cover as described in Removing and Replacing the Server Top Cover, page 3-9.

d. Use a Number 2 Phillips-head screwdriver to loosen the two captive thumbscrews that secure the PCIe riser assembly.

e. Lift straight up on both ends of the PCIe riser assembly to disengage its circuit board from the socket on the motherboard.

Step 2 Install a new PCIe riser assembly:

a. Set the PCIe riser assembly back in place over the power supply bay.

b. Align the circuit board edge with the socket on the motherboard and then push straight down on both ends of the assembly to fully engage the board with the socket.

c. Tighten the two captive thumbscrews that secure the riser assembly in place.

d. Replace the top cover.

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

Figure 3-20 Removing and Replacing the PCIe Riser Assembly

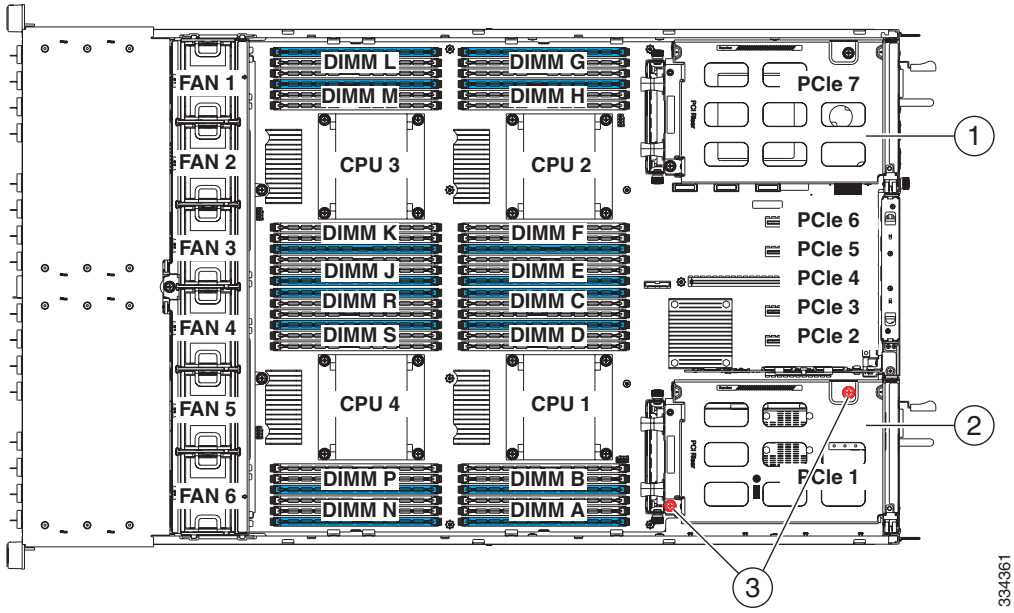

1	PCIe riser 1 (PCIe slot 7)	**3**	PCI riser captive thumbscrews (two)
2	PCIe riser 2 (PCIe slot 1)		

Replacing an I/O Riser

The I/O riser is a modular PCIe form-factor card that has a designated motherboard socket inside the server (see Figure 3-22). The module provides the external VGA video and RS-232 serial connectors for the rear panel of the server. See Figure 3-21.

The I/O riser also provides two internal slots for Cisco FlexFlash cards. For more information about Cisco FlexFlash cards, see Replacing a Cisco Flexible Flash Drive, page 3-39.

Figure 3-21 I/O Riser, Including Cisco FlexFlash Slots

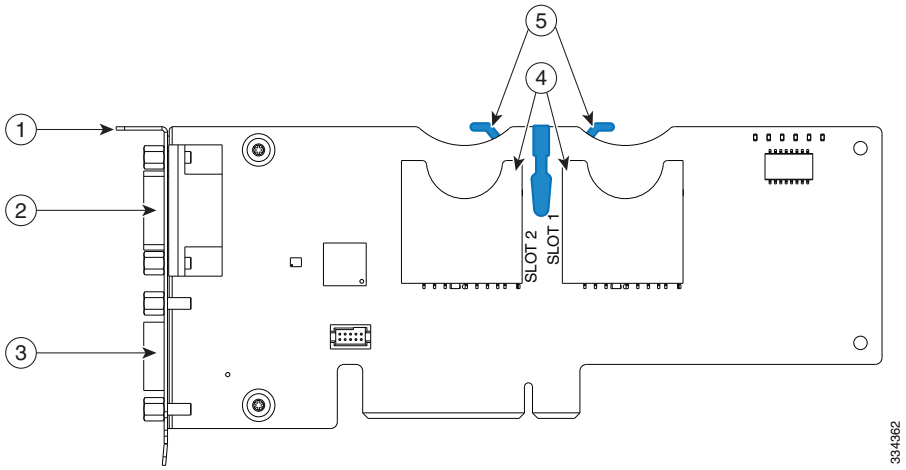

1	I/O riser rear panel tab	**4**	Internal Cisco FlexFlash card slots
2	RS-232 serial connector (9-pin)	**5**	Cisco FlexFlash card retainer
3	VGA video connector (15-pin)		

To install or replace an I/O riser, follow these steps:

Step 1 Remove the I/O riser that you are replacing (see Figure 3-22):

a. Power off the server as described in Shutting Down and Powering Off the Server, page 3-8.

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

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

c. Remove the top cover as described in Removing and Replacing the Server Top Cover, page 3-9.

d. Disconnect any external cables from the I/O riser ports.

e. Use a Number 1 Phillips-head screwdriver to remove the single screw that secures the I/O riser rear-panel tab to the chassis.

f. Lift the I/O riser straight up from the motherboard connector.

Note Lift up on both ends of the I/O riser evenly to avoid damaging its connector.

Step 2 Install a new I/O riser:

a. Align the new I/O riser with its designated socket on the motherboard. See Figure 3-22.

Note You must install the I/O riser in its designated motherboard socket.

b. Push down evenly on both ends of the I/O riser until it is fully seated in its motherboard socket.

c. Ensure that the I/O riser rear panel tab sits flat against the chassis rear panel opening.

d. Install the single screw that secures the I/O riser to the chassis.

e. Replace the top cover.

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

Figure 3-22 Removing and Replacing an I/O Riser

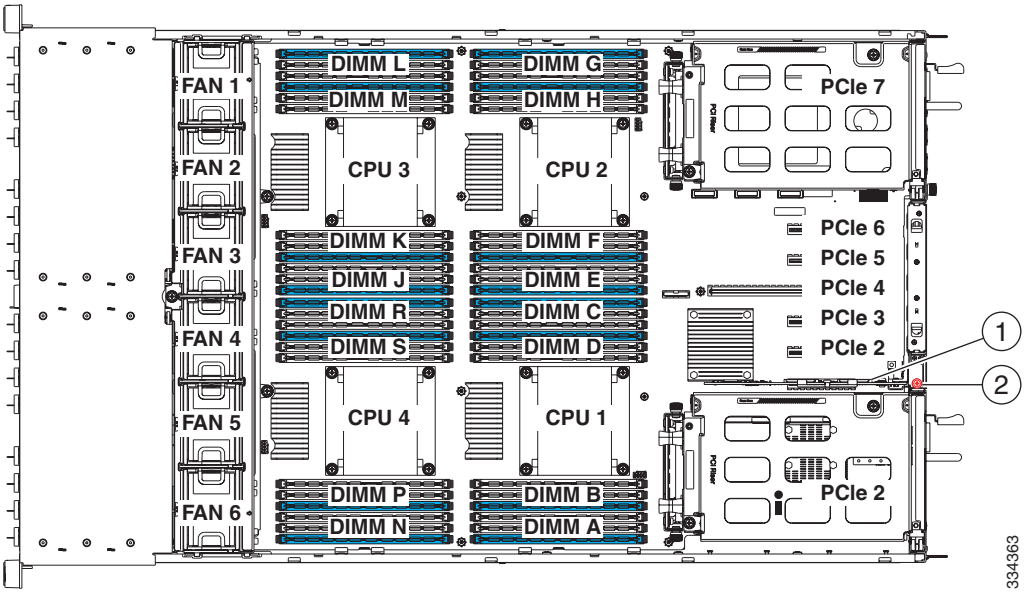

1	I/O riser in designated motherboard socket	**2**	I/O riser securing screw

Replacing a Cisco Flexible Flash Drive

This section includes the following topics:

Overview of the Preinstalled Cisco FlexFlash Drive

The server is shipped from the factory with one preinstalled Cisco FlexFlash drive. You can add a second drive for a mirrored configuration. The slots for these cards are on the I/O riser (see Replacing an I/O Riser, page 3-37).

The Cisco FlexFlash drive is preloaded with three software bundles, each on one of four preconfigured virtual drives (VDs). The fourth VD allows you to install an operating system or embedded hypervisor.

The VDs are configured as follows:

1. Cisco UCS Server Configuration Utility (SCU).
2. Hypervisor (HV). This is a VD that you can use for your own purposes.
3. Cisco Drivers (Drivers).
4. Cisco Host Upgrade Utility (HUU).

Enabling a Cisco FlexFlash Virtual Drive

Each of the VDs on the preinstalled Cisco FlexFlash drive can be separately enabled or hidden from the host. The default as shipped from the factory is for all VDs to be hidden.

To enable VDs and expose them to the host, follow these steps:

Step 1 Log in to CIMC interface for the server using the IP address of the server.

Step 2 Navigate through GUI tabs to **Server > Inventory > Storage**.

Step 3 In the Storage Adapters list, click **FlexFlash-0**.

Step 4 In the Actions area, click on **Configure Operational Profile**.

The Operational Profile dialog opens.

Step 5 Check the box for each VD that you want to enable and expose to the host, and then click **Save Changes**.

Booting a Cisco FlexFlash Virtual Drive

When you want to access the Cisco SCU or Cisco HUU software, you boot its VD with a one-time boot option. When you want to boot the hypervisor (HV) VD, you boot it with a permanent boot order selection. This section includes the following topics:

Booting the Cisco SCU or Cisco HUU Software VDs

You can access the preinstalled Cisco SCU or Cisco HUU software bundles on a Cisco FlexFlash drive by booting their respective VDs with a one-time boot option:

To boot the VD, follow these steps:

Step 1 Enable the SCU or HUU VD.

See Enabling a Cisco FlexFlash Virtual Drive, page 3-39.

Step 2 In the CIMC GUI interface, navigate through the tabs to **Server > BIOS.**

Step 3 Click **Configure Boot Override Priority**.

The Boot Override Priority dialog opens.

Step 4 From the drop-down list, choose **SCU** or **HUU**, and then click **OK**.

Step 5 Reboot the server.

The server boots the selected VD.

Note This is a one-time boot option. After running Cisco SCU or Cisco HUU, the server returns to its previously configured boot path.

Booting the Hypervisor VD

You can boot the hypervisor (HV) VD with a more permanent boot selection. To change the boot order one time, use the procedure in Booting the Cisco SCU or Cisco HUU Software VDs, page 3-40.

To permanently set the boot order for an HV VD, follow these steps:

Step 1 Boot the server and watch for the prompt to press F2 to open the BIOS Setup utility.

Step 2 When prompted, press **F2** to open the BIOS Setup utility.

Step 3 Navigate to the **Boot Options** tab.

Step 4 Use the Boot Options screen to set the HV VD to your desired boot order for the server.

Monitoring and Managing a Cisco FlexFlash Drive

You can monitor and manage your Cisco FlexFlash drives by using the CIMC GUI interface or the CLI interface. See the *Cisco UCS C-Series Rack-Mount Server Configuration Guide* or the *Cisco UCS C-Series Rack-Mount Server CLI Configuration Guide* in the documentation roadmap link below.

http://www.cisco.com/go/unifiedcomputing/c-series-doc

Synchronizing RAID After Installing a Second Cisco FlexFlash Drive

After you install or replace a second Cisco FlexFlash drive, you must synchronize the RAID partition by using the Cisco UCS Server Configuration Utility (SCU).

The SCU provides an option to synchronize the Hypervisor VD, configured as a RAID-1 disk. This feature is available only when both Cisco FlexFlash drive slots are populated.

When one member slot of the SD card is corrupt, use this option to synchronize the hypervisor data across two members of the RAID-1 virtual disk. You can initiate this synchronization only if two cards are detected and the RAID-1 group is determined as unhealthy (one member is corrupt).

Step 1 Click the **Hypervisor Sync** icon on the toolbar of the SCU interface.

A dialog prompts you to confirm that you want to synchronize the Hypervisor RAID.

Step 2 Click **Yes**.

A dialog is displayed when the synchronization is complete.

Step 3 Click **OK**.

After you click OK, the Hypervisor Sync icon on the toolbar is greyed out.

For more information about the utility, see the Cisco UCS Server Configuration Utility User Guide.

Cisco FlexFlash Drive Replacement Procedure

To install or replace a Cisco FlexFlash drive, follow these steps:

Step 1 Remove the Cisco FlexFlash drive that you are replacing (see Figure 3-22):

a. Power off the server as described in Shutting Down and Powering Off the Server, page 3-8.

b. Slide the server out the front of the rack far enough so that you can remove the top cover. You might have to detach cables from the rack.

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

c. Remove the top cover as described in Removing and Replacing the Server Top Cover, page 3-9.

d. Locate the Cisco FlexFlash drive that you are replacing on the I/O riser card. See Figure 3-22.

e. Pull back the card retainer slightly (see Figure 3-21), and then push down on the top of the Cisco FlexFlash drive and release it to allow it to spring up in the socket.

f. Remove the Cisco FlexFlash drive from the socket.

Step 2 Install a Cisco FlexFlash drive:

Note To be usable for Cisco FlexFlash, an SD card must be at least 16 GB in size.

Note Any SD card that is installed into the Cisco FlexFlash slot is configured with the VD partitioning described in Overview of the Preinstalled Cisco FlexFlash Drive, page 3-39. Data is overwritten on the SD card where Cisco stores the configuration metadata.

a. Pull the card retainer back lightly, then insert the Cisco FlexFlash drive into SD card slot 1 on the I/O riser with the label side facing outward.
b. Press down on the top of the card until it clicks in the slot and stays in place.
c. Replace the top cover.
d. Replace the server in the rack, replace cables, and then power on the server by pressing the **Power** button.

Replacing a PCIe Card

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

This server has seven PCIe expansion slots. See Figure 3-23 and Table 3-6 for information about the slots. The replacement procedures differ depending on whether the PCIe slot is in a riser or on the motherboard. This section includes the following topics:

PCIe Slots

Figure 3-23 PCIe Slot Locations

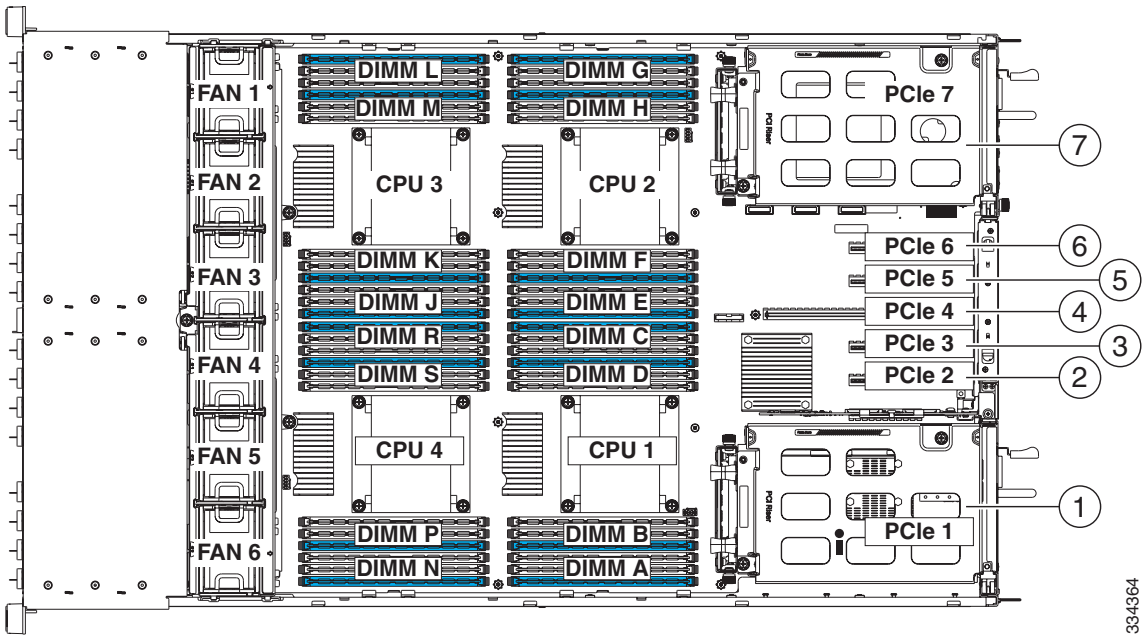

Table 3-6 PCIe Expansion Slots

Slot Number	Electrical Lane Width	Connector Length	Card Length[1]	Card Height [2]	NCSI[3] Support
1 (on riser)	Gen-3 x16	x16 connector	½ length	Full height	Yes
2	Gen-3 x8	x8 connector	½ length	½ height	No
3	Gen-3 x8	x8 connector	½ length	½ height	No
4	Gen-3 x8	x16 connector	½ length	½ height	Yes
5	Gen-3 x8	x8 connector	½ length	½ height	No
6	Gen-3 x8	x8 connector	½ length	½ height	No
7 (on riser)	Gen-3 x16	x16 connector	½ length	Full height	Yes

1. This is the supported length because of internal clearance.
2. This is the size of the rear panel opening.
3. NCSI = Network Communications Services Interface protocol. An NCSI slot is powered even when the server is in standby power mode.

PCIe Configuration Guide

For the best performance, we recommend that you populate the PCIe slots in the order shown in Table 3-7 for each type of add-on card. For each card type, populate the primary slot first, followed by the secondary slot, and then any alternate slots. See Figure 3-23 for the slot locations.

Table 3-7 PCIe Slot Population

PCIe Card Type	Primary Slot	Secondary Slot	Alternate Slots
RAID controller	Slot 3	Slot 5	—
Low-profile (half-height) network adapter	Slot 2	Slot 6	Slots 1, 4, 5, or 7[1]
Cisco UCS Virtual Interface Card	Slot 4	Slot 1 or 7	—

1. To use a half-height card in slots 1 or 7, you must have a full-height rear panel attached to the card.

Replacing a PCIe Card in a Riser Slot

Note If you are installing a Cisco UCS Virtual Interface Card (VIC), see Special Considerations for Cisco UCS Virtual Interface Cards, page 3-49.

To install or replace a PCIe card in full-height slots 1 or 7 on the PCIe risers, follow these steps:

Step 1 Remove a PCIe card (or a blank filler panel) from the PCIe riser assembly:

a. Shut down and power off the server as described in Shutting Down and Powering Off the Server, page 3-8.

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

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

c. Remove the top cover as described in Removing and Replacing the Server Top Cover, page 3-9.

d. Use a Number 2 Phillips-head screwdriver to loosen the two captive thumbscrews that secure the PCIe riser assembly. See Figure 3-20 on page 3-36.

e. Lift straight up on both ends of the PCIe riser assembly to disengage its attached circuit board from the socket on the motherboard.

Note Lift up on both ends of the PCIe riser evenly to avoid damaging its connector.

f. Pinch and lift up on the blue-plastic card-retaining clip that secures the card (or filler panel) into the riser.

g. Pull evenly on both corners of the PCIe card to remove it from the socket on the PCIe riser.

Step 2 Install a PCIe card:

a. Align the new PCIe card with the empty socket on the PCIe riser assembly.

b. Push down evenly on both ends of the card until it is fully seated in the socket.

c. Ensure that the card's rear-panel tab sits flat against the PCIe riser rear panel opening, and then close the blue-plastic card-retaining clip to secure the card in the riser (see Figure 3-24).

d. Set the PCIe riser assembly back in place over the power supply bay.

e. Align the PCIe riser circuit board edge with the socket on the motherboard and then push straight down on both ends of the PCIe riser assembly to fully engage the board with the socket.

f. Tighten the two captive thumbscrews that secure the PCIe riser assembly in place.

g. Replace the top cover.

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

Figure 3-24 PCIe Riser Assembly Side View

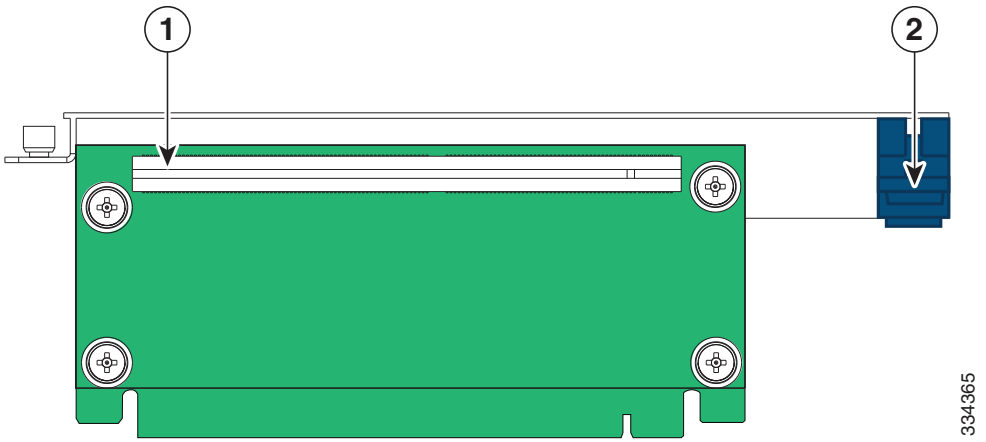

1	PCIe card socket on PCIe riser assembly	**2**	Card retaining clip

Replacing a PCIe Card in a Motherboard Slot

Note If you replace an LSI MegaRAID controller card, you must restore your RAID configuration to the new card. See Restoring RAID Configuration After Replacing a RAID Controller, page C-6.

To install or replace a PCIe card in low-profile motherboard slots 2 through 6, follow these steps:

Step 1 Remove a PCIe card (or a blank filler panel):

a. Shut down and power off the server as described in Shutting Down and Powering Off the Server, page 3-8.

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

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

c. Remove the top cover as described in Removing and Replacing the Server Top Cover, page 3-9.

d. Open the hinged PCIe card retainer. Pinch the two release latches toward the center of the retainer while you lift up on the front edge of the retainer to open it. See Figure 3-25.

e. Lift the card straight up from the motherboard connector.

Note Lift up on both ends of the card evenly to avoid damaging its connector.

Step 2 Install a PCIe card:

a. Align the PCIe card with the empty PCIe connector on the motherboard.

b. Push down evenly on both ends of the card until it is fully seated in the motherboard connector.

c. Ensure that the card rear panel sits flat against the chassis rear panel opening.

d. Close the hinged PCIe retainer and push down on it in the closed position until the latches lock in place.

e. Replace the top cover.

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

Step 3 If the card that you replaced is a mass storage controller, restore the RAID configuration on your drives to the new mass storage controller.

See Restoring RAID Configuration After Replacing a RAID Controller, page C-6.

Figure 3-25 Hinged PCIe Card Retainer on Rear of Chassis

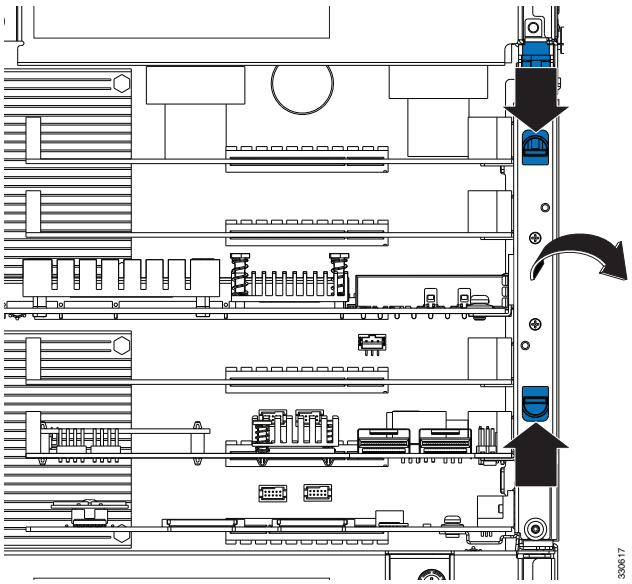

Special Considerations for Cisco UCS Virtual Interface Cards

Table 3-8 describes the requirements for the supported Cisco UCS virtual interface cards (VICs).

Table 3-8 Cisco UCS C420 Requirements for Virtual Interface Cards

Virtual Interface Card (VIC)	Number of VICs Supported in Server	Slots That Support VICs [1]	Primary Slot For UCS Integration or Cisco Card NIC Mode	Minimum CIMC Firmware	Minimum VIC Firmware for Use With UCS Integration	Minimum Cisco Nexus OS on an Upstream Nexus Fabric Interconnect
Cisco UCS VIC 1225 UCSC-PCIE-CSC-02	3	PCIE 4 PCIE 1 PCIE 7	PCIE 4	1.4(7)	2.1(0)	5.0
Cisco UCS VIC1225T UCSC-PCIE-C10T-02	3	PCIE 4 PCIE 1 PCIE 7	PCIE 4[2]	1.5(1)	2.1(1)	5.0

1. See PCIe Slots, page 3-43.
2. The Cisco UCS VIC1225T is not supported for UCS integration at this time.

Installing Multiple PCIe Cards and Resolving Limited Resources

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

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

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

- Resolving Insufficient Memory Space to Execute Option ROMs, page 3-50
- Resolving Insufficient 16-Bit I/O Space, page 3-51

Resolving Insufficient Memory Space to Execute Option ROMs

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

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

```
PCI OUT OF RESOURCES CONDITION:
ERROR: Insufficient PCI Resources Detected!!!

System is running with Insufficient PCI Resources!
In order to display this message some
PCI devices were set to disabled state!
It is strongly recommended to power off the system
and remove some PCI/PCI Express cards from the system!
To continue booting, proceed to <Save & Exit> Menu Option
and select Boot Device or <Discard Changes and Exit>.

WARNING: If you choose to continue booting some Operating
Systems might not be able to complete boot correctly!
```

To resolve this issue, disable the option ROMs that are not needed for system booting. The BIOS Setup Utility provides the setup options to enable or disable the option ROMs at the PCIe slot level for the PCIe expansion slots and at the port level for the onboard NICs. These options can be found in the BIOS Setup Utility **Advanced →PCI Configuration** page (see Figure 3-26).

Figure 3-26 BIOS Setup Utility Advanced > PCI Configuration Page

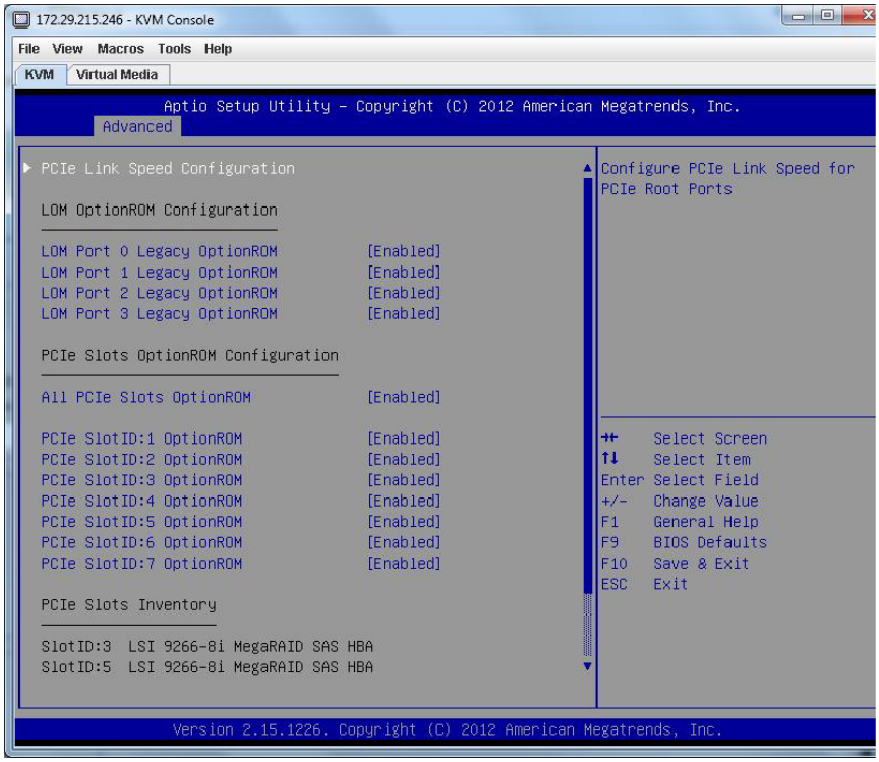

- Guidelines for RAID controller booting

  When the server is configured in the factory, the RAID controllers are installed in PCIe slot 3 or slot 5 for single RAID controller configuration, or in both slot 3 and slot 5 for dual RAID controller configuration.

  If the server is configured to boot primarily from RAID storage, make sure that the option ROMs for slot 3 or slot 5 are enabled in the BIOS, depending upon your RAID controller configuration. If the RAID controller does not appear in the system boot order even with the option ROMs for slot 3 and slot 5 enabled, the RAID controller option ROM might not have sufficient memory space to execute. In that case, disable other option ROMs that are not needed for the system configuration to free up some memory space for the RAID controller option ROM.

- Guidelines for onboard NIC PXE booting

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

Resolving Insufficient 16-Bit I/O Space

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

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

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

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

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

   Uncheck the “Enable PXE boot entry” in the vNIC properties dialog to disable the PXE boot, as shown in Figure 3-27.

Figure 3-27 vNIC Properties Dialog

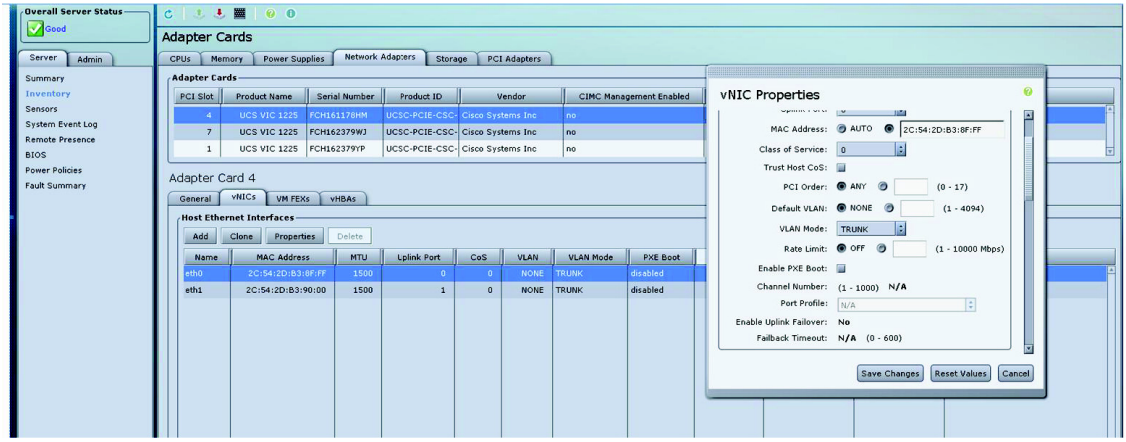

Replacing the Supercap Power Module (RAID Backup Unit)

The server supports installation of up to two RAID controllers and up to two supercap power modules (SCPMs). These SCPMs mount to trays that are on the removable air baffle.

The SCPM is available only when using the LSI MegaRAID-CV controller cards. This module provides approximately 3 years of backup for the disk write-back cache DRAM in the case of sudden power loss by offloading the cache to the NAND flash.

Warning **There is danger of explosion if the battery is replaced incorrectly. Replace the battery only with the same or equivalent type recommended by the manufacturer. Dispose of used batteries according to the manufacturer's instructions.**
Statement 1015

To replace the supercap power module, follow these steps:

Step 1 Remove an SCPM (see Figure 3-28):

a. Power off the server as described in Shutting Down and Powering Off the Server, page 3-8.

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

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

c. Remove the top cover as described in Removing and Replacing the Server Top Cover, page 3-9.

d. Disconnect the cable that is attached to the SCPM.

e. Remove the SCPM tray from the air baffle. Pinch the retaining clip on the side of the tray and slide it toward the rear to remove the tray from the air baffle.

Note There are two trays, an upper tray and a lower tray. Remove only the tray that holds the SCPM that you are replacing.

f. Remove the SCPM from the clips on the SCPM tray.

Step 2 Install a new SCPM:

a. Slide the new SCPM into the clips on the SCPM tray.

b. Replace the SCPM tray to the air baffle. Slide it forward in the air baffle guides until the tray retaining clip clicks and locks in place.

c. Connect the cable from the RAID controller to the new SCPM.

d. Replace the top cover.

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

Figure 3-28 Removing the SCPM

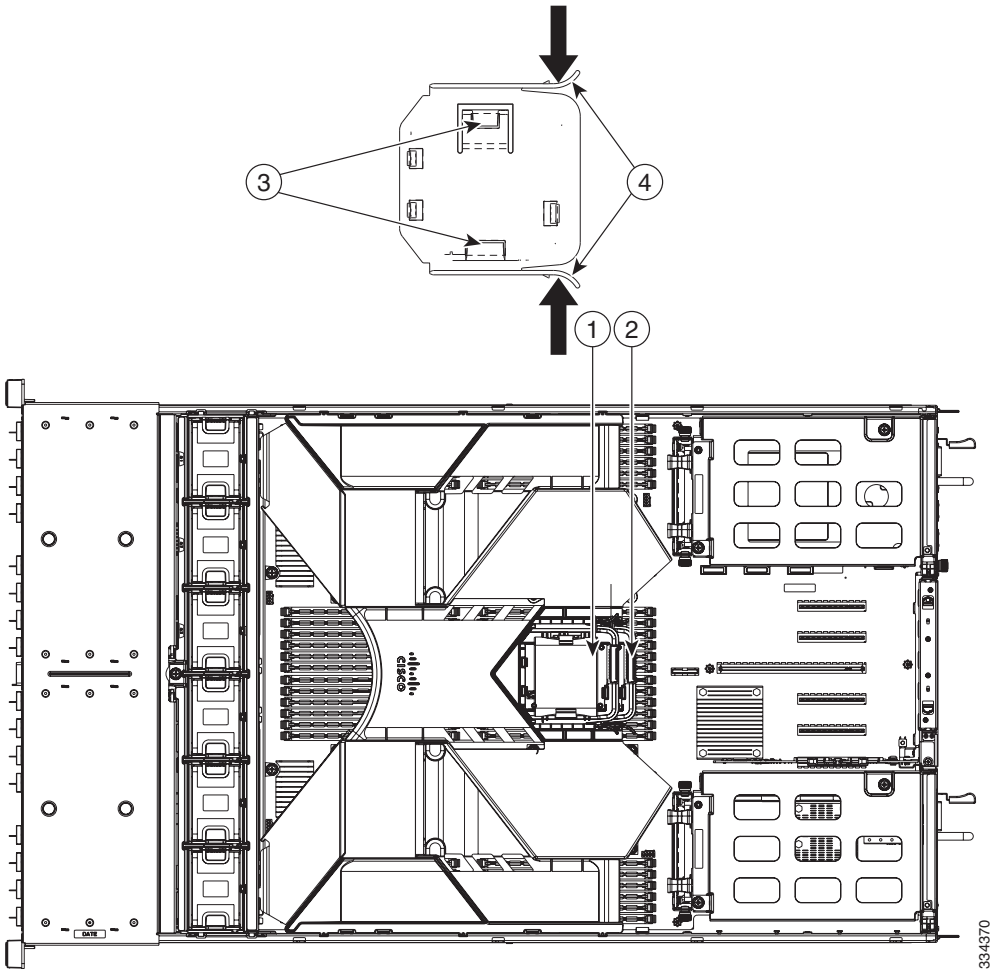

1	Upper SCPM tray on air baffle	**3**	SCPM retaining clips on tray
2	Lower SCPM tray on air baffle	**4**	SCPM tray release levers

Installing a Trusted Platform Module

The trusted platform module (TPM) is a small circuit board that attaches to a motherboard socket. The socket location is on the motherboard, underneath the cage for power supply 1.

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

To install a trusted platform module (TPM), follow these steps:

Step 1 Prepare the server for component installation:

a. Power off the server as described in Shutting Down and Powering Off the Server, page 3-8.

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

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

c. Remove the top cover as described in Removing and Replacing the Server Top Cover, page 3-9.

d. Remove the PCIe riser 1, which contains PCIe slot 1 (see Figure 3-23 on page 3-43). See Replacing a PCIe Riser Assembly, page 3-35 for details.

e. Remove power supply 1 from its cage. See Replacing Power Supplies, page 3-58 for details.

Step 2 Install a TPM (see Figure 3-29):

a. Locate the TPM socket on the motherboard, as shown in Figure 3-29.

b. Align the connector that is on the bottom of the TPM circuit board with the motherboard TPM socket. Align the screw hole and standoff on the TPM board with the screw hole that is adjacent to the TPM socket.

c. Push down evenly on the TPM to seat it in the motherboard socket.

d. Install the single one-way screw that secures the TPM to the motherboard.

e. Replace power supply 1 to the empty bay.

f. Replace the PCIe riser 1.

g. Replace the top cover.

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

Step 3 Enable the TPM:

a. Watch during bootup for the F2 prompt, and then press **F2** to enter BIOS setup.

b. Log into the BIOS Setup utility with your BIOS Administrator password.

c. On the BIOS Setup utility screen, select the **Advanced** tab.

d. Select **Trusted Computing** to open the TPM Security Device Configuration screen.

e. Change TPM SUPPORT to **Enabled**.

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

Step 4 Verify that the TPM is now enabled.

a. Watch during bootup for the F2 prompt, and then press **F2** to enter BIOS setup.

b. Log into the BIOS Setup utility with your BIOS Administrator password.

c. Select the **Advanced** tab.

d. Select **Trusted Computing** to open the TPM Security Device Configuration screen.

e. Verify that TPM SUPPORT is Enabled.

Note If you want to use the Intel Trusted Execution Technology (TXT) feature, you must enable it in the server BIOS as described in Enabling the Intel Trusted Execution Technology Feature For the TPM, page 3-56.

Figure 3-29 TPM Socket Location on Motherboard

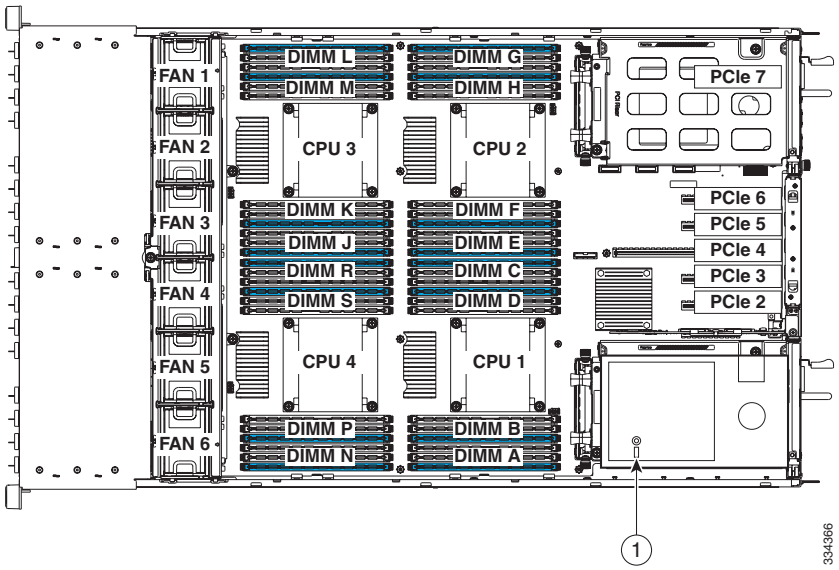

1	TPM socket and screw hole on motherboard (visible through power supply cage with PCIe riser 1 and power supply 1 removed)		–

Enabling the Intel Trusted Execution Technology Feature For the TPM

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

To enable the TXT feature, follow these steps:

Step 1 Verify that a TPM is now installed and enabled in the server:

a. Either attach a VGA monitor and USB keyboard to the server, or log in remotely to the CIMC interface of the server and open a virtual KVM console window.

b. Reboot the server.

c. Watch during bootup for the F2 prompt, and then press **F2** to enter BIOS setup.

d. Log in to the BIOS Setup utility with your BIOS Administrator password.

Note You must be logged in as the BIOS administrator to perform this procedure. If you have not done so already, set a BIOS administrator password on the Security tab of the BIOS Setup utility.

e. Select the **Advanced** tab.

f. On the Advanced tab, select **Trusted Computing** to open the TPM Security Device Configuration screen.

g. Verify that **TPM SUPPORT** is **Enabled**. If it is not, set TPM SUPPORT to Enabled.

h. Press **Escape** to return to the BIOS Setup utility **Advanced** tab.

Step 2 Enable the Intel Trusted Execution Technology (TXT) feature:

a. On the Advanced tab, select **Intel TXT(LT-SX) Configuration** to open the Intel TXT(LT-SX) Hardware Support screen.

Note The Intel Trusted Execution Technology feature can be enabled only when the server has a TPM installed on the TPM header.

b. Set **TXT Support** to **Enabled**.

Step 3 On the same screen, verify that the Intel Virtualization Technology (VT) and the Intel VT for Directed I/O (VT-d) features are enabled (the factory default).

a. On the Intel TXT(LT-SX) Hardware Support screen, verify that **VT-d Support** and **VT Support** are both listed as **Enabled**.

- If they are already enabled, skip to Step 4.
- If VT-d Support and VT Support are not enabled, continue with the next steps to enable them.

b. Press **Escape** to return to the BIOS Setup utility **Advanced** tab.

c. On the Advanced tab, select **Processor Configuration** to open the Processor Configuration screen.

d. Set **Intel (R) VT** and **Intel (R) VT-d** to Enabled.

Step 4 Press **F10** to save your changes and exit the BIOS Setup utility.

Step 5 Verify that the Intel TXT, VT, and VT-d features are enabled:

a. Reboot the server.

b. Watch during bootup for the F2 prompt, and then press **F2** to enter BIOS setup.

c. Select the **Advanced** tab.

d. Select **Intel TXT(LT-SX) Configuration** and verify that TXT Support, VT-d Support, and VT Support are Enabled.

Replacing Power Supplies

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

To replace or install a power supply, follow these steps:

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 power supplies because they are redundant as 1+1.

Step 1 Remove the power supply that you are replacing or a blank panel from an empty bay (see Figure 3-30):

a. Perform one of the following actions:
   - If your server has only one power supply, shut down and power off the server as described in the "Shutting Down and Powering Off the Server" section on page 3-8.
   - If your server has two power supplies, you do not have to shut down the server.

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

c. Grasp the power supply handle while pinching the release lever towards 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. If you shut down the server, press the **Power** button to return the server to main power mode.

Figure 3-30 Removing and Replacing Power Supplies

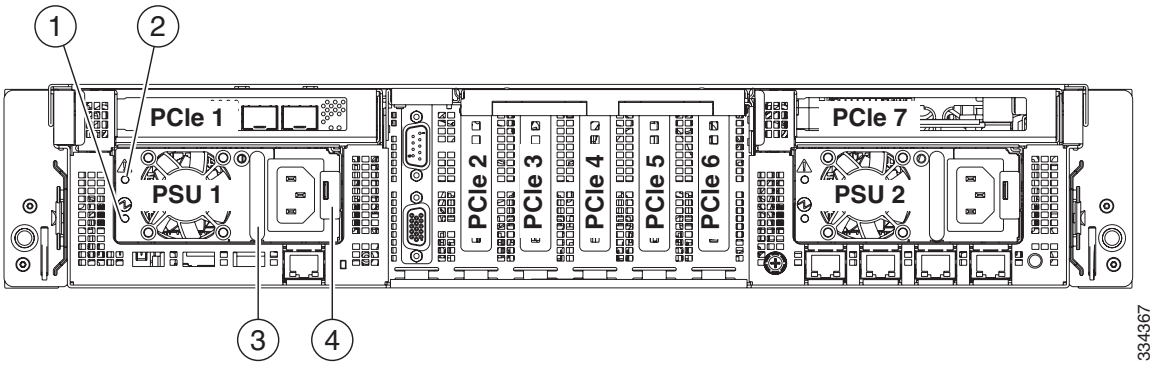

1	Power supply status LED	**3**	Power supply handle
2	Power supply fault LED	**4**	Power supply release lever

Enabling or Disabling the Internal USB Port

The factory default is for all USB ports on the server to be enabled. However, the internal USB port can be enabled or disabled in the server BIOS. To enable or disable the internal USB port, follow these steps:

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

Step 2 Navigate to the **Advanced** tab.

Step 3 On the Advanced tab, select **USB Configuration**.

Step 4 On the USB Configuration page, select **USB Ports Configuration**.

Step 5 Scroll to **USB Port: Internal**, press **Enter**, and then select either Enabled or Disabled from the pop-up menu.

Step 6 Press F10 to save and exit the utility.