

Servicing the Compute Node

This chapter contains the following topics:

- Removing and Installing the Compute Node Cover, on page 1
- Internal Components, on page 3
- Replacing a Drive, on page 4
- Replacing the Front Mezzanine Module, on page 11
- Servicing the Mini Storage Module, on page 17
- Replacing the SuperCap Module, on page 24
- Replacing CPUs and Heatsinks, on page 32
- Replacing Memory DIMMs, on page 46
- Servicing the mLOM, on page 52
- Servicing the Rear Mezzanine, on page 55
- Servicing the Bridge Card, on page 58
- Servicing the Trusted Platform Module (TPM), on page 61

Removing and Installing the Compute Node Cover

The top cover for the Cisco UCS X215c M8 compute node can be removed to allow access to internal components, some of which are field-replaceable. The green button on the top cover releases the compute node so that it can be removed from the chassis.

- Removing a Compute Node Cover, on page 1
- Installing a Compute Node Cover , on page 2

Removing a Compute Node Cover

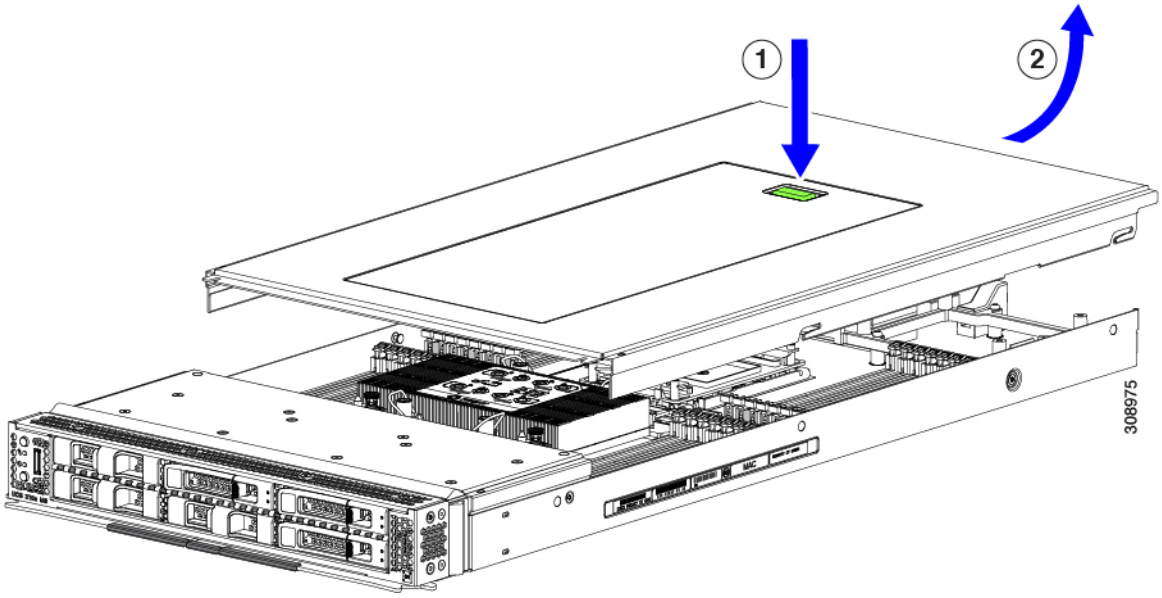
To remove the cover of the UCS X215c M8 compute node, follow these steps:

Procedure

Step 1 Press and hold the button down (1, in the figure below).

Step 2 While holding the back end of the cover, slide it back, then pull it up (2).

By sliding the cover back, you enable the front edge to clear the metal lip on the rear of the front mezzanine module.



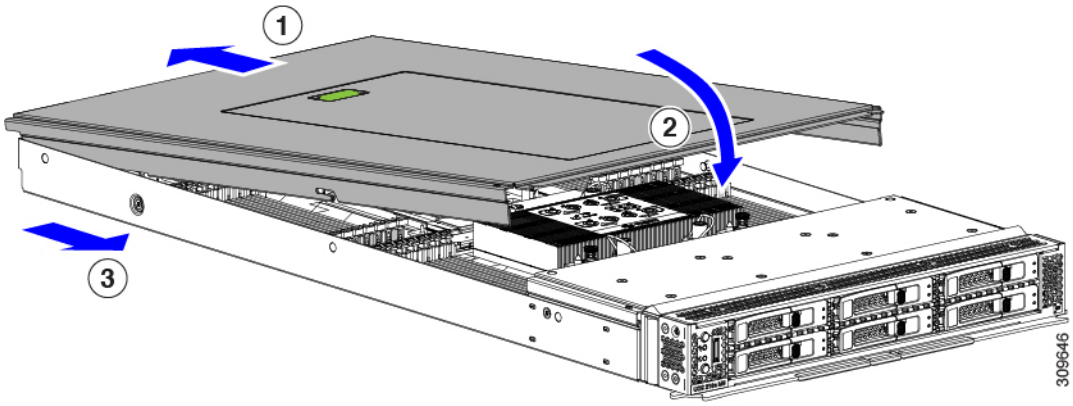
Installing a Compute Node Cover

Use this task to install a removed top cover for the UCS X215c M8 compute node.

Procedure

Step 1 Insert the cover angled so that it hits the stoppers on the base.

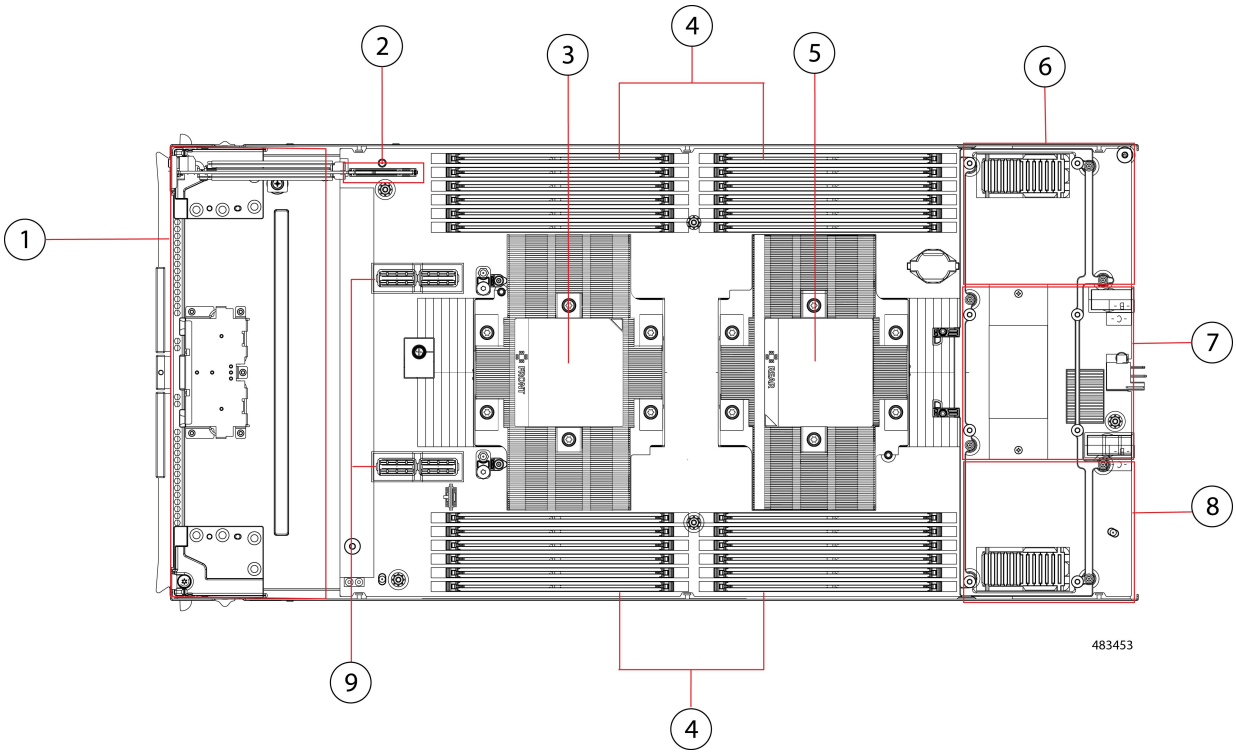
Step 2 Lower the compute node's cover until it reaches the bottom.



Step 3 Keeping the compute node's cover flat, slide it forward until the release button clicks.

Internal Components

The following illustration shows the location of internal components on the compute node.



1	Front mezzanine module slot	**2**	Mini-Storage module connector, which supports one mini-storage module with up to two M.2 SATA or M.2 NVMe drives.
3	CPU 1, which supports one Fourth Gen AMD EPYC™ Processors with up to 96 cores per processor and up to 384 MB of Level 3 cache. CPU 1 must always be populated.	**4**	DIMM Slots, which support up to 6 TB of main memory through 24 256 GB DDR5 5600 MT/s or DDR5 4800 MT/s DIMMs depending on the CPU installed.
5	CPU 2, which supports up to one Fourth Gen AMD EPYC™ Processors with up to 96 cores per processor and up to 384 MB of Level 3 cache. Although not optimal, this CPU can be unpopulated.	**6**	Rear mezzanine slot, which supports X-Series mezzanine cards, such as VIC 15422.

7	Bridge Card slot, which connects the rear mezzanine slot and the mLOM/VIC slot	**8**	mLOM/VIC slot that supports zero or one Cisco VIC or Cisco X-Series 100 Gbps mLOM, such as the Cisco UCS VIC 15420 or VIC 15230
9	Front Mezzanine Connectors, which can support a PCIe 4.0 front mezzanine module consisting of: • Up to 6 hot-pluggable, front-loading State Drives (SSDs) or Non-volatile Memory Express (NVMe) 2.5-inch drives.		

Replacing a Drive

You can remove and install some drives without removing the compute node from the chassis. All drives have front-facing access, and they can be removed and inserted by using the ejector handles.

The SAS/SATA or NVMe drives supported in this compute node come with the drive sled attached. Spare drive sleds are not available.

Before upgrading or adding a drive to a running compute node, check the service profile through Cisco UCS management software and make sure the new hardware configuration will be within the parameters allowed by the management software.



Caution To prevent ESD damage, wear grounding wrist straps during these procedures.

NVMe SSD Requirements and Restrictions

For 2.5-inch NVMe SSDs, be aware of the following:

- NVMe 2.5 SSDs support booting only in UEFI mode. Legacy boot is not supported.
- NVME U.3 SSDs connect to the RAID controller so RAID is supported for these drives.
- UEFI boot is supported in all supported operating systems.

Enabling Hot Plug Support

Only OS-informed hot plug is supported.

Removing a Drive

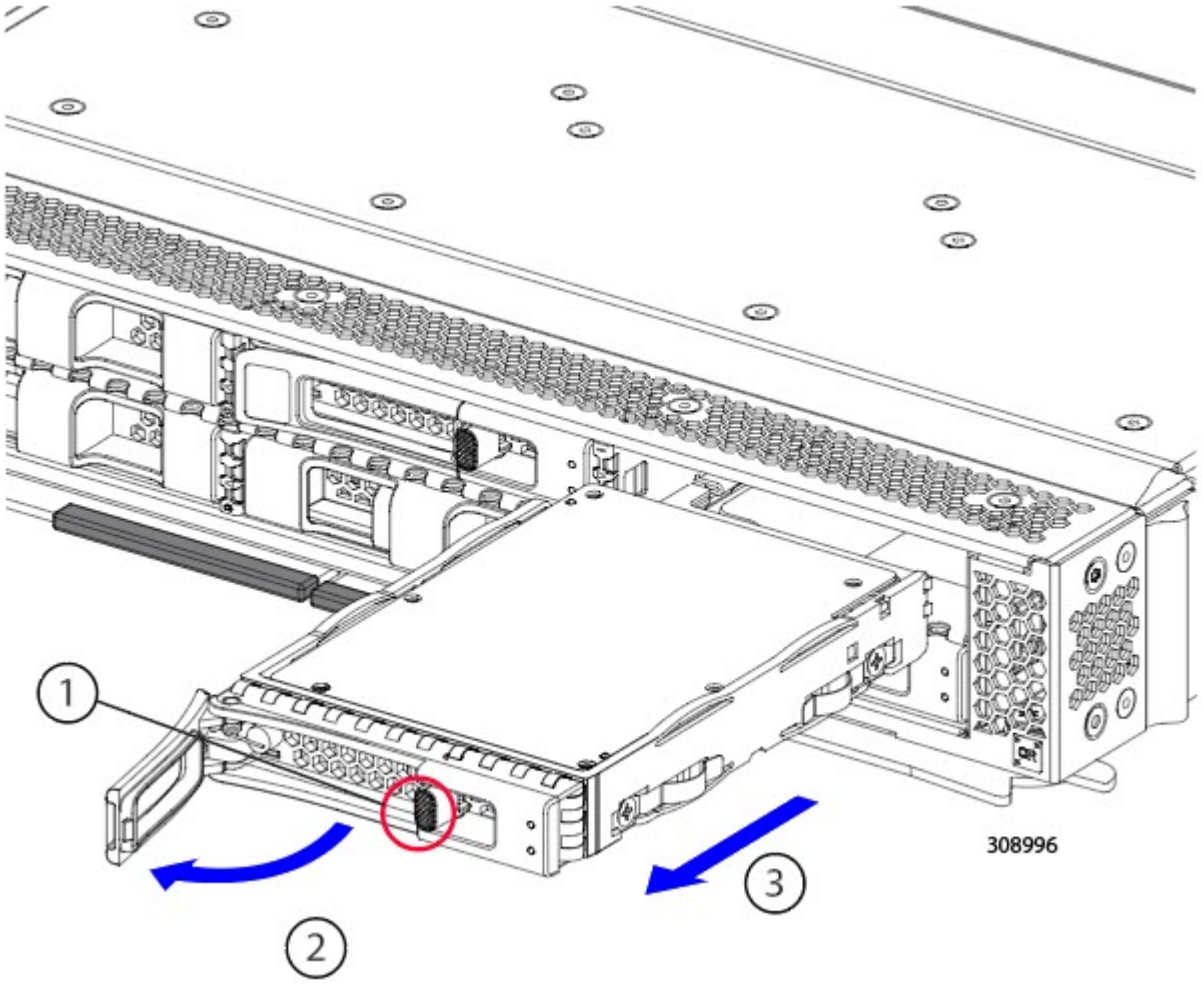
Use this task to remove a SAS/SATA or NVMe drive from the compute node.

Caution Do not operate the system with an empty drive bay. If you remove a drive, you must reinsert a drive or cover the empty drive bay with a drive blank.

Procedure

Step 1 Push the release button to open the ejector, and then pull the drive from its slot.

Caution
To prevent data loss, make sure that you know the state of the system before removing a drive.



Step 2 Place the drive on an antistatic mat or antistatic foam if you are not immediately reinstalling it in another compute node.

Step 3 Install a drive blanking panel to maintain proper airflow and keep dust out of the drive bay if it will remain empty.

What to do next

Cover the empty drive bay. Choose the appropriate option:

- Installing a Drive, on page 6
- Installing a Drive Blank, on page 10

Installing a Drive

Caution For hot installation of drives, after the original drive is removed, you must wait for 20 seconds before installing a drive. Failure to allow this 20-second wait period causes the Cisco UCS management software to display incorrect drive inventory information. If incorrect drive information is displayed, remove the affected drive(s), wait for 20 seconds, then reinstall them.

To install a SAS/SATA or NVMe drive in the compute node, follow this procedure:

Before you begin

Note When the UCS X215c M8 Compute Node hosts the UCSX-X10C-PTE3 pass-through controller, the maximum number of E3.S drives that can be installed is eight instead of nine. This difference is because slot five cannot support an E3.S drive due to electrical reasons. Slot five on the UCSX-X10C-PTE3 front mezzanine module must be covered with a drive filler panel (UCSC-E3S1T-F).

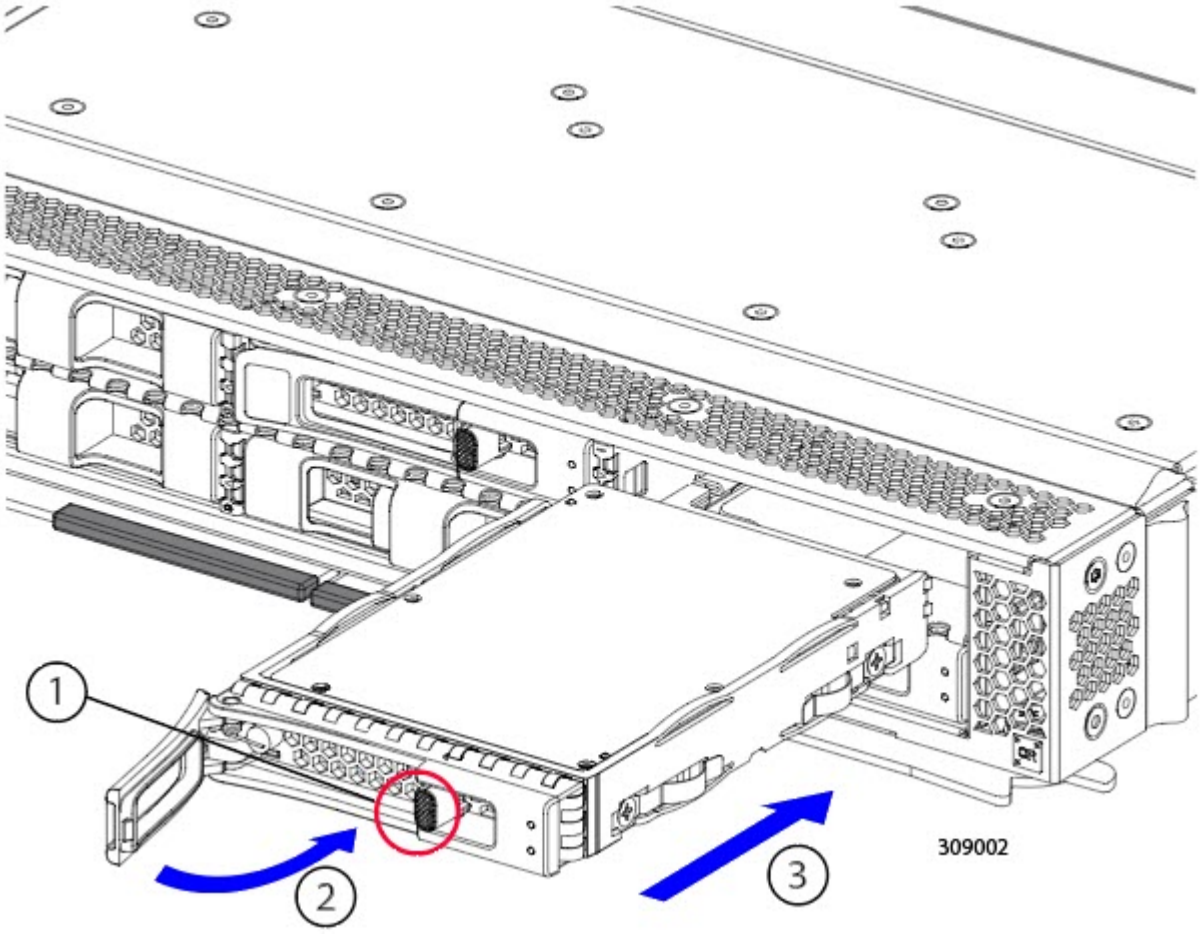
Procedure

Step 1 Place the drive ejector into the open position by pushing the release button.

Step 2 Gently slide the drive into the empty drive bay until it seats into place.

Step 3 Push the drive ejector into the closed position.

You should feel the ejector click into place when it is in the closed position.



Basic Troubleshooting: Reseating a SAS/SATA Drive

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

- Only drives that are managed by the UCS MegaRAID controller are affected.
- Both SFF and LFF form factor drives can be affected.
- Drives can be affected regardless of whether they are configured for hot plug or not.
- The UBAD error is not always terminal, so the drive is not always defective or in need of repair or replacement. However, it is also possible that the error is terminal, and the drive will need replacement.

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

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

To reseat the drive, see Reseating a SAS/SATA Drive, on page 8.

Reseating a SAS/SATA Drive

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

Use the following procedure to reseat the drive.

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

Before you begin

Before attempting this procedure, be aware of the following:

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

Procedure

Step 1 Attempt a hot reseat of the affected drive(s).

For a front-loading drive, see Removing a Drive, on page 5.

Note
While the drive is removed, it is a best practice to perform a visual inspection. Check the drive bay to ensure that no dust or debris is present. Also, check the connector on the back of the drive and the connector on the inside of the compute node for any obstructions or damage.

Also, when reseating the drive, allow 20 seconds between removal and reinsertion.

Step 2 During boot up, watch the drive's LEDs to verify correct operation.

See Interpreting LEDs.

Step 3 If the error persists, cold reseat the drive, which requires a compute node power down. Choose the appropriate option:

a) Use your server management software to gracefully power down the compute node.

See the appropriate Cisco UCS management software documentation.

b) If compute node power down through software is not available, you can power down the compute node by pressing the power button.

See Compute Node Front Panel.

c) Reseat the drive as documented in Step 1.

d) When the drive is correctly reseated, restart the compute node, and check the drive LEDs for correct operation as documented in Step 2.

Step 4 If hot and cold reseating the drive (if necessary) does not clear the `UBAD` error, choose the appropriate option:

a) Contact Cisco Systems for assistance with troubleshooting.

b) Begin an RMA of the errored drive.

Removing a Drive Blank

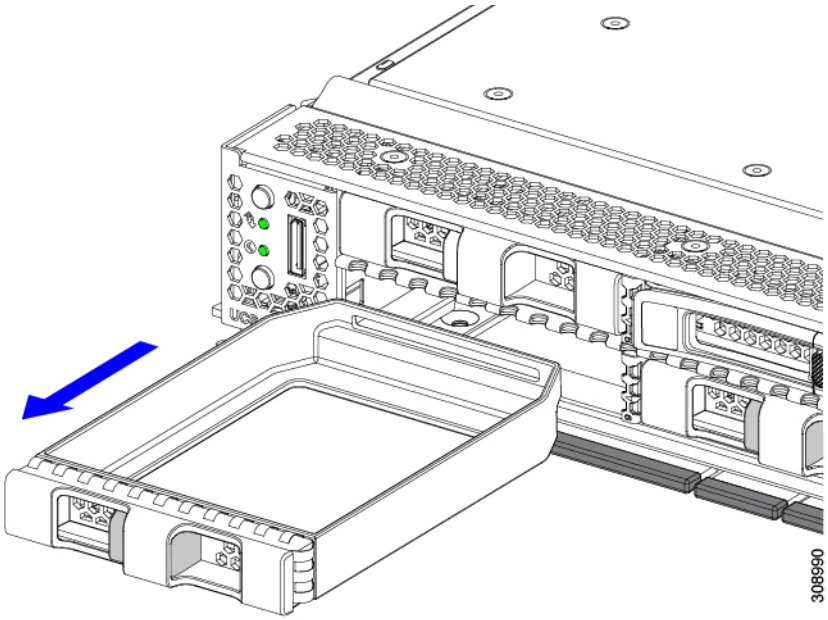
A maximum of six SAS/SATA or NVMe drives are contained in the front mezzanine storage module as part of the drive housing. The drives are front facing, so removing them does not require any disassembly.

Use this procedure to remove a drive blank from the compute node.

Procedure

Step 1 Grasp the drive blank handle.

Step 2 Slide the drive blank out of the slot.



What to do next

Cover the empty drive bay. Choose the appropriate option:

- Installing a Drive, on page 6
- Installing a Drive Blank, on page 10

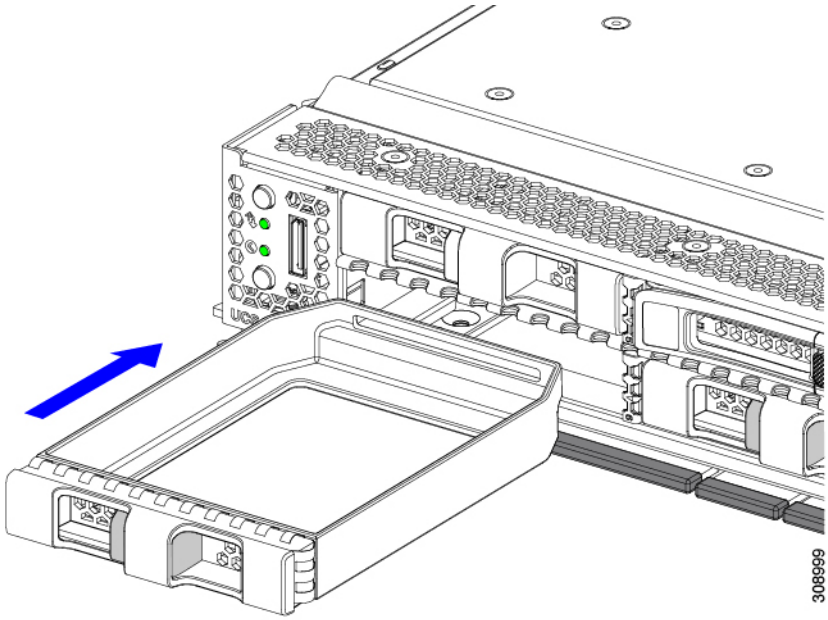
Installing a Drive Blank

Use this task to install a drive blank.

Procedure

Step 1 Align the drive blank so that the sheet metal is facing down.

Step 2 Holding the blank level, slide it into the empty drive bay.



Replacing the Front Mezzanine Module

The front mezzanine module is a steel cage that contains the compute node's storage devices or a mix of GPUs and drives. The front mezzanine storage module can contain any of the following storage configurations:

- NVMe U.3 drives
- SAS/SATA drives
- Cisco L4-MEZZ GPUs plus up to two U.3 NVMe drives
- EDSFF 1T E3.S PCIe drives
- Up to six (6) SAS/SATA/NVMe SSD drives in a Tri-Mode M1 RAID controller.
- A front mezzanine blank (UCSX-M8A-FMEZZBLK) for systems without local disk requirements.
- Compute Pass Through Controller (UCSX-X10C-PT4F): supports up to six hot pluggable 15mm NVMe drives directly connected to CPU 1.
- MRAID Storage Controller Module (UCSX-X10C-RAIDF):
  - Supports a mixed drive configuration of up to six SAS/SATA/NVMe drives. With a mix of SAS/SATA, any NVMe drives are supported in slots one through four only.
  - Provides HW RAID support for SAS/SATA drives in multiple RAID groups and levels.
  - Supports NVMe U.3 drives in slots 1 through 6 and can be configured into multiple RAID groups and levels similar to SAS/SATA drives.
  - Supports a mix of SAS/SATA and NVMe U.3 drives behind the MRAID controller. However, these NVMe drives and SAS/SATA drives cannot be combined in the same RAID group.

NVME U.3 drives can be combined to make RAID groups separately. Also, SAS/SATA drives can be formed into different RAID groups, and the different RAID groups can co-exist in the same MRAID storage setup.

- The front mezzanine module also contains the SuperCap module. For information about replacing the SuperCap module, see Replacing the SuperCap Module, on page 24.

Note The SuperCap module is only needed when the MRAID Storage Controller module (UCSX-X10C-RAIDF) or 24g Tri-Mode RAID Controller is installed.

- A compute and storage option (UCSX-X10C-GPUFM) consisting of a GPU adapter supporting zero, one, or two Cisco L4-MEZZ GPUs (UCSX-GPU-L4-MEZZ) plus zero, one, or two U.3 NVMe SSDs.
- A Front Mezzanine Pass Through Controller for E3.S drives (UCSX-X10C-PTE3) that supports up to eight E3.S PCIe drives.
- A Tri-mode M1 front mezzanine module (UCSX-RAID-M1L6). Each Tri-mode M1 front mezzanine module consists of up to six (6) SAS/SATA/NVMe SSD drives. Each drive slot supports either SAS, SATA or NVMe U.3 SSDs (RAID Controller)

The front mezzanine module can be removed and installed as a whole unit to give easier access to the storage drives that it holds. Or, you can leave the front mezzanine module installed because SAS/SATA and the NVMe drives are accessible directly through the front of the front mezzanine panel and are hot pluggable.

To replace the front mezzanine module, use the following topics:

- Removing the Front Mezzanine Module, on page 13
- Installing the Front Mezzanine Module, on page 15

Front Mezzanine Module Guidelines

Be aware of the following guidelines for the front mezzanine slot:

- For MRAID Storage Controller Module (UCSX-X10C-RAIDF), M.2 Mini Storage, and NVMe storage, only UEFI boot mode is supported.
- The compute node has a configuration option (UCSX-X10C-GPUFM) that supports up to two Cisco L4-MEZZ GPUs (UCSX-GPU-L4-MEZZ) and up to two NVMe U.3 drives in the front mezzanine slot. For information about the GPU-based front mezzanine option, see the Cisco UCS X10c Front Mezzanine GPU Module Installation and Service Guide.
- The compute node has a configuration option (UCSX-X10C-PTE3) that supports up to eight E3.S 1T PCIe5 drives. For more information about E3.S drives based front mezzanine option, see the Cisco UCS X10c Pass Through Controller for E3.S Installation and Service Guide .
- The compute node has a front mezzanine module option (UCSX-RAID-M1L6) front mezzanine that supports up to six (6) SAS/SATA/NVMe SSD drives. Each drive slot supports either SAS, SATA or NVMe U.3 SSDs through the RAID controller. For more information about the RAID controller-based front mezzanine option, see the Cisco UCS X10c Front Mezzanine GPU Module Installation and Service Guide.

Removing the Front Mezzanine Module

Use the following procedure to remove the front mezzanine module. This procedure applies to the following modules:

- Front mezzanine blank (UCSX-M8A-FMEZZBLK)
- Compute Pass Through Controller (UCSX-X10C-PT4F)
- MRAID Storage Controller Module (UCSX-X10C-RAIDF)
- Compute and storage option (UCSX-X10C-GPUFM)
- Compute Pass Through Controller for E3.S (UCSX-X10C-PTE3)
- Cisco 24G Tri-Mode M1 RAID Controller (UCSX-RAID-M1L6)

Before you begin

To remove the front mezzanine module, you need a T8 screwdriver and a #2 Phillips screwdriver.

Procedure

Step 1 If the compute node's cover is not already removed, remove it now. Remove the compute node cover.

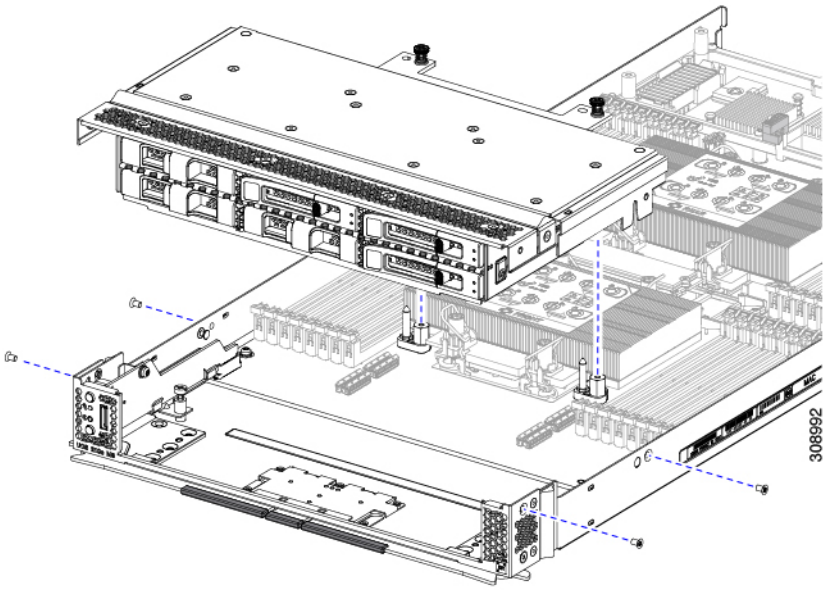
See Removing a Compute Node Cover, on page 1.

Step 2 Remove the securing screws:

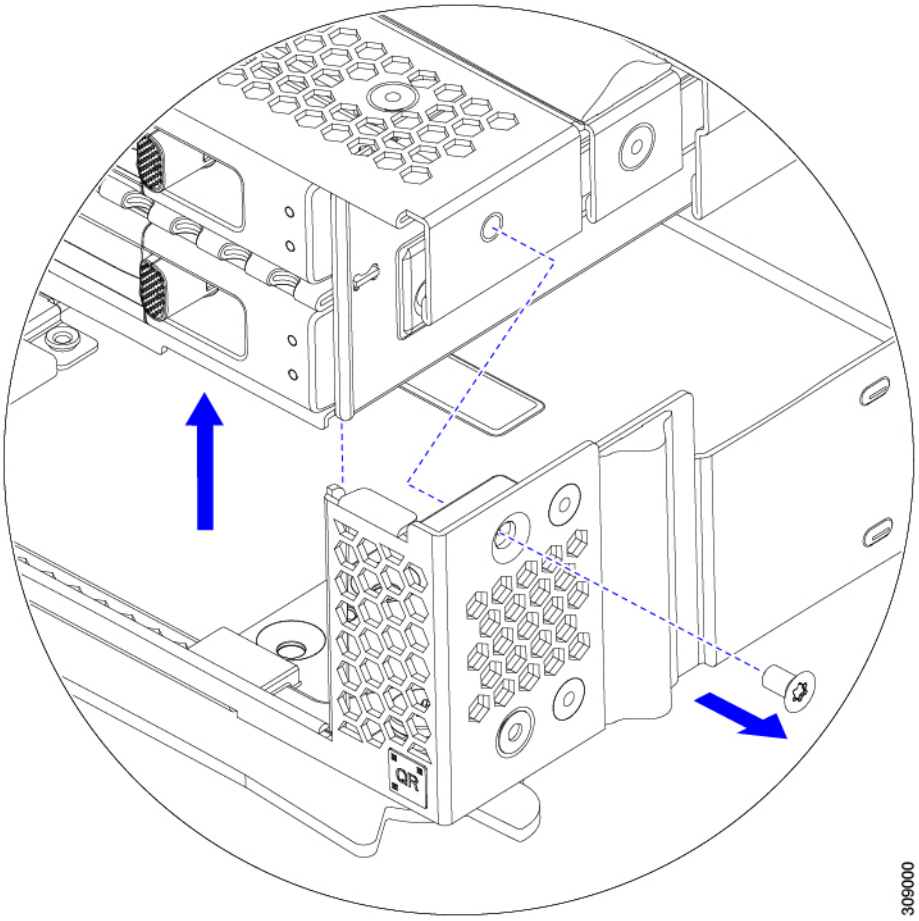
a) Using a #2 Phillips screwdriver, loosen the two captive screws on the top of the front mezzanine module.

Note
This step may be skipped if removing the front mezzanine blank (UCSX-M8A-FMEZZBLK).

b) Using a T8 screwdriver, remove the two screws on each side of the compute node that secure the front mezzanine module to the sheet metal.



Step 3 Making sure that all the screws are removed, lift the front mezzanine module to remove it from the compute node.



What to do next

To install the front mezzanine module, see Installing the Front Mezzanine Module, on page 15

Installing the Front Mezzanine Module

Use the following procedure to install the front mezzanine module. This procedure applies to the following modules:

- Front mezzanine blank (UCSX-M8A-FMEZZBLK)
- Compute Pass Through Controller (UCSX-X10C-PT4F)
- MRAID Storage Controller Module (UCSX-X10C-RAIDF)
- Compute and storage option (UCSX-X10C-GPUFM)
- Compute Pass Through Controller for E3.S (UCSX-X10C-PTE3)
- Cisco 24G Tri-Mode M1 RAID Controller (UCSX-RAID-M1L6)

Before you begin

To install the front mezzanine module, you need a T8 screwdriver and a #2 Phillips screwdriver.

Procedure

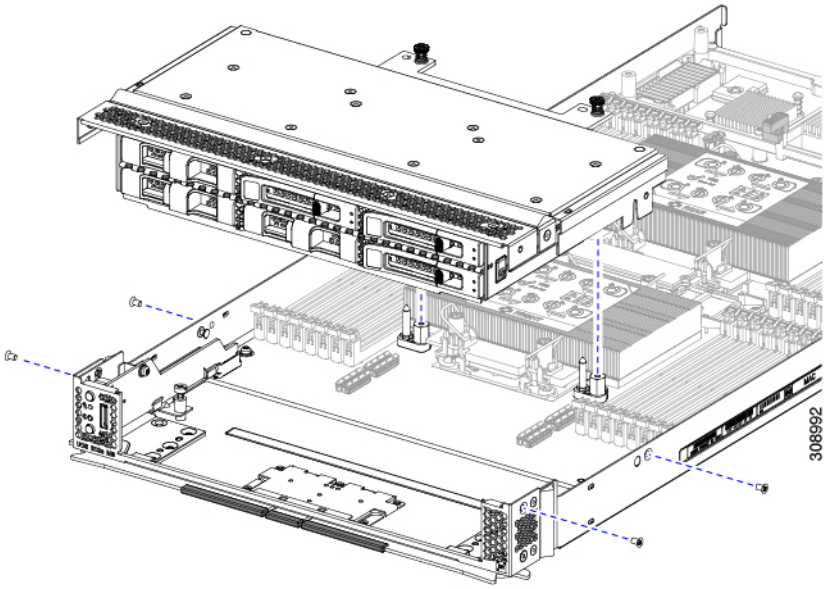
Step 1 Align the front mezzanine module with its slot on the compute node.

Step 2 Lower the front mezzanine module onto the compute node, making sure that the screws and screwholes line up.

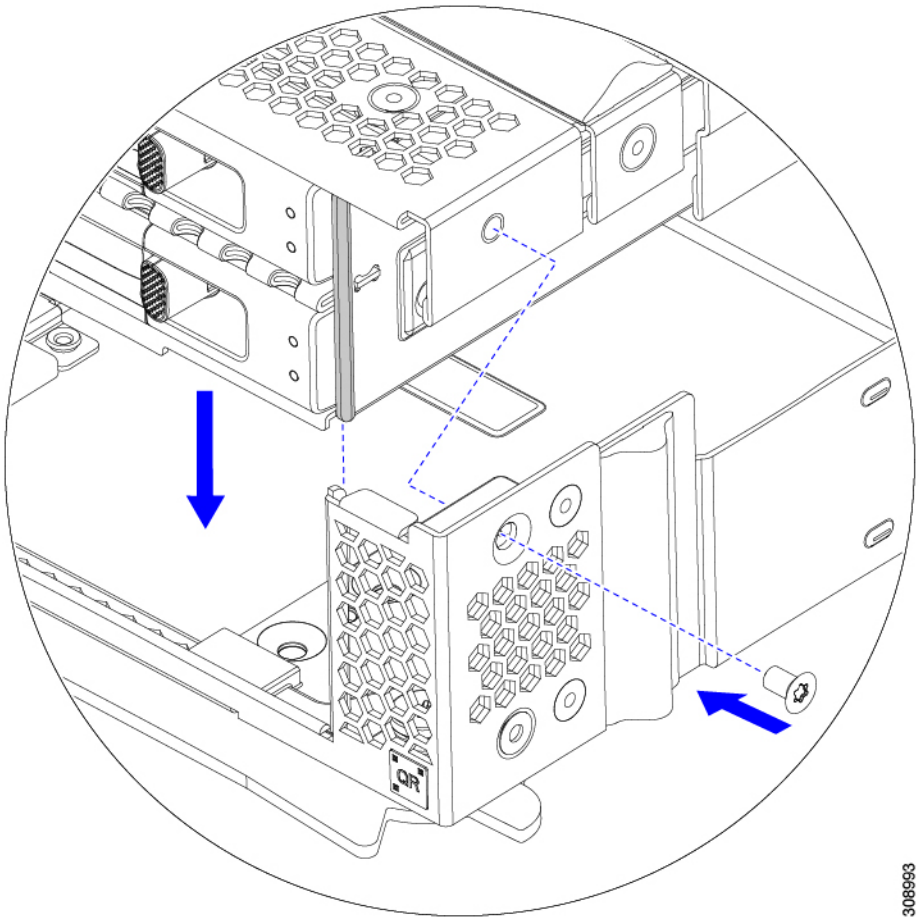
Step 3 Secure the front mezzanine module to the compute node.

a) Using a #2 Phillips screwdriver, tighten the captive screws on the top of the front mezzanine module.

Note
This step may be skipped if installing the front mezzanine blank (UCSX-M8A-FMEZZBLK).



b) Using a T8 screwdriver, insert and tighten the four screws, two on each side of the sever node.



What to do next

If you removed the drives from the front mezzanine module, reinstall them now. See Installing a Drive, on page 6.

Servicing the Mini Storage Module

The compute node has a mini-storage module option that plugs into a motherboard socket to provide additional internal storage. The module sits vertically behind the left side front panel. See Internal Components, on page 3.

Two configurations of mini storage module are supported, one with an integrated RAID controller card, and one without.

Replacing a Boot-Optimized M.2 RAID Controller Module or NVMe Pass-Through Module

The Cisco Boot-Optimized M.2 RAID Controller for M.2 SATA drives or the NVMe Pass-Through Controller for M.2 NVMe drives connects to the mini-storage module socket on the motherboard. Each of the following components contains two module slots for M.2 drives:

- The Cisco UCSX Front panel with M.2 RAID controller for SATA drives (UCSX-M2-HWRD-FPS). This component has an integrated 6-Gbps SATA RAID controller that can control the SATA M.2 drives in a RAID 1 array.
- The Cisco UCSX Front panel with M.2 Pass Through controller for NVME drives (UCSX-M2-PT-FPN). The M.2 NVMe drives are not configurable in a RAID group.

Cisco Boot-Optimized M.2 RAID Controller Considerations

Review the following considerations:

- This controller supports RAID 1 (single volume) and JBOD mode.
- A SATA M.2 drive in slot 1 is located on the right side, or front, of the module when installed. This drive faces the interior of the compute node. This drive is the first SATA device.
- A SATA M.2 drive in slot 2 is located on the left side, or back, of the module when installed. This drive faces the compute node's sheet metal wall. This drive is the second SATA device.
    - The name of the controller in the software is MSTOR-RAID.
    - A drive in slot 1 is mapped as drive 253; a drive in slot 2 is mapped as drive 254.
- When using RAID, we recommend that both SATA M.2 drives are the same capacity. If different capacities are used, the smaller capacity of the two drives is used to create a volume and the rest of the drive space is unusable.

  JBOD mode supports mixed capacity SATA M.2 drives.
- Hot-plug replacement is *not* supported. The compute node must be powered off.
- Monitoring of the controller and installed SATA M.2 drives can be done using Cisco UCS management software. They can also be monitored using other utilities such as UEFI HII, and Redfish.
- The SATA M.2 drives can boot in UEFI mode only. Legacy boot mode is not supported.
- If you replace a single SATA M.2 drive that was part of a RAID volume, rebuild of the volume is auto-initiated after the user accepts the prompt to import the configuration. If you replace both drives of a volume, you must create a RAID volume and manually reinstall any OS.
- We recommend that you erase drive contents before creating volumes on used drives from another compute node. The configuration utility in the compute node BIOS includes a SATA secure-erase function.

Removing the M.2 RAID Controller Module or NVMe Pass-Through Module

This topic describes how to remove a Cisco Boot-Optimized M.2 RAID Controller or a Cisco NVMe Pass-Through Controller:

- The Cisco UCSX Front panel with M.2 RAID controller for SATA drives (UCSX-M2-HWRD-FPS).
- The Cisco UCSX Front panel with M.2 Pass-Through module for NVME drives (UCSX-M2-PT-FPN).

Both types of controller board have two slots, one for each M.2 drive:

- one M.2 slot (Slot 1) for either a SATA drive (in UCSX-M2-HWRD-FPS) or an NVMe drive (in UCSX-M2-PT-FPN). The drive in this slot faces the interior of the compute node.
- one M.2 slot (Slot 2) for either a SATA drive (in UCSX-M2-HWRD-FPS) or an NVMe drive (in UCSX-M2-PT-FPN). The drive in this slot faces the chassis sheetmetal wall.
- Drive slot numbering differs depending on which Cisco management tool you are using and which component is being managed.

Component	Cisco Management Tool	
	Intersight (IMM)	UCS Manager (UCSM)
RAID Controller	Slot 1 contains Drive 253 Slot 2 contains Drive 254	Slot 1 contains Drive 253 Slot 2 contains Drive 254
NVMe Pass-Through Controller	Slot 1 contains Drive 253 Slot 2 contains Drive 254	Slot 1 contains Drive 32 Slot 2 contains Drive 33

Note For the NVMe Pass-Through Controller in Intersight Managed Mode (IMM), drives are shown as MSTOR-NVME-1 / MSTOR-NVME-2, which map to physical slot numbers.

Each controller can be populated with up to two M.2 drives of the correct type, either SATA for the RAID controller or NVMe for the Pass-Through controller. Single M.2 SATA or NVMe drives are supported. You cannot mix M.2 drive types in the same controller.

To remove the controller or the M.2 drives, the front mezzanine module must be removed first.

Procedure

Step 1 Remove the controller from the compute node:

a) Decommission, power off, and remove the compute node from the chassis.
b) Remove the top cover from the compute node as described in Removing and Installing the Compute Node Cover, on page 1.

Step 2 If you have not already done so, remove the front mezzanine module.

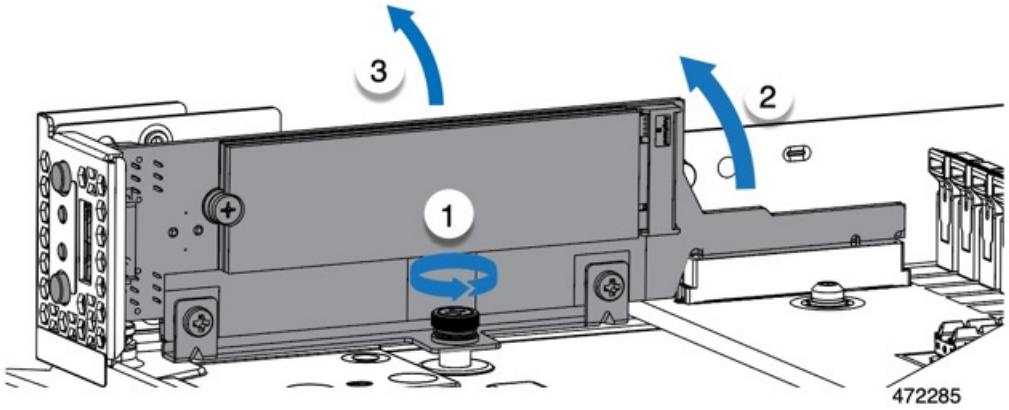
See Removing the Front Mezzanine Module, on page 13.

Step 3 Remove the controller.

a) Locate the controller in the front corner of the compute node along the compute node's sidewall.
b) Using a #2 Phillips screwdriver, loosen the captive screw that secures the module to the motherboard.

c) At the end opposite the front panel, grasp the module and pull up in an arc to disconnect the controller from its motherboard socket.
d) Holding the controller at an angle, slide it away from the front panel and lift it up to disengage the LEDs and buttons from their cutouts in the front panel.

Caution
If you feel resistance while lifting the controller, make sure that the LEDs and buttons are not still seated in the front panel.



Step 4 If you are transferring M.2 drives from the old controller to a replacement controller, do that before installing the replacement controller:

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

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

Installing the M.2 RAID Controller Module or NVMe Pass-Through Controller Module

Use this task to install the RAID controller or NVME Pass-through controller module.

Before you begin

This topic describes how to remove a Cisco Boot-Optimized M.2 RAID Controller or a Cisco NVMe Pass-Through Controller:

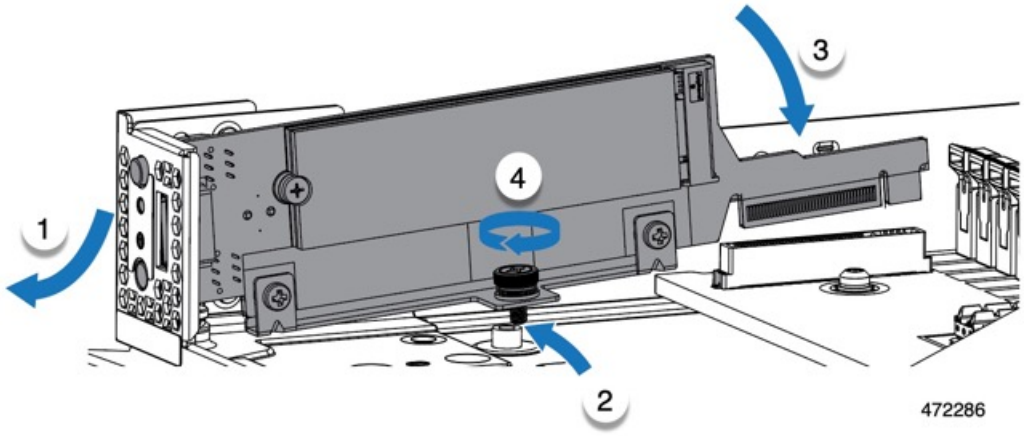
- The Cisco UCSX Front panel with M.2 RAID controller for SATA drives (UCSX-M2-HWRD-FPS).
- The Cisco UCSX Front panel with M.2 Pass-Through module for NVME drives (UCSX-M2-PT-FPN).

Each type of controller mounts vertically on the motherboard, and the M.2 drive sockets are positioned vertically on the controller.

Procedure

Step 1 Install the controller to its socket on the motherboard:

a) Position the controller over the socket, making sure the golden fingers on the connector are facing down.
b) Lower the controller into the chassis at an angle and insert the LEDs and buttons into their cutouts on the front panel.
c) Holding the controller level, align the captive screw with its screwhole and the golden fingers with their socket on the motherboard.
d) Carefully push down on the controller to seat the golden fingers into the socket.
e) Use a #2 Phillips screwdriver to tighten the controller onto the threaded standoff.



Step 2 Reinstall the front mezzanine module.

Step 3 Return the compute node to service:

a) Replace the top cover on the compute node.
b) Reinstall the compute node and allow it to power up and be automatically reacknowledged, reassociated, and recommissioned.

Replacing an M.2 SATA or M.2 NVMe SSD

M.2 SATA and NVMe SSD cards can be installed in vertical drive bays. One drive bay, or slot, is on each side of the M.2 module carrier.

There are some specific rules for populating mini-storage M.2 SSD cards:

- Each carrier supports a maximum of two M.2 cards. Do not mix SATA and NVMe SSD cards in the same mini-storage module. Replacement cards are available from Cisco as pairs.

- When installed in the compute node, the M.2 SSDs are mounted vertically.
  - M.2 slot 1 is located on the right side, or front, of the module when installed. This drive faces inward towards the interior the compute node.
  - M.2 slot 2 is located on the left side, or back, of the module when installed. This drive faces outward towards the compute node sheetmetal wall.
  - Drive slot numbering depends on the M.2 SSD type and which Cisco Management tool you are using.
    - **M.2 SATA SSD**: Slot 1 contains Drive 253 in both Intersight (IMM) and UCS Manager (UCSM).
    - **M.2 SATA SSD**: Slot 2 contains Drive 254 in both IMM and UCSM.
    - **M.2 NVMe SSD**: Slot 1 contains Drive 253 in IMM, but Slot 1 contains Drive 32 in UCSM.
    - **M.2 NVMe SSD**: Slot 2 contains Drive 254 in IMM, but Slot 2 contains Drive 33 in UCSM.
  - If your compute node contains only one M.2 SATA or NVMe SSD, it can be installed in either slot.
- Dual SATA M.2 SSDs can be configured in a RAID 1 array through the BIOS Setup Utility's embedded SATA RAID interface and configured through IMM.

Note The M.2 SSDs are managed by the MSTOR-RAID controller.

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

Removing an M.2 SATA or M.2 NVMe SSD

Each M.2 card plugs into a slot on the carrier, which mounts vertically to the motherboard.

- One slot is on the front of the carrier, which faces inwards towards the rest of the compute node.
- One slot is on the back of the carrier, which faces towards the compute node sheetmetal wall.

Each M.2 SSD is secured to the carrier by the slot at one end, and a small retaining screw at the other end. The carrier is installed on the same component that has the compute node LEDs and buttons on the node's front panel.

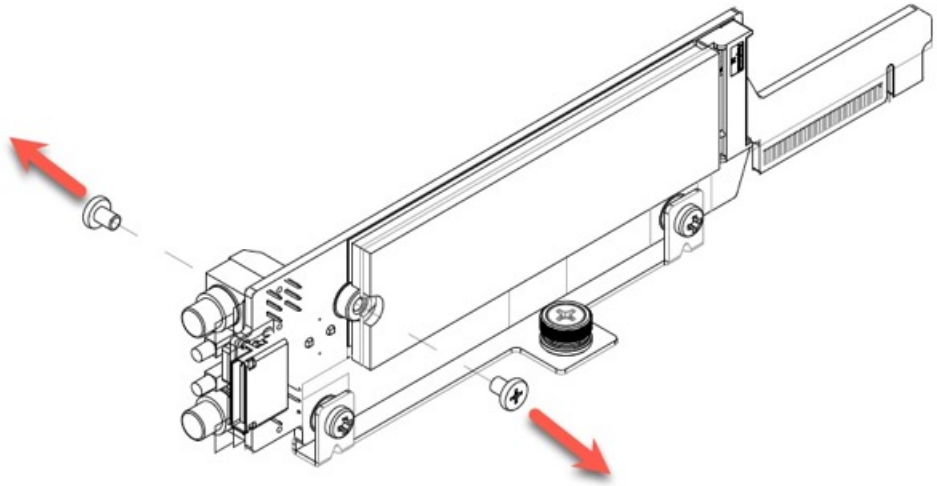
Use the following procedure for any type of mini-storage module carrier.

Procedure

Step 1 Remove the controller.

See Removing the M.2 RAID Controller Module or NVMe Pass-Through Module, on page 18.

Step 2 Using a #1 Phillips screwdriver, remove the screws that secure the M.2 SSD to the carrier.



Step 3 Grasping the M.2 card by its edges, gently lift the end that held the screws at an angle, then slide the card out of its connector.

What to do next

Installing an M.2 SATA or M.2 NVMe SSD, on page 23

Installing an M.2 SATA or M.2 NVMe SSD

Each M.2 SATA or NVMe SSD plugs into a slot on the carrier and is held in place by a retaining screw for each SSD.

Use the following procedure to install the M.2 SATA or NVMe SSD onto the carrier

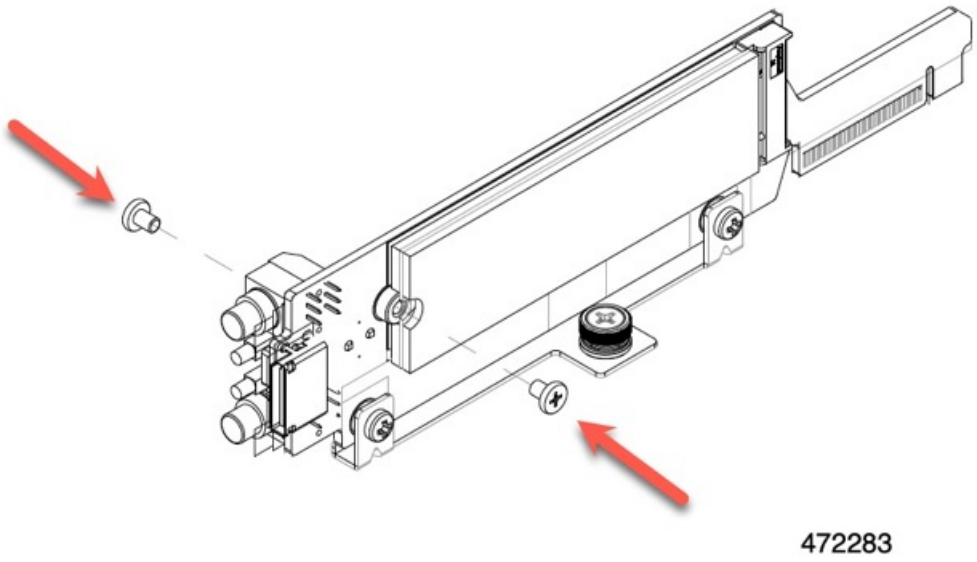
Procedure

Step 1 Install the M.2 SATA or NVMe SSD.

a) Orient the SSD correctly.

Note
When correctly oriented, the end of the SSD with two alignment holes lines up with the two alignment pins on the carrier.

b) Angle the end opposite the screw into the connector
c) Press down on the end of the SSD that holds the screws until the SSD snaps into place.
d) Reinsert and tighten the retaining screw to secure the M.2 module to the carrier.



Step 2 When you are ready, reinstall the controller onto the motherboard.

Installing the M.2 RAID Controller Module or NVMe Pass-Through Controller Module, on page 20.

Step 3 Reinstall the compute node cover

Step 4 Reapply power and return the compute node to service.

Replacing the SuperCap Module

The SuperCap module (UCSB-MRAID-SC) is a battery bank which connects to the front mezzanine storage module board and provides power to the RAID controller if facility power is interrupted. The front mezzanine with the SuperCap module installed is UCSX-X10C-RAIDF.

Note The SuperCap module is only needed when the MRAID Storage Controller module (UCSX-X10C-RAIDF) is installed.

Note To remove the SuperCap Module you must remove the front mezzanine module.

To replace the SuperCap module, use the following topics:

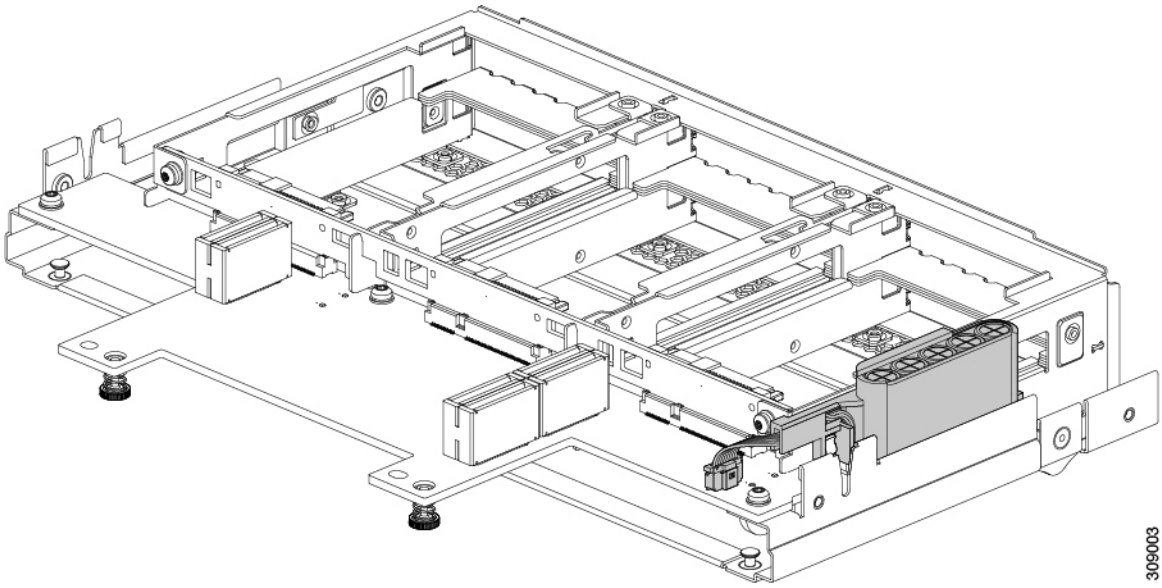
- Removing the SuperCap Module, on page 25
- Installing the SuperCap Module, on page 29

Removing the SuperCap Module

The SuperCap module is part of the Front Mezzanine Module, so the Front Mezzanine Module must be removed from the compute node to provide access to the SuperCap module.

The SuperCap module sits in a plastic tray on the underside of the front mezzanine module. The SuperCap module connects to the board through a ribbon cable with one connector to the module.

Figure 1: Location of the SuperCap Module on the UCS X215c M8 Compute Node



To replace the SuperCap module, follow these steps:

Procedure

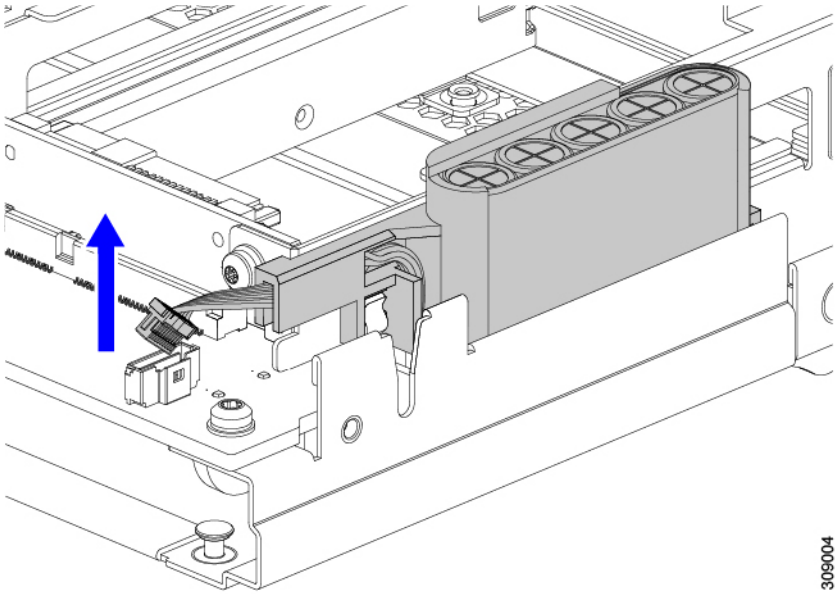
Step 1 If you have not already removed the Front Mezzanine module, do so now.

See Removing the Front Mezzanine Module, on page 13.

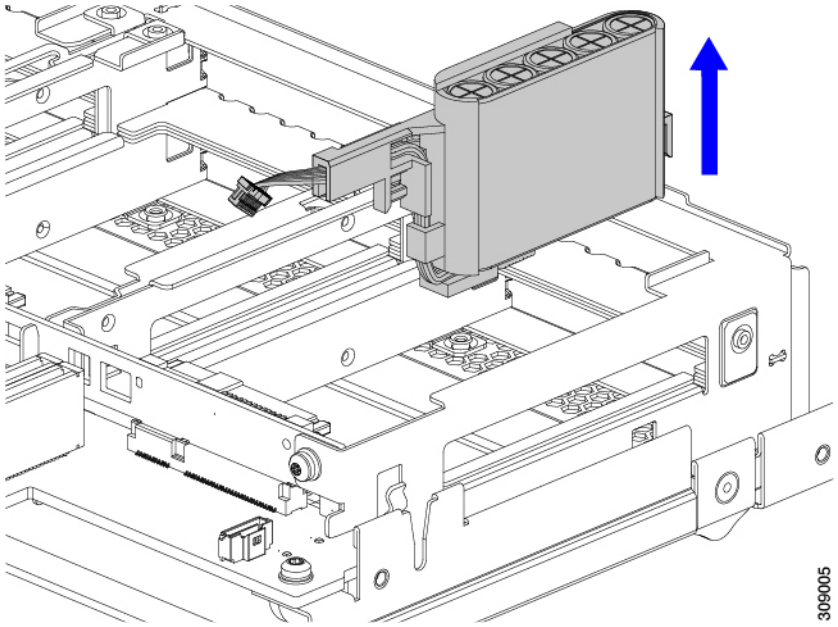
Step 2 Before removing the SuperCap module, note its orientation in the tray as shown in the previous image.

When correctly oriented, the SuperCap connection faces downward so that it easily plugs into the socket on the board. You will need to install the new SuperCap module with the same orientation.

Step 3 Grasp the cable connector at the board and gently pull to disconnect the connector.



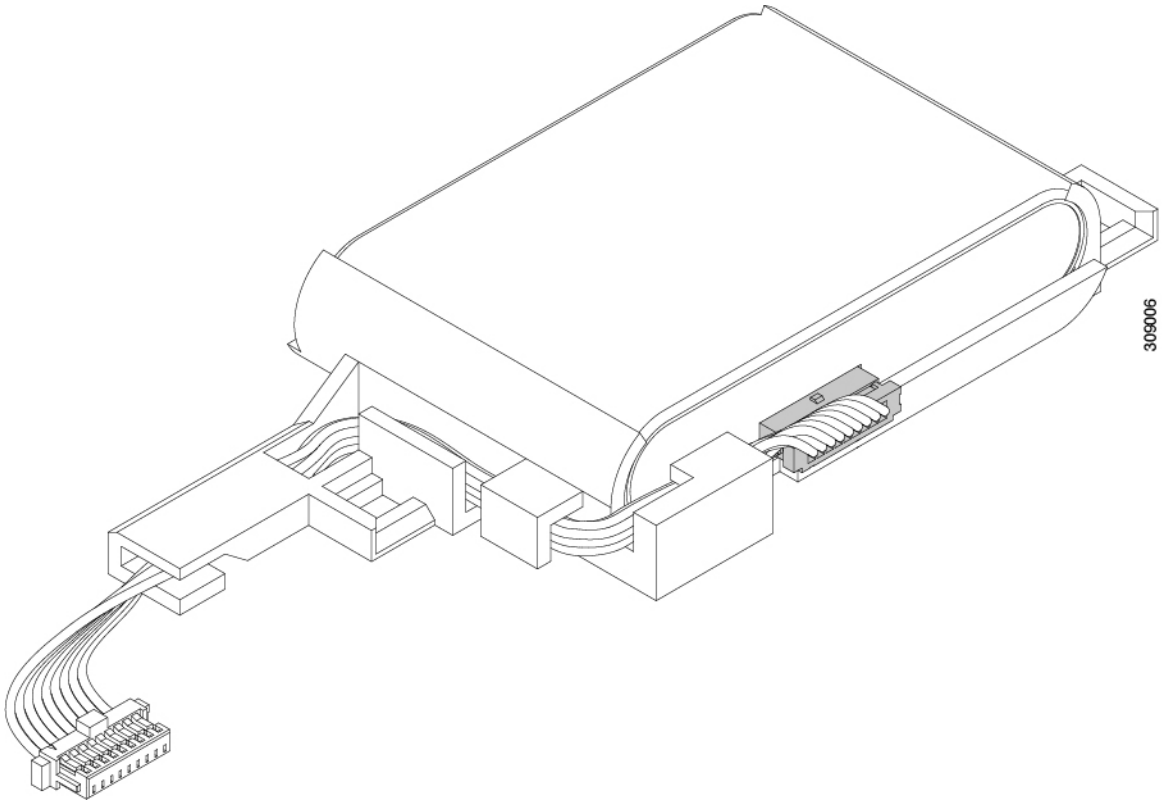
Step 4 Grasp the sides of the SuperCap module, but not the connector, and lift the SuperCap module out of the tray.



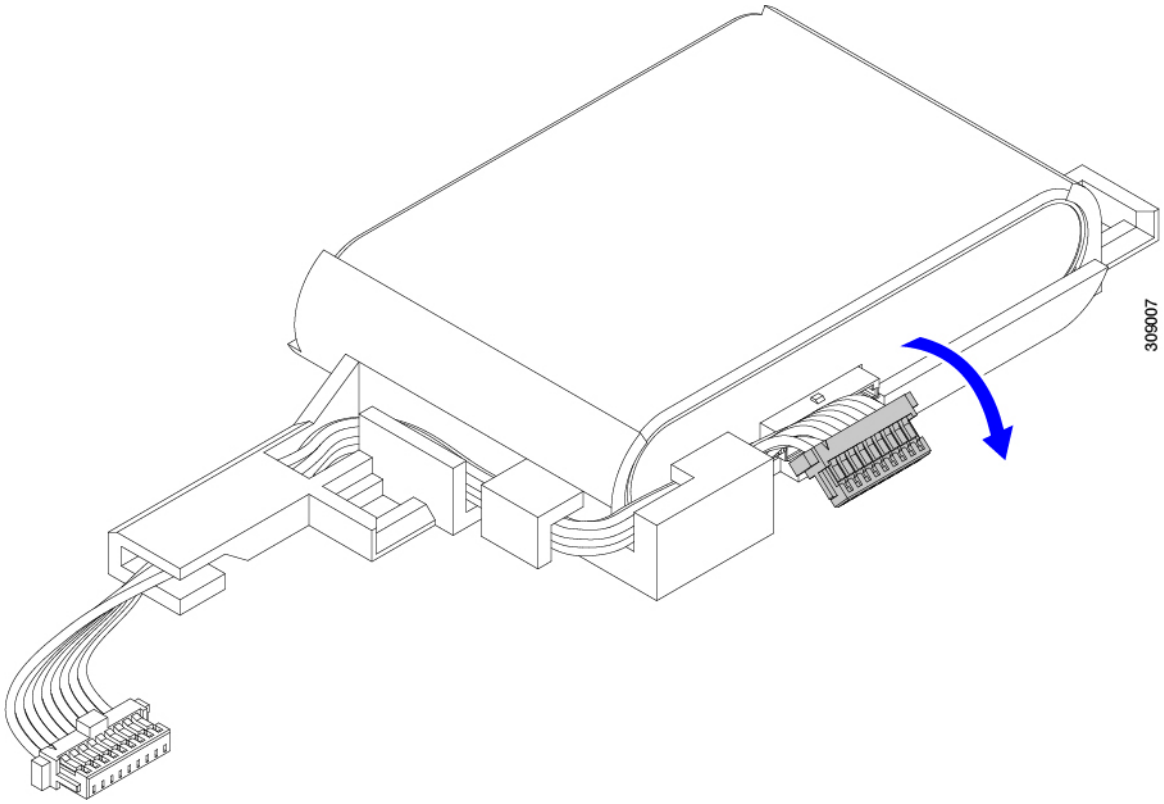
You might feel some resistance because the tray is curved to secure the module.

Step 5 Disconnect the ribbon cable from the SuperCap module:

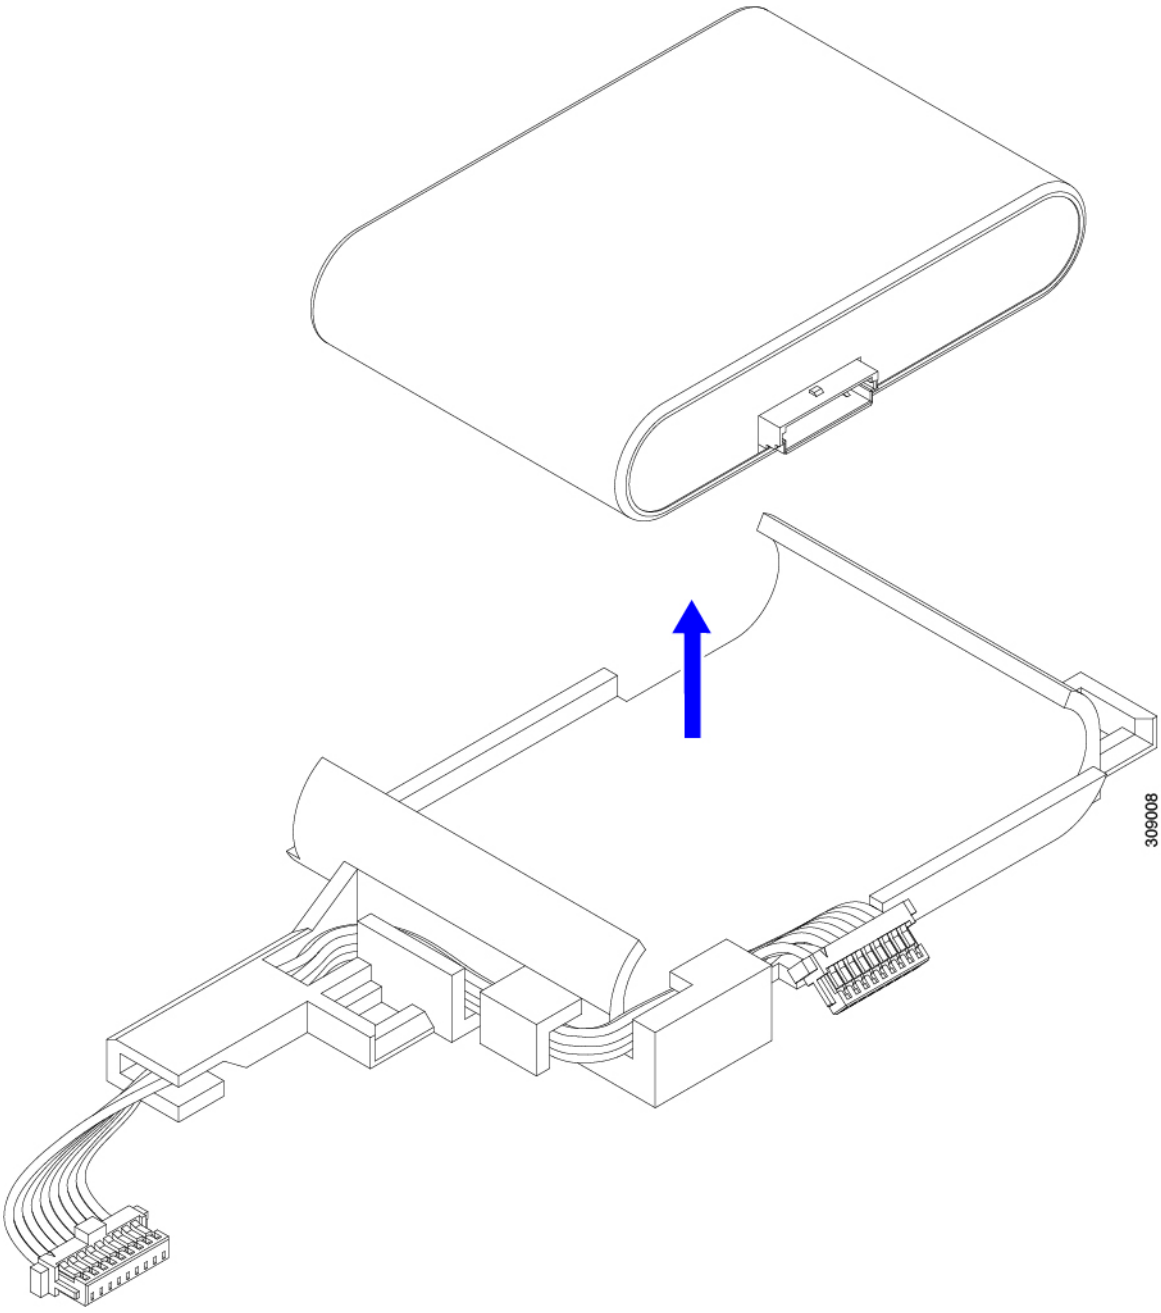
a) On the SuperCap module, locate the lever that secures the ribbon cable to the battery pack.


309006

b) Gently pivot the securing lever downward to release the ribbon cable connection from the SuperCap module.



Step 6 Remove the existing battery pack from its case, and insert a new one, making sure to align the new battery pack so that the connector aligns with the ribbon cable.



What to do next

Installing the SuperCap Module, on page 29

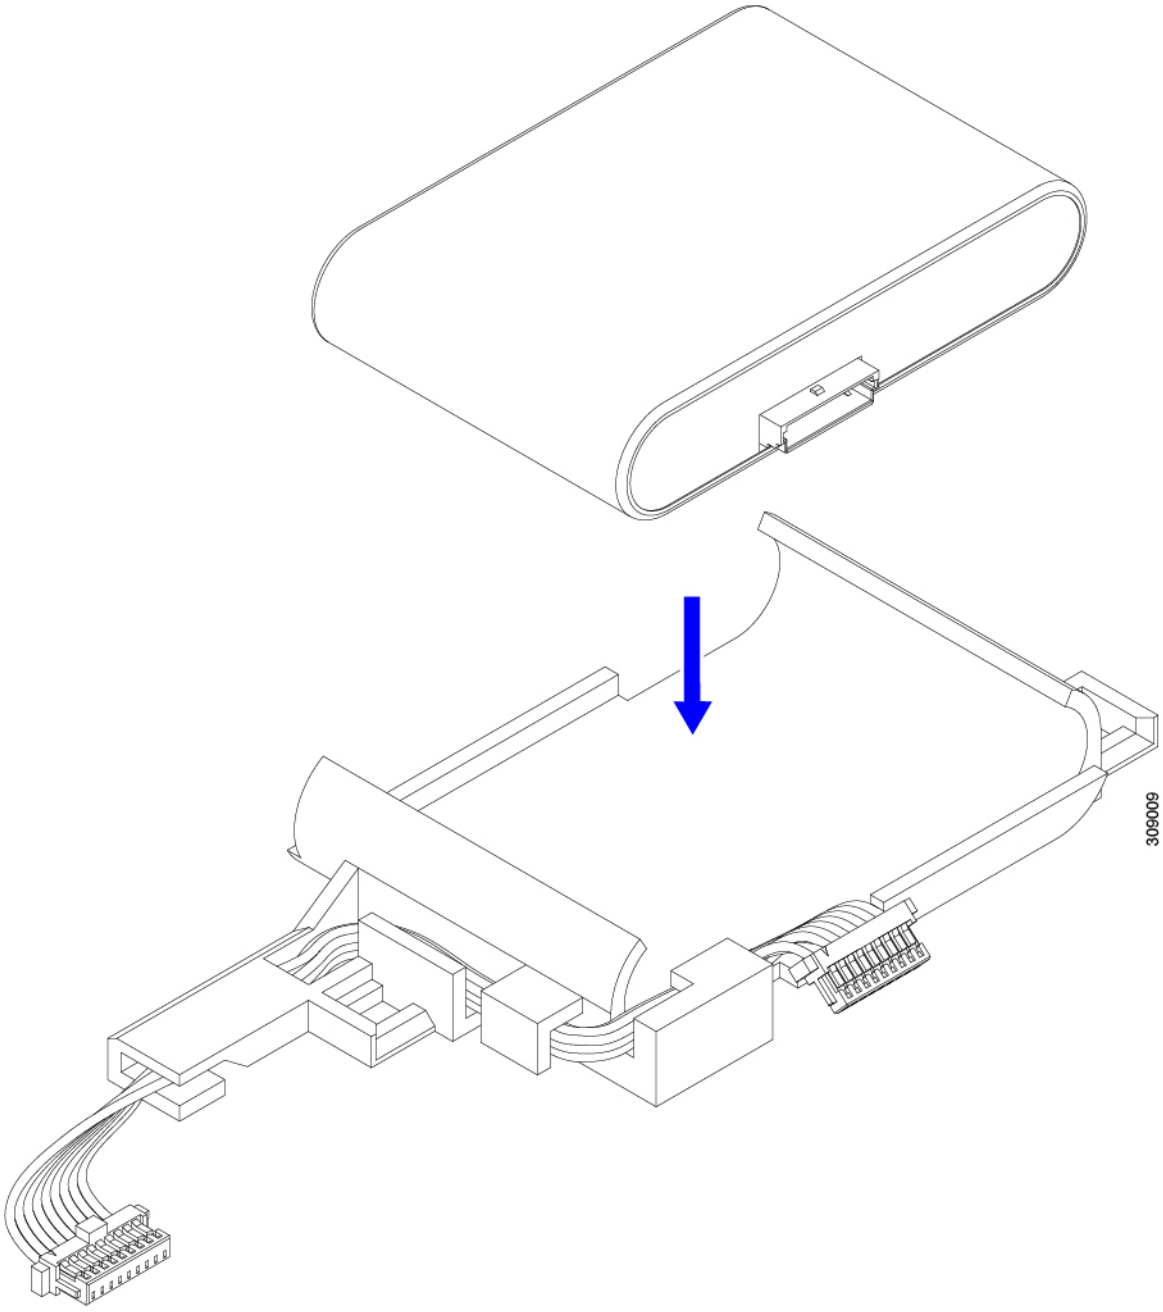
Installing the SuperCap Module

If you removed the SuperCap module, use this procedure to reinstall and reconnect it.

Procedure

Step 1 Insert the Super Cap module into its case.

a) Align the SuperCap module so that the connector will meet the connector.

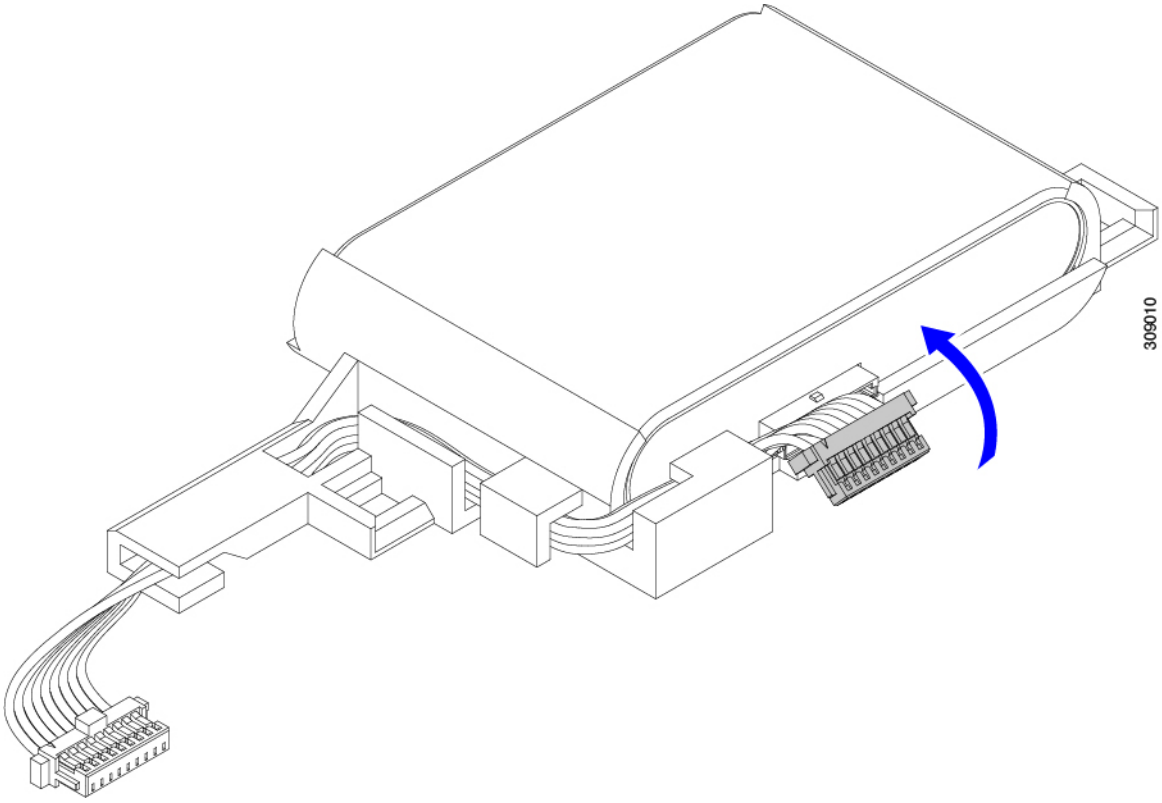


b) Before seating the SuperCap module, make sure that the ribbon cable is not in the way. You do not want to pinch the ribbon cable when you install the SuperCap.

c) When the ribbon cables are clear of the case, press the SuperCap module until it is seated in the case.

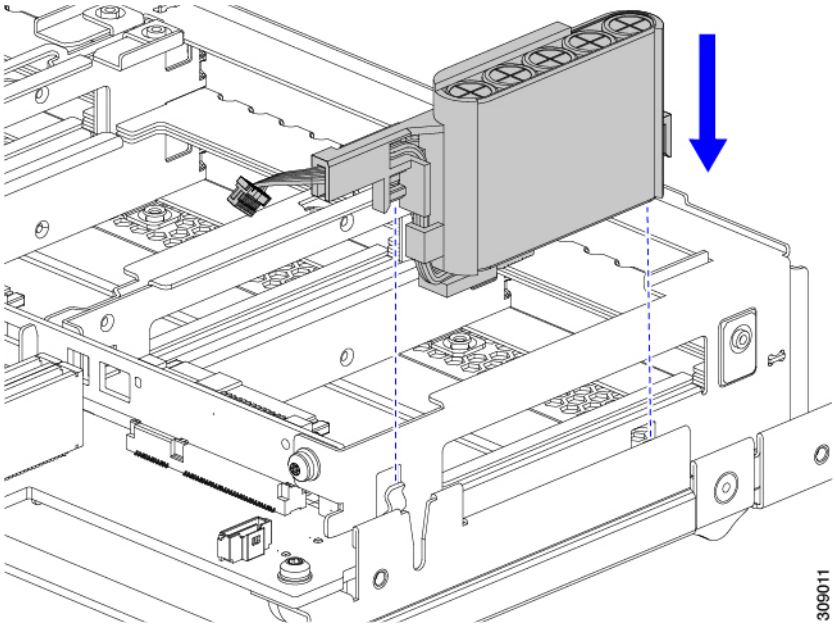
You might feel some resistance as the SuperCap snaps into place.

Step 2 When the SuperCap module is completely seated in its plastic case, pivot the securing lever to connect the ribbon cable to the SuperCap module.



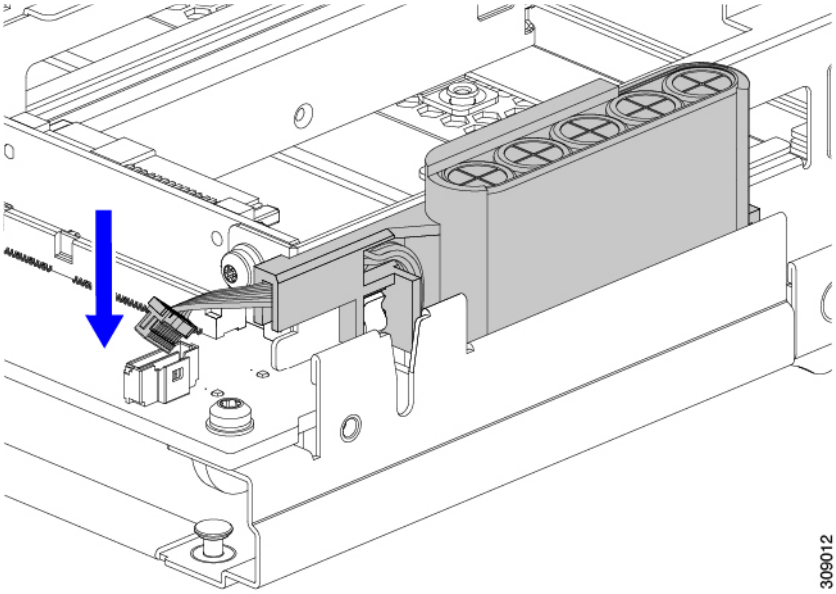
Step 3 Align the SuperCap module with its slot on the module and seat the module into the slot.

Caution
Make sure not to pinch the ribbon cable while inserting the SuperCap module into the slot.



When the SuperCap is securely seated in the slot, the module does not rock or twist.

Step 4 After the SuperCap module is seated, reconnect the ribbon cable to the board.



Replacing CPUs and Heatsinks

This topic describes the configuration rules and procedure for replacing CPUs and heatsinks.

CPU Configuration Rules

This compute node has two CPU sockets on the motherboard. Each CPU supports 12 DIMM channels (12 DIMM slots for each CPU).

- The compute node can operate with one CPU or two identical CPUs installed.
- The minimum configuration is that the compute node must have at least CPU 1 installed. Install CPU 1 first, and then CPU 2.
- Unpopulated CPU sockets must be covered with a dust cover. If you require a dust cover, contact Cisco.
- The following restrictions apply when using a single-CPU configuration:
  - Any unused CPU socket must have the socket dust cover from the factory in place.
  - The maximum number of DIMMs is 12.

Tools Required for CPU Replacement

You need the following tools and equipment for this procedure:

- T-20 Torx driver (for heatsink and CPU socket screws).
- Heatsink cleaning kit—Supplied with replacement CPU. Orderable separately as Cisco PID UCSX-HSCK=

  One cleaning kit can clean up to four CPUs.
- Thermal interface material (TIM)—Syringe supplied with replacement CPU. Use only if you are reusing your existing heatsink (new heatsinks have a pre-applied pad of TIM).

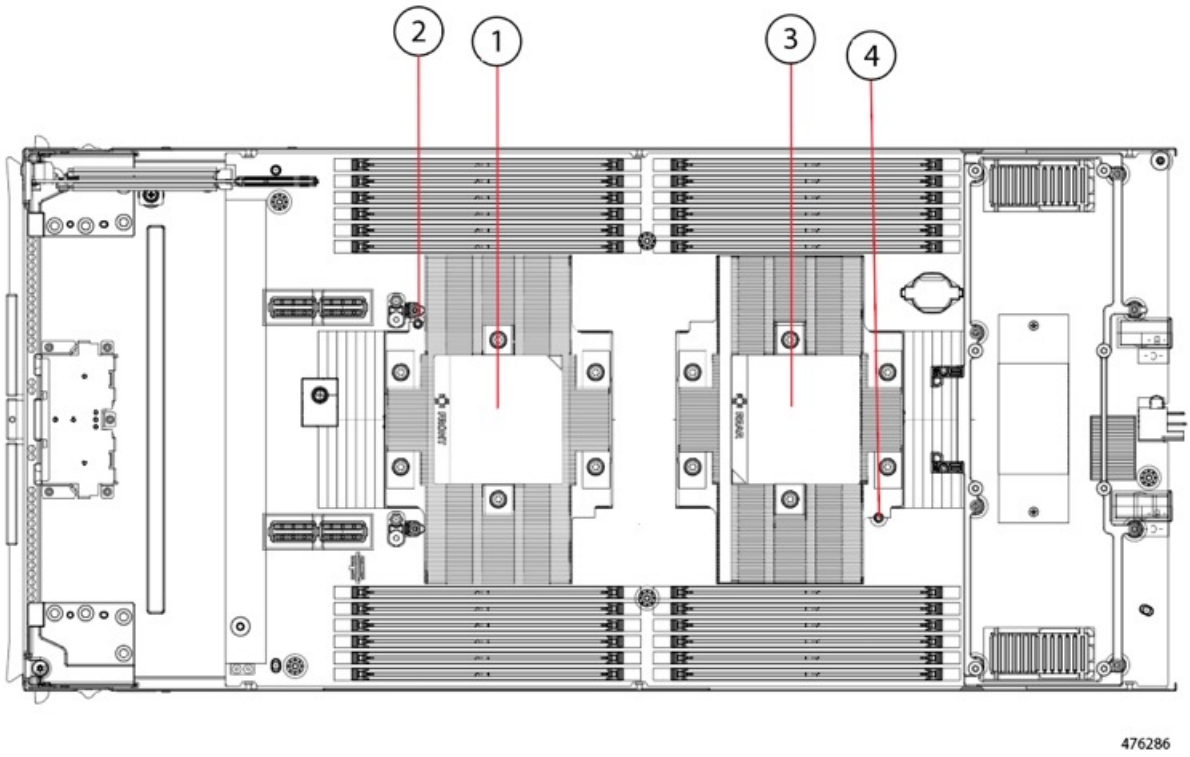
  One TIM kit covers one CPU.

Heatsink Alignment Features

For installation and field-replacement procedures, each heatsink must be properly aligned to the pin 1 location as shown in the following example.



Note The front heatsink is installed on the front CPU. The rear heatsink is installed on the rear CPU. The orientation of each CPU is different between CPU socket 1 and CPU socket 2, as indicated by the alignment pins on the motherboard.



1	Front heatsink	**2**	Front heatsink alignment pin
3	Rear heatsink	**4**	Rear heatsink alignment pin

Removing the CPU and Heatsink

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

Caution When handling the CPU, always use the handling tab. Do not hold the CPU by its edges, and do not touch the CPU top, bottom, or pins.

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

Important During the removal process, make sure thermal grease residuals or foreign material does not drop in the CPU socket.

Procedure

Step 1 Decommission the compute node by using Cisco UCS management software, such as Cisco Intersight.

Step 2 Disconnect any cables from ports on the compute node or installed cards.

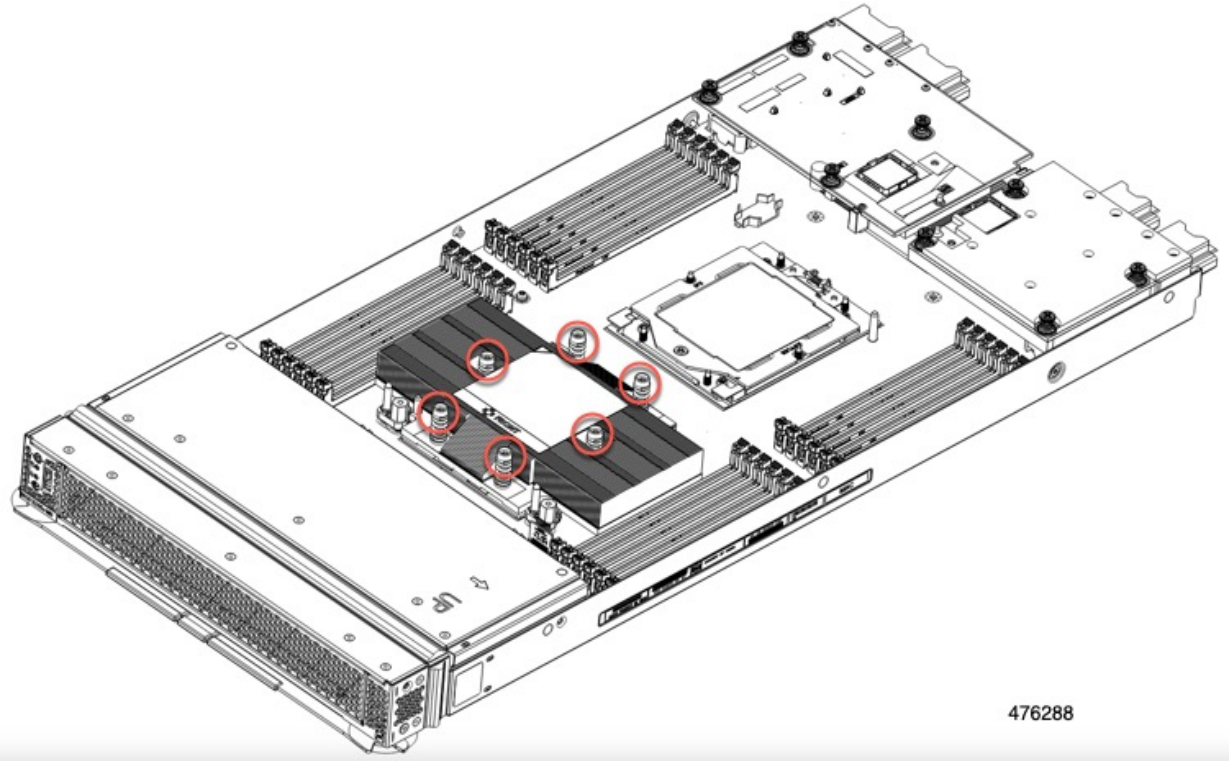
Step 3 Remove the heatsink from the CPU that you are replacing:

Caution
Before handling the heatsink, refer to the label for additional instructions.

a) Use a T-20 Torx driver and a Phillips screwdriver to loosen the six captive screws that secure the heatsink.

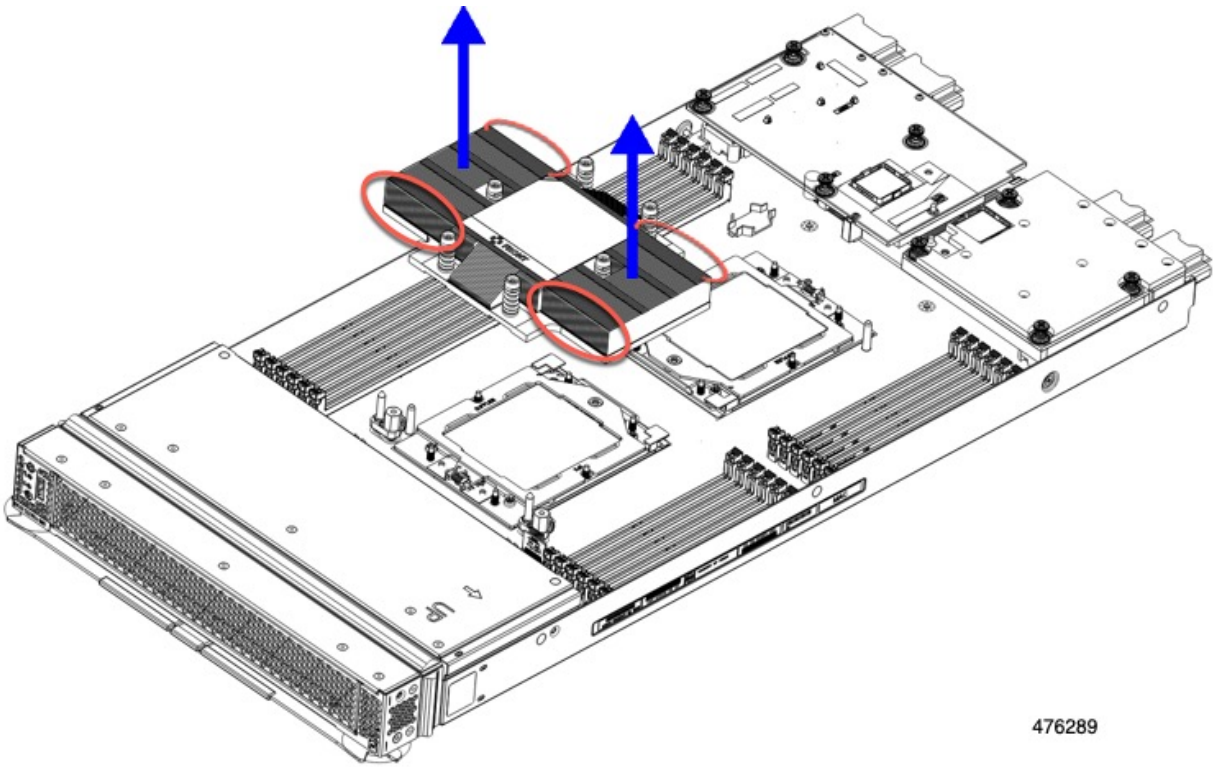
Note
Alternate loosening the heatsink screws evenly so that the heatsink remains level as it is raised. Make sure to loosen all screws in a star pattern, or loosen a screw then loosen the screw diagonally opposite of it.

Figure 2: Loosening the Heatsink Screws



b) Grasp the heatsink by the vertical edges of the fins, then lift straight up on the heatsink and set it down on an antistatic surface. Use caution to avoid damaging the heatsink-to-CPU surface.

Figure 3: Removing the Heatsink



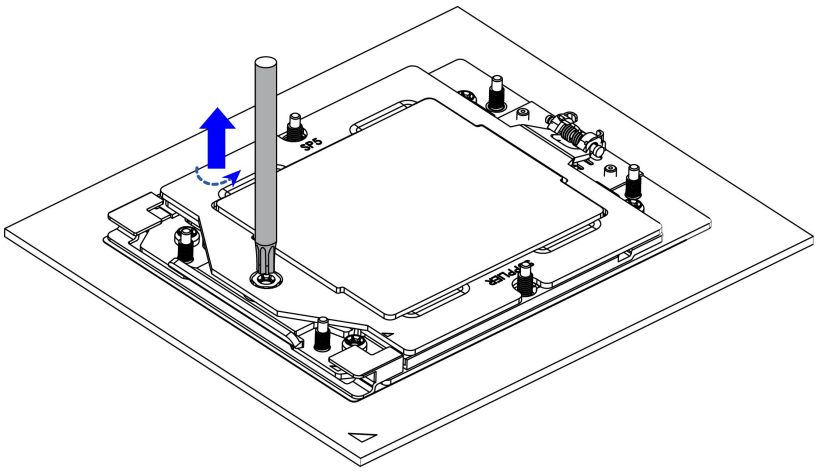
Step 4 Remove the CPU from the socket:

Caution
Before handling the CPU, refer to the heatsink label for additional instructions.

Note
Use the TIM cleaning kit in UCSX-HSCK= to clean the thermal grease on the top of the CPU and the retention frame before loosening the retention frame screw.

a) Use the T-20 Torx driver to loosen the captive socket-frame screw.

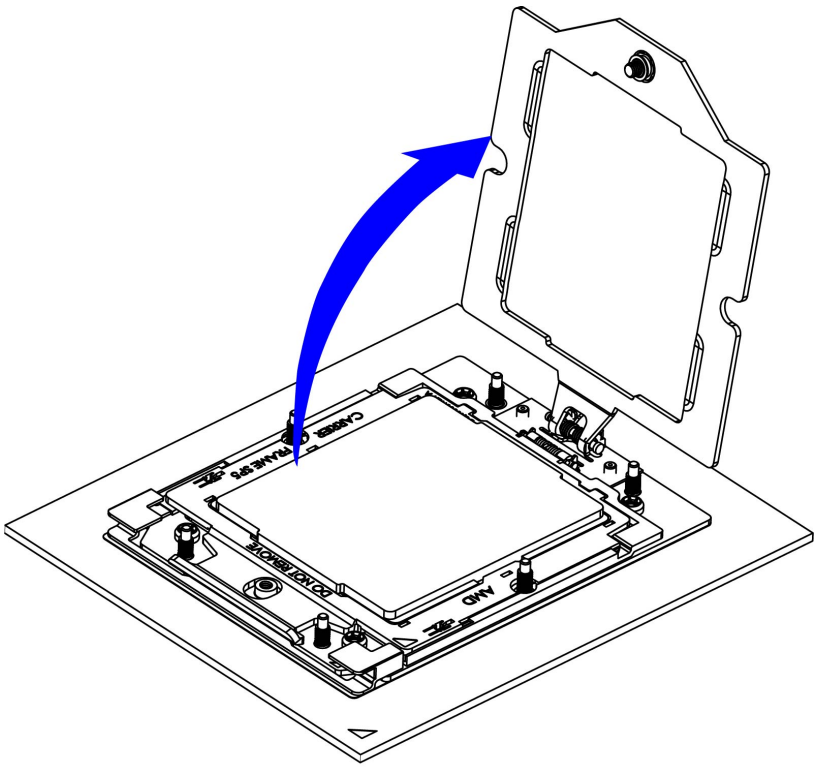
Figure 4: Loosening the Retention Frame Screw



481604

b) Pivot the hinged retention frame to the upright position.

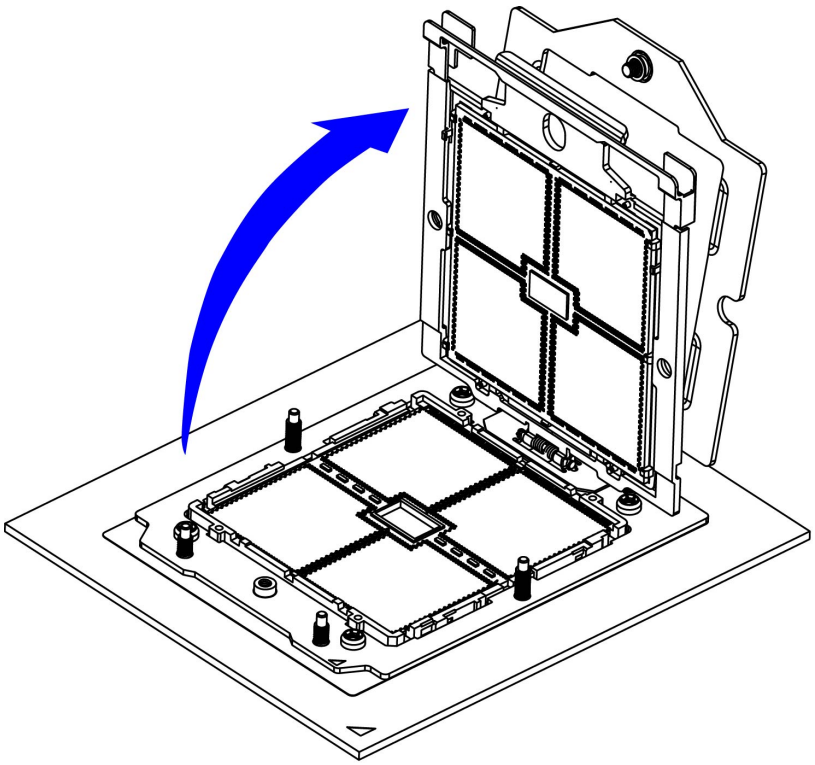
Figure 5: Opening the Retention Frame



481605

c) Pivot the rail frame to the upright position.

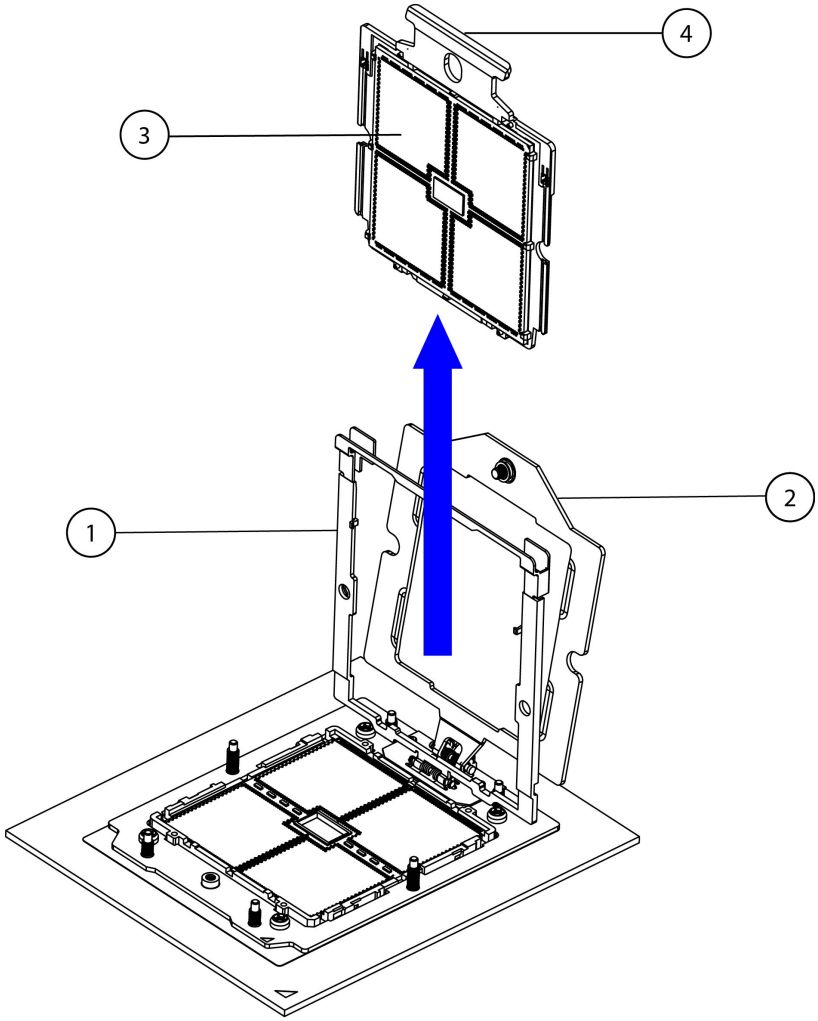
Figure 6: Opening the Rail Frame



481606

d) Grasp the CPU only by the handling tab that is on its carrier frame and pull straight up to remove the CPU from the rail frame.

Figure 7: Removing the CPU From the Socket



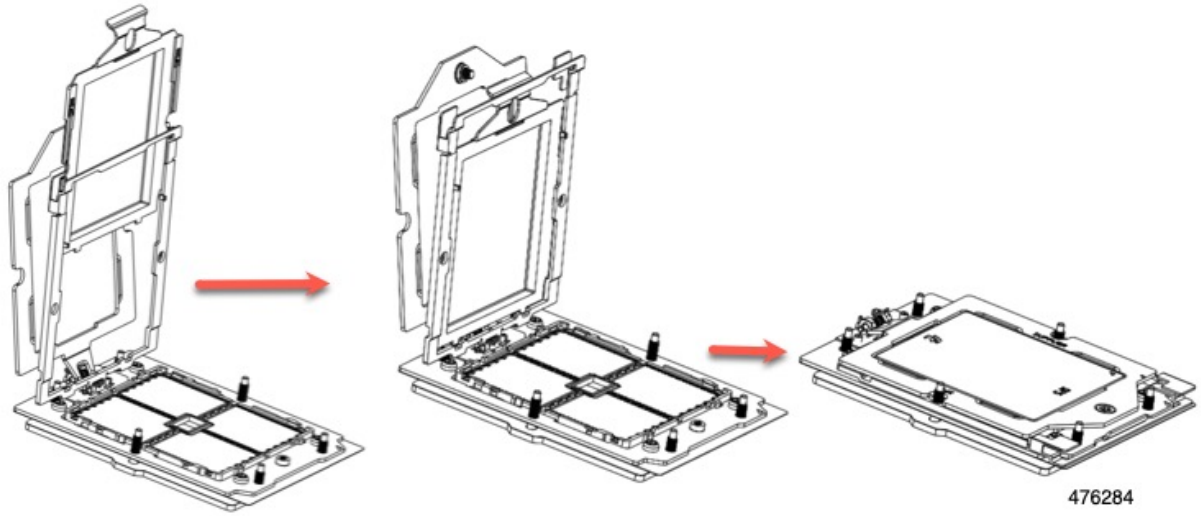
1	Rail frame in open position	**3**	CPU in carrier frame
2	Retention frame in open position	**4**	Handling tab on CPU carrier frame

What to do next

Choose the appropriate option:

- If you will be installing a CPU, go to Installing the CPU and Heatsink, on page 40.
- If you will not be installing a CPU, verify that a CPU dust cover is installed. This option is valid only for CPU socket 2 because CPU socket 1 must always be populated in a runtime deployment. If you require a dust cover, contact Cisco.

Figure 8: Installing a Dust Cover



Installing the CPU and Heatsink

Use the following procedure to install a CPU assembly consisting of the CPU and heatsink.

Procedure

Step 1 Take the new CPU assembly out of the carton.

Caution
There may be CPU SKUs that limit the maximum recommended compute node operating temperature. See CPU Configuration Rules, on page 33.

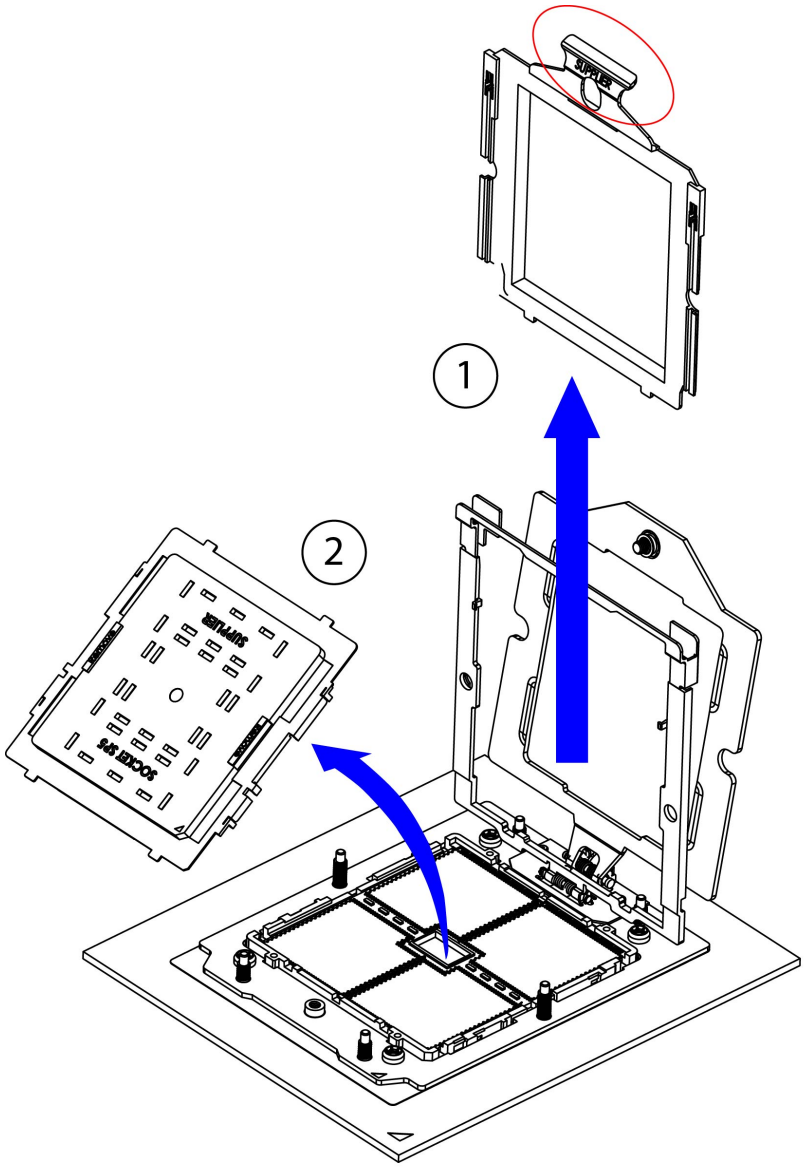
Caution
CPUs and their sockets are fragile and must be handled with extreme care to avoid damaging pins.

Step 2 Install the new CPU:

Caution
The CPU contacts and pins are extremely fragile. In this step, use extreme care to avoid touching or damaging the CPU contacts or the CPU socket pins.

Note
Ensure that you are following CPU Configuration Rules, on page 33.

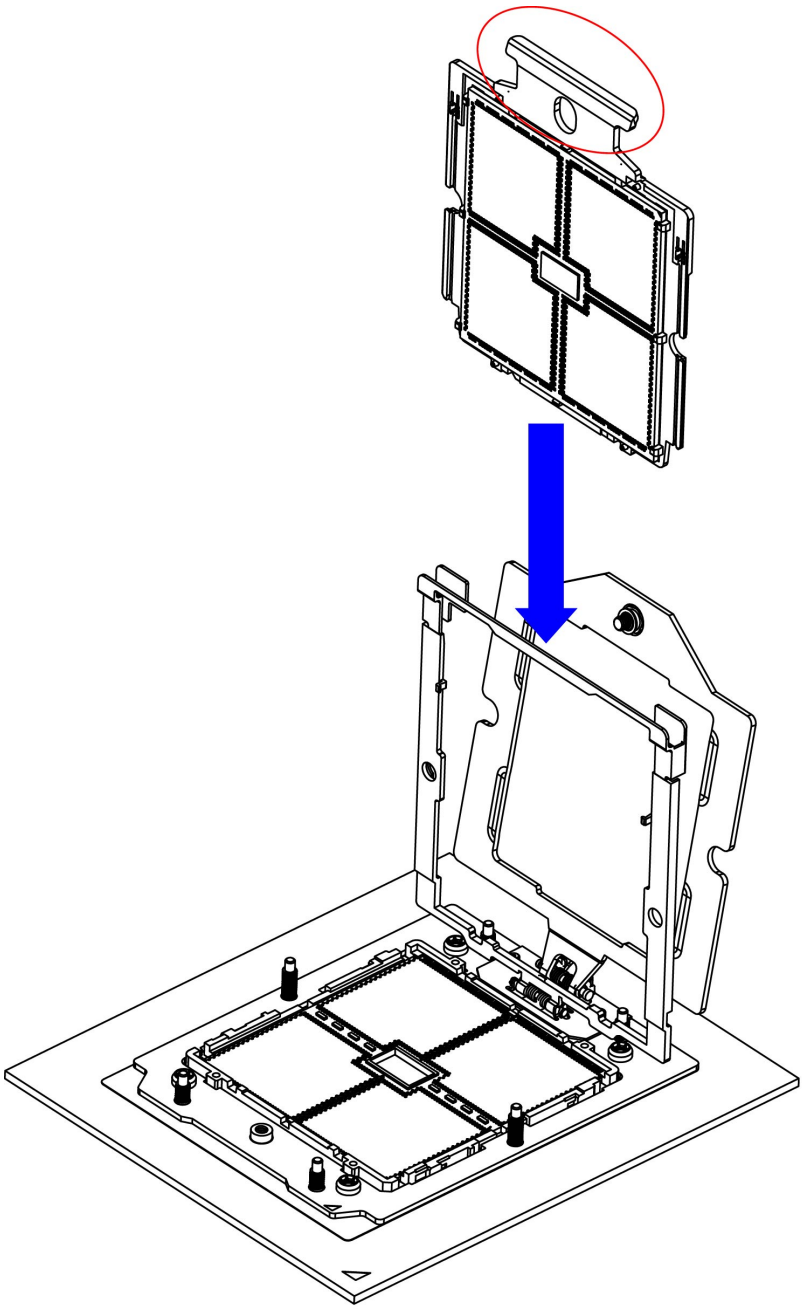
a) If the CPU socket still has a CPU and heatsink installed, remove them now. For more information, go to Removing the CPU and Heatsink, on page 34.
b) If the CPU socket has the dust cap and socket cap in place, open the retention frame, and remove the two caps now.



481613

c) Grasping the CPU only by the handling tab on its carrier frame, carefully slide it down into the open rail frame.

Figure 9: Inserting the CPU into Carrier Frame

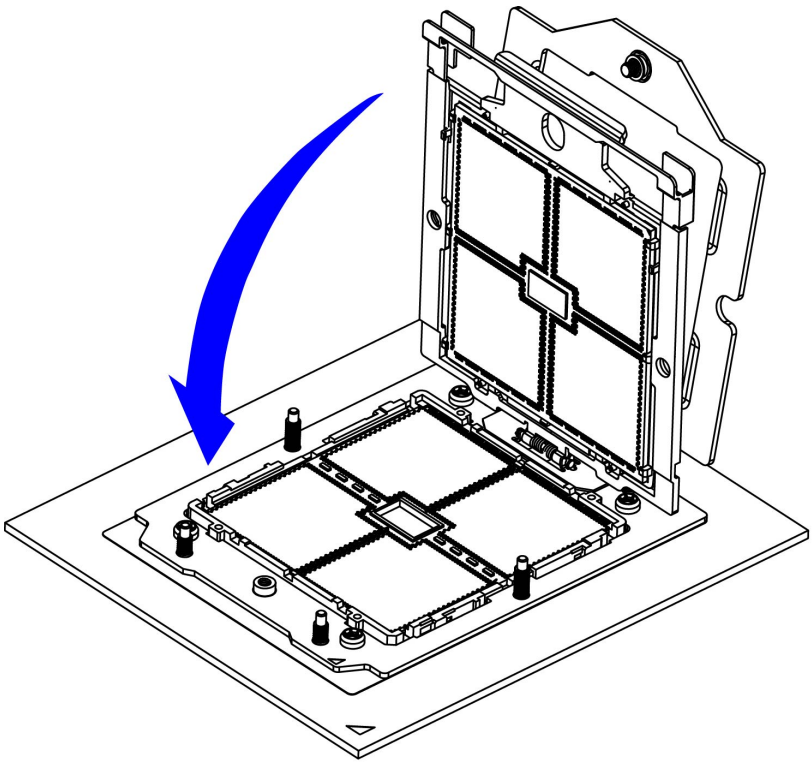


481614

Step 3 Secure the CPU into the socket.

a) Gently close the rail frame down to the flat, closed position.

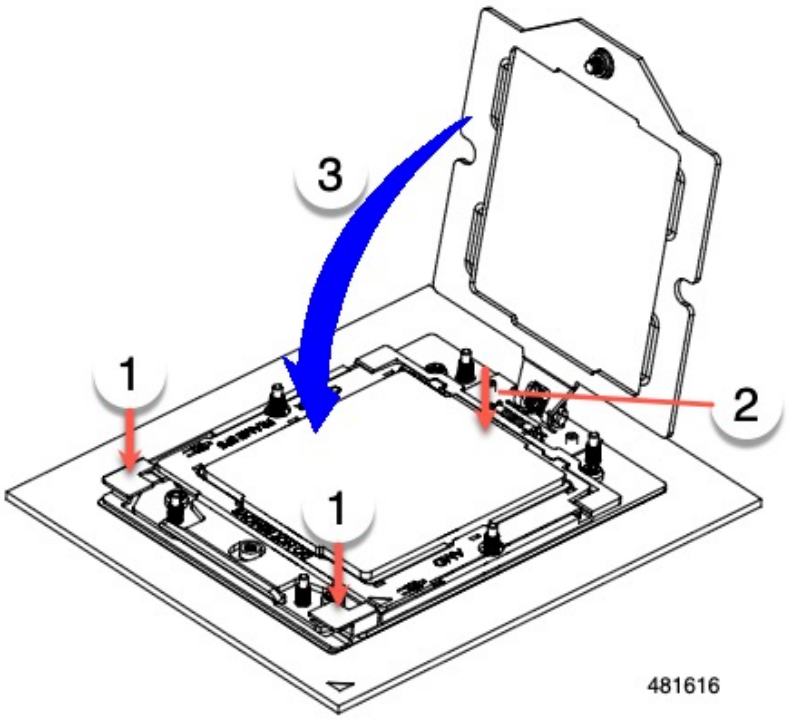
Figure 10: Closing the Rail Frame



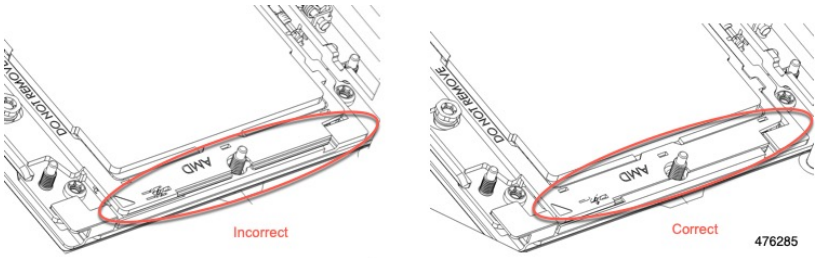
481615

b) Using your fingers: 1) Use a finger from each hand to press the two tabs of the Rail Frame on the Stiffener Frame. Slightly press down until you feel a click. 2) Slightly press the CPU package to ensure it is properly seated. 3) Gently close the retention frame down to the flat, closed position.

Figure 11: Closing the Retention Frame

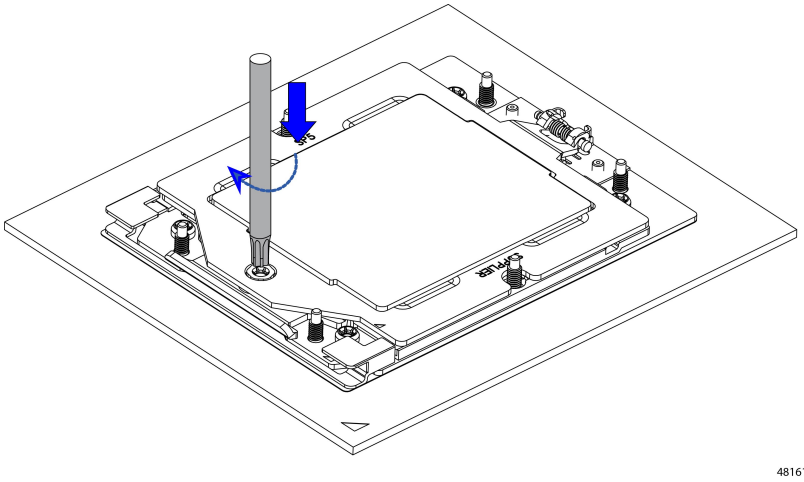


Note
Make sure the CPU is correctly inserted in the socket and does not sit on top of the retention frame. In the following example, a CPU is incorrect if it sits on top of the retention frame.



c) Tighten the screw on the retention frame.

Figure 12: Securing the Retention Frame



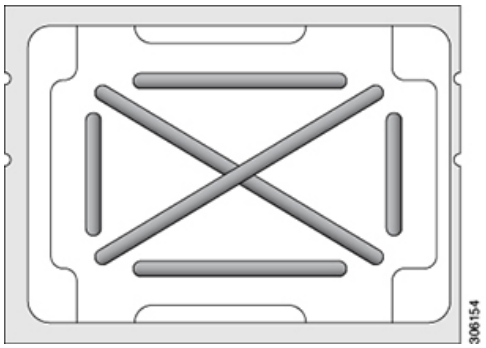
Step 4 Before installing the heatsink, apply new TIM:

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

- If you are installing a new heatsink, it is shipped with a pre-applied pad of TIM. Go to step 5.
- If you are reusing a heatsink, you must remove the old TIM from the heatsink and then apply new TIM to the CPU surface from the supplied syringe. Continue with step a below.

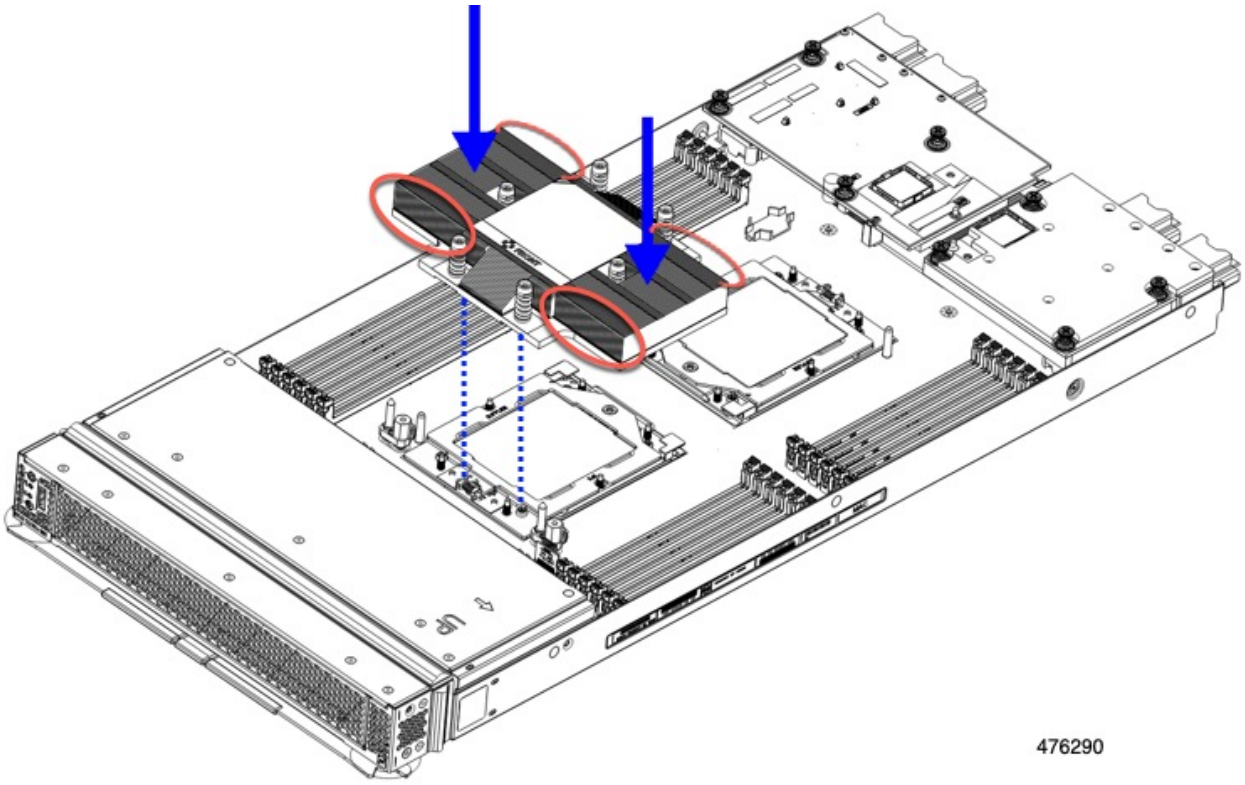
a) Apply the cleaning solution that is included with the heatsink cleaning kit (provided with replacement CPU and also orderable separately as UCSX-HSCK=) to the old TIM on the heatsink and let it soak for a least 15 seconds.
b) Wipe all of the TIM off the heatsink using the soft cloth that is included with the heatsink cleaning kit. Be careful to avoid scratching the heatsink surface.
c) Using the syringe of TIM provided with the new CPU, apply 1.5 cubic centimeters (1.5ml) of thermal interface material to the top of the CPU. Use the pattern shown below to ensure even coverage.

Figure 13: Thermal Interface Material Application Pattern



Step 5 Install the heatsink to the CPU:

a) Holding the heatsink level and by the vertical edges of the fins, align the heatsink over the CPU socket, and make sure to align the screws with their corresponding screw holes.
b) Keeping the heatsink level, lower it onto the CPU socket.



c) Use a T-20 Torx driver to tighten the six captive screws that secure the heatsink. Be sure the heatsink is installed correctly using the alignment pin.

Caution
Alternate tightening the heatsink screws evenly so that the heatsink remains level while it is lowered. Tighten the heatsink screws in the order shown on the heatsink label.

Step 6 Reconnect any cables that you removed.

Step 7 Return the compute node to service.

a) Replace the top cover to the compute node.
b) Replace the compute node in the chassis.
c) Power on the compute node.
d) Wait for Cisco Intersight, or another Cisco management platform, to complete discovery of the compute node.

Replacing Memory DIMMs

⚠

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

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

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

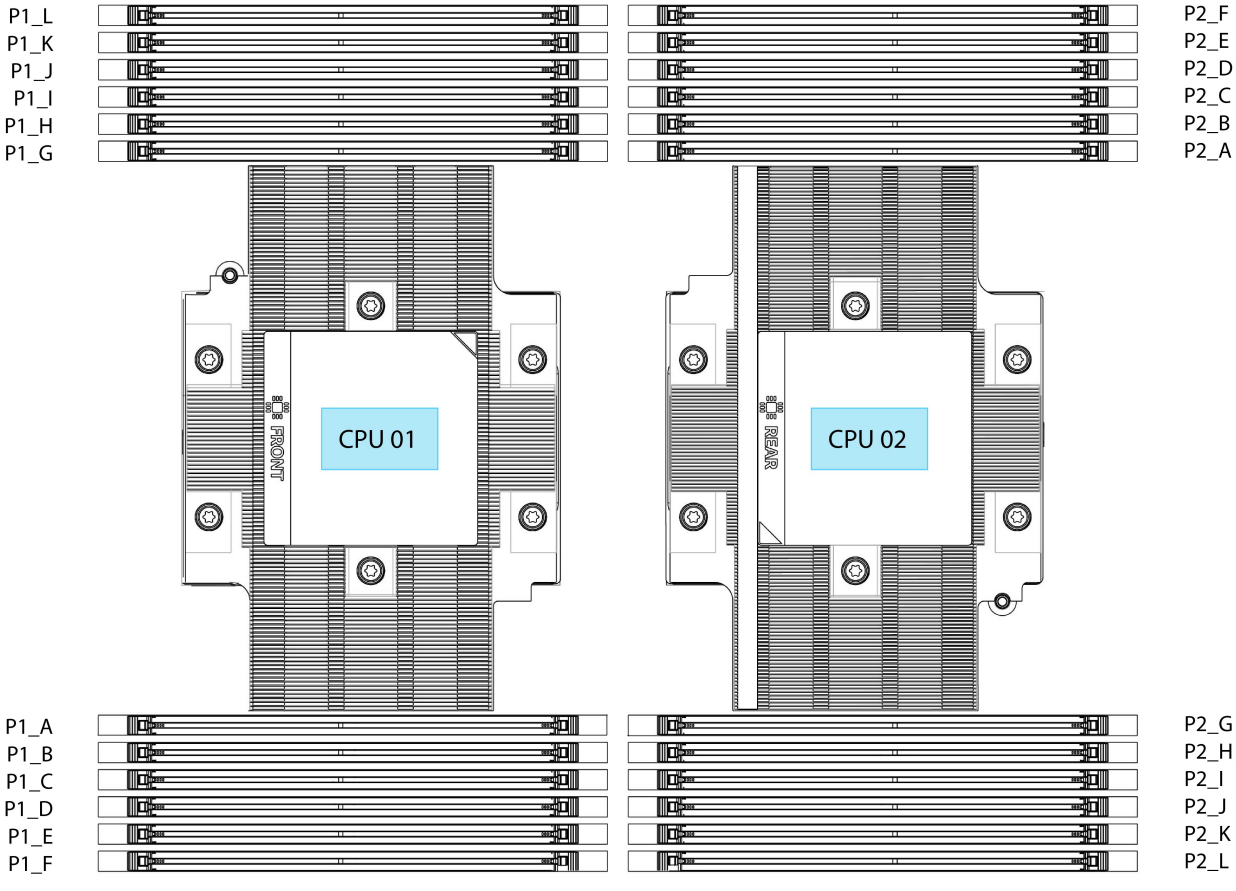
Memory Population Guidelines

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

DIMM Slot Numbering

The following figure shows the numbering of the DIMM slots on the motherboard. Because there is only one channel per DIMM, the DIMM slot number is not shown. However, you can think of each DIMM as being installed in slot 1.

Figure 14: DIMM Slot Numbering



DIMM Population Rules

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

- For a single-CPU compute node:
  - The minimum number of supported DIMMs is 1 and the maximum is 12.
  - Using 1, 2, 4, 6, 8, 10, or 12 DIMMs is supported. Using 3, 5, 7, 9, or 11 DIMMs is not supported.
- For a dual-CPU compute node:
  - The minimum number of supported DIMMs is 2 and the maximum is 24.
  - Using 2, 4, 8, 12, 16, 20, or 24 DIMMs is supported. Using 6, 10, 14, 18, or 22 DIMMs is not supported.
- Each CPU supports twelve memory channels, A through L.
  - CPU 1 supports channels P1_A, P1_B, P1_C, P1_D, P1_E, P1_F, P1_G, P1_H, P1_I, P1_J, P1_K, and P1_L.
  - CPU 2 supports channels P2_A, P2_B, P2_C, P2_D, P2_E, P2_F, P2_G, P2_H, P2_I, P2_J, P2_K, and P2_L.
- When both CPUs are installed, populate the DIMM slots of each CPU identically.
- In a single-CPU configuration, populate the channels for CPU1 only (P1_A through P1_L1).

Memory Population Order

For optimal performance, populate DIMMs in the order shown in the following table, depending on the number of CPUs and the number of DIMMs per CPU. If your compute node has two CPUs, balance DIMMs evenly across the two CPUs as shown in the table.

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

Table 1: DIMMs Population Order for 2 CPU Configuration

Number of DDR5 DIMMs (Recommended Configurations)	Populate CPU 1 Slot	Populate CPU2 Slots
2	P1_A	P2_A
4	P1_A P1_G	P2_A P2_G
8	P1_A P1_C P1_G P1_I	P2_A P2_C P2_G P2_I

12	P1_A P1_B P1_C P1_G P1_H P1_I	P2_A P2_B P2_C P2_G P2_H P2_I
16	P1_A P1_B P1_C P1_E P1_G P1_H P1_I P1_K	P2_A P2_B P2_C P2_E P2_G P2_H P2_I P2_K
20	P1_A P1_B P1_C P1_D P1_E P1_G P1_H P1_I P1_J P1_K	P2_A P2_B P2_C P2_D P2_E P2_G P2_H P2_I P2_J P2_K
24	All (P1_A) through (P1_L)	All (P1_A) through (P1_L)

Table 2: DIMMs Population Order for 1 CPU Configuration

Number of DDR5 DIMMs (Recommended Configurations)	Populate CPU 1 Slot
1	P1_A
2	P1_A P1_G

Number of DDR5 DIMMs (Recommended Configurations)	Populate CPU 1 Slot
4	P1_A P1_C P1_G P1_I
6	P1_A P1_B P1_C P1_G P1_H P1_I
8	P1_A P1_B P1_C P1_E P1_G P1_H P1_I P1_K
10	P1_A P1_B P1_C P1_D P1_E P1_G P1_H P1_I P1_J P1_K
12	All populated (P1_A) through (P1_L)

- The maximum combined memory allowed per CPU is 3TB (12 DIMM slots x 256 GB). For a dual-CPU configuration, the allowed system memory is 6TB

DIMM Mixing

Observe the DIMM mixing rules shown in the following table.

- For this compute node, all CPUs support only DDR5-5600 DIMMs, but they can run at 4800 speed.
- Some restrictions exist with 256GB DIMMs. You will be notified of the restrictions when attempting to configure and order your compute node.

Table 3: DIMM Mixing Rules

DIMM Parameter	DIMMs in the Same Bank
DIMM Capacity For example, 16GB, 32GB, 64GB, 128GB, and 256GB	You cannot mix DIMMs with different capacities and Revisions in the same bank (for example A1, B1). The Revision value depends on the manufacturers. Two DIMMs with the same PID can have different Revisions.
DIMM speed For example, 5600 GHz	You cannot mix DIMMs with different speeds and Revisions in the same bank (for example A1, B1). The Revision value depends on the manufacturers. Two DIMMs with the same PID can have different Revisions.

Installing a DIMM or DIMM Blank

To install a DIMM or a DIMM blank (UCS-DDR5-BLK=) into a slot on the compute node, follow these steps:

Procedure

Step 1 Open both DIMM connector latches.

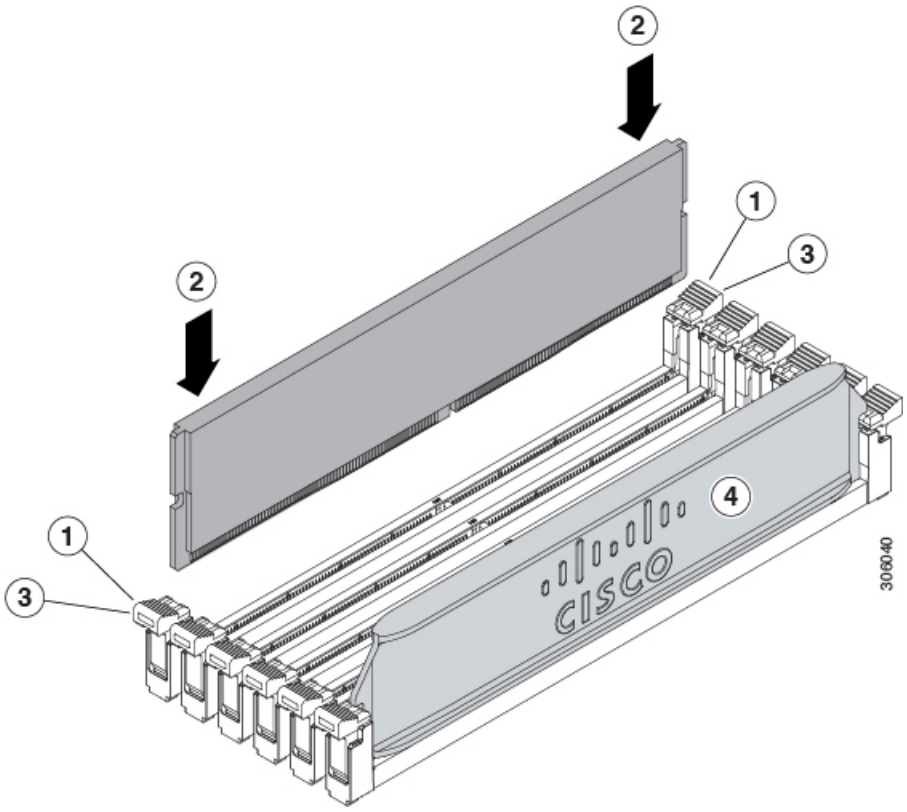
Step 2 Press evenly on both ends of the DIMM until it clicks into place in its slot.

Note
Ensure that the notch in the DIMM aligns with the slot. If the notch is misaligned, it is possible to damage the DIMM, the slot, or both.

Step 3 Press the DIMM connector latches inward slightly to seat them fully.

Step 4 Populate all slots with a DIMM or DIMM blank. A slot cannot be empty.

Figure 15: Installing Memory



Servicing the mLOM

The UCS X215c M8 compute node supports a modular LOM (mLOM) card to provide additional rear-panel connectivity. The mLOM socket is on the rear corner of the motherboard.

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

The following mLOM cards are supported on the compute node.

Table 4: Supported mLOM VICs on Cisco UCS X215c M8

UCSX-ML-V5Q50G-D	Cisco UCS Virtual Interface Card (VIC) 15420, Quad-Port 25G
UCSX-MLV5D200GV2D	Cisco UCS Virtual Interface Card (VIC) 15230, Dual-Port 40/100/200G mLOM

To service the mLOM card, use the following procedures:

- Installing an mLOM Card, on page 53

- Removing the mLOM, on page 54

Installing an mLOM Card

Use this task to install an mLOM onto the compute node.

Before you begin

If the compute node is not already removed from the chassis, power it down and remove it now. You might need to disconnect cables to remove the compute node.

Gather a torque screwdriver.

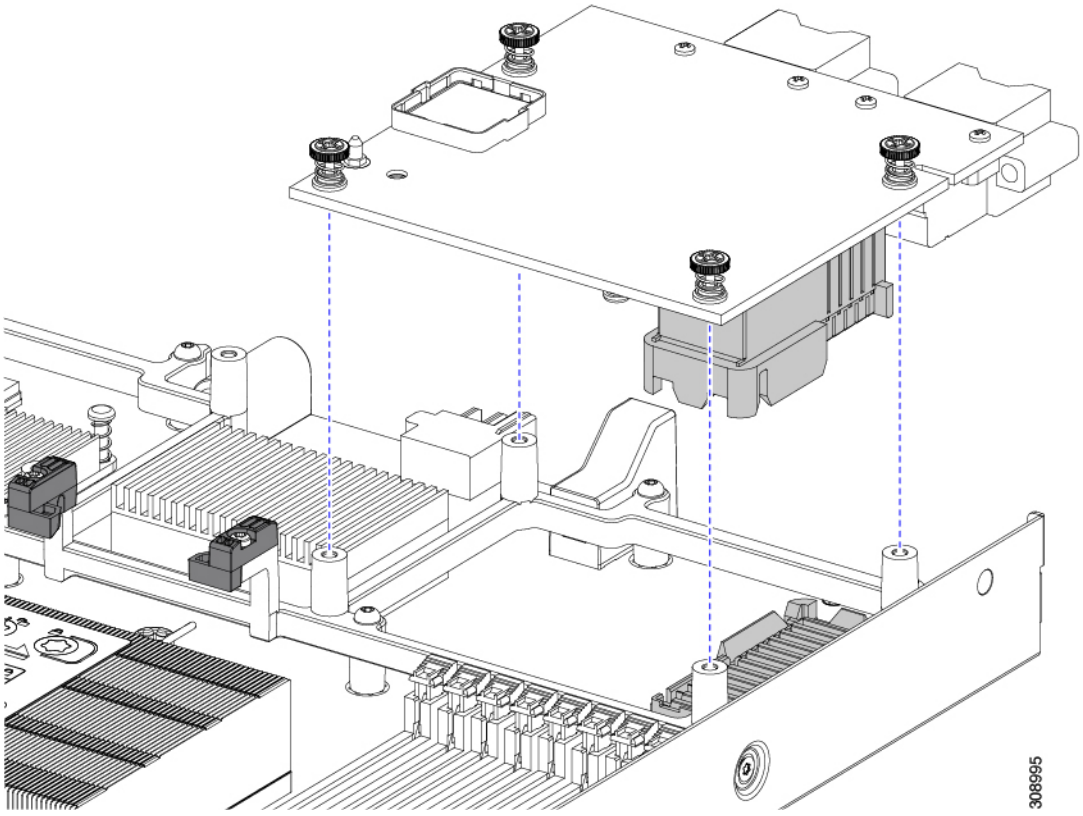
Procedure

Step 1 Remove the top cover.

See Removing a Compute Node Cover, on page 1.

Step 2 Orient the mLOM card so that the socket is facing down.

Step 3 Align the mLOM card with the motherboard socket so that the bridge connector is facing inward.



Step 4 Keeping the card level, lower it and press firmly to seat the card into the socket.

Step 5 Using a #2 Phillips torque screwdriver, tighten the captive thumbscrews to 4 in-lb of torque to secure the card.

Step 6 If your compute node has a bridge card (Cisco UCS VIC 15000 Series Bridge), reattach the bridge card.

See Installing a Bridge Card, on page 60.

Step 7 Replace the top cover of the compute node.

Step 8 Reinsert the compute node into the chassis. replace cables, and then power on the compute node by pressing the Power button.

Removing the mLOM

The compute node supports an mLOM in the rear mezzanine slot. Use this procedure to remove an mLOM.

Procedure

Step 1 Remove the compute node.

a) Decommission the compute node by using Cisco UCS management software, such as Cisco Intersight.

b) Remove the compute node from the chassis.

You might have to detach cables from the rear panel to provide clearance.

c) Remove the top cover from the compute node.

See Removing a Compute Node Cover, on page 1.

Step 2 If the compute node has a UCS VIC 15000 Series Bridge Card, remove the card.

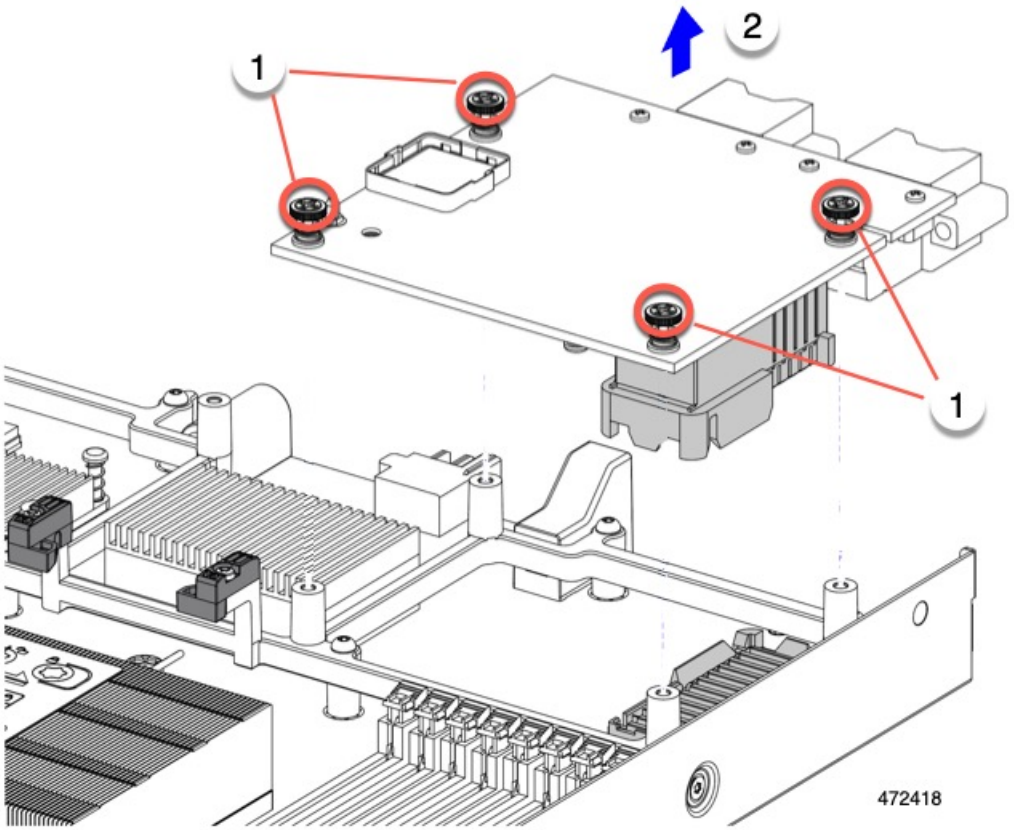
See Removing the Bridge Card, on page 59.

Step 3 Remove the MLOM.

a) Using a #2 Phillips head screwdriver, loosen the two captive thumbscrews.

b) Lift the MLOM off of its socket.

You might need to gently rock the mLOM card while lifting it to disengage it from the socket.



What to do next

After completing service, reinstall the VIC. See Installing a Rear Mezzanine Card in Addition to the mLOM VIC, on page 57.

Servicing the Rear Mezzanine

The UCS X215c M8 compute node supports a Rear Mezzanine card in the rear mezzanine slot. The VIC can be either half-slot or full-slot in size.

The following Rear Mezzanine cards are supported on the compute node.

Table 5: Supported Rear Mezzanine VICs on Cisco UCS X215c M8

UCSX-ME-V5Q50G-D	Cisco UCS Virtual Interface Card (VIC) 15422, Quad-Port 25G
UCSX-V4-PCIME	UCS PCI Mezz card for X-Fabric Connectivity

Cisco Virtual Interface Card (VIC) Considerations

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

- A blade with only one mezzanine card is an unsupported configuration. With this configuration, blade discovery does not occur through Cisco UCS management software. No error is displayed.

Removing a Rear Mezzanine

The compute node supports a VIC in the rear of the compute node. Use this procedure to remove the Rear Mezzanine VIC.

Procedure

Step 1 Remove the compute node.

a) Decommission the compute node by using Cisco UCS management software.
b) Remove the compute node from the chassis. You might have to detach cables from the rear panel to provide clearance.
c) Remove the top cover from the compute node. See Removing a Compute Node Cover, on page 1.

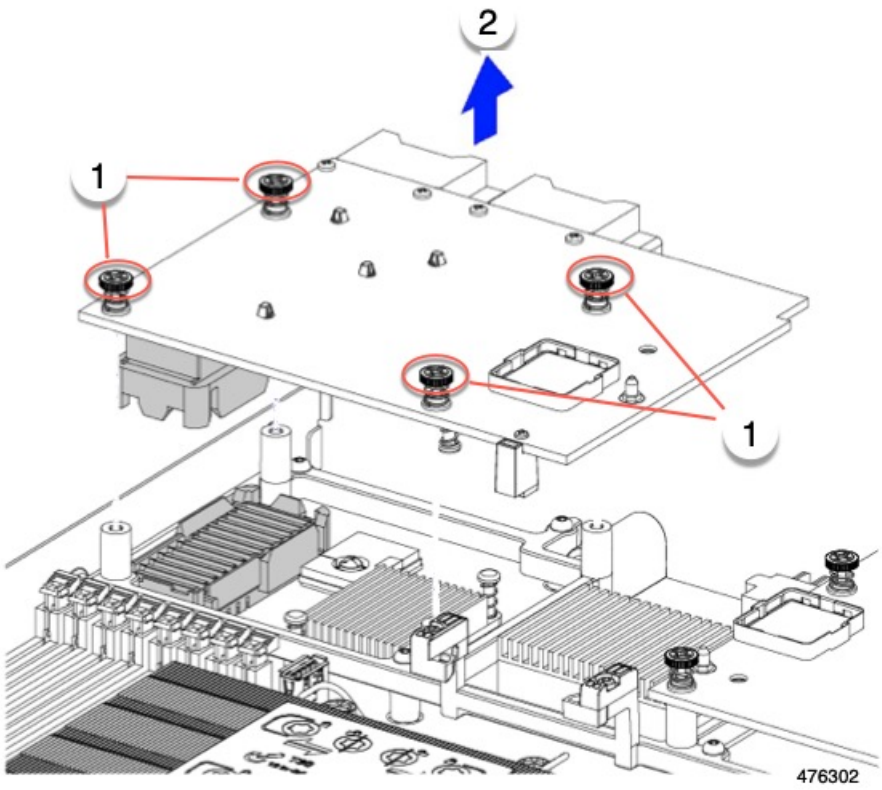
Step 2 If the compute node has a UCS VIC 15000 Series Bridge Card, remove the card.

See Removing the Bridge Card, on page 59.

Step 3 Remove the Rear Mezzanine.

a) Using a #2 Phillips head screwdriver, loosen the captive thumbscrews.
b) Lift the VIC off of its socket.

You might need to gently rock the Rear Mezzanine card while lifting it to disengage it from the socket.



Installing a Rear Mezzanine Card in Addition to the mLOM VIC

The compute node has a rear mezzanine slot which can accept a virtual interface card (VIC) unless the compute node has a full size mLOM. In the case of a separate mLOM and VIC, another component (the UCS VIC 14000 Series Bridge is required to provide data connectivity between the mLOM and VIC. See Installing a Bridge Card, on page 60.

Use this task to install a VIC in the rear mezzanine slot.

Note The VIC installs upside down so that the connectors meet with the sockets on the compute node.

Before you begin

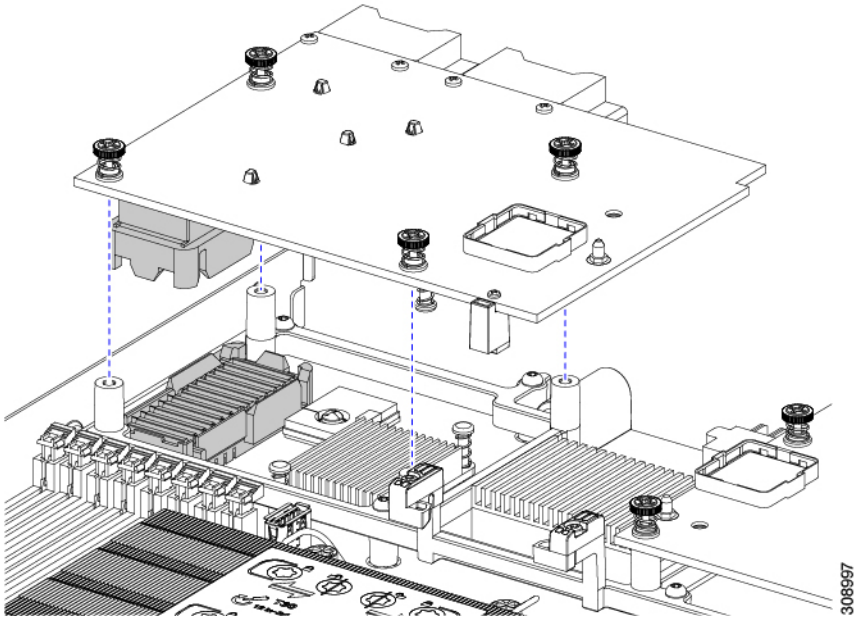
Gather a torque screwdriver.

Procedure

Step 1 Orient the VIC with the captive screws facing up and the connectors facing down.

Step 2 Align the VIC so that the captive screws line up with their threaded standoffs, and the connector for the bridge card is facing inward.

Step 3 Holding the VIC level, lower it and press firmly to seat the connectors into the sockets.



Step 4 Using a #2 Phillips torque screwdriver, tighten the captive screws to 4 in-lb of torque to secure the VIC to the compute node.

What to do next

- If the mLOM card is already installed, install a bridge card. Go to Installing a Bridge Card, on page 60.
- If not, install the mLOM, which must be installed before the bridge card can be attached. Go to Installing an mLOM Card, on page 53.

Servicing the Bridge Card

The compute node supports a Cisco UCS Series 15000 Bridge Card (UCSX-V5-BRIDGE-D) that spans between the rear mezzanine slot and the MLOM slot. The bridge card connects the UCS X-Series Compute Node to the following Intelligent Fabric Modules (IFMs) in the server chassis that contains the compute nodes:

- Cisco UCS X9108 25G Intelligent Fabric Module (UCSX-I-9108-25G)
- Cisco UCS X9108 100G Intelligent Fabric Module (UCSX-I-9108-100G)

See the following topics:

- Removing the Bridge Card, on page 59
- Installing a Bridge Card, on page 60

Removing the Bridge Card

Use the following procedure to remove the bridge card.

Procedure

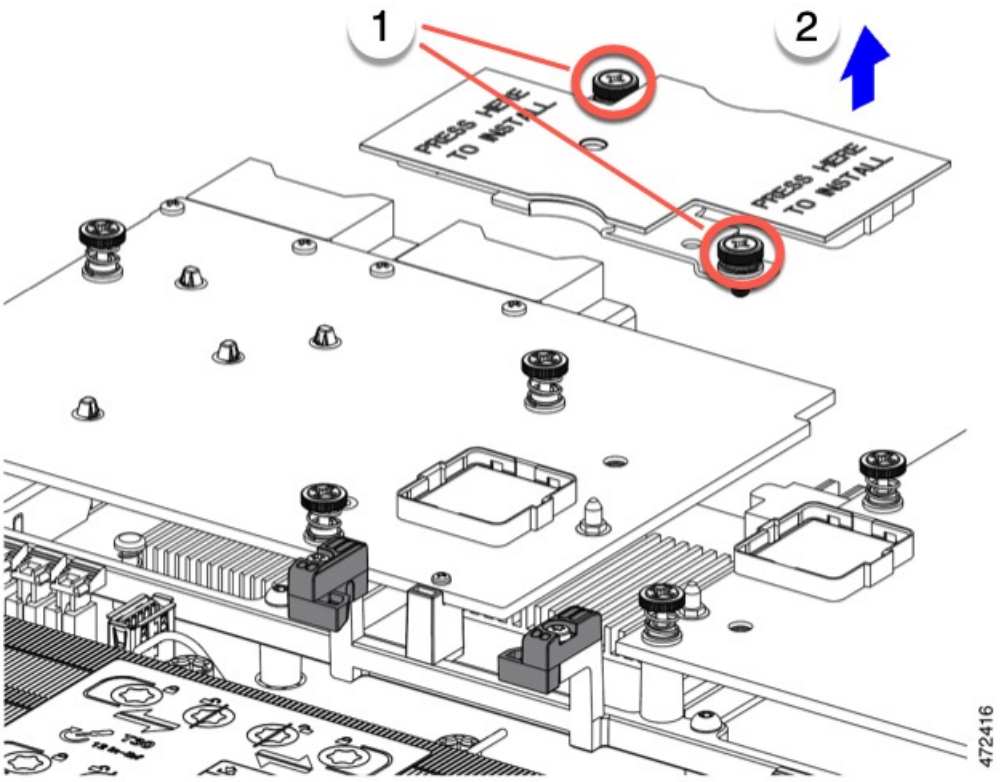
Step 1 Remove the compute node.

a) Decommission the compute node by using Cisco UCS management software.
b) Remove the compute node from the chassis. You might have to detach cables from the rear panel to provide clearance.
c) Remove the top cover from the compute node. See Removing a Compute Node Cover, on page 1.

Step 2 Remove the bridge card from the motherboard.

a) Using a #2 Phillips screwdriver, loosen the two captive screws.
b) Lift the bridge card off of the socket.

Note
You might need to gently rock the bridge card to disconnect it.



What to do next

Choose the appropriate option:

- Perform service on the MLOM. See Servicing the mLOM, on page 52.
- Perform service on the VIC. See Servicing the Rear Mezzanine, on page 55.
- Reinstall the bridge card. See Install the Bridge card.

Installing a Bridge Card

The Cisco UCS VIC 15000 Series Bridge is a physical card that provides data connection between the mLOM and VIC. Use this procedure to install the bridge card.

Note The bridge card installs upside down so that the connectors meet with the sockets on the MLOM and VIC.

Before you begin

To install the bridge card, the compute node must have an mLOM and a VIC installed. The bridge card ties these two cards together to enable communication between them.

If these components are not already installed, install them now. See:

- Installing a Rear Mezzanine Card in Addition to the mLOM VIC, on page 57

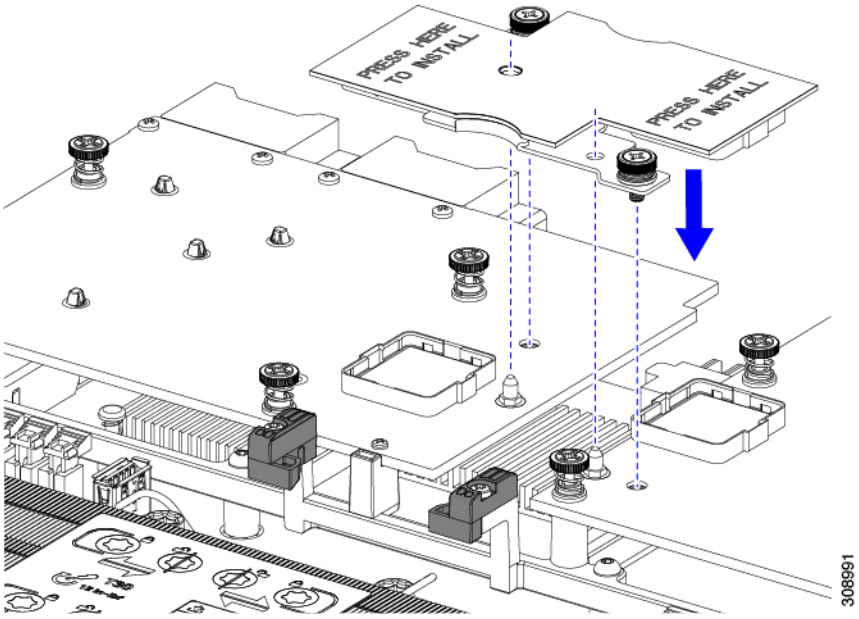
Procedure

Step 1 Orient the bridge card so that the Press Here to Install text is facing you.

Step 2 Align the bridge card so that the connectors line up with the sockets on the MLOM and VIC.

When the bridge card is correctly oriented, the hole in the part's sheet metal lines up with the alignment pin on the VIC.

Step 3 Keeping the bridge card level lower it onto the MLOM and VIC cards and press evenly on the part where the Press Here to Install text is.



Step 4 When the bridge card is correctly seated, use a #2 Phillips screwdriver to secure the captive screws.

Caution
Make sure the captive screws are snug, but do not overdrive them or you risk stripping the screw.

Servicing the Trusted Platform Module (TPM)

The Trusted Platform Module (TPM) is a component that can securely store artifacts used to authenticate the compute node. These artifacts can include passwords, certificates, or encryption keys. A TPM can also be

used to store platform measurements that help ensure that the platform remains trustworthy. Authentication (ensuring that the platform can prove that it is what it claims to be) and attestation (a process helping to prove that a platform is trustworthy and has not been breached) are necessary steps to ensure safer computing in all environments. It is a requirement for the Intel Trusted Execution Technology (TXT) security feature, which must be enabled in the BIOS settings for a compute node equipped with a TPM.

The UCS X215c M8 Compute Node supports the Trusted Platform Module 2.0, which is FIPS140-2 compliant and CC EAL4+ certified (UCSX-TPM2-002D=).

To install and enable the TPM, go to Enabling the Trusted Platform Module, on page 62.

Note Removing the TPM is supported only for recycling and e-waste purposes. Removing the TPM will destroy the part so that it cannot be reinstalled.

To remove the TPM, go to Removing the Trusted Platform Module (TPM).

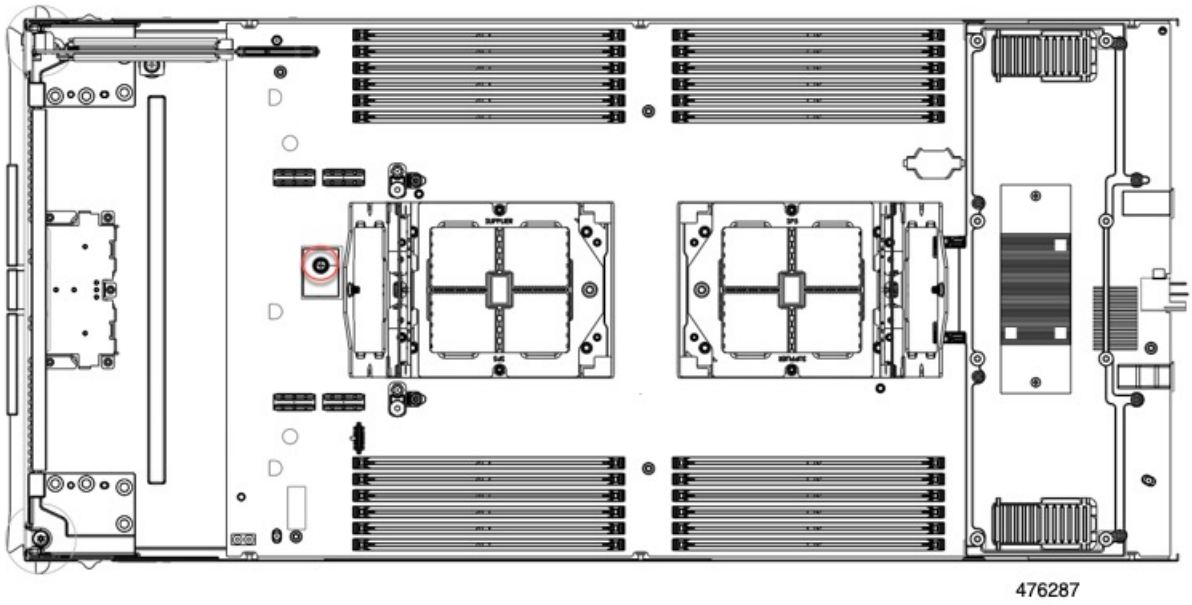
Enabling the Trusted Platform Module

Use this task to enable the TPM:

Procedure

Step 1 Install the TPM hardware.

a) Decommission, power off, and remove the compute node from the chassis.
b) Remove the top cover from the compute node as described in Removing and Installing the Compute Node Cover, on page 1.
c) Install the TPM to the TPM socket on the compute node motherboard and secure it using the one-way screw that is provided. See the figure below for the location of the TPM socket.
d) Return the compute node to the chassis and allow it to be automatically reacknowledged, reassociated, and recommissioned.
e) Continue with enabling TPM support in the compute node BIOS in the next step.



Step 2 Enable TPM Support in the BIOS.