



Cisco UCS X10c Front Mezzanine GPU Module Installation and Service Guide

First Published: 2022-07-14

Last Modified: 2023-03-24

Americas Headquarters

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



CONTENTS

CHAPTER 1

Overview 1

Overview 1

Front Panel 1

Drive Bays 2

Drives and Drive LEDs 2

CHAPTER 2

Installing the GPU Module 3

Installing and Removing the GPU Module 3

Required Equipment 3

Removing the Front Mezzanine GPU Module 4

Installing the Front Mezzanine GPU Module 6

Front Mezzanine GPU Module Configuration 8

CHAPTER 3

Servicing the GPU Module 9

Servicing the GPU Module 9

Removing a Compute Node Cover 9

Replacing a GPU Cage 9

Cabling Considerations 10

Removing the GPU Cage 12

Installing a GPU Cage 14

Replacing a GPU Card 15

GPU Card Guidelines and Restrictions 15

Removing a GPU Card 15

Installing a GPU Card 16

Removing a GPU Filler Blank 17

Installing a GPU Filler Blank 18

Replacing a Drive 19

 NVMe SSD Requirements and Restrictions 20

 Hot Plug Considerations 20

 Removing a Drive 20

 Installing a Drive 21

 Removing a Drive Blank 22

 Installing a Drive Blank 23

Recycling the PCB Assembly (PCBA) 24

Installing a Compute Node Cover 28

APPENDIX A **Technical Specifications 29**

 GPU Module Specifications 29



Preface

This preface contains the following topics:

- [Bias-Free Documentation, on page v](#)
- [Introduction, on page v](#)
- [Communications, Services, and Additional Information, on page vii](#)

Bias-Free Documentation



Note The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.

Introduction

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

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

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated

in a commercial environment. This equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

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

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

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

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

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

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

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

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

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

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be

present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at [Cisco Profile Manager](#).
- To get the business impact you're looking for with the technologies that matter, visit [Cisco Services](#).
- To submit a service request, visit [Cisco Support](#).
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit [Cisco Marketplace](#).
- To obtain general networking, training, and certification titles, visit [Cisco Press](#).
- To find warranty information for a specific product or product family, access [Cisco Warranty Finder](#).

Cisco Bug Search Tool

[Cisco Bug Search Tool](#) (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.



CHAPTER 1

Overview

This chapter contains the following topics:

- [Overview, on page 1](#)
- [Front Panel, on page 1](#)

Overview

The Cisco UCS X10c Front Mezzanine GPU module is a passive PCIe Gen 4 front mezzanine option for Cisco UCS X-Series compute nodes. The module is a mixed mode assembly that provides both storage and GPU acceleration.


Each Cisco UCS X10c front mezzanine GPU module consists of the following components:

- A GPU adapter card supporting zero, one or two, Cisco T4 GPUs (UCSX-GPU-T4-MEZZ).
Each GPU is connected directly into the GPU adapter card by a x8 Gen 4 PCI connection.
- A storage adapter and riser card supporting zero, one, or two U.2 NVMe drives. NVMe RAID is supported through Intel VROC key.

Front Panel

The front mezzanine GPU module supports up to 2 U.2 NVMe drives and up to 2 GPUs.

In the following illustration, the top row of the module is slot 1 and the bottom row of the module is slot 2.




1	U.2 NVMe drive slot 1	3	GPU slot 1
2	U.2 NVMe drive slot 2	4	GPU slot 2

Drive Bays

Each Cisco UCS X10c Front Mezzanine GPU Module has a front mezzanine slot that can support a maximum of two 2.5-inch U.2 NVMe drives. A drive blank panel (UCSC-BBLKD-S2) must cover any empty drive bays.

Drives and Drive LEDs



Each U.2 is a front-loading drive that slides into and out of the front mezzanine module's storage cage.



1	Drive Ejector button	2	Drive Ejector
3	Drive Activity LED	4	Drive Status LED

Each U.2 drive also has a pair of LEDs that indicate drive health and activity.

Table 1: Drive LEDs

LED	Color	Description
Drive Activity 	Off	Inactive.
	Flashing Green	I/O activity.
	Flashing Green	If this LED is flashing green and the Drive Health LED is simultaneously flashing amber, a drive rebuild is in progress.
Drive Health 	Off	No fault detected, the drive is not installed, or it is not receiving power.
	Amber	Fault detected
	Flashing Amber 4 Hz	If this LED is flashing amber and the Drive Activity LED is flashing green simultaneously, a drive rebuild is in progress.



CHAPTER 2

Installing the GPU Module

This chapter contains the following topics:

- [Installing and Removing the GPU Module, on page 3](#)
- [Front Mezzanine GPU Module Configuration, on page 8](#)

Installing and Removing the GPU Module

The front mezzanine module installs into the front mezzanine slot of Cisco UCS X-Series compute nodes.

- [Required Equipment, on page 3](#)
- [Removing the Front Mezzanine GPU Module, on page 4](#)
- [Installing the Front Mezzanine GPU Module, on page 6](#)

Required Equipment

Cisco provides a kit (UCSX-X10C-GPUFM) for installing the Cisco UCS X10c front mezzanine GPU module. The kit contains the following:

- Required GPU cables
- The front mezzanine cage including:
 - The storage cage which supports a maximum of two U.2 drives
 - The GPU cage which supports a maximum of two Cisco T4 GPUs (UCSX-GPU-T4-MEZZ).



Note If you order the Cisco UCS X10c front mezzanine GPU module with one GPU, the unused GPU slot is pre-populated with a filler module. You can purchase additional filler modules from Cisco

To provide flexibility for your deployment, the U.2 drives and GPUs are orderable separately. However, if you order the Cisco UCS X10c front mezzanine GPU module with one GPU, the unused GPU slot is pre-populated with a filler module. You can purchase additional filler modules from Cisco.

The following additional equipment, which is not provided by Cisco, is required to install or remove the GPU module.

- T8 Torx screwdriver
- #2 Phillips screwdriver
- Torque wrench

Removing the Front Mezzanine GPU Module

To remove an existing front mezzanine module, you will remove screws on the side of the compute node's sheet metal tray and on the node's motherboard.



Caution Do not operate the compute node without a front mezzanine module installed.

Use the following procedure to remove an existing front mezzanine module.

Before you begin

Review [Cabling Considerations, on page 10](#) for important information before beginning this task.


The compute node cables must be disconnected, and the compute node must be removed from the chassis to perform this task.

Step 1 Remove the top cover of the compute node.


[Removing a Compute Node Cover, on page 9.](#)

Step 2 Remove the screws.


a) Using a T8 Torx screwdriver, remove the four T8 Flathead screws, two screws per side.



b) Using a #2 Phillips screwdriver, loosen the two #2 Phillips captive screws at the rear of the front mezzanine module.



Step 3 Grasp the front mezzanine module at the middle of the top front edge and lift the front mezzanine module straight up to remove it from the compute node.



What to do next

Connect the front mezzanine GPU module to the compute node. [Installing the Front Mezzanine GPU Module, on page 6](#)

Installing the Front Mezzanine GPU Module

To install the front mezzanine GPU module, you will insert screws on the side of the compute node's sheet metal tray and on the node's motherboard.

Use the following task to install the X10c front mezzanine GPU module.


Before you begin

Review [Cabling Considerations, on page 10](#) for important information before beginning this task.


Step 1 Holding the front mezzanine module by the center of the top edge, keep the module level, and lower it down onto the compute node (1).

Step 2 Align the rear captive screws with the threaded standoffs (2).

Note The motherboard has guide pins that assist with installing the module onto the compute node. Use these guide pins while installing the module.



- Step 3** Press down evenly on the module to seat the module's connectors into their sockets on the motherboard.
- Step 4** Using a Phillips screwdriver, tighten the two rear captive screws to 5 in-lb of torque.
- Step 5** Using a T8 Torx driver, tighten the four T8 flat head screws on the sides of the compute node (2 screws per side) to secure the module to the compute node.



- Step 6** Replace the compute node top cover and reinstall the compute node into the chassis.

See [Installing a Compute Node Cover](#) , on page 28.

Front Mezzanine GPU Module Configuration

The Cisco UCS X10c front mezzanine GPU module can be configured and managed using the Cisco Intersight management platform in Intersight Managed Mode (Cisco Intersight Managed Mode). For details, see the *Cisco Intersight Managed Mode Configuration Guide*, which is available at the following URL: [Cisco Intersight Managed Mode Configuration Guide](#)



CHAPTER 3

Servicing the GPU Module

This chapter contains the following topics:

- [Servicing the GPU Module, on page 9](#)
- [Recycling the PCB Assembly \(PCBA\), on page 24](#)
- [Installing a Compute Node Cover , on page 28](#)

Servicing the GPU Module

The X10c front mezzanine GPU module contains U.2 drives, GPU cables, GPU cages, and GPU cards as field-replaceable components.

To service these components, see the following topics:

- [Replacing a GPU Cage, on page 9](#)
- [Replacing a GPU Card, on page 15](#)
- [Replacing a Drive, on page 19](#)

Removing a Compute Node Cover

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

-
- Step 1** Press and hold the button down.
- Step 2** While holding the back end of the cover, slide it back, then pull it up to lift the top cover off of the compute node.
- By sliding the cover back, you enable the front edge to clear the metal lip on the rear of the front mezzanine module.
-

Replacing a GPU Cage

The GPU cage is an adapter card that provides two slots for GPUs cards. GPU cages provide the mechanical housing for individual GPU cards. Removing the GPU cage is a prerequisite for removing or installing a GPU card.

Use the following procedures to replace a GPU cage.


- [Removing the GPU Cage, on page 12](#)
- [Installing a GPU Cage, on page 14](#)

Cabling Considerations

Refer to the following illustration while reviewing these cabling considerations.



Note In the following illustration, the front mezzanine GPU module is upside down to easily show the cables.



1	Front Mezzanine GPU Module Slot Identifier, 1 or 2	2	MCIO cable, GPU cage end This cable connects to front Mezzanine GPU Module Slot 1
---	--	---	--

3	MCIO cable and connector, GPU cage end This cable connects to front Mezzanine GPU Module Slot 2	4	MCIO cable, front mezzanine PCB end.
---	--	---	--------------------------------------

If you are removing the entire GPU cage, be aware of the following considerations.




Note If you are removing or installing the GPU cards themselves, these considerations do not apply.

- Each GPU slot has an MCIO cable. One end of the cable connects to the GPU cage, and the other end connects to a socket on the mezzanine PCB.
- The slots on the Front mezzanine GPU module numbered 1 or 2. Although cables and connectors are not labeled, you can use these slot numbers to help identify which cable must be connected to which front mezzanine slot.



Note The GPU card slots are also numbered 1 or 2.

- Cables are not labeled, so make sure that you don't mix up the connections, such as connect the cable for slot 2 to slot 1.
- Before disconnecting cables, you will find it helpful to mark which cable connects to GPU 1 and which cable connects to GPU 2.
- If you are working with both GPU cables, do the following to allow more room for your fingers:
 - Disconnect slot 2's cable first, then slot 1's cable.
 - Connect slot 1's cable first, then slot 2's cable.
- When you are installing cables, make sure that you route them so that they do not pinch, bind, or interfere with the removal or installation of other parts.
- When installing a GPU cage, always connect the GPU-cage end of the cable first, then connect the mezzanine PCB-end of the cable.
- When removing a GPU cage, always disconnect the PCB-end of the cable first before detaching the GPU-cage end of the cable.
- In most cases, you can leave the cables connected at the front mezzanine end and disconnect cables from just the GPU end. However, if you must disconnect the front mezzanine end of the cable, the MCIO connectors are labeled on the front mezzanine PCB to help identify how to reconnect the cables.




540687

1	Front Mezzanine PCB connector for MCIO cable 1	2	Front Mezzanine PCB connector for MCIO cable 2
---	--	---	--

Removing the GPU Cage

The GPU cage houses up to two GPUs. Use the following task to remove a GPU cage.

-
- Step 1** If you have not done so already, remove the front mezzanine GPU module from the server.
See [Removing the Front Mezzanine GPU Module, on page 4](#).
- Step 2** Place the front mezzanine GPU module upside down on an ESD-safe work surface.
- Step 3** Locate the 2 Phillips-head screws at the rear of the GPU cage.




Step 4 Using a #2 Phillips screwdriver, remove the 2 Phillips-head screws from the GPU cage.

Step 5 Slide the GPU cage towards the rear of the front mezzanine GPU board.

Step 6 Holding the front edge of the GPU cage (the end opposite the 2 Phillips-head screws you loosened), rotate the GPU cage 90 degrees so that its vertical.

The front of the cage has a "T" nut as a retention feature. If you cannot rotate the GPU cage up, it is probably obstructed by the "T" nut. Repeat the previous step until the GPU cage rotates easily.

Note When rotating the GPU cage upward, make sure not to scrape or damage other parts on the compute node or the front mezzanine GPU module.



What to do next

Choose the appropriate option.

- [Removing a GPU Card, on page 15.](#)
- [Installing a GPU Card, on page 16.](#)
- [Installing a GPU Cage, on page 14.](#)

Installing a GPU Cage

Use the following task to reinsert a GPU cage into the front mezzanine GPU module.

- Step 1** Rotate the GPU cage down until it lies flat on the sheet metal as far away from the GPU module's front panel as possible.
- Step 2** Carefully slide the GPU cage toward the front of the front mezzanine GPU module until the rear captive screws line up with their screw holes.
- When properly installed, the front of the GPU cage is captured by "T" nuts.
- Step 3** Lift up on the GPU cage to verify that it is captured by the "T" nuts.
- If you cannot lift the GPU cage, it is correctly installed.
- Step 4** Using a #2 Phillips screwdriver, tighten each of the captive screws to 5 in-lb of torque.

Step 5 Replace the compute node's top cover and reinsert the compute node into the server chassis.

Replacing a GPU Card

A maximum of two GPU cards (UCSX-GPU-T4-MEZZ) are supported in the GPU cage in slots 1 or 2. Slots are numbered on the side of the GPU cage.



Note Do not operate the compute node with an empty GPU card slot.

Use the following tasks to replace a GPU card.

- [Removing a GPU Card, on page 15](#)
- [Installing a GPU Card, on page 16.](#)
- [Removing a GPU Filler Blank, on page 17](#)
- [Installing a GPU Filler Blank, on page 18](#)

GPU Card Guidelines and Restrictions

Be aware of the following guidelines and restrictions for GPU cards:

- Do not operate the compute node with an empty GPU card slot.
- In a one GPU system, it is a best practice to have the GPU in slot 1.
- If your front mezzanine GPU module has only one GPU card, the other GPU slot must contain a GPU filler blank (UCSX-GPUFM-BLK). If you ordered your GPU module with only one GPU, the filler blank is preinstalled at the factory.
- GPU slots are keyed so that they can accept only the supported GPU (UCSX-GPU-T4-MEZZ).




Caution Cisco offers multiple models of Cisco T4 GPU, but they are not all interchangeable between the different Cisco products. For the Cisco X10c Front Mezzanine GPU Module, the only Cisco T4 GPU supported is UCSX-GPU-T4-MEZZ, which has a custom heatsink. You can identify this GPU by the M3 screw on the GPU bracket, which is a different screw from any other Cisco T4 GPU. Because of the heatsink and screw, the UCSX-GPU-T4-MEZZ is supported only on the front mezzanine GPU module. Do not attempt to install the UCSX-GPU-T4-MEZZ in any other Cisco product.

- Always follow guidance on any service label or warning sticker physically installed on the GPU and front mezzanine module.

Removing a GPU Card

Use this task to remove a GPU from the GPU cage.

- Step 1** When the GPU cage is vertical, use a #2 Phillips screwdriver to loosen the captive Phillips-head screw at the bottom of the GPU slot.
- Step 2** Grasping the bracket and the opposite edge of the GPU, slide the GPU straight up to remove it from the GPU cage.



What to do next

Choose the appropriate option.


- [Installing a GPU Card, on page 16.](#)
- [Installing a GPU Filler Blank, on page 18.](#)
- [Installing a GPU Cage, on page 14.](#)

Installing a GPU Card

Use this task to install a supported GPU card (UCSX-GPU-T4-MEZZ) in the GPU slots in the GPU cage.

- Step 1** Using the GPU card's bracket, align the GPU with its slot in the GPU cage.
Each slot in the GPU cage has a notch that catches the GPU bracket.
- Step 2** Slide the GPU into its socket, making sure that the PCI connector seats into the socket inside the GPU cage.
- Note** If the GPU does not easily slide into the socket, gently remove it and verify that it is a Cisco T4 GPU (UCSX-GPU-T4-MEZZ).
- Step 3** Using a #2 Phillips screwdriver, tighten the captive Phillips-head screw to 5 in-lbs of torque.

Step 4 If needed, repeat this process to install another GPU card in a different GPU slot.



What to do next

[Installing a GPU Cage, on page 14](#)

Removing a GPU Filler Blank


Each GPU filler blank has a retaining clip on each side of the blank. The retaining clips hook behind the sheet metal of the GPU cage to secure the filler panel in place.

Use the following task to remove a GPU filler blank.



Note Do not operate the front mezzanine GPU module with an empty GPU slot.

- Step 1** Locate each of the retaining clips on the side of the GPU cage (1).
- Step 2** Press both of the filler blank's retaining clips inward until they clear the sheet metal of the GPU cage.
- Step 3** Grasping the filler blank, slide it straight up to remove it from the slot in the GPU cage (2).



What to do next


Insert a GPU card in the empty GPU slot. See [Installing a GPU Card, on page 16](#).

Installing a GPU Filler Blank

If your front mezzanine GPU module has only one GPU, it must be installed in the GPU 1 slot. The GPU 2 slot must contain a GPU filler blank.

GPU filler blank has a flexible retention clip that meets the sheet metal of the GPU cage to securely hold the blank in the slot. Use the following task to install a GPU filler blank.

- Step 1** Locate the retaining clips on the filler blank (2, in the illustration). Each side of the filler blank has a retaining clip.
- Step 2** Grasping the filler blank by the bracket, align the filler blank with the GPU slot.
- Step 3** Holding the filler blank level, slide it into the GPU slot until the retaining clips fit into place behind the sheet metal wall of the GPU cage (1).



- Step 4** Make sure that the tab on the filler bracket sits flush in the GPU slot.

What to do next

Reinsert the GPU cage onto the front mezzanine GPU module. [Installing a GPU Cage, on page 14](#)

Replacing a Drive

The front mezzanine GPU module supports up to two U.2 NVMe drives. Each drive is front-loading into the front of the module.

Use the following tasks to replace a U.2 drive.

- [Removing a Drive, on page 20](#)
- [Installing a Drive, on page 21](#)

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.

UEFI boot mode can be configured through the Boot Order Policy setting in the Server Policy supported by Cisco Intersight Managed Mode (IMM). For instructions about setting up UEFI boot mode through Cisco IMM, go to:

[Cisco Intersight Managed Mode Configuration Guide](#)

- NVMe PCIe SSDs cannot be controlled with a SAS RAID controller because NVMe SSDs interface with the server via the PCIe bus.
- UEFI boot is supported in all supported operating systems.

Hot Plug Considerations

Enabling Hot Plug Support

Surprise and OS-informed hot plug is supported with the following conditions:

- VMD must be enabled to support hot plug. VMD must be enabled before installing an OS on the drive.
- If VMD is not enabled, surprise hot plug is not supported, and you must do OS-informed hot plug instead.
- VMD is required for both surprise hot plug and drive LED support.

Removing a Drive


Use this task to remove a U.2 NVMe drive from the front mezzanine GPU module.



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.

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 21](#)
- [Installing a Drive Blank, on page 23](#)


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 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 U.2 NVMe drive, follow this 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.



Removing a Drive Blank


A maximum of U.2 NVMe drives are contained in the front mezzanine storage module. If your front mezzanine module has fewer than two U.2 drives, you must install drive blank panels in the empty drive bays.



Note Do not operate a front mezzanine GPU module that has empty drive bays without a drive blank panel.

Use this procedure to remove a drive blank.

- 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 21](#)
- [Installing a Drive Blank, on page 23](#)

Installing a Drive Blank

Use this task to install a drive blank.

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.



Recycling the PCB Assembly (PCBA)

Each front mezzanine assembly has four printed circuit boards (PCBs) that are connected to the sheet metal cage by 15 M3 screws.

To remove the PCBs, you must:

- Remove the compute node that contains the front mezzanine module from its chassis
- Remove the front mezzanine assembly from its compute node.
- Disassemble and remove additional parts to gain access to the PCBs.
- Disconnect the PCBs from the sheet metal to recycle them.
- Recycle each front mezzanine assembly in the Cisco UCS X-Series server chassis.

Use the following procedure to recycle the PCBs from the UCS front mezzanine GPU module.

Before you begin



Note **For Recyclers Only!** This procedure is not a standard field-service option. This procedure is for recyclers who will be reclaiming the electronics for proper disposal to comply with local eco design and e-waste regulations.

Gather the following tools before you start this procedure:


- #2 Phillips screwdriver
- T10 Torx screwdriver

Step 1 Remove the SSD drives.

See [Removing a Drive, on page 20](#).


Step 2 Remove the GPU module.

- Using a #2 Phillips screwdriver, remove the M3 screws.
- Grasp the GPU module and remove it.
- Grasp each of the two cable connectors and disconnect the cables from the module.




Step 3 (Optional) Remove the GPU card(s) from the GPU module.

- Using a #2 Phillips screwdriver, remove the M3 screws.
- Grasp each GPU at both ends and slide it out of the GPU module.




Step 4 Remove the GPU adapter.

- a) Using a T10 screwdriver, remove the M3 screws.




- b) Grasp the GPU adapter and remove it.
- c) Turn the GPU cage over and repeat this step to remove the other GPU adapter.




Step 5 Remove the front mezzanine riser card from the inside of the front mezzanine assembly.

- a) Using a T10 screwdriver, remove the M3 screws.
- b) Grasp the riser card and remove it.



Step 6 Remove each front mezz PCB.

- a) Using a T10 screwdriver, remove the M3 screws.




- b) Grasp each PCBA and remove it.

Step 7

Remove additional components from front mezzanine PCB.

- a) Turn the PCBA over and release the connector and remove the cables.
- b) Grasp the lightpipe and pry it up with your fingers to remove it.



Step 8 Recycle the sheet metal and PCBs in compliance with your local recycling and e-waste regulations.

Installing a Compute Node Cover

Use this task to install a removed top cover on the compute node.

- Step 1** Notice the cutouts on the rear of the top cover.
These cutouts receive the stopper pins on the compute node.
- Step 2** Holding the top cover with the rear angled down, lower it onto the compute node.
- Step 3** Slide the compute node's cover until it hits the stopper pins.
- Step 4** Lower the front of the top cover onto the compute node.
- Step 5** Keeping the compute node's cover flat, slide it forward until the release button clicks.
-



APPENDIX **A**

Technical Specifications

This appendix contains the following topics:

- [GPU Module Specifications, on page 29](#)

GPU Module Specifications

Physical Specifications

Specification	Value
Height	1.76 inches (44.70 mm)
Width	11.28 inches (286.51 mm)
Depth	8.11 inches (205.99 mm)
Weight	The weight depends on the components installed. <ul style="list-style-type: none">• Minimally configured compute node weight: 1.96 lb (0.89 kg)• Fully configured compute node weight: 3.86 lb (1.75 kg)

Environmental Specifications

The Cisco UCS X10c Front Mezzanine GPU Module installs into a supported compute node. Therefore, the GPU module itself inherits the environmental and power specifications from the compute node that is hosting the GPU module.

For environmental specifications, see the appropriate documentation:

- For the Cisco UCS X210c M6 Compute Node, go to: [Cisco UCS X210c M6 Compute Node Installation and Service Guide, Technical Specifications](#).
- For the Cisco UCS X210c M7 Compute Node, go to: [Cisco UCS X210c M7 Compute Node Installation and Service Guide, Technical Specifications](#).

