

Overview

This chapter contains the following topics:

- Cisco UCS X215c M8 Compute Node Overview, on page 1
- Local Console, on page 5
- Front Mezzanine Options, on page 6
- mLOM and Rear Mezzanine Slot Support, on page 8
- System Health States, on page 9
- Interpreting LEDs, on page 10
- Optional Hardware Configuration, on page 12

Cisco UCS X215c M8 Compute Node Overview

The Cisco UCS X215c M8 is a single-slot compute node that has two CPU sockets that can support a maximum of one Fourth Gen AMD EPYC™ Processors with up to 96 cores per processor and up to 384 MB of Level 3 cache per CPU or Fifth Gen AMD EPYC™ Processors with up to 196 cores per processor and up to 384 MB of Level 3 cache per CPU. The minimum system configuration requires one CPU installed in the CPU1 slot.

Additionally, the compute node supports the following features with one CPU or two identical CPUs:

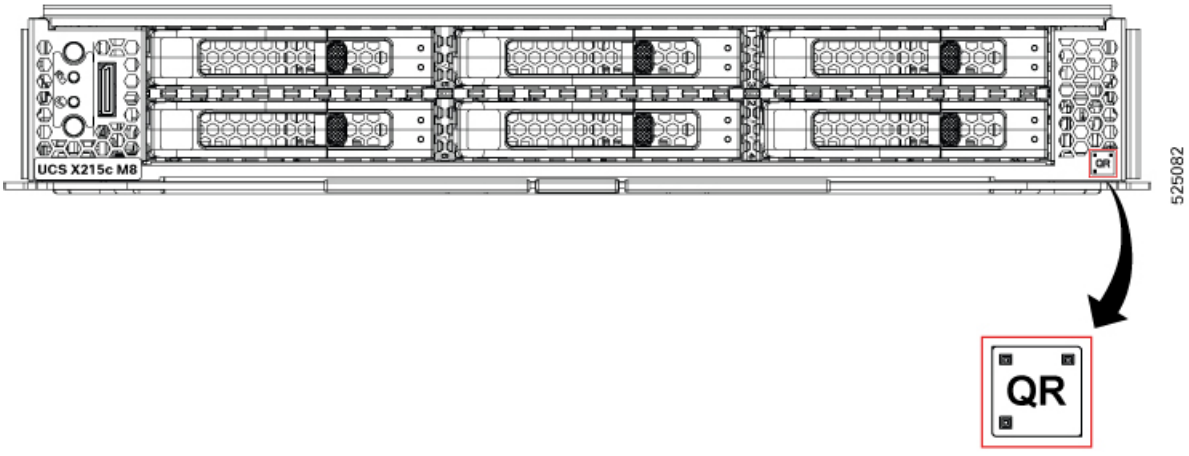
- 24 total DIMMs, 12 channels per CPU socket, 1 DIMM per channel.
- Up to 6TB of main memory with a maximum of 24 256 GB DDR5 5600 MT/s or DDR5 4800 MT/s or DDR5 6400 MT/s DIMMs, running at 6000 MT/s with Fifth Gen AMD EPYC processors.
- RAS is supported.
- One front mezzanine module can support the following:
  - A front storage module, which supports multiple different storage device configurations:
    - Up to six hot pluggable SAS/SATA/U.3 NVMe 2.5inch SSDs (slots 1-6).
    - SATA/SAS/U.3 drives can co-exist on the front mezzanine module. RAID volumes are restricted to same type of drives only. For example, RAID 1 volume need to use a set of SATA or SAS or U.3 NVMe drives.
    - Up to eight E3.S 1T PCIe Gen 5 drives through a pass-through controller.

  For additional information, see Front Mezzanine Options, on page 6.

- 1 modular LAN on motherboard (mLOM/VIC) module supporting a maximum of 200G traffic, 100G to each fabric. For more information, see mLOM and Rear Mezzanine Slot Support, on page 8.
- 1 rear mezzanine module (UCSX-V4-PCIME or UCSX-ME-V5Q50G).
- A mini-storage module with slots for up to two M.2 drives with optional hardware RAID. Two options of mini-storage exist:
  - One supporting M.2 SATA drives with a RAID controller (UCSX-M2-HWRD-FPS)
  - One supporting M.2 NVMe drives direct-attached to CPU 1 through a pass-through controller (UCSX-M2-PT-FPN).
- Local console connectivity through a OCU connector.
- Connection with a paired UCS PCIe module, such as the Cisco UCS X440p PCIe node, to support GPU offload and acceleration. For more information, see the Optional Hardware Configuration, on page 12.
- Up to eight UCS X215c M8 compute nodes can be installed in a Cisco UCS X9508 modular system.
- Through the Cisco UCS X9508 that hosts the Cisco X215c M8, connections to the following Cisco Fabric Interconnects are supported:
  - Cisco UCS Fabric Interconnect 6454
  - Cisco UCS Fabric Interconnect 64108
  - Cisco UCS Fabric Interconnect 6536

Compute Node Identification

Each Cisco UCS X215c M8 compute node features a node identification tag at the lower right corner of the primary node.



The node identification tag is a QR code that contains information that uniquely identifies the product, such as:

- The Cisco product identifier (PID) or virtual identifier (VID)
- The product serial number

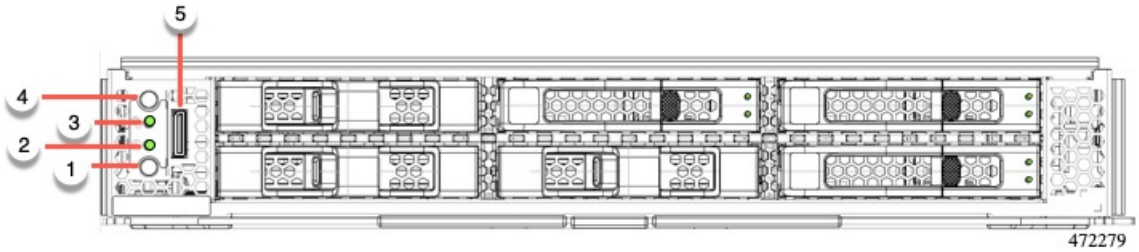
The product identification tag applies to the entire compute node, both the primary and secondary.

You will find it helpful to scan the QR code so that the information is available if you need to contact Cisco personnel.

Compute Node Front Panel

The Cisco UCS X215c M8 front panel contains system LEDs that provide visual indicators for how the overall compute node is operating. An external connector is also supported.

Compute Node Front Panel



1	Power LED and Power Switch The LED provides a visual indicator about whether the compute node is on or off. • Steady green indicates the compute node is on. • Steady Amber indicates the compute node is in Standby power mode. • Off or dark indicates that the compute node is not powered on. The switch is a push button that can power off or power on the compute node. See Front Panel Buttons, on page 4.	2	System Activity LED The LED blinks to show whether data or network traffic is written to or read from the compute node. If no traffic is detected, the LED is dark. The LED is updated every 10 seconds.

3	System Health LED A multifunction LED that indicates the state of the compute node. • Steady green indicates the compute node successfully booted to runtime and is in normal operating state. • Steady amber indicates that the compute node successfully booted but is in a degraded runtime state. • Blinking amber indicates that the compute node is in a critical state, which requires attention.	4	Locator LED/Switch The LED provides a visual indicator that glows solid blue to identify a specific compute node. The switch is a push button that toggles the Indicator LED on or off. See Front Panel Buttons, on page 4.
5	External Optical Connector (Oculink) that supports local console functionality.		

Front Panel Buttons

The front panel has some buttons that are also LEDs. See Compute Node Front Panel, on page 3.

- The front panel Power button is a multi-function button that controls system power for the compute node.
  - Immediate power up: Quickly pressing and releasing the button, but not holding it down, causes a powered down compute node to power up.
  - Immediate power down: Pressing the button and holding it down 7 seconds or longer before releasing it causes a powered-up compute node to immediately power down.
  - Graceful power down: Quickly pressing and releasing the button, but not holding it down, causes a powered-up compute node to power down in an orderly fashion.
- The front panel Locator button is a toggle that controls the Locator LED. Quickly pressing the button, but not holding it down, toggles the locator LED on (when it glows a steady blue) or off (when it is dark). The LED can also be dark if the compute node is not receiving power.

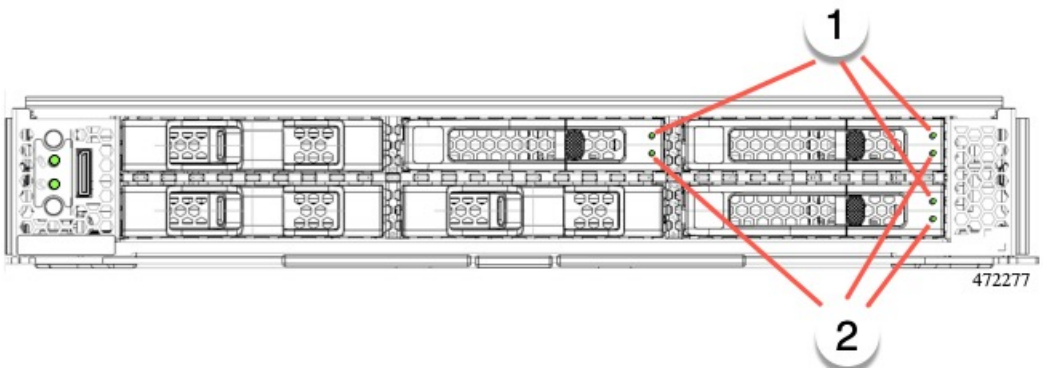
For more information, see Interpreting LEDs, on page 10.

Drive Front Panels

The front drives are installed in the front mezzanine slot of the compute node. SAS/SATA and NVMe drives are supported.

Compute Node Front Panel with SAS/SATA Drives

The compute node front panel contains the front mezzanine module, which can support a maximum of 6 SAS/SATA drives. The drives have additional LEDs that provide visual indicators about each drive's status.



1	Drive Health LED	2	Drive Activity LED

Compute Node Front Panel with NVMe Drives

The compute node front panel contains the front mezzanine module, which can support a maximum of six 2.5-inch NVMe drives.

Local Console

The local console connector is a horizontal oriented OcuLink on the compute node faceplate.

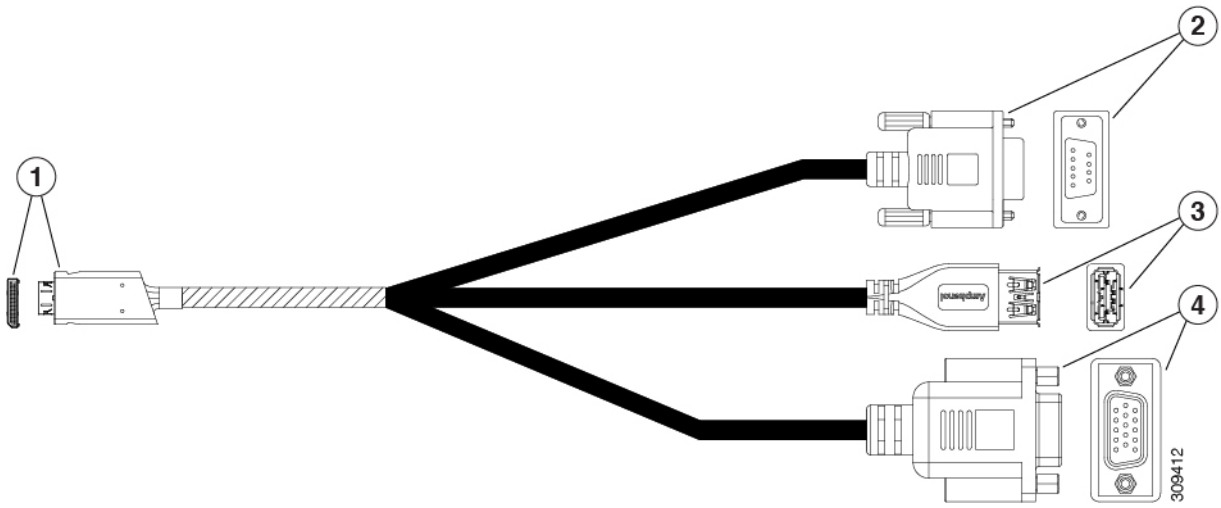
The connector allows a direct connection to a compute node to allow operating system installation directly rather than remotely.

The connector terminates to a KVM dongle cable (UCSX-C-DEBUGCBL) that provides a connection into a Cisco UCS compute node. The cable provides connection to the following:

- VGA connector for a monitor
- Host Serial Port
- USB port connector for a keyboard and mouse

With this cable, you can create a direct connection to the operating system and the BIOS running on a compute node. A KVM cable can be ordered in separately and it doesn’t come with compute node’s accessary kit.

Figure 1: KVM Cable for Compute Nodes



1	Oculink connector to compute node	**2**	Host Serial Port
3	USB connector to connect to single USB 3.0 port (keyboard or mouse)	**4**	VGA connector for a monitor

Front Mezzanine Options

The Cisco UCS X215c M8 Compute Node supports front mezzanine module storage through SAS/SATA or NVMe SSDs, and compute acceleration through GPUs. See:

- Storage Options, on page 6
- GPU Options, on page 8

Storage Options

The compute node supports the following local storage options in the front mezzanine module.

Cisco UCS 24G M1 Tri-Mode RAID Controller

The compute node supports the Cisco UCS 24G M1 Tri-Mode RAID Controller (UCSX-RAID-M1L6). This front mezzanine module supports:

- Support up to six 6 SAS/SATA SSDs, or up to four or six NVME SSDs as:
  - U.3 NVMe drives in slots 1 to 6 connected to the RAID controller at PCIe Gen4 and configurable with HW RAID.
- PCIe Gen3 and Gen4, x8 lanes
- Drive hot plug
- RAID support depends on the type of drives, and how they are configured.

- RAID is not supported in a mixture of SAS/SATA and U.3 NVMe drives in the same RAID group.
- The following RAID levels are supported across SAS/SATA and U.3 NVMe SSDs when the RAID group is either all SAS/SATA drives or all U.3 NVMe drives: RAID 0, 1, 5, 6, 00, 10, and 50.

- Support for drive slot 5 and 6 can be either Controller attached mode or direct attached mode. Only NVMe U.3 drives in drive slot 5 and 6 will become CPU attached in direct attach mode.

Cisco UCS X10c Passthrough Module

The compute node supports the Cisco FlexStorage NVMe passthrough controller, which is a passthrough controller for NVMe drives only. This module supports:

- Support up to six NVME SSDs in slots 1 through 6
- PCIe Gen3 and Gen4, x24 total lanes, partitioned as six x4 lanes
- Drive hot plug is supported
- Virtual RAID on CPU (VROC) is not supported, so RAID across NVME SSDs is not supported

Cisco UCS X10c E3.S Pass-Through Controller for E3.S Drives

The compute node can support an E3.S-drive based front mezzanine module, the Cisco UCS X10c E3.S Front Mezzanine Module ((UCSX-X10C-PTE3). This module is a passive PCIe Gen 5 front mezzanine option for Cisco UCS X-Series compute nodes, including the Cisco UCS X215c M8 compute node.

One module is supported per compute node, and each Cisco UCS X10c front mezzanine drives module supports up to eight E3.S 1T PCIe Gen5 drives.

For information about this hardware option, see the Cisco UCS X10c Pass Through Controller for E3.S Installation and Service Guide.

Cisco UCS X10c RAID Module

This storage option supports:

- Support up to six 6 SAS/SATA SSDs, or
- Up to four or six NVME SSDs as:
  - U.3 NVMe drives in slots 1 to 6 connected to the RAID controller at PCIe Gen4 and configurable with HW RAID.
- PCIe Gen3 and Gen4, x8 lanes
- Drive hot plug is supported
- RAID support depends on the type of drives and how they are configured in the RAID:
  - RAID is not supported in a mixture of SAS/SATA and U.3 NVMe drives in the same RAID group.
  - The following RAID levels are supported across SAS/SATA and U.3 NVMe SSDs when the RAID group is either all SAS/SATA drives or all U.3 NVMe drives: RAID0, 1, 5, 6, 00, 10, 50, and 60.
- RAID is not supported with a mixture of:

- SAS and SATA drives in the same RAID group
- SAS and U.3 NVMe drives in the same RAID group
- SATA and U.3 NVMe drives in the same RAID group

GPU Options

The compute node offers GPU offload and acceleration through the following optional GPU support.

Cisco UCS X10c Front Mezzanine GPU Module

As an option, the compute node can support:

- a GPU-based front mezzanine module, the Cisco UCS X10c Front Mezzanine GPU Module.

Each UCS X10c Front Mezzanine GPU Module contains:

- A GPU adapter card supporting zero, one or two, Cisco L4-MEZZ GPUs (UCSX-GPU-L4-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.3 NVMe drives.
- PCI Gen 3 and Gen4, x32 configured as one x 16 plus two x8 lanes
- Drive hot plug is supported

For information about this hardware option, see the Cisco UCS X10c Front Mezzanine GPU Module Installation and Service Guide.

mLOM and Rear Mezzanine Slot Support

The following rear mezzanine and modular LAN on motherboard (mLOM) modules and Virtual interface cards (VICs) are supported.

- Cisco UCS VIC 15422 (UCSX-ME-V5Q50G) occupies the rear mezzanine slot. This card supports:
  - Four 25G KR interfaces.
  - Occupying the compute node's rear mezzanine slot at the bottom rear of the chassis.
  - An included bridge card extends this VIC's 2x 50 Gbps of network connections through Cisco Intelligent Fabric Modules (IFMs), bringing the total bandwidth to 100 Gbps per fabric (for a total of 200 Gbps per compute node).
  - This card supports Secure Boot.
- The Cisco UCS PCI Mezz card for X-Fabric Gen5 (UCSX-V5-PCIME) is a rear mezzanine card that features:
  - Two physical ports.
  - Two PCIE Gen5 x16 electrical lanes to each CPU on the compute node.

- Two PCIE Gen5 x16 electrical lanes to each Cisco X-Fabric.
- This card is required to provide connectivity between the compute node and the Cisco UCS X580p PCIe node, if the PCIe node is present.

- The Cisco UCS PCI Mezz Card for X-Fabric (UCSX-V4-PCIME) is a rear mezzanine card that features:
  - Two physical ports.
  - Two PCIe Gen4 x16 electrical lanes to each CPU on the compute node.
  - Two PCIE Gen4 x16 electrical lanes to each Cisco X-Fabric.
  - This card is required to provide connectivity between the compute node and the Cisco UCS PCIe node for GPU access, if present.

- Cisco UCS VIC 15420 mLOM (UCSX- ML-V5Q50G), which supports:
  - Quad-Port 25G mLOM.
  - Occupies the compute node's modular LAN on motherboard (mLOM) slot.
  - Enables up to 50 Gbps of unified fabric connectivity to each of the chassis intelligent fabric modules (IFMs) for 100 Gbps connectivity per compute node.
  - This card supports Secure Boot.

- Cisco UCS VIC 15230 mLOM (UCSX-ML-V5D200GV2), which supports:
  - x16 PCIE Gen 4 host interface to UCS X215c M8 compute node
  - 4GB DDR4 DIMM, 3200MHz with ECC
  - Two or four KR interfaces that connect to Cisco UCS X Series Intelligent Fabric Modules (IFMs):
    - Two 100G KR interfaces connecting to the UCSX 100G Intelligent Fabric Module (UCSX-I-9108-100G)
    - Four 25G KR interfaces connecting to the Cisco UCSX 9108 25G Intelligent Fabric Module (UCSX-I-9108-25G)
  - This card supports Secure Boot

System Health States

The compute node's front panel has a System Health LED, which is a visual indicator that shows whether the compute node is operating in a normal runtime state (the LED glows steady green). If the System Health LED shows anything other than solid green, the compute node is not operating normally, and it requires attention.

The following System Health LED states indicate that the compute node is not operating normally.

System Health LED Color	Compute Node State	Conditions
Solid Amber	Degraded	• Power supply redundancy lost • Intelligent Fabric Module (IFM) redundancy lost • Mismatched processors in the system. This condition might prevent the system from booting. • Faulty processor in a dual processor system. This condition might prevent the system from booting. • Memory RAS failure if memory is configured for RAS • Failed drive in a compute node configured for RAID
Blinking Amber	Critical	• Boot failure • Fatal processor or bus errors detected • Fatal uncorrectable memory error detected • Lost both IFMs • Lost both drives • Excessive thermal conditions

Interpreting LEDs

Table 1: Compute Node LEDs

LED	Color	Description
Compute Node Power (callout 1 on the Chassis Front Panel) 	Off	Power off.
	Green	Normal operation.
	Amber	Standby.

LED	Color	Description
Compute Node Activity (callout 2 on the Chassis Front Panel) 	Off	None of the network links are up.
	Green	At least one network link is up.
Compute Node Health (callout 3 on the Chassis Front Panel) 	Off	Power off.
	Green	Normal operation.
	Amber	Degraded operation.
	Blinking Amber	Critical error.
Compute Node Locator LED and button (callout 4 on the Chassis Front Panel) 	Off	Locator not enabled.
	Blinking Blue 1 Hz	Locates a selected compute node—If the LED is not blinking, the compute node is not selected. You can initiate the LED through Cisco UCS management software (Cisco Intersight or Cisco UCS Manager) or by pressing the button, which toggles the LED on and off.

Table 2: Drive LEDs, SAS/SATA

Activity/Presence LED 	Status/Fault LED 	Description
Off	Off	Drive not present or drive powered off
On (glowing solid green)	Off	Drive present, but no activity or drive is a hot spare
Blinking green, 4HZ	Off	Drive present and drive activity
Blinking green, 4HZ	Blinking amber, 4HZ	Drive Locate indicator or drive prepared for physical removal
On (glowing solid green)	On (glowing solid amber)	Failed or faulty drive
Blinking green, 1HZ	Blinking amber, 1HZ	Drive rebuild or copyback operation in progress
On (glowing solid green)	Two 4HZ amber blinks with a ½ second pause	Predict Failure Analysis (PFA)

Table 3: Drive LEDs, NVMe (VMD Disabled)

Activity/Presence LED	Status/Fault LED	Description
Off	Off	Drive not present or drive powered off
On (glowing solid green)	Off	Drive present, but no activity
Blinking green, 4HZ	Off	Drive present and drive activity
N/A	N/A	Drive Locate indicator or drive prepared for physical removal
N/A	N/A	Failed or faulty drive
N/A	N/A	Drive Rebuild

Optional Hardware Configuration

The Cisco UCS X215c M8 compute node can be installed in a Cisco UCS X9508 Server Chassis either as a standalone compute node or with the following optional hardware configuration.

Cisco UCS X440p PCIe Node

As an option, the compute node can be paired with a full-slot GPU acceleration hardware module in the Cisco UCS X9508 Server Chassis. This option is supported through the Cisco X440p PCIe node. For information about this option, see the Cisco UCS X440p PCIe Node Installation and Service Guide.

Note When the compute node is paired with the Cisco UCS X440p PCIe node, the Cisco UCS PCI Mezz card for X-Fabric Connectivity (UCSX-V5-BRIDGE-D) is required. This bridge card installs on the compute node.