Spec Sheet

Cisco Compute Hyperconverged with Nutanix

HCIX210c M7 All-NVMe Node

A printed version of this document is only a copy and not necessarily the latest version. Refer to the following link for the latest released version:

https://www.cisco.com/c/en/us/products/hyperconverged-infrastructure/compute-hyperconverged/datasheet-listing.html

CISCO SYSTEMS
170 WEST TASMAN DR.
SAN JOSE, CA, 95134
WWW.CISCO.COM

PUBLICATION HISTORY

REV A.07 SEPTEMBER 16, 2024

OVERVIEW

The Cisco Compute Hyperconverged X-Series Modular System simplifies your data center, adapting to the unpredictable needs of modern applications while also providing for traditional scale-out and enterprise workloads. It reduces the number of server types to maintain, helping to improve operational efficiency and agility as it helps reduce complexity. Powered by the Cisco Intersight™ cloud operations platform, it shifts your thinking from administrative details to business outcomes with hybrid cloud infrastructure that is assembled from the cloud, shaped to your workloads, and continuously optimized.

The Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node is integrate into the Cisco Compute Hyperconverged X-Series Modular System. Up to eight Cisco Compute Hyperconverged Nodes can reside in the 7-Rack-Unit (7RU) Cisco Compute Hyperconverged 9508 Chassis, offering one of the highest densities of compute, IO, and storage per rack unit in the industry.

The Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node harnesses the power of the Intel® Xeon® Scalable Processors, and offers the following:

- **CPU**:
  - 5th Generation Intel® Xeon® Scalable Processors with up to 64 cores per processor **or**
  - 4th Generation Intel® Xeon® Scalable Processors with up to 60 cores per processor.
- **Memory**:
  - Up to 8TB with 32 x 256GB DDR5-5600 DIMMs, in a 2-socket configuration with 5th Gen. Intel® Xeon® Scalable Processors **or**
  - Up to 8TB with 32 x 256GB DDR5-4800 DIMMs, in a 2-socket configuration with 4th Gen. Intel® Xeon® Scalable Processors.
- **Storage**: Up to 6 Non-Volatile Memory Express (NVMe) 2.5-inch drives with a choice of pass-through controllers with four lanes each of PCIe Gen 4 connectivity and up to 2 M.2 SATA drives for flexible boot and local storage capabilities.
- **mLOM virtual interface cards:**
  - Cisco Virtual Interface Card (VIC) 15420 occupies the server's Modular LAN on Motherboard (mLOM) slot, enabling up to 50Gbps (2 x25Gbps) of unified fabric connectivity to each of the chassis Intelligent Fabric Modules (IFMs) for 100Gbps connectivity per server.
  - Cisco Virtual Interface Card (VIC) 15230 occupies the server's modular LAN on motherboard (mLOM) slot, enabling up to 100 Gbps of unified fabric connectivity to each of the chassis Intelligent Fabric Modules (IFMs) for 100 Gbps connectivity per server with secure boot capability.
- **Optional Mezzanine card:**
  - Cisco Virtual Interface Card (VIC) 15422 can occupy the server's mezzanine slot at the bottom rear of the chassis. An included bridge card extends this VIC's 100Gbps (4 x 25Gbps) of network connections through IFM connectors, bringing the total bandwidth to 100Gbps per VIC 15420 and 15422 (for a total of 200Gbps per server). In addition to IFM connectivity, the VIC 15422 I/O connectors link to Cisco X-Fabric technology.
  - Cisco PCI Mezz card for X-Fabric can occupy the server's mezzanine slot at the bottom rear of the chassis. This card's I/O connectors link to Cisco X-Fabric modules and enable connectivity to the X440p PCIe Node.

- **Security**: Includes secure boot silicon root of trust FPGA, ACT2 anti-counterfeit provisions, and optional Trusted Platform Model (TPM).

Figure 1 on page 5 shows a front view of the Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node.

Figure 1 Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node

Front View with Drives

DETAILED VIEWS

Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node Front View

Figure 2 is a front view of the Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node.

Figure 2 Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node Front View (Drives option)

Storage Drives Option

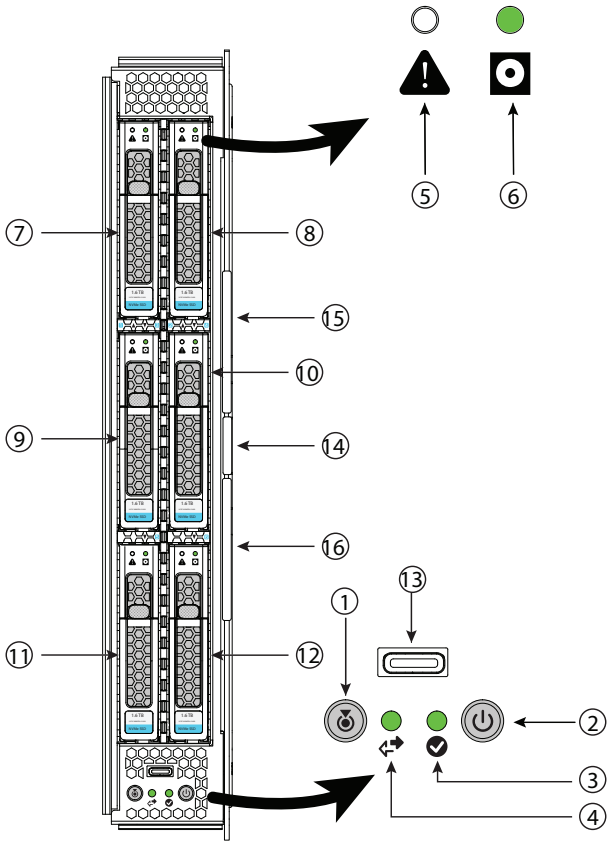

1	Locate button/LED	**9**	Drive Bay 3 (shown populated)
2	Power button/LED	**10**	Drive Bay 4 (shown populated)
3	Status LED	**11**	Drive Bay 5 (shown populated)
4	Network activity LED	**12**	Drive Bay 6 (shown populated)
5	Warning LED (one per drive)	**13**	OCuLink console port[1]
6	Disk drive activity LED (one per drive)	**14**	Ejector handle retention button
7	Drive Bay 1 (shown populated)	**15**	Upper ejector handle
8	Drive Bay 2 (shown populated)	**16**	Lower ejector handle

Notes:

1. An adapter cable (PID HCIX-C-DEBUGCBL) is required to connect the OCuLink port to the transition serial USB and video (SUV) octopus cable.

CISCO COMPUTE HYPERCONVERGED NODE STANDARD CAPABILITIES and FEATURES

Table 1 lists the capabilities and features of the base Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node. Details about how to configure the Cisco Compute Hyperconverged Node for a listed feature or capability (for example, number of processors, disk drives, or amount of memory) are provided in CONFIGURING the CISCO COMPUTE HYPERCONVERGED HCIX210C M7 ALL-NVME NODE on page 8.

Table 1 Capabilities and Features

Capability/Feature	Description
Chassis	The Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node mounts in a Cisco Compute Hyperconverged 9508 chassis.
CPU	■ 5th Gen. Intel® Xeon® Scalable Processors **or** 4th Gen. Intel® Xeon® Scalable Processors. ■ Each CPU has 8 channels with up to 2 DIMMs per channel, for up to 16 DIMMs per CPU. ■ UPI Links: Up to 4 at 20GT/s
Chipset	Intel® C741 series chipset
Memory	■ 32 total DDR5-5600 MT/s DIMM slots with 5th Gen. Intel® Xeon® Scalable Processors (16 per CPU) **or** 32 total DDR5-4800 MT/s DIMM slots with 4th Gen. Intel® Xeon® Scalable Processors ■ Up to 8TB DDR5-5600 DIMM memory capacity (32x 256GB DIMMs) with 5th Gen. Intel® Xeon® Scalable Processors **or** Up to 8TB DDR5-4800 DIMM memory capacity (32x 256GB DIMMs) with 4th Gen. Intel® Xeon® Scalable Processors ■ 75% peak bandwidth increase over DDR4-3200, with on-die ECC; all densities are Registered DIMMs (RDIMMs)
Storage	Up to 6 Non-Volatile Memory Express (NVMe) 2.5-inch drives with a choice of pass-through controllers with four lanes each of PCIe Gen 4 connectivity and up to 2 M.2 SATA or NVMe drives for flexible boot and local storage capabilities.
Additional Storage	Dual 80 mm SATA 3.0 M.2 cards (up to 480GB per card) on a boot-optimized hardware RAID controller
Mezzanine Adapters (Front)	One front mezzanine connector that supports: ■ Up to 6 x 2.5-inch NVMe PCIe drives  **Note:** Drives pass-through controller in the front mezzanine module slot.
Mezzanine Adapter (Rear)	■ An optional Cisco Virtual Interface Card 15422 can occupy the server's mezzanine slot at the bottom of the chassis. A bridge card extends this VIC's 2x 50Gbps of network connections up to the mLOM slot and out through the mLOM's IFM connectors, bringing the total bandwidth to 100Gbps per fabric—a total of 200Gbps per server. ■ An optional PCIe Mezz card for X-Fabric is also supported in the server's mezzanine slot. This card's I/O connectors link to the Cisco X-Fabric modules for Hyperconverged X-Series Gen4 PCIe node access.

Table 1 Capabilities and Features *(continued)*

Capability/Feature	Description
mLOM	The modular LAN on motherboard (mLOM) cards (the Cisco VIC 15230 and 15420) is located at the rear of the Cisco Compute Hyperconverged Node. ■ The Cisco Virtual Interface Card VIC 15420 is a Cisco designed PCI Express (PCIe) based card that supports two 2x25G-KR network interfaces to provide Ethernet communication to the network by means of the Intelligent Fabric Modules (IFMs) in the Cisco Compute Hyperconverged 9508 chassis. The Cisco VIC 15420 mLOM can connect to the rear mezzanine adapter card with a bridge connector. ■ The Cisco Virtual Interface Card (VIC) 15230 occupies the server's modular LAN on motherboard (mLOM) slot, enabling up to 100 Gbps of unified fabric connectivity to each of the chassis Intelligent Fabric Modules (IFMs) for 100 Gbps connectivity per server with secure boot capability.
Video	Video uses a Matrox G200e video/graphics controller. ■ Integrated 2D graphics core with hardware acceleration ■ DDR4 memory interface supports up to 512 MB of addressable memory (16 MB is allocated by default to video memory) ■ Supports display resolutions up to 1920 x 1200 32 bpp@ 60Hz ■ Video is available with an Oculink connector on the front panel. An adapter cable (PID HCIX-C-DEBUGCBL) is required to connect the OCuLink port to the transition serial USB and video (SUV) octopus cable.
Front Panel Interfaces	OCuLink console port. Note that an adapter cable is required to connect the OCuLink port to the transition serial USB and video (SUV) octopus cable.
Power subsystem	Power is supplied from the Cisco Compute Hyperconverged 9508 chassis power supplies. The Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node consumes a maximum of 1300 W.
Fans	Integrated in the Cisco Compute Hyperconverged 9508 chassis.
Integrated management processor	The built-in Cisco Integrated Management Controller enables monitoring of Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node inventory, health, and system event logs.
Baseboard Management Controller (BMC)	ASPEED Pilot IV
ACPI	Advanced Configuration and Power Interface (ACPI) 6.5 Standard Supported. ACPI states S0 and S5 are supported. There is no support for states S1 through S4.
Front Indicators	■ Power button and indicator ■ System activity indicator ■ Location button and indicator
Management	■ **Cisco Intersight software** (SaaS, Virtual Appliance and Private Virtual Appliance)
Fabric Interconnect	Compatible with the Cisco HCIX 6454, 64108 and 6536 fabric interconnects

CONFIGURING the CISCO COMPUTE HYPERCONVERGED HCIX210C M7 ALL-NVME NODE

Follow these steps to configure the Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node:

STEP 1 CHOOSE BASE CISCO COMPUTE HYPERCONVERGED NODE SKU

Top Level ordering product ID (PID) of the Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node as shown in*Table 2*.

Table 2 Top level ordering PID

Product ID (PID)	Description
HCIX-M7-MLB	Cisco Compute Hyperconverged X-Series M7 with Nutanix MLB

Select the product ID (PID) of the Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node as shown in *Table 3*.

Table 3 PID of the Base Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node

Product ID (PID)	Description
HCIX210C-M7SN	210cM7 All NVMe Hyperconverged Node w/o CPU,Memory,Storage
HCIX210C-M7SN-U	210cM7 All NVMe Hyperconverged Node w/o CPU,Memory,Storage

NOTE:

- A base Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node ordered in *Table 3* does not include any components or options. They must be selected during product ordering.
- Hardware orderability will ship Bare Metal Cisco Compute Hyperconverged X-series servers (HCI SKUs) in anticipation of the IMM Software support (see below). Customers will be able to place orders, ship, and stage equipment but will not be able to deploy the Nutanix cluster until IMM & Prism Central Software support for X-series is available.
  - Day 0 - Cluster Deployment with Foundation Central (FC)
  - Day 2 - Cluster Expansion and Integrated Firmware upgrades using LCM is targeted for Q4CY'24

Please follow the steps on the following pages to order components such as the following, which are required in a functional Cisco Compute Hyperconverged Node:

- CPUs
- Memory
- Storage controller
- Drives
- Cisco adapters

STEP 2 CHOOSE CPU(S)

The standard CPU features are:

- The 5th Gen Intel® Xeon® Scalable Processors are paired with Intel® C741 series chipset:
  - Up to 64 cores
  - Cache size of up to 320 MB
  - Power: Up to 350Watts
  - UPI Links: Up to 4 at 20GT/s

With 5th Gen Intel® Xeon® Processors, improve performance and efficiency while reducing costs. See *5TH GEN INTEL XEON BENEFIT PILLARS*

- The 4th Gen Intel® Xeon® Scalable Processors are paired with Intel® C741 series chipset:
  - Up to 60 cores
  - Cache size of up to 112.50 MB
  - Power: Up to 350Watts
  - UPI Links: Up to 4 at 16GT/s

With 4th Gen Intel® Xeon® processors, improve performance efficiency for critical workloads with the most built-in accelerators. See 4th gen intel xeon benefit pillars in *LEADERSHIP PERFORMANCE WITH 4TH GEN INTEL® XEON® PROCESSORS*.

Select CPUs

- The available 5th Gen Intel® Xeon® Scalable CPUs are listed in *Table 4*. See *Table 6 on page 14* for CPU suffix notations.
- The available 4th Gen Intel® Xeon® Scalable CPUs are listed in *Table 5*. See *Table 6 on page 14* for CPU suffix notations.

Table 4 Available 5th Gen. Intel® Xeon® Scalable CPUs

Product ID (PID)	Segment/Workload	Maximum Socket (S)	Cores (C)	Clock Freq (GHz)	Power (W)	Cache Size (MB)	Highest DDR5 DIMM Clock Support (MT/s)
8000 Series Processors							
HCIX-CPU-I8592V	Cloud/SaaS/IaaS	2S	64	2.00	330	320.00	4800
HCIX-CPU-I8592+	2S Performance	2S	64	1.90	350	320.00	5600
HCIX-CPU-I8580	2S Performance	2S	60	2.00	350	300.00	5600
HCIX-CPU-I8570	2S Performance	2S	56	2.10	350	300.00	5600
HCIX-CPU-I8568Y+	2S Performance	2S	48	2.30	350	300.00	5600
HCIX-CPU-I8562Y+	2S Performance	2S	32	2.80	300	60.00	5600
HCIX-CPU-I8558P	Cloud/SaaS/IaaS	2S	48	2.70	350	260.00	5600

Table 4 Available 5th Gen. Intel® Xeon® Scalable CPUs

Product ID (PID)	Segment/Workload	Maximum Socket (S)	Cores (C)	Clock Freq (GHz)	Power (W)	Cache Size (MB)	Highest DDR5 DIMM Clock Support (MT/s)
HCIX-CPU-I8558	2S Mainline	2S	48	2.10	330	260.00	5200
6000 Series Processors							
HCIX-CPU-I6554S	Storage	2S	36	2.20	270	180.00	5200
HCIX-CPU-I6548Y+	2S Performance	2S	32	2.50	250	60.00	5200
HCIX-CPU-I6548N	Networking	2S	32	2.80	250	60.00	5200
HCIX-CPU-I6544Y	2S Performance	2S	16	3.60	270	45.00	5200
HCIX-CPU-I6542Y	2S Performance	2S	24	2.90	250	60.00	5200
HCIX-CPU-I6538Y+	2S Mainline	2S	32	2.20	225	60.00	5200
HCIX-CPU-I6538N	Networking	2S	32	2.10	205	60.00	5200
HCIX-CPU-I6534	2S Performance	2S	8	3.90	195	22.50	4800
HCIX-CPU-I6530	2S Mainline	2S	32	2.10	270	160.00	4800
HCIX-CPU-I6526Y	2S Performance	2S	16	2.80	195	37.50	5200
5000 Series Processors							
HCIX-CPU-I5520+	2S Mainline	2S	28	2.20	205	52.50	4800
HCIX-CPU-I5515+	2S Performance	2S	8	3.20	165	22.50	4800
4000 Series Processors							
HCIX-CPU-I4516Y+	2S Mainline	2S	24	2.20	185	45.00	4400
HCIX-CPU-I4514Y	2S Mainline	2S	16	2.00	150	30.00	4400
HCIX-CPU-I4510T[1]	EDGE (IOT)	2S	12	2.00	115	30.00	4400
HCIX-CPU-I4510[1]	2S Mainline	2S	12	2.40	150	30.00	4400
HCIX-CPU-I4509Y[1]	2S Mainline	2S	8	2.60	125	22.50	4400

Notes:

1. 48GB and 96GB memory DIMMs not supported on HCIX-CPU-I4509Y, HCIX-CPU-I4510, HCIX-CPU-I4510T

Table 5 Available 4th Gen. Intel® Xeon® Scalable CPUs

Product ID (PID)	Segment/Workload	Maximum Socket (S)	Cores (C)	Clock Freq (GHz)	Power (W)	Cache Size (MB)	Highest DDR5 DIMM Clock Support (MT/s)
8000 Series Processors							
HCIX-CPU-I8490H	IMDB/Analytics	2S	60	1.90	350	112.50	4800
HCIX-CPU-I8480+	2S Performance	2S	56	2.00	350	105.00	4800

Table 5 Available 4th Gen. Intel® Xeon® Scalable CPUs

Product ID (PID)	Segment/Workload	Maximum Socket (S)	Cores (C)	Clock Freq (GHz)	Power (W)	Cache Size (MB)	Highest DDR5 DIMM Clock Support (MT/s)
HCIX-CPU-I8470N	5G/Networking	2S	52	1.70	300	97.50	4800
HCIX-CPU-I8470	2S Performance	2S	52	2.00	350	105.00	4800
HCIX-CPU-I8468V	Cloud/SaaS/Media	2S	48	2.40	330	97.50	4800
HCIX-CPU-I8468H	IMDB/Analytics	2S	48	2.10	330	105.00	4800
HCIX-CPU-I8468	2S Performance	2S	48	2.10	350	105.00	4800
HCIX-CPU-I8462Y+	2S Performance	2S	32	2.80	300	60.00	4800
HCIX-CPU-I8460Y+	2S Performance	2S	40	2.00	300	105.00	4800
HCIX-CPU-I8460H	IMDB/Analytics	2S	40	2.20	330	105.00	4800
HCIX-CPU-I8458P	Cloud/SaaS/Media	2S	44	2.70	350	82.50	4800
HCIX-CPU-I8454H	IMDB/Analytics	2S	32	2.10	270	82.50	4800
HCIX-CPU-I8452Y	2S Mainline	2S	36	2.00	300	67.50	4800
HCIX-CPU-I8450H	IMDB/Analytics	2S	28	2.00	250	75.00	4800
HCIX-CPU-I8444H	IMDB/Analytics	2S	16	2.90	270	45.00	4800
6000 Series Processors							
HCIX-CPU-I6454S	Storage	2S	32	2.20	270	60.00	4800
HCIX-CPU-I6448Y	2S Performance	2S	32	2.10	225	60.00	4800
HCIX-CPU-I6448H	IMDB/Analytics	2S	32	2.40	250	60.00	4800
HCIX-CPU-I6444Y	2S Performance	2S	16	3.60	270	45.00	4800
HCIX-CPU-I6442Y	2S Performance	2S	24	2.60	225	60.00	4800
HCIX-CPU-I6438Y+	2S Mainline	2S	32	2.00	205	60.00	4800
HCIX-CPU-I6438N	5G/Networking	2S	32	2.00	205	60.00	4800
HCIX-CPU-I6438M	Cloud/SaaS/Media	2S	32	2.20	205	60.00	4800
HCIX-CPU-I6434H	IMDB/Analytics	2S	8	3.70	195	22.50	4800
HCIX-CPU-I6434	2S Performance	2S	8	3.70	195	22.50	4800
HCIX-CPU-I6430	2S Mainline	2S	32	2.10	270	60.00	4400
HCIX-CPU-I6428N	5G/Networking	2S	32	1.80	185	60.00	4000
HCIX-CPU-I6426Y	2S Performance	2S	16	2.50	185	37.50	4800
HCIX-CPU-I6418H	IMDB/Analytics	2S	24	2.10	185	60.00	4800
HCIX-CPU-I6416H	IMDB/Analytics	2S	18	2.20	165	45.00	4800
5000 Series Processors							
HCIX-CPU-I5420+	2S Mainline	2S	28	2.00	205	52.50	4400

Table 5 Available 4th Gen. Intel® Xeon® Scalable CPUs

Product ID (PID)	Segment/Workload	Maximum Socket (S)	Cores (C)	Clock Freq (GHz)	Power (W)	Cache Size (MB)	Highest DDR5 DIMM Clock Support (MT/s)
HCIX-CPU-I5418Y	2S Mainline	2S	24	2.00	185	45.00	4400
HCIX-CPU-I5418N	5G/Networking	2S	24	1.80	165	45.00	4000
HCIX-CPU-I5416S	Storage	2S	16	2.00	150	30.00	4400
HCIX-CPU-I5415+	2S Performance	2S	8	2.90	150	22.50	4400
4000 Series Processors							
HCIX-CPU-I4416+	2S Mainline	2S	20	2.00	165	37.50	4000
HCIX-CPU-I4410Y	2S Mainline	2S	12	2.00	150	30.00	4000
HCIX-CPU-I4410T	IOT	2S	10	2.70	150	26.25	4000

Table 6 CPU Suffixes

CPU Suffix	Description	Features
P	Cloud (IaaS)	Designed for cloud IaaS environments to deliver higher frequencies at constrained TDPs
V	Cloud (SaaS)	Designed for high rack density, maximize VM/core, and lower power VM environment
M	Media Transcode	Designed for Media processing, AI, and HPC workloads
H	DB and Analytics	Designed for Data Analytics and Big Data usages
N	Network/5G/Edge (High TDP/Low latency)	Designed and optimized for a range of broadly-deployed network and 5G workload environments from Edge to the Data Center
S	Storage & HCI	Designed for Storage usages and workloads
T	Long-life Use/High Tcase	Designed for Network Environment-Building System (NEBS) and IoT market
U	1-Socket	Optimized for targeted platforms adequately served by the cores, memory bandwidth and IO capacity available from a single processor
Y	General SKU with SST-PP	Designator is used for general SKU stack to highlight SST-PP (Speed Select Technology Performance Profile) feature enabled
+	Feature Plus SKU	Designed to enable 1 instance of each DSA, IAA, QAT, DLB embedded accelerator

Supported Configurations

(1) Configurations with NVMe PCIe drives:

- Select two identical CPUs from *Table 5 on page 11*

(2) Two-CPU Configuration

- Choose two identical CPUs from any one of the rows of *Table 5 on page 11*

STEP 3 CHOOSE MEMORY

The *Table 7* below describes the main memory DIMM features supported on Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node.

CAUTION: When populating 256GB DIMMs, the ambient temperature shall be limited to a maximum of 32^{0}C.

Table 7 HCIX210c M7 Main Memory Features

Memory DIMM server technologies	Description
DDR5 memory clock speed	4th Gen. CPU: Up to 4800MT/s 1DPC; Up to 4400MT/s 2DPC
	5th Gen. CPU: Up to 5600MT/s 1DPC; Up to 4400MT/s 2DPC
Operational voltage	1.1 Volts
DRAM fab density	16Gb and 24Gb
DRAM DIMM type	RDIMM (Registered DDR5 DIMM with on die ECC)
Memory DIMM organization	Eight memory DIMM channels per CPU; up to 2 DIMMs per channel
Maximum number of DRAM DIMM per server	32 (2-Socket)
DRAM DIMM Densities and Ranks	16GB 1Rx8, 32GB 1Rx4, 64GB 2Rx4, 128GB 4Rx4, 256GB 8Rx4
	48GB 1Rx4, 96GB 2Rx4 - 5th Gen. only
Maximum system capacity (DRAM DIMMs only)	8TB (32x256GB)

Figure 3 Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node Memory Organization

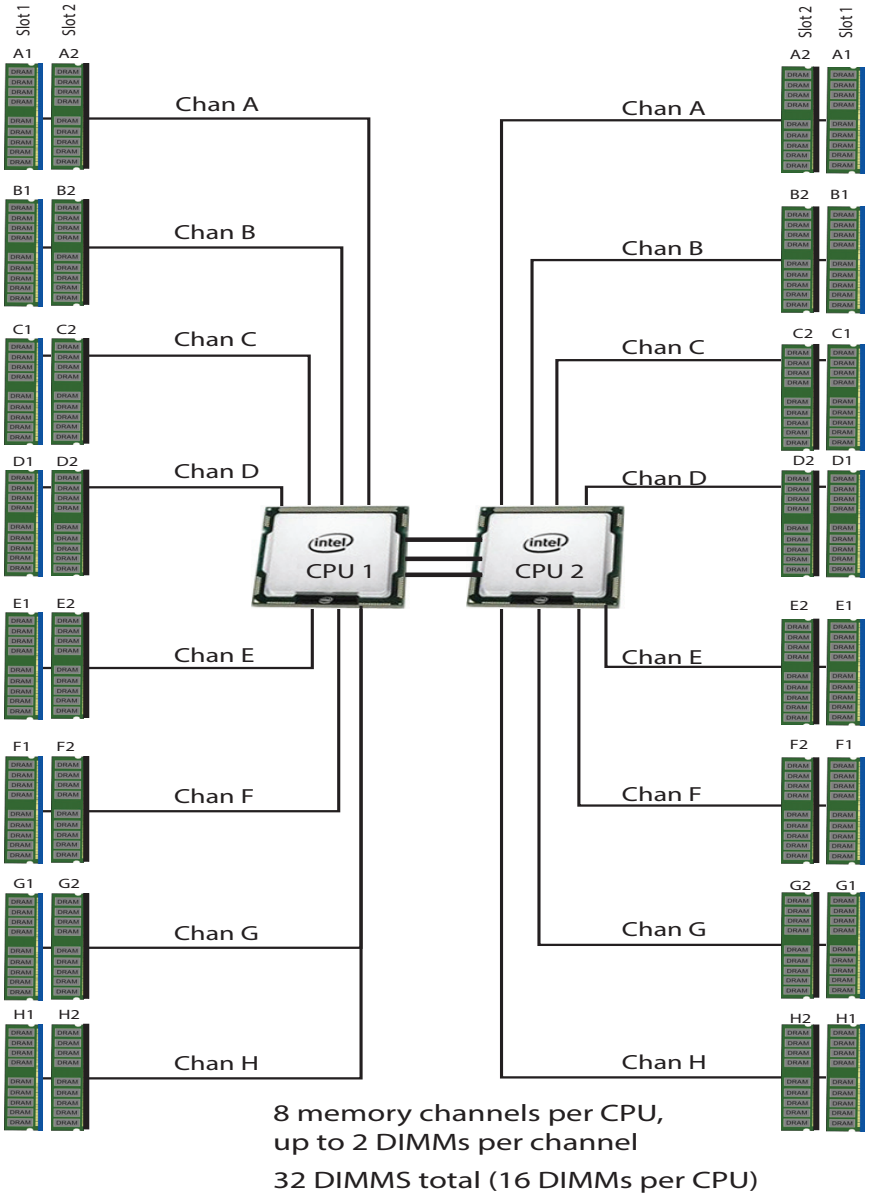

Select DIMMs

The available memory DIMMs option are listed in *Table 8*.

Table 8 Available DDR5 DIMMs

Product ID (PID)	PID Description
DDR5-4800MT/s PID list	
HCIX-MRX16G1RE1	16GB DDR5-4800 RDIMM 1Rx8 (16Gb)
HCIX-MRX32G1RE1	32GB DDR5-4800 RDIMM 1Rx4 (16Gb)
HCIX-MRX64G2RE1	64GB DDR5-4800 RDIMM 2Rx4 (16Gb)
HCIX-MR128G4RE1	128GB DDR5-4800 RDIMM 4Rx4 (16Gb)
HCIX-MR256G8RE1[1]	256GB DDR5-4800 RDIMM 8Rx4 (16Gb)
DDR5-5600MT/s PID list[2]	
HCIX-MRX16G1RE3	16GB DDR5-5600 RDIMM 1Rx8 (16Gb)
HCIX-MRX32G1RE3	32GB DDR5-5600 RDIMM 1Rx4 (16Gb)
HCIX-MRX48G1RF3[3]	48GB DDR5-5600 RDIMM 1Rx4 (24Gb)
HCIX-MRX64G2RE3	64GB DDR5-5600 RDIMM 2Rx4 (16Gb)
HCIX-MRX96G2RF3[3]	96GB DDR5-5600 RDIMM 2Rx4 (24Gb)
HCIX-MR128G4RE3	128GB DDR5-5600 RDIMM 4Rx4 (16Gb)
HCIX-MR256G8RE3[1]	256GB DDR5-5600 RDIMM 8Rx4 (16Gb)
Accessories/spare included with Memory configuration: ■ UCS-DDR5-BLK[4] is auto included for the unselected DIMMs slots	

Notes:

1. If selecting 256GB DIMMs the ambient temperature shall be limited to a maximum of 32ºC.
2. DDR5-5600 supported on Intel® Xeon® 5th generation only.
3. 48GB and 96GB memory DIMMs not supported on HCIX-CPU-I3508U, HCIX-CPU-I4509Y, HCIX-CPU-I4510, HCIX-CPU-I4510T
4. Any empty DIMM slot must be populated with a DIMM blank to maintain proper cooling airflow.

Memory configurations and mixing rules

- **Golden Rule:** Memory on every CPU socket shall be configured identically.
- System speed is dependent on the CPU DIMM speed support. Refer to Available 4th Gen. Intel® Xeon® Scalable CPUs on page 11 and Available 5th Gen. Intel® Xeon® Scalable CPUs on page 10 for DIMM speeds.
- For full details on supported memory configurations see the **M7 Memory Guide.**
- DIMM Count Rules:

Table 9 Allowed DIMM Count for 1-CPU, 2-CPU

Allowed DIMM Count rules	Minimum Count	Maximum Count	Allowed Count	Not Allowed Count
16GB, 32GB, 64GB, 128GB, 256GB (4th and 5th Gen. CPUs)				
DIMM count for 1 CPU	1	16	1,2,4,6,8,12,16	3,5,7,9,10,11,13,14,15
DIMM count for 2-CPU	2	32	2,4,8,12,16,24,32	6,10,14,18,20,22,26,28,30
48GB (5th Gen. CPUs Only)				
DIMM count for 1 CPU	1	8	1,6,8	2,3,4,5,7,9,10,11,12,13,14,15, 16
DIMM count for 2-CPU	2	16	2,12,16	4,6,8,10,14,18,20,22,24,26,28, 30,32
96GB (5th Gen. CPUs Only)				
DIMM count for 1 CPU	1	16	1,6,8,12,16	2,3,4,5,7,9, 10,11,13,14,15
DIMM count for 2-CPU	2	32	2,12,16,24,32	4,6,8,10,14,18,20,22,26,28,30

NOTE(1): 12 DIMMs count for 1-CPU, 24 DIMMs count and for 2-CPU configurations are only allowed when all DIMMs have the same density.

- DIMM Population Rules:
  - Each channel has two memory slots (for example, channel A = slots A1 and A2). See **golden rule** above.
    - A channel can operate with one or two DIMMs installed.
    - If a channel has only one DIMM, populate slot 1 first (the blue slot).
  - When both CPUs are installed, populate the memory slots of each CPU identically. Fill the blue slots (slot 1) in the memory channels first according to the recommended DIMM populations in ***Table 10*.0**, ***Table 10*.1** and ***Table 10*.2**.

Table 10.0 M7 DIMM population order for 16GB, 32GB, 64GB, 128GB, 256GB

#DIMMs per CPU	DIMM Population - 16GB, 32GB, 64GB, 128GB, 256GB (4th and 5th Gen. CPUs)[1]	
	Slot 1 (Blue)	Slot 2 (Black)
1	A1	-
2	A1, G1	-
4	A1, C1, E1, G1	-
6	A1, C1, D1, E1, F1, G1	-
8	A1, B1, C1, D1, E1, F1, G1, H1	-
12[2]	A1, B1, C1, D1, E1, F1, G1, H1	A2, C2, E2, G2
16	A1, B1, C1, D1, E1, F1, G1, H1	A2, B2, C2, D2, E2, F2, G2, H2

Notes:
1. See DIMM Mixing Rules for allowed combinations across slots 1 and 2.
2. Only valid when DIMMs in blue and black slots are the same density.

Table 10.1 M7 DIMM population order for 48GB

#DIMMs per CPU	DIMM Population - 48GB (5th Gen. CPUs only)[1,2]	
	Slot 1 (Blue)	Slot 2 (Black)
1	A1	-
6	A1, C1, D1, E1, F1, G1	-
8	A1, B1, C1, D1, E1, F1, G1, H1	-

Notes:
1. 48GB DIMMs cannot be mixed with other DIMM capacities.
2. Only 1 DPC supported by 48GB.

Table 10.2 M7 DIMM population order for 96GB

#DIMMs per CPU	DIMM Population - 96GB (5th Gen. CPUs only)[1]	
	Slot 1 (Blue)	Slot 2 (Black)
1	A1	-
6	A1, C1, D1, E1, F1, G1	-
8	A1, B1, C1, D1, E1, F1, G1, H1	-
12[2]	A1, B1, C1, D1, E1, F1, G1, H1	A2, C2, E2, G2
16[2]	A1, B1, C1, D1, E1, F1, G1, H1	A2, B2, C2, D2, E2, F2, G2, H2

Notes:

1. 96GB DIMMs cannot be mixed with other DIMM capacities.
2. Only valid when DIMMs in blue and black slots are the same density.

- DIMM Mixing Rules:
  - Higher rank DIMMs shall be populated on Slot 1.
  - Mixing different DIMM densities in the same slot across channels is not supported. All populated slots of the same color must have the same DIMM density.
  - Mixing X4 and X8 DIMMs is not allowed
  - Mixing 16Gb DRAM based and 24Gb DRAM based DIMMs is not allowed. Therefore, 48GB and 96GB cannot be mixed with any other memory DIMMs
  - 48GB supports 1 DIMM Per Channel (1DPC) only
  - The DIMM mixing rules matrix is described in the ***Table 11*** and ***Table 12***, below

Table 11 Supported DIMM mixing and population across 2 slots in each channel - 16GB, 32GB, 64GB, 128GB, and 256GB DIMMs

Channel Mixing		DIMM Slot 2 (Black)				
DIMM Slot 1 (Blue)		16GB	32GB	64GB	128GB	256GB
		1Rx8	1Rx4	2Rx4	4Rx4	8Rx4
16GB	1Rx8	Yes[1]	No	No	No	No
32GB	1Rx4	No	Yes[1]	No	No	No
64GB	2Rx4	No	Yes[2]	Yes[1]	No	No
128GB	4Rx4	No	No	No	Yes[1]	No
256GB	8Rx4	No	No	No	Yes[2]	Yes[1]

Notes:

1. For 2,4,6, 8 DIMMs count, only populate slot 1 (Blue slot). see ***Table 10.0*** for details
2. When mixing two different DIMM densities, all 8 channels per CPU must be populated. Use of fewer than 8 channels (16 slots per CPU) is not supported.

Table 12 Supported DIMM mixing and population across 2 slots in each channel - 48GB and 96GB DIMMs

Channel Mixing		DIMM Slot 2 (Black)	
DIMM Slot 1 (Blue)		48GB	96GB
		1Rx4	2Rx4
48GB	1Rx4	No	No
96GB	2Rx4	No	Yes[1]

Notes:

1. For 6, 8 DIMMs count, only populate slot 1 (Blue slot). see ***Table 10.2*** for details

- Memory Limitations:
  - Memory on every CPU socket shall be configured identically.
  - Refer to *Table 10*.0, *Table 10*.1 and *Table 10*.2. and *Table 11* and *Table 12*, for DIMM population and DIMM mixing rules.
  - Cisco memory from previous generation servers (DDR3 and DDR4) is not supported with the M7 servers.
- For best performance, observe the following:
  - For optimum performance, populate at least one DIMM per memory channel per CPU. When one DIMM per channel is used, it must be populated in DIMM slot 1 (blue slot farthest away from the CPU) of a given channel.
  - The maximum 2 DPC speed is 4400 MT/s, refer to *Table 13*.0 and *Table 13*.1 for the details.

Table 13.0 DDR5-4800 DIMM 1DPC and 2DPC max speed matrix - 4th Gen. CPU

4th Gen. CPU Shelves and Memory Speed	1DPC	2DPC
	All RDIMMs	All RDIMMs
Platinum Series 8	4800 MT/s	4400 MT/s
Gold Series 6	4800 MT/s	4400 MT/s
Gold Series 5	4400 MT/s	4400 MT/s
Silver Series 4	4000 MT/s	4000 MT/s
Bronze Series 3	4000 MT/s	4000 MT/s

Table 13.1 DDR5-5600 DIMM 1DPC and 2DPC max speed matrix - 5th Gen. CPU

5th Gen. CPU Shelves and Memory Speed	1DPC	2DPC
	All RDIMMs	All RDIMMs
Platinum Series 8	5600 MT/s	4400 MT/s
Gold Series 6	5200 MT/s	4400 MT/s
Gold Series 5	4800 MT/s	4400 MT/s
Silver Series 4	4400 MT/s	4400 MT/s
Bronze Series 3	4400 MT/s	4400 MT/s

NOTE: For full details on supported memory configurations see the M7 Memory Guide.

STEP 4 CHOOSE REAR mLOM ADAPTER

The Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node must be ordered with a Cisco VIC mLOM Adapter. The adapter is located at the back and can operate in a single-CPU or dual-CPU configuration. *Table 14* shows the mLOM adapter choices.

Table 14 mLOM Adapters

Product ID (PID)	Description	Connection type
HCIX-MLV5D200GV2	Cisco VIC 15230 modular LOM w/Secure Boot X Cisco Compute Hyperconverged Node	mLOM
HCIX-ML-V5Q50G	Cisco VIC 15420 4x25G secure boot mLOM for X Cisco Compute Hyperconverged Node	mLOM

NOTE:

- VIC 15420 are supported with both X9108-IFM-25G and X9108-IFM-100G. VIC 15420 will operate at 4x 25G with both X9108-IFM-25G and X9108-IFM-100G. While, VIC 15230 will operate at 4x 25G with X9108-IFM-25G and at 2x 100G with X9108-IFM-100G.
- The mLOM adapter is mandatory for the Ethernet connectivity to the network by means of the IFMs and has x16 PCIe Gen4 connectivity with Cisco VIC 15420 or x16 Gen4 connectivity with Cisco VIC 15230 towards the CPU1.
- There is no backplane in the Cisco Compute Hyperconverged 9508 chassis; thus, the Cisco Compute Hyperconverged Nodes directly connect to the IFMs using Orthogonal Direct connectors.
- *Figure 4* shows the location of the mLOM and rear mezzanine adapters on the Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node. The bridge adapter connects the mLOM adapter to the rear mezzanine adapter.

Figure 4 Location of mLOM and Rear Mezzanine Adapters

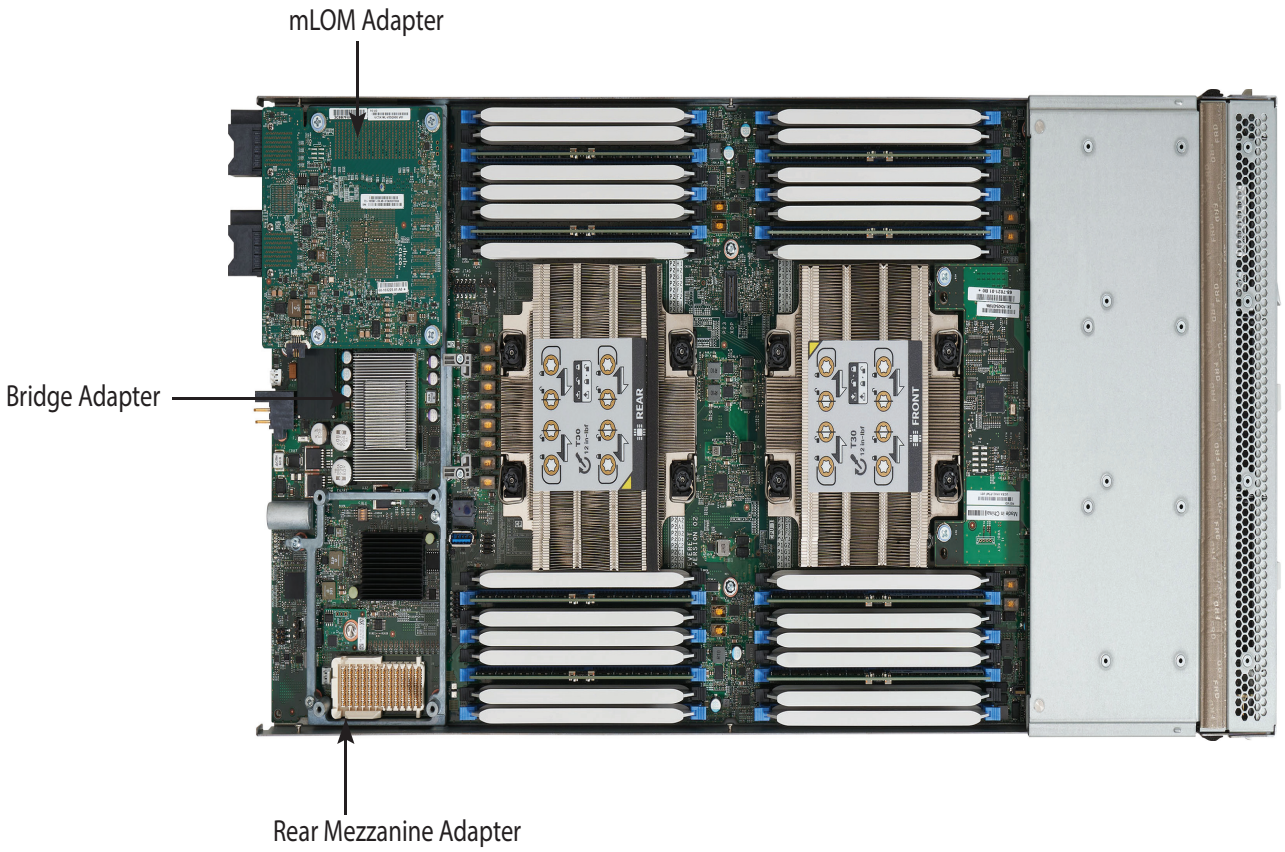

Figure 5 shows the network connectivity from the mLOM out to the 25G IFMs.

Figure 5 Network Connectivity 25G IFMs

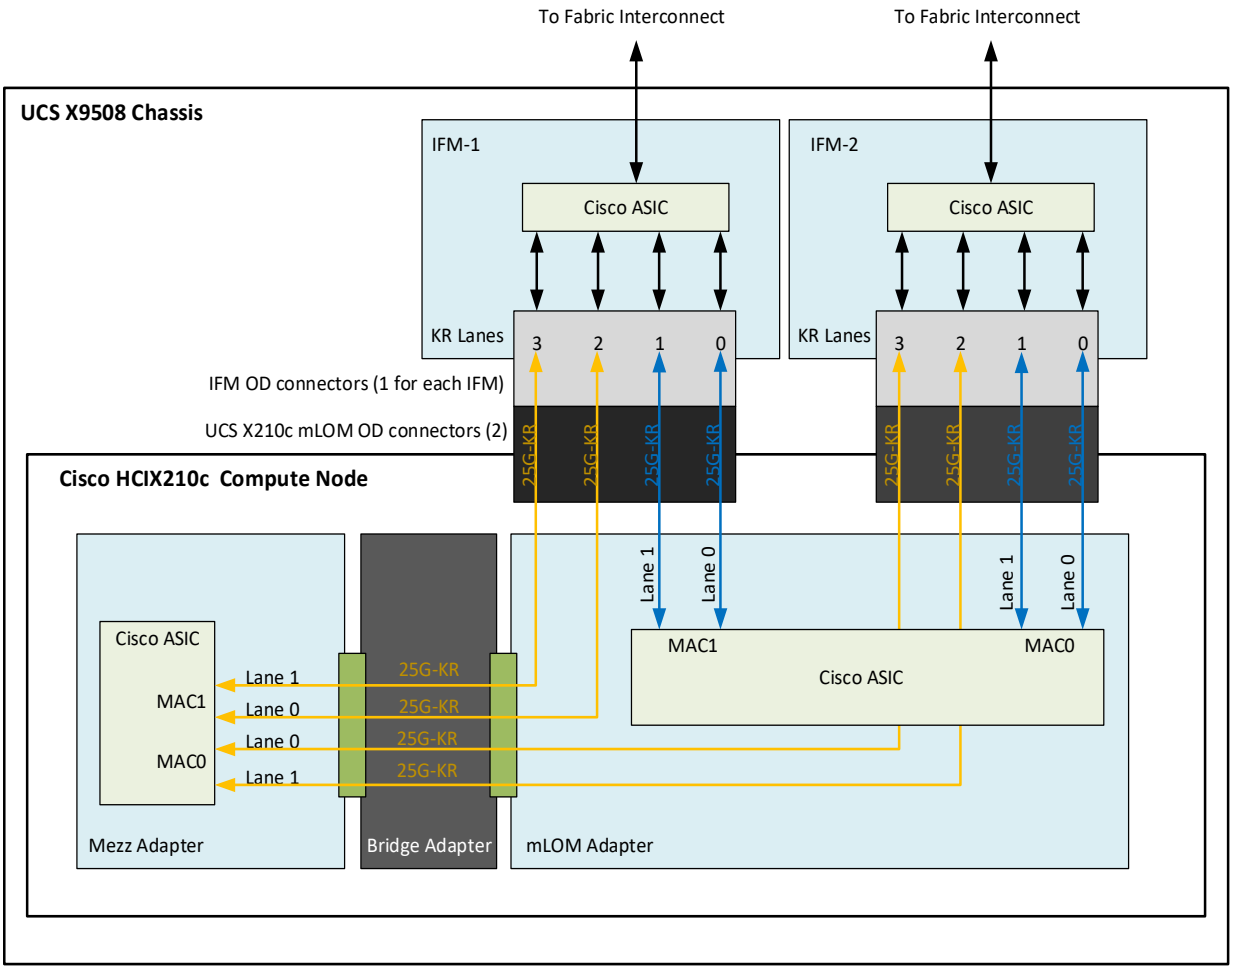

Figure 6 shows the network connectivity from the mLOM out to the 100G IFMs.

Figure 6 Network Connectivity 100G IFMs

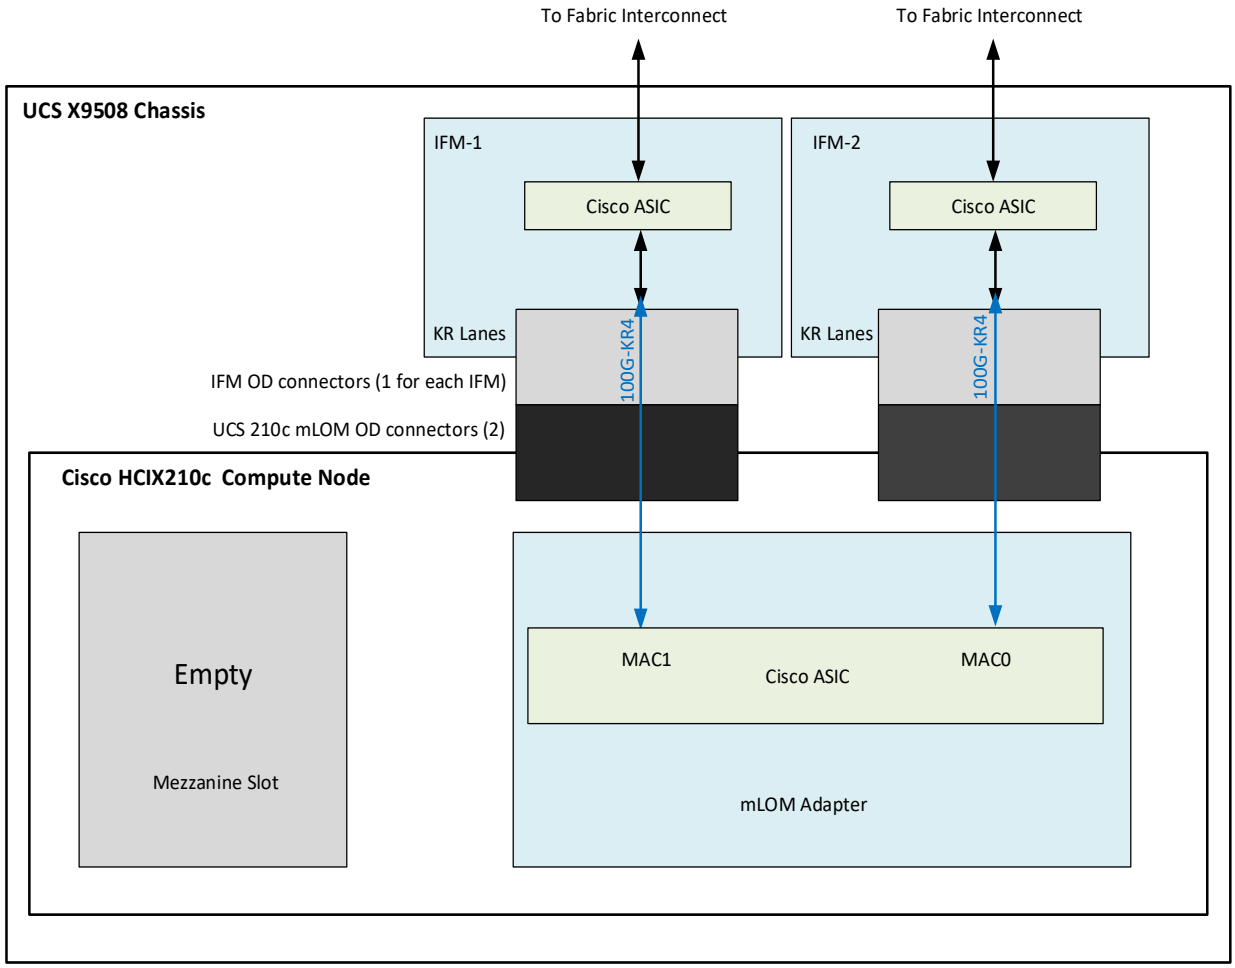

STEP 5 CHOOSE OPTIONAL REAR MEZZANINE VIC/BRIDGE ADAPTERS

The Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node has one rear mezzanine adapter connector which can have a VIC 15422 Mezz card that can be used as a second VIC card on the Cisco Compute Hyperconverged Node for network connectivity or as a connector to the X440p PCIe node via X-Fabric modules. The same mezzanine slot on the Cisco Compute Hyperconverged Node can also accommodate a pass-through mezzanine adapter for X-Fabric which enables Cisco Compute Hyperconverged Node connectivity to the X440p PCIE node. Refer to *Table 15* for supported adapters.

Table 15 Available Rear Mezzanine Adapters

Product ID (PID)	PID Description	CPUs Required	Connector Type
Cisco VIC Card			
HCIX-V4-PCIME[1]	PCI Mezz Card for X-Fabric	2 CPUs required	Rear Mezzanine connector on motherboard
HCIX-ME-V5Q50G	VIC 15422 4x25G secure boot mezz for X Cisco Compute Hyperconverged Node	2 CPUs required	Rear Mezzanine connector on motherboard
Cisco VIC Bridge Card[2]			
HCIX-V5-BRIDGE	VIC 15000 bridge to connect mLOM and mezz X Cisco Compute Hyperconverged Node (This bridge to connect the Cisco VIC 15420 mLOM and Cisco VIC 15422 Mezz for the HCIX210c M7 All-NVMe Node)	2 CPUs required	One connector on Mezz card and one connector on mLOM card

Notes:

1. If this adapter is selected, then two CPUs are required and **HCIX-ME-V5Q50G** or **HCIX-V4-PCIME** is required.
2. Included with the Cisco VIC 15422 mezzanine adapter.

NOTE: The **HCIX-V4-PCIME** rear mezzanine card for X-Fabric has PCIe Gen4 x16 connectivity towards each CPU1 and CPU2. Additionally, the **HCIX-V4-PCIME** also provides two PCIe Gen4 x16 to each X-fabric. This rear mezzanine card enables connectivity from the HCIX210c M7 All-NVMe Node to the X440p PCIe node.

Table 16 Throughput Per HCIX210c M7 All-NVMe Node

HCIX210c M7 Cisco Compute Hyperconverged Node	FI-6536 + X9108-IFM-100G	FI-6536/6400 + X9108-IFM-25G	FI-6536 + X9108-IFM-25G/100G or FI-6400 + X9108-IFM-25G	FI-6536 + X9108-IFM-25G/100G or FI-6400 + X9108-IFM-25G	
X210c configuration	VIC 15230	VIC 15230	VIC 15420	VIC 15420 + VIC 15422	
Throughput per node	200G (100G per IFM)	100G (50G per IFM)	100G (50G per IFM)	200G (100G per IFM)	
vNICs needed for max BW	2	2	2	4	
KR connectivity from VIC to each IFM	1x 100GKR	2x 25GKR	2x 25GKR	4x 25GKR	
Single vNIC throughput on VIC	100G (1x100GKR)	50G (2x25G KR)	50G (2x25G KR)	50G (2x25G KR)	50G (2x25G KR)
Max Single flow BW per vNIC	100G	25G	25G	25G	25G
Single vHBA throughput on VIC	100G	50G	50G	50G	50G

Supported Configurations

- One of mLOM VIC from *Table 14* is always required.
- If a **HCIX-ME-V5Q50G** rear mezzanine VIC card is installed, a **HCIX-V5-BRIDGE** VIC bridge card is included and connects the mLOM to the mezzanine adapter.
- The **HCIX-ME-V5Q50G** rear mezzanine card has Ethernet connectivity to the IFM using the **HCIX-V5-BRIDGE** and has a PCIE Gen4 x16 connectivity towards CPU2. Additionally, the **HCIX-ME-V5Q50G** also provides two PCIe Gen4 x16 to each X-fabric.
- All the connections to Cisco X-Fabric 1 and Cisco X-Fabric 2 are through the Molex Orthogonal Direct (OD) connector on the mezzanine card.
- The rear mezzanine card has 32 x16PCIe lanes to each Cisco X-Fabric for I/O expansion to enable resource consumption from the PCIe resource nodes.

STEP 6 CHOOSE FRONT MEZZANINE ADAPTER

The Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node has one front mezzanine connector that can accommodate one of the following mezzanine cards:

- Pass-through controller for up to 6 U.2/U.3 NVMe drives.

NOTE: The Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node can be ordered with or without the front mezzanine adapter. Refer to **Table 17 Available Front Mezzanine Adapters**

Table 17 Available Front Mezzanine Adapters

Product ID(PID)	PID Description	Connector Type
HCIX-X10C-PT4F	Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node compute pass through controller for up to 6 NVMe drives	Front Mezzanine

STEP 7 CHOOSE OPTIONAL GPU PCIe NODE

Refer to *Table 18* for GPU PCIe Node

Table 18 GPU PCIe Node

Product ID(PID)	PID Description
HCIX-440P	Cisco Compute Hyperconverged X-Series Gen4 PCIe node

NOTE:

- If HCIX-440P is selected, then rear mezzanine is required.

STEP 8 CHOOSE OPTIONAL GPUs

Select GPU Options

The available PCIe node GPU options are listed in *Table 19*.

Table 19 Available PCIe GPU Cards supported on the PCIe Node

GPU Product ID (PID)	PID Description	Maximum number of GPUs per node
HCIX-GPU-A16	NVIDIA A16 PCIE 250W 4X16GB	2
HCIX-GPU-L4	NVIDIA L4 Tensor Core, 70W, 24GB	4
HCIX-GPU-L40S	NVIDIA L40S: 350W, 48GB, 2-slot FHFL GPU	2
HCIX-GPU-H100-80	NVIDIA H100: 350W, 80GB, 2-slot FHFL GPU	2

STEP 9 CHOOSE DRIVES

The Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node can be ordered with or without drives. The drive options are:

Select Drives - HCIX210C-M7SN (All-NVMe)

The available drives are listed in *Table 20*.

Table 20 Available Capacity Drives

Product ID (PID)	PID Description	Drive Type	Capacity
Capacity Drive			
HCIX-NVME4-1920	1.9TB 2.5in U.2 15mm P5520 Hg Perf Med End NVMe	NVMe	1.9TB
HCIX-NVME4-3840	3.8TB 2.5in U.2 15mm P5520 Hg Perf Med End NVMe	NVMe	3.8TB
HCIX-NVME4-7680	7.6TB 2.5in U.2 15mm P5520 Hg Perf Med End NVMe	NVMe	7.6TB
HCIX-NVME4-15360	15.3TB 2.5in U.2 15mm P5520 Hg Perf Med End NVMe	NVMe	15.3TB
HCIX-NVMEG4-M1920	1.9TB 2.5in U.3 15mm P7450 Hg Perf Med End NVMe	NVMe	1.9TB
HCIX-NVMEG4-M3840	3.8TB 2.5in U.3 15mm P7450 Hg Perf Med End NVMe	NVMe	3.8TB
HCIX-NVMEG4-M7680	7.6TB 2.5in U.3 15mm P7450 Hg Perf Med End NVMe	NVMe	7.6TB
HCIX-NVMEG4-M1536	15.3TB 2.5in U.3 15mm P7450 Hg Perf Med End NVMe	NVMe	15.3TB

Approved Configurations

- Two to six PCIe U.2/U.3 NVMe drives

STEP 10 ORDER M.2 SATA SSDs AND RAID CONTROLLER

- **Cisco 6GB/s SATA Boot-Optimized M.2 RAID Controller (included):** Boot-Optimized RAID controller (HCIX-M2-HWRD-FPS) for hardware RAID across two SATA M.2 storage modules. The Boot-Optimized RAID controller plugs into the motherboard and the M.2 SATA drives plug into the Boot-Optimized RAID controller.

NOTE:

- The HCIX-M2-HWRD-FPS is auto included with the server configuration
- The HCIX-M2-HWRD-FPS controller supports RAID 1 and JBOD mode and is available only with 240GB, 480GB, and 960GB M.2 SATA SSDs.
- Cisco IMM is supported for configuring of volumes and monitoring of the controller and installed SATA M.2 drives
- Hot-plug replacement is not supported. The Cisco Compute Hyperconverged Node must be powered off to replace.
- The Boot-Optimized RAID controller supports VMware, Windows, and Linux Operating Systems

Table 21 Boot-Optimized RAID controller (auto included)

Product ID (PID)	PID Description
HCIX-M2-HWRD-FPS	HCIX Front panel with M.2 RAID controller for SATA drives

- **Select Cisco M.2 SATA SSDs:** Order one or two matching M.2 SATA SSDs. This connector accepts the boot-optimized RAID controller (see *Table 21*). Each boot-optimized RAID controller can accommodate up to two SATA M.2 SSDs shown in *Table 22*.

NOTE:

- Each boot-optimized RAID controller can accommodate up to two SATA M.2 SSDs shown in *Table 22*. The boot-optimized RAID controller plugs into the motherboard.
- It is recommended that M.2 SATA SSDs be used as boot-only devices.
- The SATA M.2 drives can boot in UEFI mode only. Legacy boot mode is not supported.

Table 22 M.2 SATA SSDs

Product ID (PID)	PID Description
HCIX-M2-240G	240GB 2.5in M.2 SATA Micron G2 SSD
HCIX-M2-480G	480GB 2.5in M.2 SATA Micron G2 SSD
HCIX-M2-I240GB	240GB M.2 Boot SATA Intel SSD
HCIX-M2-I480GB	480GB M.2 Boot SATA Intel SSD

STEP 11 CHOOSE OPTIONAL TRUSTED PLATFORM MODULE

Trusted Platform Module (TPM) is a computer chip or microcontroller that can securely store artifacts used to authenticate the platform or Cisco Compute Hyperconverged HCIX210c M7 All-NVMe 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.

Table 23 Available TPM Option

Product ID (PID)	Description
HCIX-TPM-002C	Trusted Platform Module 2.0, FIPS140-2 Compliant, HCIX M7 server
HCIX-TPM-OPT-OUT[1]	OPT OUT, TPM 2.0, TCG, FIPS140-2, CC EAL4+ Certified

Notes:

1. Please note Microsoft certification requires a TPM 2.0 for bare-metal or guest VM deployments. Opt-out of the TPM 2.0 voids the Microsoft certification.

NOTE:

- The TPM module used in this system conforms to TPM v2.0 as defined by the Trusted Computing Group (TCG).
- TPM installation is supported after-factory. However, a TPM installs with a one-way screw and cannot be replaced, upgraded, or moved to another Cisco Compute Hyperconverged Node. If a Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node with a TPM is returned, the replacement Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node must be ordered with a new TPM. If there is no existing TPM in the Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node, you can install a TPM 2.0. Refer to the following document for Installation location and instructions:

https://www.cisco.com/content/en/us/td/docs/unified_computing/ucs/x/hw/210c-m6/install/b-cisco-ucs-x210c-m7-install.html

SUPPLEMENTAL MATERIAL

Simplified Block Diagram

A simplified block diagram of the Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node system board is shown in *Figure 7*.

Figure 7 Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node Simplified Block Diagram (VIC 25G with Drives)

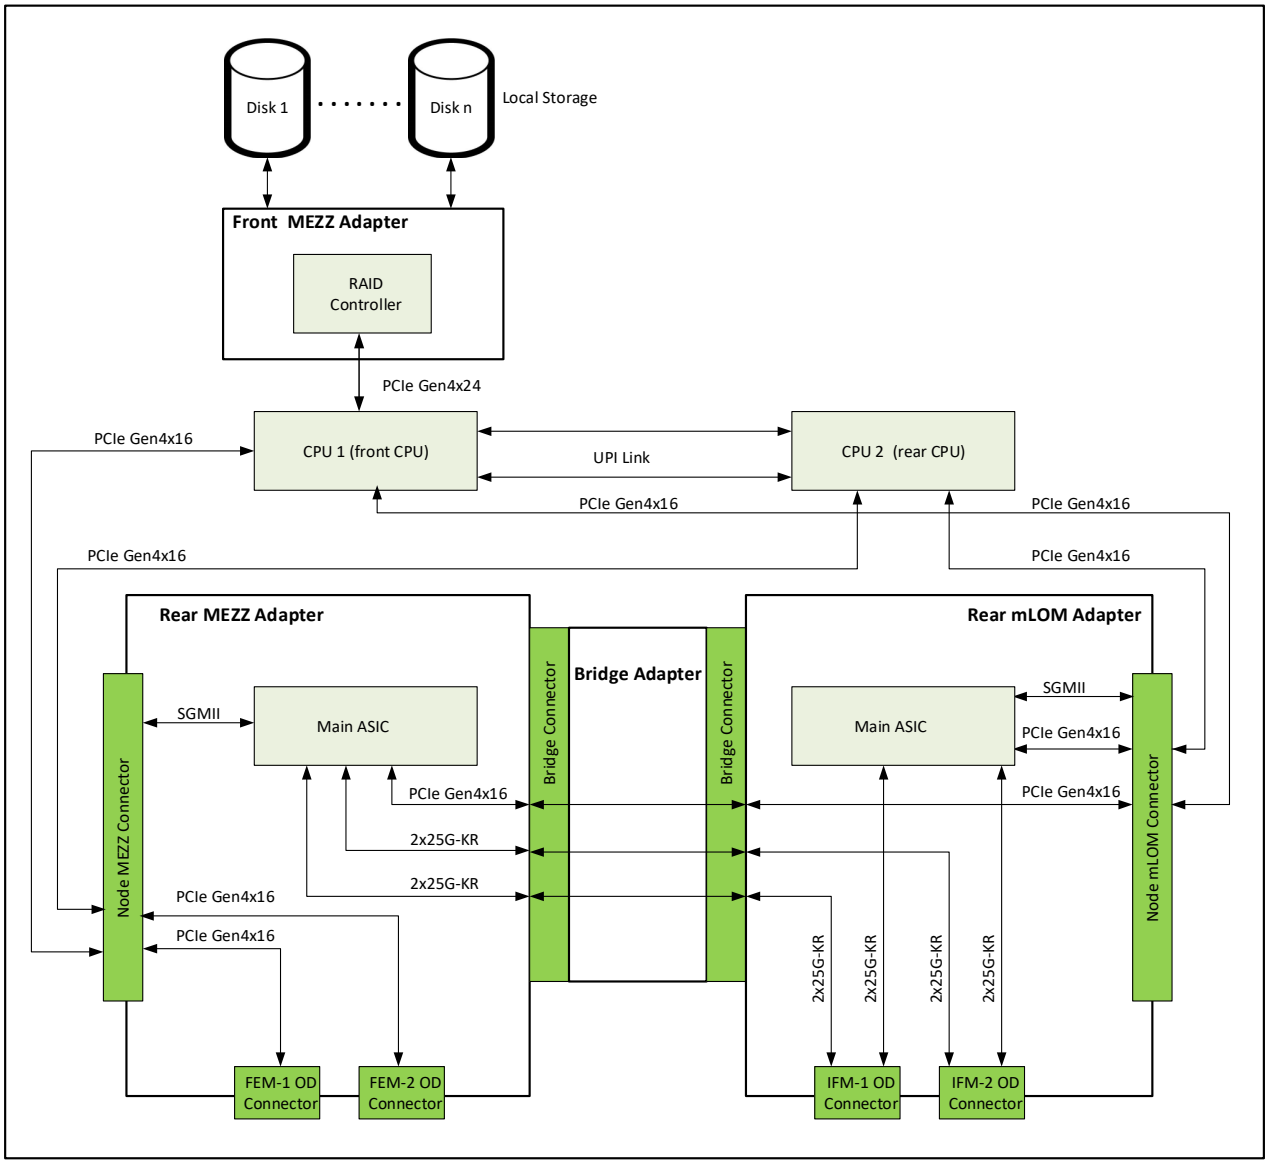

Figure 8 Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node Simplified Block Diagram (VIC 100G with Drives)

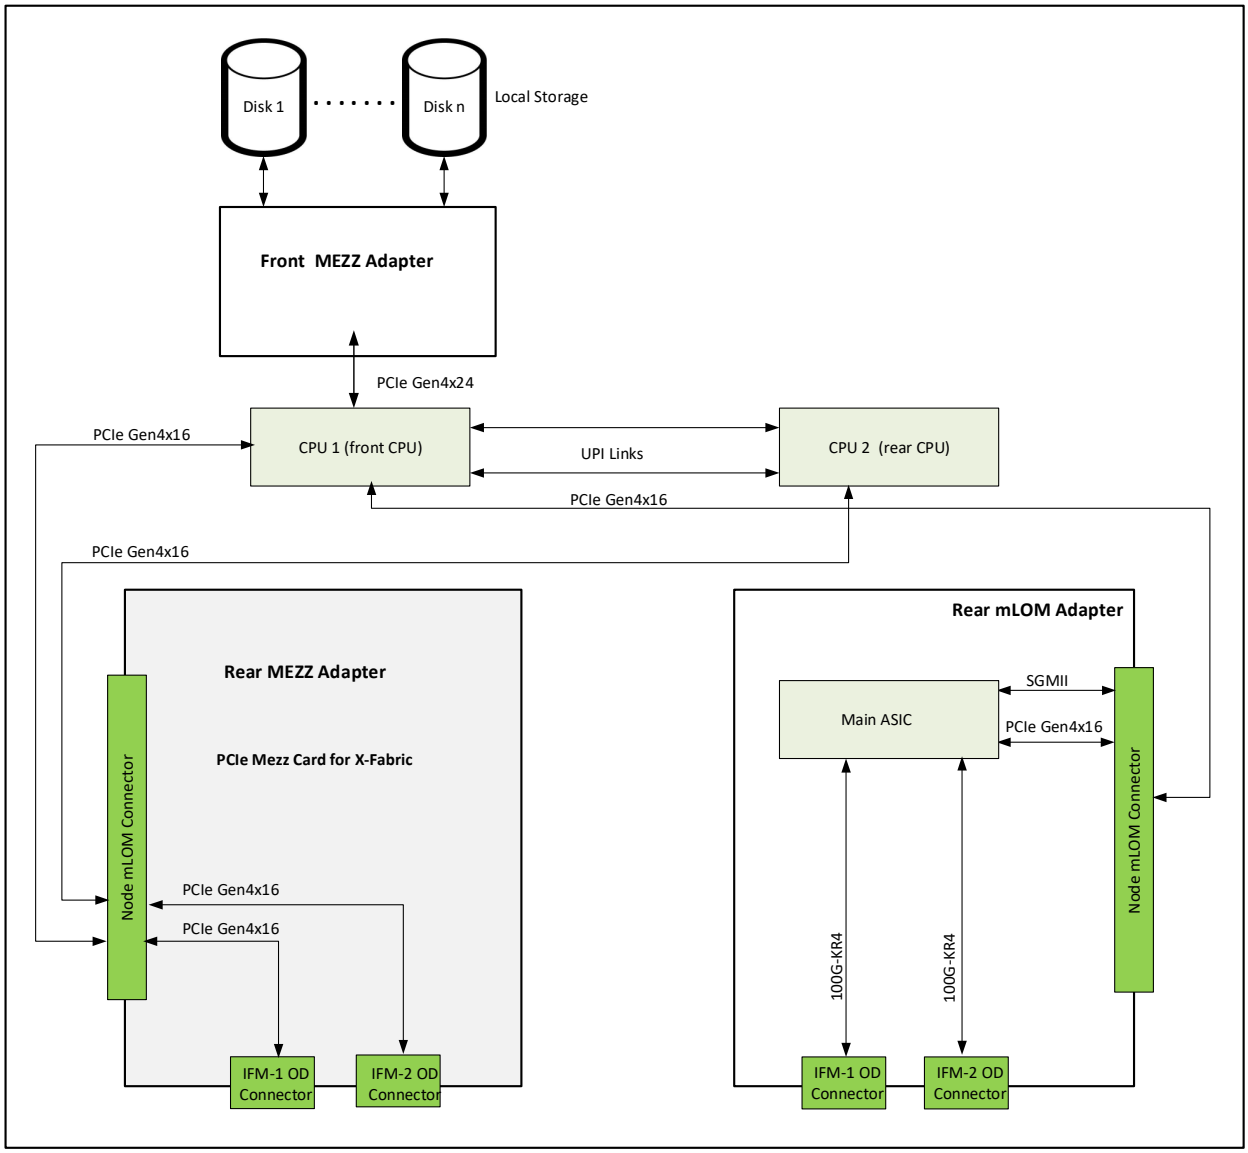

System Board

A top view of the Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node system board is shown in ***Figure 9***.

Figure 9 Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node System Board

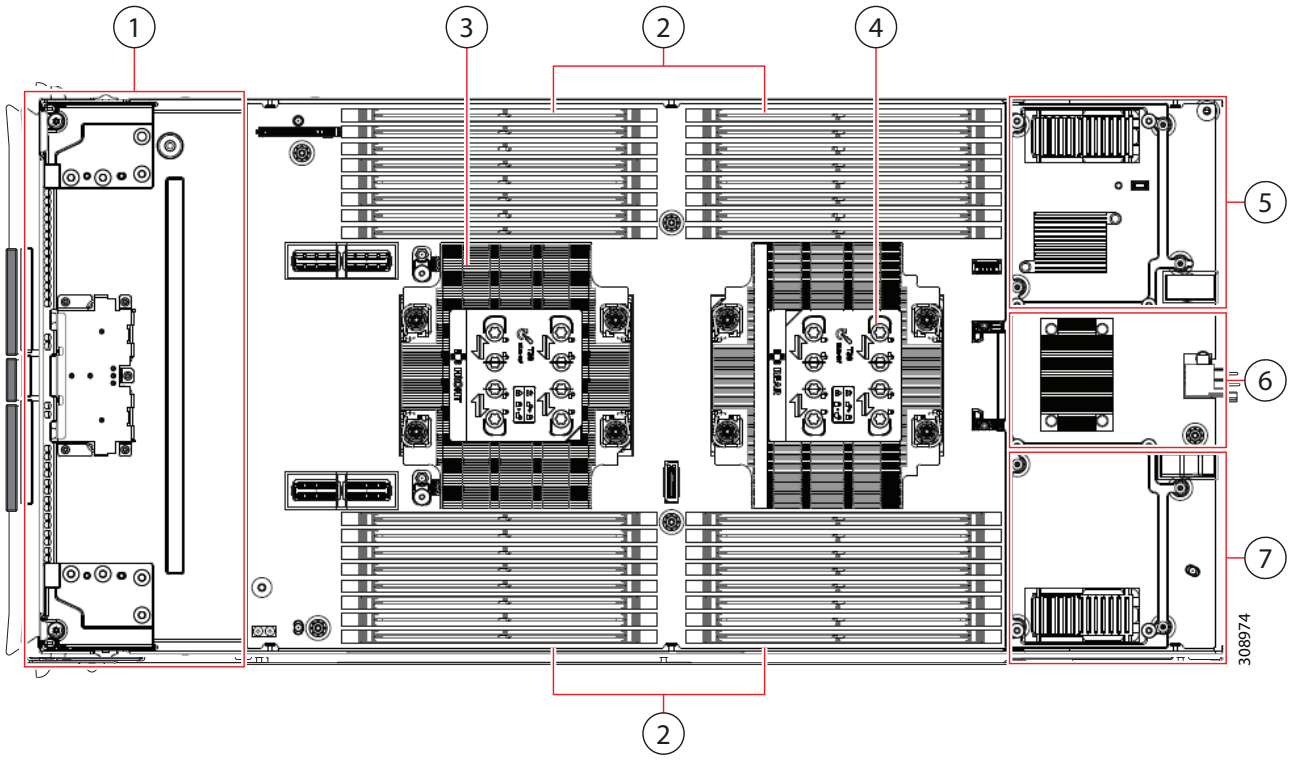

1	Front mezzanine slot for SAS/SATA or NVMe drives and M.2 Controllers.	5	Rear mezzanine slot, which supports a mezzanine card with standard or extended mLOM. If an extended mLOM slot is used, it occupies this slot, such that no rear mezzanine card can be installed.
2	DIMM slots (32 maximum)	6	Bridge adapter (for connecting the mLOM to the rear mezzanine card)
3	CPU 1 slot (shown populated)	7	mLOM slot for a standard or extended mLOM
4	CPU 2 slot (shown populated)	-	-

Please refer to the Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node Installation Guide for installation procedures.

UPGRADING or REPLACING CPUs and Memory

- Refer to Server Installation and Service Guide to upgrading or replacing the CPUs
- Refer to Server Installation and Service Guide to upgrading or replacing the Memory

5TH GEN INTEL XEON BENEFIT PILLARS

INTEL® XEON® PROCESSORS → 5TH GEN INTEL® XEON® PROCESSORS

INTEL XEON HEADLINE: Trusted performance. Exceptional efficiency.

5TH GEN INTEL XEON VALUE PROPOSITION:

Get impressive performance per watt gains across all your workloads,[1] plus outsized performance and TCO in AI, database, networking, and HPC.[2] 5th Gen Intel® Xeon® processors deliver more compute and faster memory[3] at the same TDP as the previous generation. It's software- and platform-compatible with the previous generation, so you can minimize testing and validation when deploying new systems.

AI-FOCUSED LEAD MESSAGE: The processor designed for AI

The processor designed for AI With AI acceleration in every core, 5th Gen Intel® Xeon® processors are ready to handle your demanding AI workloads—including inference and fine tuning on models up to 20 billion parameters[4]—before you need to add discrete accelerators.

5TH GEN INTEL XEON BENEFIT PILLARS:

Performance	AI	Efficiency	Quality and security
■ 21% average performance gain at the same TDP as 4th Gen Intel® Xeon® processors[1] ■ Up to 16% memory bandwidth improvement[3] and 2.7x increased last level cache[4] vs. 4th Gen Intel® Xeon® processors ■ 84% average performance gain over 3rd Gen Intel® Xeon® processors[5] ■ Out-of-box software performance using accelerators	■ Less than 100 ms second-token latency on LLMs under 20 billion parameters[6] ■ Up to 2.7x better AI inference performance vs. 4th Gen AMD EPYC processors[7] ■ Up to 14x better AI training and inference performance vs. 3rd Gen Intel® Xeon® processors[8] ■ Software tools and ecosystem to accelerate AI ■ Confidential computing to help protect AI models	■ Up to 10x higher performance per watt using built-in accelerators[9] ■ Up to 3x higher performance per watt with built-in accelerators vs. 4th Gen AMD EPYC processors[10]	■ Software- and pin-compatible with 4th Gen Intel® Xeon® processors ■ Silicon-based security features, confidential computing, and trust services ■ Leading quality and enhanced telemetry ■ Largest ecosystem of hardware and software vendors

Availability of accelerators varies depending on SKU. Visit the 5th Gen Intel Product Information page for additional product details.

See *Intel® Xeon® Processors Notices and Disclaimers* in next page.

Notes:

1. Average performance gain as measured by the geomean of SPEC CPU rate, STREAM Triad, and LINPACK compared to 4th Gen Intel® Xeon® processor. See G1 at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.

2. As measured by performance per watt on a range of AI, database, networking, and HPC workloads compared to 4th Gen Intel® Xeon® processor. See A2, A19-A25, D1, D2, D5, H1, N16 at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.
3. See G12 at intel.com/processor claims: 5th Gen Intel Xeon Scalable processors. Results may vary
4. See G11 at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.
5. Average performance gain as measured by the geomean of SPEC CPU rate, STREAM Triad, and LINPACK compared to 3rd Gen Intel® Xeon® processor. See G3 at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.
6. Based on Intel internal modelling as of December 2023.
7. Based on performance gains of 1.19x to 2.69x with Intel® Advanced Matrix Extensions (Intel® AMX) for inference on GPT-J, LLaMA-2 13B, DLRM, DistilBERT, BERT-Large, and ResNet50v1.5 compared to AMD EYPC 9654 and 9754. See A201, A202, A208-A211 at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.
8. Based on performance gains of 4.4x to 14.2x for training (ResNet50v1.5, BERT-Large, SSD-ResNet34, RNN-T, MaskRCNN, and DLRM) and 2.9x to14x for inference (ResNet50v1.5, BERT-Large, SSD-ResNet34, RNN-T (BF16 only), Resnext101 32x16d, MaskRCNN (BF16 only), DistilBERT) compared to 3rd Gen Intel® Xeon® processor. See A15-A16 at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.
9. Based on performance per watt gains of 1.46x to 10.6x with built-in accelerators on a range of AI, database, and networking workloads. See A19-A25, D1, D2, D5, N16 at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.
10. Based on performance per watt gains of 1.11x to 2.96x with built-in accelerators on a range of AI, database, networking, and HPC workloads compared to AMD EYPC 9554, 9654, and 9754. See A208-A211, D201-D204, H201, N201 at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary

Intel® Xeon® Processors Notices and Disclaimers

- Availability of accelerators varies depending on SKU. Visit the Intel Product Information page for additional product details.
- Performance varies by use, configuration and other factors.
- Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.
- Your costs and results may vary.
- Intel technologies may require enabled hardware, software or service activation.
- Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.
- © Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.
- Intel contributes to the development of benchmarks by participating in, sponsoring, and/or contributing technical support to various benchmarking groups, including the BenchmarkXPRT Development Community administered by Principled Technologies.

LEADERSHIP PERFORMANCE WITH 4TH GEN INTEL® XEON® PROCESSORS

Improve performance efficiency for critical workloads with the most built-in accelerators.

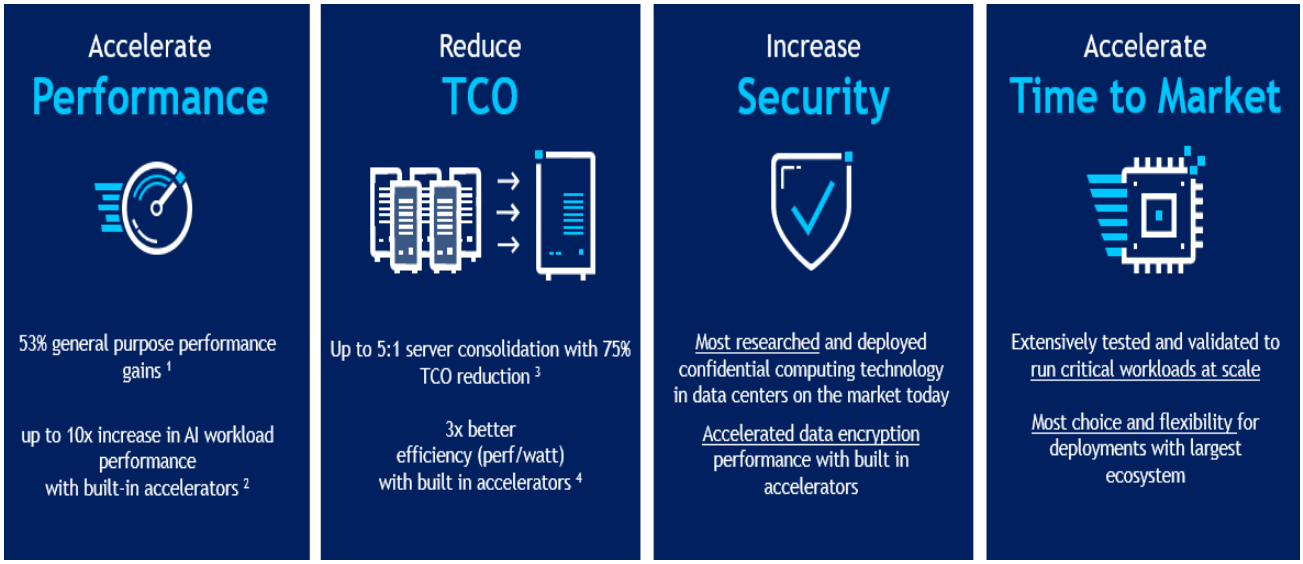

Availability of accelerators varies depending on SKU. Visit the 4th Gen Intel Product Information page for additional product details.

See *Intel® Xeon® Processors Notices and Disclaimers* in next page.

Notes:

1. Compared to prior generation Intel® Xeon® processor. See [G1] at intel.com/processorclaims: 4th Gen Intel® Xeon® Scalable processors. Results may vary.

2. Compared to prior generation Intel® Xeon® processor. See [A16] and [A17] at intel.com/processorclaims: 4th Gen Intel® Xeon® Scalable processors. Results may vary.

3. Comparing benefits transitioning from Intel® Xeon® 4110 to Intel® Xeon® 5420+. See [E11] at intel.com/processorclaims: 4th Gen Intel® Xeon® Scalable processors. Results may vary.

4. Compared to prior generation Intel® Xeon® processor. See [E1] at intel.com/processorclaims: 4th Gen Intel® Xeon® Scalable processors. Results may vary.

Intel® Xeon® Processors Notices and Disclaimers

- Performance varies by use, configuration and other factors.
- Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.
- Your costs and results may vary.
- Intel technologies may require enabled hardware, software or service activation.
- Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.
-
- Intel contributes to the development of benchmarks by participating in, sponsoring, and/or contributing technical support to various benchmarking groups, including the BenchmarkXPRT Development Community administered by Principled Technologies.

TECHNICAL SPECIFICATIONS

Dimensions and Weight

Table 24 Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node Dimensions and Weight

Parameter	Value
Height	1.80 in. (45.7 mm)
Width	11.28 in. (286.5 mm)
Depth	23.7 in. (602 mm)
Weight	■ Minimally configured node weight = 12.84 lbs. (5.83 kg) ■ Fully configured cisco compute hyperconverged node weight = 25.1 lbs. (11.39 kg)

Environmental Specifications

Table 25 Cisco Compute Hyperconverged HCIX210c M7 All-NVMe Node Environmental Specifications

Parameter	Value
Operating temperature	50° to 95°F (10° to 35°C)
Non-operating temperature	-40° to 149°F (-40° to 65°C)
Operating humidity	5% to 90% noncondensing
Non-operating humidity	5% to 93% noncondensing
Operating altitude	0 to 10,000 ft (0 to 3000m); maximum ambient temperature decreases by 1°C per 300m
Non-operating altitude	40,000 ft (12,000m)

For configuration-specific power specifications, use the Cisco Power Calculator at:

http://ucspowercalc.cisco.com

NOTE: The Cisco Compute Hyperconverged HCIX210c M7 node has a power cap of 1300 Watts for all combinations of components (CPUs, DIMMs, drives, and so on). Also, the ambient temperature must be less than 35 ºC (95 ºF).

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

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