Spec Sheet

Cisco Compute Hyperconverged with vSAN

HCIXNX215c M8 All-NVMe vSAN ReadyNode

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.01 SEPTEMBER 24, 2025

CONTENTS

CONTENTS

OVERVIEW

VMware vSAN Express Storage Architecture (ESA) is a software-defined storage solution that runs natively as part of ESXi hypervisor. It aggregates local storage from multiple hosts to create a shared storage pool for virtual machines that can then be accessed by all hosts in the vSAN cluster.

Cisco Compute Hyperconverged with vSAN solutions are purpose-built platforms that unify compute, storage, and networking into a single, software-defined infrastructure. Cisco and VMware by Broadcom have partnered to deliver a robust, scalable, and high-performance hyperconverged infrastructure (HCI) solution for modern workloads.

VMware vSAN ReadyNodes are pre-configured, tested, and jointly certified by Broadcom and Cisco to deliver enterprise-grade storage performance and reliability for IT customers. When deployed on Cisco UCS® servers qualified as vSAN ReadyNodes customers can confidently build a robust hyperconverged infrastructure stack that maximizes hardware utilization, simplifies operations, and scales linearly with business growth.

Cisco Compute Hyperconverged with Nutanix accelerates and simplifies the delivery of infrastructure and applications, at a global scale, through best-in-class cloud-operating models, industry-leading flexibility, and enhanced support and resiliency capabilities so you can power your hybrid multicloud future with the industry's most complete hyperconverged solution.

The Cisco Compute Hyperconverged HCIXVS215C M8 All-NVMe vSAN ReadyNode is integrated 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 HCIXVS215C M8 All-NVMe vSAN ReadyNode harness the power of 5th Gen. AMD EPYC™ CPUs with up to 160 cores per processor and ip to 6TB of capacity with 24 x 256GB DDR5-6400 DIMMs, in a 2-socket configuration with 5th Gen. AMD EPYC™ processors. refer to the **CISCO COMPUTE HYPERCONVERGED VSAN READY NODE STANDARD CAPABILITIES and FEATURES on page 5** for more details.

Figure 1 on page 4 shows a front view of the vSAN ReadyNode.

Figure 1 vSAN ReadyNode

Front View with Drives

DETAILED VIEWS

Front View

Figure 2 is a front view of the vSAN ReadyNode.

Figure 2 vSAN ReadyNode Front View

Storage Drives Option

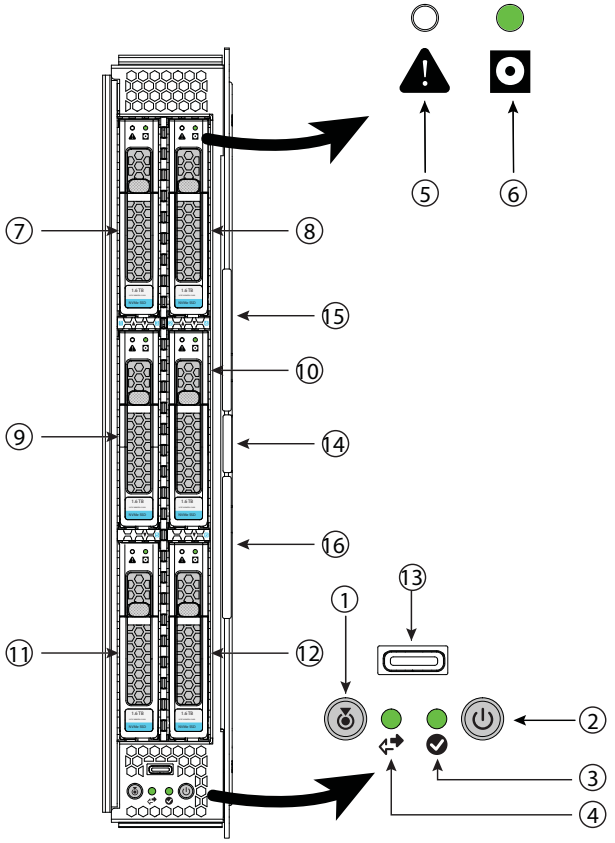

1	Locate button/LED	**9**	Drive Bay 3 (shown populated)
2	Power button/LED	**10**	Drive Bay 4 (shown populated)
3	Status LED	**11**	Drive Bay 5 (shown populated)
4	Network activity LED	**12**	Drive Bay 6 (shown populated)
5	Warning LED (one per drive)	**13**	OCuLink console port[1]
6	Disk drive activity 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 UCSX-C-DEBUGCBL) is required to connect the OCuLink port to the transition serial USB and video (SUV) octopus cable.

CISCO COMPUTE HYPERCONVERGED VSAN ReadyNode STANDARD CAPABILITIES and FEATURES

Table 1 lists the capabilities and features of the base Cisco Compute Hyperconverged HCIXVS215C M8 All-NVMe vSAN ReadyNode. 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 vSAN ReadyNode on page 7**.

Table 1 Capabilities and Features

Capability/Feature	Description
Chassis	■ The vSAN ReadyNodes mounts in a Cisco Compute Hyperconverged 9508 chassis.
CPU	■ One or two AMD EPYC™ 5th Gen Series Processors
Memory	■ 24 slots for registered DIMMs (RDIMMs)
Storage	■ Up to 6 Non-Volatile Memory Express (NVMe) 2.5-inch drives
Additional Storage	■ Dual 80 mm SATA 3.0 M.2 cards on a boot-optimized hardware RAID controller
Mezzanine Adapters (Front)	One front mezzanine connector that supports: ■ Three to six 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.
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.

Table 1 Capabilities and Features *(continued)*

Capability/Feature	Description
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 UCSX-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 vSAN ReadyNode 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 vSAN ReadyNode 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 UCS 6454, 64108 and 6536 fabric interconnects

CONFIGURING THE vSAN ReadyNode

Follow these steps to configure the Cisco Compute Hyperconverged HCIXVS215C M8 All-NVMe vSAN ReadyNode:

STEP 1 CHOOSE BASE CISCO COMPUTE HYPERCONVERGED NODE SKU

Top Level ordering product ID (PID) of the Cisco Compute Hyperconverged HCIXVS215C M8 All-NVMe vSAN ReadyNode as shown in*Table 2*.

Table 2 Top level ordering PID

Product ID (PID)	Description
HCIX-M8-VSAN-MLB	Cisco Compute Hyperconverged X-Series M8 with vSAN MLB

Select the product ID (PID) of the Cisco Compute Hyperconverged HCIXVS215C M8 All-NVMe vSAN ReadyNode as shown in *Table 3*.

Table 3 PID of the Base Cisco Compute Hyperconverged HCIXVS215C M8 All-NVMe vSAN ReadyNode

Product ID (PID)	Description
HCIXVS215C-M8SN	215cM8 All NVMe Hyperconverged Node w/o CPU,Memory,Storage
HCIXVS215C-M8SN-U	215cM8 All NVMe Hyperconverged Node w/o CPU,Memory,Storage

A base Cisco Compute Hyperconverged HCIXVS215C M8 All-NVMe vSAN ReadyNode ordered in *Table 3* does not include any components or options. They must be selected during product ordering.

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:

- CPU-to-CPU communication using Infinity Fabric Interconnect
- Cache size of up to 512 MB
- Up to 160 cores
- Power: Up to 400Watts

Select CPUs

The available 5th Gen. AMD EPYC™ processors are listed in *Table 4*.

CAUTION: For systems configured with processors operating above 28º C [82.4º F], a fan fault or executing workloads with extensive use of heavy instructions sets may assert thermal and/or performance faults with an associated event recorded in the System Event Log (SEL).

Table 4 Available 5th Gen. AMD EPYC™ CPUs

Product ID (PID)[1]	Maximum Socket	Core	CPU Base Frequency	CPU Boost Frequency	Default TDP	Cache Size	Highest DDR5 DIMM Clock
	(S)	(C)	(GHz)	(GHz)	(W)	(MB)	(MT/s)[2]
5th Gen EPYC 9005 Series Processors							
UCSX-CPU-A9845	1S	160	2.10	3.70	390	320	6000
UCSX-CPU-A9825	1S	144	2.20	3.70	390	384	6000
UCSX-CPU-A9745	1S	128	2.40	3.70	400	256	6000
UCSX-CPU-A9655	2S	96	2.60	4.50	400	384	6000
UCSX-CPU-A9645[3]	2S	96	2.30	3.70	320	256	6000
UCSX-CPU-A9565[3]	2S	72	3.15	4.30	400	384	6000
UCSX-CPU-A9555	2S	64	3.20	4.40	360	256	6000
UCSX-CPU-A9535[3]	2S	64	2.40	4.30	300	256	6000
UCSX-CPU-A9455[3]	2S	48	3.15	4.40	300	256	6000
UCSX-CPU-A9365[3]	2S	36	3.40	4.30	300	192	6000
UCSX-CPU-A9355	2S	32	3.55	4.40	280	256	6000
UCSX-CPU-A9335[3]	2S	32	3.00	4.40	210	128	6000
UCSX-CPU-A9255[3]	2S	24	3.20	4.30	200	128	6000
UCSX-CPU-A9135	2S	16	3.65	4.30	200	64	6000
UCSX-CPU-A9115[3]	2S	16	2.60	4.10	125	64	6000
UCSX-CPU-A9015[3]	2S	8	3.60	4.10	125	64	6000
UCSX-CPU-A9575F	2S	64	3.30	5.00	400	256	6000

Table 4 Available 5th Gen. AMD EPYC™ CPUs

Product ID (PID)[1]	Maximum Socket	Core	CPU Base Frequency	CPU Boost Frequency	Default TDP	Cache Size	Highest DDR5 DIMM Clock
	(S)	(C)	(GHz)	(GHz)	(W)	(MB)	(MT/s)[2]
UCSX-CPU-A9475F[3]	2S	48	3.65	4.80	400	256	6000
UCSX-CPU-A9375F[3]	2S	32	3.80	4.80	320	256	6000
UCSX-CPU-A9275F[3]	2S	24	4.10	4.80	320	256	6000
UCSX-CPU-A9175F[3]	2S	16	4.20	5.00	320	512	6000
UCSX-CPU-A9655P[3]	1S	96	2.60	4.50	400	384	6000
UCSX-CPU-A9555P[3]	1S	64	3.20	4.40	360	256	6000
UCSX-CPU-A9455P[3]	1S	48	3.15	4.40	300	256	6000
UCSX-CPU-A9355P[3]	1S	32	3.55	4.40	280	256	6000

Notes:

1. Any CPU PID ending in “P” cannot be used in a 2-CPU system. They can only be used in a 1-CPU system.The X215c M8 is IO optimized for 1 CPU configurations. All storage and network options can be used with one CPU. X-Fabric options for connecting to PCIe nodes require a two CPU configuration.
2. If higher or lower speed DIMMs are selected than what is shown in *Table 6 on page 12* for a given CPU speed, the DIMMs will be clocked at the lowest common denominator of CPU clock and DIMM clock.
3. SKU available in Q2CY25

Supported Configurations

(1) One-CPU Configuration

- Choose one CPU from any one of the rows of *Table 4 on page 9*.

(2) Two-CPU Configuration

- Choose two identical CPUs from any one of the rows of *Table 4 on page 9*.

STEP 3 CHOOSE MEMORY

The *Table 5* below describes the main memory DIMM features supported on vSAN ReadyNode.

Table 5 Server Main Memory Features

Memory DIMM server technologies	Description
DDR5 memory clock speed	5th Gen. AMD EPYC™ CPUs: Up to 6000 MT/s 1DPC
Operational voltage	1.1 Volts
DRAM fab density	16Gb, 24Gb, and 32Gb
DRAM DIMM type	RDIMM (Registered DDR5 DIMM)
Memory DIMM organization	Twelve memory DIMM channels per CPU; 1 DIMM per channel only
Maximum number of DRAM DIMM per server	Up to 24 (2-Socket)
DRAM DIMM Densities and Ranks	32GB 1Rx4, 64GB 2Rx4, 96GB 2Rx4, 128GB 2Rx4, 256GB 4Rx4
Maximum system capacity (DRAM DIMMs only)	6TB (24x256GB)

Figure 3 12-Channel Memory Organization

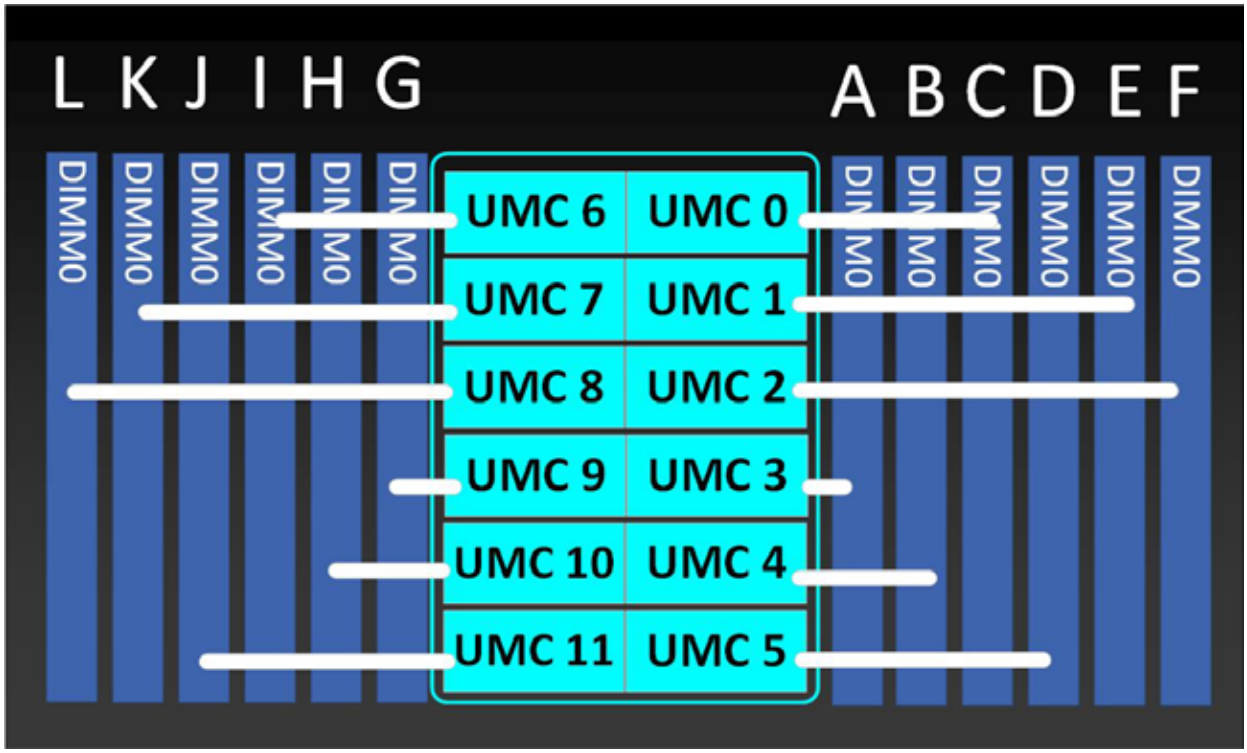

Select DIMMs

The supported memory DIMMs are listed in *Table 6*.

NOTE: All memory DIMMs must be Cisco DDR5-6400 memory PIDs, although the memory will operate at the maximum speed of the 5th Gen. AMD EPYC™ CPUs memory controller, up to 6000 MT/s.

Table 6 Available DDR5 DIMMs

Product ID (PID)	PID Description	Ranks/DIMM
DDR5-6400 MT/s PIDs list[1]		
UCSX-MRX16G1RE5	16GB DDR5-6400 RDIMM 1Rx4 (16Gb)	1
UCSX-MRX32G1RE5	32GB DDR5-6400 RDIMM 1Rx4 (16Gb)	1
UCSX-MRX48G1RF5	48GB DDR5-6400RDIMM 1Rx4 (24Gb)	1
UCSX-MRX64G2RE5	64GB DDR5-6400 RDIMM 2Rx4 (16Gb)	2
UCSX-MRX96G2RF5	96GB DDR5-6400 RDIMM 2Rx4 (24Gb)	2
UCSX-MR128G2RG5	128GB DDR5-6400 RDIMM 2Rx4 (32Gb)	2
UCSX-MR256G4RG5	256GB DDR5-6400 RDIMM 4Rx4 (32Gb)	4
DIMM Blank[2]		
UCSX-DIMM-BLK	UCS DIMM Blank	

Notes:

1. If higher or lower speed DIMMs are selected than for a given CPU speed, the DIMMs will be clocked at the lowest common denominator of CPU clock and DIMM clock. check the *Table 4* column "**Highest DDR5 DIMM Clock Support**"
2. 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.
- For full details on supported memory configurations see the **M8 Memory Guide.**

STEP 4 CHOOSE REAR mLOM ADAPTER

The vSAN ReadyNode 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 7* shows the mLOM adapter choices.

Table 7 mLOM Adapters

Product ID (PID)	Description	Connection type	Compatibility/Functionality
UCSX-MLV5D200GV2	Cisco VIC 15230 modular LOM w/Secure Boot X Cisco Compute Hyperconverged Node	mLOM	■ Supported with both IFM 25G and IFM 100G ■ Operates at 4x 25G with both IFM 25G and IFM 100G
UCSX-ML-V5Q50G	Cisco VIC 15420 4x25G secure boot mLOM for X Cisco Compute Hyperconverged Node	mLOM	■ Supported with both IFM 25G and IFM 100G ■ Operates at 4x 25G with both IFM 25G ■ Operates at 2x 25G with IFM 100G

NOTE:

- 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 vSAN ReadyNode. The bridge adapter connects the mLOM adapter to the rear mezzanine adapter.

Figure 4 Location of mLOM and Rear Mezzanine Adapters

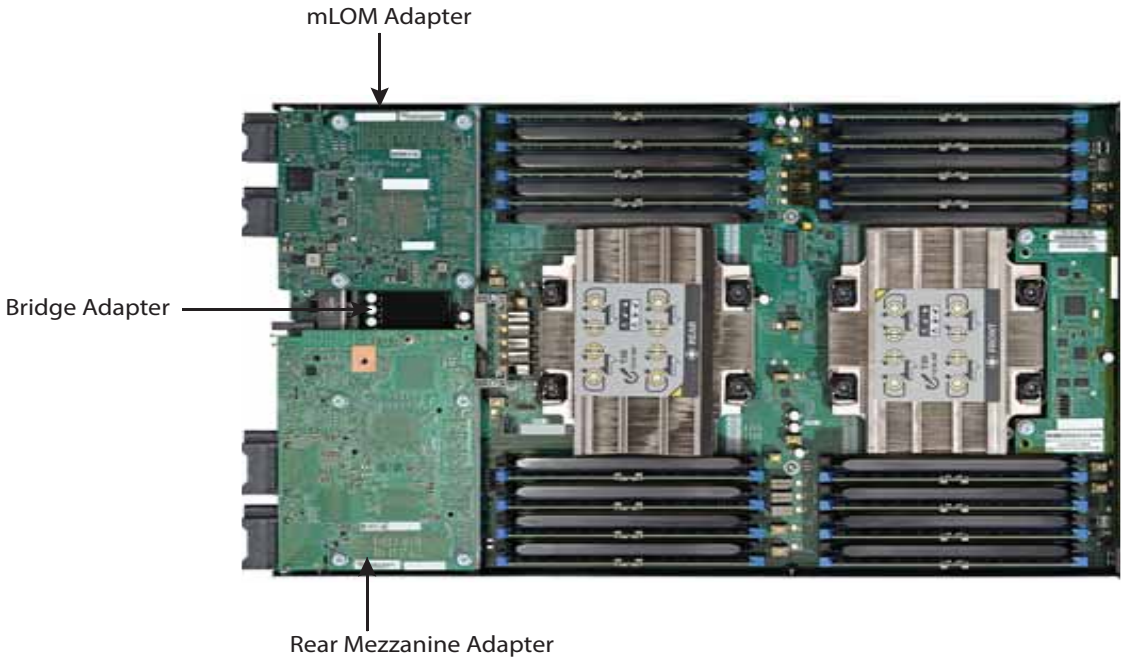

Figure 5 shows the network connectivity from the mLOM out to the 25G IFMs.

Figure 5 Network Connectivity 25G IFMs

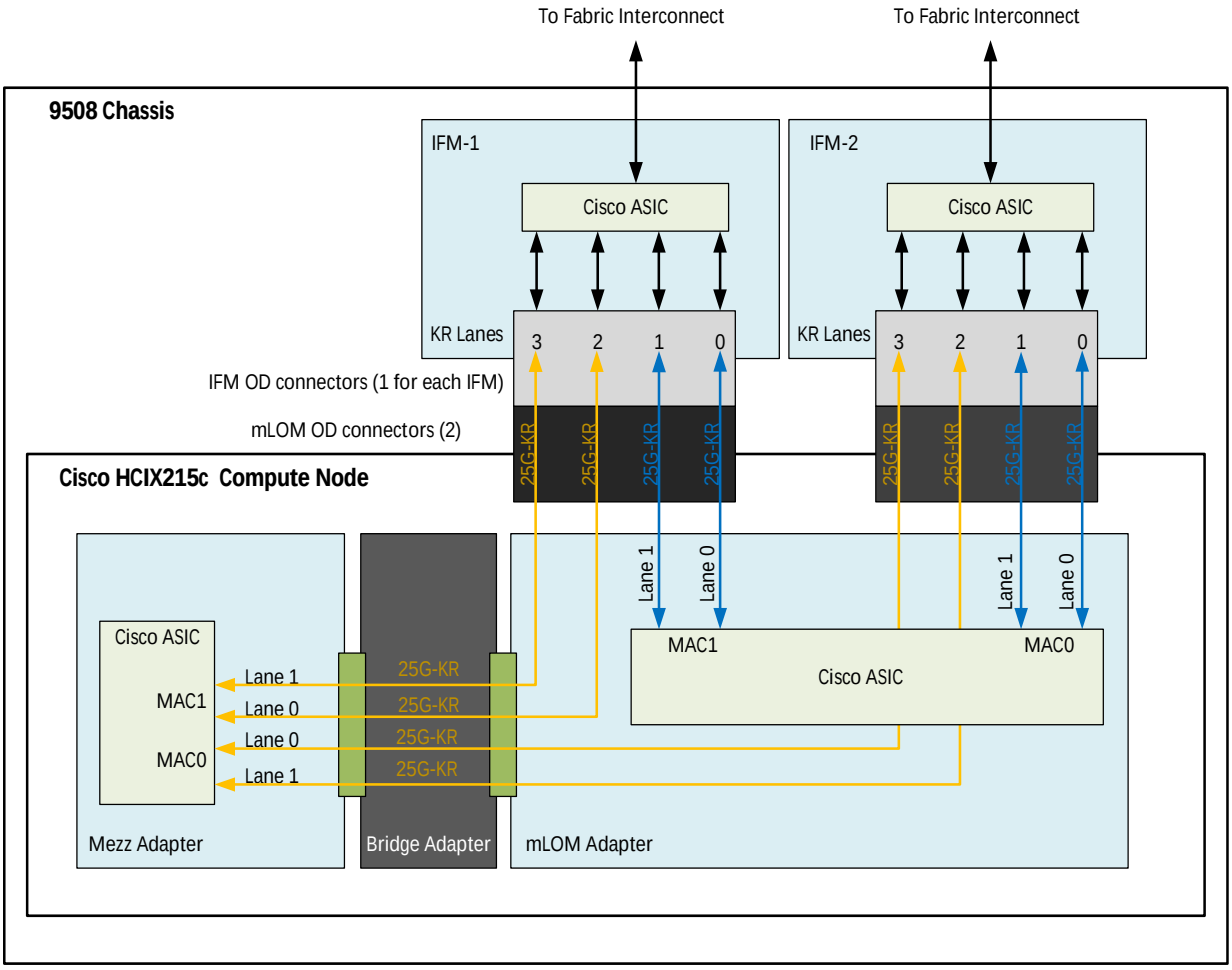

Figure 6 shows the network connectivity from the mLOM out to the 100G IFMs.

Figure 6 Network Connectivity 100G IFMs

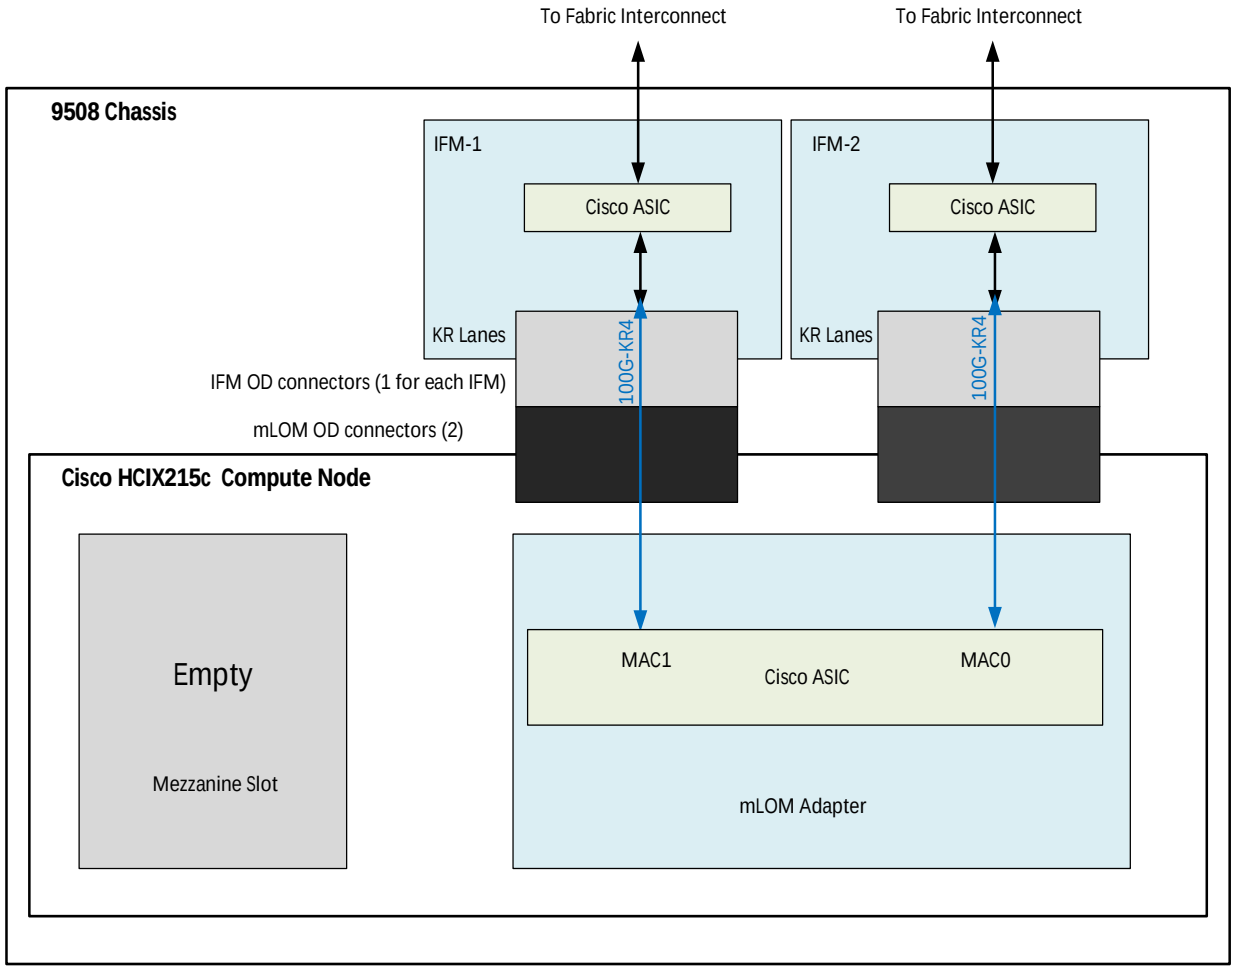

STEP 5 CHOOSE OPTIONAL REAR MEZZANINE VIC/BRIDGE ADAPTERS

The vSAN ReadyNode 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 8* for supported adapters.

Table 8 Available Rear Mezzanine Adapters

Product ID (PID)	PID Description	CPUs Required	Connector Type
Cisco VIC Card			
UCSX-V4-PCIME	PCI Mezz Card for X-Fabric	2 CPUs required	Rear Mezzanine connector on motherboard
UCSX-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[1,2]			
UCSX-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 HCIXVS215C M8 AII-NVMe vSAN ReadyNode)	2 CPUs required	One connector on Mezz card and one connector on mLOM card

Notes:

1. Included with the Cisco VIC 15422 mezzanine adapter.
2. This bridge connects the Cisco VIC 15420 mLOM and Cisco VIC 15422 Mezzanine adapters for the vSAN ReadyNode

NOTE: The rear mezzanine card for X-Fabric has PCIe Gen4 x16 connectivity towards each CPU1 and CPU2. Additionally, the rear mezzanine card also provides two PCIe Gen4 x16 to each X-fabric. This rear mezzanine card enables connectivity from the vSAN ReadyNode to the X440p PCIe node.

Table 9 Throughput Per HCIXVS215C M8 All-NVMe vSAN ReadyNode

HCIX215c M8 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	
X215c 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

- A mLOM VIC from *Table 7* is always required.
- If a rear mezzanine VIC card is installed, a VIC bridge card is included and connects the mLOM to the mezzanine adapter.
- The rear mezzanine card has Ethernet connectivity to the IFM using the VIC bridge card and has a PCIE Gen4 x16 connectivity towards CPU2. Additionally, the rear mezzanine card 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 vSAN ReadyNode has one front mezzanine connector that can accommodate one of the following mezzanine cards:

- Pass-through controller for up to six U.3 NVMe drives.

NOTE: The node can be ordered with or without the front mezzanine adapter. Refer to **Table 10 Available Front Mezzanine Adapters**

Table 10 Available Front Mezzanine Adapters

Product ID(PID)	PID Description	Connector Type
UCSX-X10C-PT4F-D	Cisco Compute Hyperconverged HCIXVS215C M8 All-NVMe vSAN ReadyNode compute pass through controller for up to 6 NVMe drives	Front Mezzanine

STEP 7 CHOOSE OPTIONAL GPU PCIe NODE

Refer to *Table 11* for GPU PCIe Node

Table 11 GPU PCIe Node

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

NOTE: If UCSX-440P is selected, then rear mezzanine is required.

Select GPU Options

The available PCIe node GPU options are listed in *Table 12*.

Table 12 Available PCIe GPU Cards supported on the PCIe Node

GPU Product ID (PID)	PID Description	Maximum number of GPUs per node
UCSX-GPU-A16	NVIDIA A16 PCIE 250W 4X16GB	2
UCSX-GPU-L4	NVIDIA L4 Tensor Core, 70W, 24GB	4
UCSX-GPU-L40	NVIDIA L40 300W, 48GB	2
UCSX-GPU-L40S	NVIDIA L40S: 350W, 48GB, 2-slot FHFL GPU	2
UCSX-GPU-H100-80	NVIDIA H100: 350W, 80GB, 2-slot FHFL GPU	2
UCSX-GPU-H100-NVL	TESLA H100, PASSIVE, 350W, 80GB	2

STEP 8 CHOOSE DRIVES

The available drives are listed in *Table 13*.

Table 13 Available Capacity Drives

Product ID (PID)	PID Description	Drive Type	Capacity
Capacity Drive			
UCSX-NVMEG4M1600D	1.6TB 2.5in U.3 15mm P7450 Hg Perf Hg End NVMe (3X)	NVMe	1.6 TB
UCSX-NVMEG4M1920D	1.9TB 2.5in U.3 15mm P7450 Hg Perf Med End NVMe	NVMe	1.9 TB
UCSX-NVMEG4M3840D	3.8TB 2.5in U.3 15mm P7450 Hg Perf Med End NVMe	NVMe	3.8 TB
UCSX-NVMEG4M7680D	7.6TB 2.5in U.3 15mm P7450 Hg Perf Med End NVMe	NVMe	7.6 TB
UCSX-NVMEG4M1536D	15.3TB 2.5in U.3 15mm P7450 Hg Perf Med End NVMe	NVMe	15.3 TB
UCSX-NVMEG4M3200D	3.2TB 2.5in U.3 15mm Micron P7450 Hg Perf Hg End NVMe (3X)	NVMe	3.2 TB
UCSX-NVMEG4M6400D	6.4TB 2.5in U.3 Micron 7450 NVMe High Perf High Endurance	NVMe	6.4 TB

Approved Configurations

- Three to six U.3 NVMe drives

STEP 9 ORDER M.2 BOOT RAID CONTROLLER and SATA SSDs

- **Cisco 6GB/s SATA Boot-Optimized M.2 RAID Controller (included):** Boot-Optimized RAID controller (UCSX-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 UCSX-M2-HWRD-FPS is auto included with the server configuration
- The UCSX-M2-HWRD-FPS controller supports RAID 1 and JBOD mode
- 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 Windows and Linux Operating Systems

Table 14 Boot-Optimized RAID controller (auto included)

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

- **Select Cisco M.2 SATA SSDs:** Order two matching M.2 SATA SSDs. This connector accepts the boot-optimized RAID controller (see *Table 14*). Each boot-optimized RAID controller can accommodate two SATA M.2 SSDs shown in *Table 15*.

NOTE:

- Each boot-optimized RAID controller can accommodate up to two SATA M.2 SSDs shown in *Table 15*. 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 15 M.2 SATA SSDs

Product ID (PID)	PID Description	Drive Type	Capacity
UCSX-M2-480G	480GB M.2 SATA SSD	SATA	480GB
UCSX-M2-960G-D	960GB M.2 SATA SSD	SATA	960GB
UCSX-M2480OA1V	480GB M.2 Boot Solidigm S4520 SATA 1X SSD	SATA	480GB

STEP 10 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 vSAN ReadyNode. 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 16 Available TPM Option

Product ID (PID)	Description
UCSX-TPM2-002D	TPM 2.0 FIPS 140-2 MSW2022 compliant AMD M8 servers
HCI-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 vSAN ReadyNode with a TPM is returned, the replacement vSAN ReadyNode must be ordered with a new TPM. If there is no existing TPM in the vSAN ReadyNode, you can install a TPM 2.0. Refer to the following document for Installation location and instructions: Cisco UCS X215c M8 Server Installation and Service Guide

SUPPLEMENTAL MATERIAL

Simplified Block Diagram

A simplified block diagram of the vSAN ReadyNode system board is shown in *Figure 7*.

Figure 7 vSAN ReadyNode Simplified Block Diagram (VIC 25G with Drives)

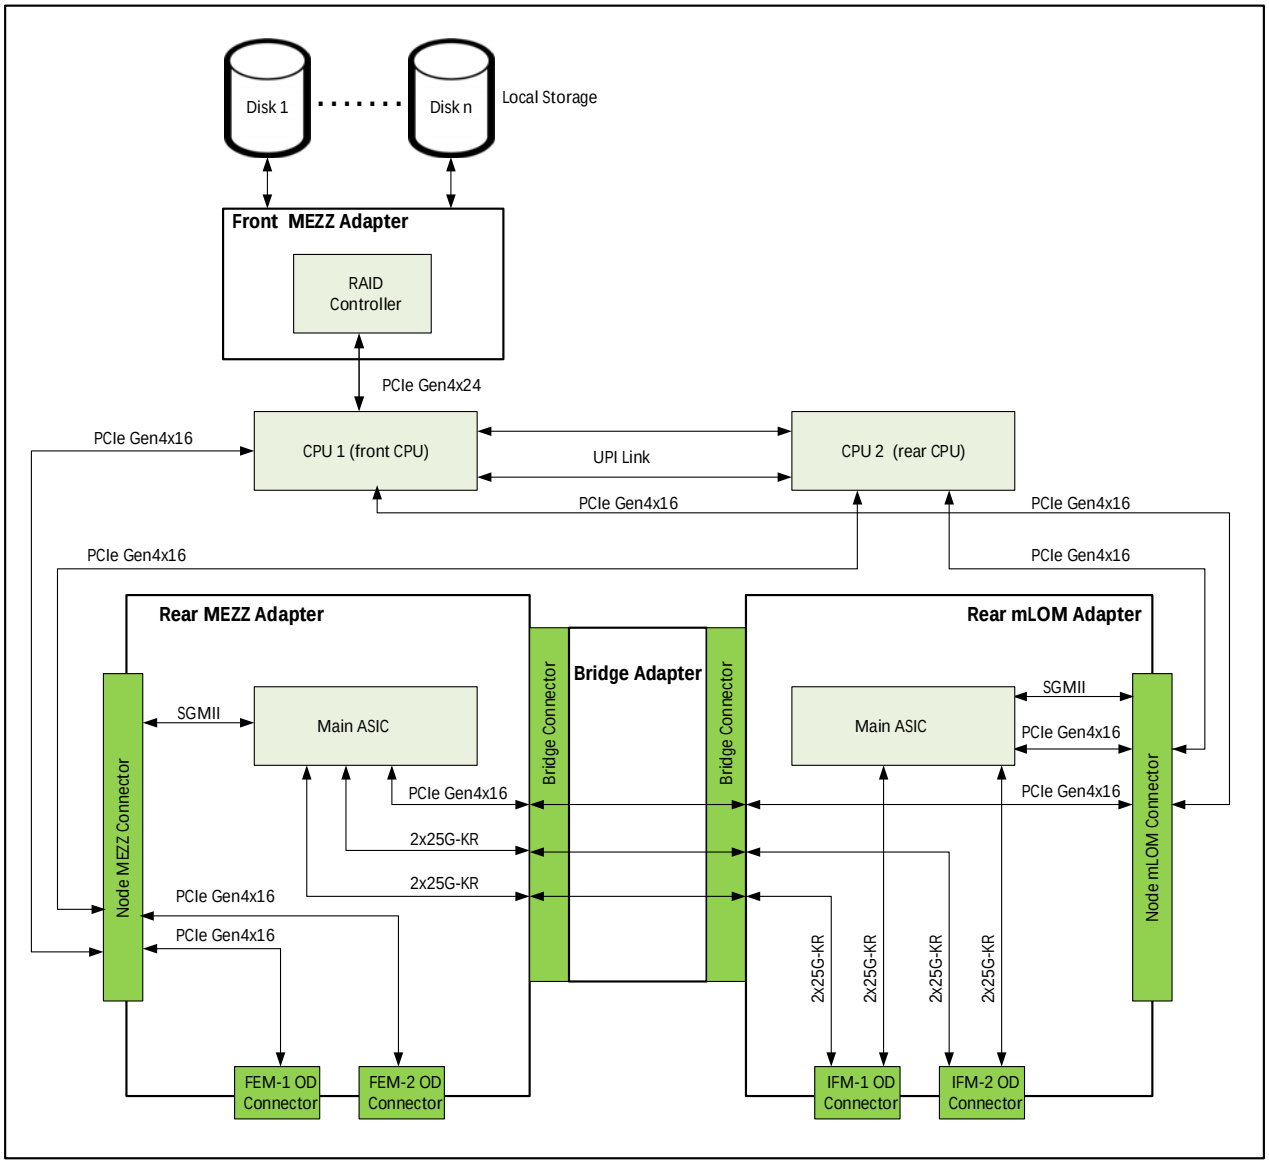

Figure 8 vSAN ReadyNode Simplified Block Diagram (VIC 100G with Drives)

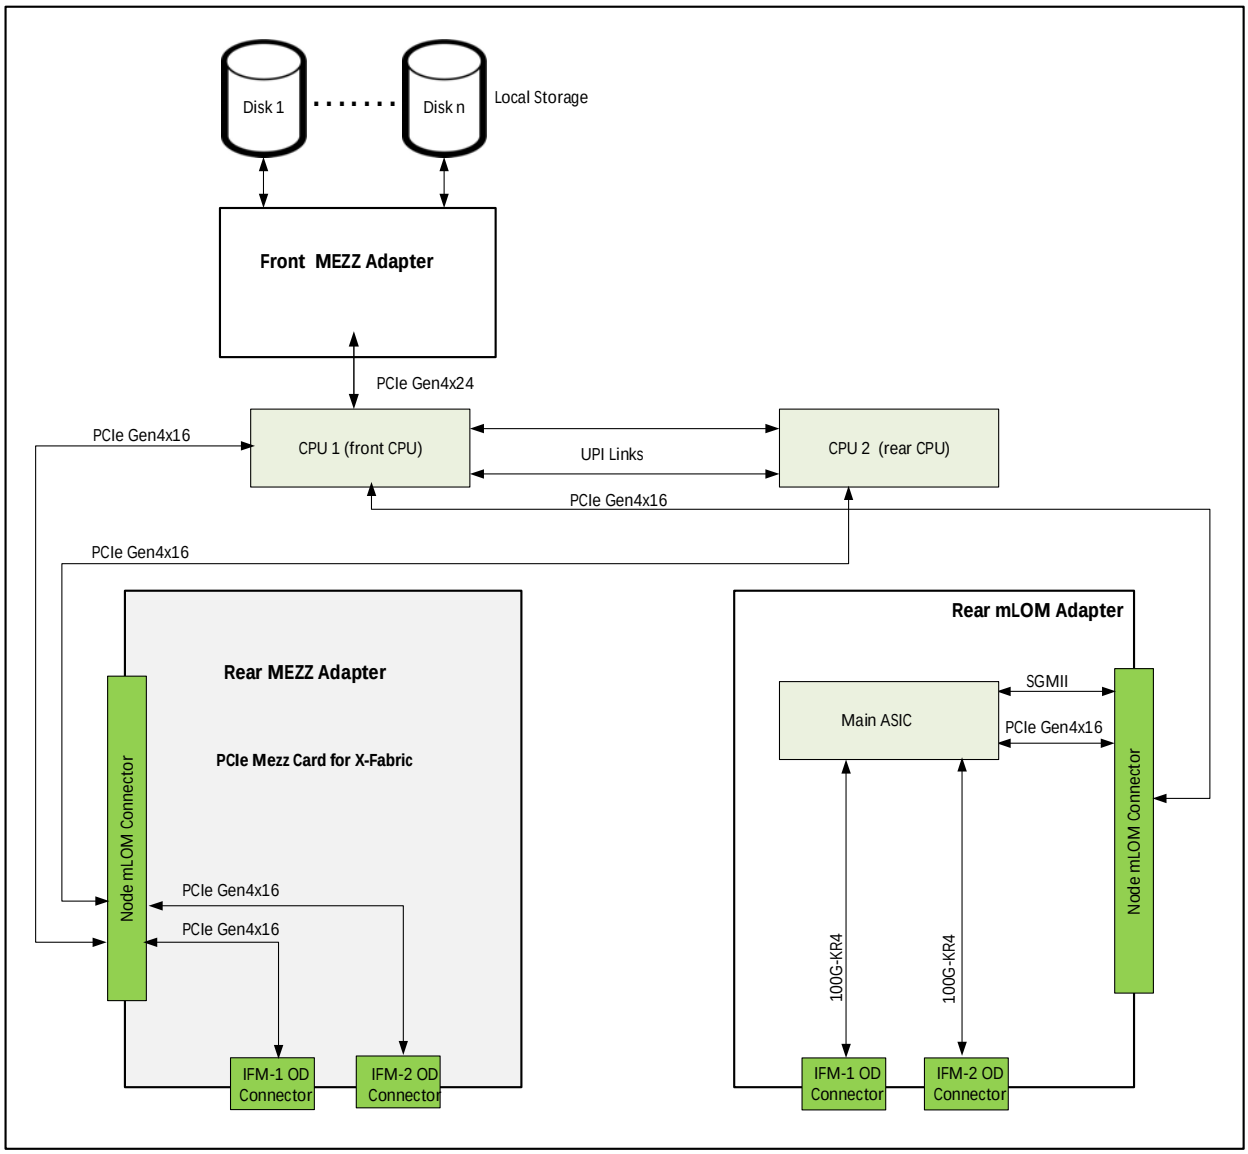

UPGRADING or REPLACING CPUs and Memory

Refer to **Cisco UCS X215c M8 Server Installation and Service Guide** for upgrading or replacing the CPUs and Memory Devices.

TECHNICAL SPECIFICATIONS

Dimensions and Weight

Table 17 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 vSAN ReadyNode weight = 25.1 lbs. (11.39 kg)

Environmental Specifications

Table 18 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 HCIX215c M8 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.**