

Cisco UCS X580p PCIe 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/servers-unified-computing/ucs-x-series-modular-system/datasheet-listing.html>



OVERVIEW	3
Cisco UCS X580p PCIe Node	3
Cisco UCS X9516 X-Fabric	3
PCIe Node STANDARD CAPABILITIES and FEATURES	5
CONFIGURING the Cisco UCS X580p PCIe Node	6
STEP 1 CHOOSE BASE CISCO UCS X580p PCIe Node SKU	7
STEP 2 ORDER GPU CARDS	8
X580p GPU Assignment And Population Rules:	13
Supported GPU Assignment Configurations:	13
Unsupported GPU Assignment Configurations:	13
STEP 3 ORDER CISCO UCS X9516 X-FABRIC MODULES	14
STEP 4 SELECT OPTICS OPTIONS (REQUIRED)	15
SUPPLEMENTAL MATERIAL	17
System Board	17
Solution Topology	19
X580p PCIe Node Configuration - 2 Server 4x GPU + 2 NIC each	20
X580p PCIe Node Configuration - 2 Server 4x GPU + 2 NIC each Chassis PCIe Mapping Policy	20
X580p PCIe Node Configuration	21
X580p PCIe Node Configuration - PCIe Mapping Policy	21
Spare Parts	23
TECHNICAL SPECIFICATIONS	24
Dimensions and Weight	24
Environmental Specifications	25

OVERVIEW

The Cisco UCS 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.

Cisco UCS X580p PCIe Node

The Cisco UCS X580p PCIe Node extends the modular capabilities of the Cisco UCS X-Series Modular System by delivering high-performance GPU support for a wide range of workloads including generative and agentic AI fine-tuning, inference, and legacy applications that require GPU acceleration. As the second generation of PCIe nodes, the X580p supports up to four PCIe GPUs and can be paired with the Cisco UCS X210c M8 Compute Node with Intel® Xeon® 6 processors and Cisco UCS X210c M7 Compute Node with 5th Generation Intel® Xeon® Scalable Processors, as well as the UCS X215c M8 compute node with EPYC processors offering Compute flexibility and choice, while leveraging Cisco UCS X-Fabric Technology for GPU connectivity.

The Cisco UCS X580p PCIe Node is the latest PCIe resource node to integrate into the Cisco UCS X-Series Modular System. The X580p supports up to four FHFL PCIe GPUs and can be paired with the X215c M8 and X210c M8/M7 compute nodes offering flexibility and choice, while also leveraging Cisco UCS X-Fabric Technology. Occupying two slots in the X-series Chassis, the X580p offers significantly greater flexibility than the first-generation X440p, allowing users to assign up to four GPUs. Up to two PCIe Nodes can reside in the 7-Rack-Unit (7RU) Cisco UCS X9508 Chassis and can be paired with one or two compute node each, offering up to four GPUs to a Cisco X-Series Compute Node with Cisco UCS X-Fabric Technology.

The UCS 2nd generation X-Fabric Technology solution is a combination of two products: the Cisco UCS X9516 X-Fabric Module which provides a PCIe Gen5 fabric and the Cisco UCS X580p PCIe Node which hosts the GPUs.

The Cisco UCS X9508 Chassis has eight node slots, up to four of which can be X580p PCIe Nodes when paired with a Cisco X-Series compute node. This provides up to 8 GPUs per chassis to accelerate your applications.

Cisco UCS X580p supports several GPUs please refer to [STEP 2 ORDER GPU CARDS, page 8](#) for the available GPUs

Cisco UCS X9516 X-Fabric

The Cisco UCS X9516 X-Fabric is the second-generation X-Fabric module for the UCS X9508 chassis. It brings ultra-low-latency PCIe Gen5 switching and SmartNIC support, enabling disaggregated and composable IT resources. Managed via Cisco Intersight®, the X9516 allows GPUs and NICs to be dynamically provisioned across servers using policy-based management.

Complementary to the X9516, the UCS X580p PCIe Node expands GPU resources by supporting up to 4 dual-slot GPUs per node with NVLink Bridge support. This combination makes the X9516 the control point for modular GPU compositability in AI-driven infrastructures.

Figure 1 Front and rear views of Cisco UCS X580p PCIe Node

Front View



Rear View



Figure 2 Front and rear views of UCS X9516 X-Fabric

Front View



Rear View



PCIe Node STANDARD CAPABILITIES and FEATURES

Table 1 lists the capabilities and features of the base Cisco UCS X580p PCIe Node. Details about how to configure the PCIe Node for a listed feature or capability are provided in [CONFIGURING the Cisco UCS X580p PCIe Node on page 6](#).

Table 1 Capabilities and Features

Capability/Feature	Description
Chassis	The Cisco UCS X580p PCIe Node mounts in a Cisco UCS X9508 chassis.
GPU Slots	2x Riser Cages with 2x GPU slots per cage
Available GPUs	<ul style="list-style-type: none"> ■ NVIDIA Please refer Table 3 for the complete GPU Cards PID lists
Power subsystem	Power is supplied from the Cisco UCS X9508 chassis power supplies. The Cisco UCS X580p PCIe Node consumes a maximum of 2400 W.
Fans	Integrated in the Cisco UCS X9508 chassis.
Integrated management processor	The built-in Cisco Integrated Management Controller enables monitoring of Cisco UCS X580p PCIe Node inventory, health, and system event logs.
ACPI	Advanced Configuration and Power Interface (ACPI) 4.0 Standard Supported. ACPI states S0 and S5 are supported. There is no support for states S1 through S4.
Front Indicators	<ul style="list-style-type: none"> ■ Status indicator ■ Location indicator
Management	<ul style="list-style-type: none"> ■ Managed PCIe node with dedicated BMC support. ■ Policy-based GPU management via Cisco Intersight.
Chassis	Compatible with the Cisco UCS 9508 X-Series Server Chassis

CONFIGURING the Cisco UCS X580p PCIe Node

Follow these steps to configure the Cisco UCS X580p PCIe Node:

- [STEP 1 CHOOSE BASE CISCO UCS X580p PCIe Node SKU, page 7](#)
- [STEP 2 ORDER GPU CARDS, page 8](#)
- [STEP 3 ORDER CISCO UCS X9516 X-FABRIC MODULES, page 14](#)
- [STEP 4 SELECT OPTICS OPTIONS \(REQUIRED\), page 15](#)

STEP 1 CHOOSE BASE CISCO UCS X580p PCIe Node SKU

Verify the product ID (PID) of the Cisco UCS X580p PCIe Node as shown in [Table 2](#).

Table 2 PIDs of the Base Cisco UCS X580p PCIe Node

Product ID (PID)	Description
UCSX-580P	UCS X-series PCIe Node

A base Cisco PCIe Node ordered in [Table 2](#) 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 PCIe Node:

- GPUs
- Cisco UCS X9516 X-Fabric Modules

STEP 2 ORDER GPU CARDS

The Cisco UCS X-Series solution, managed by Intersight, offers significant flexibility in how GPUs are assigned and utilized, enabling dynamic resource allocation based on workload needs.

- **Variable GPU Allocation per Compute Node:**
 - Supports assigning up to 4 GPUs per compute node (e.g., X210c or X215c).
 - Also supports assigning 1 or 2 GPUs per compute node, allowing for granular control over GPU resources.
- **GPU Sharing Capability:** The X580p PCIe Node design allows for the sharing of GPUs across two compute nodes, maximizing resource utilization and flexibility.
- **Intersight Policy-Based Management:**
 - **Automated Assignment:** Intersight handles the assignment of PCIe Nodes and their contained GPUs to specific compute nodes based on defined policies.
 - **NIC Assignment:** Intersight also assigns necessary NICs (e.g., East-West NICs for AI Fabric) within the compute node's profile.
 - **PCIe Lane Allocation:** The system automatically allocates the necessary PCIe lanes on the internal switches to ensure proper connectivity and performance.
- **Physical Slot Configuration for Connectivity:**
 - PCIe Nodes (X580p) can be placed in chassis slots 1/2, 3/4, 5/6, or 7/8.
 - Compute Nodes establish connectivity with PCIe Nodes located on the same side of the chassis.
 - Example: A compute node in Slot 1/2 connects to a PCIe Node in Slot 3-4.
 - Example: A compute node in Slot 7/8 connects to a PCIe Node in Slot 5-6.
- **Profile Based Management:** The existing Intersight server profile will be used, with an added PCIe device policy to specify GPU and NIC mappings for compute nodes.
- **PCIe Device Connection Policies:** Policies can be defined to filter and assign GPUs based on criteria such as the number of GPUs required or specific references types, streamlining deployment.

Select GPU Options

**NOTE:**

- Mixing different GPU types is allowed, but mixed GPUs cannot be installed in the same “Cage.” For example, a UCSX-GPU-L40S in Cage A Slot 1 and a UCSX-GPU-H200-NVL in Cage B Slot 3 is allowed, but they cannot both be installed in Cage A.
- **Server Node and GPU Quantity:** Mixing of GPUs on a Compute node is not allowed, all the GPUs assigned must have the same Product ID (PID).
- **Server Memory Recommendation:** If x580P is selected, Cisco recommends the Server Memory quantity to be 3 times per GPU's memory size.
- If UCSX-580P is selected in the 9508 Chassis, then a quantity of 1 PID (UCSX-V5-PCIME) is required for each server node mapped to a UCSX-580P. For example, if there are 1 or 2 UCSX-580P units and 2 to 4 server nodes, each server node must include a UCSX-V5-PCIME to access the GPUs.

The available GPU PCIe options and their riser slot compatibilities are listed in [Table 3](#).

Table 3 Available PCIe GPU Cards

GPU Product ID (PID)	PID Description	Riser Slot Compatibility
Supported GPUs on X210c/X215c M8 Compute Node		
UCSX-GPU-L40S	NVIDIA L40S: 350W, 48GB, 2-slot FHFL GPU	Cage A & B, Slot 1, 2, 3, 4
UCSX-GPU-H200-NVL ^{1,2}	NVIDIA H200 NVL: 600W, 141GB, 2-Slot NVL2 FHFL GPU	Cage A & B, Slot 1, 2, 3, 4
UCSX-GPU-RTXP6000	NVIDIA RTX Pro 6000: 600W, 96GB, 2-Slot FHFL GPU	Cage A & B, Slot 1, 2, 3, 4
UCSX-GPU-RTXP4500 ³	Nvidia RTX Pro 4500 GPU 160W 32GB Gen5 FHFL 1-slot	Cage A & B, Slot 1, 2, 3, 4
Supported GPUs on X210c M7 Compute Node		
UCSX-GPU-H200-NVL ^{1,2}	NVIDIA H200 NVL: 600W, 141GB, 2-Slot NVL2 FHFL GPU	Cage A & B, Slot 1, 2, 3, 4
UCSX-GPU-RTXP6000	NVIDIA RTX Pro 6000: 600W, 96GB, 2-Slot FHFL GPU	Cage A & B, Slot 1, 2, 3, 4
UCSX-GPU-RTXP4500 ⁶	Nvidia RTX Pro 4500 GPU 160W 32GB Gen5 FHFL 1-slot	Cage A & B, Slot 1, 2, 3, 4

Notes:

1. Please check the [Select NVL Bridge \(optional\)](#), page 11 section for the full details
2. CBL-X580p-GPU-N is auto included if this GPU is selected
3. CBL-X580P-GPU-N is auto included if this GPU is selected



NOTE: Following [STEP 3 ORDER CISCO UCS X9516 X-FABRIC MODULES](#) is optional only if the Cisco UCS X9508 Chassis already has the UCS X9516 X-Fabric modules installed and the X-Series compute node has one of the supported mezzanine adapters to connect to Cisco UCS X580p PCIe Node

Select NVL Bridge (optional)

The available GPU bridges options are listed in [Table 4](#)



NOTE:

- **Availability:** The NVL Bridge (UCSX-NVL2-H200) option shows only when 2 NVL-H200 GPUs are selected.
- **Recommendation:** If selected, Cisco recommends a minimum quantity of 2 UCSX-GPU-H200-NVL GPUs.
- **Placement:** A maximum of 2 GPUs per Bridge and a maximum of 2 NVL Bridges per system are allowed.
- **Side-by-Side Requirement:** If a Bridge is selected, NVL-H200 GPUs need to be installed side-by-side (e.g., Slots 1 & 2 or 3 & 4).
- **Removal Tool Kit:** If the NVL Bridge PID UCSX-NVL2-H200 is selected, 1 UCSX-GPU-RKIT-NV (GPU/NVLINK Bridge Removal Tool Kit NVIDIA) is required.

Table 4 PCIe NVL Bridge

Product ID (PID)	PID Description	Supported GPUs
UCSX-NVL2-H200	NVIDIA NVL-2way Bridge for H200 GPU; NVPN 900-23945-0000-000	UCSX-GPU-H200-NVL
Accessories/spare included: <ul style="list-style-type: none"> ■ UCSX-GPU-RKIT-NV: NVIDIA GPU and NVLINK Bridge Removal Tool Kit 		

Select NVL GPU LICENSE

The available NVL GPU License options are listed in [Table 4](#).



NOTE:

- Requirement: Selection of an NVIDIA LICENSE is required when UCSX-GPU-L40S or UCSX-GPU-RTXP6000 GPUs are selected.
- Opt-Out: Users can select PID NV-GRID-OPT-OUT-D to opt out of the NVIDIA License.
- License Quantity: The NVIDIA License quantity must equal the quantity of GPUs selected.
- Available Licenses:
 - Standard License: NV-AIE-P, NV-AIE-S, NV-GRID-PCS, NV-GRID-VAS, NV-GRID-WKS.
 - Education License: NV-AIE-EDP, NV-AIE-EDS, NV-GRID-EDS.
 - NVIDIA Optout: NV-GRID-OPT-OUT-D.

Table 5 NVL GPU License

Product ID (PID)	PID Description
Standard License	
NV-AIE-P	NVIDIA AIE Essentials Perp Lic & Support per GPU
NV-AIE-S	NVIDIA AI Enterprise Essentials Subscription per GPU
NV-GRID-PCS	NVIDIA GRID Software Subscription - VDI PC 1CCU
NV-GRID-VAS	NVIDIA GRID Software Subscription - VDI Apps 1CCU
NV-GRID-WKS	NVIDIA Quadro SW Subscription - vDWS 1CCU
Education License	
NV-AIE-EDP	NVIDIA AIE Essentials Perp Lic & Support per GPU, EDU
NV-AIE-EDS	NVIDIA AIE Essentials Subscription per GPU, EDU
NV-GRID-EDS	EDU - NVIDIA Quadro vDWS SW Subscription - 1CCU
NVIDIA Optout	
NV-GRID-OPT-OUT-D	NVIDIA GRID SW OPT-OUT

X580p GPU Assignment And Population Rules:

Supported GPU Assignment Configurations:

Table 6 Supported GPU Assignment Configurations

Scenario	GPU Assignment	Compute Node	Notes
Single Host	4x GPUs (Same type)	Node 1	All 4 GPUs must be identical (e.g., all H200 or all L40S)
Single Host	1x GPU	Node 1	Can be from any Cage/Riser
Single Host	2x GPUs	Node 1	Both GPUs must be from Cage 1
Single Host	2x GPUs	Node 1	Both GPUs must be from Cage 2
Multi-Host	Cage 1 (e.g., 2x H200-NVL)	Node 1	Same PCIe Switch
Multi-Host	Cage 2 (e.g., 2x RTX PRO 6k)	Node 2	Same PCIe Switch

Unsupported GPU Assignment Configurations:

Table 7 Unsupported GPU Assignment Configurations

Scenario	GPU Assignment	Compute Node	Notes
Split Assignment	1x Cage 1 (GPU 1) + 1x Cage 2 (GPU 3)	Node 1	Not Supported
Split Assignment	1x Cage 1 (GPU 2) + 1x Cage 2 (GPU 4)	Node 2	Not Supported

STEP 3 ORDER CISCO UCS X9516 X-FABRIC MODULES

The Cisco UCS X580p connectivity to the Cisco UCS X-Series compute node is enabled with the X-Fabric Module. When a compute node is inserted into the chassis, the compute node's mezzanine card plugs directly into the two Fabric Module slots (with no midplane) for PCIe connectivity to the Cisco UCS X580p PCIe Node.

Select X-Fabric Module (Required)

Select X-Fabric Modules from the [Table 8](#).

Table 8 Base PID of X-Fabric Modules

Product ID (PID) ¹	Description
UCSX-FS-X9516	UCS X9516 X-Fabric PCIe Gen5 switch module for 9508 chassis

Notes:

1. The X-Fabric modules are required on the X9508 chassis

Select X-Fabric Module Adapter (Optional)

Select X-Fabric Modules adapter from the [Table 9](#).

Table 9 X-Fabric Modules Adapter

Product ID (PID) ^{1,2}	Description	Slot compatibility
UCSX-P-N7S400GFO	NVIDIA OEM MCX715105AS-WEAT 1x400GbE QSFP112 PCIe Gen5 NIC	Slot 1 & 2
UCSX-P-N7D200GFO	NVIDIA OEM MCX755106AS-HEAT 2x200GbE QSFP112 PCIe Gen5 NIC	Slot 1 & 2

Notes:

1. The X-Fabric modules are required on the X9508 chassis
2. The UCS X440p is not compatible with the UCS X9516 and the UCS X580p is not compatible with the UCS 9416 modules

STEP 4 SELECT OPTICS OPTIONS (REQUIRED)**Select Optic Cables**

The recommended Cisco Optics for the server are listed in [Table 10](#).

Table 10 Supported Optic Cables¹

Product ID (PID)	Description
Optics	
QSFP-200G-SL4	200GBASE SL4 QSFP56 Transceiver, MPO, 30m over OM4 MMF
QSFP-200G-SR4-S	200GBASE SR4 QSFP56 Transceiver, MPO, 100m over OM4 MMF
QSFP-400G-DR4	400G QSFP112 Transceiver, 400GBASE-DR4, MPO-12, 500m parallel
Single Mode Patch Cable	
CB-M12-M12-SMF10M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, SMF, 10M
CB-M12-M12-SMF15M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, SMF, 15M
CB-M12-M12-SMF1M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, SMF, 1M
CB-M12-M12-SMF20M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, SMF, 20M
CB-M12-M12-SMF25M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, SMF, 25M
CB-M12-M12-SMF2M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, SMF, 2M
CB-M12-M12-SMF30M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, SMF, 30M
CB-M12-M12-SMF3M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, SMF, 3M
CB-M12-M12-SMF5M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, SMF, 5M
CB-M12-M12-SMF7M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, SMF, 7M
Multi-mode Patch Cable	
CB-M12-M12-MMF1.5	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, MMF, 1.5M
CB-M12-M12-MMF10M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, MMF, 10M
CB-M12-M12-MMF15M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, MMF, 15M
CB-M12-M12-MMF1M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, MMF, 1M
CB-M12-M12-MMF20M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, MMF, 20M
CB-M12-M12-MMF25M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, MMF, 25M
CB-M12-M12-MMF2M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, MMF, 2M
CB-M12-M12-MMF30M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, MMF, 30M
CB-M12-M12-MMF3M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, MMF, 3M
CB-M12-M12-MMF4M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, MMF, 4M
CB-M12-M12-MMF5M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, MMF, 5M
CB-M12-M12-MMF7M	CABLE, MPO12-MPO12, TRUNK CABLE, TYPE B, MMF, 7M


Notes:

1. Please make sure to refer TMG Compatability matric tool <https://tmgmatrix.cisco.com/home> to determine what transceivers are supported with the adapters selected on UCS 9516. The above list is the current supported list but its subject to change.

SUPPLEMENTAL MATERIAL

System Board

Figure 3 A Top View Of The Cisco UCS X580p PCIe Node System Board

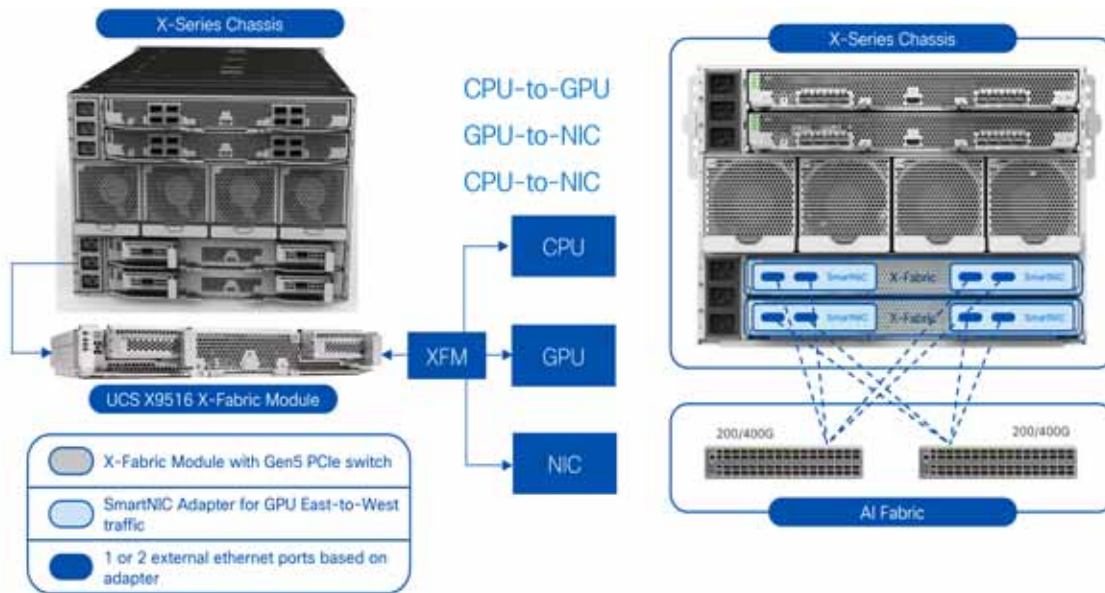


1	Cage slot 1 Supports all type of GPUs.	2	Cage slot 2 Supports all type of GPUs.
---	---	---	---

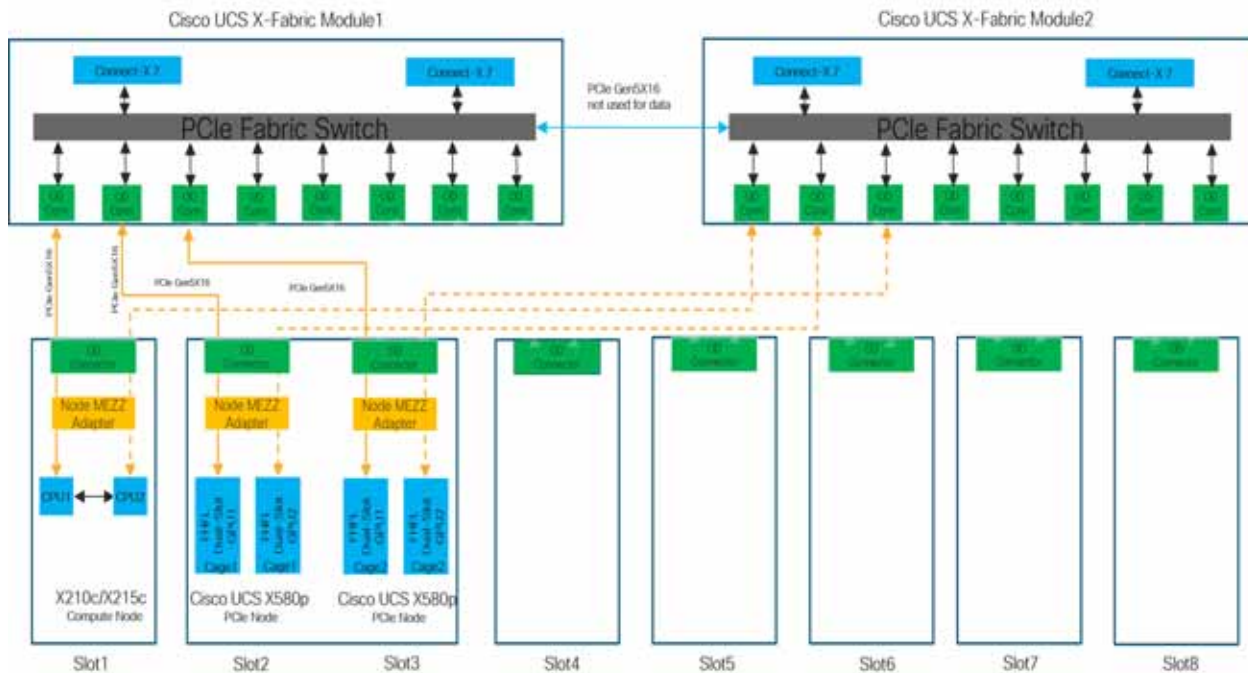
3	Cage slot 1 Supports all type of GPUs.	4	Cage slot 2 Supports all type of GPUs.
5	mezzanine connector (included)	-	-

Solution Topology

Figure 4 Solution Topology




PCIe Connectivity



X580p PCIe Node Configuration - 2 Server 4x GPU + 2 NIC each

Figure 5 Two servers with 4x GPU + 2 NIC each



X580p PCIe Node Configuration - 2 Server 4x GPU + 2 NIC each Chassis PCIe Mapping Policy

Figure 6 Two servers with 4x GPU + 2 NIC each


Two servers with 4 GPU + 2 NICs each

Server 1	CPU-1	GPU-3/1	GPU-4/1	NIC-1/2	Server 5	CPU-1	GPU-3/1	GPU-4/1	NIC-1/1
		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>			<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	
Server 2	CPU-2	GPU-3/2	GPU-4/2	NIC-2/2	Server 6	CPU-2	GPU-3/2	GPU-4/2	NIC-2/1
		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>			<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Server 2	CPU-1	GPU-3/1	GPU-4/1	NIC-1/2	Server 6	CPU-1	GPU-3/1	GPU-4/1	NIC-1/1
		<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Server 2	CPU-2	GPU-3/2	GPU-4/2	NIC-2/2	Server 6	CPU-2	GPU-3/2	GPU-4/2	NIC-2/1
		<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

X580p PCIe Node Configuration

- 1 Server 4x GPU + 2 NIC each
- 2 Server 2 GPU on 1 CPU + 1 NIC

Figure 7 One servers with 4x GPU + 2 NIC each & Two servers with 2x GPU + 1 NIC



X580p PCIe Node Configuration - PCIe Mapping Policy

- 1 Server 4x GPU + 2 NIC each
- 2 Server 2x GPU on 1 CPU

Table 11 One server with 4x GPU + 2 NIC each & 2 server with 2x GPU + 1 NIC

X580P 1: One Server, 4 GPU + 2 NIC

X580P 2: Two server, 2 GPU on 1 CPU + 1 NIC (AI use case)

Server	CPU	GPU-3/1	GPU-4/1	NIC-1/2	Server	CPU	GPU-3/1	GPU-4/1	NIC-1/1
		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>			<input checked="" type="checkbox"/>	Server 5	CPU-1
Server 1	CPU-1	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Server 5	CPU-2	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
	CPU-2	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		Server 6	CPU-1	<input type="checkbox"/>	<input type="checkbox"/>
Server 2	CPU-1	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Server 6		CPU-2	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	CPU-2	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>					

Spare Parts

This section lists the upgrade and service-related parts for the Cisco UCS X580p PCIe Node.

Table 12 Spare Parts

Product ID (PID)	PID Description
X-Fabric Module	
UCSX-F-X9516=	UCS X9516 X-Fabric module for 9508 chassis
GPU Cards	
UCSX-GPU-L40S=	NVIDIA L40S: 350W, 48GB, 2-slot FHFL GPU
UCSX-GPU-H200-NVL=	NVIDIA H200 NVL: 600W, 141GB, 2-Slot NVL2 FHFL GPU
UCSX-GPU-RTXP6000=	NVIDIA RTX Pro 6000: 600W, 96GB, 2-Slot FHFL GPU
CBL-X580P-GPU-N=	UCS X580P NVIDIA GPU Power Cable
NVL Bridge	
UCSX-NVL2-H200=	NVIDIA NVL-2way Bridge for H200 GPU; NVPN 900-23945-0000-000
UCSX-GPU-RKIT-NV=	UCSX GPU/NVLINK Bridge Removal Tool Kit NVIDIA
PCI Mezz card	
UCSX-V5-PCIME ¹	UCS PCI Mezz card for X-Fabric Gen5

Notes:

1. If UCSX-580P is selected in the 9508 Chassis, then a quantity of 1 PID (UCSX-V5-PCIME) is required for each server node mapped to a UCSX-580P. For example, if there are 1 or 2 UCSX-580P units and 2 to 4 server nodes, each server node must include a UCSX-V5-PCIME to access the GPUs.

TECHNICAL SPECIFICATIONS

Dimensions and Weight

Table 13 Cisco UCS X580p PCIe Node Dimensions and Weight

Parameter	Value
Height	3.7 in. (94 mm)
Width	11.28 in.(286.5 mm)
Depth	24 in. (602 mm)
Weight	<ul style="list-style-type: none"> ■ Minimally configured node weight = 18.51 lbs (8.4 kg) ■ Fully loaded PCIe Node with L40S GPU = 28.66 lb (13 kg); minimum config with 1x L40S GPU = 21.05 lb (9.55 kg) ■ Fully loaded PCIe Node with RTXPRO 6000 GPU = 29.98 lb (13.6 kg); minimum config with 1X RTXPRO 6000 GPU = 21.38 lb (9.7 kg) ■ Fully loaded PCIe Node with H200-NVL GPU (with 2x NVL bridge cards) = 30.42 lb (13.8 kg); minimum config with 1X H200-NVL GPU = 21.43 lb (9.72 kg)

Table 14 Cisco UCS X9516 X-Fabric module Dimensions and Weight

Parameter	Value
Height	1.80 in. (45.7 mm)
Width	11.28 in.(286.5 mm)
Depth	24 in. (602 mm)
Weight	<ul style="list-style-type: none"> ■ Minimally configured node weight = 12.84 lbs (5.83 kg) ■ Fully loaded X-Fabric Module with 2x Connect-X 7 Adapters = 14.9 lb; minimum config with 1x Connect-X 7 = 12.9 lb

Environmental Specifications

Table 15 Cisco UCS X580p PCIe Node and Cisco UCS X9516 X-Fabric Environmental Specifications

Parameter	Value
Operating temperature	50° to 81° F (10° to 27° 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 UCS Power Calculator at:

<http://ucspowercalc.cisco.com>



NOTE: The ambient temperature must be less than 27 °C (80.6 °F).



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.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

