



Cisco Compute Hyperconverged 440p 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/hyperconverged-infrastructure/compute-hyperconverged/datasheet-listing.html>

OVERVIEW	3
DETAILED VIEWS	4
Cisco UCS X440p PCIe Node Front View	4
PCIe Node STANDARD CAPABILITIES and FEATURES	5
CONFIGURING the Cisco UCS X440p PCIe Node	7
STEP 1 CHOOSE BASE CISCO UCS X440p PCIe Node SKU	8
STEP 2 SELECT RISER CARDS (REQUIRED)	9
STEP 3 ORDER GPU CARDS	10
STEP 4 ORDER CISCO UCS X9416 X-FABRIC MODULES	11
STEP 5 CHOOSE REAR MEZZANINE VIC/BRIDGE ADAPTERS	12
Cisco UCS X440p PCIe Node under UCSX-M6-MLB:	12
Cisco UCS X440p PCIe Node under UCSX-M7-MLB:	12
SUPPLEMENTAL MATERIAL	14
Simplified Block Diagram	14
System Board	15
SPARE PARTS	16
TECHNICAL SPECIFICATIONS	17
Dimensions and Weight	17
Environmental Specifications	17

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 440p PCIe Node is a new node type that is now supported in the Cisco Compute Hyperconverged 9508 chassis. This can be attached to Compute Hyperconverged node in the Cisco Compute Hyperconverged 9508 chassis to provide GPU accelerators support using the 9416 X-Fabric modules for Cisco Compute Hyperconverged 9508 chassis.

The Cisco Compute Hyperconverged 440p PCIe Node is the first PCIe resource node to integrate into the Cisco Compute Hyperconverged X-Series Modular System. Up to four PCIe Nodes can reside in the 7-Rack-Unit (7RU) Cisco Compute Hyperconverged 9508 Chassis and can be paired with one compute node each, offering up to four GPUs to a Cisco Compute Hyperconverged node with Cisco X-Fabric Technology.

The X-Fabric Technology solution is a combination of two products: the Cisco 9416 X-Fabric Module which provides a PCIe Gen 4 fabric and the Cisco Compute Hyperconverged 440p PCIe Node which hosts the GPUs.

The Cisco Compute Hyperconverged 9508 Chassis has eight node slots, up to four of which can be 440p PCIe Nodes when paired with a Cisco Compute Hyperconverged node. This provides up to 16 GPUs per chassis to accelerate your applications. If your application needs even more GPU acceleration, up to two additional GPUs can be added on each compute node using optional GPU front mezz on X210c and X410c compute node

Cisco Compute Hyperconverged 440p PCIe Node supports several GPUs please refer to [STEP 3 ORDER GPU CARDS, page 9](#) for the available GPUs

Figure 1 Front views of Cisco Compute Hyperconverged 440p PCIe Node

Front View



Rear View




DETAILED VIEWS

Cisco Compute Hyperconverged 440p PCIe Node Front View

Figure 2 is a front view of the Cisco Compute Hyperconverged 440p PCIe Node.

Figure 2 Cisco Compute Hyperconverged 440p PCIe Node Front View

GPUs Option




1	Locate LED & Status LED	3	PCI Node Ejector Button
2	PCI Node Ejector Handles	-	-

PCIe Node STANDARD CAPABILITIES and FEATURES

Table 1 lists the capabilities and features of the base Cisco Compute Hyperconverged 440p PCIe Node. Details about how to configure the PCIe Node for a listed feature or capability (for example, number of processors, disk drives, or amount of memory) are provided in [CONFIGURING the Cisco X440p PCIe Node on page 6](#).

Table 1 Capabilities and Features

Capability/Feature	Description
Chassis	The Cisco Compute Hyperconverged 440p PCIe Node mounts in a Cisco Compute Hyperconverged 9508 chassis.
GPU slots	<ul style="list-style-type: none"> ■ Riser Type A (1 PCIe slots) for 1x dual slot GPU per riser ■ Riser Type B (2 PCIe slots) for 2x single slot GPU per riser  <p>Note: Not all risers are available in every server configuration option.</p>
Available GPUs	<ul style="list-style-type: none"> ■ NVIDIA A16 PCIE 250W 4X16GB ■ NVIDIA L4 Tensor Core, 70W, 24GB ■ NVIDIA L40S: 350W, 48GB, 2-slot FHFL GPU
Power subsystem	Power is supplied from the Cisco Compute Hyperconverged 9508 chassis power supplies. The Cisco Compute Hyperconverged 440p PCIe 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 440p 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"> ■ Cisco Intersight software (SaaS, Virtual Appliance and Private Virtual Appliance)
Chassis	Compatible with the Cisco Compute Hyperconverged 9508 X-Series Server Chassis

CONFIGURING the Cisco Compute Hyperconverged 440p PCIe Node

Follow these steps to configure the Cisco Compute Hyperconverged 440p PCIe Node:

- *STEP 1 CHOOSE BASE CISCO X440P PCIE NODE SKU, page 7*
- *STEP 2 SELECT RISER CARDS (REQUIRED), page 8*
- *STEP 3 ORDER GPU CARDS, page 9*
- *STEP 4 ORDER CISCO 9416 X-FABRIC MODULES, page 10*
- *STEP 5 CHOOSE REAR MEZZANINE VIC/BRIDGE ADAPTERS, page 11*

STEP 1 CHOOSE BASE Cisco Compute Hyperconverged 440p PCIe Node SKU

Verify the product ID (PID) of the Cisco Compute Hyperconverged 440p PCIe Node as shown in [Table 2](#).

Table 2 PIDs of the Base Cisco Compute Hyperconverged 440p PCIe Node

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

A base Cisco Gen4 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
- Riser Cards
- Cisco 9416 X-Fabric Modules

STEP 2 SELECT RISER CARDS (REQUIRED)

Select risers from [Table 3](#).

Table 3 PIDs of the Risers

Product ID (PID)	Description
HCIX-RIS-A-440P	Riser A for 1x dual slot GPU per riser, 440P PCIe Node <ul style="list-style-type: none">■ Riser1A (controlled with CPU1 on Hyperconverged Node)■ Riser2A (controlled with CPU2 on Hyperconverged Node)
HCIX-RIS-B-440P	Riser B for 2x single slot GPUs per riser, 440P PCIe Node <ul style="list-style-type: none">■ Riser1B (controlled with CPU1 on Hyperconverged Node)■ Riser2B (controlled with CPU2 on Hyperconverged Node)



NOTE: The PCIe Node requires both the risers to be configured and doesn't support orderability without both risers included. Riser cards include all required power cables for supported GPUs.

STEP 3 ORDER GPU CARDS

Select GPU Options

The available GPU PCIe options and their riser slot compatibilities are listed in [Table 4](#).

Table 4 Available PCIe GPU Cards

GPU Product ID (PID)	PID Description	Riser Slot Compatibility	Maximum Number of GPUs Per node
HCIX-GPU-A16	NVIDIA A16 PCIE 250W 4X16GB	Riser 1A (Gen 4), Riser 2A (Gen 4)	2
HCIX-GPU-L4	NVIDIA L4 Tensor Core, 70W, 24GB	Riser 1B (Gen 4), Riser 2B (Gen 4)	4
HCIX-GPU-L40S	NVIDIA L40S: 350W, 48GB, 2-slot FHFL GPU	Riser 1A (Gen 4), Riser 2A (Gen 4)	2

Caveats

Riser cards and GPUs cannot be mixed.



NOTE: Following [Step 4](#) and [Step 5](#) are optional only if the Cisco Compute Hyperconverged 9508 Chassis already has the 9416 X-Fabric modules installed and the Compute Hyperconverged node has one of the supported mezzanine adapters to connect to Cisco Compute Hyperconverged 440p PCIe Node

STEP 4 ORDER CISCO 9416 X-FABRIC MODULES

The Cisco Compute Hyperconverged 440p connectivity to the Cisco Compute Hyperconverged 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 Compute Hyperconverged 440p PCIe Node.

Select X-Fabric Modules on the Cisco Compute Hyperconverged 9508 chassis [Table 5](#).

Table 5 PIDs of the Risers

Product ID (PID) ¹	Description
HCIX-F-9416	9416 X-Fabric module for 9508 chassis

Notes:

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

STEP 5 CHOOSE REAR MEZZANINE VIC/BRIDGE ADAPTERS

The Cisco Compute Hyperconverged 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 compute node for network connectivity or as a connector to the 440p PCIe Node via X-Fabric modules. The same mezzanine slot on the compute node can also accommodate a pass-through mezzanine adapter for X-Fabric which enables compute node connectivity to the 440p PCIe Node. Refer to [Table 6](#) for supported adapters.

Table 6 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	Cisco VIC 15422 mezzanine adapter for Cisco Compute Hyperconverged Node	2 CPUs required	Rear Mezzanine connector on motherboard
Cisco VIC Bridge Card²			
HCIX-V5-BRIDGE	VIC 15000 bridge to connect mLOM and mezz X Compute Node (This bridge to connect the Cisco VIC 15420 mLOM and Cisco VIC 15422 Mezz for the Cisco Compute Hyperconverged Node)	2 CPUs required	One connector on Mezz card and one connector on mLOM card

Notes:

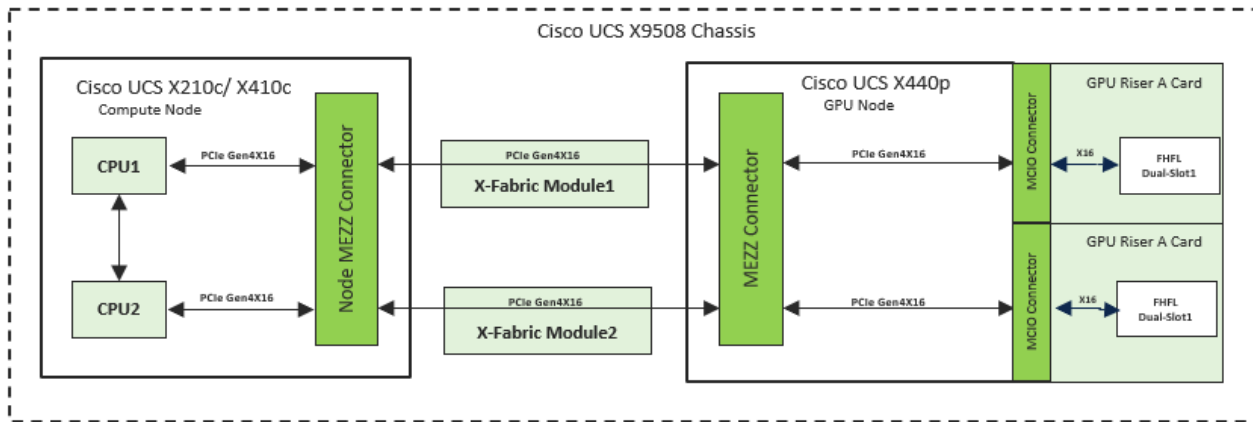
1. If this adapter is selected, then two CPUs are required and HCIX-ME-V5Q50G or HCIX-V4-PCIME is required.
2. Included with the Cisco VIC 15422 mezzanine adapter.

SUPPLEMENTAL MATERIAL

Simplified Block Diagram

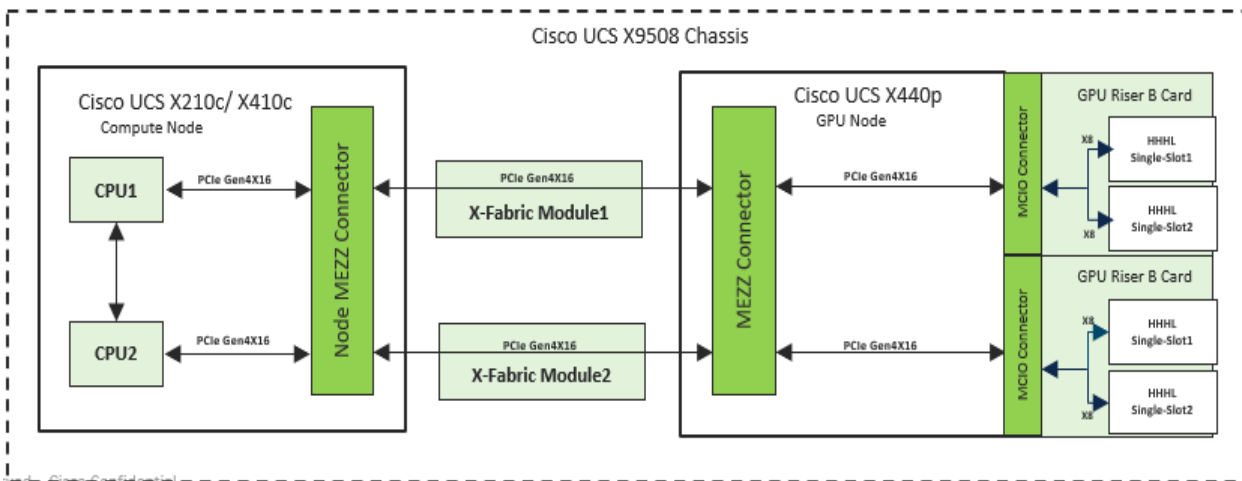
A simplified block diagram of the Cisco Compute Hyperconverged 440p PCIe Node system board is shown in [Figure 3](#) with riser A.

Figure 3 Cisco Compute Hyperconverged 440p PCIe Node Simplified Block Diagram with Riser A



A simplified block diagram of the Cisco Compute Hyperconverged 440p PCIe Node system board is shown in [Figure 4](#) with riser B.


Figure 4 Cisco Compute Hyperconverged 440p PCIe Node Simplified Block Diagram with Riser B



System Board

A top view of the Cisco Compute Hyperconverged 440p PCIe Node system board is shown in [Figure 5](#).

Figure 5 Cisco Compute Hyperconverged 440p PCIe Node System Board



1	Riser slot 1 Supports both Type A and Type B risers.	2	Riser slot 2 Supports both Type A and Type B risers.
3	GPU slot 1 (FHFL GPU shown) Supports either FHFL or HHHL GPU depending on the riser type.	4	GPU slot 2 (FHFL GPU shown) Supports either FHFL or HHHL GPU depending on the riser type.
5	mezzanine connector (included)	-	-

TECHNICAL SPECIFICATIONS

Dimensions and Weight

Table 7 Cisco Compute Hyperconverged 440p PCIe Node 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 PCIe Node with T4 GPU = 14.9 lb; minimum config with 1x T4 GPU = 12.9 lb ■ Fully loaded PCIe Node with A16 GPU = 17.1 lb; minimum config with 1X A16 GPU = 14.6 lb ■ Fully loaded PCIe Node with A40 GPU = 16.6 lb; minimum config with 1X A40 GPU = 14.4 lb ■ Fully loaded PCIe Node with A100 GPU = 17.9 lb; minimum config with 1X A100 GPU = 15 lb

Environmental Specifications

Table 8 Cisco Compute Hyperconverged 440p PCIe 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 440p PCIe Node has a power cap of 1300 Watts for all combinations of components. Also, the ambient temperature must be less than 35 °C (95 °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)

