

Cisco UCS X580p PCIe Node



Benefits

- Provides PCIe Gen 5 for fast connectivity
- Add up to four AMD or NVIDIA GPUs to Cisco UCS® X-Series Compute Nodes for AI/ML, VDI, and other GPU-intensive workloads
- Supports NVIDIA NVLink for larger workloads

Learn more

For more information about modernizing your infrastructure with Cisco UCS X-Fabric Technology and the Cisco UCS X580p PCIe Node, go to <https://cisco.com/go/ucsx>.

For all Cisco UCS servers, please visit <https://www.cisco.com/go/ucs>.

Product overview

Applications such as accelerated AI/ML, VDI, rendering, and visualization perform best with GPUs. In the past, higher performing GPUs necessitated a rack-server form factor because blade servers had limited GPU support. This created infrastructure silos based on applications. Not with the Cisco UCS X-Series Modular System.

The Cisco UCS X580p PCIe Node allows you to add up to four GPUs to a Cisco UCS X210c and X215c Compute Nodes with Cisco UCS X-Fabric Technology. Now you can easily and independently manage the different lifecycles of CPU and GPU components.



Figure 1. Cisco UCS X580p PCIe Node

What it does

A Cisco UCS X580p PCIe Node hosts GPUs and connects to the Cisco UCS X9516 X-Fabric Module, which extends the PCIe Gen 5 fabric to the compute node. This combination of the two products is the second-generation Cisco UCS X-Fabric Technology solution.

The Cisco UCS X9508 Chassis has eight node slots, up to two of which can be X580p PCIe Nodes. The GPUs in the X580p PCIe node can be shared by two servers. 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 using the optional GPU front mezzanine on the Cisco UCS X210c or Cisco UCS X215c compute nodes.

Supported GPUs include options from NVIDIA and AMD which can be shared across multiple compute nodes but note that only one GPU model may be used in a single Cisco UCS Compute Node.