

# Reach for the Edge with Cisco UCS X-Series Direct

Everything and everyone are distributed

The edge is where the action is. It's where valuable Internet-of-Things (IoT) data is generated, where AI inferencing can shape customer experiences, where robotic manufacturing takes place, and where 5G networks must deliver near-instant response times to requests from self-driving cars and dynamic applications. The edge is where employees live and—increasingly—work, needing seamless access to IT resources. Because all this activity takes place outside of the domain of enterprise data centers, applications and data need to follow. By mid-decade, the computing and storage resources at the edge could

equal those in data centers. International Data Corporation (IDC) also predicts a major shift in spending, with the edge making up 28 percent of IT spending by 2026.<sup>1</sup>

The edge presents challenges. You need excellent visibility over operations to minimize cost and risk. You need specialized systems that can support GPU-accelerated operations, big data analysis, and straightforward applications like database management systems. Ideally, all of these needs are supported in a single system so that you can deploy without the time, cost, and risk of do-it-yourself integration.

That's why we created Cisco UCS X-Series Direct

We designed the Cisco UCS® X-Series Modular System to be ready for the future, to accommodate a wide set of diverse


applications, and to be globally managed by the Cisco Intersight® IT operations platform or the time-proven Cisco UCS Manager. The Cisco

Standalone Cisco UCS  
X-Series Modular System  
for any location, anywhere

- Data center in a box with built-in fabric interconnects with 100-Gbps unified fabric connectivity throughout
- Deploy in any location—data center, medium-size business, remote, and edge locations
- Adapt to support any workload, from online transaction processing to AI inferencing

1. IDC Worldwide Edge Enterprise Infrastructure Forecast, 2022–2026: New Demands as Edge Supports Digital-First Business, document US49617922 2022

## Flexible configurations to support all your workloads



Cisco UCS X-Series Modular System

UCS X-Series is a next-generation platform designed to take our customers into the next decade of computing as our original blade server chassis has done for more than ten years.

The Cisco UCS X-Series is designed to shift your thinking from administrative details to business outcomes—with infrastructure that is assembled from the cloud, shaped to your workloads, and continuously optimized. Now that hardware can think and change like software, you are free to think like tomorrow. This brings about immense simplification to edge applications:

- **Simplify with cloud-operated infrastructure:** By shaping system resources to workload requirements with Cisco Intersight, you can deliver intelligent

### Self-contained modular solution

Cisco UCS X-Series Direct transforms the Cisco UCS X-Series Modular System into a self-contained system with pair of internal fabric interconnects that integrate each of up to eight server nodes with Cisco® unified fabric connectivity, upstream IP and Fibre Channel connectivity, all managed through Cisco Intersight or the time-proven Cisco UCS Manager.

visualization, optimization, and orchestration to all your applications and infrastructure.

- **Simplify with an adaptable platform:** Edge applications have widely varying requirements. At the edge, you can use a single platform that combines the density and efficiency of blade servers with the expandability of rack servers for better performance, automation, and efficiency.
- **Simplify with a system engineered for the future:** Embrace emerging technology and reduce risk with a modular system engineered to support future generations of processor, accelerator, and interconnects with management enabled by a constant stream of new capabilities delivered as software-as-a-service (SaaS).

The platform supports two- and four-socket server nodes of which up to eight and four (respectively) can populate a chassis. Each node can support up to six NVMe drives, two M.2 drives, and up to two half-height, half-length GPU accelerators. The platform's capabilities can be extended through the Cisco UCS X-Fabric.

## Cisco UCS X-Fabric

We named it so because X is the variable. With a midplane-free physical design, the Cisco UCS X-Fabric can augment each blade's capabilities by plugging into a connector which, today, extends each server's PCIe connectivity to optional PCIe nodes. These can support up to two full-height, full-length GPU accelerators or four half-height, half-length devices. We envision the X-Fabric evolving over time, with future capabilities including Compute Express Link (CXL)-connected memory.

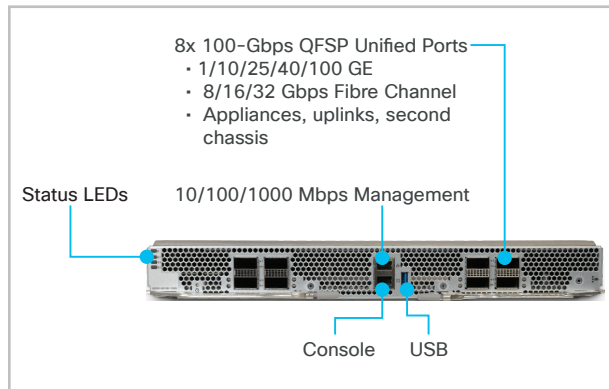


Figure 1. Cisco UCS Fabric Interconnect 9108 100G ports

## Designed to support your use cases

Cisco UCS X-Series Direct can be configured to support all of your edge and remote-office use cases today, with the capability to incorporate future technologies so that you can avoid forklift upgrades. You can configure servers, GPU accelerators, networking, and storage to propel use cases including:

- **AI/ML inferencing**, with up to six compact GPUs per server, or up to two dual-slot, full-length, full-height GPUs
- **Data protection** onto NVMe storage to guard mission-critical data and enable rapid restores in the event of a failure or ransomware event
- **Hyperconverged infrastructure** that requires large amounts of local storage
- **Virtual desktop infrastructure (VDI)** that you can deploy close to users, and ensure a smooth user experience with GPU rendering

## Solution Architecture

In the data center, X-Series modular systems typically use intelligent fabric modules to pass IP, storage, and management traffic up to a pair of fabric interconnects that integrate a number

- **Cloud-native, containerized environments** with support from the Red Hat OpenShift Container Platform
- **Virtualized environments** such as those supported by VMware vSphere 8.0
- **Converged infrastructure** solutions that integrate enterprise shared storage from our leading storage partners

One of the benefits of choosing Cisco® solutions is that our engineering team works with partners to create Cisco Validated Designs that can guide your deployment in a cookbook-like fashion. These designs help reduce cost and risk by eliminating the cost of do-it-yourself integration. They also help reduce time-to-value because we select solution components from best-of-breed products from our software and storage partners, eliminating the delay of evaluating and undertaking test integrations of various options.

of chassis and rack servers into a single system. Cisco UCS X-Series Direct replaces the pair of intelligent fabric modules with a pair of Cisco UCS Fabric Interconnects 9108 100G (Figure 1).

## Architectural Options

Cisco UCS X-Direct architectural options include connectivity to data center switches and SANs, direct Fibre Channel connectivity to storage, and to a second chassis (future capability)




Figure 2. Cisco UCS X-Direct architectural options

## IP, storage, and management connectivity

The fabric interconnects create a self-contained system connecting IP, storage, and management networks with two 100-Gbps connections to each of eight slots in the chassis. Within each compute node, Cisco virtual interface cards provide the number and type of I/O devices that you create depending on your application.

The fabric interconnects connect upstream to switches or security appliances with eight 100-Gbps unified ports that can carry 1, 10, 25, 40, and 100 Gigabit Ethernet to upstream switches (Figure 2). It can translate internal Fibre Channel over Ethernet (FCoE) to 8-, 16-, or 32-Gbps Fibre Channel storage that can connect directly to storage devices or to a SAN with Cisco Multilayer Director Switch (Cisco MDS) switches. In the future, these ports can be used to connect to a second Cisco UCS X9508 chassis equipped with intelligent fabric modules, with the primary chassis serving as a single point of connectivity and management for the edge location.

## Flexible expansion with Cisco UCS X-Fabric connectivity

With the addition of a pair of Cisco UCS X9416 X-Fabric Modules to the rear of the chassis, you can add 32 lanes of 4th Gen PCIe connectivity between each server and a Cisco UCS X440p PCIe Node that can contain up to four NVIDIA

or Intel GPU accelerators. Now, without even opening a server, you can slide in a module to propel AI/ML and VDI workloads.

## Powerful Cisco servers

The X-Series chassis has 8 slots that can hold single-width, 2-socket servers, double-width 4-socket servers, and X-Fabric modules:

- Cisco UCS X210c M6 and M7 Compute Nodes:** These can be equipped with two 3rd, 4th, or 5th Gen Intel® Xeon® processors (respectively), up to 8 TB of main memory, two M.2 drives, and up to 6 SAS, SATA, or NVMe drives. An optional mezzanine module supports up to two NVMe drives and up to two GPU accelerators.
- Cisco UCS X410c M7 Compute Nodes:** These are double-width servers that host up to four 4th Gen Intel Xeon processors with up to 60 cores per processor and up to 16 TB of main memory. These nodes enable your modular system to support the most challenging workloads, including performance to support SAP HANA and entire medical centers running EPIC Electronic Health Record (EHR) software. This compute node offers the same storage options as the 2-socket X210c node: Two M.2 drives, and up to 6 SAS, SATA, or NVMe drives.

## Flexible management options

One of the challenges of deploying computing and storage capabilities to the edge is being able to replicate installations at scale with a minimum of hands-on interaction. The Cisco UCS X-Series simplifies deployment with a modular architecture that reduces cabling and configuration overhead along with the points of failure that they can pose.

### Cisco Intersight

The other challenge is managing edge locations as if they are local, and this is where the Cisco Intersight IT operations platform comes in. Cisco Intersight is designed to operate Cisco UCS X-Series Direct from the cloud, so you can manage your infrastructure from initial provisioning through ongoing maintenance from a single software-as-a-service interface that is always up to date and always supports the latest platform features. Cisco Intersight is designed to connect to each fabric interconnect you deploy globally through secure connections negotiated

between the interconnects and the cloud-based software, eliminating an entire class of security concerns. If you prefer to run Cisco Intersight on site, we offer private and virtual appliance options.

All of your edge locations appear in the dashboard, so you have central control over distributed resources. Global policies enforce consistent deployments and eliminate configuration drift. You can monitor infrastructure health with real-time and historical data to troubleshoot problems, forecast and budget resources, and find intrinsic correlations in dispersed data. Cisco Intersight is directly connected to the Cisco Technical Assistance Center (Cisco TAC) so anomalies can be reported and service requests prepared even before you might notice a problem developing. The software's energy management features enable you to monitor and track power consumption and energy use through power

policies that apply to individual nodes and chassis.

### Cisco UCS Manager

Cisco UCS manager is the traditional role- and policy-based management platform embedded in Cisco UCS fabric interconnects from the very beginning. Organizations that wish to continue to use this familiar interface to manage their X-Series Direct systems can use the software's built-in graphical user interface or application programming interface (API) to deploy resources with automation and policy-based service profiles that have been hallmark features of Cisco UCS since its initial offering in 2009. Cisco UCS Manager can be used remotely over secure virtual private networks to edge locations. For a global view of inventory, policy, and operations, Cisco UCS Central software can aggregate multiple Cisco UCS Manager-based systems and apply global policies to UCS Manager instances.

## Why choose Cisco?

Whether yours is a small business or a global enterprise, you know that automation is the key to consistent and accurate lifecycle

management. From the very beginning we designed Cisco Unified Computing System™ to make it as easy to deploy a thousand servers

as it is to deploy a single one. You determine policies that enforce compliance with your standards, and then through automation your

## SEAL Sustainable Product of the Year Award for 2023



The Cisco UCS X-Series Modular System has achieved [sustainable product of the year](#) status as a system that is purpose built for a sustainable future. We

designed the X-Series with best-in-class energy efficiency in mind, helping balance performance needs with new sustainability demands on today's data centers. In a scenario where UCS X-Series is replacing 64 previous-generation servers:

- Customers can use 3.3x less hardware overall, saving precious rack space.
- They can also reduce almost 100,000 kilowatt hours (KwH) of energy, or the equivalent of powering 10 residential homes for a full year.
- The result is that we can help save almost 40 tons of carbon dioxide (tCO<sub>2</sub>e) emissions.

deployments—in the data center, remote or branch locations, or at the edge—are executed without the time-consuming, error-prone, manual configuration and multiple management interfaces of the past. A single management interface, whether Cisco Intersight or Cisco UCS Manager, can help you reach for the edge with minimal time and risk.

We build our infrastructure for longevity and to prepare customers for the future. Our original blade-server chassis, the Cisco UCS 5108, is still in operation at many sites and has served customers for six generations of computing technology. The Cisco UCS Mini configuration similarly put fabric interconnects into the chassis for remote and edge operations, just as Cisco UCS X-Series Direct does today.

Our experience and our approach of designing infrastructure that endures the test of time has created this platform designed to take you in to the next decade. Based on customer case studies and industry benchmarks, compared to our prior solutions, Cisco UCS X-Series Direct can yield up to:

- **68% CapEx savings** and 78 percent less cabling due to the Cisco unified fabric
- **64% better performance** than our previous-generation servers
- **52% lower power consumption** after updating to a Cisco UCS solution

Sustainability is important, especially at the edge where power and cooling resources are at a premium. Sustainability begins with a product that is designed for longevity and reuse, so it can grow and change with your requirements without having to waste existing chassis, power supplies, and cooling fans. This results in a product that consumes roughly 50 percent less raw material over the course of three generations compared to using traditional rack servers. This flexibility and reusability earned the Cisco UCS X-Series the [Silverlining Best Cloud Sustainability Initiative](#) award for 2023.

With shared power supplies, zoned cooling, and high-voltage power distribution within the chassis, the Cisco UCS X-Series can help reduce resource consumption, which is why it has achieved recognition through the [SEAL award](#) for its outstanding innovations.

When you choose Cisco when you reach to the edge, you can be assured of excellent performance, flexibility to support an ever wider range of workloads, and a path that will support you today and into the future.

### Learn more

- [X-Series Direct data sheet](#)
- [X-Series Direct at-a-glance](#)
- [X-Series Direct FAQ](#)