

Cisco UCS Server Hardware Architecture Guide - Official Technical Overview & Hardware Datasheet

CISCO UCS SERVER HARDWARE ARCHITECTURE GUIDE

OFFICIAL TECHNICAL OVERVIEW & HARDWARE DATASHEET

EXECUTIVE SUMMARY

The Cisco Unified Computing System (UCS) represents a paradigm shift in data center architecture, unifying computing, networking, management, virtualization, and storage access into a single, cohesive, and highly scalable system. This document provides a comprehensive technical overview of the Cisco UCS hardware architecture, detailing the foundational components, design principles, and key technical specifications that enable a simplified, high-performance, and cost-effective infrastructure.

At its core, Cisco UCS is designed to address the limitations of traditional, siloed IT infrastructure. By integrating network and compute resources with intelligent, policy-based management, Cisco UCS dramatically reduces total cost of ownership (TCO) and increases business agility. The architecture is built around a high-bandwidth, low-latency unified fabric that consolidates LAN, SAN, and management traffic onto a single set of cables, significantly reducing the number of switches, cables, and management points.



ARCHITECTURE & CHASSIS DESIGN

Cisco UCS is a modular and hierarchical system, anchored by a pair of Cisco UCS Fabric Interconnects that serve as the central nervous system for the entire UCS domain . This architecture delivers a single, highly available management domain controlled by the embedded Cisco UCS Manager software.

Core Architectural Components:

- Cisco UCS Fabric Interconnects: These are the core components of the system, providing single points of connectivity and management. Typically deployed as an active-active pair, they integrate all components into a unified management domain. They provide low-latency, line-rate, lossless Ethernet and Fibre Channel over Ethernet (FCoE) connectivity .

- Cisco UCS Fabric Extenders (FEX): These are zero-management, low-power devices that extend the system's connectivity and management plane from the Fabric Interconnects to blade chassis and rack servers. They eliminate the need for top-of-rack switches and blade-resident switches, simplifying the architecture and reducing costs .

- Cisco UCS 5100 Series Blade Server Chassis: A crucial building block, the chassis is a highly simplified device with no management points, acting as a logical extension of the Fabric Interconnects. The 6RU Cisco UCS 5108 chassis can support up to eight half-width, four full-width, or two full-width double-height blade servers, offering flexible bay configurations to match diverse workload requirements .

- Cisco UCS Servers: The system supports a wide range of server form factors, including:

- B-Series Blade Servers: High-density, high-performance blade servers designed for computation-intensive and enterprise application environments .

- C-Series Rack Servers: Standalone rack servers that can be integrated into a UCS domain for unified management, providing flexibility for specialized workloads like AI and high-performance computing (HPC) .

- Cisco UCS Virtual Interface Cards (VICs): These cards extend the network fabric directly to servers and virtual machines. They provide complete programmability of the I/O infrastructure, allowing the number and type of interfaces to be configured on demand .

HARDWARE FEATURES

Cisco UCS hardware is engineered for performance, scalability, and investment protection.

Key Features:

- Unified Fabric & SingleConnect Technology: This technology dramatically simplifies connectivity by allowing all traffic (LAN, SAN, and management) to traverse a single set of cables. This reduces cabling complexity, infrastructure cost per server, and management overhead .
- Centralized, Policy-Based Management: Cisco UCS Manager, embedded in the Fabric Interconnects, provides a single-pane-of-glass management interface. It uses a "create once, deploy many" approach, automating system configuration through policies and service profiles. This abstracts server identity from hardware, enabling stateless computing and rapid bare-metal provisioning .
- High Availability & Redundancy: The system is designed for high availability. Redundant Fabric Interconnects, I/O modules, power supplies, and fans ensure continuous operation. The service profile model also enhances application availability by allowing workloads to be easily moved between servers .
- Modularity & Investment Protection: The architecture is modular, allowing independent upgrades of servers and network components. The blade chassis

is designed with sufficient power and cooling headroom to support multiple future generations of processors and network technologies, protecting the customer's investment .

- High-Performance I/O: The latest Cisco UCS servers leverage PCIe Gen4 and Gen5 technologies, offering immense bandwidth to support demanding applications. For instance, the Cisco UCS VIC 15400 series provides up to 200G connectivity, and the C880A M8 server is designed for massive AI workloads with 8 NVIDIA HGX B300 GPUs .

COMPLIANCE & STANDARDS

Cisco UCS hardware is designed and tested to meet rigorous industry standards and compliance requirements for enterprise and carrier-grade deployments. This ensures interoperability, reliability, and safety. Key certifications typically include:

- Safety & EMC: Certified to international safety standards (e.g., UL, CSA, CE) and electromagnetic compatibility standards (e.g., FCC, VCCI).
- Environmental: Compliant with RoHS and other environmental directives.
- Network & Storage Protocols: Fully supports industry-standard protocols including Ethernet, Fibre Channel, FCoE, and iSCSI, ensuring seamless integration with existing data center infrastructure.
- Management: Adheres to standards like Intelligent Platform Management

Interface (IPMI) for hardware monitoring and management .

TECHNICAL SPECIFICATIONS

Parameter	Specification
System Architecture	Unified Fabric, Centralized Management (UCS Manager), Stateless Computing
Key Components	Fabric Interconnects, Fabric Extenders, Blade/Rack Servers, VICs
Server Form Factors	Blade (B-Series), Rack (C-Series)
Chassis Form Factor (Blade)	6 RU, up to 8 half-width or 4 full-width blades
CPU Support	Latest Intel Xeon Scalable Processors (Multiple Generations)
System Memory	DDR4, DDR5 DIMMs (Capacity varies by model, e.g., up to 12TB in C220 M6)
I/O Bandwidth	10GbE, 25GbE, 40GbE, 100GbE, 200GbE (Network & FCoE)
Management Interface	Cisco UCS Manager (GUI, CLI, XML API)
Power Supply Redundancy	N, N+1, N+N (Grid) Redundant, 80

	PLUS Titanium
Storage Options	SAS, SATA, NVMe SSDs (Hot-swappable, E1.S, M.2)

ORDERING OPTIONS

Cisco offers a comprehensive portfolio of UCS hardware components to build a solution tailored to specific workload requirements. The primary ordering categories include:

Fabric Interconnects

Various models are available, such as the Cisco UCS 6300 and 6400 Series, offering different port densities and bandwidth capacities (e.g., 10/25/40/100 Gbps) to support domains of varying sizes .

Blade Server Chassis & Extenders

The Cisco UCS 5108 Blade Server Chassis is ordered with a selection of Fabric Extenders (e.g., Cisco UCS 2300 and 2400 Series) that provide the I/O connectivity from the chassis to the Fabric Interconnects .

Blade Servers (B-Series)

A variety of blade models are available with different CPU, memory, and storage options to suit different workloads, from entry-level to mission-critical applications .

Rack Servers (C-Series)

An extensive lineup of rack servers offers flexibility for numerous applications, from general-purpose compute (C220 M6) to GPU-accelerated AI and HPC workloads (C880A M8) .

Virtual Interface Cards (VICs)

Available in mLOM and PCIe form factors, these cards provide the networking and storage connectivity, with options ranging from 10G to 200G interfaces .

