



Overview

- [Cisco UCS Manager User CLI Documentation](#), on page 1
- [Infrastructure Management Guide Overview](#), on page 2
- [Cisco Unified Computing System Overview](#), on page 3
- [Cisco UCS Building Blocks and Connectivity](#), on page 5

Cisco UCS Manager User CLI Documentation

Cisco UCS Manager offers you a new set of smaller, use-case based documentation described in the following table:

Guide	Description
Cisco UCS Manager Getting Started Guide	Discusses Cisco UCS architecture and Day 0 operations, including Cisco UCS Manager initial configuration, and configuration best practices.
Cisco UCS Manager Administration Guide	Discusses password management, role-based access configuration, remote authentication, communication services, CIMC session management, organizations, backup and restore, scheduling options, BIOS tokens and deferred deployments.
Cisco UCS Manager Infrastructure Management Guide	Discusses physical and virtual infrastructure components used and managed by Cisco UCS Manager.
Cisco UCS Manager Firmware Management Guide	Discusses downloading and managing firmware, upgrading through Auto Install, upgrading through service profiles, directly upgrading at endpoints using firmware auto sync, managing the capability catalog, deployment scenarios, and troubleshooting.
Cisco UCS Manager Server Management Guide	Discusses the new licenses, registering Cisco UCS domains with Cisco UCS Central, power capping, server boot, server profiles and server-related policies.

Guide	Description
Cisco UCS Manager Storage Management Guide	Discusses all aspects of storage management such as SAN and VSAN in Cisco UCS Manager.
Cisco UCS Manager Network Management Guide	Discusses all aspects of network management such as LAN and VLAN connectivity in Cisco UCS Manager.
Cisco UCS Manager System Monitoring Guide	Discusses all aspects of system and health monitoring including system statistics in Cisco UCS Manager.
Cisco UCS S3260 Server Integration with Cisco UCS Manager	Discusses all aspects of management of UCS S-Series servers that are managed through Cisco UCS Manager.

Infrastructure Management Guide Overview

This guide provides an overview of the physical and virtual infrastructure used in Cisco Unified Computing System (UCS) and managed through Cisco UCS Manager. It also provides detailed information about managing these infrastructure components. The following table summarizes the overall organization of the guide.

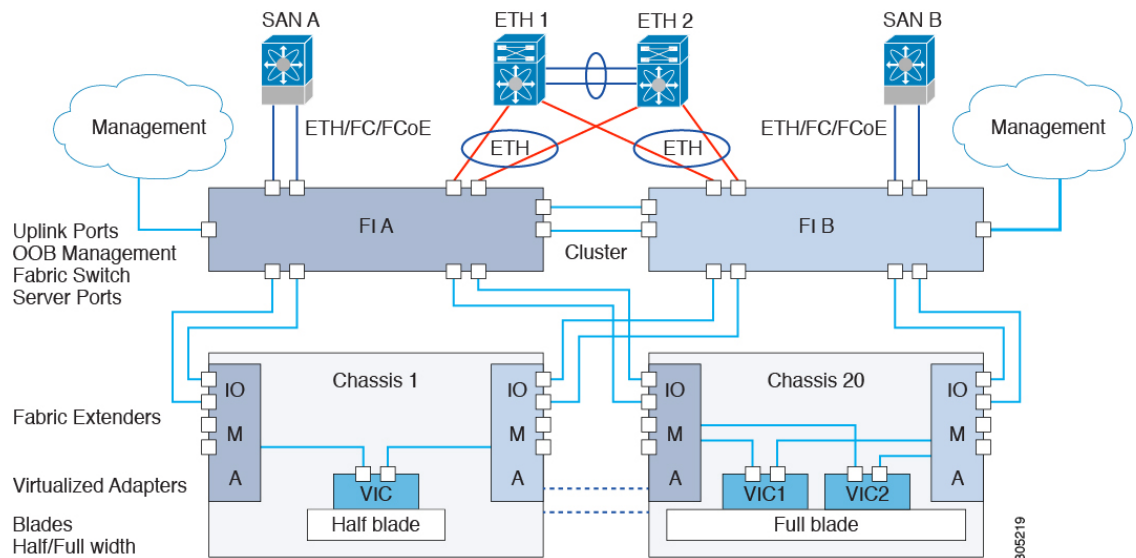
Topic	Description
Overview	Conceptual overview of Cisco UCS architecture including Cisco Fabric Interconnects, I/O Module, Chassis, Servers, and Virtualization in Cisco UCS.
Equipment Policies	Equipment policies such as Chassis/FEX discovery policy, Chassis connectivity policy, and Rack Server discovery policy.
Chassis Management	Overview of the chassis supported and procedures to manage them.
I/O Module Management	Overview of I/O Modules and procedures to manage them.
Power Management in Cisco UCS	Overview of UCS Power Management policies, Global Power policies, and Power Capping.
Blade Server Management	Overview of Blade Servers and procedures to manage them.
Rack-Mount Server Management	Overview of Rack-Mount Servers and procedures to manage them.
S3X60 Server Node Management	Overview of S3X60 Server Node and procedures to manage them.
Virtual Interface Management	Overview of Virtualization in Cisco UCS, Virtual Interfaces and procedures to manage them.
Server Troubleshooting	Common troubleshooting scenarios for Servers.

Cisco Unified Computing System Overview

Cisco UCS has a unique architecture that integrates compute, data network access, and storage network access into a common set of components under a single-pane-of-glass management interface.

Cisco UCS fuses access layer networking and servers. This high-performance, next-generation server system provides a data center with a high degree of workload agility and scalability. The hardware and software components support Cisco's unified fabric, which runs multiple types of data center traffic over a single converged network adapter.

Figure 1: Cisco Unified Computing System Architecture



Architectural Simplification

The simplified architecture of Cisco UCS reduces the number of required devices and centralizes switching resources. By eliminating switching inside a chassis, network access-layer fragmentation is significantly reduced. Cisco UCS implements Cisco unified fabric within racks and groups of racks, supporting Ethernet and Fibre Channel protocols over 10 Gigabit Cisco Data Center Ethernet and Fibre Channel over Ethernet (FCoE) links. This radical simplification reduces the number of switches, cables, adapters, and management points by up to two-thirds. All devices in a Cisco UCS domain remain under a single management domain, which remains highly available through the use of redundant components.

High Availability

The management and data plane of Cisco UCS is designed for high availability and redundant access layer fabric interconnects. In addition, Cisco UCS supports existing high availability and disaster recovery solutions for the data center, such as data replication and application-level clustering technologies.

Scalability

A single Cisco UCS domain supports multiple chassis and their servers, all of which are administered through one Cisco UCS Manager. For more detailed information about the scalability, speak to your Cisco representative.

Flexibility

A Cisco UCS domain allows you to quickly align computing resources in the data center with rapidly changing business requirements. This built-in flexibility is determined by whether you choose to fully implement the stateless computing feature. Pools of servers and other system resources can be applied as necessary to respond to workload fluctuations, support new applications, scale existing software and business services, and accommodate both scheduled and unscheduled downtime. Server identity can be abstracted into a mobile service profile that can be moved from server to server with minimal downtime and no need for additional network configuration.

With this level of flexibility, you can quickly and easily scale server capacity without having to change the server identity or reconfigure the server, LAN, or SAN. During a maintenance window, you can quickly do the following:

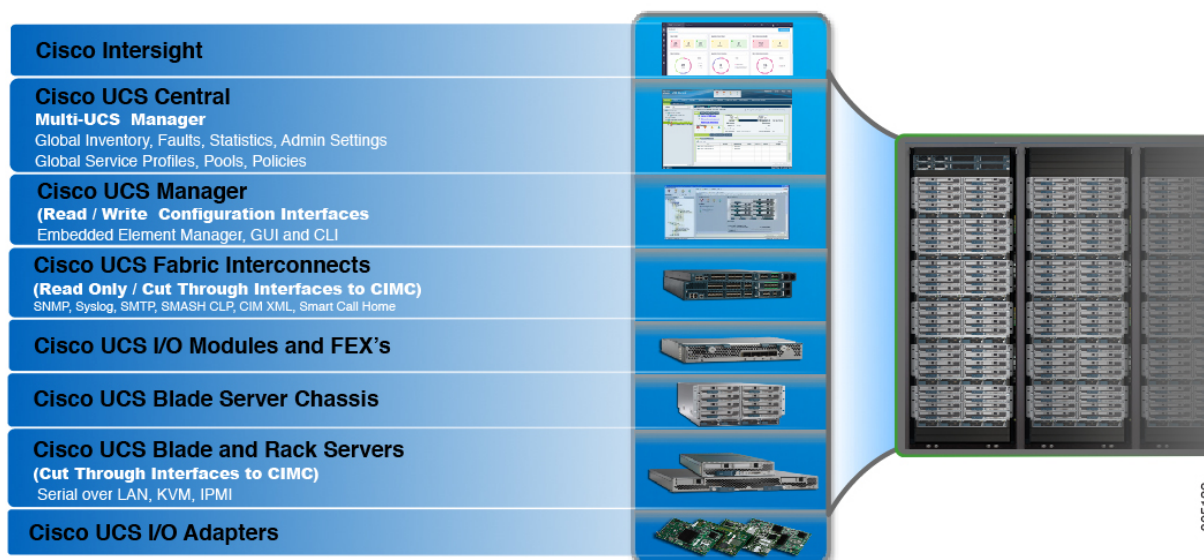
- Deploy new servers to meet unexpected workload demand and rebalance resources and traffic.
- Shut down an application, such as a database management system, on one server and then boot it up again on another server with increased I/O capacity and memory resources.

Optimized for Server Virtualization

Cisco UCS has been optimized to implement VM-FEX technology. This technology provides improved support for server virtualization, including better policy-based configuration and security, conformance with a company's operational model, and accommodation for VMware's VMotion.

Cisco UCS Building Blocks and Connectivity

Figure 2: Cisco UCS Building Blocks and Connectivity



As shown in the figure above, the primary components included within Cisco UCS are as follows:

- **Cisco UCS Manager**—Cisco UCS Manager is the centralized management interface for Cisco UCS. For more information on Cisco UCS Manager, see *Introduction to Cisco UCS manager* in *Cisco UCS Manager Getting Started Guide*
- **Cisco UCS Fabric Interconnects**—The Cisco UCS Fabric Interconnect is the core component of Cisco UCS deployments, providing both network connectivity and management capabilities for the Cisco UCS system. The Cisco UCS Fabric Interconnects run the Cisco UCS Manager control software and consist of the following components:
 - Cisco UCS 6200 series Fabric Interconnects, Cisco UCS 6332 series Fabric Interconnects, and Cisco UCS Mini
 - Transreceivers for network and storage connectivity
 - Expansion modules for the various Fabric Interconnects
 - Cisco UCS Manager software

For more information on Cisco UCS Fabric Interconnects, see [Cisco UCS Fabric Infrastructure Portfolio](#), on page 6.

- **Cisco UCS I/O Modules and Cisco UCS Fabric Extender**—IOM modules are also known as Cisco FEXs or simply FEX modules. These modules serve as line cards to the FIs in the same way that Nexus series switches can have remote line cards. IOM modules also provide interface connections to blade servers. They multiplex data from blade servers and provide this data to FIs and do the same in the reverse direction. In production environments, IOM modules are always used in pairs to provide redundancy and failover.



Important The 40G backplane setting is not applicable for 22xx IOMs.

- **Cisco UCS Blade Server Chassis**—The Cisco UCS 5100 Series Blade Server Chassis is a crucial building block of Cisco UCS, delivering a scalable and flexible architecture for current and future data center needs, while helping reduce total cost of ownership.
- **Cisco UCS Blade and Rack Servers**—Cisco UCS Blade servers are at the heart of the UCS solution. They come in various system resource configurations in terms of CPU, memory, and hard disk capacity. All blade servers are based on Intel Xeon processors. There is no AMD option available. The Cisco UCS rack-mount servers are standalone servers that can be installed and controlled individually. Cisco provides Fabric Extenders (FEXs) for the rack-mount servers. FEXs can be used to connect and manage rack-mount servers from FIs. Rack-mount servers can also be directly attached to the fabric interconnect.

Small and Medium Businesses (SMBs) can choose from different blade configurations as per business needs

- **Cisco UCS I/O Adapters**—Cisco UCS B-Series Blade Servers are designed to support up to two network adapters. This design can reduce the number of adapters, cables, and access-layer switches by as much as half because it eliminates the need for multiple parallel infrastructure for both LAN and SAN at the server, chassis, and rack levels.

Cisco UCS Fabric Infrastructure Portfolio

The Cisco UCS fabric interconnects are top-of-rack devices and provide unified access to the Cisco UCS domain. The following illustration shows the evolution of the Cisco UCS fabric interconnects product family. The Cisco UCS Infrastructure hardware is now in its third generation.



Note The Cisco UCS 6100 Series Fabric Interconnects and Cisco UCS 2104 I/O Modules have reached end of life.



Cisco UCS I/O Modules and Cisco UCS Fabric Extenders

Cisco UCS 2200 and 2300 Series Fabric Extenders, also known as I/O modules, bring the unified fabric into the blade server enclosure, providing multiple 10 Gigabit Ethernet connections between servers and the fabric interconnect, simplifying diagnostics, cabling, and management.

The Cisco UCS 2200 Series extends the I/O fabric between the Cisco UCS 6200 Series fabric interconnects and the Cisco UCS 5100 Series blade server chassis. The Cisco UCS 2300 Series extends the I/O fabric between the Cisco UCS 6300 Series fabric interconnects and the Cisco UCS 5100 Series blade server chassis. These fabric extenders enable a lossless and deterministic Fibre Channel over Ethernet (FCoE) fabric to

connect all servers and chassis together. Because the fabric extender is similar to a distributed line card, it does not perform any switching, and is managed as an extension of the fabric interconnects.

The Cisco UCS 2200 Series manages the fabric interconnect and the chassis environment, including the power supply, fans, and blade servers. Therefore, separate chassis management modules are not required. These fabric extenders fit into the back of the Cisco UCS 5100 Series chassis. Each Cisco UCS 5100 Series chassis can support up to two fabric extenders, allowing increased capacity and redundancy.

Cisco UCS 2232PP and Cisco UCS 2232TM-E extend the fabric between the Cisco UCS 6200 Series fabric interconnects and rack mount servers. Cisco UCS 2348UPQ extends the fabric between the Cisco UCS 6300 Series fabric interconnects and rack mount servers.

[I/O Module Management](#) provides more details about managing I/O Modules.

Cisco UCS Chassis

Cisco UCS Manager Release 3.1(1) and later releases provide support for Cisco UCS 5108 Blade Server Chassis

[Chassis Management](#) provides details on managing the chassis through Cisco UCS Manager.

Cisco UCS 5108 Blade Server Chassis

The Cisco UCS 5108 Blade Server Chassis, is six rack units (6RU) high, can mount in an industry-standard 19-inch rack, and uses standard front-to-back cooling. A chassis can accommodate up to eight half-width, or four full-width Cisco UCS B-Series Blade Servers form factors within the same chassis. By incorporating unified fabric and fabric-extender technology, the Cisco Unified Computing System enables the chassis to:

- Have fewer physical components
- Require no independent management
- Be more energy efficient than a traditional blade-server chassis

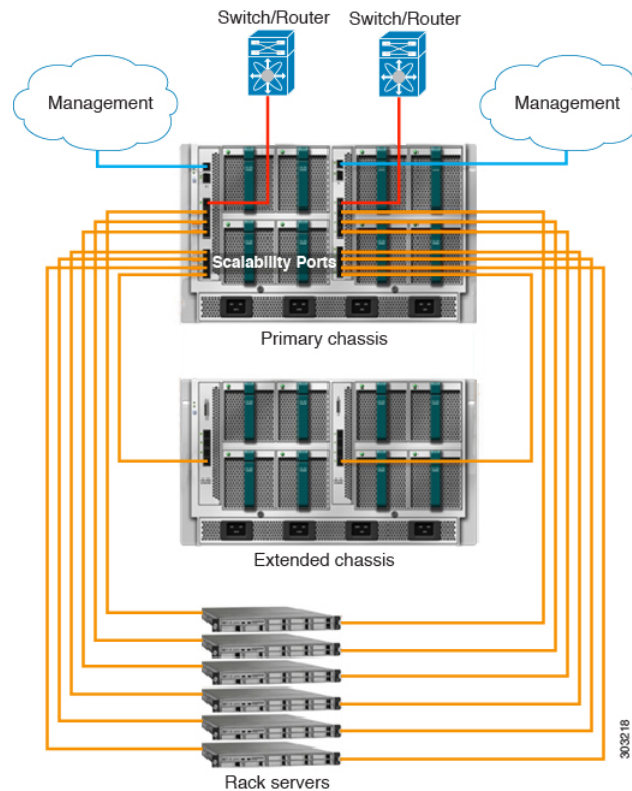
The Cisco UCS 5108 Blade Server Chassis is supported with all generations of fabric interconnects.

Cisco UCS Mini Infrastructure

The Cisco UCS Mini solution extends the Cisco UCS architecture into environments with requirements for smaller domains, including branch and remote offices, point-of-sale locations, and smaller IT environments. Cisco UCS Mini has three main infrastructure components:

- Cisco UCS 6324 fabric interconnect
- Cisco UCS blade server chassis
- Cisco UCS blade or rack mount servers

Figure 3: Cisco UCS Mini



In the Cisco UCS Mini solution, the Cisco UCS 6324 fabric interconnect is collapsed into the IO Module form factor, and is inserted into the IOM slot of the blade server chassis. The Cisco UCS 6324 fabric interconnect has 24 10G ports available on it. Sixteen of these ports are server facing, two 10G ports are dedicated to each of the eight half width blade slots. The remaining eight ports are divided into groups of four 1/10G Enhanced Small Form-Factor Pluggable (SFP+) ports and one 40G Quad Small Form-factor Pluggable (QSFP) port, which is called the 'scalability port'.

Cisco UCS Manager Release 3.1(1) introduces support for a second UCS 5108 chassis to an existing single-chassis Cisco UCS 6324 fabric interconnect setup. This extended chassis enables you to configure an additional 8 servers. Unlike the primary chassis, the extended chassis supports IOMs. Currently, it supports UCS-IOM-2204XP and UCS-IOM-2208XP IOMs. The extended chassis can only be connected through the scalability port on the FI-IOM.

**Important**

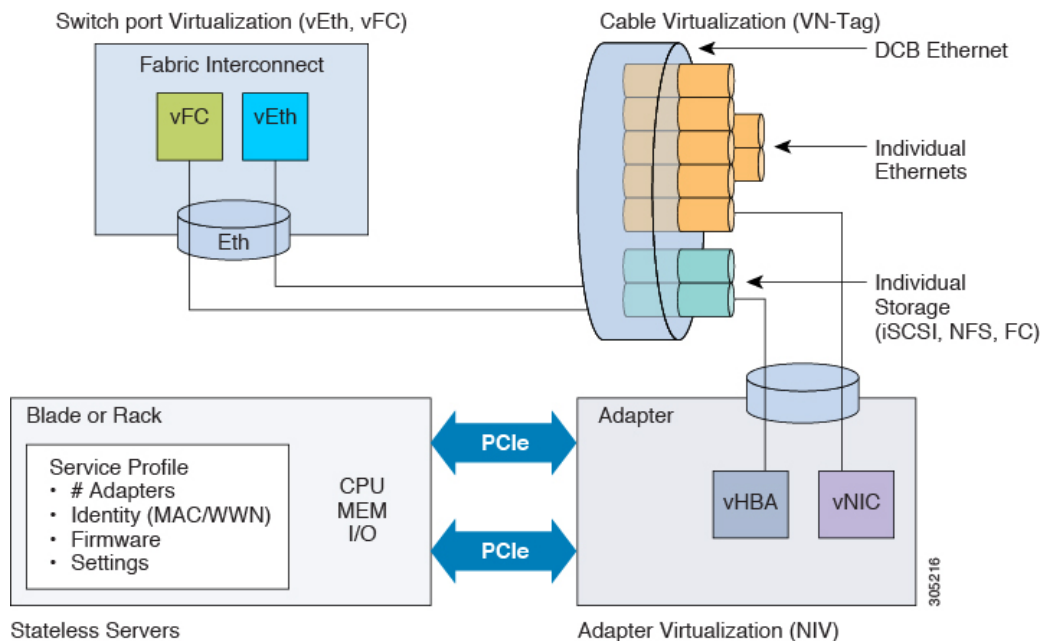
Currently, Cisco UCS Manager supports only one extended chassis for UCS Mini.

Cisco UCS Infrastructure Virtualization

Cisco UCS is a single integrated system with switches, cables, adapters, and servers that are all tied together and managed by unified management software. One capability that enables this unification is the ability to virtualize every component of the system at every level. Switch port, cables, adapter, and servers can all be virtualized. Because of the virtualization capabilities at every component of the system, you have the unique

capability to provide rapid provisioning of any service on any server on any blade through a system that is wired once. The following image illustrates these virtualization capabilities.

Figure 4: Virtualization Capabilities of Cisco UCS



Switch Port Virtualization

The physical interfaces provide physical connectivity for what are logical virtual interfaces on the fabric interconnects—virtual Fibre Channel interfaces (vFC) and virtual Ethernet interfaces (vEth). The logical connectivity to a server is provided through these virtual interfaces.

Cable Virtualization

The physical cables that connect to physical switch ports provide the infrastructure for logical and virtual cables. These virtual cables connect to virtual adapters on any given server in the system.

Adapter Virtualization

On the server, you have physical adapters, which provide physical infrastructure for virtual adapters. A virtual network interface card (vNIC) or virtual host bus adapter (vHBA) logically connects a host to a virtual interface on the fabric interconnect and allows the host to send and receive traffic through that interface. Each virtual interface in the fabric interconnect corresponds to a vNIC.

An adapter that is installed on the server appears to the server as multiple adapters through standard PCIe virtualization. When the server scans the PCIe bus, the virtual adapters that are provisioned appear to be physically plugged into the PCIe bus.

Server Virtualization

Server virtualization provides you with the ability of stateless servers. As part of the physical infrastructure, you have physical servers. However, the configuration of a server is derived from the service profile to which it is associated. All service profiles are centrally managed and stored in a database on the fabric interconnect.

A service profile defines all the settings of the server, for example, the number of adapters, virtual adapters, the identity of these adapters, the firmware of the adapters, and the firmware of the server. It contains all the settings of the server that you typically configure on a physical machine. Because the service profile is abstracted from the physical infrastructure, you can apply it to any physical server and the physical server will be configured according to the configuration defined in the service profile. *Cisco UCS Manager Server Management Guide* provides detailed information about managing service profiles.