Overview of Cisco Unified Computing System

This chapter includes the following sections:

About Cisco Unified Computing System

Cisco Unified Computing System (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.

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.

Unified Fabric

With unified fabric, multiple types of data center traffic can run over a single Data Center Ethernet (DCE) network. Instead of having a series of different host bus adapters (HBAs) and network interface cards (NICs) present in a server, unified fabric uses a single converged network adapter. This type of adapter can carry LAN and SAN traffic on the same cable.

Cisco UCS uses Fibre Channel over Ethernet (FCoE) to carry Fibre Channel and Ethernet traffic on the same physical Ethernet connection between the fabric interconnect and the server. This connection terminates at a converged network adapter on the server, and the unified fabric terminates on the uplink ports of the fabric interconnect. On the core network, the LAN and SAN traffic remains separated. Cisco UCS does not require that you implement unified fabric across the data center.

The converged network adapter presents an Ethernet interface and Fibre Channel interface to the operating system. At the server, the operating system is not aware of the FCoE encapsulation because it sees a standard Fibre Channel HBA.

At the fabric interconnect, the server-facing Ethernet port receives the Ethernet and Fibre Channel traffic. The fabric interconnect (using Ethertype to differentiate the frames) separates the two traffic types. Ethernet frames and Fibre Channel frames are switched to their respective uplink interfaces.

Fibre Channel over Ethernet

Cisco UCS leverages Fibre Channel over Ethernet (FCoE) standard protocol to deliver Fibre Channel. The upper Fibre Channel layers are unchanged, so the Fibre Channel operational model is maintained. FCoE network management and configuration is similar to a native Fibre Channel network.

FCoE encapsulates Fibre Channel traffic over a physical Ethernet link. FCoE is encapsulated over Ethernet with the use of a dedicated Ethertype, 0x8906, so that FCoE traffic and standard Ethernet traffic can be carried on the same link. FCoE has been standardized by the ANSI T11 Standards Committee.

Fibre Channel traffic requires a lossless transport layer. Instead of the buffer-to-buffer credit system used by native Fibre Channel, FCoE depends upon the Ethernet link to implement lossless service.

Ethernet links on the fabric interconnect provide two mechanisms to ensure lossless transport for FCoE traffic:

- Link-level flow control
- Priority flow control

Link-Level Flow Control

IEEE 802.3x link-level flow control allows a congested receiver to signal the endpoint to pause data transmission for a short time. This link-level flow control pauses all traffic on the link.

The transmit and receive directions are separately configurable. By default, link-level flow control is disabled for both directions.

On each Ethernet interface, the fabric interconnect can enable either priority flow control or link-level flow control (but not both).

Priority Flow Control

The priority flow control (PFC) feature applies pause functionality to specific classes of traffic on the Ethernet link. For example, PFC can provide lossless service for the FCoE traffic, and best-effort service for the standard Ethernet traffic. PFC can provide different levels of service to specific classes of Ethernet traffic (using IEEE 802.1p traffic classes).

PFC decides whether to apply pause based on the IEEE 802.1p CoS value. When the fabric interconnect enables PFC, it configures the connected adapter to apply the pause functionality to packets with specific CoS values.

By default, the fabric interconnect negotiates to enable the PFC capability. If the negotiation succeeds, PFC is enabled and link-level flow control remains disabled (regardless of its configuration settings). If the PFC negotiation fails, you can either force PFC to be enabled on the interface or you can enable IEEE 802.x link-level flow control.

IPv6 Compliance

Cisco UCS Manager supports IPv6 addressing. This is important for the following reasons:

- IPv4 addresses have a shorter address space than IPv6 addresses.

- The number of unique IPv4 addresses is finite, and the allocation scheme used by the Internet addressing body has exacerbated the decline of available addresses.

- IPv6 addresses have a larger address space, and the pool of available IPv6 addresses is much greater than the pool of IPv4 addresses.

- Some customers require that all networking software they purchase be IPv6 standards compliant.

All features in Cisco UCS Manager that support IPv4 addressing also support IPv6.

Note   Only public global unicast IPv6 addresses are supported.

IPv6 addresses can be used to configure inband access to management interfaces, the Cisco Cisco UCS Manager GUI, the KVM Console, and SSH over SoL.

Note   IPv6 addresses are not supported for out-of-band access to CIMC.

Services Supported

Services that support IPv6 addresses include:

- HTTP and HTTPS

- SSH

- Telnet

- CIM XML

- SNMP

- Flash policy server

Client Support

External clients that support IPv6 addresses include:

- NTP

- DNS

- DHCP

- LDAP

- RADIUS

- TACACS+

- SSH

- Syslog

- vCenter

- Call Home

- NFS

Fabric Interconnects

Initial setup of the fabric interconnects supports the use of IPv6 addresses for the management IP address, default gateway and DNS servers.
In a cluster setup, if Fabric A is configured using IPv6 addresses and a cluster configuration is enabled, when Fabric B is subsequently configured, the setup process retrieves the address type from Fabric A, and prompts you to use IPv6 addresses. IPv4 addresses then need to be configured for both fabric interconnects for out-of-band (OOB) access after initial setup is complete.

Cisco UCS Manager and the fabric interconnects support OOB access over both IPv4 and IPv6 addresses.

Configurations that Support IPv6 Addressing

IPv6 addresses can be used to configure key ring certificate requests, SNMP traps, management IP pools and address blocks, service profiles, service profile templates, VLAN groups, backup and restore operations, the core file exporter, the Cisco UCS Manager Syslog, NTP servers, ARP targets in the Management Interface Monitoring policy, System Event Log (SEL) management, license management, firmware download, Call Home, and vCenter.

LDAP, RADIUS and TACACS+ authentication service provider configurations all support IPv6 addressing.

Servers

Cisco UCS blade and rack servers can be configured to use static IPv6 addresses. Inband access to the server Cisco Integrated Management Controller (CIMC) is possible using IPv6 addresses. Inband access is faster because management traffic flows between the fabric interconnects and the servers using the higher-bandwidth uplink port.

Note   Only Cisco UCS M3 and M4 servers support IPv6 addresses. IPv6 addressing for Cisco UCS M1 and M2 servers is not supported.

Server Architecture and Connectivity

Overview of Service Profiles

Service profiles are the central concept of Cisco UCS. Each service profile serves a specific purpose: ensuring that the associated server hardware has the configuration required to support the applications it will host.

The service profile maintains configuration information about the server hardware, interfaces, fabric connectivity, and server and network identity. This information is stored in a format that you can manage

through Cisco UCS Manager. All service profiles are centrally managed and stored in a database on the fabric interconnect.

Every server must be associated with a service profile.

☞

Important   At any given time, each server can be associated with only one service profile. Similarly, each service profile can be associated with only one server at a time.

After you associate a service profile with a server, the server is ready to have an operating system and applications installed, and you can use the service profile to review the configuration of the server. If the server associated with a service profile fails, the service profile does not automatically fail over to another server.

When a service profile is disassociated from a server, the identity and connectivity information for the server is reset to factory defaults.

Network Connectivity through Service Profiles

Each service profile specifies the LAN and SAN network connections for the server through the Cisco UCS infrastructure and out to the external network. You do not need to manually configure the network connections for Cisco UCS servers and other components. All network configuration is performed through the service profile.

When you associate a service profile with a server, the Cisco UCS internal fabric is configured with the information in the service profile. If the profile was previously associated with a different server, the network infrastructure reconfigures to support identical network connectivity to the new server.

Configuration through Service Profiles

A service profile can take advantage of resource pools and policies to handle server and connectivity configuration.

Hardware Components Configured by Service Profiles

When a service profile is associated with a server, the following components are configured according to the data in the profile:

- Server, including BIOS and CIMC
- Adapters
- Fabric interconnects

You do not need to configure these hardware components directly.

Server Identity Management through Service Profiles

You can use the network and device identities burned into the server hardware at manufacture or you can use identities that you specify in the associated service profile either directly or through identity pools, such as MAC, WWN, and UUID.

The following are examples of configuration information that you can include in a service profile:

- Profile name and description

- Unique server identity (UUID)

- LAN connectivity attributes, such as the MAC address

- SAN connectivity attributes, such as the WWN

Operational Aspects configured by Service Profiles

You can configure some of the operational functions for a server in a service profile, such as the following:

- Firmware packages and versions

- Operating system boot order and configuration

- IPMI and KVM access

vNIC Configuration by Service Profiles

A vNIC is a virtualized network interface that is configured on a physical network adapter and appears to be a physical NIC to the operating system of the server. The type of adapter in the system determines how many vNICs you can create. For example, a converged network adapter has two NICs, which means you can create a maximum of two vNICs for each adapter.

A vNIC communicates over Ethernet and handles LAN traffic. At a minimum, each vNIC must be configured with a name and with fabric and network connectivity.

vHBA Configuration by Service Profiles

A vHBA is a virtualized host bus adapter that is configured on a physical network adapter and appears to be a physical HBA to the operating system of the server. The type of adapter in the system determines how many vHBAs you can create. For example, a converged network adapter has two HBAs, which means you can create a maximum of two vHBAs for each of those adapters. In contrast, a network interface card does not have any HBAs, which means you cannot create any vHBAs for those adapters.

A vHBA communicates over FCoE and handles SAN traffic. At a minimum, each vHBA must be configured with a name and fabric connectivity.

Service Profiles that Override Server Identity

This type of service profile provides the maximum amount of flexibility and control. This profile allows you to override the identity values that are on the server at the time of association and use the resource pools and policies set up in Cisco UCS Manager to automate some administration tasks.

You can disassociate this service profile from one server, then associate it with another server. This re-association can be done either manually or through an automated server pool policy. The burned-in settings, such as UUID and MAC address on the new server are overwritten with the configuration in the service profile. As a result, the change in the server is transparent to your network. You do not need to reconfigure any component or application on your network to begin using the new server.

This profile allows you to take advantage of and manage system resources through resource pools and policies, such as the following:

- Virtualized identity information, including pools of MAC addresses, WWN addresses, and UUIDs

- Ethernet and Fibre Channel adapter profile policies

- Firmware package policies

• Operating system boot order policies

Unless the service profile contains power management policies, a server pool qualification policy, or another policy that requires a specific hardware configuration, you can use the profile for any type of server in the Cisco UCS domain.

You can associate these service profiles with either a rack-mount server or a blade server. The ability to migrate the service profile depends upon whether you choose to restrict migration of the service profile.

Note If you choose not to restrict migration, Cisco UCS Manager does not perform any compatibility checks on the new server before migrating the existing service profile. If the hardware of both servers are not similar, the association might fail.

Service Profiles that Inherit Server Identity

This hardware-based service profile is the simplest to use and create. This profile uses the default values in the server and mimics the management of a rack-mounted server. It is tied to a specific server and cannot be moved or migrated to another server.

You do not need to create pools or configuration policies to use this service profile.

This service profile inherits and applies the identity and configuration information that is present at the time of association, such as the following:

• MAC addresses for the two NICs

• For a converged network adapter or a virtual interface card, the WWN addresses for the two HBAs

• BIOS versions

• Server UUID

Important The server identity and configuration information inherited through this service profile might not have the values burned into the server hardware at the manufacturer if those values were changed before this profile is associated with the server.

Initial and Existing Templates

With a service profile template, you can quickly create several service profiles with the same basic parameters, such as the number of vNICs and vHBAs, and with identity information drawn from the same pools.

Tip If you need only one service profile with similar values to an existing service profile, you can clone a service profile in the Cisco UCS Manager GUI.

For example, if you need several service profiles with similar values to configure servers to host database software, you can create a service profile template, either manually or from an existing service profile. You then use the template to create the service profiles.

Cisco UCS supports the following types of service profile templates:

Initial template

Service profiles created from an initial template inherit all the properties of the template. Service profiles created from an initial service profile template are bound to the template. However, changes to the initial template do not *automatically* propagate to the bound service profiles. If you want to propagate changes to bound service profiles, unbind and rebind the service profile to the initial template.

Updating template

Service profiles created from an updating template inherit all the properties of the template and remain connected to the template. Any changes to the template automatically update the service profiles created from the template.

> **Note**
>
> Service profiles that are created from the initial template and normal service profiles fetch the lowest available IDs in the sequential pool when you press **Reset**.
>
> Service profiles created from updating template might attempt to retain the same ID when you press **Reset** even when lower IDs of sequential pool are free.

Policies

Policies determine how Cisco UCS components will act in specific circumstances. You can create multiple instances of most policies. For example, you might want different boot policies, so that some servers can PXE boot, some can SAN boot, and others can boot from local storage.

Policies allow separation of functions within the system. A subject matter expert can define policies that are used in a service profile, which is created by someone without that subject matter expertise. For example, a LAN administrator can create adapter policies and quality of service policies for the system. These policies can then be used in a service profile that is created by someone who has limited or no subject matter expertise with LAN administration.

You can create and use two types of policies in Cisco UCS Manager:

- Configuration policies that configure the servers and other components

- Operational policies that control certain management, monitoring, and access control functions

Pools

Pools are collections of identities, or physical or logical resources, that are available in the system. All pools increase the flexibility of service profiles and allow you to centrally manage your system resources.

You can use pools to segment unconfigured servers or available ranges of server identity information into groupings that make sense for the data center. For example, if you create a pool of unconfigured servers with similar characteristics and include that pool in a service profile, you can use a policy to associate that service profile with an available, unconfigured server.

If you pool identifying information, such as MAC addresses, you can preassign ranges for servers that host specific applications. For example, you can configure all database servers within the same range of MAC addresses, UUIDs, and WWNs.

Domain Pools

Domain Pools are defined locally in a Cisco UCS domain, and can only be used in that Cisco UCS domain.

Global Pools

Global Pools are defined in Cisco UCS Central, and can be shared between Cisco UCS domains. If a Cisco UCS domain is registered with Cisco UCS Central, you can assign **Global Pools** in Cisco UCS Manager.

CIMC Inband Management

A driving factor for providing inband management access to Cisco Integrated Management Controller (CIMC) is the desire to separate tenant traffic from provider traffic in multi-tenant, public or private service provider cloud deployments. Out-of-band (OOB) management traffic moves in and out of the fabric interconnects and traverses the management plane via the management port. This has the potential to cause bottlenecks and affect the CPU bandwidth in the management ports.

Inband management allows CIMC traffic to take the same path as the data traffic, entering and exiting the fabric interconnects via the uplink ports. The higher bandwidth available to the uplink ports means that inband access greatly speeds up management traffic, and reduces the risk of traffic bottlenecks and CPU stress. Both out-of-band (OOB) and inband address pools can be configured for management access in Cisco UCS Manager. Out-of-band access only supports IPv4 addresses. Inband access supports both IPv4 and IPv6 addresses, which allows for single or dual stack management.

The two OOB management interface addresses that can be configured in Cisco UCS Manager blade and rack servers are:

- An OOB IPv4 address assigned to the physical server via the global ext-mgmt pool
- An OOB IPv4 address derived from a service profile associated with the physical server

In addition, up to four inband management interface addresses can be configured:

- An inband IPv4 address assigned to the physical server
- An inband IPv4 address derived from a service profile associated with the physical server
- An inband IPv6 address assigned to the physical server
- An inband IPv6 address derived from a service profile associated with the physical server

Multiple inband management IP addresses for each server support additional CIMC sessions. When you configure both OOB and inband addresses, users can choose from a list of those addresses in the KVM Console dialog box when they launch KVM from a server, SSH to SoL, a service profile, the KVM Launch Manager, or from the Cisco UCS Manager GUI web URL.

CIMC inband access supports the following services:

- KVM Console
- SSH to CIMC for SoL

• vMedia for ISO, virtual CD/DVD, removable disk, and floppy

Note Only Cisco UCS M3 and M4 servers support inband CIMC access. Inband CIMC access for Cisco UCS M1 and M2 servers is not supported.

You can configure inband IP pools of IPv4 or IPv6 addresses and use them to assign addresses to servers. You can configure inband VLAN groups and assign them to servers using service profiles.

You need to configure an Inband Profile with an Inband VLAN group to select an Inband Network (VLAN) in Service Profiles and Service Profile templates.

You can configure the network and IP pool name in an Inband profile to assign Inband CIMC addresses to Cisco UCS M3 and M4 servers.

Inband Management Support

Inband management access is supported in Cisco UCS Manager for the following external services:

• KVM

• vMedia for ISO, virtual CD/DVD, removable disk, and floppy

• SSH to SoL

You can configure inband IP pools of IPv4 or IPv6 addresses and use them to assign addresses to servers. You can configure inband VLAN groups and assign them to servers using service profiles.

Traffic Management

Oversubscription

Oversubscription occurs when multiple network devices are connected to the same fabric interconnect port. This practice optimizes fabric interconnect use, since ports rarely run at maximum speed for any length of time. As a result, when configured correctly, oversubscription allows you to take advantage of unused bandwidth. However, incorrectly configured oversubscription can result in contention for bandwidth and a lower quality of service to all services that use the oversubscribed port.

For example, oversubscription can occur if four servers share a single uplink port, and all four servers attempt to send data at a cumulative rate higher than available bandwidth of uplink port.

Oversubscription Considerations

The following elements can impact how you configure oversubscription in a Cisco UCS domain:

Ratio of Server-Facing Ports to Uplink Ports

You need to know what how many server-facing ports and uplink ports are in the system, because that ratio can impact performance. For example, if your system has twenty ports that can communicate down to the

servers and only two ports that can communicate up to the network, your uplink ports will be oversubscribed. In this situation, the amount of traffic created by the servers can also affect performance.

Number of Uplink Ports from Fabric Interconnect to Network

You can choose to add more uplink ports between the Cisco UCS fabric interconnect and the upper layers of the LAN to increase bandwidth. In Cisco UCS, you must have at least one uplink port per fabric interconnect to ensure that all servers and NICs to have access to the LAN. The number of LAN uplinks should be determined by the aggregate bandwidth needed by all Cisco UCS servers.

For the 6100 series fabric interconnects, Fibre Channel uplink ports are available on the expansion slots only. You must add more expansion slots to increase number of available Fibre Channel uplinks. Ethernet uplink ports can exist on the fixed slot and on expansion slots.

For the 6200 series fabric interconnects running Cisco UCS Manager, version 2.0 and higher, Ethernet uplink ports and Fibre Channel uplink ports are both configurable on the base module, as well as on the expansion module.

For example, if you have two Cisco UCS 5100 series chassis that are fully populated with half width Cisco UCS B200-M1 servers, you have 16 servers. In a cluster configuration, with one LAN uplink per fabric interconnect, these 16 servers share 20GbE of LAN bandwidth. If more capacity is needed, more uplinks from the fabric interconnect should be added. We recommend that you have symmetric configuration of the uplink in cluster configurations. In the same example, if 4 uplinks are used in each fabric interconnect, the 16 servers are sharing 80 GB of bandwidth, so each has approximately 5 GB of capacity. When multiple uplinks are used on a Cisco UCS fabric interconnect the network design team should consider using a port channel to make best use of the capacity.

Number of Uplink Ports from I/O Module to Fabric Interconnect

You can choose to add more bandwidth between I/O module and fabric interconnect by using more uplink ports and increasing the number of cables. In Cisco UCS, you can have one, two, or four cables connecting a I/O module to a Cisco UCS 6100 series fabric interconnect. You can have up to eight cables if you're connecting a 2208 I/O module and a 6248 fabric interconnect. The number of cables determines the number of active uplink ports and the oversubscription ratio.

Number of Active Links from Server to Fabric Interconnect

The amount of non-oversubscribed bandwidth available to each server depends on the number of I/O modules used and the number of cables used to connect those I/O modules to the fabric interconnects. Having a second I/O module in place provides additional bandwidth and redundancy to the servers. This level of flexibility in design ensures that you can provide anywhere from 80 Gbps (two I/O modules with four links each) to 10 Gbps (one I/O module with one link) to the chassis.

With 80 Gbps to the chassis, each half-width server in the Cisco UCS domain can get up to 10 Gbps in a non-oversubscribed configuration, with an ability to use up to 20 Gbps with 2:1 oversubscription.

Guidelines for Estimating Oversubscription

When you estimate the optimal oversubscription ratio for a fabric interconnect port, consider the following guidelines:

Cost/Performance Slider

The prioritization of cost and performance is different for each data center and has a direct impact on the configuration of oversubscription. When you plan hardware usage for oversubscription, you need to know where the data center is located on this slider. For example, oversubscription can be minimized if the data center is more concerned with performance than cost. However, cost is a significant factor in most data centers, and oversubscription requires careful planning.

Bandwidth Usage

The estimated bandwidth that you expect each server to actually use is important when you determine the assignment of each server to a fabric interconnect port and, as a result, the oversubscription ratio of the ports. For oversubscription, you must consider how many GBs of traffic the server will consume on average, the ratio of configured bandwidth to used bandwidth, and the times when high bandwidth use will occur.

Network Type

The network type is only relevant to traffic on uplink ports, because FCoE does not exist outside Cisco UCS. The rest of the data center network only differentiates between LAN and SAN traffic. Therefore, you do not need to take the network type into consideration when you estimate oversubscription of a fabric interconnect port.

Pinning

Pinning in Cisco UCS is only relevant to uplink ports. You can pin Ethernet or FCoE traffic from a given server to a specific uplink Ethernet port or uplink FC port.

When you pin the NIC and HBA of both physical and virtual servers to uplink ports, you give the fabric interconnect greater control over the unified fabric. This control ensures more optimal utilization of uplink port bandwidth.

Cisco UCS uses pin groups to manage which NICs, vNICs, HBAs, and vHBAs are pinned to an uplink port. To configure pinning for a server, you can either assign a pin group directly, or include a pin group in a vNIC policy, and then add that vNIC policy to the service profile assigned to that server. All traffic from the vNIC or vHBA on the server travels through the I/O module to the same uplink port.

Guidelines for Pinning

When you determine the optimal configuration for pin groups and pinning for an uplink port, consider the estimated bandwidth usage for the servers. If you know that some servers in the system will use a lot of bandwidth, ensure that you pin these servers to different uplink ports.

Quality of Service

Cisco UCS provides the following methods to implement quality of service:

- System classes that specify the global configuration for certain types of traffic across the entire system
- QoS policies that assign system classes for individual vNICs
- Flow control policies that determine how uplink Ethernet ports handle pause frames

Global QoS changes made to the QoS system class may result in brief data-plane interruptions for all traffic. Some examples of such changes are:

- Changing the MTU size for an enabled class

- Changing packet drop for an enabled class

- Changing the CoS value for an enabled class

Guidelines and Limitations for Quality of Service on Cisco UCS 6300 Series Fabric Interconnect

- Cisco UCS 6300 Series Fabric Interconnect uses a shared buffer for all system classes.

- Multicast optimization is not supported.

- When you change the QoS parameters for any class causes traffic disruption to all classes. The following table lists the changes in the QoS system class and the conditions that trigger a system reboot.

QoS System class status	Condition	FI Reboot Status
Enabled	Change between drop and no drop	Yes
No-drop	Change between enable and disable	Yes
Enable and no-drop	Change in MTU size	Yes

- The subordinate FI reboots first as a result of the change in the QoS system class. The primary FI reboots only after you acknowledge it in **Pending Activities**.

Guidelines and Limitations for Quality of Service on Cisco UCS Mini

- Cisco UCS Mini uses a shared buffer for all system classes.

- The bronze class shares the buffer with SPAN. We recommend using either SPAN or the bronze class.

- Multicast optimization is not supported.

- Changing the QoS parameters for any class causes traffic disruption to all classes.

- When mixing Ethernet and FC or FCoE traffic, the bandwidth distribution is not equal.

- Multiple streams of traffic from the same class may not be distributed equally.

- Use the same CoS values for all no-drop policies to avoid any FC or FCoE performance issues.

- Only the platinum and gold classes support no-drop policies.

System Classes

Cisco UCS uses Data Center Ethernet (DCE) to handle all traffic inside a Cisco UCS domain. This industry standard enhancement to Ethernet divides the bandwidth of the Ethernet pipe into eight virtual lanes. Two virtual lanes are reserved for internal system and management traffic. You can configure quality of service

(QoS) for the other six virtual lanes. System classes determine how the DCE bandwidth in these six virtual lanes is allocated across the entire Cisco UCS domain.

Each system class reserves a specific segment of the bandwidth for a specific type of traffic, which provides a level of traffic management, even in an oversubscribed system. For example, you can configure the Fibre Channel Priority system class to determine the percentage of DCE bandwidth allocated to FCoE traffic.

The following table describes the system classes that you can configure.

Table 1: System Classes

System Class	Description
Platinum Gold Silver Bronze	A configurable set of system classes that you can include in the QoS policy for a service profile. Each system class manages one lane of traffic. All properties of these system classes are available for you to assign custom settings and policies. For Cisco UCS Mini, packet drop can only be disabled on the platinum and gold classes. Only one platinum and one gold class can be configured as a no drop class at a time.
Best Effort	A system class that sets the quality of service for the lane reserved for basic Ethernet traffic. Some properties of this system class are preset and cannot be modified. For example, this class has a drop policy that allows it to drop data packets if required. You cannot disable this system class.
Fibre Channel	A system class that sets the quality of service for the lane reserved for Fibre Channel over Ethernet traffic. Some properties of this system class are preset and cannot be modified. For example, this class has a no-drop policy that ensures it never drops data packets. You cannot disable this system class. **Note** FCoE traffic has a reserved QoS system class that should not be used by any other type of traffic. If any other type of traffic has a CoS value that is used by FCoE, the value is remarked to 0.

Quality of Service Policy

A quality of service (QoS) policy assigns a system class to the outgoing traffic for a vNIC or vHBA. This system class determines the quality of service for that traffic. For certain adapters, you can also specify additional controls on the outgoing traffic, such as burst and rate.

You must include a QoS policy in a vNIC policy or vHBA policy and then include that policy in a service profile to configure the vNIC or vHBA.

Flow Control Policy

Flow control policies determine whether the uplink Ethernet ports in a Cisco UCS domain send and receive IEEE 802.3x pause frames when the receive buffer for a port fills. These pause frames request that the transmitting port stop sending data for a few milliseconds until the buffer clears.

For flow control to work between a LAN port and an uplink Ethernet port, you must enable the corresponding receive and send flow control parameters for both ports. For Cisco UCS, the flow control policies configure these parameters.

When you enable the send function, the uplink Ethernet port sends a pause request to the network port if the incoming packet rate becomes too high. The pause remains in effect for a few milliseconds before traffic is reset to normal levels. If you enable the receive function, the uplink Ethernet port honors all pause requests from the network port. All traffic is halted on that uplink port until the network port cancels the pause request.

Because you assign the flow control policy to the port, changes to the policy have an immediate effect on how the port reacts to a pause frame or a full receive buffer.

Opt-In Features

Each Cisco UCS domain is licensed for all functionality. Depending upon how the system is configured, you can decide to opt in to some features or opt out of them for easier integration into existing environment. If a process change happens, you can change your system configuration and include one or both of the opt-in features.

The opt-in features are as follows:

- Stateless computing, which takes advantage of mobile service profiles with pools and policies where each component, such as a server or an adapter, is stateless.

- Multi-tenancy, which uses organizations and role-based access control to divide the system into smaller logical segments.

Stateless Computing

Stateless computing allows you to use a service profile to apply the personality of one server to a different server in the same Cisco UCS domain. The personality of the server includes the elements that identify that server and make it unique in the Cisco UCS domain. If you change any of these elements, the server could lose its ability to access, use, or even achieve booted status.

The elements that make up a server's personality include the following:

- Firmware versions

- UUID (used for server identification)

- MAC address (used for LAN connectivity)

- World Wide Names (used for SAN connectivity)

- Boot settings

Stateless computing creates a dynamic server environment with highly flexible servers. Every physical server in a Cisco UCS domain remains anonymous until you associate a service profile with it, then the server gets

the identity configured in the service profile. If you no longer need a business service on that server, you can shut it down, disassociate the service profile, and then associate another service profile to create a different identity for the same physical server. The "new" server can then host another business service.

To take full advantage of the flexibility of statelessness, the optional local disks on the servers should only be used for swap or temp space and not to store operating system or application data.

You can choose to fully implement stateless computing for all physical servers in a Cisco UCS domain, to not have any stateless servers, or to have a mix of the two types.

If You Opt In to Stateless Computing

Each physical server in the Cisco UCS domain is defined through a service profile. Any server can be used to host one set of applications, then reassigned to another set of applications or business services, if required by the needs of the data center.

You create service profiles that point to policies and pools of resources that are defined in the Cisco UCS domain. The server pools, WWN pools, and MAC pools ensure that all unassigned resources are available on an as-needed basis. For example, if a physical server fails, you can immediately assign the service profile to another server. Because the service profile provides the new server with the same identity as the original server, including WWN and MAC address, the rest of the data center infrastructure sees it as the same server and you do not need to make any configuration changes in the LAN or SAN.

If You Opt Out of Stateless Computing

Each server in the Cisco UCS domain is treated as a traditional rack mount server.

You create service profiles that inherit the identify information burned into the hardware and use these profiles to configure LAN or SAN connectivity for the server. However, if the server hardware fails, you cannot reassign the service profile to a new server.

Multitenancy

Multi-tenancy allows you to divide the large physical infrastructure of an Cisco UCS domain into logical entities known as organizations. As a result, you can achieve a logical isolation between organizations without providing a dedicated physical infrastructure for each organization.

You can assign unique resources to each tenant through the related organization in the multi-tenant environment. These resources can include different policies, pools, and quality of service definitions. You can also implement locales to assign or restrict user privileges and roles by organization, if you do not want all users to have access to all organizations.

If you set up a multi-tenant environment, all organizations are hierarchical. The top-level organization is always root. The policies and pools that you create in root are system-wide and are available to all organizations in the system. However, any policies and pools created in other organizations are only available to organizations that are above it in the same hierarchy. For example, if a system has organizations named Finance and HR that are not in the same hierarchy, Finance cannot use any policies in the HR organization, and HR cannot access any policies in the Finance organization. However, both Finance and HR can use policies and pools in the root organization.

If you create organizations in a mult-itenant environment, you can also set up one or more of the following for each organization or for a sub-organization in the same hierarchy:

- Resource pools
- Policies

- Service profiles
- Service profile templates

If You Opt In to Multitenancy

Each Cisco UCS domain is divided into several distinct organizations. The types of organizations you create in a multitenancy implementation depends upon the business needs of the company. Examples include organizations that represent the following:

- Enterprise groups or divisions within a company, such as marketing, finance, engineering, or human resources
- Different customers or name service domains, for service providers

You can create locales to ensure that users have access only to those organizations that they are authorized to administer.

If You Opt Out of Multitenancy

The Cisco UCS domain remains a single logical entity with everything in the root organization. All policies and resource pools can be assigned to any server in the Cisco UCS domain.

Virtualization in Cisco UCS

Overview of Virtualization

Virtualization allows you to create multiple Virtual Machines (VMs) to run in isolation, side by side on the same physical machine.

Each virtual machine has its own set of virtual hardware (RAM, CPU, NIC) upon which an operating system and fully configured applications are loaded. The operating system sees a consistent, normalized set of hardware regardless of the actual physical hardware components.

In a virtual machine, both hardware and software are encapsulated in a single file for rapid provisioning and moving between physical servers. You can move a virtual machine, within seconds, from one physical server to another for zero-downtime maintenance and continuous workload consolidation.

The virtual hardware makes it possible for many servers, each running in an independent virtual machine, to run on a single physical server. The advantages of virtualization include better use of computing resources, greater server density, and seamless server migration.

Overview of Cisco Virtual Machine Fabric Extender

A virtualized server implementation consists of one or more VMs that run as guests on a single physical server. The guest VMs are hosted and managed by a software layer called the hypervisor or virtual machine manager (VMM). Typically, the hypervisor presents a virtual network interface to each VM and performs Layer 2 switching of traffic from a VM to other local VMs or to another interface to the external network.

Working with a Cisco virtual interface card (VIC) adapter, the Cisco Virtual Machine Fabric Extender (VM-FEX) bypasses software-based switching of VM traffic by the hypervisor for external hardware-based

switching in the fabric interconnect. This method reduces the load on the server CPU, provides faster switching, and enables you to apply a rich set of network management features to local and remote traffic.

VM-FEX extends the IEEE 802.1Qbh port extender architecture to the VMs by providing each VM interface with a virtual Peripheral Component Interconnect Express (PCIe) device and a virtual port on a switch. This solution allows precise rate limiting and quality of service (QoS) guarantees on the VM interface.

Virtualization with Network Interface Cards and Converged Network Adapters

Network interface card (NIC) and converged network adapters support virtualized environments with the standard VMware integration with ESX installed on the server and all virtual machine management performed through the VC.

Portability of Virtual Machines

If you implement service profiles you retain the ability to easily move a server identity from one server to another. After you image the new server, the ESX treats that server as if it were the original.

Communication between Virtual Machines on the Same Server

These adapters implement the standard communications between virtual machines on the same server. If an ESX host includes multiple virtual machines, all communications must go through the virtual switch on the server.

If the system uses the native VMware drivers, the virtual switch is out of the network administrator's domain and is not subject to any network policies. As a result, for example, QoS policies on the network are not applied to any data packets traveling from VM1 to VM2 through the virtual switch.

If the system includes another virtual switch, such as the Nexus 1000, that virtual switch is subject to the network policies configured on that switch by the network administrator.

Virtualization with a Virtual Interface Card Adapter

A Cisco VIC adapter is a converged network adapter (CNA) that is designed for both bare metal and VM-based deployments. The VIC adapter supports static or dynamic virtualized interfaces, which includes up to 128 virtual network interface cards (vNICs).

There are two types of vNICs used with the VIC adapter—static and dynamic. A static vNIC is a device that is visible to the OS or hypervisor. Dynamic vNICs are used for VM-FEX by which a VM is connected to a veth port on the Fabric Interconnect.

VIC adapters support VM-FEX to provide hardware-based switching of traffic to and from virtual machine interfaces.