



Overview of Cisco Unified Computing System

This chapter includes the following sections:

- [About Cisco Unified Computing System](#) , page 1
- [Unified Fabric](#), page 2
- [Server Architecture and Connectivity](#), page 4
- [Traffic Management](#), page 16
- [Opt-In Features](#), page 20
- [Overview of Virtualization](#), page 23

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.

The result of this radical simplification is a reduction by up to two-thirds of the switches, cables, adapters, and management points. All devices in a Cisco UCS instance 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 instance 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 instance 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:

- 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 VN-Link 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 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.

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 BMC
- 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:

- 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 Cisco UCS CNA M71KR 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 Cisco UCS CNA M71KR has two HBAs, which means you can create a maximum of two vHBAs for each of those adapters. In contrast, a Cisco UCS 82598KR-CI 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 and 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 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:

- 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

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 to another server.

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

This service profile inherits and automatically 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 the Cisco UCS CNA M71KR adapters, the WWN addresses for the two HBAs
- BIOS versions
- Server UUID



Important

The server identity and configuration information inherited through this service profile may not be the values burned into the server hardware at manufacture if those values were changed before this profile is associated with the server.

Service Profile 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. However, after you create the profile, it is no longer connected to the template. If you need to make changes to one or more profiles created from this template, you must change each profile individually.
- 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.

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

Configuration Policies

Boot Policy

The boot policy determines the following:

- Configuration of the boot device
- Location from which the server boots
- Order in which boot devices are invoked

For example, you can choose to have associated servers boot from a local device, such as a local disk or CD-ROM (VMedia), or you can select a SAN boot or a LAN (PXE) boot.

You must include this policy in a service profile, and that service profile must be associated with a server for it to take effect. If you do not include a boot policy in a service profile, the server uses the default settings in the BIOS to determine the boot order.



Important

Changes to a boot policy may be propagated to all servers created with an updating service profile template that includes that boot policy. Reassociation of the service profile with the server to rewrite the boot order information in the BIOS is auto-triggered.

Guidelines

When you create a boot policy, you can add one or more of the following to the boot policy and specify their boot order:

Boot type	Description
SAN boot	Boots from an operating system image on the SAN. You can specify a primary and a secondary SAN boot. If the primary boot fails, the server attempts to boot from the secondary. We recommend that you use a SAN boot, because it offers the most service profile mobility within the system. If you boot from the SAN, when you move a service profile from one server to another, the new server boots from the exact same operating system image. Therefore, the new server appears to be the exact same server to the network.
LAN boot	Boots from a centralized provisioning server. It is frequently used to install operating systems on a server from that server.
Local disk boot	If the server has a local drive, boots from that drive.
Virtual media boot	Mimics the insertion of a physical CD-ROM disk (read-only) or floppy disk (read-write) into a server. It is typically used to manually install operating systems on a server.



Note The default boot order is as follows:

- 1 Local disk boot
- 2 LAN boot
- 3 Virtual media read-only boot
- 4 Virtual media read-write boot

Chassis Discovery Policy

This discovery policy determines how the system reacts when you add a new chassis. If you create a chassis discovery policy, the system does the following:

- Automatically configures the chassis for the number of links between the chassis and the fabric interconnect specified in the policy.
- Specifies the power policy to be used by the chassis.

Ethernet and Fibre Channel Adapter Policies

These policies govern the host-side behavior of the adapter, including how the adapter handles traffic. For example, you can use these policies to change default settings for the following:

- Queues

- Interrupt handling
- Performance enhancement
- RSS hash
- Failover in an cluster configuration with two fabric interconnects

By default, Cisco UCS provides a set of Ethernet adapter policies and Fibre Channel adapter policies. These policies include the recommended settings for each supported server operating system. Operating systems are sensitive to the settings in these policies. Storage vendors typically require non-default adapter settings. You can find the details of these required settings on the support list provided by those vendors.

**Note**

For Fibre Channel adapter policies, the values displayed by Cisco UCS Manager may not match those displayed by applications such as QLogic SANsurfer. For example, the following values may result in an apparent mismatch between SANsurfer and Cisco UCS Manager:

- Max LUNs Per Target—SANsurfer has a maximum of 256 LUNs and does not display more than that number. Cisco UCS Manager supports a higher maximum number of LUNs.
 - Link Down Timeout—In SANsurfer, you configure the timeout threshold for link down in seconds. In Cisco UCS Manager, you configure this value in milliseconds. Therefore, a value of 5500 ms in Cisco UCS Manager displays as 5s in SANsurfer.
 - Max Data Field Size—SANsurfer has allowed values of 512, 1024, and 2048. Cisco UCS Manager allows you to set values of any size. Therefore, a value of 900 in Cisco UCS Manager displays as 512 in SANsurfer.
-

Host Firmware Pack

This policy enables you to specify a common set of firmware versions that make up the host firmware pack. The host firmware includes the following server and adapter components:

- BIOS
- SAS controller
- Emulex Option ROM (applicable only to Emulex-based Converged Network Adapters [CNAs])
- Emulex firmware (applicable only to Emulex-based CNAs)
- QLogic option ROM (applicable only to QLogic-based CNAs)
- Adapter firmware

The firmware pack is pushed to all servers associated with service profiles that include this policy.

This policy ensures that the host firmware is identical on all servers associated with service profiles which use the same policy. Therefore, if you move the service profile from one server to another, the firmware versions are maintained. Also, if you change the firmware version of the component in the firmware pack, new versions are applied to all the affected service profiles immediately, which could cause server reboots.

You must include this policy in a service profile, and that service profile must be associated with a server for it to take effect.

Prerequisites

This policy is not dependent upon any other policies. However, you must ensure that the appropriate firmware has been downloaded to the fabric interconnect. If the firmware image is not available when Cisco UCS Manager is associating a server with a service profile, Cisco UCS Manager ignores the firmware update and completes the association.

IPMI Access Profile

This policy allows you to determine whether IPMI commands can be sent directly to the server, using the IP address. For example, you can send commands to retrieve sensor data from the BMC. This policy defines the IPMI access, including a username and password that can be authenticated locally on the server, and whether the access is read-only or read-write.

You must include this policy in a service profile and that service profile must be associated with a server for it to take effect.

Local Disk Configuration Policy

This policy configures any optional SAS local drives that have been installed on a server through the onboard RAID controller of the local drive. This policy enables you to set a local disk mode for all servers that are associated with a service profile that includes the local disk configuration policy. The local disk modes include the following:

- **Any Configuration**—For a server configuration that carries forward the local disk configuration without any changes.
- **No Local Storage**—For a diskless workstation or a SAN only configuration. If you select this option, you cannot associate any service profile which uses this policy with a server that has a local disk.
- **No RAID**—For a server configuration that removes the RAID and leaves the disk MBR and payload unaltered.
- **RAID Mirrored**—For a 2-disk RAID 1 server configuration.
- **RAID Stripes**—For a 2-disk RAID 0 server configuration.

You must include this policy in a service profile, and that service profile must be associated with a server for it to take effect.

Management Firmware Pack

This policy enables you to specify a common set of firmware versions that make up the management firmware pack. The management firmware includes the server controller (BMC) on the server.

The firmware pack is pushed to all servers associated with service profiles that include this policy.

This policy ensures that the BMC firmware is identical on all servers associated with service profiles which use the same policy. Therefore, if you move the service profile from one server to another, the firmware versions are maintained.

You must include this policy in a service profile, and that service profile must be associated with a server for it to take effect.

Prerequisites

This policy is not dependent upon any other policies. However, you must ensure that the appropriate firmware has been downloaded to the fabric interconnect.

Network Control Policy

This policy configures the network control settings for the Cisco UCS instance, including the following:

- Whether the Cisco Discovery Protocol (CDP) is enabled or disabled
- How the VIF behaves if no uplink port is available in end-host mode

Quality of Service Policies

QoS policies assign a system class to the outgoing traffic for a vNIC or vHBA. This system class determines the quality of service for that traffic.

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.

Server Autoconfiguration Policy

This policy determines whether one or more of the following is automatically applied to a new server:

- A server pool policy qualification that qualifies the server for one or more server pools
- An organization
- A service profile template that associates the server with a service profile created from that template

Server Discovery Policy

This discovery policy determines how the system reacts when you add a new server. If you create a server discovery policy, you can control whether the system conducts a deep discovery when a server is added to a chassis, or whether a user must first acknowledge the new server. By default, the system conducts a full discovery.

With this policy, an inventory of the server is conducted, then server pool policy qualifications are run to determine whether the new server qualifies for one or more server pools.

Server Inheritance Policy

This policy is invoked during the server discovery process to create a service profile for the server. All service profiles created from this policy use the values burned into the blade at manufacture. The policy performs the following:

- Analyzes the inventory of the server
- If configured, assigns the server to the selected organization
- Creates a service profile for the server with the identity burned into the server at manufacture

You cannot migrate a service profile created with this policy to another server.

Server Pool Policy

This policy is invoked during the server discovery process. It determines what happens if server pool policy qualifications match a server to the target pool specified in the policy.

If a server qualifies for more than one pool and those pools have server pool policies, the server is added to all those pools.

Server Pool Policy Qualifications

This policy qualifies servers based on the inventory of a server conducted during the discovery process. The qualifications are individual rules that you configure in the policy to determine whether a server meets the selection criteria. For example, you can create a rule that specifies the minimum memory capacity for servers in a data center pool.

Qualifications are used in other policies to place servers, not just by the server pool policies. For example, if a server meets the criteria in a qualification policy, it can be added to one or more server pools or have a service profile automatically associated with it.

Depending upon the implementation, you may include server pool policy qualifications in the following policies:

- Autoconfiguration policy
- Chassis discovery policy
- Server discovery policy
- Server inheritance policy
- Server pool policy

vHBA Template

This template is a policy that defines how a vHBA on a server connects to the SAN. It is also referred to as a vHBA SAN connectivity template.

You need to include this policy in a service profile for it to take effect.

vNIC Template

This policy defines how a vNIC on a server connects to the LAN. This policy is also referred to as a vNIC LAN connectivity policy.

You need to include this policy in a service profile for it to take effect.

Operational Policies

Fault Collection Policy

The fault collection policy controls the lifecycle of a fault in a Cisco UCS instance, including when faults are cleared, the flapping interval (the length of time between the fault being raised and the condition being cleared), and the retention interval (the length of time a fault is retained in the system).

A fault in Cisco UCS has the following lifecycle:

- 1 A condition occurs in the system and Cisco UCS Manager raises a fault. This is the active state.
- 2 When the fault is alleviated, it is cleared if the time between the fault being raised and the condition being cleared is greater than the flapping interval, otherwise, the fault remains raised but its status changes to soaking-clear. Flapping occurs when a fault is raised and cleared several times in rapid succession. During the flapping interval the fault retains its severity for the length of time specified in the fault collection policy.
- 3 If the condition reoccurs during the flapping interval, the fault remains raised and its status changes to flapping. If the condition does not reoccur during the flapping interval, the fault is cleared.
- 4 When a fault is cleared, it is deleted if the clear action is set to delete, or if the fault was previously acknowledged, otherwise, it is retained until either the retention interval expires, or if the fault is acknowledged.
- 5 If the condition reoccurs during the retention interval, the fault returns to the active state. If the condition does not reoccur, the fault is deleted.

Scrub Policy

This policy determines what happens to local data on a server during the discovery process and when the server is disassociated from a service profile. This policy can ensure that the data on local drives is erased at those times.

Serial over LAN Policy

This policy sets the configuration for the serial over LAN connection for all servers associated with service profiles that use the policy. By default, the serial over LAN connection is disabled.

If you implement a serial over LAN policy, we recommend that you also create an IPMI profile.

You must include this policy in a service profile and that service profile must be associated with a server for it to take effect.

Statistics Collection Policy

A statistics collection policy defines how frequently statistics are to be collected (collection interval), and how frequently the statistics are to be reported (reporting interval). Reporting intervals are longer than collection intervals so that multiple statistical data points can be collected during the reporting interval, which provides Cisco UCS Manager with sufficient data to calculate and report minimum, maximum, and average values.

Statistics can be collected and reported for the following five functional areas of the Cisco UCS system:

- Adapter—statistics related to the adapters
- Chassis—statistics related to the blade chassis
- Host—this policy is a placeholder for future support
- Port—statistics related to the ports, including server ports, uplink Ethernet ports, and uplink Fibre Channel ports
- Server—statistics related to servers

**Note**

Cisco UCS Manager has one default statistics collection policy for each of the five functional areas. You cannot create additional statistics collection policies and you cannot delete the existing default policies. You can only modify the default policies.

Statistics Threshold Policy

A statistics threshold policy monitors statistics about certain aspects of the system and generates an event if the threshold is crossed. You can set both minimum and maximum thresholds. For example, you can configure the policy to raise an alarm if the CPU temperature exceeds a certain value, or if a server is overutilized or underutilized.

These threshold policies do not control the hardware or device-level thresholds enforced by endpoints, such as the BMC. Those thresholds are burned in to the hardware components at manufacture.

Cisco UCS enables you to configure statistics threshold policies for the following components:

- Servers and server components
- Uplink Ethernet ports
- Ethernet server ports, chassis, and fabric interconnects
- Fibre Channel port

**Note**

You cannot create or delete a statistics threshold policy for Ethernet server ports, uplink Ethernet ports, or uplink Fibre Channel ports. You can only configure the existing default policy.

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 pre-assign ranges for servers that will host specific applications. For example, all database servers could be configured within the same range of MAC addresses, UUIDs, and WWNs.

Server Pools

A server pool contains a set of servers. These servers typically share the same characteristics. Those characteristics can be their location in the chassis, or an attribute such as server type, amount of memory, local storage, type of CPU, or local drive configuration. You can manually assign a server to a server pool, or use server pool policies and server pool policy qualifications to automate the assignment.

If your system implements multi-tenancy through organizations, you can designate one or more server pools to be used by a specific organization. For example, a pool that includes all servers with two CPUs could be assigned to the Marketing organization, while all servers with 64 GB memory could be assigned to the Finance organization.

A server pool can include servers from any chassis in the system. A given server can belong to multiple server pools.

MAC Pools

A MAC pool is a collection of network identities, or MAC addresses, that are unique in their layer 2 environment and are available to be assigned to vNICs on a server. If you use MAC pools in service profiles, you do not have to manually configure the MAC addresses to be used by the server associated with the service profile.

In a system that implements multi-tenancy, you can use the organizational hierarchy to ensure that MAC pools can only be used by specific applications or business services. Cisco UCS Manager uses the name resolution policy to assign MAC addresses from the pool.

To assign a MAC address to a server, you must include the MAC pool in a vNIC policy. The vNIC policy is then included in the service profile assigned to that server.

You can specify your own MAC addresses or use a group of MAC addresses provided by Cisco.

UUID Suffix Pools

A UUID suffix pool is a collection of SMBIOS UUIDs that are available to be assigned to servers. The first number of digits that constitute the prefix of the UUID are fixed. The remaining digits, the UUID suffix, is variable. A UUID suffix pool ensures that these variable values are unique for each server associated with a service profile which uses that particular pool to avoid conflicts.

If you use UUID suffix pools in service profiles, you do not have to manually configure the UUID of the server associated with the service profile.

WWN Pools

A WWN pool is a collection of WWNs for use by the Fibre Channel vHBAs in a Cisco UCS instance. You create separate pools for the following:

- WW node names assigned to the server
- WW port names assigned to the vHBA



Important

A WWN pool can include only WWNNs or WWPNs in the ranges from 20:00:00:00:00:00:00:00 to 20:FF:FF:FF:FF:FF:FF:FF or from 50:00:00:00:00:00:00:00 to 5F:FF:FF:FF:FF:FF:FF:FF. All other WWN ranges are reserved.

If you use WWN pools in service profiles, you do not have to manually configure the WWNs that will be used by the server associated with the service profile. In a system that implements multi-tenancy, you can use a WWN pool to control the WWNs used by each organization.

You assign WWNs to pools in blocks. For each block or individual WWN, you can assign a boot target.

WWNN Pools

A WWNN pool is a WWN pool which contains only WW node names. If you include a pool of WWNNs in a service profile, the associated server is assigned a WWNN from that pool.

WWPN Pools

A WWPN pool is a WWN pool which contains only WW port names. If you include a pool of WWPNs in a service profile, the port on each vHBA of the associated server is assigned a WWPN from that pool.

Management IP Pool

The management IP pool is a collection of external IP addresses. Cisco UCS Manager reserves each block of IP addresses in the management IP pool for external access that terminates in the server controller (BMC) in a server.

Cisco UCS Manager uses the IP addresses in a management IP pool for external access to a server through the following:

- KVM console
- Serial over LAN
- IPMI

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:

The 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.

The number of uplink ports from the fabric interconnect to the 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.

FC uplink ports are available on the expansion slots only. You must add more expansion slots to increase number of available FC uplinks. Ethernet uplink ports can exist on the fixed slot and on expansion slots.

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.

The number of uplink ports from the I/O module to the 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 fabric interconnect. The number of cables determines the number of active uplink ports and the oversubscription ratio. For example, one cable results in 8:1 oversubscription for one I/O module. If two I/O modules are in place, each with one cable, and you have 8 half-width blades, the 8 blades will be sharing two uplinks (one left IOM and one right IOM). This results in 8 blades sharing an aggregate bandwidth of 20 GB of Unified Fabric capacity. If two cables are used, this results in 4:1 oversubscription per IOM (assuming all slots populated with half width blades), and four cables result in 2:1 oversubscription. The lower oversubscription ratio gives you higher performance, but is also more costly as you consume more fabric interconnect ports.

The number of active links from the server to the fabric interconnect

Oversubscription is affected by how many servers are in a particular chassis and how bandwidth intensive those servers are. The oversubscription ratio will be reduced if the servers which generate a large amount of traffic are not in the same chassis, but are shared between the chassis in the system. The number of cables between chassis and fabric interconnect determines the oversubscription ratio. For example, one cable results in 8:1 oversubscription, two cables result in 4:1 oversubscription, and four cables result in 2:1 oversubscription. The lower oversubscription ratio gives you higher performance, but it is also more costly.

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.

Pinning Server Traffic to Server Ports

All server traffic travels through the I/O module to server ports on the fabric interconnect. The number of links for which the chassis is configured determines how this traffic is pinned.

The pinning determines which server traffic goes to which server port on the fabric interconnect. This pinning is fixed. You cannot modify it. As a result, you must consider the server location when you determine the appropriate allocation of bandwidth for a chassis.



Note

You must review the allocation of ports to links before you allocate servers to slots. The cabled ports are not necessarily port 1 and port 2 on the I/O module. If you change the number of links between the fabric interconnect and the I/O module, you must reacknowledge the chassis to have the traffic rerouted.

All port numbers refer to the fabric interconnect-side ports on the I/O module.

Chassis with One I/O Module

Links on Chassis	Servers Pinned to Link 1	Servers Pinned to Link 2	Servers Pinned to Link 3	Servers Pinned to Link 4
1 link	All server slots	None	None	None
2 links	Slots 1, 3, 5, and 9	Slots 2, 4, 6, and 8	None	None
4 links	Slots 1 and 5	Slots 2 and 6	Slots 3 and 7	Slots 4 and 8

Chassis with Two I/O Modules

If a chassis has two I/O modules, traffic from one I/O module goes to one of the fabric interconnects and traffic from the other I/O module goes to the second fabric interconnect. You cannot connect two I/O modules to a single fabric interconnect.

Adding a second I/O module to a chassis does not improve oversubscription. The server port pinning is the same for a single I/O module. The second I/O module improves the high availability of the system through the vNIC binding to the fabric interconnect.

Fabric Interconnect Configured in vNIC	Server Traffic Path
A	Server traffic goes to fabric interconnect A. If A fails, the server traffic does not fail over to B.
B	All server traffic goes to fabric interconnect B. If B fails, the server traffic does not fail over to A.
A-B	All server traffic goes to fabric interconnect A. If A fails, the server traffic fails over to B.
B-A	All server traffic goes to fabric interconnect B. If B fails, the server traffic fails over to A.

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

System Classes

Cisco UCS uses Data Center Ethernet (DCE) to handle all traffic inside a Cisco UCS instance. This industry standard enhancement to Ethernet divides the bandwidth of the Ethernet pipe into eight virtual lanes. System classes determine how the DCE bandwidth in these virtual lanes is allocated across the entire Cisco UCS instance.

Each system class reserves a specific segment of the bandwidth for a specific type of traffic. This 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:

System Class	Description
Platinum Priority Gold Priority Silver Priority Bronze Priority	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.
Best Effort Priority	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.
Fibre Channel Priority	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.

Quality of Service Policies

QoS policies assign a system class to the outgoing traffic for a vNIC or vHBA. This system class determines the quality of service for that traffic.

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 Policies

Flow control policies determine whether the uplink Ethernet ports in a Cisco UCS instance 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 instance 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 instance. The personality of the server includes the elements that identify that server and make it unique in the instance. 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 instance 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 instance, 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 instance 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 instance. 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 instance 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.

Multi-Tenancy

In Cisco UCS, you can use multi-tenancy to divide up the large physical infrastructure of an instance 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 Cisco UCS 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 multi-tenant 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 Multi-Tenancy

The Cisco UCS instance is divided into several distinct organizations. The types of organizations you create in a multi-tenancy 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 Multi-Tenancy

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

Overview of Virtualization

Virtualization allows the creation of multiple virtual machines 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 copying, 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.

Virtualization with the Cisco UCS CNA M71KR and Cisco UCS 82598KR-CI Adapters

The Cisco UCS 82598KR-CI 10-Gigabit Ethernet Adapter, Cisco UCS M71KR - E Emulex Converged Network Adapter, and Cisco UCS M71KR - Q QLogic Converged Network Adapter support virtualized environments with the following VMware versions:

- VMware 3.5 update 4
- VMware 4.0

These environments support 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, quality of service 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.

