Cisco UCS Manager GUI Configuration Guide, Release 2.2

First Published: 2013-12-11

Last Modified: 2017-04-17

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

CONTENTS

Preface

Audience

This guide is intended primarily for data center administrators with responsibilities and expertise in one or more of the following:

- Server administration
- Storage administration
- Network administration
- Network security

Conventions

Text Type	Indication
GUI elements	GUI elements such as tab titles, area names, and field labels appear in **this font**. Main titles such as window, dialog box, and wizard titles appear in **this font**.
Document titles	Document titles appear in *this font*.
TUI elements	In a Text-based User Interface, text the system displays appears in `this font`.
System output	Terminal sessions and information that the system displays appear in `this font`.

Text Type	Indication
CLI commands	CLI command keywords appear in **this font**. Variables in a CLI command appear in *this font*.
[ ]	Elements in square brackets are optional.
{x | y | z}	Required alternative keywords are grouped in braces and separated by vertical bars.
[x | y | z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.
< >	Nonprinting characters such as passwords are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

Note Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the document.

Tip Means *the following information will help you solve a problem*. The tips information might not be troubleshooting or even an action, but could be useful information, similar to a Timesaver.

Timesaver Means *the described action saves time*. You can save time by performing the action described in the paragraph.

Caution Means *reader be careful*. In this situation, you might perform an action that could result in equipment damage or loss of data.

Warning IMPORTANT SAFETY INSTRUCTIONS

This warning symbol means danger. You are in a situation that could cause bodily injury. Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents. Use the statement number provided at the end of each warning to locate its translation in the translated safety warnings that accompanied this device.

SAVE THESE INSTRUCTIONS

Related Cisco UCS Documentation

Documentation Roadmaps

For a complete list of all B-Series documentation, see the *Cisco UCS B-Series Servers Documentation Roadmap* available at the following URL: http://www.cisco.com/go/unifiedcomputing/b-series-doc.

For a complete list of all C-Series documentation, see the *Cisco UCS C-Series Servers Documentation Roadmap* available at the following URL: http://www.cisco.com/go/unifiedcomputing/c-series-doc.

For information on supported firmware versions and supported UCS Manager versions for the rack servers that are integrated with the UCS Manager for management, refer to Release Bundle Contents for Cisco UCS Software.

Other Documentation Resources

Follow Cisco UCS Docs on Twitter to receive document update notifications.

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to ucs-docfeedback@cisco.com. We appreciate your feedback.

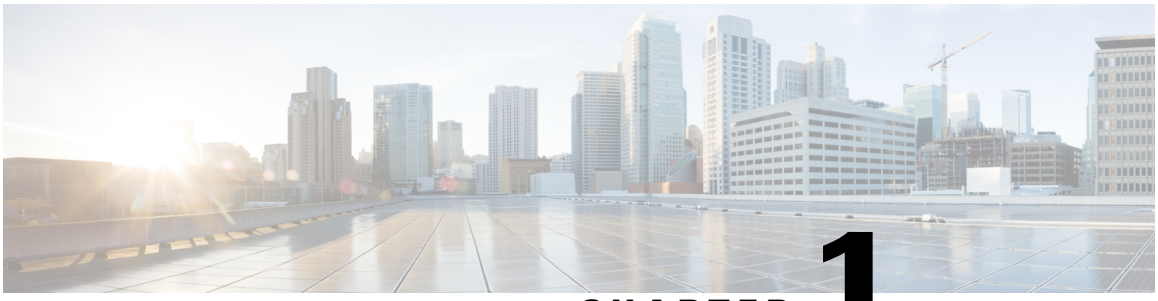

CHAPTER 1

New and Changed Information

This chapter includes the following sections:

New and Changed Information for this Release

The following table provides an overview of the significant changes to this guide for this current release. The table does not provide an exhaustive list of all changes made to the configuration guides or of the new features in this release. For information about new supported hardware in this release, see the *Cisco UCS B-Series Servers Documentation Roadmap* available at the following URL: http://www.cisco.com/go/unifiedcomputing/b-series-doc.

Table 1: New Features and Changed Behavior in Cisco UCS, Release 2.2(1)

Feature	Description	Where Documented
Cisco UCS C-Series Server Integration through Direct Connect	Enables you to directly connect Cisco UCS C-Series rack servers to the fabric interconnect without using a FEX.	This feature is documented in *Cisco UCS C-Series Server Integration with Cisco UCS Manager 2.2*. The C-Series integration guides can be found here: http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-c-series-rack-mount-ucs-manager-software/products-installation-and-configuration-guides-list.html

Feature	Description	Where Documented
Activating board controller firmware on Cisco UCS C-Series M3 Rack Servers	Cisco UCS C-Series M3 rack servers have board controller firmware which controls many server functions such as eUSBs and I/O connectors. You can activate board controller firmware on these rack servers.	This feature is documented in the following configuration guides: • *Cisco UCS B-Series GUI Firmware Management Configuration Guide* • *Cisco UCS B-Series CLI Firmware Management Configuration Guide* The firmware configuration guides can be found here: http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html
Firmware Automatic Synchronization Server policy	Enables you to determine when and how firmware versions on recently discovered servers must be upgraded. With this policy, you can upgrade firmware versions of recently discovered unassociated servers to match with the firmware version specified in the default host firmware pack. You can also specify if the firmware upgrade process should run immediately after the server is discovered or at a later point in time.	This feature is documented in the following configuration guides: • *Cisco UCS B-Series GUI Firmware Management Configuration Guide* • *Cisco UCS B-Series CLI Firmware Management Configuration Guide* The firmware configuration guides can be found here: http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html

Feature	Description	Where Documented
IPv6 Addressing	Enables you to configure the following with IPv6 addresses: • Access services • Fabric Interconnects • External clients • SNMP traps • Certificate requests • IP pools and address blocks • VLAN groups • LDAP, RADIUS, and TACACS+ provider groups • Blade and rack server static IPs • Service profiles and service profile templates • Backup and restore operations	IPv6 Compliance, on page 20
Inband Management Support	Enables you to configure inband addresses, two IPv4 and two IPv6, for each CIMC. Enables you to configure inband VLAN groups, service profiles and service profile templates.	Inband Management Support, on page 27
Enhanced Boot Order	Enables you to select either legacy or UEFI boot mode. With UEFI boot mode on Cisco UCS M3 blade and rack servers, you can add second-level devices to your boot order, and enable secure boot.	Configuring Server Boot, on page 523

Feature	Description	Where Documented
Local Storage Monitoring	Enables you to monitor status information on local storage that is physically attached to a blade or rack server. This includes RAID controllers, physical drives and drive groups, virtual drives, RAID controller batteries (BBU), Transportable Flash modules (TFM) and super-capacitors, FlexFlash controllers, and SD cards.	This feature is documented in the following configuration guides: • *Cisco UCS B-Series GUI System Monitoring Guide* • *Cisco UCS B-Series CLI System Monitoring Guide* The system monitoring guides can be found here: http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html
Local SD Card Monitoring and FlexFlash Support	Enables support for internal Secure Digital (SD) memory cards. You can enable FlexFlash in a local disk policy, configure new SD cards in a RAID pair with a FlexFlash scrub policy, and boot from the HV partition on an SD card.	FlexFlash Support, on page 477
TPM Inventory	Enables monitoring of Trusted Platform Module (TPM), including whether TPM is present, enabled, or activated.	This feature is documented in the following configuration guides: • *Cisco UCS B-Series GUI System Monitoring Guide* • *Cisco UCS B-Series CLI System Monitoring Guide* The system monitoring guides can be found here: http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html
UDLD Support	Enables UniDirectional Link Detection (UDLD) that monitors the physical configuration of the cables and detects when a unidirectional link exists.	Understanding UDLD, on page 314

Feature	Description	Where Documented
VMQ Connection Policy	Enables you to configure a VMQ connection policy for a vNIC. VMQ provides improved network performance to the entire management operating system.	VMQ Connection Policy, on page 319
DIMM Blacklisting Support	Monitors memory test execution messages, and blacklists any DIMMs that encounter memory errors in the DIMM SPD data. This allows the host to map out any DIMMs that encounter uncorrectable ECC errors.	DIMM Blacklisting, on page 485
DIMM Correctable Error Handling	Enables you to reset the correctable and uncorrectable memory errors on all the DIMMs in a server.	DIMM Correctable Error Handling, on page 484
Two-factor Authentication	Enables you to configure more secure access to Cisco UCS Manager by implementing a password plus token login scheme for RADIUS and TACACS+ provider groups when coupled with third-party authentication software.	Two-Factor Authentication, on page 144
KVM Direct Access	Enables users to directly access blade and rack servers using a web browser. Only out-of-band IPv4 management interface addresses are supported for KVM Direct access.	KVM Direct Access, on page 722

Feature	Description	Where Documented
VM-FEX Integration for Hyper-V SRIOV	Cisco UCS Manager and Microsoft SCVMM integration extends the VM-FEX technology to the Microsoft virtualization platform. The architecture allows Cisco UCS Manager to configure the networking objects that Microsoft SCVMM uses to set up its networking stacks. Microsoft SCVMM consumes the networking objects that are created by Cisco UCSM and deploys them on the Microsoft Hyper-V host that hosts the VMs. Hyper-V uses SR-IOV technology to deploy virtual connections. SR-IOV support in Release 2.2 enriches the management plane integration with Microsoft SCVMM and provides a centralized VM network management for the Hyper-V hosts. The deployment leverages the SR-IOV technology that is available on the Cisco UCS VIC adapters and enables the fabric interconnects to be VM-aware.	This feature is documented in the following configuration guides: • *Cisco UCS Manager VM-FEX for Hyper-V GUI Configuration Guide* • *Cisco UCS Manager VM-FEX for Hyper-V CLI Configuration Guide* The VM-FEX configuration guides can be found here: http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html

Table 2: New Features and Changed Behavior in Cisco UCS, Release 2.2(2)

Feature	Description	Where Documented
Netflow Monitoring	Enables you to collect IP traffic data from netflow capable routers.	This feature is documented in the following configuration guides: • *Cisco UCS B-Series GUI System Monitoring Guide* • *Cisco UCS B-Series CLI System Monitoring Guide* The system monitoring guides can be found here: http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html

Feature	Description	Where Documented
LACP	Enables you to use the link aggregation control protocol (LACP) policy to provide additional control over link aggregation groups.	LACP Policy, on page 313
Wear Level Monitoring	Enables you to monitor the life span of solid state drives on certain Cisco UCS blade servers.	This feature is documented in the following configuration guides: • *Cisco UCS B-Series GUI System Monitoring Guide* • *Cisco UCS B-Series CLI System Monitoring Guide* The system monitoring guides can be found here: http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html
CIMC Security Enhancements	Enables you to restrict remote connectivity and use vMedia encryption.	CIMC Security Policies, on page 467
Graphics Card Monitoring	Enables you to view the properties of graphics cards and controllers on certain Cisco UCS rack servers.	This feature is documented in the following configuration guides: • *Cisco UCS B-Series GUI System Monitoring Guide* • *Cisco UCS B-Series CLI System Monitoring Guide* The system monitoring guides can be found here: http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html

Feature	Description	Where Documented
Auto Install Firmware Enhancements	Enables you to view a list of warnings and potential issues before upgrading.	This feature is documented in the following configuration guides: • *Cisco UCS Manager VM-FEX for Hyper-V GUI Configuration Guide* • *Cisco UCS Manager VM-FEX for Hyper-V CLI Configuration Guide* The VM-FEX configuration guides can be found here: http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html
KVM Enhancements	Redesign of virtual KVM console now allows menu access to all functions, including virtual media.	Virtual KVM Console , on page 718
Scriptable vMedia Support	Enables you to mount virtual media from a remote file location and is available through CLI and WebGUI CIMC interfaces. This feature supports multiple share types including NFS, CIFS, HTTP and HTTPS.	CIMC Mounted vMedia, on page 514
Community VLAN Support	Enables you to create community VLANs on Fabric Interconnects and maintain multiple community VLANs under a given primary VLAN. Also provides support for promiscuous access and truck modes along with community access for server vNICs.	Community VLANs , on page 251

Table 3: New Features and Changed Behavior in Cisco UCS, Release 2.2(3)

Feature	Description	Where Documented
Anonymous Reporting	Enables you to retrieve anonymous reports from the SMTP server. This feature can be enabled even when call home is disabled.	This feature is documented in the following configuration guides: • *Cisco UCS B-Series GUI System Monitoring Guide* • *Cisco UCS B-Series CLI System Monitoring Guide* The system monitoring guides can be found here: http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html
DIMM Blacklisting Support	Now supports Cisco C-Series rack server.	DIMM Blacklisting, on page 485
C-Direct Rack Licenses	Now C-Direct rack licenses are supported on the ports that are connected to the rack servers.	C-Direct Rack Licensing Support, on page 213
CIMC Secure Boot	Allows only Cisco signed firmware images to be installed and run on the Cisco C-Series rack servers.	CIMC Secure Boot, on page 525
UEFI Secure Boot	Now supports Cisco C-Series rack server.	Configuring Server Boot, on page 523
FX3S Controller	Updated controller for SD cards allows you to reset the controller, format the SD cards, and enable automatic synchronization of your paired SD cards.	FlexFlash FX3S Support, on page 479
Stateless offload for Overlay Networks (NVGRE)	Enables you to create Ethernet adapter policies for checksum offloads using NVGRE.	Configuring an Ethernet Adapter Policy to Enable Stateless Offloads with NVGRE, on page 299
Stateless offloads for Overlay Networks (VXLAN)	Enables you to create Ethernet adapter policies for checksum offloads using VXLAN.	Configuring an Ethernet Adapter Policy to Enable Stateless Offloads with VXLAN, on page 299

Table 4: New Features and Changed Behavior in Cisco UCS, Release 2.2(4)

Feature	Description	Where Documented
RDMA over Converged Ethernet (RoCE) Support for Microsoft SMB Direct.	Enables communication between any two hosts in the same Ethernet broadcast domain. RoCE delivers superior performance compared to traditional network socket implementations because of lower latency, lower CPU utilization and higher utilization of network bandwidth. Microsoft SMB Direct with RoCE is supported only on Windows 2012 R2.	RDMA Over Converged Ethernet for SMB Direct, on page 293
LLDP Support for Fabric Interconnect vEthernet Interfaces	When using Cisco ACI, LAN uplinks are connected to ACI leaf nodes. Cisco UCS Manager Release 2.2.4 allows you to enable and disable LLDP on a vEthernet interface.	Configuring Link Layer Discovery Protocol for Fabric Interconnect vEthernet Interfaces, on page 309
Policy-Based Port Error Handling	If Cisco UCS Manager detects any errors on active NI ports, and if the error-disable feature is enabled, Cisco UCS Manager automatically disables the respective FI port that is connected to the NI port that had errors.	Policy-Based Port Error Handling, on page 109
Advanced Local Storage Configuration	This enhancement enabled you to do the following: • Configure multiple virtual drives. • Create and use storage profiles to allow flexibility in defining the number of storage disks, roles and usage of these disks, and other storage parameters. • Specify a local LUN or a JBOD disk as the primary boot device for a storage controller. • Configure local storage on multiple controllers. • Configure out-of-band local storage.	Storage Profiles, on page 619

Feature	Description	Where Documented
Automatic Internal Backup	Creates an automatic, full state, internal backup file when the infrastructure firmware is being upgraded.	This feature is documented in the following configuration guide: • *Cisco UCS B-Series CLI Firmware Management Configuration Guide* The firmware configuration guides can be found here: http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html
Fabric Interconnect Traffic Evacuation	During upgrade, enables you to evacuate a Fabric Interconnect to ensure that there is no traffic flowing through the Fabric Interconnect from all servers attached to it through an IOM or FEX.	Fabric Evacuation, on page 72
IOM Acknowledgment	Enables you to acknowledge specific IOMs in a chassis.	Acknowledging an IO Module, on page 733
Trusted Platform Module (TPM) Support	Allows you to enable and disable TPM and TXT. Cisco UCS M4 blade and rack-mount servers include support for TPM and TXT.	Trusted Platform Module, on page 462
Consistent Device Naming (CDN) Support	Allows Ethernet interfaces to be named in a consistent manner. CDN is supported only on Windows 2012 R2.	Consistent Device Naming, on page 463
Scriptable vMedia Enhancements	Variable mapped file name—Enables you either manually specify the file name of the vMedia mount image or automatically assign it the name of the service profile with which the vMedia policy is associated. CIFS authentication protocol support—Enables you to select the protocol to be used for authentication when you use CIFS as the protocol for communicating with the remote server.	Creating a vMedia Policy, on page 515

Feature	Description	Where Documented
SNMP Host Name Enhancement	Enables the use of a fully qualified domain name of an IPv4 address as the SNMP host name while creating an SNMP trap.	Creating an SNMP Trap, on page 135
Server Pool Policy Qualification Enhancement	Enables you to specify the storage disk type as HDD, SSD or unspecified when creating a server pool policy qualification.	Creating Server Pool Policy Qualifications, on page 496
Support for Service Profile Migration with UEFI Boot Mode	When a service profile is migrated from one server to another server, the BIOS on the destination server continues to load the boot loader information and boot in UEFI boot mode.	UEFI Boot Parameters, on page 563
NVGRE with IPv6 and VMQ Support	Provides the ability to enable VMQ and NVGRE offloading on the same vNIC.	Enabling VMQ and NVGRE Offloading on the same vNIC, on page 321
usNIC Support with Intel MPI	Provides the environment setup and installation guidance to use Cisco usNIC with Open and Intel® MPI technologies.	This feature is documented in the following configuration guides: • *Cisco usNIC Deployment Guide for Cisco UCS B-Series Blade Servers* • *Cisco usNIC Deployment Guide for Cisco UCS C-Series Rack-Mount Standalone Servers* The usNIC deployment guides can be found here: • http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html • http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-c-series-integrated-management-controller/products-installation-and-configuration-guides-list.html

Table 5: New Features and Changed Behavior in Cisco UCS Release 2.2(7)

Feature	Description	Where Documented
Firmware Upgrade Checks the VIF/Interface Status After Fabric Interconnect Reboot	Cisco UCS Manager displays any service that is not re-established after the last reboot of a fabric interconnect.	This feature is documented in the following configuration guides: • *Cisco UCS B-Series GUI Firmware Management Configuration Guide* • *Cisco UCS B-Series CLI Firmware Management Configuration Guide* The firmware configuration guides can be found here: http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html
vNIC Redundancy Pair	Supports two vNICs/vHBAs that are being configured with a common set of parameters through the vNIC/vHBA template pair.	vNIC Template, on page 283
Locator LED support for server hard-disks	Identifies the location a specific disk inserted in a blade or rack server using the local locator LED.	Local Disk Locator LED Status, on page 679
Reset Peer I/O Modules to Factory Defaults	Enables reboot of an I/O module that is unreachable through its peer I/O module.	Resetting an I/O Module, on page 734
vNIC template CDN Source	Enables selection of the Consistent Naming Device (CDN) Source as the vNIC Name, which in turn can either be customized or derived from the vNIC instance.	vNIC Template, on page 283
PCH SSD Controller Definition	Cisco UCS Manager Platform Controller Hub (PCH) Solid State Drive (SSD) Controller Definition provides a local storage configuration in storage profiles where you can configure all the disks in a single RAID or in a JBOD disk array.	PCH SSD Controller Definition, on page 633

Feature	Description	Where Documented
Host Firmware Package Enhancement	Allows exclusion of the firmware of specific components from a host firmware package either when creating a new host firmware package or when modifying an existing host firmware package.	This feature is documented in the following configuration guides: • *Cisco UCS B-Series GUI Firmware Management Configuration Guide* • *Cisco UCS B-Series CLI Firmware Management Configuration Guide* The firmware configuration guides can be found here: http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html
EFI Shell as a Boot Device	Enables creation of a boot policy with an EFI Shell as the boot device. Booting from an EFI Shell prevents loss of data and provides more options to script, debug, and control various booting scenarios. EFI Shell is supported as a boot device only in the Uefi boot mode.	Configuring an EFI Shell Boot for a Boot Policy, on page 562
Multicast Hardware Hash	Enables all links between the IOM and the fabric interconnect in a port channel to be used for multicast traffic when multicast hardware hashing is enabled.	Chassis/FEX Discovery Policy, on page 203

Table 6: New Features and Changed Behavior in Cisco UCS Release 2.2(8)

Feature	Description	Where Documented
Next Boot	The maintenance policy now provides an On Next Boot option. This option is used in combination with either User Ack or Timer Automatic.	Maintenance Policy, on page 568

Feature	Description	Where Documented
Graceful shutdown	When you acknowledge a server reboot using the graceful shut down options or a change in the service profile that requires the server reboot, Cisco UCS Manager waits until the time specified time in the maintenance policy before performing a hard shut down.	Maintenance Policy, on page 568
Health monitoring of end points	Provides enhanced health monitoring of end points for fabric interconnects, IOMs, FEXes, blade servers, and rack servers.	This feature is documented in the following configuration guides: • *Cisco UCS B-Series GUI System Monitoring Guide* • *Cisco UCS B-Series CLI System Monitoring Guide* The system monitoring guides can be found here: http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html
Power synchronization between servers and their associated service profiles	Provides a global (default) power sync policy to synchronize the power state between the associated service profiles and the servers when the desired power state of the service profile differs from the actual power state of the server.	Power Sync Policy, on page 668
Factory reset of blade servers	You can now reset a blade server to its factory settings. By default, the factory reset operation does not affect storage drives and flexflash drives. This is to prevent any loss of data.	Resetting a Blade Server to Factory Default Settings, on page 692
Support for 160 LDAP group maps	Cisco UCS Manager now supports a maximum of 160 LDAP group maps.	LDAP Group Mapping, on page 153

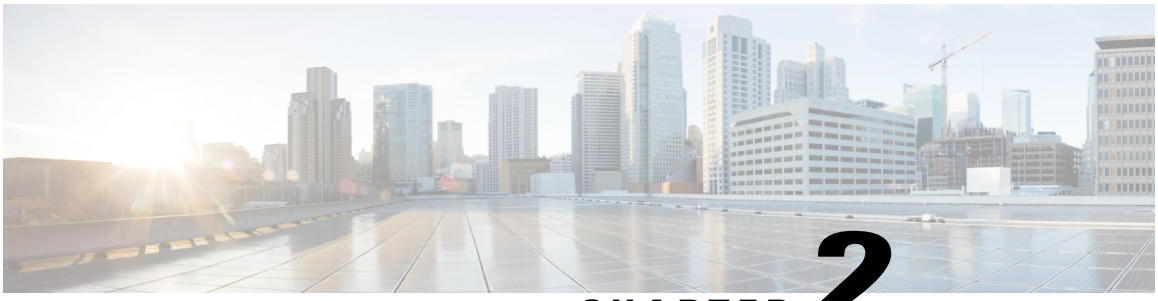

CHAPTER 2

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 pres**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 7: 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.

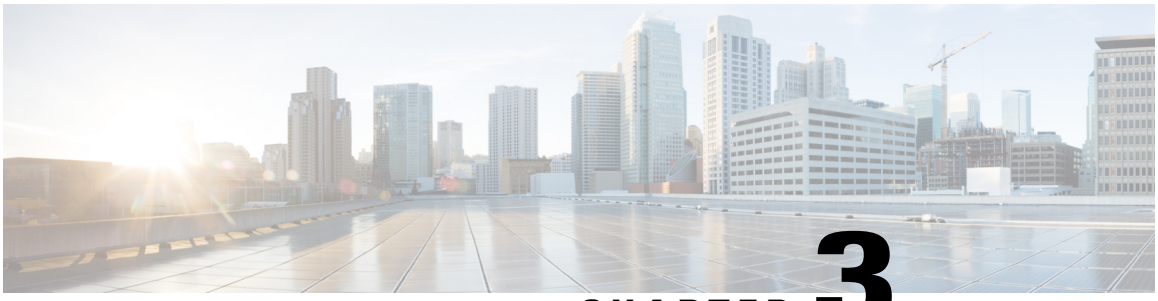

CHAPTER 3

Overview of Cisco UCS Manager

This chapter includes the following sections:

About Cisco UCS Manager

Cisco UCS Manager is the management system for all components in a Cisco UCS domain. Cisco UCS Manager runs within the fabric interconnect. You can use any of the interfaces available with this management service to access, configure, administer, and monitor the network and server resources for all chassis connected to the fabric interconnect.

Multiple Management Interfaces

Cisco UCS Manager includes the following interfaces you can use to manage a Cisco UCS domain:

- Cisco UCS Manager GUI
- Cisco UCS Manager CLI
- XML API
- KVM
- IPMI

Almost all tasks can be performed in any of the interfaces, and the results of tasks performed in one interface are automatically displayed in another.

However, you cannot do the following:

- Use Cisco UCS Manager GUI to invoke Cisco UCS Manager CLI.
- View the results of a command invoked through Cisco UCS Manager CLI in Cisco UCS Manager GUI.

- Generate CLI output from Cisco UCS Manager GUI.

Centralized Management

Cisco UCS Manager centralizes the management of resources and devices, rather than using multiple management points. This centralized management includes management of the following devices in a Cisco UCS domain:

- Fabric interconnects.
- Software switches for virtual servers.
- Power and environmental management for chassis and servers.
- Configuration and firmware updates for server network interfaces (Ethernet NICs and converged network adapters).
- Firmware and BIOS settings for servers.

Support for Virtual and Physical Servers

Cisco UCS Manager abstracts server state information—including server identity, I/O configuration, MAC addresses and World Wide Names, firmware revision, and network profiles—into a service profile. You can apply the service profile to any server resource in the system, providing the same flexibility and support to physical servers, virtual servers, and virtual machines connected to a virtual device provided by a VIC adapter.

Role-Based Administration and Multi-Tenancy Support

Cisco UCS Manager supports flexibly defined roles so that data centers can use the same best practices with which they manage discrete servers, storage, and networks to operate a Cisco UCS domain. You can create user roles with privileges that reflect user responsibilities in the data center. For example, you can create the following:

- Server administrator roles with control over server-related configurations.
- Storage administrator roles with control over tasks related to the SAN.
- Network administrator roles with control over tasks related to the LAN.

Cisco UCS is multi-tenancy ready, exposing primitives that allow systems management software using the API to get controlled access to Cisco UCS resources. In a multi-tenancy environment, Cisco UCS Manager enables you to create locales for user roles that can limit the scope of a user to a particular organization.

Tasks You Can Perform in Cisco UCS Manager

You can use Cisco UCS Manager to perform management tasks for all physical and virtual devices within a Cisco UCS domain.

Cisco UCS Hardware Management

You can use Cisco UCS Manager to manage all hardware within a Cisco UCS domain, including the following:

- Chassis
- Servers

- Fabric interconnects
- Fans
- Ports
- Interface cards
- I/O modules

Cisco UCS Resource Management

You can use Cisco UCS Manager to create and manage all resources within a Cisco UCS domain, including the following:

- Servers
- WWN addresses
- MAC addresses
- UUIDs
- Bandwidth

Server Administration

A server administrator can use Cisco UCS Manager to perform server management tasks within a Cisco UCS domain, including the following:

- Create server pools and policies related to those pools, such as qualification policies
- Create policies for the servers, such as discovery policies, scrub policies, and IPMI policies
- Create service profiles and, if desired, service profile templates
- Apply service profiles to servers
- Monitor faults, alarms, and the status of equipment

Network Administration

A network administrator can use Cisco UCS Manager to perform tasks required to create LAN configuration for a Cisco UCS domain, including the following:

- Configure uplink ports, port channels, and LAN PIN groups
- Create VLANs
- Configure the quality of service classes and definitions
- Create the pools and policies related to network configuration, such as MAC address pools and Ethernet adapter profiles

Storage Administration

A storage administrator can use Cisco UCS Manager to perform tasks required to create SAN configuration for a Cisco UCS domain, including the following:

- Configure ports, port channels, and SAN PIN groups

- Create VSANs
- Configure the quality of service classes and definitions
- Create the pools and policies related to the network configuration, such as WWN pools and Fibre Channel adapter profiles

Tasks You Cannot Perform in Cisco UCS Manager

You cannot use Cisco UCS Manager to perform certain system management tasks that are not specifically related to device management within a Cisco UCS domain.

No Cross-System Management

You cannot use Cisco UCS Manager to manage systems or devices that are outside the Cisco UCS domain where Cisco UCS Manager is located. For example, you cannot manage heterogeneous environments, such as non-Cisco UCS x86 systems, SPARC systems, or PowerPC systems.

No Operating System or Application Provisioning or Management

Cisco UCS Manager provisions servers and, as a result, exists below the operating system on a server. Therefore, you cannot use it to provision or manage operating systems or applications on servers. For example, you cannot do the following:

- Deploy an OS, such as Windows or Linux
- Deploy patches for software, such as an OS or an application
- Install base software components, such as anti-virus software, monitoring agents, or backup clients
- Install software applications, such as databases, application server software, or web servers
- Perform operator actions, including restarting an Oracle database, restarting printer queues, or handling non-Cisco UCS user accounts
- Configure or manage external storage on the SAN or NAS storage

Cisco UCS Manager in a High Availability Environment

In a high availability environment with two fabric interconnects, you can run a separate instance of Cisco UCS Manager on each fabric interconnect. The Cisco UCS Manager on the primary fabric interconnect acts as the primary management instance, and the Cisco UCS Manager on the other fabric interconnect is the subordinate management instance.

The two instances of Cisco UCS Manager communicate across a private network between the L1 and L2 Ethernet ports on the fabric interconnects. Configuration and status information is communicated across this private network to ensure that all management information is replicated. This ongoing communication ensures that the management information for Cisco UCS persists even if the primary fabric interconnect fails. In addition, the "floating" management IP address that runs on the primary Cisco UCS Manager ensures a smooth transition in the event of a failover to the subordinate fabric interconnect.

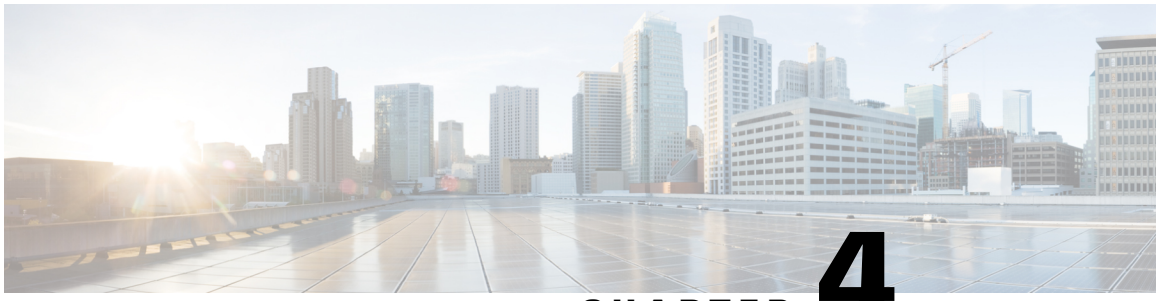

CHAPTER 4

Overview of Cisco UCS Manager GUI

This chapter includes the following sections:

Overview of Cisco UCS Manager GUI

Cisco UCS Manager GUI is the Java application that provides a GUI interface to Cisco UCS Manager. You can start and access Cisco UCS Manager GUI from any computer that meets the requirements listed in the System Requirements section of the Cisco UCS Software Release Notes.

Each time you start Cisco UCS Manager GUI, Cisco UCS Manager uses Java Web Start technology to cache the current version of the application on your computer. As a result, you do not have to download the application every time you log in. You only have to download the application the first time that you log in from a computer after the Cisco UCS Manager software has been updated on a system.

Tip The title bar displays the name of the Cisco UCS domain to which you are connected.

Fault Summary Area

The **Fault Summary** area displays in the upper left of Cisco UCS Manager GUI. This area displays a summary of all faults that have occurred in the Cisco UCS domain.

Each type of fault is represented by a different icon. The number below each icon indicates how many faults of that type have occurred in the system. If you click an icon, Cisco UCS Manager GUI opens the **Faults** tab in the **Work** area and displays the details of all faults of that type.

The following table describes the types of faults each icon in the **Fault Summary** area represents:

Fault Type	Description
Critical Alarms	Critical problems exist with one or more components. Ensure that you fix the issues before moving on to other tasks.
Major Alarms	Serious problems exist with one or more components. Ensure that you fix the issues before moving on to other tasks.
Minor Alarms	Problems exist with one or more components that might adversely affect the system performance. These issues should be researched and fixed as soon as possible before they become major or critical issues.
Warning Alarms	Potential problems exist with one or more components that might adversely affect the system performance if the problem persists. Ensure that you fix the issues before moving on to other tasks.
Suppression Status	The state of fault suppression tasks on the component. Click **Suppression Task Properties** in the **Actions** area to view all fault suppression tasks. **Note** This field displays only if the component supports fault suppression.

Tip If you only want to see faults for a specific object, navigate to that object and then review the **Faults** tab for that object.

Navigation Pane

The **Navigation** pane displays on the left side of Cisco UCS Manager GUI below the **Fault Summary** area. This pane provides centralized navigation to all equipment and other components in the Cisco UCS domain. When you select a component in the **Navigation** pane, the object displays in the **Work** area.

The **Navigation** pane has five tabs. Each tab includes the following elements:

- A **Filter** combo box that you can use to filter the navigation tree to view all nodes or only one node.
- An expandable navigation tree that you can use to access all components on that tab. An icon next to an folder indicates that the node or folder has subcomponents.

Equipment Tab

This tab contains a basic inventory of the equipment in the Cisco UCS domain. A system or server administrator can use this tab to access and manage the chassis, fabric interconnects, servers, and other hardware. A red, orange, or yellow rectangle around a device name indicate that the device has a fault.

The major nodes in this tab are the following:

- **Equipment—**An overview of the entire Cisco UCS domain, including active and decommissioned hardware, firmware management, equipment-related policies, power groups, and an aggregated list of faults.
- **Chassis—**The fans, I/O modules, power supply units (PSUs), and Cisco UCS B-Series blade servers for each chassis in the Cisco UCS domain.
- **Rack-Mounts—**The FEXes and Cisco UCS C-Series rack servers integrated with the Cisco UCS domain.
- **Fabric Interconnects—**The fixed and expansion modules, fans, and PSUs associated with the fabric interconnects in the Cisco UCS domain.

Servers Tab

This tab contains the server-related components, such as service profiles, polices, and pools. A server administrator typically accesses and manages the components on this tab.

The major nodes below the **Servers** node in this tab are the following:

- **Servers—**Service profiles and the relationship between the defined organizations and the service profiles.
- **Service Profiles—**The service profiles defined in the system divided by organization.
- **Service Profile Templates—**The service profile templates defined in the system divided by organization.
- **Policies—**Server-related policies for adapters, BIOS, firmware, IPMI access, local disk configuration, maintenance, power, disk scrubbing, Serial over LAN, server pools, iSCSI authentication, vNIC/vHBA placement, and fault thresholds.
- **Pools—**Server pools and UUID suffix pools.
- **Schedules—**Maintenance and fault suppression schedules.

LAN Tab

This tab contains the components related to LAN configuration, such as LAN pin groups, quality of service classes, VLANs, policies, pools, and the internal domain. A network administrator typically accesses and manages the components on this tab.

The major nodes below the **LAN** node in this tab are the following:

- **LAN Cloud—**Quality of service settings, port channels, pin groups, VLANs, VLAN optimization sets, threshold policies.
- **Appliances—**Interfaces, port channels, and VLANs.

- **Internal LAN—**Ports and threshold polices associated with the internal fabric.
- **Policies—**Policies governing flow control, adapters, vNICs, vNIC templates, quality of services, and fault thresholds.
- **Pools—**The IP pools and MAC pools defined in the system.
- **Traffic Monitoring Sessions—**The port traffic monitoring sessions defined in the system.
- **Netflow Monitoring—**Flow record definitions, flow exporters, flow monitors, and flow monitor sessions.

SAN Tab

This tab contains the components related to SAN configuration, such as pin groups, VSANs, policies, and pools. A storage administrator typically accesses and manages the components on this tab.

The major nodes in this tab are the following:

- **SAN—**SAN uplinks, fibre channel address assignment, SAN-related pools, and VSANs.
- **SAN Cloud—**SAN uplinks, fibre channel address assignment, SAN-related pools, and VSANs.
- **Storage Cloud—**Storage ports and VSANs.
- **Policies—**Fibre Channel adapter policies, default vHBA behavior, SAN connectivity policies, storage connection policies, vHBA templates, and fault thresholds.
- **Pools—**The iSCSI Qualified Name (IQN) pools and World Wide Name pools defined in the system.
- **Traffic Monitoring Sessions—**The port traffic monitoring sessions defined in the system.

VM Tab

This tab contains the components required to configure VM-FEX for servers with a VIC adapter. For example, you use components on this tab to configure the connection between Cisco UCS Manager and VMware vCenter, to configure distributed virtual switches, port profiles, and to view the virtual machines hosted on servers in the Cisco UCS domain.

The major nodes in this tab are the following:

- **All–**Port profiles, virtual machines, virtual switches, certificates, the lifecycle policy, VM-related events and FSM tasks.
- **Clusters—**Clusters, including the associated virtual machines and port profiles.
- **Fabric Network Sets—** Fabric Network Definitions
- **Port Profiles—** Port Profiles
- **VM Networks—** Virtual Networks
- **Microsoft—** Microsoft Networking Components
- **VMware—**vCenters, including folders, Datacenters, virtual machines, and virtual switches

Admin Tab

This tab contains system-wide settings, such as user manager and communication services, and troubleshooting components, such as faults and events. The system administrator typically accesses and manages the components on this tab.

The major nodes in this tab are the following:

- **All—**Management interfaces, backup and import configuration, tech support file creation, the full state backup policy, and the all configuration export policy.
- **Faults, Events and Audit Log—**System-wide faults, events, audit logs, syslog entries, core files, tech support files, and global fault policies.
- **User Management—**Authentication methods, remote access methods, local users, locales, and user roles.
- **Key Management—**SSH key and trusted point settings.
- **Communication Management—**Communication service settings for SSH, Telnet, HTTP, HTTPS, SNMP, web session limits, Call Home settings, DNS management, and management interfaces, and Cisco UCS Central settings.
- **Stats Management—**Threshold statistics settings that control when faults are issued by the system.
- **Time Zone Management—**NTP server settings to establish time zone synchronization.
- **Capability Catalog—**The capability catalog, a set of tunable parameters, strings, and rules.
- **Management Extension—**Management extensions, which allow you add support for previously unsupported servers and other hardware to Cisco UCS Manager.
- **License Management—**The feature and port licenses installed on this system.

Toolbar

The toolbar displays on the right side of Cisco UCS Manager GUI above the **Work** pane. You can use the menu buttons in the toolbar to perform common actions, including the following actions:

- Navigate between previously viewed items in the **Work** pane
- Create elements for the Cisco UCS domain
- Set options for Cisco UCS Manager GUI
- Access online help for Cisco UCS Manager GUI

Work Pane

The **Work** pane displays on the right side of Cisco UCS Manager GUI. This pane displays details about the component selected in the **Navigation** pane.

The **Work** pane includes the following elements:

- A navigation bar that displays the path from the main node of the tab in the **Navigation** pane to the selected element. You can click any component in this path to display that component in the **Work** pane.

- A content area that displays tabs with information related to the component selected in the **Navigation** pane. The tabs displayed in the content area depends upon the selected component. You can use these tabs to view information about the component, create components, modify properties of the component, and examine a selected object.

Status Bar

The status bar displays across the bottom of Cisco UCS Manager GUI. The status bar provides information about the state of the application.

On the left, the status bar displays the following information about your current session in Cisco UCS Manager GUI:

- A lock icon that indicates the protocol you used to log in. If the icon is locked, you connected with HTTPS and if the icon is unlocked, you connected with HTTP.
- The username you used to log in.
- The IP address of the server where you logged in.

If your Cisco UCS domain is registered with Cisco UCS Central, the status bar also displays the following information:

- A lock icon that indicates a secure connection.
- The IP address of the server for Cisco UCS Central.

On the right, the status bar displays the system time.

Table Customization

Cisco UCS Manager GUI enables you to customize the tables on each tab. You can change the type of content that you view and filter the content.

Table Customization Menu Button

This menu button in the upper right of every table enables you to control and customize your view of the table. The drop-down menu for this button includes the following options:

Menu Item	Description
Column Name	The menu contains an entry for each column in the table. Click a column name to display or hide the column.
Horizontal Scroll	If selected, adds a horizontal scroll bar to the table. If not selected, when you widen one of the columns, all columns to the right narrow and do not scroll.
Pack All Columns	Resizes all columns to their default width.
Pack Selected Column	Resizes only the selected column to its default width.

Table Content Filtering

The **Filter** button above each table enables you to filter the content in the table according to the criteria that you set in the **Filter** dialog box. The dialog box includes the following filtering options:

Name	Description
Disable option	No filtering criteria is used on the content of the column. This is the default setting.
Equal option	Displays only that content in the column which exactly matches the value specified.
Not Equal option	Displays only that content in the column which does not exactly match the value specified.
Wildcard option	The criteria you enter can include one of the following wildcards: • _ (underscore) or ? (question mark)—replaces a single character • % (percent sign) or * (asterisk)—replaces any sequence of characters
Less Than option	Displays only that content in the column which is less than the value specified.
Less Than Or Equal option	Displays only that content in the column which is less than or equal to the value specified.
Greater Than option	Displays only that content in the column which is greater than the value specified.
Greater Than Or Equal option	Displays only that content in the column which is greater than or equal to the value specified.

LAN Uplinks Manager

The LAN Uplinks Manager provides a single interface where you can configure the connections between Cisco UCS and the LAN. You can use the LAN Uplinks Manager to create and configure the following:

- Ethernet switching mode
- Uplink Ethernet ports
- Port channels
- LAN pin groups
- Named VLANs

- Server ports
- QoS system classes

Some of the configuration that you can do in the LAN Uplinks Manager can also be done in nodes on other tabs, such as the **Equipment** tab or the **LAN** tab.

Internal Fabric Manager

The Internal Fabric Manager provides a single interface where you can configure server ports for a fabric interconnect in a Cisco UCS domain. The Internal Fabric Manager is accessible from the **General** tab for that fabric interconnect.

Some of the configuration that you can do in the Internal Fabric Manager can also be done in nodes on the **Equipment** tab, on the **LAN** tab, or in the LAN Uplinks Manager.

Hybrid Display

For each chassis in a Cisco UCS domain, Cisco UCS Manager GUI provides a hybrid display that includes both physical components and connections between the chassis and the fabric interconnects.

This tab displays detailed information about the connections between the selected chassis or rack server and the fabric interconnects. It has an icon for the following:

- Each fabric interconnect in the system
- The I/O module (IOM) in the selected component, which is shown as an independent unit to make the connection paths easier to see
- The selected chassis showing the servers and PSUs, or the selected rack server.

The lines between the icons represent the connections between the following:

- DCE interface on each server and the associated server port on the I/O module. These connections are created by Cisco and cannot be changed.
- Server port on the I/O module and the associated port on the fabric interconnect. You can change these connections if desired.

You can mouse over the icons and lines to view tooltips identifying each component or connection, and you can double-click any component to view properties for that component.

If there is a fault associated with the component or any of its subcomponents, Cisco UCS Manager GUI displays a fault icon on top of the appropriate component. If there are multiple fault messages, Cisco UCS Manager GUI displays the icon associated with the most serious fault message in the system.

Logging in to the Cisco UCS Manager GUI through HTTPS

The default HTTP web link for the Cisco UCS Manager GUI is `http://UCSManager_IPv4`, or `http://UCSManager_IPv6`, where *UCSManager_IPv4* or *UCSManager_IPv6* represents the IPv4 or IPv6 address, respectively, assigned to Cisco UCS Manager. This IP address can be one of the following:

- Cluster configuration: *UCSManager_IPv4* or *UCSManager_IPv6* represents the virtual or cluster IPv4 address or IPv6 address, respectively, assigned to Cisco UCS Manager. Do not use the IP addresses assigned to the management port on the fabric interconnects.
- Standalone configuration: *UCSManager_IPv4* or *UCSManager_IPv6* represents the IPv4 or IPv6 address, respectively, of the management port on the fabric interconnect

Note Some browsers do not support HTTPS access using an IPv6 address.

Procedure

Step 1 In your web browser, type the Cisco UCS Manager GUI web link or select the bookmark in your browser.

Step 2 If a **Security Alert** dialog box appears, click **Yes** to accept the security certificate and continue.

Step 3 In the Cisco UCS Manager launch page, click **Launch UCS Manager**.
Based on the web browser you use to log in, a prompt might display to download or save the .JNLP file.

Step 4 If Cisco UCS Manager displays a pre-login banner, review the message and click **OK** to close the dialog box.

Step 5 If a **Security** dialog box displays, do the following:

a) (Optional) Check the check box to accept all content from Cisco.
b) Click **Yes** to accept the certificate and continue.

Step 6 In the **Login** dialog box, do the following:

a) Enter your username and password.
b) If your Cisco UCS implementation includes multiple domains, select the appropriate domain from the **Domain** drop-down list.
c) Click **Login**.

Logging in to the Cisco UCS Manager GUI through HTTP

The default HTTP web link for the Cisco UCS Manager GUI is `http://`*UCSManager_IPv4*, or `http://`*UCSManager_IPv6*, where *UCSManager_IPv4* or *UCSManager_IPv6* represents the IPv4 or IPv6 address, respectively, assigned to Cisco UCS Manager. This IP address can be one of the following:

- Cluster configuration: *UCSManager_IPv4* or *UCSManager_IPv6* represents the virtual or cluster IPv4 address or IPv6 address, respectively, assigned to Cisco UCS Manager. Do not use the IP addresses assigned to the management port on the fabric interconnects.
- Standalone configuration: *UCSManager_IPv4* or *UCSManager_IPv6* represents the IPv4 or IPv6 address, respectively, of the management port on the fabric interconnect

Procedure

Step 1 In your web browser, type the Cisco UCS Manager GUI web link or select the bookmark in your browser.

Step 2 If Cisco UCS Manager displays a pre-login banner, review the message and click **OK** to close the dialog box.

Step 3 In the Cisco UCS Manager launch page, click **Launch UCS Manager**.
Based on the web browser you use to log in, a prompt might display to download or save the .JNLP file.

Step 4 In the **Login** dialog box, do the following:

a) Enter your username and password.

b) If your Cisco UCS implementation includes multiple domains, select the appropriate domain from the **Domain** drop-down list.

c) Click **Login**.

Logging Out of the Cisco UCS Manager GUI

Procedure

Step 1 In the Cisco UCS Manager GUI, click **Exit** in the upper right.
The Cisco UCS Manager GUI blurs on your screen to indicate that you cannot use it and displays the **Exit** dialog box.

Step 2 From the drop-down list, select one of the following:

- **Exit** to log out and shut down the Cisco UCS Manager GUI.
- **Log Off** to log out of the Cisco UCS Manager GUI and log in a different user.

Step 3 Click **OK**.

Web Session Limits

Web session limits are used by Cisco UCS Manager to restrict the number of web sessions (both GUI and XML) permitted access to the system at any one time.

By default, the number of concurrent web sessions allowed by Cisco UCS Manager is set to the maximum value: 256.

Setting the Web Session Limit for Cisco UCS Manager

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **Communication Management** > **Communication Services**.
Step 3 In the **Work** pane, click the **Communication Services** tab.
Step 4 In the **Web Session Limits** area, complete the required fields.
Step 5 Click **Save Changes**.

Pre-Login Banner

With a pre-login banner, when a user logs into Cisco UCS Manager GUI, Cisco UCS Manager displays the banner text in the **Create Pre-Login Banner** dialog box and waits until the user dismisses that dialog box before it prompts for the username and password. When a user logs into Cisco UCS Manager CLI, Cisco UCS Manager displays the banner text in a dialog box and waits for the user to dismiss that dialog box before it prompts for the password. It then repeats the banner text above the copyright block that it displays to the user.

Creating the Pre-Login Banner

If the **Pre-Login Banner** area does not appear on the **Banners** tab, Cisco UCS Manager does not display a pre-login banner when users log in. If the **Pre-Login Banner** area does appear, you cannot create a second pre-login banner. You can only delete or modify the existing banner.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **User Management**.
Step 3 Click the **User Services** node.
Step 4 In the **Work** pane, click the **Banners** tab.
Step 5 In the **Actions** area, click **Create Pre-Login Banner**.
Step 6 In the **Create Pre-Login Banner** dialog box, click in the text field and enter the message that you want users to see when they log in to Cisco UCS Manager.
You can enter any standard ASCII character in this field.
Step 7 Click **OK**.

Modifying the Pre-Login Banner

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management**.

Step 3 Click the **User Services** node.

Step 4 In the **Work** pane, click the **Banners** tab.

Step 5 Click in the text field in the **Pre-Login Banner** area and make the necessary changes to the text. You can enter any standard ASCII character in this field.

Step 6 Click **Save Changes**.

Deleting the Pre-Login Banner

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management**.

Step 3 Click the **User Services** node.

Step 4 In the **Work** pane, click the **Banners** tab.

Step 5 In the **Actions** area, click **Delete**.

Step 6 If a confirmation dialog box displays, click **Yes**.

Cisco UCS Manager GUI Properties

Configuring the Cisco UCS Manager GUI Session and Log Properties

These properties determine how Cisco UCS Manager GUI reacts to session interruptions and inactivity, and configures the Cisco UCS Manager GUI Java message logging.

Procedure

Step 1 In the toolbar, click **Options** to open the **Properties** dialog box.
Step 2 In the right pane, click **Session**.
Step 3 In the **Session** tab, update the fields as needed.
Step 4 Click **OK**.

Configuring Properties for Confirmation Messages

These properties determine whether or not Cisco UCS Manager GUI displays a confirmation message after configuration changes and other operations.

Procedure

Step 1 In the toolbar, click **Options** to open the **Properties** dialog box.
Step 2 In the right pane, click **Confirmation Messages**.
Step 3 In the **Confirmation Messages** tab, updated the fields as needed.
Step 4 Click **OK**.

Configuring Properties for External Applications

Cisco UCS Manager GUI uses these properties to connect with external applications, such as SSH.

Procedure

Step 1 In the toolbar, click **Options** to open the **Properties** dialog box.
Step 2 In the right pane, click **External Applications**.
Step 3 In the **External Applications** tab, specify values for the **SSH** and **SSH Parameters** fields.
Step 4 Click **OK**.

Customizing the Appearance of Cisco UCS Manager GUI

These properties allow you to customize the some of the visual properties of Cisco UCS Manager GUI.

Procedure

Step 1 In the toolbar, click **Options** to open the **Properties** dialog box.

Step 2 In the right pane, click **Visual Enhancements**.

Step 3 In the **Visual Enhancements** tab, update the fields as needed.

Step 4 Click **OK**.

Determining the Acceptable Range of Values for a Field

Some properties have a restricted range of values that you can enter. You can use this procedure to determine that acceptable range for fields in a dialog box, window, or tab. You cannot use this procedure to determine the acceptable range of values for properties listed in a table or tree.

Procedure

Step 1 Place your cursor in the field for which you want to check the range to give focus to that field.

Step 2 Press Alt + Shift + **R**.
Cisco UCS Manager GUI displays the acceptable range of values for a few seconds. The range disappears if you click anywhere on the screen.

Determining Where a Policy Is Used

You can use this procedure to determine which service profiles and service profile templates are associated with the selected policy.

Procedure

Step 1 In the **Navigation** pane, click the policy whose usage you want to view.

Step 2 In the **Work** pane, click the **General** tab.

Step 3 In the **Actions** area, click **Show Policy Usage**.
Cisco UCS Manager GUI displays the **Service Profiles/Templates** dialog box that shows the associated service profiles and service profile templates.

Determining Where a Pool Is Used

You can use this procedure to determine which service profiles and service profile templates are associated with the selected pool.

Procedure

Step 1 In the **Navigation** pane, click the pool whose usage you want to view.

Step 2 In the **Work** pane, click the **General** tab.

Step 3 In the **Actions** area, click **Show Pool Usage**.
Cisco UCS Manager GUI displays the **Service Profiles/Templates** dialog box that shows the associated service profiles and service profile templates.

Deleting a Pool, Policy, or Other Object

The method you use to delete a pool, policy, or other object, such as a VLAN, is the same for all objects.

Note Before you delete an object, ensure that it is not being used or referenced by another object in the system. For example, before you delete a network policy, ensure that a service profile does not reference that policy.

Procedure

Step 1 In the **Navigation** pane, expand the pod to view the objects.

Step 2 In the **Work** pane, click the appropriate tab and navigate to where the object is located.
For example, if you want to delete a VLAN, click the **VLANs** tab.

Step 3 Choose the object that you want to delete and click **Delete**.

Step 4 Click **Delete**.

Copying the XML

To assist you in developing scripts or creating applications with the XML API for Cisco UCS, Cisco UCS Manager GUI includes an option to copy the XML used to create an object in Cisco UCS Manager. This option is available on the right-click menu for most object nodes in the **Navigation** pane, such as the **Port Profiles** node or the node for a specific service profile.

Procedure

Step 1 In the **Navigation** pane, navigate to the object for which you want to copy the XML.

Step 2 Right-click on that object and choose **Copy XML**.

Step 3 Paste the XML into an XML editor, Notepad, or another application.

HTML5 GUI for Cisco UCS Manager

Overview of Cisco UCS Manager HTML5 GUI

Beginning with Release 3.0(2), you can access Cisco UCS Manager using a browser-based GUI interface. This allows you to run Cisco UCS Manager from any supported HTML5 web browser, without the need for Java. The HTML5 GUI has a similar look and feel to the Java-based GUI. See the following sections for descriptions:

- Fault Summary Area, on page 42
- Navigation Pane, on page 42
- Toolbar, on page 45
- Work Pane, on page 45
- Status Bar, on page 46
- Table Customization, on page 46
- LAN Uplinks Manager, on page 47
- Internal Fabric Manager, on page 48
- Hybrid Display, on page 48

The tasks for Web Session Limits, on page 50 and Pre-Login Banner, on page 51 are performed the same. However, Configuring Properties for External Applications, on page 53 is not supported.

Note Java is still required to launch the KVM client.

Logging in to the Cisco UCS Manager HTML5 GUI through HTTPS

The default HTTP web link for the Cisco UCS Manager GUI is `http://`*`UCSManager_IPv4`*, or `http://`*`UCSManager_IPv6`*, where *UCSManager_IPv4* or *UCSManager_IPv6* represents the IPv4 or IPv6 address, respectively, assigned to Cisco UCS Manager. This IP address can be one of the following:

- Cluster configuration: *UCSManager_IPv4* or *UCSManager_IPv6* represents the virtual or cluster IPv4 address or IPv6 address, respectively, assigned to Cisco UCS Manager. Do not use the IP addresses assigned to the management port on the fabric interconnects.
- Standalone configuration: *UCSManager_IPv4* or *UCSManager_IPv6* represents the IPv4 or IPv6 address, respectively, of the management port on the fabric interconnect

Note Some browsers do not support HTTPS access using an IPv6 address.

Procedure

Step 1 In your web browser, type the Cisco UCS Manager GUI web link or select the bookmark in your browser.

Step 2 If a **Security Alert** dialog box appears, click **Yes** to accept the security certificate and continue.

Step 3 In the HTML area of the Cisco UCS Manager launch page, click **Launch UCS Manager**.

Step 4 If Cisco UCS Manager displays a pre-login banner, review the message and click **OK** to close the dialog box.

Step 5 If a **Security** dialog box displays, do the following:

a) (Optional) Check the check box to accept all content from Cisco.
b) Click **Yes** to accept the certificate and continue.

Step 6 In the **Login** dialog box, do the following:

a) Enter your username and password.
b) If your Cisco UCS implementation includes multiple domains, select the appropriate domain from the **Domain** drop-down list.
c) Click **Login**.

Logging in to the Cisco UCS Manager HTML5 GUI through HTTP

The default HTTP web link for the Cisco UCS Manager GUI is `http://UCSManager_IPv4`, or `http://UCSManager_IPv6`, where *UCSManager_IPv4* or *UCSManager_IPv6* represents the IPv4 or IPv6 address, respectively, assigned to Cisco UCS Manager. This IP address can be one of the following:

- Cluster configuration: *UCSManager_IPv4* or *UCSManager_IPv6* represents the virtual or cluster IPv4 address or IPv6 address, respectively, assigned to Cisco UCS Manager. Do not use the IP addresses assigned to the management port on the fabric interconnects.
- Standalone configuration: *UCSManager_IPv4* or *UCSManager_IPv6* represents the IPv4 or IPv6 address, respectively, of the management port on the fabric interconnect

Procedure

Step 1 In your web browser, type the Cisco UCS Manager GUI web link or select the bookmark in your browser.

Step 2 If Cisco UCS Manager displays a pre-login banner, review the message and click **OK** to close the dialog box.

Step 3 In the HTML area of the Cisco UCS Manager launch page, click **Launch UCS Manager**.

Step 4 In the **Login** dialog box, do the following:

a) Enter your username and password.

b) If your Cisco UCS implementation includes multiple domains, select the appropriate domain from the **Domain** drop-down list.

c) Click **Login**.

Logging Out of the Cisco UCS Manager HTML5 GUI

Procedure

Step 1 In the Cisco UCS Manager GUI, click **Exit** in the upper right.

Step 2 From the drop-down list, select one of the following:

- **Yes** to log out.

  After logging out, the Cisco UCS Manager launch page displays.

- **No** to continue to use the Cisco UCS Manager GUI.

Behavior Changes in the HTML5 GUI

The Cisco UCS Manager HTML5 GUI is designed to match the look and feel of the existing java-based application. Some of the changes include:

- The **Show Navigator** command is not available on the **Navigation** pane tree view. However, all of the information is displayed in the **Work** pane.
- The HTML5 GUI is a single page web application, and multiple windows are not available. For example, if you are working on a web page that launches a child window, you must complete the tasks in the child window, and then close the child window before the original window is available. This behavior affects the following pages:
  - LAN Uplinks Manager
  - SAN Uplinks Manager
  - NAS Appliance Manager

- SAN Storage Manager
- Internal Fabric Manager

- Clicking the online help **?** button immediately launches the online help for the currently active tab.
- In most cases, table row operations are applied immediately. However, creating vNICs, vHBAs, and iSCSI vNICs requires that you click the **Save Changes** button.
- The default value for confirmation boxes is **No** or **Cancel**.
- All file downloads are determined by your local browser settings.

HTML5 Supported Browsers

The following web browsers are supported for the Cisco UCS Manager HTML5 GUI:

- Microsoft Internet Explorer version 10 or higher
- Mozilla Firefox version 26 or higher
- Google Chrome version 30 or higher
- Apple Safari version 7 or higher

Note You may need to accept the certificate the first time you log in from a Safari browser.

- Opera version 19 or higher

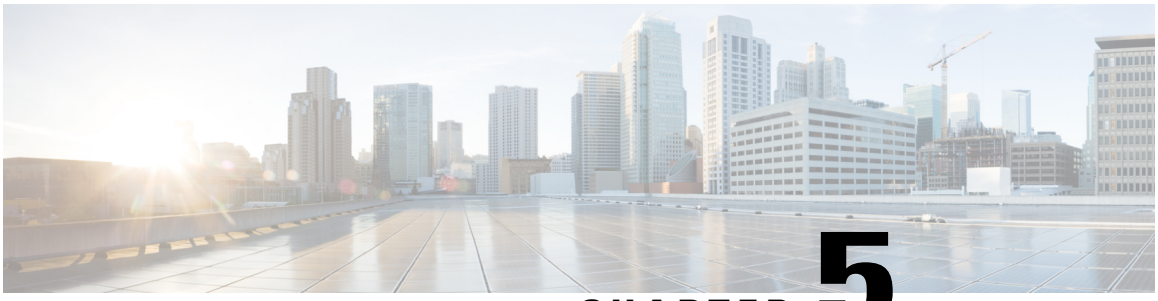

CHAPTER 5

Configuring the Fabric Interconnects

This chapter includes the following sections:

Initial System Setup

The first time that you access a fabric interconnect in a Cisco UCS domain, a setup wizard prompts you for the following information required to configure the system:

- Installation method (GUI or CLI)
- Setup mode (restore from full system backup or initial setup)
- System configuration type (standalone or cluster configuration)
- System name
- Admin password
- Management port IPv4 address and subnet mask, or IPv6 address and prefix

- Default gateway IPv4 or IPv6 address
- DNS Server IPv4 or IPv6 address
- Default domain name

Setup Mode

You can choose to either restore the system configuration from an existing backup file, or manually set up the system by going through the Setup wizard. If you choose to restore the system, the backup file must be reachable from the management network.

System Configuration Type

You can configure a Cisco UCS domain to use a single fabric interconnect in a standalone configuration or to use a redundant pair of fabric interconnects in a cluster configuration.

A cluster configuration provides high availability. If one fabric interconnect becomes unavailable, the other takes over. Only one management port (Mgmt0) connection is required to support a cluster configuration; however, both Mgmt0 ports should be connected to provide link-level redundancy.

In addition, a cluster configuration actively enhances failover recovery time for redundant virtual interface (VIF) connections. When an adapter has an active VIF connection to one fabric interconnect and a standby VIF connection to the second, the learned MAC addresses of the active VIF are replicated but not installed on the second fabric interconnect. If the active VIF fails, the second fabric interconnect installs the replicated MAC addresses and broadcasts them to the network through gratuitous ARP messages, shortening the switchover time.

Note The cluster configuration provides redundancy only for the management plane. Data redundancy is dependent on the user configuration and might require a third-party tool to support data redundancy.

To use the cluster configuration, you must directly connect the two fabric interconnects together using Ethernet cables between the L1 (L1-to-L1) and L2 (L2-to-L2) high-availability ports, with no other fabric interconnects in between. Also you can connect the fabric interconnects directly through a patch panel to allow the two fabric interconnects to continuously monitor the status of each other and quickly know when one has failed.

Both fabric interconnects in a cluster configuration must go through the initial setup process. You must enable the first fabric interconnect that you set up for a cluster configuration. When you set up the second fabric interconnect, it detects the first fabric interconnect as a peer fabric interconnect in the cluster.

For more information, see to the *Cisco UCS 6100 Series Fabric Interconnect Hardware Installation Guide*.

Management Port IP Address

In a standalone configuration, you must specify only one IPv4 address, gateway, and subnet mask, or only one IPv6 address, gateway, and network prefix for the single management port on the fabric interconnect. You can configure either an IPv4 or an IPv6 address for the management port IP address.

In a cluster configuration, you must specify the following three IPv4 addresses in the same subnet, or three IPv6 addresses with the same prefix:

- Management port IP address for fabric interconnect A
- Management port IP address for fabric interconnect B
- Cluster IP address

Note In a cluster configuration, the management port for both fabric interconnects must be configured with the same address type, either IPv4 or IPv6. If you configure the first FI with an IPv4 address then attempt to configure the second FI with an IPv6 address, the configuration will fail.

Performing an Initial System Setup for a Standalone Configuration

Before You Begin

1. Verify the following physical connections on the fabric interconnect:
   - The console port is physically connected to a computer terminal or console server
   - The management Ethernet port (mgmt0) is connected to an external hub, switch, or router

   For more information, refer to the *Cisco UCS Hardware Installation Guide* for your fabric interconnect.
2. Verify that the console port parameters on the computer terminal (or console server) attached to the console port are as follows:
   - 9600 baud
   - 8 data bits
   - No parity
   - 1 stop bit
3. Collect the following information that you will need to supply during the initial setup:
   - System name
   - Password for the admin account. Choose a strong password that meets the guidelines for Cisco UCS Manager passwords. This password cannot be blank.
   - Management port IPv4 and subnet mask, or IPv6 address and prefix.
   - Default gateway IPv4 or IPv6 address.
   - DNS server IPv4 or IPv6 address (optional).
   - Domain name for the system (optional).

Procedure

Step 1 Power on the fabric interconnect.
You will see the power on self-test messages as the fabric interconnect boots. The system will run a DHCP client to check for a lease.

Step 2 If the system obtains a lease go to step 6, otherwise, continue to the next step.

Step 3 Connect to the console port.

Step 4 At the installation method prompt, enter gui.

Step 5 If the system cannot access a DHCP server, you are prompted to enter the following information:

- IPv4 or IPv6 address for the management port on the fabric interconnect
- IPv4 subnet mask or IPv6 prefix for the management port on the fabric interconnect
- IPv4 or IPv6 address for the default gateway assigned to the fabric interconnect

Note Cisco UCS Manager does not support auto configuration from IPv6 DHCP servers, or IPv6 router advertisements.

Step 6 Copy the web link from the prompt into a supported web browser and go to the Cisco UCS Manager GUI launch page.

Step 7 On the Cisco UCS Manager GUI launch page, select **Express Setup**.

Step 8 On the **Express Setup** page, select **Initial Setup** and click **Submit**.

Step 9 In the **Cluster and Fabric Setup** Area, select the **Standalone Mode** option.

Step 10 In the **System Setup** Area, complete the following fields:

Field	Description
System Name field	The name assigned to the Cisco UCS domain. In a standalone configuration, the system adds "-A" to the system name. In a cluster configuration, the system adds "-A" to the fabric interconnect assigned to fabric A, and "-B" to the fabric interconnect assigned to fabric B.
Admin Password field	The password used for the Admin account on the fabric interconnect. Choose a strong password that meets the guidelines for Cisco UCS Manager passwords. This password cannot be blank.
Confirm Admin Password field	The password used for the Admin account on the fabric interconnect.
Mgmt IP Address field	The static IPv4 or IPv6 address for the management port on the fabric interconnect.

Field	Description
Mgmt IP Netmask field or **Mgmt IP Prefix** field	The IPv4 subnet mask or IPv6 prefix for the management port on the fabric interconnect. **Note** The system prompts for a **Mgmt IP Netmask** or a **Mgmt IP Prefix** based on what address type you entered in the **Mgmt IP Address** field.
Default Gateway field	The IPv4 or IPv6 address for the default gateway assigned to the management port on the fabric interconnect. **Note** The system prompts for a **Default Gateway** address type based on what address type you entered in the **Mgmt IP Address** field
DNS Server IP field	The IPv4 or IPv6 address for the DNS server assigned to the fabric interconnect.
Domain Name field	The name of the domain in which the fabric interconnect resides.

Step 11 Click **Submit.**
A page displays the results of your setup operation.

Initial System Setup for a Cluster Configuration

Performing an Initial System Setup on the First Fabric Interconnect

Before You Begin

1 Verify the following physical connections on the fabric interconnect:

- A console port on the first fabric interconnect is physically connected to a computer terminal or console server
- The management Ethernet port (mgmt0) is connected to an external hub, switch, or router
- The L1 ports on both fabric interconnects are directly connected to each other
- The L2 ports on both fabric interconnects are directly connected to each other

For more information, refer to the *Cisco UCS Hardware Installation Guide* for your fabric interconnect.

2 Verify that the console port parameters on the computer terminal (or console server) attached to the console port are as follows:

- 9600 baud

- 8 data bits
- No parity
- 1 stop bit

3 Collect the following information that you will need to supply during the initial setup:

- System name.
- Password for the admin account. Choose a strong password that meets the guidelines for Cisco UCS Manager passwords. This password cannot be blank.
- Three static IPv4 or IPv6 addresses: two for the management port on both fabric interconnects (one per fabric interconnect) and one for the cluster IP address used by Cisco UCS Manager.
- Subnet mask for the three static IPv4 addresses, or network prefix for the three static IPv6 addresses.
- Default gateway IPv4 or IPv6 address.
- DNS server IPv4 or IPv6 address (optional).
- Domain name for the system (optional).

Procedure

Step 1 Power on the fabric interconnect.
You will see the power on self-test messages as the fabric interconnect boots. The system will run a DHCP client to check for a lease.

Step 2 If the system obtains a lease go to step 6, otherwise, continue to the next step.

Step 3 Connect to the console port.

Step 4 At the installation method prompt, enter gui.

Step 5 If the system cannot access a DHCP server, you are prompted to enter the following information:

- IPv4 or IPv6 address for the management port on the fabric interconnect
- IPv4 subnet mask or IPv6 prefix for the management port on the fabric interconnect
- IPv4 or IPv6 address for the default gateway assigned to the fabric interconnect

Note In a cluster configuration, both fabric interconnects must be assigned the same management interface address type during setup.

Step 6 Copy the web link from the prompt into a web browser and go to the Cisco UCS Manager GUI launch page.

Step 7 On the Cisco UCS Manager GUI launch page, select **Express Setup**.

Step 8 On the **Express Setup** page, select **Initial Setup** and click **Submit**.

Step 9 In the **Cluster and Fabric Setup** Area:

a) Click the **Enable Clustering** option.
b) For the **Fabric Setup** option, select **Fabric A**.
c) In the **Cluster IP Address** field, enter the IPv4 or IPv6 address that Cisco UCS Manager will use.

Step 10 In the **System Setup** Area, complete the following fields:

Field	Description
System Name field	The name assigned to the Cisco UCS domain. In a standalone configuration, the system adds "-A" to the system name. In a cluster configuration, the system adds "-A" to the fabric interconnect assigned to fabric A, and "-B" to the fabric interconnect assigned to fabric B.
Admin Password field	The password used for the Admin account on the fabric interconnect. Choose a strong password that meets the guidelines for Cisco UCS Manager passwords. This password cannot be blank.
Confirm Admin Password field	The password used for the Admin account on the fabric interconnect.
Mgmt IP Address field	The static IPv4 or IPv6 address for the management port on the fabric interconnect.
Mgmt IP Netmask field or **Mgmt IP Prefix** field	The IPv4 subnet mask or IPv6 prefix for the management port on the fabric interconnect. **Note** The system prompts for a **Mgmt IP Netmask** or a **Mgmt IP Prefix** based on what address type you entered in the **Mgmt IP Address** field.
Default Gateway field	The IPv4 or IPv6 address for the default gateway assigned to the management port on the fabric interconnect. **Note** The system prompts for a **Default Gateway** address type based on what address type you entered in the **Mgmt IP Address** field
DNS Server IP field	The IPv4 or IPv6 address for the DNS server assigned to the fabric interconnect.
Domain Name field	The name of the domain in which the fabric interconnect resides.

Step 11 Click **Submit**.
A page displays the results of your setup operation.

Performing an Initial System Setup on the Second Fabric Interconnect

Before You Begin

You must ensure the following:

- A console port on the second fabric interconnect is physically connected to a computer terminal or console server
- You know the password for the admin account on the first fabric interconnect that you configured.

Procedure

Step 1 Power on the fabric interconnect.
You will see the power on self-test messages as the fabric interconnect boots. The system will run a DHCP client to check for a lease.

Step 2 If the system obtains a lease go to step 6, otherwise, continue to the next step.

Step 3 Connect to the console port.

Step 4 At the installation method prompt, enter gui.

Step 5 If the system cannot access a DHCP server, you are prompted to enter the following information:

- IPv4 or IPv6 address for the management port on the fabric interconnect
- IPv4 subnet mask or IPv6 prefix for the management port on the fabric interconnect
- IPv4 or IPv6 address for the default gateway assigned to the fabric interconnect

Note In a cluster configuration, both fabric interconnects must be assigned the same management interface address type during setup.

Step 6 Copy the web link from the prompt into a web browser and go to the Cisco UCS Manager GUI launch page.

Step 7 On the Cisco UCS Manager GUI launch page, select **Express Setup**.

Step 8 On the **Express Setup** page, select **Initial Setup** and click **Submit**.
The fabric interconnect should detect the configuration information for the first fabric interconnect.

Step 9 In the **Cluster and Fabric Setup** Area:

a) Select the **Enable Clustering** option.
b) For the **Fabric Setup** option, make sure **Fabric B** is selected.

Step 10 In the **System Setup** Area, enter the password for the Admin account into the **Admin Password of Master** field.
The **Manager Initial Setup** Area is displayed.

Step 11 In the **Manager Initial Setup** Area, the field that is displayed depends on whether you configured the first fabric interconnect with an IPv4 or IPv6 management address. Complete the field that is appropriate for your configuration as follows:

Field	Description
Peer FI is IPv4 Cluster enabled. Please Provide Local Fabric Interconnect Mgmt0 IPv4 Address field	Enter an IPv4 address for the Mgmt0 interface on the local fabric interconnect.
Peer FI is IPv6 Cluster Enabled. Please Provide Local Fabric Interconnect Mgmt0 IPv6 Address field	Enter an IPv6 address for the Mgmt0 interface on the local fabric interconnect.

Step 12 Click **Submit**.
A page displays the results of your setup operation.

Adding Out-of-band IPv4 Addresses to a Fabric Interconnect

All fabric interconnects require an OOB IPv4 address, network mask and gateway. This procedure describes how to configure an OOB IPv4 address for a fabric interconnect that was set up with static IPv6 addresses.

Before You Begin

Collect the out-of-band (OOB) IPv4 address you want to assign to the fabric interconnect.

Procedure

	Command or Action	Purpose
Step 1	UCS-A # **scope fabric interconnect a**	Enters fabric configuration mode for Fabric A.
Step 2	UCS-A/fabric-interconnect # **set out-of-band ip** *ip-addr* **netmask** *ip-addr***gw** *ip-addr*	Sets the OOB IPv4 address, network mask and gateway address. The system warns that the console session change may be disconnected when the change is committed.
Step 3	UCS-A/fabric-interconnect # **commit-buffer**	Commits the transaction to the system configuration.

The following example shows configuring an OOB IPv4 address for fabric interconnect A:

```
UCS-A# scope fabric-interconnect a
UCS-A /fabric-interconnect # set out-of-band ip 10.105.214.107 netmask 255.255.255.0 gw
10.105.214.1
Warning: When committed, this change may disconnect the current CLI session
UCS-A /fabric-interconnect* # commit-buffer
```

Enabling a Standalone Fabric Interconnect for Cluster Configuration

You can add a second fabric interconnect to an existing Cisco UCS domain that uses a single standalone fabric interconnect. To do this, you must enable the standalone fabric interconnect for cluster operation by configuring it with the virtual IP or IPv6 address of the cluster, and then add the second fabric interconnect to the cluster.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# **connect local-mgmt**	Enters local management mode.
Step 2	UCS-A(local-mgmt) # **enable cluster** {*virtual-ip-addr*|*virtual-ip6-addr*}	Enables cluster operation on the standalone fabric interconnect with the specified IPv4 or IPv6 address. When you enter this command, you are prompted to confirm that you want to enable cluster operation. Type **yes** to confirm. The IP address must be the virtual IPv4 or IPv6 address for the cluster configuration, not the IP address assigned to the fabric interconnect that you are adding to the cluster.

The following example enables a standalone fabric interconnect with a virtual IPv4 address of 192.168.1.101 for cluster operation:

```
UCS-A# connect local-mgmt
UCS-A(local-mgmt)# enable cluster 192.168.1.101
This command will enable cluster mode on this setup. You cannot change it
back to stand-alone. Also, any GUI or KVM sessions may be terminated. Are you sure you want
 to continue? (yes/no): yes
UCS-A(local-mgmt)#
```

The following example enables a standalone fabric interconnect with a virtual IPv6 address of 192.168.1.101 for cluster operation:

```
UCS-A# connect local-mgmt
UCS-A(local-mgmt)# enable cluster ipv6 2001::109
This command will enable IPv6 cluster mode on this setup. You cannot change it
back to stand-alone. Also, any GUI or KVM sessions may be terminated. Are you sure you want
 to continue? (yes/no): yes
UCS-A(local-mgmt)#
```

What to Do Next

Add the second fabric interconnect to the cluster.

Configuring the Information Policy on the Fabric Interconnect

You must configure the information policy to display the uplink switches that are connected to Cisco UCS.

Important You must enable the information policy on the fabric interconnect to view the SAN, LAN, and LLDP neighbors of the fabric interconnect.

Enabling or Disabling the Information Policy on the Fabric Interconnect

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Click the **Equipment** node.

Step 3 In the **Work** pane, click the **Policies** tab.

Step 4 Click the **Global Policies** subtab.

Step 5 In the **Info Policy** area, select one of the following:

Option	Description
Disabled	Disable the information policy
Enabled	Enable the information policy

Step 6 Click **Save Changes**.

Viewing the LAN Neighbors of a Fabric Interconnect

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 In the **Equipment** tab, expand **Equipment** > **Fabric Interconnects**.

Step 3 Click the fabric interconnect for which you want to view the LAN neighbors.

Step 4 In the **Work** pane, click the **Neighbors** tab.

Step 5 Click the **LAN** subtab.
This subtab lists all the LAN neighbors of the specified Fabric Interconnect.

Viewing the SAN Neighbors of a Fabric Interconnect

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 In the **Equipment** tab, expand **Equipment** > **Fabric Interconnects**.

Step 3 Click the fabric interconnect for which you want to view the SAN neighbors.

Step 4 In the **Work** pane, click the **Neighbors** tab.

Step 5 Click the **SAN** subtab.
This subtab lists all the SAN neighbors of the specified Fabric Interconnect.

Viewing the LLDP Neighbors of a Fabric Interconnect

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 In the **Equipment** tab, expand **Equipment** > **Fabric Interconnects**.

Step 3 Click the fabric interconnect for which you want to view the LLDP neighbors.

Step 4 In the **Work** pane, click the **Neighbors** tab.

Step 5 Click the **LLDP** subtab.
This subtab lists all the LLDP neighbors of the specified Fabric Interconnect.

Fabric Evacuation

Cisco UCS Manager 2.2(4) introduces fabric evacuation, which is the ability to evacuate all traffic that flows through a Fabric Interconnect from all servers attached to it through an IOM or FEX while upgrading a system.

Upgrading the secondary Fabric Interconnect in a system disrupts the traffic that is active on the Fabric Interconnect. This traffic fails over to the primary Fabric Interconnect. You can use fabric evacuation as follows during the upgrade process:

1. Stop all the traffic that is active through a Fabric Interconnect.
2. For vNICs configured with failover, verify that the traffic has failed over by using Cisco UCS Manager or tools such as vCenter.
3. Upgrade the secondary Fabric Interconnect.
4. Restart all the stopped traffic flows.

5 Change the cluster lead to the secondary Fabric Interconnect.

6 Repeat steps 1 to 4 and upgrade the other Fabric Interconnect.

Note Fabric evacuation is supported only with the following:

- Manual install
- Cluster configuration

Configuring Fabric Evacuation

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click **Equipment**.	
Step 2	Expand **Equipment** > **Fabric Interconnects** > ***Fabric_Interconnect_Name***.	
Step 3	In the **Work** pane, click the **General** tab.	
Step 4	In the **Actions** area of the **General** tab, click **Configure Evacuation**.	The Configure Evacuation dialog box appears.
Step 5	To configure fabric evacuation on the specified Fabric Interconnect, click one of the following radio buttons in the **Admin Evac Mode** field: • **On—**Stops all the traffic that is active through the specified Fabric Interconnect. • **Off—**Restarts traffic through the specified Fabric Interconnect.	
Step 6	To evacuate a Fabric Interconnect irrespective of its current evacuation state, check the **Force** check box.	(Optional)

Step 7	Click **Apply**.	A warning dialog box appears.
		`Enabling fabric evacuation will stop all traffic through this Fabric Interconnect from servers attached through IOM/FEX.`
		`The traffic will fail over to the Primary Fabric Interconnect for fail over vnics.`
		`Are you sure you want to continue?`

Step 8	Click **OK** to confirm fabric evacuation and continue.	

Ethernet Switching Mode

The Ethernet switching mode determines how the fabric interconnect behaves as a switching device between the servers and the network. The fabric interconnect operates in either of the following Ethernet switching modes:

End-Host Mode

End-host mode allows the fabric interconnect to act as an end host to the network, representing all servers (hosts) connected to it through vNICs. This behavior is achieved by pinning (either dynamically pinned or hard pinned) vNICs to uplink ports, which provides redundancy to the network, and makes the uplink ports appear as server ports to the rest of the fabric. In end-host mode, the fabric interconnect does not run the Spanning Tree Protocol (STP) but it avoids loops by denying uplink ports from forwarding traffic to each other and by denying egress server traffic on more than one uplink port at a time. End-host mode is the default Ethernet switching mode and should be used if either of the following are used upstream:

- Layer 2 switching for Layer 2 aggregation
- Virtual Switching System (VSS) aggregation layer

Note When you enable end-host mode, if a vNIC is hard pinned to an uplink port and this uplink port goes down, the system cannot repin the vNIC, and the vNIC remains down.

Switch Mode

Switch mode is the traditional Ethernet switching mode. The fabric interconnect runs STP to avoid loops, and broadcast and multicast packets are handled in the traditional way. Switch mode is not the default Ethernet switching mode, and should be used only if the fabric interconnect is directly connected to a router, or if either of the following are used upstream:

- Layer 3 aggregation
- VLAN in a box

Note For both Ethernet switching modes, even when vNICs are hard pinned to uplink ports, all server-to-server unicast traffic in the server array is sent only through the fabric interconnect and is never sent through uplink ports. Server-to-server multicast and broadcast traffic is sent through all uplink ports in the same VLAN.

Configuring Ethernet Switching Mode

Important When you change the Ethernet switching mode, Cisco UCS Manager logs you out and restarts the fabric interconnect. For a cluster configuration, Cisco UCS Manager restarts both fabric interconnects. The subordinate fabric interconnect reboots first as a result of the change in switching mode. The primary fabric interconnect reboots only after you acknowledge it in **Pending Activities**. The primary fabric interconnect can take several minutes to complete the change in Ethernet switching mode and become system ready. The existing configuration is retained.

While the fabric interconnects are rebooting, all blade servers lose LAN and SAN connectivity, causing a complete outage of all services on the blades. This might cause the operating system to fail.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.

Step 3 In the **Work** pane, click the **General** tab.

Step 4 In the **Actions** area of the **General** tab, click one of the following links:

- **Set Ethernet Switching Mode**
- **Set Ethernet End-Host Mode**

The link for the current mode is dimmed.

Step 5 In the dialog box, click **Yes**.
Cisco UCS Manager restarts the fabric interconnect, logs you out, and disconnects Cisco UCS Manager GUI.

Fibre Channel Switching Mode

The Fibre Channel switching mode determines how the fabric interconnect behaves as a switching device between the servers and storage devices. The fabric interconnect operates in either of the following Fibre Channel switching modes:

End-Host Mode

End-host mode allows the fabric interconnect to act as an end host to the connected fibre channel networks, representing all servers (hosts) connected to it through virtual host bus adapters (vHBAs). This behavior is achieved by pinning (either dynamically pinned or hard pinned) vHBAs to Fibre Channel uplink ports, which makes the Fibre Channel ports appear as server ports (N-ports) to the rest of the fabric. When in end-host mode, the fabric interconnect avoids loops by denying uplink ports from receiving traffic from one another.

End-host mode is synonymous with N Port Virtualization (NPV) mode. This mode is the default Fibre Channel Switching mode.

Note When you enable end-host mode, if a vHBA is hard pinned to an uplink Fibre Channel port and this uplink port goes down, the system cannot repin the vHBA, and the vHBA remains down.

Switch Mode

Switch mode is the traditional Fibre Channel switching mode. Switch mode allows the fabric interconnect to connect directly to a storage device. Enabling Fibre Channel switch mode is useful in Pod models where there is no SAN (for example, a single Cisco UCS domain that is connected directly to storage), or where a SAN exists (with an upstream MDS).

Switch mode is not the default Fibre Channel switching mode.

Note In Fibre Channel switch mode, SAN pin groups are irrelevant. Any existing SAN pin groups are ignored.

Cisco UCS Fabric Interconnect in Switch Mode with Cisco MDS 9000 Family Fibre Channel Switching Modules

While creating a port channel between a Cisco MDS 9000 family FC switching module and a Cisco UCS Fabric Interconnect in switch mode, use the following order:

1. Create the port channel on the MDS side.
2. Add the port channel member ports.
3. Create the port channel on the Fabric Interconnect side.
4. Add the port channel member ports.

If you create the port channel on the Fabric Interconnect side first, the ports will go into a suspended state.

When the Cisco UCS Fabric Interconnect is in switch mode, the port channel mode can only be in **ON** mode and not **Active**. However, to get the peer wwn information for the Fabric Interconnect, the port channel must be in **Active** mode.

Configuring Fibre Channel Switching Mode

Important When you change the Fibre Channel switching mode, Cisco UCS Manager logs you out and restarts the fabric interconnect. For a cluster configuration, Cisco UCS Manager restarts both fabric interconnects simultaneously in Cisco UCS Manager Release 3.1(1) and earlier releases. In Cisco UCS Manager Release 3.1(2), when the Fibre Channel switching mode is changed, the UCS fabric interconnects reload sequentially. In Cisco UCS Manager Release 3.1(3), and later releases, the subordinate fabric interconnect reboots first as a result of the change in switching mode. The primary fabric interconnect reboots only after you acknowledge it in **Pending Activities**.The primary fabric interconnect can take several minutes to complete the change in Fibre Channel switching mode and become system ready.

Note When the Fibre Channel switching mode is changed, both UCS fabric interconnects will reload simultaneously. Reloading of fabric interconnects will cause a system-wide downtime for approximately 10-15 minutes.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Fabric Interconnects** > ***Fabric_Interconnect_Name***.

Step 3 In the **Work** pane, click the **General** tab.

Step 4 In the **Actions** area of the **General** tab, click one of the following links:

- **Set Fibre Channel Switching Mode**
- **Set Fibre Channel End-Host Mode**

The link for the current mode is dimmed.

Step 5 In the dialog box, click **Yes**.
Cisco UCS Manager restarts the fabric interconnect, logs you out, and disconnects Cisco UCS Manager GUI.

Changing the Properties of the Fabric Interconnects

Note To change the subnet or network prefix for a Cisco UCS domain, you must simultaneously change all subnets or prefixes, the virtual IPv4 or IPv6 address used to access Cisco UCS Manager, and the IPv4 or IPv6 addresses for both fabric interconnects.

Both fabric interconnects must maintain the same management address type, either IPv4 or IPv6. You cannot change the management address type for Fabric A without changing the management address type for Fabric B.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 On the **Admin** tab, click **All**.

Step 3 In the **Work** pane, click the **General** tab.

Step 4 In the **Actions** area, click **Management Interfaces** to open the **Management Interfaces** dialog box.

Step 5 In the **Management Interfaces** dialog box, modify the values as necessary.

Step 6 To change only the virtual IP address that you use to access Cisco UCS Manager, enter the desired IP address in either the **IPv4 Address** or the **IPv6 Address** field in the **Virtual IP** area.

Step 7 To change only the name assigned to the Cisco UCS domain, enter the desired name in the **Name** field in the **Virtual IP** area.

Step 8 To change the subnet and IPv4 address, or the network prefix and IPv6 address, and default gateway assigned to the fabric interconnects, update the following fields:

a) In the **Virtual IP** area, change the IP address used to access Cisco UCS Manager in the **IPv4 Address** or **IPv6 Address** field.

b) In the **Fabric Interconnect** area for each fabric interconnect, click either the IPv4 or IPv6 tab.

c) On the IPv4 tab, update the IP address, subnet mask, and default gateway.

d) On the IPv6 tab, update the IP address, prefix, and default gateway.

Step 9 Click **OK**.

Step 10 Log out of Cisco UCS Manager GUI and log back in again to see your changes.

Determining the Leadership Role of a Fabric Interconnect

Important To determine the role of the fabric interconnects in a cluster when the admin password is lost, open the Cisco UCS Manager GUI from the IP addresses of both fabric interconnects. The subordinate fabric interconnect fails with the following message:

```
UCSM GUI is not available on secondary node.
```

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 In the **Equipment** tab, expand **Equipment** > **Fabric Interconnects**.

Step 3 Click the fabric interconnect for which you want to identify the role.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **General** tab, click the down arrows on the **High Availability Details** bar to expand that area.

Step 6 View the **Leadership** field to determine whether the fabric interconnect is the primary or subordinate.

CHAPTER 6

Configuring Ports and Port Channels

This chapter includes the following sections:

- Fabric Port Channels, page 115
- Configuring Server Ports with the Internal Fabric Manager, page 118

Server and Uplink Ports on the 6100 Series Fabric Interconnect

Each Cisco UCS 6100 Series Fabric Interconnect has a set of ports in a fixed port module that you can configure as either server ports or uplink Ethernet ports. These ports are not reserved. They cannot be used by a Cisco UCS domain until you configure them. You can add expansion modules to increase the number of uplink ports on the fabric interconnect or to add uplink Fibre Channel ports to the fabric interconnect.

Note When you configure a port on a fabric interconnect, the administrative state is automatically set to enabled. If the port is connected to another device, this may cause traffic disruption. You can disable the port after configuring it.

You need to create LAN pin groups and SAN pin groups to pin traffic from servers to an uplink port.

Note Ports on the Cisco UCS 6100 Series Fabric Interconnect are not unified. For more information on Unified Ports, see Unified Ports on the Fabric Interconnect.

Each fabric interconnect can include the following port types:

Server Ports

Server ports handle data traffic between the fabric interconnect and the adapter cards on the servers.

You can only configure server ports on the fixed port module. Expansion modules do not include server ports.

Uplink Ethernet Ports

Uplink Ethernet ports handle Ethernet traffic between the fabric interconnect and the next layer of the network. All network-bound Ethernet traffic is pinned to one of these ports.

By default, Ethernet ports are unconfigured. However, you can configure them to function in the following ways:

- Uplink
- FCoE
- Appliance

You can configure uplink Ethernet ports on either the fixed module or an expansion module.

Uplink Fibre Channel Ports

Uplink Fibre Channel ports handle FCoE traffic between the fabric interconnect and the next layer of the storage area network. All network-bound FCoE traffic is pinned to one of these ports.

By default, Fibre Channel ports are uplink. However, you can configure them to function as Fibre Channel storage ports. This is useful in cases where Cisco UCS requires a connection to a Direct-Attached Storage (DAS) device.

You can only configure uplink Fibre Channel ports on an expansion module. The fixed module does not include uplink Fibre Channel ports.

Unified Ports on the Fabric Interconnect

Unified ports are ports on the fabric interconnect that can be configured to carry either Ethernet or Fibre Channel traffic. These ports are not reserved. A Cisco UCS domain cannot use these ports until you configure them.

Note When you configure a port on a fabric interconnect, the administrative state is automatically set to enabled. If the port is connected to another device, this may cause traffic disruption. You can disable the port after configuring it.

Configurable beacon LEDs indicate which unified ports are configured for the selected port mode.

Port Modes

The port mode determines whether a unified port on the fabric interconnect is configured to carry Ethernet or Fibre Channel traffic. You configure the port mode in Cisco UCS Manager. However, the fabric interconnect does not automatically discover the port mode.

Changing the port mode deletes the existing port configuration and replaces it with a new logical port. Any objects associated with that port configuration, such as VLANs and VSANS, are also removed. There is no restriction on the number of times you can change the port mode for a unified port.

Port Types

The port type defines the type of traffic carried over a unified port connection.

By default, unified ports changed to Ethernet port mode are set to the Ethernet uplink port type. Unified ports changed to Fibre Channel port mode are set to the Fibre Channel uplink port type. You cannot unconfigure Fibre Channel ports.

Changing the port type does not require a reboot.

Ethernet Port Mode

When you set the port mode to Ethernet, you can configure the following port types:

- Server ports
- Ethernet uplink ports

- Ethernet port channel members
- FCoE ports
- Appliance ports
- Appliance port channel members
- SPAN destination ports
- SPAN source ports

Note For SPAN source ports, configure one of the port types and then configure the port as SPAN source.

Fibre Channel Port Mode

When you set the port mode to Fibre Channel, you can configure the following port types:

- Fibre Channel uplink ports
- Fibre Channel port channel members
- Fibre Channel storage ports
- FCoE Uplink ports
- SPAN source ports

Note For SPAN source ports, configure one of the port types and then configure the port as SPAN source.

TCP and UDP Ports

The tables below list the incoming and outgoing TCP and UDP ports used in Cisco UCS for management access.

Table 8: Incoming ports

Port	Interface	Protocol	Traffic type	Usage
23	CLI	Telnet	TCP	Cisco UCS Manager CLI access
22	CLI	SSH	TCP	Cisco UCS Manager CLI access
443	Static HTML	HTTPS	TCP	Cisco UCS Manager login page access
80	Static HTML	HTTP	TCP	Client download

Port	Interface	Protocol	Traffic type	Usage
443	XML	HTTPS	TCP	Cisco UCS Manager XML API access
80	XML	HTTP	TCP	Ports used by Cisco UCS Manager GUI and third party management stations.
23	Serial-over-LAN	Telnet	TCP	COM1 port access on a specified server
22	Serial-over-LAN	SSH	TCP	COM1 port access on a specified server
161	SNMP	SNMP	UDP	SNMP MIBs exposed for monitoring
623	IPMI-over-LAN	RMCP	UDP	IPMI access to BMCs
2068	KVM	Avocent Video Session	TCP	Data path for the BMCs
843	xmlPolicy		TCP	Adobe port used by KVM launcher
5988	CIM XML		TCP	Send CIM messages over HTTP

Table 9: Outgoing ports

Port	Service	Protocol	Traffic type	Usage
1812	AAA	RADIUS	UDP	AAA server authentication requests
1813	AAA	RADIUS	UDP	AAA server authentication requests
49	AAA	TACACS	TCP	AAA server authentication requests
389	AAA	LDAP	UDP	
123	Time Sync	NTP	UDP	Synchronize the time with global time servers

Port	Service	Protocol	Traffic type	Usage
162	SNMP Traps	SNMP	UDP	Send traps to a remote network management system.
25	Call Home	SMTP	TCP	Email-based and web-based notifications for critical system events
514	Syslog	SYSLOG	UDP	Cisco UCS Manager generated Syslog messages
53	Name resolution	DNS	UDP	DNS queries
69	TFTP	TFTP	UDP	File transfers
115	SFTP	SFTP	TCP	File transfers
20-21	FTP	FTP	TCP	File transfers
21	SCP	SCP	TCP	File transfers

Beacon LEDs for Unified Ports

Each port on the 6200 series fabric interconnect has a corresponding beacon LED. When the **Beacon LED** property is configured, the beacon LEDs illuminate, showing you which ports are configured in a given port mode.

You can configure the **Beacon LED** property to show you which ports are grouped in one port mode: either Ethernet or Fibre Channel. By default, the Beacon LED property is set to Off.

Note For unified ports on the expansion module, you can reset the **Beacon LED** property to the default value of **Off** during expansion module reboot.

Guidelines for Configuring Unified Ports

Consider the following guidelines and restrictions when configuring unified ports:

Hardware and Software Requirements

Unified ports are supported on the 6200 series fabric interconnect with Cisco UCS Manager, version 2.0.

Unified ports are not supported on 6100 series fabric interconnects, even if they are running Cisco UCS Manager, version 2.0.

Port Mode Placement

Because the Cisco UCS Manager GUI interface uses a slider to configure the port mode for unified ports on a fixed or expansion module, it automatically enforces the following restrictions which limits how port modes can be assigned to unified ports. When using the Cisco UCS Manager CLI interface, these restrictions are enforced when you commit the transaction to the system configuration. If the port mode configuration violates any of the following restrictions, the Cisco UCS Manager CLI displays an error:

- Ethernet ports must be grouped together in a block. For each module (fixed or expansion), the Ethernet port block must start with the first port and end with an even numbered port.
- Fibre Channel ports must be grouped together in a block. For each module (fixed or expansion), the first port in the Fibre Channel port block must follow the last Ethernet port and extend to include the rest of the ports in the module. For configurations that include only Fibre Channel ports, the Fibre Channel block must start with the first port on the fixed or expansion module.
- Alternating Ethernet and Fibre Channel ports is not supported on a single module.

Example of a valid configuration— Might include unified ports 1–16 on the fixed module configured in Ethernet port mode and ports 17–32 in Fibre Channel port mode. On the expansion module you could configure ports 1–4 in Ethernet port mode and then configure ports 5–16 in Fibre Channel mode. The rule about alternating Ethernet and Fibre Channel port types is not violated because this port arrangement complies with the rules on each individual module.

Example of an invalid configuration— Might include a block of Fibre Channel ports starting with port 16. Because each block of ports has to start with an odd-numbered port, you would have to start the block with port 17.

Note The total number of uplink Ethernet ports and uplink Ethernet port channel members that can be configured on each fabric interconnect is limited to 31. This limitation includes uplink Ethernet ports and uplink Ethernet port channel members configured on the expansion module.

Cautions and Guidelines for Configuring Unified Uplink Ports and Unified Storage Ports

The following are cautions and guidelines to follow while working with unified uplink ports and unified storage ports:

- In an unified uplink port, if you enable one component as a SPAN source, the other component will automatically become a SPAN source.

Note If you create or delete a SPAN source under the Ethernet uplink port, Cisco UCS Manager automatically creates or deletes a SPAN source under the FCoE uplink port. The same happens when you create a SPAN source on the FCOE uplink port.

- You must configure a non default native VLAN on FCoE and unified uplink ports. This VLAN is not used for any traffic. Cisco UCS Manager will reuse an existing fcoe-storage-native-vlan for this purpose. This fcoe-storage-native-vlan will be used as a native VLAN on FCoE and unified uplinks.
- In an unified uplink port, if you do not specify a non default VLAN for the Ethernet uplink port the fcoe-storage-native-vlan will be assigned as the native VLAN on the unified uplink port. If the Ethernet port has a non default native VLAN specified as native VLAN, this will be assigned as the native VLAN for unified uplink port.
- When you create or delete a member port under an Ethernet port channel, Cisco UCS Manager automatically creates or deletes the member port under FCoE port channel. The same happens when you create or delete a member port in FCoE port channel.
- When you configure an Ethernet port as a standalone port, such as server port, Ethernet uplink, FCoE uplink or FCoE storage and make it as a member port for an Ethernet or FCOE port channel, Cisco UCS Manager automatically makes this port as a member of both Ethernet and FCoE port channels.
- When you remove the membership for a member port from being a member of server uplink, Ethernet uplink, FCoE uplink or FCoE storage, Cisco UCS Manager deletes the corresponding members ports from Ethernet port channel and FCoE port channel and creates a new standalone port.
- If you downgrade Cisco UCS Manager from release 2.1 to any of the prior releases, all unified uplink ports and port channels will be converted to Ethernet ports and Ethernet port channels when the downgrade is complete. Similarly, all the unified storage ports will be converted to appliance ports.
- For unified uplink ports and unified storage ports, when you create two interfaces, only one license is checked out. As long as either interface is enabled, the license remains checked out. The license will be released only if both the interfaces are disabled for a unified uplink port or a unified storage port.
- Cisco UCS 6100 series fabric interconnect switch can only support 1VF or 1VF-PO facing same downstream NPV switch.

Effect of Port Mode Changes on Data Traffic

Port mode changes can cause an interruption to the data traffic for the Cisco UCS domain. The length of the interruption and the traffic that is affected depend upon the configuration of the Cisco UCS domain and the module on which you made the port mode changes.

Tip To minimize the traffic disruption during system changes, form a Fibre Channel uplink port-channel across the fixed and expansion modules.

Impact of Port Mode Changes on an Expansion Module

After you make port mode changes on an expansion module, the module reboots. All traffic through ports on the expansion module is interrupted for approximately one minute while the module reboots.

Impact of Port Mode Changes on the Fixed Module in a Cluster Configuration

A cluster configuration has two fabric interconnects. After you make port changes to the fixed module, the fabric interconnect reboots. The impact on the data traffic depends upon whether or not you have configured the server vNICs to failover to the other fabric interconnect when one fails.

If you change the port modes on the expansion module of one fabric interconnect and then wait for that to reboot before changing the port modes on the second fabric interconnect, the following occurs:

- With server vNIC failover, traffic fails over to the other fabric interconnect and no interruption occurs.
- Without server vNIC failover, all data traffic through the fabric interconnect on which you changed the port modes is interrupted for approximately eight minutes while the fabric interconnect reboots.

If you change the port modes on the fixed modules of both fabric interconnects simultaneously, all data traffic through the fabric interconnects are interrupted for approximately eight minutes while the fabric interconnects reboot.

Impact of Port Mode Changes on the Fixed Module in a Standalone Configuration

A standalone configuration has only one fabric interconnect. After you make port changes to the fixed module, the fabric interconnect reboots. All data traffic through the fabric interconnect is interrupted for approximately eight minutes while the fabric interconnect reboots.

Configuring Port Modes for a 6248 Fabric Interconnect

Caution Changing the port mode on either module can cause an interruption in data traffic because changes to the fixed module require a reboot of the fabric interconnect and changes on an expansion module require a reboot of that module .

If the Cisco UCS domain has a cluster configuration that is set up for high availability and servers with service profiles that are configured for failover, traffic fails over to the other fabric interconnect and data traffic is not interrupted when the port mode is changed on the fixed module.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.

Step 3 In the **Work** pane, click the **General** tab.

Step 4 In the **Actions** area of the **General** tab, click **Configure Unified Ports**.

Step 5 Review the confirmation message and click one of the following:

- **Yes—**To continue with configuring the port mode.
- **No—**To exit without configuring the port mode and wait for an appropriate maintenance window.

Step 6 Click one of the following buttons to choose the module for which you want to configure the port modes:

- **Configure Fixed Module**
- **Configure Expansion Module**

Step 7 Use your mouse to drag the slider along the bar until the displays shows the port mode configuration that you want for the module.
If you change the port mode for a previously configured port, the port returns to an unconfigured state.

Step 8 If you need to configure port modes for the other module, repeat Steps 6 and 7.

Step 9 Click **Finish** to save your port mode configuration.
Based on the module for which you configured the port modes, data traffic for the Cisco UCS domain is interrupted as follows:

- Fixed module—The fabric interconnect reboots. All data traffic through that fabric interconnect is interrupted. In a cluster configuration that provides high availability and includes servers with vNICs that are configured for failover, traffic fails over to the other fabric interconnect and no interruption occurs.

  It takes about 8 minutes for the fixed module to reboot.

- Expansion module—The module reboots. All data traffic through ports in that module is interrupted.

  It takes about 1 minute for the expansion module to reboot.

What to Do Next

Configure the port types for the ports. You can right-click on any port in the module display above the slider and configure that port for an available port type.

Configuring Port Modes for a 6296 Fabric Interconnect

Caution Changing the port mode on either module can cause an interruption in data traffic because changes to the fixed module require a reboot of the fabric interconnect and changes on an expansion module require a reboot of that module .

If the Cisco UCS domain has a cluster configuration that is set up for high availability and servers with service profiles that are configured for failover, traffic fails over to the other fabric interconnect and data traffic is not interrupted when the port mode is changed on the fixed module.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.

Step 3 In the **Work** pane, click the **General** tab.

Step 4 In the **Actions** area of the **General** tab, click **Configure Unified Ports**.

Step 5 Review the confirmation message and click one of the following:

- **Yes**—To open the **Configure Unified Ports** wizard and continue with configuring the port mode.
- **No**—To exit without configuring the port mode and wait for an appropriate maintenance window.

Step 6 On the **Configure Fixed Module Ports** page, do the following:

a) Use your mouse to drag the slider along the bar until the displays shows the port mode configuration that you want for the fixed module.

b) If you want to configure the port type for a port, right-click on any port in the module display above the slider and configure that port for an available port type.
c) Do one of the following:

- Click **Next** to configure the port mode for ports in expansion module 1.
- If you do not wish to configure the port mode for ports on the expansion modules, continue with Step 9.

If you change the port mode for a previously configured port, the port returns to an unconfigured state.

Step 7 On the **Configure Expansion Module 1 Ports** page, do the following:

a) Use your mouse to drag the slider along the bar until the displays shows the port mode configuration that you want for the expansion module.
b) If you want to configure the port type for a port, right-click on any port in the module display above the slider and configure that port for an available port type.
c) Do one of the following:

- Click **Next** to configure the port mode for ports in expansion module 2.
- If you do not wish to configure the port mode for ports on the remaining expansion modules, continue with Step 9.

If you change the port mode for a previously configured port, the port returns to an unconfigured state.

Step 8 If you need to configure port modes for expansion module 3, repeat Step 7.

Step 9 Click **Finish** to save your port mode configuration.
Based on the module for which you configured the port modes, data traffic for the Cisco UCS domain is interrupted as follows:

- Fixed module—The fabric interconnect reboots. All data traffic through that fabric interconnect is interrupted. In a cluster configuration that provides high availability and includes servers with vNICs that are configured for failover, traffic fails over to the other fabric interconnect and no interruption occurs.

  It takes about 8 minutes for the fixed module to reboot.
- Expansion module—The module reboots. All data traffic through ports in that module is interrupted.

  It takes about 1 minute for the expansion module to reboot.

Configuring the Beacon LEDs for Unified Ports

Complete the following task for each module for which you want to configure beacon LEDs.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.

Step 3 Depending upon the location of the unified ports for which you want to configure the beacon LEDs, click on one of the following:

- **Fixed Module**
- **Expansion Module**

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Properties** area, click one of the following radio buttons in the **Beacon LED** field:

- **Off—**All physical LEDs are off.
- **Eth—**The physical LEDs next to all Ethernet ports are on.
- **Fc—**The physical LEDs next to all Fibre Channel ports are on.

Step 6 Click **Save Changes**.

Server Ports

Configuring Server Ports

All of the port types listed are configurable on both the fixed and expansion module, including server ports, which are not configurable on the 6100 series fabric interconnect expansion module, but are configurable on the 6200 series fabric interconnect expansion module.

This task describes only one method of configuring ports. You can also configure ports from a right-click menu or in the LAN Uplinks Manager.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 In the **Equipment** tab, expand **Fabric Interconnects** > *Fabric_Interconnect_Name* > **Fixed Module** > **Ethernet Ports**.

Step 3 Click on a port under the **Ethernet Ports** node.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Reconfigure**.

Step 6 From the drop-down list choose **Configure as Server Port**.

Uplink Ethernet Ports

Configuring Uplink Ethernet Ports

You can configure uplink Ethernet ports on either the fixed module or an expansion module.

This task describes only one method of configuring uplink Ethernet ports. You can also configure uplink Ethernet ports from a right-click menu.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Fabric Interconnects** > ***Fabric_Interconnect_Name***.

Step 3 Expand the node for the ports that you want to configure.

Step 4 Click on one of the ports under the **Ethernet Ports** node.
If you want to reconfigure a server port, appliance port, or FCoE storage port, expand the appropriate node.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 In the **Actions** area, click **Reconfigure**.

Step 7 From the drop-down list choose **Configure as Uplink Port**.

What to Do Next

If desired, change the properties for the default flow control policy and admin speed of the uplink Ethernet port.

Changing the Properties of an Uplink Ethernet Port

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Fabric Interconnects** > ***Fabric_Interconnect_Name***.

Step 3 Expand the node for the ports that you want to configure.

Step 4 In the **Ethernet Ports** node, click the uplink Ethernet port that you want to change.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 In the **Actions** area, click **Show Interface**.

Step 7 In the **Properties** dialog box, complete the following fields:

a) (Optional) In the **User Label** field, enter a label to identify the port.

b) From the **Flow Control Policy** drop-down list, select a flow control policy to determine how the port sends and receives IEEE 802.3x pause frames when the receive buffer fills.

c) In the **Admin Speed** field, click one of the following radio buttons:

- 1Gbps
- 10Gbps

Step 8 Click **OK**.

Reconfiguring a Port on a Fabric Interconnect

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
Step 3 Expand the node for the ports that you want to reconfigure.
Step 4 Click the port or ports that you want to reconfigure.
Step 5 In the **Work** pane, click the **General** tab.
Step 6 In the **Actions** area, click **Reconfigure**.
Step 7 From the drop-down list choose which way you want the port reconfigured.

Example: Reconfiguring an Uplink Ethernet Port as a Server Port

1. Expand the **Ethernet Ports** node and select the port you want to reconfigure.
2. Follow steps 5 and 6 above.
3. From the drop-down list choose **Configure as Server Port**.

Enabling or Disabling a Port on a Fabric Interconnect

After you enable or disable a port on a fabric interconnect, wait for at least 1 minute before you re-acknowledge the chassis. If you re-acknowledge the chassis too soon, the pinning of server traffic from the chassis might not get updated with the changes to the port that you enabled or disabled.

You can enable or disable a port only when it is configured. If the port is unconfigured, the enable disable option is not active.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
Step 3 Expand the node for the ports that you want to enable or disable.
Step 4 Under the **Ethernet Ports** node, select a port.
Step 5 In the **Work** pane, click the **General** tab.
Step 6 In the **Actions** area, click **Enable Port** or **Disable Port**.
Step 7 If a confirmation dialog box displays, click **Yes**.
Step 8 Click **OK**.

Unconfiguring a Port on a Fabric Interconnect

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
Step 3 Expand the node for the ports that you want to unconfigure.
Step 4 Under the **Ethernet Ports** node, select a port.
Step 5 In the **Work** pane, click the **General** tab.
Step 6 In the **Actions** area, click **Unconfigure**.
Step 7 If a confirmation dialog box displays, click **Yes**.
Step 8 Click **OK**.

Appliance Ports

Appliance ports are only used to connect fabric interconnects to directly attached NFS storage.

Note When you create a new appliance VLAN, its IEEE VLAN ID is not added to the LAN Cloud. Therefore, appliance ports that are configured with the new VLAN remain down, by default, due to a pinning failure. To bring up these appliance ports, you have to configure a VLAN in the LAN Cloud with the same IEEE VLAN ID.

Cisco UCS Manager supports up to four appliance ports per fabric interconnect.

Configuring an Appliance Port

You can configure Appliance ports on either the fixed module or an expansion module.

This task describes only one method of configuring appliance ports. You can also configure appliance ports from the **General** tab for the port.

Note If you configure an appliance port when the uplink port is down, Cisco UCS Manager may display an error message stating that the appliance port has failed. This message is controlled by the **Action on Uplink Fail** option in the associated Network Control Policy. For details about this option, see Network Control Policy, on page 308.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Fabric Interconnects** > ***Fabric_Interconnect_Name***.

Step 3 Expand the node for the ports that you want to configure.

Step 4 Under the **Ethernet Ports** node, select a port.
If you want to reconfigure a server port, uplink Ethernet port, or FCoE storage port, expand the appropriate node.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 In the **Actions** area, click **Reconfigure**.

Step 7 From the drop-down list, click **Configure as Appliance Port**.

Step 8 If a confirmation dialog box displays, click **Yes**.

Step 9 In the **Configure as Appliance Port** dialog box, complete the required fields.

Step 10 In the **VLANs** area, do the following:

a) In the **Port Mode** field, click one of the following radio buttons to select the mode you want to use for the port channel:

- **Trunk—**Cisco UCS Manager GUI displays the VLANs Table that lets you choose the VLANs you want to use.
- **Access—**Cisco UCS Manager GUI displays the **Select VLAN** drop-down list that allows you to choose a VLAN to associate with this port or port channel.

With either mode, you can click the **Create VLAN** link to create a new VLAN.

Note If traffic for the appliance port needs to traverse the uplink ports, you must also define each VLAN used by this port in the LAN cloud. For example, you need the traffic to traverse the uplink ports if the storage is also used by other servers, or if you want to ensure that traffic fails over to the secondary fabric interconnect if the storage controller for the primary fabric interconnect fails.

b) If you clicked the **Trunk** radio button, complete the required fields in the VLANs table.

c) If you clicked the **Access** radio button, choose a VLAN from the **Select VLAN** drop-down list.

Step 11 (Optional) If you want to add an endpoint, check the **Ethernet Target Endpoint** check box and specify the name and MAC address.

Step 12 Click **OK**.

Modifying the Properties of an Appliance Port

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.

Step 3 Expand the node for the appliance port that you want to modify.

Step 4 Expand **Ethernet Ports**.

Step 5 Click the appliance port for which you want to modify the properties.

Step 6 In the **Work** pane, click the **General** tab.

Step 7 In the **Actions** area, click **Show Interface**.
You may need to expand or use the scroll bars in the **Properties** dialog box to see all the fields.

Step 8 In the **Properties** dialog box, modify the values as needed.

Step 9 Click **OK**.

FCoE and Fibre Channel Storage Ports

Configuring an FCoE Storage Port

You can configure FCoE storage ports on either the fixed module or an expansion module.

This task describes only one method of configuring FCoE storage ports. You can also configure FCoE storage ports from the **General** tab for the port.

Before You Begin

The Fibre Channel switching mode must be set to Switching for these ports to be valid. The storage ports cannot function in end-host mode.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.

Step 3 Depending upon the location of the ports you want to configure, expand one of the following:

- **Fixed Module**
- **Expansion Module**

Step 4 Click one or more of the ports under the **Ethernet Ports** node.
If you want to reconfigure an uplink Ethernet port, server port, or appliance port, expand the appropriate node.

Step 5 Right-click the selected port or ports and choose **Configure as FCoE Storage Port**.

Step 6 If a confirmation dialog box displays, click **Yes**.

Step 7 Click **OK**.

Configuring a Fibre Channel Storage Port

This task describes only one method of configuring FC storage ports. You can also configure FC storage ports from the **General** tab for the port.

Before You Begin

The Fibre Channel switching mode must be set to Switching for these ports to be valid. The storage ports cannot function in end-host mode.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.

Step 3 Expand the **Expansion Module** node.

Step 4 Click one or more of the ports under the **FC Ports** node.

Step 5 Right-click the selected port or ports and choose **Configure as FC Storage Port**.

Step 6 If a confirmation dialog box displays, click **Yes**.

Step 7 Click **OK**.

Restoring an Uplink Fibre Channel Port

This task describes only one method of restoring an FC storage port to function as an uplink FC port. You can also reconfigure FC storage ports from the **General** tab for the port.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.

Step 3 Expand the **Expansion Module** node.

Step 4 Click one or more of the ports under the **FC Ports** node.

Step 5 Right-click the selected port or ports and choose **Configure as Uplink Port**.

Step 6 If a confirmation dialog box displays, click **Yes**.

Step 7 Click **OK**.

FC Links Rebalancing

The FC uplinks balance automatically when FC Port Channels are utilized. To create FC Port Channels, refer to Creating an Uplink Fibre Channel Port Channel, on page 111.

For the FC uplinks that are not members of the Port Channels (Individual ISLs), load balancing is done according to the FC uplinks balancing algorithm. For a vHBA of a host or service profile to choose an available FC uplink, when FC uplink trunking is disabled, the uplink and vHBA must belong to the same VSAN

For each vHBA, the algorithm searches for an FC uplink in the following order:

1 Least used FC uplink based on the number of vHBAs currently bound to the uplink.

2 If FC uplinks are equally balanced, then round robin is used.

This process continues for all the other vHBAs. The algorithm also considers other parameters such as pre-fip/fip adapters and number of flogis. You may not see the least-used component when there are less than six flogis.

After a port configuration or any other uplink state changes, if the traffic passing through the FC uplinks is no longer balanced, you can re-balance the traffic by resetting the vHBA(s) on each adapter and allow the load balancing algorithm to evaluate for the current state of the FC uplinks.

Configuring FC Uplink Ports

You can configure FC Uplink port on either a fixed module or an expansion module.

This task describes only one method of configuring FC Uplink ports. You can also configure FC uplink ports from a right-click menu for the port.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.

Step 3 Expand the node for the ports that you want to configure.

Step 4 Under the **FC Ports** node, select any **Storage** port.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 From the **Actions** area, select **Configure as Uplink Port**.

Step 7 If a confirmation dialog box displays, click **Yes**.

Step 8 The Cisco UCS Manager GUI displays a success message.
In the **Actions** area, **Configure as Uplink Port** becomes greyed out and **Configure as FC Storage Port** becomes active.

FCoE Uplink Ports

FCoE uplink ports are physical Ethernet interfaces between the fabric interconnects and the upstream Ethernet switch, used for carrying FCoE traffic. With this support the same physical Ethernet port can carry both Ethernet traffic and Fibre Channel traffic.

FCoE uplink ports connect to upstream Ethernet switches using the FCoE protocol for Fibre Channel traffic. This allows both the Fibre Channel traffic and Ethernet traffic to flow on the same physical Ethernet link.

Note FCoE uplinks and unified uplinks enable the multi-hop FCoE feature, by extending the unified fabric up to the distribution layer switch.

You can configure the same Ethernet port as any of the following:

- **FCoE uplink port—**As an FCoE uplink port for only Fibre Channel traffic.
- **Uplink port—**As an Ethernet port for only Ethernet traffic.
- **Unified uplink port—**As a unified uplink port to carry both Ethernet and Fibre Channel traffic.

Configuring FCoE Uplink Ports

You can configure FCoE Uplink port on either a fixed module or an expansion module.

This task describes only one method of configuring FCoE Uplink ports. You can also configure FCoE uplink ports from a right-click menu or from the General tab for the port.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
Step 3 Expand the node for the ports that you want to configure.
Step 4 Under the **Ethernet Ports** node, select any **Unconfigured** port.
Step 5 In the **Work** pane, click the **General** tab.
Step 6 In the **Actions** area, click **Reconfigure**.
Step 7 From the drop down options, select **Configure as FCoE Uplink Port**.
Step 8 If a confirmation dialog box displays, click **Yes**.
Step 9 The Cisco UCS Manager GUI displays a success message.
In the **Properties** area, the **Role** changes to **Fcoe Uplink**.

Unified Storage Ports

Unified storage involves configuring the same physical port as both an Ethernet storage interface and an FCoE storage interface. You can configure any appliance port or FCoE storage port as a unified storage port, on either a fixed module or an expansion module. To configure a unified storage port, you must have the fabric interconnect in Fibre Channel switching mode.

In a unified storage port, you can enable or disable individual FCoE storage or appliance interfaces.

- In an unified storage port, if you do not specify a non-default VLAN for the appliance port, the FCoE-storage-native-vlan will be assigned as the native VLAN on the unified storage port. If the appliance port has a non-default native VLAN specified as native VLAN, this will be assigned as the native VLAN for the unified storage port.
- When you enable or disable the appliance interface, the corresponding physical port is enabled or disabled. So when you disable the appliance interface in unified storage, even if the FCoE storage is enabled, it goes down with the physical port.
- When you enable or disable the FCoE storage interface, the corresponding VFC is enabled or disabled. So when the FCoE storage interface is disabled in a unified storage port, the appliance interface will continue to function normally.

Configuring an Appliance Port as a Unified Storage Port

You can configure a unified storage port either from an appliance port or an FCoE storage port. You can also configure the unified storage port from an unconfigured port. If you start from an unconfigured port, you will assign either appliance or FCoE storage configuration to the port and then add another configuration to enable it as a unified storage port.

Important Make sure the FI is in FC switching mode.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.

Step 3 Depending upon the location of the ports you want to configure, expand one of the following:

- **Fixed Module**
- **Expansion Module**

Step 4 Under the **Ethernet Ports** node, select any the port that is already configured as an appliance port. In the **Work** pane, under **General** tab, in **Properties** area, the **Role** will show as **Appliance Storage**.

Step 5 In the **Actions** area, click **Reconfigure**.

Step 6 From the pop-up menu, select **Configure as FCoE Storage** Port.

Step 7 If a confirmation dialog box displays, click **Yes**.

Step 8 The Cisco UCS Manager GUI displays a success message. In the **Properties** area, the **Role** changes to **Unified Storage**.

Unconfiguring a Unified Storage Port

You can unconfigure and remove both configurations from the unified connect port. Or you can unconfigure either one of them and retain the other one on the port.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.

Step 3 Expand the node for the ports that you want to unconfigure.

Step 4 Under the **Ethernet Ports** node, select the port you want to unconfigure.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 In the **Actions** area, click **Unconfigure**. You will see the following options:

- **Unconfigure FCoE Storage Port**
- **Unconfigure Appliance Port**
- **Unconfigure both**

Step 7 Select one of the unconfigure options.

Step 8 If a confirmation dialog box displays, click **Yes**.

Step 9 The Cisco UCS Manager GUIdisplays a success message. In the **Properties** area, the **Role** changes to based on your unconfigure selection.

Unified Uplink Ports

When you configure an Ethernet uplink and an FCoE uplink on the same physical Ethernet port, it is called a unified uplink port. You can individually enable or disable either the FCoE or Ethernet interfaces independently.

- Enabling or disabling the FCoE uplink results in the corresponding VFC being enabled or disabled.
- Enabling or disabling an Ethernet uplink results in the corresponding physical port being enabled or disabled.

If you disable an Ethernet uplink, it disables the underlying physical port in a unified uplink. Therefore, even when the FCoE uplink is enabled, the FCoE uplink also goes down. But if you disable an FCoE uplink, only the VFC goes down. If the Ethernet uplink is enabled, it can still function properly in the unified uplink port.

Configuring Unified Uplink Ports

You can configure the unified uplink port from either one of the following:

- From existing FCoE uplink or Ethernet Uplink port
- From an unconfigured uplink port

This process describes one method to configure an unified uplink port from an existing FCoE uplink port. You can configure the unified uplink port on either a fixed module or an expansion module. suppress ucs-mini

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.

Step 3 Expand the node for the ports that you want to configure.

Step 4 Under the **Ethernet Ports** node, select a port.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 In the **Properties** area, make sure the **Role** shows as **Fcoe Uplink**.

Step 7 In the **Actions** area, click **Reconfigure**.

Step 8 From the drop down options, select **Configure as Uplink Port**.

Step 9 If a confirmation dialog box displays, click **Yes**.

Step 10 The Cisco UCS Manager GUI displays a success message.

In the **Properties** area, the **Role** changes to **Unified Uplink**.

Step 11 (Optional) In the **Properties** area, specify the **VSAN** in the **VSAN** field.

Unconfiguring Unified Uplink Port

You can unconfigure and remove both configurations from the unified uplink port. Or you can unconfigure either one of the FCoE or Ethernet port configuration and retain the other one on the port.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.

Step 3 Expand the node for the ports that you want to unconfigure.

Step 4 Under the **Ethernet Ports** node, select the port you want to unconfigure.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 In the **Actions** area, click **Unconfigure**. Select one of the following options:

- **Unconfigure FCoE Uplink Port**
- **Unconfigure Uplink Port**
- **Unconfigure both**

Step 7 If a confirmation dialog box displays, click **Yes**.

Step 8 The Cisco UCS Manager GUI displays a success message. In the **Properties** area, the **Role** changes based on your unconfigure selection.

Step 9 Click **Save Changes**.

Uplink Ethernet Port Channels

An uplink Ethernet port channel allows you to group several physical uplink Ethernet ports (link aggregation) to create one logical Ethernet link to provide fault-tolerance and high-speed connectivity. In Cisco UCS Manager, you create a port channel first and then add uplink Ethernet ports to the port channel. You can add up to 16 uplink Ethernet ports to a port channel.

Important The state of a configured port changes to unconfigured in the following scenarios:

- The port is deleted or removed from a port channel. The port channel can be of any type, such as, uplink or storage.
- A port channel is deleted.

Note Cisco UCS uses Link Aggregation Control Protocol (LACP), not Port Aggregation Protocol (PAgP), to group the uplink Ethernet ports into a port channel. If the ports on the upstream switch are not configured for LACP, the fabric interconnects treat all ports in an uplink Ethernet port channel as individual ports, and therefore forward packets.

Creating an Uplink Ethernet Port Channel

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **LAN Cloud**.

Step 3 Expand the node for the fabric interconnect where you want to add the port channel.

Step 4 Right-click the **Port Channels** node and choose **Create Port Channel**.

Step 5 In the **Set Port Channel Name** panel, specify the ID and name, then click **Next**.

Step 6 In the **Add Ports** panel, specify the ports that you want to add.

Note Cisco UCS Manager warns you if you select a port that has been configured as a server port. You can click **Yes** in the dialog box to reconfigure that port as an uplink Ethernet port and include it in the port channel.

Step 7 Click **Finish**.

Enabling an Uplink Ethernet Port Channel

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 Expand **LAN** > **LAN Cloud**.
Step 3 Expand the node for the fabric interconnect that includes the port channel you want to enable.
Step 4 Expand the **Port Channels** node.
Step 5 Right-click the port channel you want to enable and choose **Enable Port Channel**.
Step 6 If a confirmation dialog box displays, click **Yes**.

Disabling an Uplink Ethernet Port Channel

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 Expand **LAN** > **LAN Cloud**.
Step 3 Expand the node for the fabric interconnect that includes the port channel you want to disable.
Step 4 Expand the **Port Channels** node.
Step 5 Right-click the port channel you want to disable and choose **Enable Port Channel**.

Adding Ports to and Removing Ports from an Uplink Ethernet Port Channel

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 Expand **LAN** > **LAN Cloud** > *Fabric* > **Port Channels**.
Step 3 Click the port channel to which you want to add or remove ports.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Add Ports**.
Step 6 In the **Add Ports** dialog box, do one of the following:

- To add ports, choose one or more ports in the **Ports** table, and then click the **>>** button to add the ports to the **Ports in the port channel** table.

- To remove ports, choose one or more ports in the **Ports in the port channel** table, and then click the << button to remove the ports from the port channel and add them to the **Ports** table.

Step 7 Click **OK**.

Deleting an Uplink Ethernet Port Channel

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 Expand **LAN** > **LAN Cloud**.
Step 3 Expand the node for the fabric interconnect where you want to delete the port channel.
Step 4 Click the **Port Channels** node.
Step 5 In the **General** tab for the **Port Channels** node, choose the port channel you want to delete.
Step 6 Right-click the port channel and choose **Delete**.

Appliance Port Channels

An appliance port channel allows you to group several physical appliance ports to create one logical Ethernet storage link for the purpose of providing fault-tolerance and high-speed connectivity. In Cisco UCS Manager, you create a port channel first and then add appliance ports to the port channel. You can add up to eight appliance ports to a port channel.

Creating an Appliance Port Channel

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 Expand **LAN** > **Appliances**.
Step 3 Expand the node for the fabric interconnect where you want to add the port channel.
Step 4 Right-click the **Port Channels** node and choose **Create Port Channel**.
Step 5 In the **Set Port Channel Name** panel of the **Create Port Channel** wizard, complete the required fields to specify the identity and other properties of the port channel.
You can create a LAN pin group, network control policy, and flow control policy from this panel.
Step 6 In the **VLANs** area, specify the **Port Mode** and other information for the VLANs.
You can create a VLAN from this panel.

Step 7 (Optional) If you want to add an endpoint, check the **Ethernet Target Endpoint** check box specify the name and MAC address.

Step 8 Click **Next**.

Step 9 In the **Add Ports** panel of the **Create Port Channel** wizard, specify the ports that you want to add.

Note Cisco UCS Manager warns you if your configuration could cause issues with service profiles or port configurations. You can click **Yes** in the dialog box if you want to create the port channel despite those potential issues.

Step 10 Click **Finish**.

Enabling an Appliance Port Channel

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Appliances**.

Step 3 Expand the node for the fabric interconnect that includes the port channel you want to enable.

Step 4 Expand the **Port Channels** node.

Step 5 Right-click the port channel you want to enable and choose **Enable Port Channel**.

Step 6 If a confirmation dialog box displays, click **Yes**.

Disabling an Appliance Port Channel

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Appliances**.

Step 3 Expand the node for the fabric interconnect that includes the port channel you want to disable.

Step 4 Expand the **Port Channels** node.

Step 5 Right-click the port channel you want to disable and choose **Disable Port Channel**.

Step 6 If a confirmation dialog box displays, click **Yes**.

Adding Ports to and Removing Ports from an Appliance Port Channel

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Appliances** > *Fabric* > **Port Channels**.

Step 3 Click the port channel to which you want to add or remove ports.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Add Ports**.

Step 6 In the **Add Ports** dialog box, do one of the following:

- To add ports, choose one or more ports in the **Ports** table, and then click the **>>** button to add the ports to the **Ports in the port channel** table.
- To remove ports, choose one or more ports in the **Ports in the port channel** table, and then click the **<<** button to remove the ports from the port channel and add them to the **Ports** table.

Step 7 Click **OK**.

Deleting an Appliance Port Channel

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Appliances**.

Step 3 Expand the node for the fabric interconnect that includes the port channel you want to delete.

Step 4 Expand the **Port Channels** node.

Step 5 Right-click the port channel you want to enable and choose **Delete**.

Step 6 If a confirmation dialog box displays, click **Yes**.

Creating a Threshold Condition

Procedure

Step 1 In the **Navigation** pane, click the **Admin** tab.

Step 2 In the **Admin** tab, expand **All** > **Stats Management** > **fabric** > **Internal LAN** > **thr-policy-default**.

Step 3 Click **Create Threshold Class**.

Step 4 In the **Choose Statistics Class** screen of the **Create Threshold Class** wizard, choose **NI Ether Error Stats** statistics class to monitor the network interface ports for which you want to configure a custom threshold from the **Stat Class** drop-down list.

Step 5 Click **Next**.

Step 6 In the **Threshold Definitions** screen of the **Create Threshold Class** wizard, click **Add**.
The **Create Threshold Definition** dialog box opens.

a) From the **Property Type** field, choose the threshold property that you want to define for the class.
b) In the **Normal Value** field, enter the desired value for the property type.
c) In the **Alarm Triggers (Above Normal Value)** fields, check one or more of the following check boxes:

- **Critical**
- **Major**
- **Minor**
- **Warning**
- **Condition**
- **Info**

d) In the **Up and Down** fields, enter the range of values that should trigger the alarm.
e) In the **Alarm Triggers (Below Normal Value)** fields, check one or more of the following check boxes:

- **Critical**
- **Major**
- **Minor**
- **Warning**
- **Condition**
- **Info**

f) In the **Up and Down** fields, enter the range of values that should trigger the alarm.
g) Click **Ok**.

Monitoring a Fabric Port

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 On the **Equipment** tab, expand **Chassis** > **IO Modules** > **IO Module 1** > **Fabric Ports**.

Step 3 Click the fabric port that you want to monitor.

Step 4 Click one of the following tabs to view the status of the fabric:

Option	Description
General	Provides an overview of the status of the fabric, including a summary of any faults, a summary of the fabric properties, and a physical display of the fabric and its components.
Faults	Provides details of the faults generated by the fabric.
Events	Provides details of the events generated by the fabric.
Statistics	Provides statistics about the fabric and its components. You can view these statistics in tabular or chart format.

Policy-Based Port Error Handling

If Cisco UCS Manager detects any errors on active NI ports, and if the error-disable feature is enabled, Cisco UCS Manager automatically disables the respective FI port that is connected to the NI port that had errors. When a FI port is error disabled, it is effectively shut down and no traffic is sent or received on that port.

The error-disable function serves two purposes:

- It lets you know which FI port is error-disabled and that the connected NI Port has errors.
- It eliminates the possibility that this port can cause other ports, which are connected to the same Chassis/FEX, to fail. Such a failure can occur when the NI port has errors, which can ultimately cause serious network issues. The error-disable function helps prevent these situations.

Configuring Error-Based Action

Procedure

Step 1 In the **Navigation** pane, click the **Admin** tab.

Step 2 In the **Admin** tab, expand **All** > **Stats Management** > **fabric** > **Internal LAN** > **thr-policy-default** > **etherNiErrStats**.

Step 3 Select a delta property.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 To enable error disable on the FI port, check the **Disable FI port when fault is raised** check box.

Step 6 To enable auto recovery, in the **Enable Auto Recovery** field, select **Enable**.

Step 7 To specify the time after which the port can automatically be re-enabled, in the **Time (in minutes)** field, type the time in minutes.

Step 8 Click **Save Changes**.

Fibre Channel Port Channels

A Fibre Channel port channel allows you to group several physical Fibre Channel ports (link aggregation) to create one logical Fibre Channel link to provide fault-tolerance and high-speed connectivity. In Cisco UCS Manager, you create a port channel first and then add Fibre Channel ports to the port channel.

Note Fibre Channel port channels are not compatible with non-Cisco technology.

You can create up to four Fibre Channel port channels in each Cisco UCS domain with Cisco UCS 6200 and 6300 Series fabric interconnects. Each Fibre Channel port channel can include a maximum of 16 uplink Fibre Channel ports.

You can create up to two Fibre Channel port channels in each Cisco UCS domain with Cisco UCS 6324 fabric interconnects. Each Fibre Channel port channel can include a maximum of four uplink Fibre Channel ports.

Ensure that the Fibre Channel port channel on the upstream NPIV switch is configured with its channel mode as **active**. If both the member port(s) and peer port(s) do not have the same channel mode configured, the port channel will not come up. When the channel mode is configured as **active**, the member ports initiate port channel protocol negotiation with the peer port(s) regardless of the channel group mode of the peer port. If the peer port, while configured in a channel group, does not support the port channel protocol, or responds with a nonnegotiable status, it defaults to the On mode behavior. The **active** port channel mode allows automatic recovery without explicitly enabling and disabling the port channel member ports at either end.

This example shows how to configure channel mode as active:

```
switch(config)# int po114
switch(config-if)# channel mode active
```

Creating an Uplink Fibre Channel Port Channel

Procedure

Step 1 In the **Navigation** pane, click **SAN**.
Step 2 Expand **SAN** > **SAN Cloud**.
Step 3 Expand the node for the fabric where you want to create the port channel.
Step 4 Right-click the **FC Port Channels** node and choose **Create Port Channel**.
Step 5 In the **Set Port Channel Name** panel, specify the ID and name, then click **Next**.
Step 6 In the **Add Ports** panel, specify the port channel admin speed, and add ports to the port channel.
Step 7 Click **Finish**.

Enabling a Fibre Channel Port Channel

Procedure

Step 1 In the **Navigation** pane, click **SAN**.
Step 2 Expand **SAN** > **SAN Cloud** > *Fabric* > **FC Port Channels**.
Step 3 Click the port channel you want to enable.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Enable Port Channel**.
Step 6 If a confirmation dialog box displays, click **Yes**.

Disabling a Fibre Channel Port Channel

Procedure

Step 1 In the **Navigation** pane, click **SAN**.
Step 2 Expand **SAN** > **SAN Cloud** > *Fabric* > **FC Port Channels**.
Step 3 Click the port channel you want to disable.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Disable Port Channel**.
Step 6 If a confirmation dialog box displays, click **Yes**.

Adding Ports to and Removing Ports from a Fibre Channel Port Channel

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **SAN Cloud** > *Fabric* > **FC Port Channels**.

Step 3 Click the port channel to which you want to add or remove ports.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Add Ports**.

Step 6 In the **Add Ports** dialog box, do one of the following:

- To add ports, choose one or more ports in the **Ports** table, and then click the **>>** button to add the ports to the **Ports in the port channel** table.
- To remove ports, choose one or more ports in the **Ports in the port channel** table, and then click the **<<** button to remove the ports from the port channel and add them to the **Ports** table.

Step 7 Click **OK**.

Modifying the Properties of a Fibre Channel Port Channel

Note If you are connecting two Fibre Channel port channels, the admin speed for both port channels must match for the link to operate. If the admin speed for one or both of the Fibre Channel port channels is set to auto, Cisco UCS adjusts the admin speed automatically.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **SAN Cloud** > *Fabric* > **FC Port Channels**.

Step 3 Click the port channel that you want to modify.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Properties** area, change the values in one or more of the following fields:

Name	Description
Name field	The user-defined name given to the port channel. This name can be between 1 and 16 alphanumeric characters.
VSAN drop-down list	The VSAN associated with the port channel.

Name	Description
Port Channel Admin Speed drop-down list	The admin speed of the port channel. This can be: • **1 Gbps** • **2 Gbps** • **4 Gbps** • **8 Gbps** • **auto**

Step 6 Click **Save Changes**.

Deleting a Fibre Channel Port Channel

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 Expand **SAN** > **SAN Cloud** > *Fabric* > **FC Port Channels**.
Step 3 Right-click the port channel you want to delete and choose **Delete**.
Step 4 If a confirmation dialog box displays, click **Yes**.

FCoE Port Channels

An FCoE port channel allows you to group several physical FCoE ports to create one logical FCoE port channel. At a physical level, the FCoE port channel carries FCoE traffic over an Ethernet port channel. So an FCoE port channel with a set of members is essentially an ethernet port channel with the same members. This Ethernet port channel is used as a physical transport for FCoE traffic.

For each FCoE port channel, Cisco UCS Manager creates a VFC internally and binds it to an Ethernet port channel. FCoE traffic received from the hosts is sent over the VFC the same way as the FCoE traffic is sent over Fibre Channel uplinks.

Creating an FCoE Port Channel

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **SAN Cloud**.

Step 3 Expand the node for the fabric where you want to create the port channel.

Step 4 Right-click the **FCoE Port Channels** node and choose **Create FCoE Port Channel**.

Step 5 In the **Set Port Channel Name** panel of the **Create FCoE Port Channel** wizard, specify the ID and name, then click **Next**.

Step 6 In the **Add Ports** panel of the **Create FCoE Port Channel** wizard, specify the ports that you want to add.

Step 7 Click **Finish**.

Deleting an FCoE Port Channel

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 On the **SAN** tab, expand **SAN** > **SAN Cloud** > ***Fabric*** > **FCoE Port Channels**.

Step 3 Right-click the port channel you want to delete and choose **Delete**.

Step 4 If a confirmation dialog box displays, click **Yes**.

Unified Uplink Port Channel

When you create an Ethernet port channel and an FCoE port channel with the same ID, it is called a unified uplink port channel. When the unified port channel is created, a physical Ethernet port channel and a VFC are created on the fabric interconnect with the specified members. The physical Ethernet port channel is used to carry both Ethernet and FCoE traffic. The VFC binds FCoE traffic to the Ethernet port channel.

The following rules will apply to the member port sets of the unified uplink port channel:

- The Ethernet port channel and FCoE port channel on the same ID, must have the same set of member ports.
- When you add a member port channel to the Ethernet port channel, Cisco UCS Manager adds the same port channel to FCoE port channel as well. Similarly, adding a member to the FCoE port channel adds the member port to the Ethernet port channel.
- When you delete a member port from one of the port channels, Cisco UCS Manager automatically deletes the member port from the other port channel.

If you disable an Ethernet uplink port channel, it disables the underlying physical port channel in a unified uplink port channel. Therefore, even when the FCoE uplink is enabled, the FCoE uplink port channel also goes down. If you disable an FCoE uplink port channel, only the VFC goes down. If the Ethernet uplink port channel is enabled, it can still function properly in the unified uplink port channel.

Adapter Port Channels

An adapter port channel groups into one logical link all the physical links going from a Cisco UCS Virtual Interface Card (VIC) into an I/O.

Adapter port channels are created and managed internally by Cisco UCS Manager when it detects that the correct hardware is present. Adapter port channels cannot be configured manually. Adapter port channels are viewable using the Cisco UCS Manager GUI or the Cisco UCS Manager CLI.

Viewing Adapter Port Channels

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 On the **Equipment** tab, expand **Equipment** > **Chassis** > ***Chassis_Number*** > **Servers** > ***Server_Number*** > **Interface Cards**

Step 3 Click the adapter for which you want to view the adapter port channels.

Step 4 In the **Work** pane, click the **DCE Interfaces** tab.

Step 5 To view details of the adapter port channel, click the link in the **Port Channel** column.

Fabric Port Channels

Fabric port channels allow you to group several of the physical links from an IOM to a fabric interconnect into one logical link for redundancy and bandwidth sharing. As long as one link in the fabric port channel remains active, the fabric port channel continues to operate.

If the correct hardware is connected, fabric port channels are created by Cisco UCS Manager in the following ways:

- During chassis discovery according to the settings configured in the chassis discovery policy.
- After chassis discovery according to the settings configured in the chassis connectivity policy for a specific chassis.

For each IOM there is a single fabric port channel. Each uplink connecting an IOM to a fabric interconnect can be configured as a discrete link or included in the port channel, but an uplink cannot belong to more than one fabric port channel. For example, if a chassis with two IOMs is discovered and the chassis discovery policy is configured to create fabric port channels, Cisco UCS Manager creates two separate fabric port channels: one for the uplinks connecting IOM-1 and another for the uplinks connecting IOM-2. No other

chassis can join these fabric port channels. Similarly, uplinks belonging to the fabric port channel for IOM-1 cannot join the fabric port channel for IOM-2.

Load Balancing Over Ports

Load balancing traffic among ports between IOMs and fabric interconnects uses the following criteria for hashing.

- For Ethernet traffic:

  Layer 2 source and destination address

  Layer 3 source and destination address

  Layer 4 source and destination ports

- For FCoE traffic:

  Layer 2 source and destination address

  Source and destination IDs (SID and DID) and Originator Exchange ID (OXID)

In this example, a 2200 Series IOM module is verified by connecting iom *X* (where *X* is the chassis number).

```
show platform software fwmctrl nifport
(....)
 Hash Parameters:
   l2_da: 1 l2_sa: 1 l2_vlan: 0
   l3_da: 1 l3_sa: 1
   l4_da: 1 l4_sa: 1
   FCoE l2_da: 1 l2_sa: 1 l2_vlan: 0
   FCoE l3_did: 1 l3_sid: 1 l3_oxid: 1
```

Cabling Considerations for Fabric Port Channels

When you configure the links between the Cisco UCS 2200 Series FEX and a Cisco UCS 6200 series fabric interconnect in fabric port channel mode, the available virtual interface namespace (VIF) on the adapter varies depending on where the FEX uplinks are connected to the fabric interconnect ports.

Inside the 6248 fabric interconnect there are six sets of eight contiguous ports, with each set of ports managed by a single chip. When all uplinks from an FEX are connected to a set of ports managed by a single chip, Cisco UCS Manager maximizes the number of VIFs used in service profiles deployed on the blades in the chassis. If uplink connections from an IOM are distributed across ports managed by separate chips, the VIF count is decreased.

Figure 1: Port Groups for Fabric Port Channels

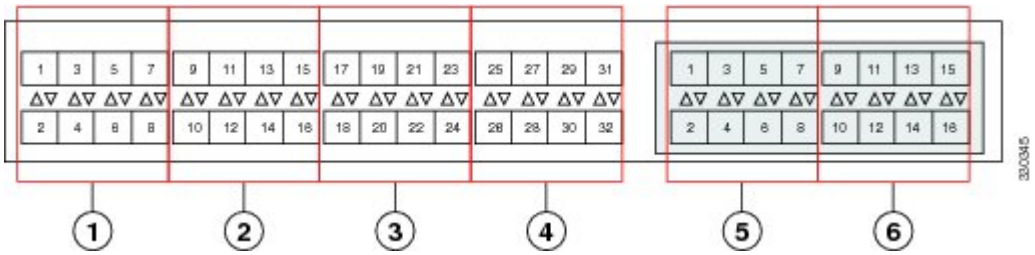

Caution Adding a second link to a fabric-port-channel port group is disruptive and will automatically increase the available amount of VIF namespace from 63 to 118. Adding further links is not disruptive and the VIF namespace stays at 118.

Caution Linking a chassis to two fabric-port-channel port groups does not affect the VIF namespace unless it is manually acknowledged. The VIF namespace is then automatically set to the smaller size fabric port-channel port group usage (either 63 or 118 VIFs) of the two groups.

For high availability cluster-mode applications, we strongly recommend symmetric cabling configurations. If the cabling is asymmetric, the maximum number of VIFs available is the smaller of the two cabling configurations.

For more information on the maximum number of VIFs for your Cisco UCS environment, see the Configuration Limits document for your hardware and software configuration.

Configuring a Fabric Port Channel

Procedure

Step 1 To include all links from the IOM to the fabric interconnect in a fabric port channel during chassis discovery, set the link grouping preference in the chassis discovery policy to port channel.

Step 2 To include links from individual chassis in a fabric port channel during chassis discovery, set the link grouping preference in the chassis connectivity policy to port channel.

Step 3 After chassis discovery, enable or disable additional fabric port channel member ports.

What to Do Next

To add or remove chassis links from a fabric port channel after making a change to the chassis discovery policy or the chassis connectivity policy, reacknowledge the chassis. Chassis reacknowledgement is not required to enable or disable chassis member ports from a fabric port channel

Viewing Fabric Port Channels

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **IO Modules**.

Step 3 Click the IOM for which you want to view the fabric port channels.

Step 4 In the **Work** pane, click the **Fabric Ports** tab.

Step 5 To view details of the fabric port channel, click the link in the **Port Channel** column.

Enabling or Disabling a Fabric Port Channel Member Port

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 Expand **LAN** > **Internal LAN** > *Fabric* > **Port Channels**.
Step 3 Expand the port channel for which you want to enable or disable a member port.
Step 4 Click the ethernet interface for the member port you want to enable or disable.
Step 5 In the **Work** pane, click the **General** tab.
Step 6 In the **Actions** area, click one of the following:

- **Enable Interface**
- **Disable Interface**

Step 7 If a confirmation dialog box displays, click **Yes**.

Configuring Server Ports with the Internal Fabric Manager

Internal Fabric Manager

The Internal Fabric Manager provides a single interface where you can configure server ports for a fabric interconnect in a Cisco UCS domain. The Internal Fabric Manager is accessible from the **General** tab for that fabric interconnect.

Some of the configuration that you can do in the Internal Fabric Manager can also be done in nodes on the **Equipment** tab, on the **LAN** tab, or in the LAN Uplinks Manager.

Launching the Internal Fabric Manager

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Fabric Interconnects** > *Fabric_Interconnect_Name*.
Step 3 Click **Fixed Module**.
Step 4 In the **Work** pane, click **Internal Fabric Manager** in the **Actions** area.
The Internal Fabric Manager opens in a separate window.

Configuring a Server Port with the Internal Fabric Manager

Procedure

Step 1 In the Internal Fabric Manager, click the down arrows to expand the **Unconfigured Ports** area.
Step 2 Right-click the port that you want to configure and choose **Configure as Server Port**.
Step 3 If a confirmation dialog box displays, click **Yes**.
Step 4 If you complete all of the tasks in the Internal Fabric Manager, click **OK**.

Unconfiguring a Server Port with the Internal Fabric Manager

Procedure

Step 1 In the Internal Fabric Manager, click the server port in the **Server Ports** table.
Step 2 Click **Unconfigure Port**.
Step 3 If a confirmation dialog box displays, click **Yes**.
Step 4 If you complete all of the tasks in the Internal Fabric Manager, click **OK**.

Enabling a Server Port with the Internal Fabric Manager

Procedure

Step 1 In the Internal Fabric Manager, click the server port in the **Server Ports** table.
Step 2 Click **Enable Port**.
Step 3 If a confirmation dialog box displays, click **Yes**.
Step 4 If you complete all of the tasks in the Internal Fabric Manager, click **OK**.

Disabling a Server Port with the Internal Fabric Manager

Procedure

Step 1 In the Internal Fabric Manager, click the server port in the **Server Ports** table.

Step 2 Click **Disable Port**.

Step 3 If a confirmation dialog box displays, click **Yes**.

Step 4 If you complete all of the tasks in the Internal Fabric Manager, click **OK**.

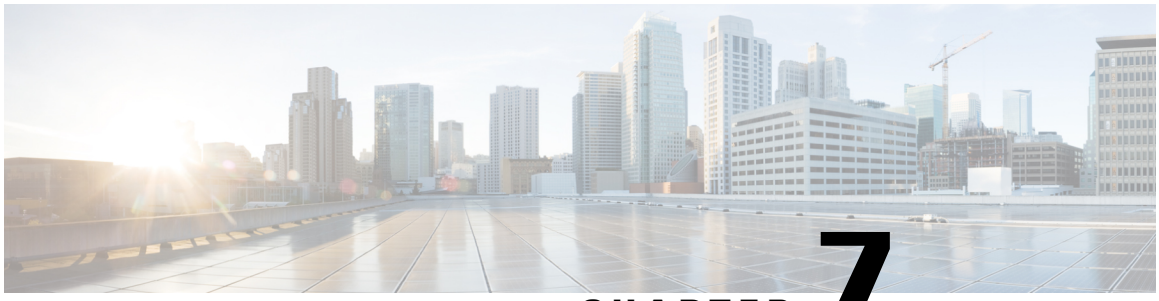

CHAPTER 7

Configuring Communication Services

This chapter includes the following sections:

Communication Services

You can use the communication services defined below to interface third-party applications with Cisco UCS.

Cisco UCS Manager supports IPv4 and IPv6 address access for the following services:

- CIM XML
- HTTP
- HTTPS
- SNMP
- SSH
- Telnet

Cisco UCS Manager supports out-of-band IPv4 address access to the **Cisco UCS KVM Direct** launch page from a web browser. To provide this access, you must enable the following service:

- CIMC Web Service

Communication Service	Description
CIM XML	The Common Information Model (CIM) XML) service is disabled by default and is only available in read-only mode. The default port is 5988. The CIM XML is a standards-based protocol for exchanging CIM information that the Distributed Management Task Force defines.
CIMC Web Service	This service is disabled by default. When this service is enabled, users can directly access a server CIMC using one of the out-of-band management IP addresses assigned directly to the server, or associated with the server through a service profile. **Note** CIMC Web Service can only be enabled or disabled globally. You cannot configure KVM direct access for individual CIMC IP addresses.
HTTP	By default, HTTP is enabled on port 80. You can run the Cisco UCS Manager GUI in an HTTP or HTTPS browser. If you select HTTP, all data is exchanged in clear text mode. For a secure browser session, we recommend that you enable HTTPS and disable HTTP. By default, Cisco UCS implements a browser redirects to an HTTPS equivalent and recommends that you do not change this behavior. **Note** If you are upgrading to Cisco UCS, version 1.4(1), the browser redirect to a secure browser does not occur by default. To redirect the HTTP browser to an HTTPS equivalent, enable the **Redirect HTTP to HTTPS** in Cisco UCS Manager.
HTTPS	By default, HTTPS is enabled on port. With HTTPS, all data is exchanged in encrypted mode through a secure server. For a secure browser session, We recommend that you only use HTTPS and either disable or redirect HTTP communications.
SMASH CLP	This service is enabled for read-only access and supports a limited subset of the protocols, such as the **show** command. You cannot disable it. This shell service is one of the standards that the Distributed Management Task Force defines.
SNMP	By default, this service is disabled. If enabled, the default port is 161. You must configure the community and at least one SNMP trap. Enable this service only if your system includes integration with an SNMP server.
SSH	This service is enabled on port 22. You cannot disable it, and you cannot change the default port. This service provides access to the Cisco UCS Manager CLI.
Telnet	By default, this service is disabled. This service provides access to the Cisco UCS Manager CLI.

Configuring CIM-XML

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **Communication Management** > **Communication Services**.

Step 3 Select the **Communication Services** tab.

Step 4 In the **CIM-XML** area, click the **Enabled** radio button.
The **CIM-XML** area expands to display the default **Port** number, 5988. You cannot change this port number.

Step 5 Click **Save Changes**.

Configuring HTTP

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **Communication Management** > **Communication Services**.

Step 3 Click the **Communication Services** tab.

Step 4 In the **HTTP** area, click the **Enabled** radio button.
The **HTTP** area expands to display the available configuration options.

Step 5 (Optional) In the **Port** field, change the default port that Cisco UCS Manager GUI uses for HTTP.
The default port is 80.

Step 6 (Optional) In the **Redirect HTTP to HTTPS** field, click the **Enabled** radio button.
You must also configure and enable HTTPS to enable redirection of HTTP logins to the HTTPS login. Once enabled, you cannot disable the redirection until you have disabled HTTPS.

Note If you redirect HTTP to HTTPS, you cannot use HTTP to access Cisco UCS Manager GUI. Redirection disables HTTP as it automatically redirects to HTTPS.

Step 7 Click **Save Changes**.

Configuring HTTPS

Certificates, Key Rings, and Trusted Points

HTTPS uses components of the Public Key Infrastructure (PKI) to establish secure communications between two devices, such as a client's browser and Cisco UCS Manager.

Encryption Keys and Key Rings

Each PKI device holds a pair of asymmetric Rivest-Shamir-Adleman (RSA) encryption keys, one kept private and one made public, stored in an internal key ring. A message encrypted with either key can be decrypted with the other key. To send an encrypted message, the sender encrypts the message with the receiver's public key, and the receiver decrypts the message using its own private key. A sender can also prove its ownership of a public key by encrypting (also called 'signing') a known message with its own private key. If a receiver can successfully decrypt the message using the public key in question, the sender's possession of the corresponding private key is proven. Encryption keys can vary in length, with typical lengths from 512 bits to 2048 bits. In general, a longer key is more secure than a shorter key. Cisco UCS Manager provides a default key ring with an initial 1024-bit key pair, and allows you to create additional key rings.

The default key ring certificate must be manually regenerated if the cluster name changes or the certificate expires.

This operation is only available in the UCS Manager CLI.

Certificates

To prepare for secure communications, two devices first exchange their digital certificates. A certificate is a file containing a device's public key along with signed information about the device's identity. To merely support encrypted communications, a device can generate its own key pair and its own self-signed certificate. When a remote user connects to a device that presents a self-signed certificate, the user has no easy method to verify the identity of the device, and the user's browser will initially display an authentication warning. By default, Cisco UCS Manager contains a built-in self-signed certificate containing the public key from the default key ring.

Trusted Points

To provide stronger authentication for Cisco UCS Manager, you can obtain and install a third-party certificate from a trusted source, or trusted point, that affirms the identity of your device. The third-party certificate is signed by the issuing trusted point, which can be a root certificate authority (CA) or an intermediate CA or trust anchor that is part of a trust chain that leads to a root CA. To obtain a new certificate, you must generate a certificate request through Cisco UCS Manager and submit the request to a trusted point.

Important The certificate must be in Base64 encoded X.509 (CER) format.

Creating a Key Ring

Cisco UCS Manager supports a maximum of 8 key rings, including the default key ring.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **Key Management**.

Step 3 Right-click **Key Management** and choose **Create Key Ring**.

Step 4 In the **Create Key Ring** dialog box, do the following:

a) In the **Name** field, enter a unique name for the key ring.

b) In the **Modulus** field, select one of the following radio buttons to specify the SSL key length in bits:

- **Mod512**
- **Mod1024**
- **Mod1536**
- **Mod2048**
- **Mod2560**
- **Mod3072**
- **Mod3584**
- **Mod4096**

c) Click **OK**.

What to Do Next

Create a certificate request for this key ring.

Creating a Certificate Request for a Key Ring

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **Key Management**.

Step 3 Click the key ring for which you want to create a certificate request.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **General** tab, click **Create Certificate Request**.

Step 6 In the **Create Certificate Request** dialog box, complete the following fields:

Name	Description
DNS field	The domain name assigned to the network that is common to all host names.

Name	Description
Locality field	The city or town in which the company requesting the certificate is headquartered. Enter up to 64 characters. You can use any letters, numbers, or spaces, as well as the following special characters: , (comma), . (period), @ (at sign), ^ (carat), ((open parenthesis),) (close parenthesis), - (dash), _ (underscore), + (plus sign), : (colon), / (forward slash).
State field	The state or province in which the company requesting the certificate is headquartered. Enter up to 64 characters. You can use any letters, numbers, or spaces, as well as the following special characters: , (comma), . (period), @ (at sign), ^ (carat), ((open parenthesis),) (close parenthesis), - (dash), _ (underscore), + (plus sign), : (colon), / (forward slash).
Country field	The country code corresponding to the country in which the company resides. Enter two alphabetic characters.
Organization Name field	The organization requesting the certificate. Enter up to 64 characters. You can use any letters, numbers, or spaces, as well as the following special characters: , (comma), . (period), @ (at sign), ^ (carat), ((open parenthesis),) (close parenthesis), - (dash), _ (underscore), + (plus sign), : (colon), / (forward slash).
Organization Unit Name field	The organizational unit. Enter up to 64 characters. You can use any letters, numbers, or spaces, as well as the following special characters: , (comma), . (period), @ (at sign), ^ (carat), ((open parenthesis),) (close parenthesis), - (dash), _ (underscore), + (plus sign), : (colon), / (forward slash).
Email field	The email address associated with the request.
Password field	An optional password for this request.
Confirm Password field	If you specified a password, enter it again for confirmation.
Subject field	The fully qualified domain name of the fabric interconnect.

Step 7 To assign IP addresses, click the **IPv4** or **IPv6** tab. The choice you make depends upon how the fabric interconnects were configured when you set up Cisco UCS Manager.

- Click the IPv4 tab, and complete the following fields:

Name	Description
IP Address field	The IPv4 address of the Cisco UCS domain.
FI-A IP field	The IPv4 address of fabric interconnect A.
FI-B IP field	The IPv4 address of fabric interconnect B.

- Click the IPv6 tab, and complete the following fields:

Name	Description
IP Address field	The IPv6 address of the Cisco UCS domain.
FI-A IP field	The IPv6 address of fabric interconnect A.
FI-B IP field	The IPv6 address of fabric interconnect B.

Step 8 Click **OK**.

Step 9 Copy the text of the certificate request from the **Request** field and save in a file.

Step 10 Send the file with the certificate request to the trust anchor or certificate authority.

What to Do Next

Create a trusted point and set the certificate chain for the certificate of trust received from the trust anchor.

Creating a Trusted Point

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **Key Management**.

Step 3 Right-click **Key Management** and choose **Create Trusted Point**.

Step 4 In the **Create Trusted Point** dialog box, complete the following fields:

Name	Description
Name field	The name of the trusted point. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Name	Description
Certificate Chain field	The certificate information for this trusted point. **Important** The certificate must be in Base64 encoded X.509 (CER) format. For windows 2012 server, using RSASSA-PSS returns the following error occurs: Trustpoint's cert-chain is invalid, reason: unknown. UCS Manager does not support this algorithm.

Step 5 Click **OK**.

What to Do Next

When you receive the certificate from the trust anchor or certificate authority, import it into the key ring.

Importing a Certificate into a Key Ring

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **Key Management**.

Step 3 Click the key ring into which you want to import the certificate.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Certificate** area, complete the following fields:

a) From the **Trusted Point** drop-down list, select the trusted point for the trust anchor that granted this certificate.

b) In the **Certificate** field, paste the text from the certificate you received from the trust anchor or certificate authority.

Important The certificate must be in Base64 encoded X.509 (CER) format.

Tip If the fields in an area do not display, click the **Expand** icon to the right of the heading.

Step 6 Click **Save Changes**.

What to Do Next

Configure your HTTPS service with the key ring.

Configuring HTTPS

Caution After you complete the HTTPS configuration, including changing the port and key ring for the HTTPS to use, all current HTTP and HTTPS sessions are closed without warning as soon as you save or commit the transaction.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **Communication Management** > **Communication Services**.

Step 3 Select the **Communication Services** tab.

Step 4 In the **HTTPS** area, click the **Enabled** radio button.
The **HTTPS** area expands to display the available configuration options.

Step 5 Complete the following fields:

Name	Description
Admin State field	This can be one of the following: • **Enabled** • **Disabled** If **Admin State** is enabled, Cisco UCS Manager GUI displays the remaining fields in this section.
Port field	The port to use for HTTPS connections. Specify an integer between 1 and 65535. By default, HTTPS is enabled on port.
Operational Port field	The port Cisco UCS Manager requires for system-level HTTPS communication. You cannot change this port.
Key Ring drop-down list	The key ring for HTTPS connections.
Cipher Suite Mode field	The level of Cipher Suite security used by the Cisco UCS domain. This can be one of the following: • **High Strength** • **Medium Strength** • **Low Strength** • **Custom—**Allows you to specify a user-defined Cipher Suite specification string.

Name	Description
Cipher Suite field	If you select **Custom** in the **Cipher Suite Mode** field, specify the user-defined Cipher Suite specification string in this field. The Cipher Suite specification string can contain up to 256 characters and must conform to the OpenSSL Cipher Suite specifications. You cannot use any spaces or special characters except ! (exclamation point), + (plus sign), - (hyphen), and : (colon). For details, see http://httpd.apache.org/docs/2.0/mod/mod_ssl.html#sslciphersuite. For example, the medium strength specification string Cisco UCS Manager uses as the default is: ALL:!ADH:!EXPORT56:!LOW:RC4+RSA:+HIGH:+MEDIUM:+EXP:+eNULL

Step 6 Click **Save Changes**.

Deleting a Key Ring

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **Key Management**.
Step 3 Right-click the key ring you want to delete and choose **Delete**.
Step 4 If a confirmation dialog box displays, click **Yes**.

Deleting a Trusted Point

Before You Begin

Ensure that the trusted point is not used by a key ring.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **Key Management**.
Step 3 Right-click the trusted point you want to delete and choose **Delete**.
Step 4 If a confirmation dialog box displays, click **Yes**.
Step 5 Click **OK**.

Enabling SNMP

SNMP Overview

The Simple Network Management Protocol (SNMP) is an application-layer protocol that provides a message format for communication between SNMP managers and agents. SNMP provides a standardized framework and a common language for monitoring and managing devices in a network.

SNMP Functional Overview

The SNMP framework consists of three parts:

- An SNMP manager—The system used to control and monitor the activities of network devices using SNMP.
- An SNMP agent—The software component within Cisco UCS, the managed device that maintains the data for Cisco UCS, and reports the data as needed to the SNMP manager. Cisco UCS includes the agent and a collection of MIBs. To enable the SNMP agent and create the relationship between the manager and agent, enable and configure SNMP in Cisco UCS Manager.
- A managed information base (MIB)—The collection of managed objects on the SNMP agent. Cisco UCS release 1.4(1) and higher supports a larger number of MIBs than earlier releases.

Cisco UCS supports SNMPv1, SNMPv2c and SNMPv3. Both SNMPv1 and SNMPv2c use a community-based form of security. SNMP is defined in the following:

- RFC 3410 (http://tools.ietf.org/html/rfc3410)
- RFC 3411 (http://tools.ietf.org/html/rfc3411)
- RFC 3412 (http://tools.ietf.org/html/rfc3412)
- RFC 3413 (http://tools.ietf.org/html/rfc3413)
- RFC 3414 (http://tools.ietf.org/html/rfc3414)
- RFC 3415 (http://tools.ietf.org/html/rfc3415)
- RFC 3416 (http://tools.ietf.org/html/rfc3416)
- RFC 3417 (http://tools.ietf.org/html/rfc3417)
- RFC 3418 (http://tools.ietf.org/html/rfc3418)
- RFC 3584 (http://tools.ietf.org/html/rfc3584)

SNMP Notifications

A key feature of SNMP is the ability to generate notifications from an SNMP agent. These notifications do not require that requests be sent from the SNMP manager. Notifications can indicate improper user

authentication, restarts, the closing of a connection, loss of connection to a neighbor router, or other significant events.

Cisco UCS Manager generates SNMP notifications as either traps or informs. Traps are less reliable than informs because the SNMP manager does not send any acknowledgment when it receives a trap, and Cisco UCS Manager cannot determine if the trap was received. An SNMP manager that receives an inform request acknowledges the message with an SNMP response Protocol Data Unit (PDU). If the Cisco UCS Manager does not receive the PDU, it can send the inform request again.

SNMP Security Levels and Privileges

SNMPv1, SNMPv2c, and SNMPv3 each represent a different security model. The security model combines with the selected security level to determine the security mechanism applied when the SNMP message is processed.

The security level determines the privileges required to view the message associated with an SNMP trap. The privilege level determines whether the message requires protection from disclosure or whether the message is authenticated. The supported security level depends on which security model is implemented. SNMP security levels support one or more of the following privileges:

- noAuthNoPriv—No authentication or encryption
- authNoPriv—Authentication but no encryption
- authPriv—Authentication and encryption

SNMPv3 provides for both security models and security levels. A security model is an authentication strategy that is set up for a user and the role in which the user resides. A security level is the permitted level of security within a security model. A combination of a security model and a security level determines which security mechanism is employed when handling an SNMP packet.

Supported Combinations of SNMP Security Models and Levels

The following table identifies the combinations of security models and levels.

Table 10: SNMP Security Models and Levels

Model	Level	Authentication	Encryption	What Happens
v1	noAuthNoPriv	Community string	No	Uses a community string match for authentication.
v2c	noAuthNoPriv	Community string	No	Uses a community string match for authentication.
v3	noAuthNoPriv	Username	No	Uses a username match for authentication.

Model	Level	Authentication	Encryption	What Happens
v3	authNoPriv	HMAC-MD5 or HMAC-SHA	No	Provides authentication based on the Hash-Based Message Authentication Code (HMAC) Message Digest 5 (MD5) algorithm or the HMAC Secure Hash Algorithm (SHA).
v3	authPriv	HMAC-MD5 or HMAC-SHA	DES	Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms. Provides Data Encryption Standard (DES) 56-bit encryption in addition to authentication based on the Cipher Block Chaining (CBC) DES (DES-56) standard.

SNMPv3 Security Features

SNMPv3 provides secure access to devices through a combination of authenticating and encrypting frames over the network. SNMPv3 authorizes only configured users to perform management operations and encrypts SNMP messages. The SNMPv3 User-Based Security Model (USM) refers to SNMP message-level security and offers the following services:

- Message integrity—Ensures that messages are not altered or destroyed in an unauthorized manner, and that data sequences are not altered beyond what can occur non-maliciously.
- Message origin authentication—Ensures that the identity of a message originator is verifiable.
- Message confidentiality and encryption—Ensures that information is not made available or disclosed to unauthorized individuals, entities, or processes.

SNMP Support in Cisco UCS

Cisco UCS provides the following support for SNMP:

Support for MIBs

Cisco UCS supports read-only access to MIBs.

For information about the specific MIBs available for Cisco UCS and where you can obtain them, see the http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/mib/b-series/b_UCS_MIBRef.html for B-series servers, and http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/mib/c-series/b_UCS_Standalone_C-Series_MIBRef.html C-series servers.

Authentication Protocols for SNMPv3 Users

Cisco UCS supports the following authentication protocols for SNMPv3 users:

- HMAC-MD5-96 (MD5)
- HMAC-SHA-96 (SHA)

AES Privacy Protocol for SNMPv3 Users

Cisco UCS uses Advanced Encryption Standard (AES) as one of the privacy protocols for SNMPv3 message encryption and conforms with RFC 3826.

The privacy password, or priv option, offers a choice of DES or 128-bit AES encryption for SNMP security encryption. If you enable AES-128 configuration and include a privacy password for an SNMPv3 user, Cisco UCS Manager uses the privacy password to generate a 128-bit AES key. The AES privacy password can have a minimum of eight characters. If the passphrases are specified in clear text, you can specify a maximum of 64 characters.

Enabling SNMP and Configuring SNMP Properties

SNMP messages from a Cisco UCS domain display the fabric interconnect name rather than the system name.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **Communication Management** > **Communication Services**.

Step 3 Select the **Communication Services** tab.

Step 4 In the **SNMP** area, complete the following fields:

Name	Description
Admin State field	This can be one of the following: • **Enabled** • **Disabled** Enable this service only if your system includes integration with an SNMP server. If **Admin State** is enabled, Cisco UCS Manager GUI displays the remaining fields in this section.

Step 5 Click **Save Changes**.

What to Do Next

Create SNMP traps and users.

Creating an SNMP Trap

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **Communication Management** > **Communication Services**.

Step 3 Select the **Communication Services** tab.

Step 4 In the **SNMP Traps** area, click +.

Step 5 In the **Create SNMP Trap** dialog box, complete the following fields:

Name	Description
Hostname (or IP Address) field	The host name or IP address of the SNMP host to which Cisco UCS Manager should send the trap. You can use an IPv4 address, or an IPv6 address for the SNMP host. The host name can also be a fully qualified domain name of an IPv4 address.
Community/Username field	The SNMP v1 or v2c community name or the SNMP v3 username Cisco UCS Manager includes when it sends the trap to the SNMP host. This must be the same as the community or username that is configured for the SNMP service. Enter an alphanumeric string between 1 and 32 characters. Do not use @ (at sign), \ (backslash), " (double quote), ? (question mark) or an empty space.
Port field	The port on which Cisco UCS Manager communicates with the SNMP host for the trap. Enter an integer between 1 and 65535. The default port is 162.
Version field	The SNMP version and model used for the trap. This can be one of the following: • **V1** • **V2c** • **V3**

Name	Description
Type field	The type of trap to send. This can be one of the following: • **Traps**, if you select **V2c** or **V3** for the version. • **Informs**, if you select V2c for the version. **Note** An inform notification can be send only if you select v2c for the version.
v3 Privilege field	If you select **V3** for the version, the privilege associated with the trap. This can be one of the following: • **Auth—**Authentication but no encryption • **Noauth—**No authentication or encryption • **Priv—**Authentication and encryption

Step 6 Click **OK**.

Step 7 Click **Save Changes**.

Deleting an SNMP Trap

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **Communication Management** > **Communication Services**.

Step 3 Select the **Communication Services** tab.

Step 4 In the **SNMP Traps** area, click the row in the table that corresponds to the user you want to delete.

Step 5 Click the **Delete** icon to the right of the table.

Step 6 If a confirmation dialog box displays, click **Yes**.

Step 7 Click **Save Changes**.

Creating an SNMPv3 user

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **Communication Management** > **Communication Services**.

Step 3 Select the **Communication Services** tab.

Step 4 In the **SNMP Users** area, click +.

Step 5 In the **Create SNMP User** dialog box, complete the following fields:

Name	Description
Name field	The username assigned to the SNMP user. Enter up to 32 letters or numbers. The name must begin with a letter and you can also specify _ (underscore), . (period), @ (at sign), and - (hyphen). **Note** You cannot create an SNMP username that is identical to a locally authenticated username.
Auth Type field	The authorization type. This can be one of the following: • **MD5** • **SHA**
Use AES-128 check box	If checked, this user uses AES-128 encryption.
Password field	The password for this user.
Confirm Password field	The password again for confirmation purposes.
Privacy Password field	The privacy password for this user.
Confirm Privacy Password field	The privacy password again for confirmation purposes.

Step 6 Click **OK**.

Step 7 Click **Save Changes**.

Deleting an SNMPv3 User

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **Communication Management** > **Communication Services**.
Step 3 Select the **Communication Services** tab.
Step 4 In the **SNMP Users** area, click the row in the table that corresponds to the user you want to delete.
Step 5 Click the **Delete** icon to the right of the table.
Step 6 If a confirmation dialog box displays, click **Yes**.
Step 7 Click **Save Changes**.

Enabling Telnet

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **Communication Management** > **Communication Services**.
Step 3 Click the **Communication Services** tab.
Step 4 In the **Telnet** area, click the **Enabled** radio button.
Step 5 Click **Save Changes**.

Enabling the CIMC Web Service

The CIMC web service is enabled by default. Follow the steps below to enable the service if you have disabled it.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **Communication Management** > **Communication Services**.
Step 3 Select the **Communication Services** tab.
Step 4 In the **CIMC Web Service** area, click the **Enabled** radio button.
Step 5 Click **Save Changes**.

Disabling Communication Services

Note We recommend that you disable all communication services that are not required to interface with other network applications.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **Communication Management** > **Communication Services**.

Step 3 On the **Communication Services** tab, click the **disable** radio button for each service that you want to disable.

Step 4 Click **Save Changes**.

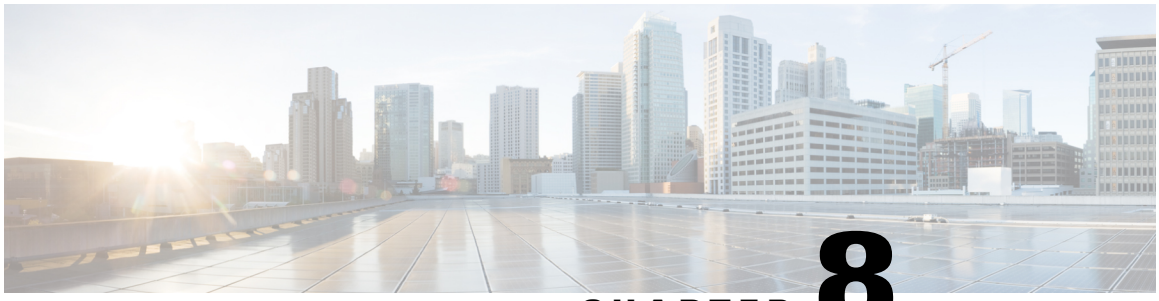

CHAPTER 8

Configuring Authentication

This chapter includes the following sections:

Authentication Services

Cisco UCS supports the following two methods to authenticate user logins:

- Local user authentication - uses user accounts that exist locally in the Cisco UCS Manager
- Remote user authentication - uses one of the following protocols:
  - LDAP
  - RADIUS
  - TACACS+

Guidelines and Recommendations for Remote Authentication Providers

If a system is configured for one of the supported remote authentication services, you must create a provider for that service to ensure that Cisco UCS Manager can communicate with the system. The following guidelines impact user authorization:

User Accounts in Remote Authentication Services

User accounts can exist locally in Cisco UCS Manager or in the remote authentication server.

You can view the temporary sessions for users who log in through remote authentication services from the Cisco UCS Manager GUI and from the Cisco UCS Manager CLI.

User Roles in Remote Authentication Services

If you create user accounts in the remote authentication server, you must ensure that the accounts include the roles those users require for working in Cisco UCS Manager and that the names of those roles match the names used in Cisco UCS Manager. Based on the role policy, a user might not be allowed to log in, or is granted only read-only privileges.

User Attributes in Remote Authentication Providers

For RADIUS and TACACS+ configurations, you must configure a user attribute for Cisco UCS in each remote authentication provider through which users log in to Cisco UCS Manager. This user attribute holds the roles and locales assigned to each user.

Note This step is not required for LDAP configurations that use the LDAP Group Mapping to assign roles and locales.

When a user logs in, Cisco UCS Manager does the following:

1. Queries the remote authentication service.
2. Validates the user.
3. If the user is validated, checks for the roles and locales assigned to that user.

The following table contains a comparison of the user attribute requirements for the remote authentication providers supported by Cisco UCS.

Table 11: Comparison of User Attributes by Remote Authentication Provider

Authentication Provider	Custom Attribute	Schema Extension	Attribute ID Requirements
LDAP	Not required if group mapping is used Optional if group mapping is not used	Optional. You can choose to do one of the following: • Do not extend the LDAP schema and configure an existing, unused attribute that meets the requirements. • Extend the LDAP schema and create a custom attribute with a unique name, such as CiscoAVPair.	The Cisco LDAP implementation requires a unicode type attribute. If you choose to create the CiscoAVPair custom attribute, use the following attribute ID: 1.3.6.1.4.1.9.287247.1 A sample OID is provided in the following section.
RADIUS	Optional	Optional. You can choose to do one of the following: • Do not extend the RADIUS schema and use an existing unused attribute that meets the requirements. • Extend the RADIUS schema and create a custom attribute with a unique name, such as cisco-avpair.	The vendor ID for the Cisco RADIUS implementation is 009 and the vendor ID for the attribute is 001. The following syntax example shows how to specify multiples user roles and locales if you choose to create the cisco-avpair attribute: `shell:roles="admin,aaa" shell:locales="L1,abc"`. Use a comma "," as the delimiter to separate multiple values.
TACACS+	Required	Required. You must extend the schema and create a custom attribute with the name cisco-av-pair.	The cisco-av-pair name is the string that provides the attribute ID for the TACACS+ provider. The following syntax example shows how to specify multiples user roles and locales when you create the cisco-av-pair attribute: `cisco-av-pair=shell:roles="admin aaa" shell:locales*"L1 abc"`. Using an asterisk (*) in the cisco-av-pair attribute syntax flags the locale as optional, preventing authentication failures for other Cisco devices that use the same authorization profile. Use a space as the delimiter to separate multiple values.

Sample OID for LDAP User Attribute

The following is a sample OID for a custom CiscoAVPair attribute:

```
CN=CiscoAVPair,CN=Schema,
CN=Configuration,CN=X
objectClass: top
objectClass: attributeSchema
cn: CiscoAVPair
distinguishedName: CN=CiscoAVPair,CN=Schema,CN=Configuration,CN=X
instanceType: 0x4
uSNCreated: 26318654
attributeID: 1.3.6.1.4.1.9.287247.1
attributeSyntax: 2.5.5.12
isSingleValued: TRUE
showInAdvancedViewOnly: TRUE
adminDisplayName: CiscoAVPair
adminDescription: UCS User Authorization Field
oMSyntax: 64
lDAPDisplayName: CiscoAVPair
name: CiscoAVPair
objectCategory: CN=Attribute-Schema,CN=Schema,CN=Configuration,CN=X
```

Two-Factor Authentication

Cisco UCS Manager uses two-factor authentication for remote user logins, which adds a level of security to account logins. Two-factor authentication login requires a username, a token, and a password combination in the password field. You can provide a PIN, a certificate, or a token.

Two-factor authentication uses authentication applications that maintain token servers to generate one-time tokens for users during the login process and store passwords in the AAA server. Requests are sent to the token server to retrieve a vendor-specific attribute. Cisco UCS Manager expects the token server to integrate with the AAA server, therefore it forwards the request to the AAA server. The password and token are validated at the same time by the AAA server. Users must enter the token and password sequence in the same order as it is configured in the AAA server.

Two-factor authentication is supported by associating RADIUS or TACACS+ provider groups with designated authentication domains and enabling two-factor authentication for those domains. Two-factor authentication does not support IPM and is not supported when the authentication realm is set to LDAP, local, or none.

Web Session Refresh and Web Session Timeout Period

The **Web Session Refresh Period** is the maximum amount of time allowed between refresh requests for a Cisco UCS Manager GUI web session. The **Web Session Timeout** is the maximum amount of time that can elapse after the last refresh request before a Cisco UCS Manager GUI web session becomes inactive.

You can increase the **Web Session Refresh Period** to a value greater than 60 seconds up 172800 seconds to avoid frequent session timeouts that requires regenerating and re-entering a token and password multiple times. The default value is 7200 seconds when two-factor authentication is enabled, and is 600 seconds when two-factor authentication is not enabled.

You can specify a value between 300 and 172800 for the **Web Session Timeout Period**. The default is 8000 seconds when two-factor authentication is enabled, and 7200 seconds when two-factor authentication is not enabled.

LDAP Group Rule

The LDAP group rule determines whether Cisco UCS should use LDAP groups when assigning user roles and locales to a remote user.

Nested LDAP Groups

You can add an LDAP group as a member of another group and nest groups to consolidate member accounts and to reduce the replication of traffic. Cisco UCS Manager release 2.1(2) and higher enables you to search LDAP groups that are nested within another group defined in an LDAP group map.

Note Nested LDAP search support is supported only for Microsoft Active Directory servers. The supported versions are Microsoft Windows 2003 SP3, Microsoft Windows 2008 R2, and Microsoft Windows 2012.

By default, user rights are inherited when you nest an LDAP group within another group. For example, if you make Group_1 a member of Group_2, the users in Group_1 have the same permissions as the members of Group_2. You can then search users that are members of Group_1 by choosing only Group_2 in the LDAP group map, instead of having to search Group_1 and Group_2 separately.

You do not always need to create subgroups in a group map in Cisco UCS Manager.

Configuring LDAP Providers

Configuring Properties for LDAP Providers

The properties that you configure in this task are the default settings for all provider connections of this type defined in Cisco UCS Manager. If an individual provider includes a setting for any of these properties, Cisco UCS uses that setting and ignores the default setting.

Before You Begin

If you are using Active Directory as your LDAP server, create a user account in the Active Directory server to bind with Cisco UCS. Give this account a non-expiring password.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management** > **LDAP**.

Step 3 In the **Properties** area, complete all fields.

Name	Description
Timeout field	The length of time in seconds the system spends trying to contact the LDAP database before it times out. Enter an integer from 1 to 60 seconds. The default value is 30 seconds. This property is required.
Vendor field	This selection identifies the vendor that is providing the LDAP provider or server details. If the LDAP provider is Microsoft Active Directory, select **MS-AD**. If the LDAP provider is not Microsoft Active Directory, select **Open Ldap**. The default is **Open Ldap**. **Note** If the vendor selection is **MS-AD** and the ldap-group-rule is enabled and set for recursive search, then Cisco UCS Manager can search through nested LDAP groups. Nested LDAP search is supported only with Active Directory. The server versions supported are Windows 2003 Sp2, Windows 2008 R2, and Windows 2012.
Attribute field	An LDAP attribute that stores the values for the user roles and locales. This property is always a name-value pair. The system queries the user record for the value that matches this attribute name. If you do not want to extend your LDAP schema, you can configure an existing, unused LDAP attribute with the Cisco UCS roles and locales. Alternatively, you can create an attribute named CiscoAVPair in the remote authentication service with the following attribute ID: 1.3.6.1.4.1.9.287247.1

Name	Description
Base DN field	The specific distinguished name in the LDAP hierarchy where the server begins a search when a remote user logs in and the system attempts to obtain the user's DN based on their username. You can set the length of the base DN to a maximum of 255 characters minus the length of CN=username, where username identifies the remote user attempting to access Cisco UCS Manager using LDAP authentication. This property is required. If you do not specify a base DN on this tab then you must specify one on the **General** tab for every LDAP provider defined in this Cisco UCS domain.
Filter field	The LDAP search is restricted to those user names that match the defined filter. This property is required. If you do not specify a filter on this tab then you must specify one on the **General** tab for every LDAP provider defined in this Cisco UCS domain.

Note User login fails if the userDn for an LDAP user exceeds 255 characters.

Step 4 Click **Save Changes**.

What to Do Next

Create an LDAP provider.

Creating an LDAP Provider

Cisco UCS Manager supports a maximum of 16 LDAP providers.

Before You Begin

If you are using Active Directory as your LDAP server, create a user account in the Active Directory server to bind with Cisco UCS. Give this account a non-expiring password.

- In the LDAP server, perform one of the following configurations:
  - Configure LDAP groups. LDAP groups contain user role and locale information.
  - Configure users with the attribute that holds the user role and locale information for Cisco UCS Manager. You can choose whether to extend the LDAP schema for this attribute. If you do not want to extend the schema, use an existing LDAP attribute to hold the Cisco UCS user roles and locales. If you prefer to extend the schema, create a custom attribute, such as the CiscoAVPair attribute.

The Cisco LDAP implementation requires a unicode type attribute.

If you choose to create the CiscoAVPair custom attribute, use the following attribute ID: 1.3.6.1.4.1.9.287247.1

- For a cluster configuration, add the management port IPv4 or IPv6 addresses for both fabric interconnects. This configuration ensures that remote users can continue to log in if the first fabric interconnect fails and the system fails over to the second fabric interconnect. All login requests are sourced from these IP addresses, not the virtual IPv4 or IPv6 address used by Cisco UCS Manager.

- If you want to use secure communications, create a trusted point containing the certificate of the root certificate authority (CA) of the LDAP server in Cisco UCS Manager.
- If you need to change the LDAP providers or add or delete them, change the authentication realm for the domain to local, make the changes to the providers, then change the domain authentication realm back to LDAP.
- If you want to use the special characters listed in the following table for defining the attributes of an Active Directory bind distinguished name, you must replace the special character with an escape, by using a backslash (\) followed by the corresponding hexadecimal value of the character.

Special Character	Description	Hexadecimal Value
,	comma	0x2C
+	plus sign	0x2B
"	double quote	0x22
\	backslash	0x5C
<	left angle bracket	0x3C
>	right angle bracket	0x3E
;	semicolon	0x3B
LF	line feed	0x0A
CR	carriage return	0x0D
=	equals sign	0x3D
/	forwards slash	0x2F

https://msdn.microsoft.com/en-us/library/aa366101 provides more details on replacing special characters with its escape and hexadecimal equivalent.

Attention LDAP remote usernames that include special characters cannot log in to systems that are running versions 2.2(3a) and later. The user cannot log in because of the Nexus OS limitations where special characters, !,%,^, are not supported in the username.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management** > **LDAP**.

Step 3 In the **Work** pane, click the **General** tab.

Step 4 In the **Actions** area, click **Create LDAP Provider**.

Step 5 On the **Create LDAP Provider** page of the wizard, complete all fields with appropriate LDAP service information.

a) Complete the following fields with information about the LDAP service you want to use:

Name	Description
Hostname/FDQN (or IP Address) field	The hostname, or IPv4 or IPv6 address on which the LDAP provider resides. If SSL is enabled, this field must exactly match a Common Name (CN) in the security certificate of the LDAP database. **Note** If you use a hostname rather than an IPv4 or IPv6 address, you must configure a DNS server. If the Cisco UCS domain is not registered with Cisco UCS Central or DNS management is set to **local**, configure a DNS server in Cisco UCS Manager. If the Cisco UCS domain is registered with Cisco UCS Central and DNS management is set to **global**, configure a DNS server in Cisco UCS Central.
Order field	The order that the Cisco UCS uses this provider to authenticate users. Enter an integer between 1 and 16, or enter lowest-available or 0 (zero) if you want Cisco UCS to assign the next available order based on the other providers defined in this Cisco UCS domain.
Bind DN field	The distinguished name (DN) for an LDAP database account that has read and search permissions for all objects under the base DN. The maximum supported string length is 255 ASCII characters.
Base DN field	The specific distinguished name in the LDAP hierarchy where the server begins a search when a remote user logs in and the system attempts to obtain the user's DN based on their username. You can set the length of the base DN to a maximum of 255 characters minus the length of CN=username, where username identifies the remote user attempting to access Cisco UCS Manager using LDAP authentication. This value is required unless a default base DN has been set on the LDAP **General** tab.

Name	Description
Port field	The port through which Cisco UCS communicates with the LDAP database. The standard port number is 389.
Enable SSL check box	If checked, encryption is required for communications with the LDAP database. If unchecked, authentication information will be sent as clear text. LDAP uses STARTTLS. This allows encrypted communication using port 389. If checked, do not change the port to 636, leave it as 389. Cisco UCS negotiates a TLS session on port 636 for SSL, but initial connection starts unencrypted on 389.
Filter field	The LDAP search is restricted to those user names that match the defined filter. This value is required unless a default filter has been set on the LDAP **General** tab.
Attribute field	An LDAP attribute that stores the values for the user roles and locales. This property is always a name-value pair. The system queries the user record for the value that matches this attribute name. If you do not want to extend your LDAP schema, you can configure an existing, unused LDAP attribute with the Cisco UCS roles and locales. Alternatively, you can create an attribute named CiscoAVPair in the remote authentication service with the following attribute ID: 1.3.6.1.4.1.9.287247.1 This value is required unless a default attribute has been set on the LDAP **General** tab.
Password field	The password for the LDAP database account specified in the **Bind DN** field. You can enter any standard ASCII characters except for space, § (section sign), ? (question mark), or = (equal sign).
Confirm Password field	The LDAP database password repeated for confirmation purposes.
Timeout field	The length of time in seconds the system spends trying to contact the LDAP database before it times out. Enter an integer from 1 to 60 seconds, or enter 0 (zero) to use the global timeout value specified on the LDAP **General** tab. The default is 30 seconds.
Vendor radio button	The LDAP vendor that you want to use. This can be one of the following: • Open Ldap—The open source implementation of the LDAP protocol. • MS AD—Microsoft Active Directory.

b) Click **Next**.

Step 6 On the **LDAP Group Rule** page of the wizard, complete all fields with appropriate LDAP group rule information.

a) Complete the following fields:

Name	Description
Group Authorization field	Whether Cisco UCS also searches LDAP groups when authenticating and assigning user roles and locales to remote users. This can be one of the following: • **Disable—**Cisco UCS does not access any LDAP groups. • **Enable—**Cisco UCS searches all LDAP groups mapped in this Cisco UCS domain. If the remote user is found, Cisco UCS assigns the user roles and locales defined for that LDAP group in the associated LDAP group map. **Note** Role and locale assignment is cumulative. If a user is included in multiple groups, or has a role or locale specified in the LDAP attribute, Cisco UCS assigns that user all the roles and locales mapped to any of those groups or attributes.
Group Recursion field	Whether Cisco UCS searches both the mapped groups and their parent groups. This can be one of the following: • **Non Recursive—**Cisco UCS searches only the groups mapped in this Cisco UCS domain. If none of the groups containing the user explicitly set the user's authorization properties, Cisco UCS uses the default settings. • **Recursive—**Cisco UCS searches each mapped group and all its parent groups for the user's authorization properties. These properties are cumulative, so for each group Cisco UCS finds with explicit authorization property settings, it applies those settings to the current user. Otherwise it uses the default settings.
Target Attribute field	The attribute Cisco UCS uses to determine group membership in the LDAP database. The supported string length is 63 characters. The default string is memberOf.
Use Primary Group field	The attribute Cisco UCS uses to determine if the primary group can be configured as an LDAP group map for membership validation. With this option Cisco UCS Manager can download and verify the primary-group membership of the user.

b) Click **Finish**.

What to Do Next

For implementations involving a single LDAP database, select LDAP as the authentication service.

For implementations involving multiple LDAP databases, configure an LDAP provider group.

Changing the LDAP Group Rule for an LDAP Provider

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management** > **LDAP**.

Step 3 Expand **LDAP Providers** and choose the LDAP provider for which you want to change the group rule.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **LDAP Group Rules** area, complete the following fields:

Name	Description
Group Authorization field	Whether Cisco UCS also searches LDAP groups when authenticating and assigning user roles and locales to remote users. This can be one of the following: • **Disable—**Cisco UCS does not access any LDAP groups. • **Enable—**Cisco UCS searches all LDAP groups mapped in this Cisco UCS domain. If the remote user is found, Cisco UCS assigns the user roles and locales defined for that LDAP group in the associated LDAP group map. **Note** Role and locale assignment is cumulative. If a user is included in multiple groups, or has a role or locale specified in the LDAP attribute, Cisco UCS assigns that user all the roles and locales mapped to any of those groups or attributes.
Group Recursion field	Whether Cisco UCS searches both the mapped groups and their parent groups. This can be one of the following: • **Non Recursive—**Cisco UCS searches only the groups mapped in this Cisco UCS domain. If none of the groups containing the user explicitly set the user's authorization properties, Cisco UCS uses the default settings. • **Recursive—**Cisco UCS searches each mapped group and all its parent groups for the user's authorization properties. These properties are cumulative, so for each group Cisco UCS finds with explicit authorization property settings, it applies those settings to the current user. Otherwise it uses the default settings.

Name	Description
Target Attribute field	The attribute Cisco UCS uses to determine group membership in the LDAP database. The supported string length is 63 characters. The default string is memberOf.
Use Primary Group field	The attribute Cisco UCS uses to determine if the primary group can be configured as an LDAP group map for membership validation. With this option Cisco UCS Manager can download and verify the primary-group membership of the user.

Step 6 Click **Save Changes**.

Deleting an LDAP Provider

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **User Management** > **LDAP**.
Step 3 Expand **LDAP Providers**.
Step 4 Right-click the LDAP provider that you want to delete and choose **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

LDAP Group Mapping

LDAP group mapping eliminates having to define role or locale information in the LDAP user object. UCSM can use group membership information to assign a role or locale to an LDAP user during login for organizations using LDAP groups to restrict access to LDAP databases.

When a user logs in to Cisco UCS Manager, the LDAP group map pulls information about the user's role and locale. If the role and locale criteria match the information in the policy, access is granted. Cisco UCS Manager supports a maximum of 28, 128, or 160 LDAP group maps depending on the release version.

Note Cisco UCS Manager Release 3.1(1) supports a maximum of 128 LDAP group maps, and Release 3.1(2) and later releases support a maximum of 160 LDAP group maps.

The role and locale definitions that you configure locally in the Cisco UCS Manager do not update automatically based on changes to an LDAP directory. When deleting or renaming LDAP groups in an LDAP directory, you must also update the Cisco UCS Manager with the change.

You can configure an LDAP group map to include any of the following combinations of roles and locales:

- Roles only
- Locales only
- Both roles and locales

For example, consider an LDAP group representing a group of server administrators at a specific location. The LDAP group map might include user roles such as server profile and server equipment. To restrict access to server administrators at a specific location, you can set the locale to a particular site name.

Note Cisco UCS Manager includes out-of-the-box user roles, but does not include any locales. Mapping an LDAP provider group to a locale requires that you create a custom locale.

Creating an LDAP Group Map

Before You Begin

- Create an LDAP group in the LDAP server.
- Configure the distinguished name for the LDAP group in the LDAP server.
- Create locales in Cisco UCS Manager (optional).
- Create custom roles in Cisco UCS Manager (optional).

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management** > **LDAP**.

Step 3 Right-click **LDAP Group Maps** and choose **Create LDAP Group Map**.

Step 4 In the **Create LDAP Group Map** dialog box, specify all LDAP group map information, as appropriate.

Important The name that you specify in the **LDAP Group DN** field must match the name in the LDAP database.

Note If you use a special character in the **LDAP Group DN** field, you must prefix the special character with an escape character \ (single back slash).

What to Do Next

Set the LDAP group rule.

Deleting an LDAP Group Map

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **User Management** > **LDAP**.
Step 3 Expand **LDAP Group Maps**.
Step 4 Right-click the LDAP group map that you want to delete and choose **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Configuring RADIUS Providers

Configuring Properties for RADIUS Providers

The properties that you configure in this task are the default settings for all provider connections of this type defined in Cisco UCS Manager. If an individual provider includes a setting for any of these properties, Cisco UCS uses that setting and ignores the default setting.

Note RADIUS authentication uses Password Authentication Protocol (PAP).

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 In the **Admin** tab, expand **User Management** > **RADIUS**.
Step 3 In the **Properties** area, complete all fields.
Step 4 Click **Save Changes**.

What to Do Next

Create a RADIUS provider.

Creating a RADIUS Provider

Cisco UCS Manager supports a maximum of 16 RADIUS providers.

Before You Begin

Perform the following configuration in the RADIUS server:

- Configure users with the attribute that holds the user role and locale information for Cisco UCS Manager. You can choose whether to extend the RADIUS schema for this attribute. If you do not want to extend the schema, use an existing RADIUS attribute to hold the Cisco UCS user roles and locales. If you prefer to extend the schema, create a custom attribute, such as the cisco-avpair attribute.

  The vendor ID for the Cisco RADIUS implementation is 009 and the vendor ID for the attribute is 001.

  The following syntax example shows how to specify multiples user roles and locales if you choose to create the cisco-avpair attribute: `shell:roles="admin,aaa" shell:locales="L1,abc"`. Use a comma "," as the delimiter to separate multiple values.

- For a cluster configuration, add the management port IPv4 or IPv6 addresses for both fabric interconnects. This configuration ensures that remote users can continue to log in if the first fabric interconnect fails and the system fails over to the second fabric interconnect. All login requests are sourced from these IP addresses, not the virtual IP address used by Cisco UCS Manager.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management** > **RADIUS**.

Step 3 In the **Create RADIUS Provider** dialog box, specify all appropriate RADIUS service information.

Note If you use a hostname rather than an IPv4 or IPv6 address, you must ensure a DNS server is configured for the hostname.

Step 4 Click **Save Changes**.

What to Do Next

For implementations involving a single RADIUS database, select RADIUS as the primary authentication service.

For implementations involving multiple RADIUS databases, configure a RADIUS provider group.

Deleting a RADIUS Provider

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 In the **Admin** tab, expand **User Management** > **RADIUS**.

Step 3 Right-click the RADIUS provider that you want to delete and choose **Delete**.

Step 4 If a confirmation dialog box displays, click **Yes**.

Configuring TACACS+ Providers

Configuring Properties for TACACS+ Providers

The properties that you configure in this task are the default settings for all provider connections of this type defined in Cisco UCS Manager. If an individual provider includes a setting for any of these properties, Cisco UCS uses that setting and ignores the default setting.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 In the **Admin** tab, expand **User Management** > **TACACS+**.

Step 3 In the **Properties** area, complete the **Timeout** field.

Step 4 Click **Save Changes**.

What to Do Next

Create a TACACS+ provider.

Creating a TACACS+ Provider

Cisco UCS Manager supports a maximum of 16 TACACS+ providers.

Before You Begin

Perform the following configuration in the TACACS+ server:

- Create the cisco-av-pair attribute. You cannot use an existing TACACS+ attribute.

  The cisco-av-pair name is the string that provides the attribute ID for the TACACS+ provider.

  The following syntax example shows how to specify multiples user roles and locales when you create the cisco-av-pair attribute: `cisco-av-pair=shell:roles="admin aaa" shell:locales*"L1 abc".` Using an asterisk (*) in the cisco-av-pair attribute syntax flags the locale as optional, preventing authentication failures for other Cisco devices that use the same authorization profile. Use a space as the delimiter to separate multiple values.

- For a cluster configuration, add the management port IPv4 or IPv6 addresses for both fabric interconnects. This configuration ensures that remote users can continue to log in if the first fabric interconnect fails and the system fails over to the second fabric interconnect. All login requests are sourced from these IP addresses, not the virtual IP address used by Cisco UCS Manager.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management** > **TACACS+**.

Step 3 In the **Actions** area of the **General** tab, click **Create TACACS+ Provider**.

Step 4 In the **Create TACACS+ Provider** dialog box:

a) Complete all fields with TACACS+ service information, as appropriate.

Note If you use a hostname rather than an IPv4 or IPv6 address, you must ensure a DNS server is configured for the hostname.

b) Click **OK**.

Step 5 Click **Save Changes**.

What to Do Next

For implementations involving a single TACACS+ database, select TACACS+ as the primary authentication service.

For implementations involving multiple TACACS+ databases, configure a TACACS+ provider group.

Deleting a TACACS+ Provider

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 In the **Admin** tab, expand **User Management** > **TACACS+**.

Step 3 Right-click the TACACS+ provider that you want to delete and choose **Delete**.

Step 4 If a confirmation dialog box displays, click **Yes**.

Multiple Authentication Services Configuration

Multiple Authentication Services

You can configure Cisco UCS to use multiple authentication services by configuring the following features:

- Provider groups
- Authentication domains

Provider Groups

A provider group is a set of providers that the Cisco UCS accesses during the authentication process. All of the providers within a provider group are accessed in the order that the Cisco UCS provider uses to authenticate users. If all of the configured servers are unavailable or unreachable, Cisco UCS Manager automatically falls back to the local authentication method using the local username and password.

Cisco UCS Manager allows you to create a maximum of 16 provider groups, with a maximum of eight providers allowed per group.

Creating an LDAP Provider Group

Creating an LDAP provider group allows you to authenticate using multiple LDAP databases.

Before You Begin

Create one or more LDAP providers.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management** > **LDAP**.

Step 3 Right-click **LDAP Provider Groups** and choose **Create LDAP Provider Group**.

Note If you use a hostname rather than an IPv4 or IPv6 address, you must ensure a DNS server is configured for the hostname.

Step 4 In the **Create LDAP Provider Group** dialog box, specify all of the appropriate LDAP provider group information.

What to Do Next

Configure an authentication domain or select a default authentication service.

Deleting an LDAP Provider Group

Before You Begin

Remove the provider group from an authentication configuration.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **User Management** > **LDAP**.
Step 3 Expand **LDAP Provider Groups**.
Step 4 Right-click the LDAP provider group that you want to delete and choose **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Creating a RADIUS Provider Group

Creating a RADIUS provider group allows you to authenticate using multiple RADIUS databases.

Before You Begin

Create one or more RADIUS providers.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **User Management** > **RADIUS**.
Step 3 Right-click **RADIUS Provider Groups** and choose **Create RADIUS Provider Group**.
Step 4 In the **Create RADIUS Provider Group** dialog box, do the following:
a) In the **Name** field, enter a unique name for the group.
This name can be between 1 and 127 ASCII characters.
b) In the **RADIUS Providers** table, choose one or more providers to include in the group.
c) Click the **>>** button to add the providers to the **Included Providers** table.
You can use the **<<** button to remove providers from the group.
d) (Optional) Use the **Move Up** or **Move Down** arrows in the **Included Providers** list to change the order in which the RADIUS providers authenticate providers.
e) After you add all of the required providers to the provider group, click **OK**.

What to Do Next

Configure an authentication domain or select a default authentication service.

Deleting a RADIUS Provider Group

You cannot delete a provider group if another authentication configuration is using that provider group.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **User Management** > **RADIUS**.
Step 3 Expand **RADIUS Provider Groups**.
Step 4 Right-click the RADIUS provider group you want to delete and choose **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Creating a TACACS+ Provider Group

Creating a TACACS+ provider group allows you to authenticate using multiple TACACS+ databases.

Before You Begin

Create one or more TACACS+ providers.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **User Management** > **TACACS+**.
Step 3 Right-click **TACACS+ Provider Groups** and choose **Create TACACS+ Provider Group**.
Step 4 In the **Create TACACS+ Provider Group** dialog box, specify all TACACS+ provider group information, as appropriate.

Deleting a TACACS+ Provider Group

You cannot delete a provider group if another authentication configuration is using that provider group.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **User Management** > **TACACS+**.
Step 3 Expand **TACACS+ Provider Groups**.
Step 4 Right-click the TACACS+ provider group that you want to delete and choose **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Authentication Domains

The Cisco UCS Manager uses Authentication Domains to leverage multiple authentication systems. You can specify and configure each authentication domain during login; otherwise, Cisco UCS Manager uses the default authentication service configuration.

You can create up to eight authentication domains. Each authentication domain is associated with a provider group and a realm in the Cisco UCS Manager. The Cisco UCS Manager uses all servers within the realm if you do not specify a provider group.

Creating an Authentication Domain

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management** > **Authentication**.

Step 3 Right-click **Authentication Domains** and choose **Create a Domain**.

Step 4 In the **Create a Domain** dialog box, complete the following fields:

Name	Description
Name	The name of the domain. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved. **Note** For systems using the remote authentication protocol, the authentication domain name is considered part of the username and counts toward the 32-character limit for locally created usernames. Because Cisco UCS inserts 5 characters for formatting, authentication fails if the domain name and username combined characters total exceeds 27.

Name	Description
Web Session Refresh Period (sec)	When a web client connects to Cisco UCS Manager, the client must send refresh requests to Cisco UCS Manager to keep the web session active. This option specifies the maximum amount of time allowed between refresh requests for a user in this domain. If this time limit is exceeded, Cisco UCS Manager considers the web session inactive, but it does not terminate the session. Specify an integer between 60 and 172800. The default is 600 seconds when Two-Factor Authentication is not enabled and 7200 seconds when it is enabled. **Note** The number of seconds set for the **Web Session Refresh Period** must be less than the number of seconds set for the **Web Session Timeout**. Do not set the **Web Session Refresh Period** to the same value as the **Web Session Timeout**.
Web Session Timeout (sec)	The maximum amount of time that can elapse after the last refresh request before Cisco UCS Manager considers a web session as inactive. If this time limit is exceeded, Cisco UCS Manager automatically terminates the web session. Specify an integer between 300 and 172800. The default is 7200 seconds when Two-Factor Authentication is not enabled and 8000 seconds when it is enabled.
Realm	The authentication protocol to apply to users in this domain. This can be one of the following: • **Local—**The user account must be defined locally in this Cisco UCS domain. • **Radius—**The user must be defined on the RADIUS server specified for this Cisco UCS domain. • **Tacacs—**The user must be defined on the TACACS+ server specified for this Cisco UCS domain. • **Ldap—**The user must be defined on the LDAP server specified for this Cisco UCS domain.
Provider Group	The default provider group to use to authenticate users during remote login. **Note** The **Provider Group** drop-down list displays when you select Ldap Radius, or Tacacs as the method to authenticate users.

Name	Description
Two Factor Authentication	Two-Factor Authentication is available only when the **Realm** is set to **Radius** or **Tacacs**. When you select this check box, Cisco UCS Manager and the KVM launch manager require users whose accounts are authenticated by Radius or Tacacs servers to enter a token plus a password to log in. When 60 seconds remain for the **Web Session Refresh Period** to expire, users must generate a new token and enter the token plus their password to continue the session.

Step 5 Click **OK**.

Selecting a Primary Authentication Service

Selecting the Console Authentication Service

Before You Begin

If the system uses a remote authentication service, create a provider for that authentication service. If the system uses only local authentication through Cisco UCS, you do not need to create a provider first.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **User Management** > **Authentication**.
Step 3 Click **Native Authentication**.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Console Authentication** area, complete the following fields:

Name	Description
Realm field	The method by which a user logging into the console is authenticated. This can be one of the following: • **Local—**The user account must be defined locally in this Cisco UCS domain. • **Radius—**The user must be defined on the RADIUS server specified for this Cisco UCS domain. • **Tacacs—**The user must be defined on the TACACS+ server specified for this Cisco UCS domain. • **Ldap—**The user must be defined on the LDAP server specified for this Cisco UCS domain. • **None—**If the user account is local to this Cisco UCS domain, no password is required when the user logs into the console.
Provider Group drop-down list	The provider group to be used to authenticate a user logging into the console. **Note** The**Provider Group** drop-down list is displayed when you select Ldap, Radius, or Tacacs as the method by which a user is authenticated.
Two Factor Authentication	Two-factor authentication is available only when the **Realm** is set to **Radius** or **Tacacs**. When this checkbox is selected, the Console requires users whose accounts are authenticated by Radius or Tacacs servers to enter a token plus a password to log in.

Step 6 Click **Save Changes**.

Selecting the Default Authentication Service

Before You Begin

If the system uses a remote authentication service, create a provider for that authentication service. If the system uses only local authentication through Cisco UCS, you do not need to create a provider first.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management** > **Authentication**.

Step 3 Click **Native Authentication**.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Default Authentication** area, complete the following fields:

Name	Description
Realm drop-down list	The default method by which a user is authenticated during remote login. This can be one of the following: • **Local—**The user account must be defined locally in this Cisco UCS domain. • **Radius—**The user account must be defined on the RADIUS server specified for this Cisco UCS domain. • **Tacacs—**The user account must be defined on the TACACS+ server specified for this Cisco UCS domain. • **Ldap—**The user account must be defined on the LDAP server specified for this Cisco UCS domain. • **None—**If the user account is local to this Cisco UCS domain, no password is required when the user logs in remotely.
Provider Group drop-down list	The default provider group to be used to authenticate the user during remote login. **Note** The **Provider Group** drop-down is displayed when you select Ldap, Radius, or Tacacs as the method by which a user is authenticated.
Web Session Refresh Period (sec)	When a web client connects to Cisco UCS Manager, the client must send refresh requests to Cisco UCS Manager to keep the web session active. This option specifies the maximum amount of time allowed between refresh requests for a user in this domain. If this time limit is exceeded, Cisco UCS Manager considers the web session inactive, but it does not terminate the session. Specify an integer between 60 and 172800. The default is 600 seconds when Two-Factor Authentication is not enabled and 7200 seconds when it is enabled.

Name	Description
Web Session Timeout (sec)	The maximum amount of time that can elapse after the last refresh request before Cisco UCS Manager considers a web session as inactive. If this time limit is exceeded, Cisco UCS Manager automatically terminates the web session. Specify an integer between 300 and 172800. The default is 7200 seconds when Two-Factor Authentication is not enabled and 8000 seconds when it is enabled.
Two Factor Authentication checkbox	Two-Factor Authentication is available only when the **Realm** is set to **Radius** or **Tacacs**. When you select this check box, Cisco UCS Manager and the KVM launch manager require users whose accounts are authenticated by Radius or Tacacs servers to enter a token plus a password to log in. When 60 seconds remain for the **Web Session Refresh Period** to expire, users must generate a new token and enter the token plus their password to continue the session. **Note** After you enable two factor authentication and save the configuration, the default **Web Session Refresh Period (sec)** changes to 7200, and the default **Web Session Timeout (sec)** changes to 8000.

Step 6 Click **Save Changes**.

Role Policy for Remote Users

By default, if user roles are not configured in Cisco UCS Manager read-only access is granted to all users logging in to Cisco UCS Manager from a remote server using the LDAP, RADIUS, or TACACS protocols. For security reasons, it might be desirable to restrict access to those users matching an established user role in Cisco UCS Manager.

You can configure the role policy for remote users in the following ways:

assign-default-role

Does not restrict user access to Cisco UCS Manager based on user roles. Read-only access is granted to all users unless other user roles have been defined in Cisco UCS Manager.

This is the default behavior.

no-login

Restricts user access to Cisco UCS Manager based on user roles. If user roles have not been assigned for the remote authentication system, access is denied.

Configuring the Role Policy for Remote Users

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management** > **Authentication**.

Step 3 Click **Native Authentication**.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Role Policy for Remote Users** field, click one of the following radio buttons to determine what happens when a user attempts to log in and the remote authentication provider does not supply a user role with the authentication information:

- **No Login—**The user is not allowed to log in to the system, even if the username and password are correct.
- **Assign Default Role—**The user is allowed to log in with a read-only user role.

Step 6 Click **Save Changes**.

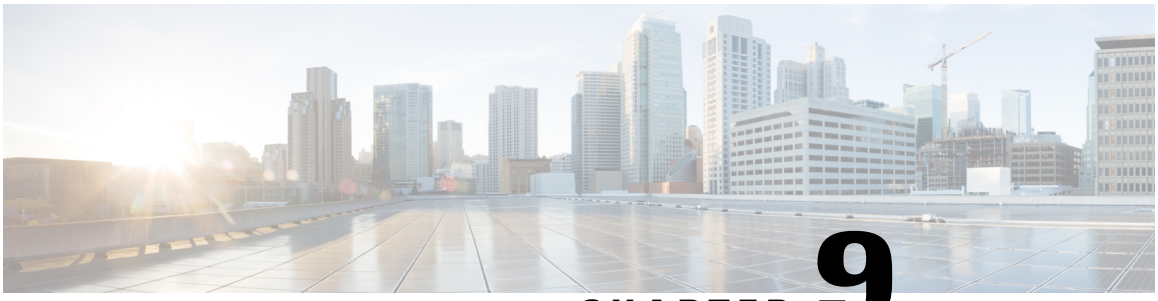

CHAPTER 9

Configuring Organizations

This chapter includes the following sections:

Organizations in a Multitenancy Environment

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

The root organization is always the top level organization.

Hierarchical Name Resolution in a Multi-Tenancy Environment

In a multi-tenant environment, Cisco UCS uses the hierarchy of an organization to resolve the names of policies and resource pools. When Cisco UCS Manager searches for details of a policy or a resource assigned to a pool, the following occurs:

1. Cisco UCS Manager checks for policies and pools with the specified name within the organization assigned to the service profile or policy.
2. If a policy is found or an available resource is inside a pool, Cisco UCS Manager uses that policy or resource. If the pool does not have any available resources at the local level, Cisco UCS Manager moves up in the hierarchy to the parent organization and searches for a pool with the same name. Cisco UCS Manager repeats this step until the search reaches the root organization.
3. If the search reaches the root organization and has not found an available resource or policy, Cisco UCS Manager returns to the local organization and begins to search for a default policy or available resource in the default pool.
4. If an applicable default policy or available resource in a default pool is found, Cisco UCS Manager uses that policy or resource. If the pool does not have any available resources, Cisco UCS Manager moves up in the hierarchy to the parent organization and searches for a default pool. Cisco UCS Manager repeats this step until the search reaches the root organization.
5. If Cisco UCS Manager cannot find an applicable policy or available resource in the hierarchy, it returns an allocation error.

Example: Server Pool Name Resolution in a Single-Level Hierarchy

In this example, all organizations are at the same level below the root organization. For example, a service provider creates separate organizations for each customer. In this configuration, organizations only have access to the policies and resource pools assigned to that organization and to the root organization.

In this example, a service profile in the XYZcustomer organization is configured to use servers from the XYZcustomer server pool. When resource pools and policies are assigned to the service profile, the following occurs:

1. Cisco UCS Manager checks for an available server in the XYZcustomer server pool.
2. If the XYZcustomer server pool has an available server, Cisco UCS Manager associates that server with the service profile and discontinues the search. If the pool does not have an available server, Cisco UCS Manager checks the root organization for a server pool with the same name.
3. If the root organization includes an XYZcustomer server pool and that pool has an available server, Cisco UCS Manager associates that server with the service profile and discontinues the search. If the pool does not have an available server, Cisco UCS Manager returns to the XYZcustomer organization to check the default server pool.
4. If the default pool in the XYZcustomer organization has an available server, Cisco UCS Manager associates that server with the service profile and discontinues the search. If the default pool does not have an available server, Cisco UCS Manager checks the default server pool in the root organization.

5 If the default server pool in the root organization has an available server, Cisco UCS Manager associates that server with the service profile and discontinues the search. If the default pool does not have an available server, Cisco UCS Manager returns an allocation error.

Example: Server Pool Name Resolution in a Multi-Level Hierarchy

In this example, each organization includes at least one suborganization. For example, a company could create organizations for each major division in the company and for subdivisions of those divisions. In this configuration, each organization has access to its local policies and resource pools and to the resource pools in the parent hierarchy.

In this example, the Finance organization includes two sub-organizations, AccountsPayable and AccountsReceivable. A service profile in the AccountsPayable organization is configured to use servers from the AP server pool. When resource pools and policies are assigned to the service profile, the following occurs:

1 Cisco UCS Manager checks for an available server in the AP server pool defined in the service profile.

2 If the AP server pool has an available server, Cisco UCS Manager associates that server with the service profile and discontinues the search. If the pool does not have an available server, Cisco UCS Manager moves one level up the hierarchy and checks the Finance organization for a pool with the same name.

3 If the Finance organization includes a pool with the same name and that pool has an available server, Cisco UCS Manager associates that server with the service profile and discontinues the search. If the pool does not have an available server, Cisco UCS Manager moves one level up in the hierarchy and checks the root organization for a pool with the same name.

4 If the root organization includes a pool with the same name and that pool has an available server, Cisco UCS Manager associates that server with the service profile and discontinues the search. If the pool does not have an available server, Cisco UCS Manager returns to the AccountsPayable organization to check the default server pool.

5 If the default pool in the AccountsPayable organization has an available server, Cisco UCS Manager associates that server with the service profile and discontinues the search. If the default pool does not have an available server, Cisco UCS Manager moves one level up in the hierarchy and checks the default server pool in the Finance organization.

6 If the default pool in the Finance organization has an available server, Cisco UCS Manager associates that server with the service profile and discontinues the search. If the default pool does not have an available server, Cisco UCS Manager moves one level up in the hierarchy and checks the default server pool in the root organization.

7 If the default server pool in the root organization has an available server, Cisco UCS Manager associates that server with the service profile and discontinues the search. If the default pool does not have an available server, Cisco UCS Manager returns an allocation error.

Creating an Organization under the Root Organization

Procedure

Step 1 On the toolbar, choose **New** > **Create Organization**.

Step 2 In the **Name** field of the **Create Organization** dialog box, enter a unique name for the organization.

This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Step 3 In the **Description** field, enter a description for the organization.

Step 4 Click **OK**.

Creating an Organization under a Sub-Organization

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 In the **Servers** tab, expand **Service Profiles** > **root**.
You can also access the **Sub-Organizations** node under the **Policies** or **Pools** nodes.

Step 3 Expand the **Sub-Organizations** node and do one of the following:

- To create an organization directly under root, right-click **Sub-Organizations** and choose **Create Organization**.
- To create an organization under a lower-level sub-organization, expand the sub-organization nodes in the hierarchy and then right-click the sub-organization under which you want to create the new organization and choose **Create Organization**.

Step 4 In the **Name** field of the **Create Organization** dialog box, enter a unique name for the organization.
This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Step 5 In the **Description** field, enter a description for the organization.

Step 6 Click **OK**.

Deleting an Organization

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Navigate to the organization that you want to delete.

Step 3 Right-click the organization and choose **Delete**.

Step 4 If a confirmation dialog box displays, click **Yes**.

CHAPTER 10

Configuring Role-Based Access Control

This chapter includes the following sections:

Role-Based Access Control Overview

Role-Based Access Control (RBAC) is a method of restricting or authorizing system access for users based on user roles and locales. A role defines the privileges of a user in the system and a locale defines the organizations (domains) that a user is allowed access. Because users are not directly assigned privileges, you can manage individual user privileges by assigning the appropriate roles and locales.

A user is granted write access to the required system resources only if the assigned role grants the access privileges and the assigned locale allows access. For example, a user with the Server Administrator role in the engineering organization can update server configurations in the Engineering organization. They cannot, however, update server configurations in the Finance organization, unless the locales assigned to the user include the Finance organization.

User Accounts for Cisco UCS

User accounts access the system. You can configure up to 48 local user accounts in each Cisco UCS Manager domain. Each user account requires a unique username and password.

You can set user accounts with an SSH public key. The public key can be set in either of the two formats: OpenSSH or SECSH.

Admin Account

An admin account comes with each Cisco UCS domain. The admin account is a default user account and cannot be modified or deleted. This account is the system administrator or superuser account s full privileges. There is no default password assigned to the admin account; you must choose the password during the initial system setup.

The admin account is always active and does not expire. You cannot configure the admin account as inactive.

Locally Authenticated User Accounts

A locally authenticated user account is authenticated directly through the fabric interconnect and can be enabled or disabled by anyone with admin or aaa privileges. After a local user account is disabled, the user cannot log in. The database does not delete the configuration details for disabled local user accounts. If you re-enable a disabled local user account, the account becomes active with the existing configuration, including the username and password.

Remotely Authenticated User Accounts

A remotely authenticated user account is any user account that is authenticated through LDAP, RADIUS, or TACACS+.

If a user maintains a local user account and a remote user account simultaneously, the roles defined in the local user account override those maintained in the remote user account.

Expiration of User Accounts

You can configure user accounts to expire at a predefined time. When the expiration time is reached, the user account is disabled.

By default, user accounts do not expire.

Note After you configure a user account with an expiration date, you cannot reconfigure the account to not expire. However, you can configure the account to use the latest expiration date available.

Guidelines for Cisco UCS Usernames

The username is also used as the login ID for Cisco UCS Manager. When you assign login IDs to Cisco UCS user accounts, consider the following guidelines and restrictions:

- The login ID can contain between 1 and 32 characters, including the following:
  - Any alphabetic character
  - Any digit
  - _ (underscore)
  - - (dash)
  - . (dot)

- The login ID must be unique within Cisco UCS Manager.
- The login ID must start with an alphabetic character. It cannot start with a number or a special character, such as an underscore.
- The login ID is case-sensitive.
- You cannot create an all-numeric login ID.
- After you create a user account, you cannot change the login ID. You must delete the user account and create a new one.

Reserved Words: Locally Authenticated User Accounts

You cannot use the following words when creating a local user account in Cisco UCS.

- root
- bin
- daemon
- adm
- lp
- sync
- shutdown
- halt
- news
- uucp
- operator
- games
- gopher
- nobody
- nscd
- mailnull
- mail
- rpcuser
- rpc
- mtsuser
- ftpuser
- ftp
- man
- sys

- samdme
- debug

Guidelines for Cisco UCS Passwords

Each locally authenticated user account requires a password. A user with admin or aaa privileges can configure Cisco UCS Manager to perform a password strength check on user passwords.

Cisco recommends using a strong password; otherwise, the password strength check for locally authenticated users, Cisco UCS Manager rejects any password that does not meet the following requirements:

- Must contain a minimum of eight characters and a maximum of 80 characters.
- If the password strength check is turned on, the minimum password length is variable and can be set from a minimum of 6 to a maximum of 80 characters.

Note The default is 8 characters.

- Must contain at least three of the following:
  - Lower case letters
  - Upper case letters
  - Digits
  - Special characters
- Must not contain a character that is repeated more than three times consecutively, such as aaabbb.
- Must not be identical to the username or the reverse of the username.
- Must pass a password dictionary check. For example, the password must not be based on a standard dictionary word.
- Must not contain the following symbols: $ (dollar sign), ? (question mark), and = (equals sign).
- Should not be blank for local user and admin accounts.

Web Session Limits for User Accounts

Cisco UCS Manager uses web session limits to restrict the number of web sessions (both GUI and XML) that a given user account is permitted to access at any one time.

Each Cisco UCS Manager domain supports a maximum of 32 concurrent web sessions per user and 256 total user sessions. By default, the number of concurrent web sessions allowed by Cisco UCS Manager is set to 32 per user, but you can configure this value up to the system maximum of 256.

User Roles

User roles contain one or more privileges that define the operations that are allowed for a user. You can assign one or more roles to each user. Users with multiple roles have the combined privileges of all assigned roles. For example, if Role1 has storage-related privileges, and Role 2 has server-related privileges, users with Role1 and Role 2 have both storage-related and server-related privileges.

A Cisco UCS domain can contain up to 48 user roles, including the default user roles. Any user roles configured after the first 48 are accepted, but they are inactive with faults raised.

All roles include read access to all configuration settings in the Cisco UCS domain. Users with read-only roles cannot modify the system state.

You can create, modify or remove existing privileges, and delete roles. When you modify a role, the new privileges apply to all users with that role. Privilege assignment is not restricted to the privileges defined for the default roles. Meaning, you can use a custom set of privileges to create a unique role. For example, the default Server Administrator and Storage Administrator roles have a different set of privileges. However, you can create a Server and Storage Administrator role that combines the privileges of both roles.

Note If you delete a role after it was assigned to users, it is also deleted from those user accounts.

Modify the user profiles on AAA servers (RADIUS or TACACS+) to add the roles corresponding to the privileges granted to that user. The attribute stores the role information. The AAA servers return this attribute with the request and parse it to obtain the roles. LDAP servers return the roles in the user profile attributes.

Note If a local and a remote user account have the same username, Cisco UCS Manager overrides any roles assigned to the remote user with those assigned to the local user.

Default User Roles

The system contains the following default user roles:

AAA Administrator

Read-and-write access to users, roles, and AAA configuration. Read access to the remaining system.

Administrator

Complete read-and-write access to the entire system. Assigns this role to the default administrator account by default. You cannot change it.

Facility Manager

Read-and-write access to power management operations through the power management privilege. Read access to the remaining system.

Network Administrator

Read-and-write access to fabric interconnect infrastructure and network security operations. Read access to the remaining system.

Operations

Read-and-write access to systems logs, including the syslog servers, and faults. Read access to the remaining system.

Read-Only

Read-only access to system configuration with no privileges to modify the system state.

Server Compute

Read and write access to most aspects of service profiles. However, the user cannot create, modify or delete vNICs or vHBAs.

Server Equipment Administrator

Read-and-write access to physical server-related operations. Read access to the remaining system.

Server Profile Administrator

Read-and-write access to logical server-related operations. Read access to the remaining system.

Server Security Administrator

Read-and-write access to server security-related operations. Read access to the remaining system.

Storage Administrator

Read-and-write access to storage operations. Read access to the remaining system.

Reserved Words: User Roles

You cannot use the following words when creating custom roles in Cisco UCS.

- network-admin
- network-operator
- vdc-admin
- vdc-operator
- server-admin

Privileges

Privileges give users, assigned to user roles, access to specific system resources and permission to perform specific tasks. The following table lists each privilege and the user role given that privilege by default.

Tip Detailed information about these privileges and the tasks that they enable users to perform is available in *Privileges in Cisco UCS* available at the following URL: http://www.cisco.com/en/US/products/ps10281/prod_technical_reference_list.html.

Table 12: User Privileges

Privilege	Description	Default Role Assignment
aaa	System security and AAA	AAA Administrator
admin	System administration	Administrator
ext-lan-config	External LAN configuration	Network Administrator
ext-lan-policy	External LAN policy	Network Administrator
ext-lan-qos	External LAN QoS	Network Administrator
ext-lan-security	External LAN security	Network Administrator
ext-san-config	External SAN configuration	Storage Administrator
ext-san-policy	External SAN policy	Storage Administrator
ext-san-qos	External SAN QoS	Storage Administrator
ext-san-security	External SAN security	Storage Administrator
fault	Alarms and alarm policies	Operations
operations	Logs and Smart Call Home	Operations
org-management	Organization management	Operations
pod-config	Pod configuration	Network Administrator
pod-policy	Pod policy	Network Administrator
pod-qos	Pod QoS	Network Administrator
pod-security	Pod security	Network Administrator
power-mgmt	Read-and-write access to power management operations	Facility Manager

Privilege	Description	Default Role Assignment
read-only	Read-only access Read-only cannot be selected as a privilege; it is assigned to every user role.	Read-Only
server-equipment	Server hardware management	Server Equipment Administrator
server-maintenance	Server maintenance	Server Equipment Administrator
server-policy	Server policy	Server Equipment Administrator
server-security	Server security	Server Security Administrator
service-profile-compute	Service profile compute	Server Compute Administrator
service-profile-config	Service profile configuration	Server Profile Administrator
service-profile-config-policy	Service profile configuration policy	Server Profile Administrator
service-profile-ext-access	Service profile endpoint access	Server Profile Administrator
service-profile-network	Service profile network	Network Administrator
service-profile-network-policy	Service profile network policy	Network Administrator
service-profile-qos	Service profile QoS	Network Administrator
service-profile-qos-policy	Service profile QoS policy	Network Administrator
service-profile-security	Service profile security	Server Security Administrator
service-profile-security-policy	Service profile security policy	Server Security Administrator
service-profile-server	Service profile server management	Server Profile Administrator
service-profile-server-oper	Service profile consumer	Server Profile Administrator
service-profile-server-policy	Service profile pool policy	Server Security Administrator
service-profile-storage	Service profile storage	Storage Administrator
service-profile-storage-policy	Service profile storage policy	Storage Administrator

User Locales

You can assign a user to one or more locales. Each locale defines one or more organizations (domains) to which a user can access. Access is usually limited to the organizations specified in the locale. An exception is a locale without any organizations. It provides unrestricted access to system resources in all organizations.

A Cisco UCS domain can contain up to 48 user locales. Any user locales configured after the first 48 are accepted, but are inactive with faults raised.

Users with admin or aaa privileges can assign organizations to the locale of other users. The assignment of organizations is restricted to only those in the locale of the user assigning the organizations. For example, if a locale contains only the Engineering organization, a user assigned to that locale can only assign the Engineering organization to other users.

Note You cannot assign a locale to users with one or more of the following privileges:

- aaa
- admin
- fault
- operations

You can hierarchically manage organizations. A user who is assigned to a top-level organization has automatic access to all organizations below it. For example, an Engineering organization can contain a Software Engineering organization and a Hardware Engineering organization. A locale containing only the Software Engineering organization has access to system resources only within that organization. However, a locale that contains the Engineering organization has access to the resources for both the Software Engineering and Hardware Engineering organizations.

Configuring User Roles

Creating a User Role

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management** > **User Services**.

Step 3 Right-click **User Services** and choose **Create Role**.
You can also right-click **Roles** to access that option.

Step 4 In the **Create Role** dialog box, complete the following fields:

Name	Description
Name field	A user-defined name for this user role. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Privileges list box	A list of the privileges defined in the system. Click a privilege to view a description of that privilege. Check the check box to assign that privilege to the selected user.
Help Section	
Description field	A description of the most recent privilege you clicked in the **Privileges** list box.

Step 5 Click **OK**.

Adding Privileges to a User Role

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **User Management** > **User Services**.
Step 3 Expand the **Roles** node.
Step 4 Choose the role to which you want to add privileges.
Step 5 In the **General** tab, check the boxes for the privileges you want to add to the role.
Step 6 Click **Save Changes**.

Removing Privileges from a User Role

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **User Management** > **User Services**.
Step 3 Expand the **Roles** node.
Step 4 Choose the role from which you want to remove privileges.
Step 5 In the **General** tab, uncheck the boxes for the privileges you want to remove from the role.
Step 6 Click **Save Changes**.

Deleting a User Role

When you delete a user role, Cisco UCS Manager removes that role from all user accounts to which the role was assigned.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **User Management** > **User Services**.
Step 3 Expand the **Roles** node.
Step 4 Right-click the role you want to delete and choose **Delete**.
Step 5 In the **Delete** dialog box, click **Yes**.

Configuring Locales

Creating a Locale

Before You Begin

One or more organizations must exist before you create a locale.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management** > **User Services**.

Step 3 Right-click **Locales** and choose **Create a Locale**.

Step 4 In the **Create Locale** page, do the following:

a) In the **Name** field, enter a unique name for the locale.
This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

b) Click **Next**.

Step 5 In the **Assign Organizations** dialog box, do the following:

a) Expand the **Organizations** area to view the organizations in the Cisco UCS domain.
b) Expand the **root** node to see the sub-organizations.
c) Click an organization that you want to assign to the locale.
d) Drag the organization from the **Organizations** area and drop it into the design area on the right.
e) Repeat Steps b and c until you have assigned all desired organizations to the locale.

Step 6 Click **Finish**.

What to Do Next

Add the locale to one or more user accounts. For more information, see Changing the Locales Assigned to a Locally Authenticated User Account, on page 192.

Assigning an Organization to a Locale

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management** > **User Services**.

Step 3 Expand the **Locales** node and click the locale to which you want to add an organization.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Organizations** area, click + on the table icon bar.

Step 6 In the **Assign Organizations** dialog box, do the following:

a) Expand the **Organizations** area to view the organizations in the Cisco UCS domain.
b) Expand the **root** node to see the sub-organizations.
c) Click an organization that you want to assign to the locale.
d) Drag the organization from the **Organizations** area and drop it into the design area on the right.

e) Repeat Steps b and c until you have assigned all desired organizations to the locale.

Step 7 Click **OK**.

Deleting an Organization from a Locale

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management** > **User Services**.

Step 3 Expand the **Locales** node and click the locale from which you want to delete an organization.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Organizations** area, right-click the organization that you want to delete from the locale and choose **Delete**.

Step 6 Click **Save Changes**.

Deleting a Locale

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management** > **User Services**.

Step 3 Expand the **Locales** node.

Step 4 Right-click the locale you want to delete and choose **Delete**.

Step 5 If a confirmation dialog box displays, click **Yes**.

Configuring Locally Authenticated User Accounts

Creating a User Account

At a minimum, Cisco recommends that you create the following users:

- Server administrator account
- Network administrator account
- Storage administrator

Note After you create the user account, if you make any changes to any of the user account fields from the Cisco UCS Manager GUI, make sure to enter the password again.

Before You Begin

Perform the following tasks, if the system includes any of the following:

- Remote authentication services—Ensures that the users exist in the remote authentication server with the appropriate roles and privileges.
- Multitenancy with organizations—Creates one or more locales. If you do not have any locales, all users are created in root and are assigned roles and privileges in all organizations.
- SSH authentication—Obtains the SSH key.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management** > **User Services**.

Step 3 Right-click **User Services** and choose **Create User** to open the **User Properties** dialog box.
You can also right-click **Locally Authenticated Users** to access that option.

Step 4 Complete the following fields with the required information about the user:

Name	Description
Login ID field	The account name that is used when logging into this account. This account must be unique and meet the following guidelines and restrictions for Cisco UCS Manager user accounts: • The login ID can contain between 1 and 32 characters, including the following: ◦ Any alphabetic character ◦ Any digit ◦ _ (underscore) ◦ - (dash) ◦ . (dot) • The login ID must be unique within Cisco UCS Manager. • The login ID must start with an alphabetic character. It cannot start with a number or a special character, such as an underscore. • The login ID is case-sensitive. • You cannot create an all-numeric login ID. • After you create a user account, you cannot change the login ID. You must delete the user account and create a new one. After you save the user, the login ID cannot be changed. You must delete the user account and create a new one.
First Name field	The first name of the user. This field can contain up to 32 characters.
Last Name field	The last name of the user. This field can contain up to 32 characters.
Email field	The email address for the user.
Phone field	The telephone number for the user.

Name	Description
Password field	The password associated with this account. If password strength check is enabled, a user's password must be strong and Cisco UCS Manager rejects any password that does not meet the following requirements: • Must contain a minimum of eight characters and a maximum of 80 characters. • If the password strength check is turned on, the minimum password length is variable and can be set from a minimum of 6 to a maximum of 80 characters. **Note** The default is 8 characters. • Must contain at least three of the following: ◦ Lower case letters ◦ Upper case letters ◦ Digits ◦ Special characters • Must not contain a character that is repeated more than three times consecutively, such as aaabbb. • Must not be identical to the username or the reverse of the username. • Must pass a password dictionary check. For example, the password must not be based on a standard dictionary word. • Must not contain the following symbols: $ (dollar sign), ? (question mark), and = (equals sign). • Should not be blank for local user and admin accounts.
Confirm Password field	The password a second time for confirmation purposes.
Account Status field	If the status is set to **Active**, a user can log into Cisco UCS Manager with this login ID and password.
Account Expires check box	If checked, this account expires and cannot be used after the date specified in the **Expiration Date** field. **Note** After you configure a user account with an expiration date, you cannot reconfigure the account to not expire. However, you can configure the account to use the latest expiration date available.

Name	Description
Expiration Date field	The date on which the account expires. The date should be in the format yyyy-mm-dd. Click the down arrow at the end of this field to view a calendar that you can use to select the expiration date. **Note** Cisco UCS Manager GUI displays this field when you check the **Account Expires** check box.

Step 5 In the **Roles** area, check one or more boxes to assign roles and privileges to the user account.

Note Do not assign locales to users with an admin or aaa role.

Step 6 (Optional) If the system includes organizations, check one or more check boxes in the **Locales** area to assign the user to the appropriate locales.

Step 7 In the **SSH** area, complete the following fields:

a) In the **Type** field, click the following:

- **Password Required—**The user must enter a password when they log in.
- **Key—**SSH encryption is used when this user logs in.

b) If you chose **Key**, enter the SSH key in the **SSH data** field.

Step 8 Click **OK**.

Enabling the Password Strength Check for Locally Authenticated Users

You must have admin or aaa privileges to enable the password strength check. If enabled, Cisco UCS Manager does not permit a user to choose a password that does not meet the guidelines for a strong password.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management** > **User Services**.

Step 3 Click the **Locally Authenticated Users** node.

Step 4 In the **Work** pane, check the **Password Strength Check** check box in the **Properties** area.

Step 5 Click **Save Changes**.

Setting the Web Session Limits for Cisco UCS Manager GUI Users

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **Communication Management** > **Communication Services**.

Step 3 Click the **Communication Services** tab.

Step 4 In the **Web Session Limits** area, complete the following fields:

Name	Description
Maximum Sessions Per User	The maximum number of concurrent HTTP and HTTPS sessions allowed for each user. Enter an integer between 1 and 256.
Maximum Sessions	The maximum number of concurrent HTTP and HTTPS sessions allowed for all users within the system. Enter an integer between 1 and 256.
Maximum Event Interval (in seconds)	The maximum time interval between two events. Tracks various types of event change notifications, such as responses to any user requests from the UI. If the interval expires, the UI session is terminated. Enter and integer between 120-3600

Step 5 Click **Save Changes**.

Changing the Locales Assigned to a Locally Authenticated User Account

Note Do not assign locales to users with an admin or aaa role.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 On the **Admin** tab, expand **All** > **User Management** > **User Services** > **Locally Authenticated Users**.

Step 3 Click the user account that you want to modify.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Locales** area, do the following:

- To assign a new locale to the user account, check the appropriate check boxes.
- To remove a locale from the user account, uncheck the appropriate check boxes.

Step 6 Click **Save Changes**.

Changing the Roles Assigned to a Locally Authenticated User Account

Changes in user roles and privileges do not take effect until the next time the user logs in. If a user is logged in when you assign a new role to or remove an existing role from a user account, the active session continues with the previous roles and privileges.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 On the **Admin** tab, expand **All** > **User Management** > **User Services** > **Locally Authenticated Users**.

Step 3 Click the user account that you want to modify.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Roles** area, do the following:

- To assign a new role to the user account, check the appropriate check boxes.
- To remove a role from the user account, uncheck the appropriate check boxes.

Step 6 Click **Save Changes**.

Enabling a User Account

You must have admin or aaa privileges to enable or disable a local user account.

Before You Begin

Create a local user account.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **User Management** > **User Services** > **Locally Authenticated Users**.
Step 3 Click the user that you want to enable.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Account Status** field, click the **active** radio button.
Step 6 Click **Save Changes**.

Disabling a User Account

You must have admin or aaa privileges to enable or disable a local user account.

Note If you change the password on a disabled account through the Cisco UCS Manager GUI, the user cannot use this changed password after you enable the account and make it active. The user must enter the required password again after the account is enabled and made active.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **User Management** > **User Services** > **Locally Authenticated Users**.
Step 3 Click the user that you want to disable.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Account Status** field, click the **inactive** radio button.
The admin user account is always set to active. It cannot be modified.
Step 6 Click **Save Changes**.

Clearing the Password History for a Locally Authenticated User

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **User Management** > **User Services** > **Locally Authenticated Users**.
Step 3 Click the user for whom you want to clear the password history.
Step 4 In the **Actions** area, click **Clear Password History**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Deleting a Locally Authenticated User Account

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **User Management** > **User Services**.
Step 3 Expand the **Locally Authenticated Users** node.
Step 4 Right-click the user account you want to delete and choose **Delete**.
Step 5 In the **Delete** dialog box, click **Yes**.

Password Profile for Locally Authenticated Users

The password profile contains the password history and the password change interval properties for all locally authenticated users of Cisco UCS Manager. You cannot specify a different password profile for locally authenticated users.

Note You must have admin or aaa privileges to change the password profile properties. Except for password history, these properties do not apply to users with admin or aaa privileges.

Password History Count

The password history count prevents locally authenticated users from reusing the same password. When you configure the password history count, Cisco UCS Manager stores up to a maximum of 15 previously used passwords. The password history count stores the passwords in reverse chronological order with the most recent password first. This ensures that the user can only reuse the oldest password when the history count reaches its threshold.

A user can create and use the number of passwords configured in the password history count before reusing a password. For example, if you set the password history count to 8, a user cannot reuse the first password until the ninth password expires.

By default, the password history is set to 0. This value disables the history count and allows users to reuse previously used passwords at any time.

You can clear the password history count for a locally authenticated user and enable reuse of previous passwords.

Password Change Interval

The password change interval restricts the number of password changes that a locally authenticated user can make within a specific number of hours. The following table describes the two interval configuration options for the password change interval.

Interval Configuration	Description	Example
No password change allowed	Does not allow changing passwords for locally authenticated user within a specified number of hours after a password change. You can specify a no change interval between 1 and 745 hours. By default, the no change interval is 24 hours.	To prevent the user from changing passwords within 48 hours after a password change: • Set **Change during interval** to disable • Set **No change interval** to 48
Password changes allowed within change interval	Specifies the maximum number of times that a locally authenticated user password change can occur within a pre-defined interval. You can specify a change interval between 1 and 745 hours and a maximum number of password changes between 0 and 10. By default, a locally authenticated user is permitted a maximum of two password changes within a 48-hour interval.	To allow a password change for a maximum of one time within 24 hours after a password change: • Set **Change during interval** to enable • Set **Change count** to 1 • Set **Change interval** to 24

Configuring the Maximum Number of Password Changes for a Change Interval

You must have admin or aaa privileges to change the password profile properties. Except for password history, these properties do not apply to users with admin or aaa privileges.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management** > **User Services**.

Step 3 Click the **Locally Authenticated Users** node.

Step 4 In the **Password Profile** area, do the following:

a) In the **Change During Interval** field, click **Enable**.

b) In the **Change Count** field, enter the maximum number of times a locally authenticated user can change his or her password during the Change Interval.
This value can be anywhere from 0 to 10.

c) In the **Change Interval** field, enter the maximum number of hours over which the number of password changes specified in the **Change Count** field are enforced.
This value can be anywhere from 1 to 745 hours.

For example, if this field is set to 48 and the **Change Count** field is set to 2, a locally authenticated user can make no more than 2 password changes within a 48 hour period.

Step 5 Click **Save Changes**.

Configuring a No Change Interval for Passwords

You must have admin or aaa privileges to change the password profile properties. Except for password history, these properties do not apply to users with admin or aaa privileges.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management** > **User Services**.

Step 3 Click the **Locally Authenticated Users** node.

Step 4 In the **Password Profile** area, do the following:

a) In the **Change During Interval** field, click **Disable**.

b) In the **No Change Interval** field, enter the minimum number of hours that a locally authenticated user must wait before changing a newly created password.
This value can be anywhere from 1 to 745 hours.

This interval is ignored if the **Change During Interval** property is set to **Disable**.

Step 5 Click **Save Changes**.

Configuring the Password History Count

You must have admin or aaa privileges to change the password profile properties.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **User Management** > **User Services**.

Step 3 Click the **Locally Authenticated Users** node.

Step 4 In the **Password Profile** area, enter the number of unique passwords that a locally authenticated user must create before that user can reuse a previously used password in the **History Count** field.
This value can be anywhere from 0 to 15.

By default, the **History Count** field is set to 0, which disables the history count and allows users to reuse previously used passwords at any time.

Step 5 Click **Save Changes**.

Monitoring User Sessions from the GUI

You can monitor Cisco UCS Manager sessions for both locally authenticated users and remotely authenticated users, whether they logged in through the CLI or the GUI.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 In the **Admin** tab, expand **All** > **User Management**.

Step 3 Click the **User Services** node.

Step 4 In the **Work** pane, click the **Sessions** tab.
The tab displays the following details of user sessions:

Name	Description
Name column	The name for the session.
User column	The username that is involved in the session.
Fabric ID column	The fabric interconnect that the user logged in to for the session.
Login Time column	The date and time the session started.

Name	Description
Refresh Period column	When a web client connects to Cisco UCS Manager, the client must send refresh requests to Cisco UCS Manager to keep the web session active. This option specifies the maximum amount of time allowed between refresh requests for a user in this domain. If this time limit is exceeded, Cisco UCS Manager considers the web session inactive, but it does not terminate the session.
Session Timeout column	The maximum amount of time that can elapse after the last refresh request before Cisco UCS Manager considers a web session as inactive. If this time limit is exceeded, Cisco UCS Manager automatically terminates the web session.
Terminal Type column	The kind of terminal the user is logged in through.
Host column	The IP address from which the user is logged in.
Current Session column	If this column displays **Y**, the associated user session is currently active.

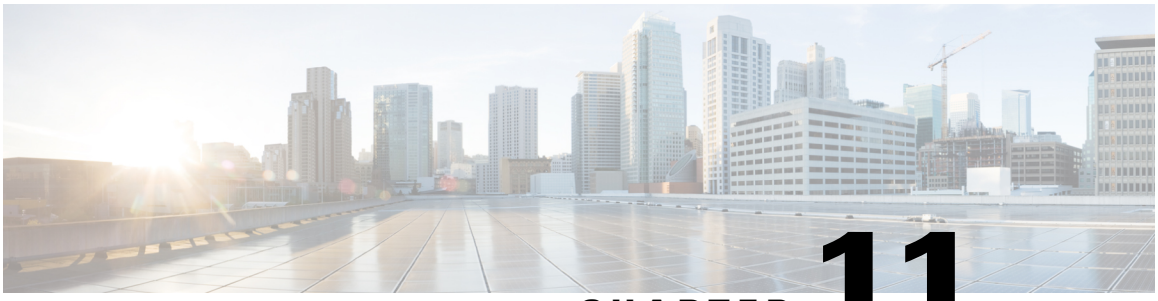

CHAPTER 11

Configuring DNS Servers

This chapter includes the following sections:

DNS Servers in Cisco UCS

You need to specify an external DNS server for each Cisco UCS domain to use if the system requires name resolution of hostnames. For example, you cannot use a name such as www.cisco.com when you are configuring a setting on a fabric interconnect if you do not configure a DNS server. You would need to use the IP address of the server, which can be either an IPv4 or an IPv6 address. You can configure up to four DNS servers for each Cisco UCS domain.

Note When you configure multiple DNS servers, the system searches for the servers only in any random order. If a local management command requires DNS server lookup, it can only search for three DNS servers in random order.

Adding a DNS Server

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 In the **Admin** tab, expand **All** > **Communication Services**.
Step 3 Click **DNS Management**.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **DNS Server** area, click +.
Step 6 In the **Specify DNS Server** dialog box, enter the IPv4 or IPv6 address of the DNS server.
Step 7 Click **OK**.

Deleting a DNS Server

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 In the **Admin** tab, expand **All** > **Communication Services**.
Step 3 Click **DNS Management**.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **DNS Server** area, right-click the DNS server you want to delete and choose **Delete**.
Step 6 If a confirmation dialog box displays, click **Yes**.
Step 7 Click **Save Changes**.

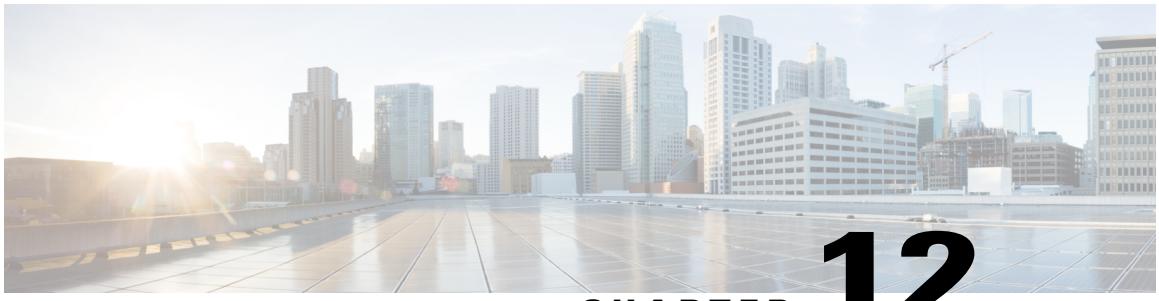

CHAPTER 12

Configuring System-Related Policies

This chapter includes the following sections:

Configuring the Chassis/FEX Discovery Policy

Chassis/FEX Discovery Policy

The chassis/FEX discovery policy determines how the system reacts when you add a new chassis or FEX. Cisco UCS Manager uses the settings in the chassis/FEX discovery policy to determine the minimum threshold for the number of links between the chassis or FEX and the fabric interconnect and whether to group links from the IOM to the fabric interconnect in a fabric port channel.

Chassis Links

If you have a Cisco UCS domain with some of the chassis' wired with one link, some with two links, some with four links, and some with eight links, Cisco recommends configuring the chassis/FEX discovery policy for the minimum number links in the domain so that Cisco UCS Manager can discover all chassis.

Tip To establish the highest available chassis connectivity in a Cisco UCS domain where Fabric Interconnect is connected to different types of IO Modules supporting different max number of uplinks, select platform max value. Setting the platform max ensures that Cisco UCS Manager discovers the chassis including the connections and servers only when the maximum supported IOM uplinks are connected per IO Module.

After the initial discovery, re-acknowledge the chassis' that are wired for a greater number of links and Cisco UCS Manager configures the chassis to use all available links.

Cisco UCS Manager cannot discover any chassis that is wired for fewer links than are configured in the chassis/FEX discovery policy. For example, if the chassis/FEX discovery policy is configured for four links,

Cisco UCS Manager cannot discover any chassis that is wired for one link or two links. Re-acknowledgement of the chassis resolves this issue.

The following table provides an overview of how the chassis/FEX discovery policy works in a multi-chassis Cisco UCS domain:

Table 13: Chassis/FEX Discovery Policy and Chassis Links

Number of Links Wired for the Chassis	1-Link Discovery Policy	2-Link Discovery Policy	4-Link Discovery Policy	8-Link Discovery Policy	Platform-Max Discovery Policy
1 link between IOM and fabric interconnects	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 1 link.	Chassis connections and servers cannot be discovered by Cisco UCS Manager and are not added to the Cisco UCS domain.	Chassis connections and servers cannot be discovered by Cisco UCS Manager and are not added to the Cisco UCS domain.	Chassis connections and servers cannot be discovered by Cisco UCS Manager and are not added to the Cisco UCS domain.	Chassis connections and servers cannot be discovered by Cisco UCS Manager and are not added to the Cisco UCS domain.
2 links between IOM and fabric interconnects	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 1 link. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 2 link.	Chassis connections and servers cannot be discovered by Cisco UCS Manager and are not added to the Cisco UCS domain.	Chassis connections and servers cannot be discovered by Cisco UCS Manager and are not added to the Cisco UCS domain.	Chassis connections and servers cannot be discovered by Cisco UCS Manager and are not added to the Cisco UCS domain.

Number of Links Wired for the Chassis	1-Link Discovery Policy	2-Link Discovery Policy	4-Link Discovery Policy	8-Link Discovery Policy	Platform-Max Discovery Policy
4 links between IOM and fabric interconnects	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 1 link. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 2 links. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 4 link.	Chassis connections and servers cannot be discovered by Cisco UCS Manager and are not added to the Cisco UCS domain.	If the IOM has 4 links, the chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 4 links. If the IOM has 8 links, the chassis is not fully discovered by Cisco UCS Manager.
8 links between IOM and fabric interconnects	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 1 link. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 2 links. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 4 links. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 8 links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 8 links.

Link Grouping

For hardware configurations that support fabric port channels, link grouping determines whether all of the links from the IOM to the fabric interconnect are grouped in to a fabric port channel during chassis discovery. If the link grouping preference is set to port channel, all of the links from the IOM to the fabric interconnect

are grouped in a fabric port channel. If set to no group, links from the IOM to the fabric interconnect are not grouped in a fabric port channel.

After you create a fabric port channel, you can add or remove links by changing the link group preference and re-acknowledging the chassis, or by enabling or disabling the chassis from the port channel.

Note The link grouping preference only takes effect if both sides of the links between an IOM or FEX and the fabric interconnect support fabric port channels. If one side of the links does not support fabric port channels, this preference is ignored and the links are not grouped in a port channel.

Configuring the Chassis/FEX Discovery Policy

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Click the **Equipment** node.

Step 3 In the **Work** pane, click the **Policies** tab.

Step 4 Click the **Global Policies** subtab.

Step 5 In the **Chassis/FEX Discovery Policy** area, specify the action and the link grouping preference.

a) In the **Action** field, specify the minimum threshold for the number of links between the chassis or FEX and the fabric interconnect.

b) In the **Link Grouping Preference** field, specify whether the links from the IOMs or FEXes to the fabric interconnects are grouped in a port channel.

c) In the **Multicast Hardware Hash** field, specify whether all the links from the IOMs or FEXes to the fabric interconnects in a port channel can be used for multicast traffic.

Step 6 Click **Save Changes**.

What to Do Next

To customize fabric port channel connectivity for a specific chassis, configure the chassis connectivity policy.

Configuring the Chassis Connectivity Policy

Chassis Connectivity Policy

The chassis connectivity policy determines the whether a specific chassis is included in a fabric port channel after chassis discovery. This policy is helpful for users who want to configure one or more chassis differently from what is specified in the global chassis discovery policy. The chassis connectivity policy also allows for different connectivity modes per fabric interconnect, further expanding the level of control offered with regards to chassis connectivity.

By default, the chassis connectivity policy is set to global. This means that connectivity control is configured when the chassis is newly discovered, using the settings configured in the chassis discovery policy. Once the chassis is discovered, the chassis connectivity policy controls whether the connectivity control is set to none or port channel.

Note The chassis connectivity policy is created by Cisco UCS Manager only when the hardware configuration supports fabric port channels. At this time, only the 6200 series fabric interconnects and the 2200 series IOMs support this feature. For all other hardware combinations, Cisco UCS Manager does not create a chassis connectivity policy.

Configuring a Chassis Connectivity Policy

Changing the connectivity mode for a chassis might result in decreased VIF namespace.

Caution Changing the connectivity mode for a chassis results in chassis re-acknowledgement. Traffic might be disrupted during this time.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis**.

Step 3 Click the chassis for which you want to configure the connectivity between the IOMs and fabric interconnects.

Step 4 In the **Work** pane, click the **Connectivity Policy** tab.

Step 5 For each IOM in the chassis, choose one of the following values in the **Admin State** field for the chassis and fabric connectivity:

- **None—**No links are grouped in a port channel
- **Port Channel—**All links from an IOM to a fabric interconnect are grouped in a port channel.
- **Global—**The chassis inherits this configuration from the chassis discovery policy. This is the default value.

Step 6 Click **Save Changes**.

Configuring the Rack Server Discovery Policy

Rack Server Discovery Policy

The rack server discovery policy determines how the system reacts when you add a new rack-mount server. Cisco UCS Manager uses the settings in the rack server discovery policy to determine whether any data on

the hard disks are scrubbed and whether server discovery occurs immediately or needs to wait for explicit user acknowledgement.

Cisco UCS Manager cannot discover any rack-mount server that has not been correctly cabled and connected to the fabric interconnects. For information about how to integrate a supported Cisco UCS rack-mount server with Cisco UCS Manager, see the appropriate rack-mount server integration guide.

Configuring the Rack Server Discovery Policy

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Click the **Equipment** node.

Step 3 In the **Work** pane, click the **Policies** tab.

Step 4 Click the **Global Policies** subtab.

Step 5 In the **Rack Server Discovery Policy** area, specify the action that you want to occur when a new rack server is added and the scrub policy.

Step 6 Click **Save Changes**.

Configuring the Aging Time for the MAC Address Table

Aging Time for the MAC Address Table

To efficiently switch packets between ports, the fabric interconnect maintains a MAC address table. It dynamically builds the MAC address table by using the MAC source address from the packets received and the associated port on which the packets were learned. The fabric interconnect uses an aging mechanism, defined by a configurable aging timer, to determine how long an entry remains in the MAC address table. If an address remains inactive for a specified number of seconds, it is removed from the MAC address table.

You can configure the amount of time (age) that a MAC address entry (MAC address and associated port) remains in the MAC address table.

Configuring the Aging Time for the MAC Address Table

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Click the **Equipment** node.
Step 3 In the **Work** pane, click the **Policies** tab.
Step 4 Click the **Global Policies** subtab.
Step 5 In the **MAC Address Table Aging** area, specify the aging time and the length of time.
Step 6 Click **Save Changes**.

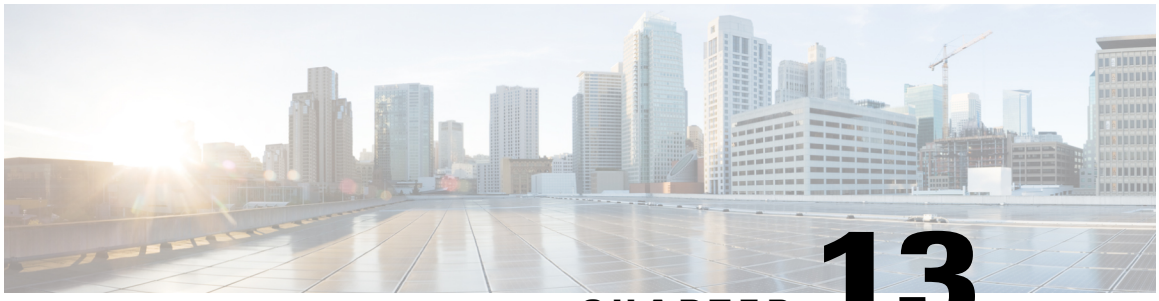

CHAPTER 13

Managing Licenses

This chapter includes the following sections:

Licenses

Each Cisco UCS fabric interconnect comes with several port licenses that are factory installed and shipped with the hardware. You can purchase fabric interconnects fully licensed or partially licensed. You can also purchase additional licenses after delivery.

The following four new licenses are added for the 6300 Series FI and are only valid on the 6332 and 6332-16UP Fis.

- 40G_ETH_PORT_ACTIVATION_PKG – Licenses used for 40 GB Ethernet ports
- 40G_ETH_C_PORT_ACTIVATION_PKG – Licenses used for 40 GB Ethernet ports directly connected to rack servers (C-Direct)
- 10G_C_PORT_ACTIVATION_PKG – Licenses used for the first 16 10 GB unified ports on the 6332-16UP that are directly connected to rack servers (C-Direct)

- 10G_PORT_ACTIVATION_PKG – Licenses used for the first 16 10 GB unified ports on the 6332-16UP

Note The 10G_PORT_ACTIVATION_PKG and 10G_C_PORT_ACTIVATION_PKG licenses are only valid for the 6332-16UP FIs, and can only be installed on them.

At a minimum, each fabric interconnect ships with the following counted licenses pre-installed:

Fabric Interconnect	Default Base Licenses
Cisco UCS 6248 (unified ports)	For the 12 first enabled Ethernet ports and any Fibre Channel ports in the expansion module.
Cisco UCS 6296 (unified ports)	For the first 18 enabled Ethernet ports and any Fibre Channel ports in the expansion module.
Cisco UCS 6324	For 4 non-breakout ports only. The fifth port, which does not include a license, is further broken in to four 10 GB ports.
Cisco UCS 6332 16UP	For four 40 GB ports and eight 10 GB ports. **Note** The first 16 ports are 10 GB. The remaining are 40 GB.
Cisco UCS 6332	For eight 40 GB ports.

Port License Consumption

Port licenses are not bound to physical ports. When you disable a licensed port, that license is retained for use with the next enabled port. To use additional fixed ports, you must purchase and install licenses for those ports. All ports, regardless of their type (fibre, ethernet) consume licenses if they are enabled.

For breakout capable ports available in the 6332 and the 6332-16UP platforms, 40 GB licenses remain applied to the main port even if that port is a breakout port, and that port continues to consume only one 40 GB license.

Note The initial configuration of a port will enable it, and consume a license.

Important Licenses are not portable across product generations. Licenses purchased for 6200 series fabric interconnects cannot be used to enable ports on 6300 series fabric interconnects or vice-versa.

Each Cisco UCS 6324 Fabric Interconnect comes with a factory installed port license that is shipped with the hardware. This license is for the eight 40 GB unified ports, and can be used for any supported purpose. The C-direct port license is factory installed with a grace period, and can be used for Cisco UCS rack servers.

Grace Period

If you attempt to use a port that does not have an installed license, Cisco UCS initiates a 120 day grace period. The grace period is measured from the first use of the port without a license and is paused when a valid license file is installed. The amount of time used in the grace period is retained by the system.

Note Each physical port has its own grace period. Initiating the grace period on a single port does not initiate the grace period for all ports.

If a licensed port is unconfigured, that license is transferred to a port functioning within a grace period. If multiple ports are acting within grace periods, the license is moved to the port whose grace period is closest to expiring.

High Availability Configurations

To avoid inconsistencies during failover, we recommend that both fabric interconnects in the cluster have the same number of ports licensed. If symmetry is not maintained and failover occurs, Cisco UCS enables the missing licenses and initiates the grace period for each port being used on the failover node.

C-Direct Rack Licensing Support

Each Cisco UCS fabric interconnect is shipped with a default number of port licenses that are factory licensed and shipped with the hardware. C-direct support is only applicable on ports that are connected to the rack servers. The 10G_C_PORT_ACTIVATION_PKG and the 40G_ETH_C_PORT_ACTIVATION _PKG are added to the existing license package with all the same properties as the existing licensing feature. The **Subordinate Quantity** property is added to the 10G_PORT_ACTIVATION_PKG and 40G_ETH_PORT_ACTIVATION_PKG to track ports connected to rack servers.

The License Tab in the Cisco UCS Manager GUI displays the new license and the **Subordinate Quantity** for the license. You can also use the **show feature** and **show usage** commands under **scope license** to view the license feature, the vendor version type, and the grace period for each license.

Ports connected to rack servers can use existing 10G_PORT_ACTIVATION_PKG and 40G_ETH_PORT_ACTIVATION_PKG if the license is available or if the license is not in use. Otherwise, you must purchase a 10G_C_PORT_ACTIVATION_PKG and 40G_ETH_C_PORT_ACTIVATION to avoid the license grace period.

There is no change in the 10 GB ports. The 10G_PORT_ACTIVATION_PKG and 10G_C_PORT_ACTIVATION_PKG license packages include all of the same properties as the existing the ETH_PORT_ACTIVATION_PKG and the ETH_PORT_C_ACTIVATION_PKG license features.

Configuration and Restrictions

- The C-Direct rack licensing feature accounts for the rack server ports that are directly connected to the FI, but not to a CIMC port. The default quantity for the 10G_C_PORT_ACTIVATION_PKG and 40G_ETH_C_PORT_ACTIVATION_PKG is always 0.
- When a 40 GB port, or a breakout port under a 40 GB breakout port is enabled without any connections, this port is allotted a license under the 40G_ETH_PORT_ACTIVATION_PKG, if available. If this port is connected to a Direct-Connect rack server after a time lag, it triggers a complete re-allocation of licenses, then this port passes through one of the following license allocation scenarios occurs:

When you enable a breakout port under a 40 GB breakout port, if that port is connected to a Direct-Connect rack server, and the 40G_C_PORT_ACTIVATION_PKG license files are installed on the FI, the following license allocation occurs:

- If no other ports under the breakout port are enabled, the parent 40 GB port is allotted a license under the 40G_C_PORT_ACTIVATION_PKG, and the used quantity is incremented for this instance.
- If other ports are enabled, and if at least one port is not connected to a Direct Connect rack server, even if the port is not being used, the parent 40 GB port is allotted a license under the 40G_ETH_PORT_ACTIVATION_PKG, and the used quantity is incremented for this instance.

- When you enable a breakout port under a 40 GB breakout port and that port is connected to a Direct-Connect rack server, and the 40G_C_PORT_ACTIVATION_PKG license files are not installed on the FI, the following license allocation occurs:
  - If no ports under the breakout port are enabled, the parent 40 GB port is allotted a license under the 40G_ETH_PORT_ACTIVATION_PKG. The subordinate quantity is increased if the licenses are available in the 40G_ETH_PORT_ACTIVATION_PKG. If the licenses are not available, the used quantity under this feature is increased and the entire port goes in to the grace period.
  - If other ports are enabled and at least one port is not connected to a Direct Connect rack server, even if the port is not being used, the parent 40 GB port is allotted a license under the 40G_ETH_PORT_ACTIVATION_PKG, and the used quantity is incremented for this instance.

Obtaining the Host ID for a Fabric Interconnect

The host ID is also known as the serial number.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Fabric Interconnects**.
Step 3 Click the node for the fabric interconnect for which you want to obtain the host ID.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Properties** area, the host ID is listed in the **Serial Number (SN)** field.

What to Do Next

Obtain the required licenses from Cisco.

Obtaining a License

Note This process may change after the release of this document. If one or more of these steps no longer applies, contact your Cisco representative for information on how to obtain a license file.

Before You Begin

Obtain the following:

- Host ID or serial number for the fabric interconnect
- Claim certificate or other proof of purchase document for the fabric interconnect or expansion module

Procedure

Step 1 Obtain the product authorization key (PAK) from the claim certificate or other proof of purchase document.

Step 2 Locate the website URL in the claim certificate or proof of purchase document.

Step 3 Access the website URL for the fabric interconnect and enter the serial number and the PAK.
Cisco sends you the license file by email. The license file is digitally signed to authorize use on only the requested fabric interconnect. The requested features are also enabled once Cisco UCS Manager accesses the license file.

What to Do Next

Install the license on the fabric interconnect.

Downloading Licenses to the Fabric Interconnect from the Local File System

Note In a cluster setup, Cisco recommends that you download and install licenses to both fabric interconnects in matching pairs. An individual license is only downloaded to the fabric interconnect that is used to initiate the download.

Before You Begin

Obtain the required licenses from Cisco.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **License Management**.

Step 3 Click the node for the fabric interconnect to which you want to download the license.

Step 4 In the **Work** pane, click the **Download Tasks** tab.

Step 5 Click **Download License**.

Step 6 In the **Download License** dialog box, click the **Local File System** radio button in the **Location of the Image File** field.

Step 7 In the **Filename** field, type the full path and name of the license file.
You cannot have spaces anywhere in the path name or the file name. For example, `c:\Path\Folder_Name\License.lic` is a valid path, but `c:\Path\Folder Name\License.lic` is invalid due to the space in "`Folder Name`".

If you do not know the exact path to the folder where the license file is located, click **Browse** and navigate to the file.

Step 8 Click **OK**.
Cisco UCS Manager GUI begins downloading the license to the fabric interconnect.

Step 9 (Optional) Monitor the status of the download on the **Download Tasks** tab.

Note If Cisco UCS Manager reports that the bootflash is out of space, delete obsolete bundles on the **Packages** tab to free up space. To view the available space in bootflash, navigate to the fabric interconnect, click **Equipment**, and expand the **Local Storage Information** area on the **General** tab.

Step 10 Repeat this task until all the required licenses are downloaded to the fabric interconnect.

What to Do Next

After all of the download tasks complete, install the licenses.

Downloading Licenses to the Fabric Interconnect from a Remote Location

Note In a cluster setup, Cisco recommends that you download and install licenses to both fabric interconnects in matching pairs. An individual license is only downloaded to the fabric interconnect that is used to initiate the download.

Before You Begin

Obtain the required licenses from Cisco.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **License Management**.

Step 3 Click the node for the fabric interconnect to which you want to download the license.

Step 4 In the **Work** pane, click the **Download Tasks** tab.

Step 5 Click **Download License**.

Step 6 In the **Download License** dialog box, click the **Remote File System** radio button in the **Location of the Image File** field.

Step 7 Specify the protocol, and enter the required information.
You cannot have spaces anywhere in the path name or the file name. For example, `c:\Path\Folder_Name\License.lic` is a valid path, but `c:\Path\Folder Name\License.lic` is invalid due to the space in "`Folder Name`".

Note If you use a hostname rather than an IPv4 or IPv6 address, you must configure a DNS server. If the Cisco UCS domain is not registered with Cisco UCS Central or DNS management is set to **local**, configure a DNS server in Cisco UCS Manager. If the Cisco UCS domain is registered with Cisco UCS Central and DNS management is set to **global**, configure a DNS server in Cisco UCS Central.

Step 8 Click **OK**.
Cisco UCS Manager GUI begins downloading the license to the fabric interconnect.

Step 9 (Optional) Monitor the status of the download on the **Download Tasks** tab.

Note If Cisco UCS Manager reports that the bootflash is out of space, delete obsolete bundles on the **Packages** tab to free up space. To view the available space in bootflash, navigate to the fabric interconnect, click **Equipment**, and expand the **Local Storage Information** area on the **General** tab.

Step 10 Repeat this task until all the required licenses are downloaded to the fabric interconnect.

What to Do Next

After all of the download tasks complete, install the licenses.

Installing a License

Before You Begin

Obtain the required licenses from Cisco.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **License Management**.
Step 3 In the **Work** pane, click the **Downloaded License Files** tab.
Step 4 Choose the license you want to install from the table.
Step 5 Click the **Install License** button.
Step 6 In the **Install License** dialog box, click **Yes**.
Cisco UCS Manager GUI installs the license and activates the unlicensed port or feature.

Viewing the Licenses Installed on a Fabric Interconnect

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **License Management**.
Step 3 In the **Work** pane, click the **Installed Licenses** tab to view the details of all licenses installed on the fabric interconnect.
Step 4 Click a license in the table to view the details of that license in the **Contents** tab.
You may need to expand the license file to view the details of individual licenses in the file.

Determining the Grace Period Available for a Port or Feature

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **License Management**.
Step 3 In the **Work** pane, click the **General** tab.
Step 4 Click a feature in the table to view details for that feature including the operational state and used grace period.

Determining the Expiry Date of a License

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **License Management**.
Step 3 In the **Work** pane, click the **Installed Licenses** tab.
Step 4 Click a license in the table to view the details of that license in the **Contents** tab below.
Step 5 In the **Contents** tab, expand the license file to view all licenses in the file.
Step 6 In the **Expiry** column, view the expiry date of the license.

Uninstalling a License

Note Permanent licenses cannot be uninstalled if they are in use. You can only uninstall a permanent license that is not in use. If you try to delete a permanent license that is being used, Cisco UCS Manager rejects the request and display an error message.

Before You Begin

Back up the Cisco UCS Manager configuration.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **License Management**.
Step 3 In the **Work** pane, click the **Installed Licenses** tab.
Step 4 Choose the license you want to uninstall from the table.
Step 5 Click the **Clear License** button.
Step 6 If a confirmation dialog box displays, click **Yes**.

Cisco UCS Manager deactivates the license, removes the license from the list of licenses, and deletes the license from the fabric interconnect. The port is moved into unlicensed mode. In a cluster setup, you must uninstall the license from the other fabric interconnect.

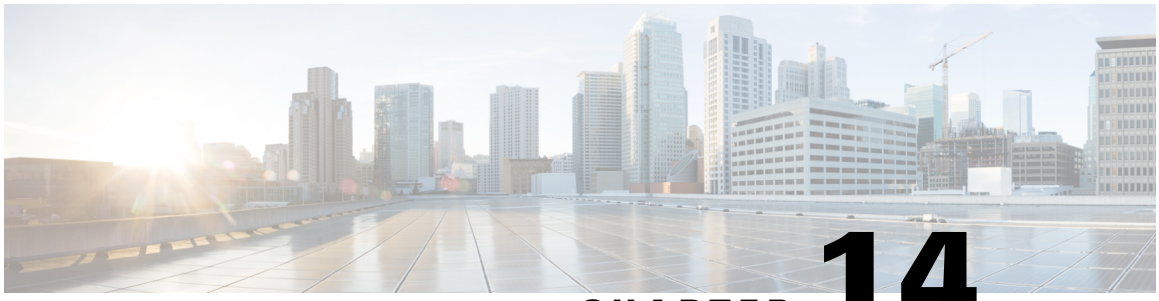

CHAPTER 14

Managing Virtual Interfaces

This chapter includes the following sections:

Virtual Interfaces

In a blade server environment, the number of vNICs and vHBAs configurable for a service profile is determined by adapter capability and the amount of virtual interface (VIF) namespace available on the adapter. In Cisco UCS, portions of VIF namespace are allotted in chunks called VIFs. Depending on your hardware, the maxiumum number of VIFs are allocated on a predefined, per-port basis.
The maximum number of VIFs varies based on hardware capability and port connectivity. For each configured vNIC or vHBA, one or two VIFs are allocated. Stand-alone vNICs and vHBAs use one VIF and failover vNICs and vHBAs use two.

The following variables affect the number of VIFs available to a blade server, and therefore, how many vNICs and vHBAs you can configure for a service profile.

- Maximum number of VIFs supported on your fabric interconnect
- How the fabric interconnects are cabled
- If your fabric interconnect and IOM are configured in fabric port channel mode

For more information about the maximum number of VIFs supported by your hardware configuration, see the appropriate *Cisco UCS Configuration Limits for Cisco UCS Manager* for your software release.

Virtual Interface Subscription Management and Error Handling

For fabric interconnects grouped in a port-channel, changes to the way you connect the fabric interconnect to the I/O module could result in a drastic change to the number of VIFs available to a blade server. To help you track the effect of these changes, Cisco UCS Manager maintains the following metrics:

- Maximum number of VIFs supported by hardware

- Connectivity type

If you change your configuration in a way that decreases the number of VIFs available to a blade, UCS Manager will display a warning and ask you if you want to proceed. This includes several scenarios, including times where adding or moving a connection decreases the number of VIFs.

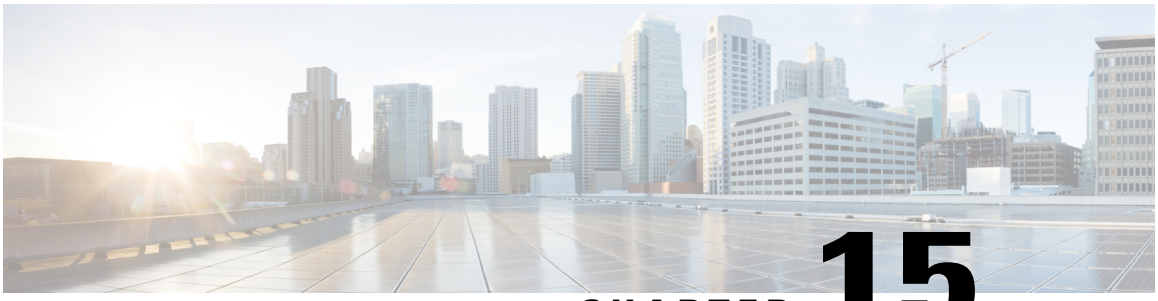

CHAPTER 15

Registering Cisco UCS Domains with Cisco UCS Central

This chapter includes the following sections:

Registration of Cisco UCS Domains

You can have Cisco UCS Central manage some or all of the Cisco UCS domains in your data center.

If you want to have Cisco UCS Central manage a Cisco UCS domain, you need to register that domain. When you register, you need to choose which types of policies and other configurations, such as backups and firmware, will be managed by Cisco UCS Central and which by Cisco UCS Manager. You can have Cisco UCS Central manage the same types of policies and configurations for all registered Cisco UCS domains or you can choose to have different settings for each registered Cisco UCS domain.

Before you register a Cisco UCS domain with Cisco UCS Central, do the following:

- Configure an NTP server and the correct time zone in both Cisco UCS Manager and Cisco UCS Central to ensure that they are in sync. If the time and date in the Cisco UCS domain and Cisco UCS Central are out of sync, the registration might fail.
- Obtain the hostname or IP address of Cisco UCS Central
- Obtain the shared secret that you configured when you deployed Cisco UCS Central

Policy Resolution between Cisco UCS Manager and Cisco UCS Central

For each Cisco UCS domain that you register with Cisco UCS Central, you can choose which application will manage certain policies and configuration settings. This policy resolution does not have to be the same for every Cisco UCS domain that you register with the same Cisco UCS Central.

You have the following options for resolving these policies and configuration settings:

- **Local—**The policy or configuration is determined and managed by Cisco UCS Manager.
- **Global—**The policy or configuration is determined and managed by Cisco UCS Central.

Note The policy resolution options in Cisco UCS Central are not supported on all versions of Cisco UCS Manager. If your Cisco UCS Manager version is earlier than the earliest supported release, the policy resolution screen may display the value as global even if it is not applicable.

The following table contains a list of the policies and configuration settings that you can choose to have managed by either Cisco UCS Manager or Cisco UCS Central:

Name	Earliest Supported Release	Description
Infrastructure & Catalog Firmware	2.1(2)	Determines whether the Capability Catalog and infrastructure firmware policy are defined locally in Cisco UCS Manager or come from Cisco UCS Central.
Time Zone Management	2.1(2)	Determines whether the time zone and NTP server settings are defined locally in Cisco UCS Manager or comes from Cisco UCS Central.
Communication Services	2.1(2)	Determines whether HTTP, CIM XML, Telnet, SNMP, web session limits, and Management Interfaces Monitoring Policy settings are defined locally in Cisco UCS Manager or in Cisco UCS Central.
Global Fault Policy	2.1(2)	Determines whether the Global Fault Policy is defined locally in Cisco UCS Manager or in Cisco UCS Central.
User Management	2.1(2)	Determines whether authentication and native domains, LDAP, RADIUS, TACACS+, trusted points, locales, and user roles are defined locally in Cisco UCS Manager or in Cisco UCS Central.
DNS Management	2.1(2)	Determines whether DNS servers are defined locally in Cisco UCS Manager or in Cisco UCS Central.

Name	Earliest Supported Release	Description
Backup & Export Policies	2.1(2)	Determines whether the Full State Backup Policy and All Configuration Export Policy are defined locally in Cisco UCS Manager or in Cisco UCS Central.
Monitoring	2.1(2)	Determines whether Call Home, Syslog, and TFTP Core Exporter settings are defined locally in Cisco UCS Manager or in Cisco UCS Central.
SEL Policy	2.1(2)	Determines whether the SEL Policy is defined locally in Cisco UCS Manager or in Cisco UCS Central.
Power Allocation Policy	2.1(2)	Determines whether the Power Allocation Policy is defined locally in Cisco UCS Manager or in Cisco UCS Central.
Power Policy	2.1(2)	Determines whether the Power Policy is defined locally in Cisco UCS Manager or in Cisco UCS Central.
Equipment Policy	2.2(7)	Determines whether the Equipment Policy is defined locally in Cisco UCS Manager or in Cisco UCS Central.
Port Configuration	2.2(7)	Determines whether port configuration is defined locally in Cisco UCS Manager or in Cisco UCS Central.
Quality of Service (QoS) Configuration	2.2(7)	Determines whether QoS configuration is defined locally in Cisco UCS Manager or in Cisco UCS Central.

Registering a Cisco UCS Domain with Cisco UCS Central

Before You Begin

Configure an NTP server and the correct time zone in both Cisco UCS Manager and Cisco UCS Central to ensure that they are in sync. If the time and date in the Cisco UCS domain and Cisco UCS Central are out of sync, the registration might fail.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **Communication Management**.

Step 3 Click the **UCS Central** node.

Step 4 In the **Actions** area, click **UCS Central**.

Step 5 In the **Actions** area, click **Register With UCS Central**.

Step 6 In the **Register with UCS Central** dialog box, do the following:

a) Complete the following fields:

Name	Description
Hostname/IP Address field	The hostname or IP address of the virtual machine where Cisco UCS Central is deployed. **Note** If you use a hostname rather than an IPv4 or IPv6 address, you must configure a DNS server. If the Cisco UCS domain is not registered with Cisco UCS Central or DNS management is set to **local**, configure a DNS server in Cisco UCS Manager. If the Cisco UCS domain is registered with Cisco UCS Central and DNS management is set to **global**, configure a DNS server in Cisco UCS Central.
Shared Secret field	The shared secret (or password) that was configured when Cisco UCS Central was deployed.

b) In the **Policy Resolution Control** area, click one of the following radio buttons for each of the fields:

- **Local—**The policy or configuration is determined and managed by Cisco UCS Manager.
- **Global—**The policy or configuration is determined and managed by Cisco UCS Central.

c) Click **OK**.

Modifying Policy Resolutions between Cisco UCS Manager and Cisco UCS Central

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **Communication Management**.

Step 3 Click the **UCS Central** node.

Step 4 In the **Actions** area, click **UCS Central**.

Step 5 In the **Policy Resolution Control** area, click one of the following radio buttons for each of the fields:

- **Local—**The policy or configuration is determined and managed by Cisco UCS Manager.
- **Global—**The policy or configuration is determined and managed by Cisco UCS Central.

Step 6 Click **Save Changes**.

Setting Cisco UCS Central Registration Properties in Cisco UCS Manager

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Expand **All** > **Communication Management**.

Step 3 Click the **UCS Central** node.

Step 4 In the **Actions** area, click **UCS Central**.

Step 5 In the **Status** area, complete the following as appropriate:

a) Click the radio button for the **Cleanup Mode** that you want to use.
This can be one of the following:

- **Localize Global—**When a Cisco UCS domain is unregistered, all global policies in the Cisco UCS domain will be localized to Cisco UCS Manager. The policies remain in the Cisco UCS domain, policy ownership is now local to Cisco UCS Manager, and Cisco UCS Manager admin users can make changes.

  Note If you reregister the Cisco UCS domain with Cisco UCS Central, there can be policy conflicts due to the policies existing both in Cisco UCS Central and in Cisco UCS Manager. Either delete the local policies, or set the local policies to global before you try to create and associate a global service profile.

- **Deep Remove Global—**This option should only be used after careful consideration. When a Cisco UCS domain is unregistered, all global policies in the Cisco UCS domain are removed. If there are global service profiles, they will now refer to Cisco UCS Manager local default policies, and one of the following occurs:
  - If there are local default policies present, the server will reboot.
  - If there are no local default policies, the service profile association fails with a configuration error.

  Note The deep remove global cleanup mode does not remove global VSANs and VLANs when you unregister from Cisco UCS Central. Those must be removed manually if desired.

b) Optionally check the **Suspend State** check box.
If checked, the Cisco UCS domain is temporarily removed from Cisco UCS Central, and all global policies revert to their local counterparts. All service profiles maintain their current identities. However, global pools are no longer visible and cannot be accessible by new service profiles.

c) Optionally check the **Acknowledge State** check box.
If the event ID stream that represents time and consistency between Cisco UCS Manager and Cisco UCS Central becomes skewed or inconsistent, Cisco UCS Manager places itself in a Suspended State and disconnects itself from Cisco UCS Central.

If you check this check box, you acknowledge that inconsistencies exist between Cisco UCS Manager and Cisco UCS Central and are still willing to reconnect the Cisco UCS domain with Cisco UCS Central.

Step 6 Click **Save Changes**.

Unregistering a Cisco UCS Domain from Cisco UCS Central

When you unregister a Cisco UCS domain from Cisco UCS Central, Cisco UCS Manager no longer receives updates to global policies.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.
Step 2 Expand **All** > **Communication Management**.
Step 3 Click the **UCS Central** node.
Step 4 In the **Actions** area, click **UCS Central**.
Step 5 In the **Actions** area, click **Unregister From UCS Central**.
Step 6 If a confirmation dialog box displays, click **Yes**.
Step 7 Click **OK**.

For more information on the impact of unregistering and registering a Cisco UCS Domain with Cisco UCS Central, see Policy Resolution between Cisco UCS Manager and Cisco UCS Central.

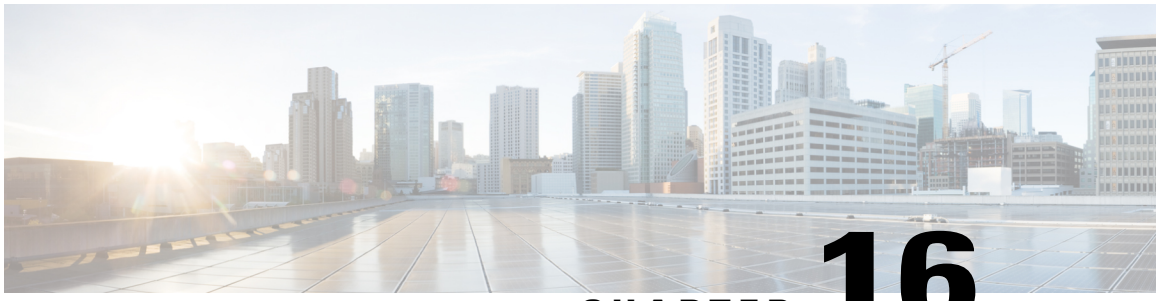

CHAPTER 16

LAN Uplinks Manager

This chapter includes the following sections:

LAN Uplinks Manager

The LAN Uplinks Manager provides a single interface where you can configure the connections between Cisco UCS and the LAN. You can use the LAN Uplinks Manager to create and configure the following:

- Ethernet switching mode
- Uplink Ethernet ports
- Port channels
- LAN pin groups
- Named VLANs
- Server ports
- QoS system classes

Some of the configuration that you can do in the LAN Uplinks Manager can also be done in nodes on other tabs, such as the **Equipment** tab or the **LAN** tab.

Launching the LAN Uplinks Manager

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 On the **LAN** tab, click the **LAN** node.

Step 3 In the **Work** pane, click the **LAN Uplinks Manager** link on the **LAN Uplinks** tab.
The LAN Uplinks Manager opens in a separate window.

Changing the Ethernet Switching Mode with the LAN Uplinks Manager

Important When you change the Ethernet switching mode, Cisco UCS Manager logs you out and restarts the fabric interconnect. For a cluster configuration, Cisco UCS Manager restarts both fabric interconnects. The subordinate fabric interconnect reboots first as a result of the change in switching mode. The primary fabric interconnect reboots only after you acknowledge it in **Pending Activities**. The primary fabric interconnect can take several minutes to complete the change in Ethernet switching mode and become system ready. The existing configuration is retained.

While the fabric interconnects are rebooting, all blade servers lose LAN and SAN connectivity, causing a complete outage of all services on the blades. This might cause the operating system to fail.

Procedure

Step 1 In the LAN Uplinks Manager, click the **LAN Uplinks** tab.

Step 2 In the **Uplink Mode** area, click one of the following buttons:

- **Set Ethernet Switching Mode**
- **Set Ethernet End-Host Mode**

The button for the current switching mode is dimmed.

Step 3 In the dialog box, click **Yes**.
Cisco UCS Manager restarts the fabric interconnect, logs you out, and disconnects Cisco UCS Manager GUI.

Configuring a Port with the LAN Uplinks Manager

All of the port types listed are configurable on both the fixed and expansion module, including server ports, which are not configurable on the 6100 series fabric interconnect expansion module, but are configurable on the 6200 series fabric interconnect expansion module.

Procedure

Step 1 In the LAN Uplinks Manager, click the **LAN Uplinks** tab.

Step 2 In the **Ports** area, click the down arrows to expand the **Unconfigured Ports** section.

Step 3 Expand **Fabric Interconnects** > ***Fabric_Interconnect_Name*** .

Step 4 Expand the node where you want to configure ports.
If no ports are listed below the node that you expanded, all ports in that module have already been configured.

Step 5 Right-click the port that you want to configure and choose one of the following:

- **Configure as Server Port**
- **Configure as Uplink Port**

Step 6 If a confirmation dialog box displays, click **Yes**.

Configuring Server Ports

Enabling a Server Port with the LAN Uplinks Manager

This procedure assumes that the port has been configured as a server port, but is disabled.

Procedure

Step 1 In the LAN Uplinks Manager, click the **LAN Uplinks** tab.

Step 2 In the **Ports** area, click the down arrows to expand the **Server Ports** section.

Step 3 Expand **Fabric Interconnects** > ***Fabric_Interconnect_Name*** .

Step 4 Right-click the port that you want to enable and choose **Enable**.

Disabling a Server Port with the LAN Uplinks Manager

Procedure

Step 1 In the LAN Uplinks Manager, click the **LAN Uplinks** tab.
Step 2 In the **Ports** area, click the down arrows to expand the **Server Ports** section.
Step 3 Expand **Fabric Interconnects** > ***Fabric_Interconnect_Name*** .
Step 4 Right-click the port that you want to disable and choose **Disable**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Configuring Uplink Ethernet Ports

Enabling an Uplink Ethernet Port with the LAN Uplinks Manager

This procedure assumes that the port has been configured as an uplink Ethernet port, but is disabled.

Procedure

Step 1 In the LAN Uplinks Manager, click the **LAN Uplinks** tab.
Step 2 In the **Port Channels and Uplinks** area, expand **Interfaces** > **Fabric Interconnects** > ***Fabric_Interconnect_Name*** .
Step 3 Right-click the port that you want to enable and choose **Enable Interface**.
Step 4 If a confirmation dialog box displays, click **Yes**.

Disabling an Uplink Ethernet Port with the LAN Uplinks Manager

Procedure

Step 1 In the LAN Uplinks Manager, click the **LAN Uplinks** tab.
Step 2 In the **Port Channels and Uplinks** area, expand **Interfaces** > **Fabric Interconnects** > ***Fabric_Interconnect_Name*** .
Step 3 Right-click the port that you want to disable and choose **Disable Interfaces**.
You can select multiple ports if you want to disable more than one uplink Ethernet port.
Step 4 If a confirmation dialog box displays, click **Yes**.

The disabled port is removed from the list of enabled interfaces and returned to the **Unconfigured Ports** list.

Configuring Uplink Ethernet Port Channels

Creating a Port Channel with the LAN Uplinks Manager

Procedure

Step 1 In the LAN Uplinks Manager, click the **LAN Uplinks** tab.

Step 2 In the **Port Channels and Uplinks** area, click **Create Port Channel**.

Step 3 From the pop-up menu, select one of the following fabric interconnects where you want to create the port channel:

- **Fabric Interconnect A**
- **Fabric Interconnect B**

Step 4 In the **Set Port Channel Name** panel, specify the ID and name, then click **Next**.

Step 5 In the **Add Ports** panel, specify the ports you want to add.

Note Cisco UCS Manager warns you if you select a port that has been configured as a sever port. You can reconfigure that port as an uplink Ethernet port and include it in the port channel by clicking **Yes** in the dialog box.

Step 6 In the **Add Ports** panel, specify the ports that you want to add.

Note Cisco UCS Manager warns you if you select a port that has been configured as a server port. You can click **Yes** in the dialog box to reconfigure that port as an uplink Ethernet port and include it in the port channel.

Step 7 Click **Finish**.

Enabling a Port Channel with the LAN Uplinks Manager

Procedure

Step 1 In the LAN Uplinks Manager, click the **LAN Uplinks** tab.

Step 2 In the **Port Channels and Uplinks** area, expand **Port Channels** > **Fabric Interconnects** > ***Fabric_Interconnect_Name***.

Step 3 Right-click the port channel that you want to enable and choose **Enable Port Channel**.

Step 4 If a confirmation dialog box displays, click **Yes**.

Disabling a Port Channel with the LAN Uplinks Manager

Procedure

Step 1 In the LAN Uplinks Manager, click the **LAN Uplinks** tab.

Step 2 In the **Port Channels and Uplinks** area, expand **Port Channels** > **Fabric Interconnects** > ***Fabric_Interconnect_Name***.

Step 3 Right-click the port channel that you want to disable and choose **Disable Port Channel**.

Step 4 If a confirmation dialog box displays, click **Yes**.

Adding Ports to a Port Channel with the LAN Uplinks Manager

Procedure

Step 1 In the LAN Uplinks Manager, click the **LAN Uplinks** tab.

Step 2 In the **Port Channels and Uplinks** area, expand **Port Channels** > **Fabric Interconnects** > ***Fabric_Interconnect_Name***.

Step 3 Right-click the port channel to which you want to add ports and choose **Add Ports**.

Step 4 In the **Add Ports** dialog box, specify the ports that you want to add.

Note Cisco UCS Manager warns you if you select a port that has been configured as a server port. You can click **Yes** in the dialog box to reconfigure that port as an uplink Ethernet port and include it in the port channel.

Step 5 Click **OK**.

Removing Ports from a Port Channel with the LAN Uplinks Manager

Procedure

Step 1 In the LAN Uplinks Manager, click the **LAN Uplinks** tab.

Step 2 In the **Port Channels and Uplinks** area, expand **Port Channels** > **Fabric Interconnects** > ***Fabric_Interconnect_Name***.

Step 3 Expand the port channel from which you want to remove ports.

Step 4 Right-click the port you want to remove from the port channel and choose **Delete**.

Step 5 If a confirmation dialog box displays, click **Yes**.

Deleting a Port Channel with the LAN Uplinks Manager

Procedure

Step 1 In the LAN Uplinks Manager, click the **LAN Uplinks** tab.

Step 2 In the **Port Channels and Uplinks** area, expand **Port Channels** > **Fabric Interconnects** > ***Fabric_Interconnect_Name***.

Step 3 Right-click the port channel you want to delete and choose **Delete**.

Step 4 If a confirmation dialog box displays, click **Yes**.

Configuring LAN Pin Groups

Creating a Pin Group with the LAN Uplinks Manager

In a system with two fabric interconnects, you can associate the pin group with only one fabric interconnect or with both fabric interconnects.

Before You Begin

Configure the ports and port channels with which you want to configure the pin group. You can only include ports and port channels configured as uplink ports in a LAN pin group.

Procedure

Step 1 In the LAN Uplinks Manager, click the **LAN Uplinks** tab.

Step 2 In the **Port Channels and Uplinks** area, click **Create Pin Group**.

Step 3 In the **Create LAN Pin Group** dialog box, enter a unique name and description for the pin group.

Step 4 To pin traffic for fabric interconnect A, do the following in the **Targets** area:

a) Check the **Fabric Interconnect A** check box.

b) Click the drop-down arrow on the **Interface** field and navigate through the tree-style browser to select the port or port channel you want to associate with the pin group.

Step 5 To pin traffic for fabric interconnect B, do the following in the **Targets** area:

a) Check the **Fabric Interconnect B** check box.

b) Click the drop-down arrow on the **Interface** field and navigate through the tree-style browser to select the port or port channel you want to associate with the pin group.

Step 6 Click **OK**.

What to Do Next

Include the pin group in a vNIC template.

Deleting a Pin Group with the LAN Uplinks Manager

Procedure

Step 1 In the LAN Uplinks Manager, click the **LAN Uplinks** tab.

Step 2 In the **Pin Groups** area, right-click the pin group you want to delete and choose **Delete**.

Step 3 If a confirmation dialog box displays, click **Yes**.

Configuring Named VLANs

Creating a Named VLAN with the LAN Uplinks Manager

In a Cisco UCS domain with two switches, you can create a named VLAN that is accessible to both switches or to only one switch.

Important You cannot create VLANs with IDs from 4030 to 4047. This range of VLAN IDs is reserved.

The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch.

VLANs in the LAN cloud and FCoE VLANs in the SAN cloud must have different IDs. Using the same ID for a VLAN and an FCoE VLAN in a VSAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Procedure

Step 1 In the LAN Uplinks Manager, click the **VLANs** tab.

Step 2 On the icon bar to the right of the table, click +.
If the + icon is disabled, click an entry in the table to enable it.

Step 3 In the **Create VLANs** dialog box, specify the required fields and then click **OK**.

Step 4 Click **OK**.
Cisco UCS Manager adds the VLAN to one of the following **VLANs** nodes:

- The **LAN Cloud** > **VLANs** node for a VLAN accessible to both fabric interconnects.

- The *Fabric_Interconnect_Name* > **VLANs** node for a VLAN accessible to only one fabric interconnect.

Deleting a Named VLAN with the LAN Uplinks Manager

If Cisco UCS Manager includes a named VLAN with the same VLAN ID as the one you delete, the VLAN is not removed from the fabric interconnect configuration until all named VLANs with that ID are deleted.

Procedure

Step 1 In the LAN Uplinks Manager, click the **VLANs** tab.

Step 2 Click one of the following subtabs, based on the VLAN that you want to delete:

Subtab	Description
All	Displays all VLANs in the Cisco UCS domain.
Dual Mode	Displays the VLANs that are accessible to both fabric interconnects.
Fabric A	Displays the VLANs that are accessible to only fabric interconnect A.
Fabric B	Displays the VLANs that are accessible to only fabric interconnect B.

Step 3 In the table, click the VLAN you want to delete.
You can use the Shift key or Ctrl key to select multiple entries.

Step 4 Right-click the highlighted VLAN or VLANs and select **Delete**.

Step 5 If a confirmation dialog box displays, click **Yes**.

Configuring QoS System Classes with the LAN Uplinks Manager

The type of adapter in a server might limit the maximum MTU supported. For example, network MTU above the maximums might cause the packet to be dropped for the following adapters:

- The Cisco UCS M71KR CNA adapter, which supports a maximum MTU of 9216.
- The Cisco UCS 82598KR-CI adapter, which supports a maximum MTU of 14000.

Procedure

Step 1 In the LAN Uplinks Manager, click the **QoS** tab.

Step 2 Update the following properties for the system class you want to configure to meet the traffic management needs of the system:

Note Some properties may not be configurable for all system classes.

Name	Description
Enabled check box	If checked, the associated QoS class is configured on the fabric interconnect and can be assigned to a QoS policy. If unchecked, the class is not configured on the fabric interconnect and any QoS policies associated with this class default to **Best Effort** or, if a system class is configured with a Cos of 0, to the Cos 0 system class. **Note** This field is always checked for **Best Effort** and **Fibre Channel**.
CoS field	The class of service. You can enter an integer value between 0 and 6, with 0 being the lowest priority and 6 being the highest priority. We recommend that you do not set the value to 0, unless you want that system class to be the default system class for traffic if the QoS policy is deleted or the assigned system class is disabled. **Note** This field is set to 7 for internal traffic and to **any** for **Best Effort**. Both of these values are reserved and cannot be assigned to any other priority.
Packet Drop check box	If checked, packet drop is allowed for this class. If unchecked, packets cannot be dropped during transmission. This field is always unchecked for the **Fibre Channel** class, which never allows dropped packets, and always checked for **Best Effort**, which always allows dropped packets. **Note** When you save changes to the Packet Drop, the following warning message displays: You are making changes to the QOS system class, which may cause momentary disruption to traffic forwarding. Are you sure you want to apply the changes?
Weight drop-down list	This can be one of the following: • An integer between 1 and 10. If you enter an integer, Cisco UCS determines the percentage of network bandwidth assigned to the priority level as described in the **Weight (%)** field. • **best-effort**. • **none**.
Weight (%) field	To determine the bandwidth allocated to a channel, Cisco UCS: **1** Adds the weights for all the channels **2** Divides the channel weight by the sum of all weights to get a percentage **3** Allocates that percentage of the bandwidth to the channel

Name	Description
MTU drop-down list	The maximum transmission unit for the channel. This can be one of the following: • An integer between 1500 and 9216. This value corresponds to the maximum packet size. **Note** When you save changes to the MTU, the following warning message displays: You are making changes to the QOS system class, which may cause momentary disruption to traffic forwarding. Are you sure you want to apply the changes? • **fc—**A predefined packet size of 2240. • **normal—**A predefined packet size of 1500. **Note** This field is always set to **fc** for **Fibre Channel**.
Multicast Optimized check box	If checked, the class is optimized to send packets to multiple destinations simultaneously. **Note** This option is not applicable to the **Fibre Channel**.

Step 3 Do one of the following:

- Click **OK** to save your changes and exit from the LAN Uplinks Manager.
- Click **Apply** to save your changes without exiting from the LAN Uplinks Manager.

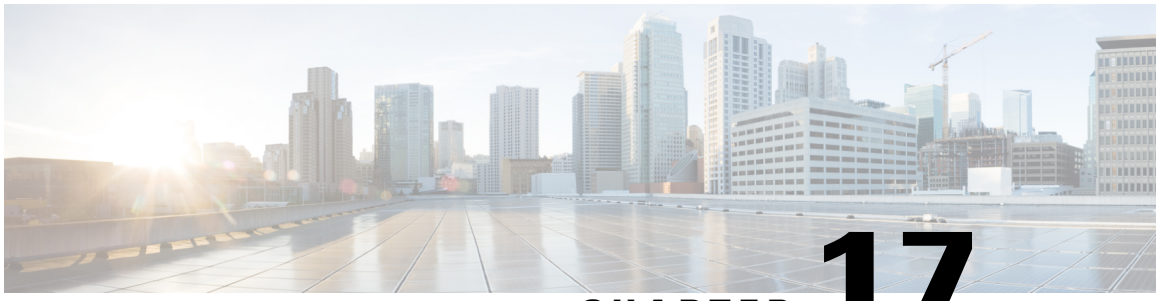

CHAPTER 17

VLANs

About VLANs

A VLAN is a switched network that is logically segmented by function, project team, or application, without regard to the physical locations of the users. VLANs have the same attributes as physical LANs, but you can group end stations even if they are not physically located on the same LAN segment.

Any switch port can belong to a VLAN. Unicast, broadcast, and multicast packets are forwarded and flooded only to end stations in the VLAN. Each VLAN is considered a logical network, and packets destined for stations that do not belong to the VLAN must be forwarded through a router or bridge.

VLANs are typically associated with IP subnetworks. For example, all of the end stations in a particular IP subnet belong to the same VLAN. To communicate between VLANs, you must route the traffic. By default, a newly created VLAN is operational. Additionally, you can configure VLANs to be in the active state, which

is passing traffic, or in the suspended state, in which the VLANs are not passing packets. By default, the VLANs are in the active state and pass traffic.

You can use the Cisco UCS Manager to manage VLANS. You can do the following:

- Configure named VLANs and Private VLANs (PVLANs).
- Assign VLANS to an access or trunk port.
- Create, delete and modify VLANs.

Guidelines for Creating, Deleting, and Modifying VLANs

VLANs are numbered from 1 to 4094. All configured ports belong to the default VLAN when you first bring up a switch. The default VLAN (VLAN1) uses only default values. You cannot create, delete, or suspend activity in the default VLAN.

You configure a VLAN by assigning a number to it. You can delete VLANs or move them from the active operational state to the suspended operational state. If you attempt to create a VLAN with an existing VLAN ID, the switch goes into the VLAN submode, but does not create the same VLAN again. Newly created VLANs remain unused until you assign ports to the specific VLAN. All of the ports are assigned to VLAN1 by default. Depending on the range of the VLAN, you can configure the following parameters for VLANs (except for the default VLAN):

- VLAN name
- Shutdown or not shutdown

When you delete a specified VLAN, the ports associated to that VLAN are shut down and no traffic flows. However, the system retains all of the VLAN-to-port mappings for that VLAN. When you re-enable or recreate the specified VLAN, the system automatically re-instates all of the original ports to that VLAN.

About the Native VLAN

The native VLAN and the default VLAN are not the same. Native refers to VLAN traffic without an 802.1q header and can be assigned or not. The native VLAN is the only VLAN that is not tagged in a trunk, and the frames are transmitted unchanged.

You can tag everything and not use a native VLAN throughout your network, and the VLAN or devices are reachable because switches use VLAN 1 as the native by default.

The UCS Manager LAN Uplink Manager enables you to configure VLANs and to change the native VLAN setting. Changing the native VLAN setting requires a port flap for the change to take effect; otherwise, the port flap is continuous. When you change the native VLAN, there is a loss of connectivity for approximately 20-40 seconds.

Native VLAN Guidelines

- You can only configure native VLANs on trunk ports.
- You can change the native VLAN on a UCS vNIC; however, the port flaps and can lead to traffic interruptions.

- Cisco recommends using the native VLAN 1 setting to prevent traffic interruptions if using the Cisco Nexus 1000v switches. The native VLAN must be the same for the Nexus 1000v port profiles and your UCS vNIC definition.
- If the native VLAN 1 setting is configured, and traffic routes to an incorrect interface, there is an outage, or the switch interface flaps continuously, there might be incorrect settings in your disjoint layer 2 network configuration.
- Using the native VLAN 1 for management access to all of your devices can potentially cause problems if someone connects another switch on the same VLAN as your management devices.

About the Access and Trunk Ports

Access Ports on a Cisco Switch

Access ports only sends untagged frames and belongs to and carries the traffic of only one VLAN. Traffic is received and sent in native formats with no VLAN tagging. Anything arriving on an access port is assumed to belong to the VLAN assigned to the port.

You can configure a port in access mode and specify the VLAN to carry the traffic for that interface. If you do not configure the VLAN for a port in access mode, or an access port, the interface carries the traffic for the default VLAN, which is VLAN 1. You can change the access port membership in a VLAN by configuring the VLAN. You must create the VLAN before you can assign it as an access VLAN for an access port. If you change the access VLAN on an access port to a VLAN that is not yet created, the UCS Manager shuts down that access port.

If an access port receives a packet with an 802.1Q tag in the header other than the access VLAN value, that port drops the packet without learning its MAC source address. If you assign an access VLAN that is also a primary VLAN for a private VLAN, all access ports with that access VLAN receives all the broadcast traffic for the primary VLAN in the private VLAN mode.

Trunk Ports on a Cisco Switch

Trunk ports allow multiple VLANs to transport between switches over that trunk link. A trunk port can carry untagged packets simultaneously with the 802.1Q tagged packets. When you assign a default port VLAN ID to the trunk port, all untagged traffic travels on the default port VLAN ID for the trunk port, and all untagged traffic is assumed to belong to this VLAN. This VLAN is referred to as the native VLAN ID for a trunk port. The native VLAN ID is the VLAN that carries untagged traffic on trunk ports.

The trunk port sends an egressing packet with a VLAN that is equal to the default port VLAN ID as untagged; all the other egressing packets are tagged by the trunk port. If you do not configure a native VLAN ID, the trunk port uses the default VLAN.

Note Changing the native VLAN on a trunk port, or an access VLAN of an access port flaps the switch interface.

Named VLANs

A named VLAN creates a connection to a specific external LAN. The VLAN isolates traffic to that external LAN, including broadcast traffic.

The name that you assign to a VLAN ID adds a layer of abstraction that allows you to globally update all servers associated with service profiles that use the named VLAN. You do not need to reconfigure the servers individually to maintain communication with the external LAN.

You can create more than one named VLAN with the same VLAN ID. For example, if servers that host business services for HR and Finance need to access the same external LAN, you can create VLANs named HR and Finance with the same VLAN ID. Then, if the network is reconfigured and Finance is assigned to a different LAN, you only have to change the VLAN ID for the named VLAN for Finance.

In a cluster configuration, you can configure a named VLAN to be accessible only to one fabric interconnect or to both fabric interconnects.

Guidelines for VLAN IDs

Important You cannot create VLANs with IDs from 4030 to 4047. This range of VLAN IDs is reserved.

The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch.

VLANs in the LAN cloud and FCoE VLANs in the SAN cloud must have different IDs. Using the same ID for a VLAN and an FCoE VLAN in a VSAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

VLAN 4048 is user configurable. However, Cisco UCS Manager uses VLAN 4048 for the following default values. If you want to assign 4048 to a VLAN, you must reconfigure these values:

- After an upgrade to Cisco UCS, Release 2.0—The FCoE storage port native VLAN uses VLAN 4048 by default. If the default FCoE VSAN was set to use VLAN 1 before the upgrade, you must change it to a VLAN ID that is not used or reserved. For example, consider changing the default to 4049 if that VLAN ID is not in use.
- After a fresh install of Cisco UCS, Release 2.0—The FCoE VLAN for the default VSAN uses VLAN 4048 by default. The FCoE storage port native VLAN uses VLAN 4049.

The VLAN name is case sensitive.

Private VLANs

A private VLAN (PVLAN) partitions the Ethernet broadcast domain of a VLAN into subdomains, and allows you to isolate some ports. Each subdomain in a PVLAN includes a primary VLAN and one or more secondary VLANs. All secondary VLANs in a PVLAN must share the same primary VLAN. The secondary VLAN ID differentiates one subdomain from another.

Isolated and Community VLANs

All secondary VLANs in a Cisco UCS domain can be Isolated or Community VLANs.

Note You cannot configure an isolated VLAN to use with a regular VLAN.

Ports on Isolated VLANs

Communications on an isolated VLAN can only use the associated port in the primary VLAN. These ports are isolated ports and are not configurable in Cisco UCS Manager. A primary VLAN can have only one isolated VLAN, but multiple isolated ports on the same isolated VLAN are allowed. These isolated ports cannot communicate with each other. The isolated ports can communicate only with a regular trunk port or promiscuous port that allows the isolated VLAN.

An isolated port is a host port that belongs to an isolated secondary VLAN. This port has complete isolation from other ports within the same private VLAN domain. PVLANs block all traffic to isolated ports except traffic from promiscuous ports. Traffic received from an isolated port is forwarded only to promiscuous ports. You can have more than one isolated port in a specified isolated VLAN. Each port is completely isolated from all other ports in the isolated VLAN.

Guidelines for Uplink Ports

When you create PVLANs, use the following guidelines:

- The uplink Ethernet port channel cannot be in promiscuous mode.
- Each primary VLAN can have only one isolated VLAN.
- VIFs on VNTAG adapters can have only one isolated VLAN.

Guidelines for VLAN IDs

Important You cannot create VLANs with IDs from 4030 to 4047. This range of VLAN IDs is reserved.

The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch.

VLANs in the LAN cloud and FCoE VLANs in the SAN cloud must have different IDs. Using the same ID for a VLAN and an FCoE VLAN in a VSAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

VLAN 4048 is user configurable. However, Cisco UCS Manager uses VLAN 4048 for the following default values. If you want to assign 4048 to a VLAN, you must reconfigure these values:

- After an upgrade to Cisco UCS, Release 2.0—The FCoE storage port native VLAN uses VLAN 4048 by default. If the default FCoE VSAN was set to use VLAN 1 before the upgrade, you must change it to a VLAN ID that is not used or reserved. For example, consider changing the default to 4049 if that VLAN ID is not in use.
- After a fresh install of Cisco UCS, Release 2.0—The FCoE VLAN for the default VSAN uses VLAN 4048 by default. The FCoE storage port native VLAN uses VLAN 4049.

The VLAN name is case sensitive.

VLAN Port Limitations

Cisco UCS Manager limits the number of VLAN port instances that you can configure under border and server domains on a fabric interconnect.

Types of Ports Included in the VLAN Port Count

The following types of ports are counted in the VLAN port calculation:

- Border uplink Ethernet ports
- Border uplink Ether-channel member ports
- FCoE ports in a SAN cloud
- Ethernet ports in a NAS cloud
- Static and dynamic vNICs created through service profiles
- VM vNICs created as part of a port profile in a hypervisor in hypervisor domain

Based on the number of VLANs configured for these ports, Cisco UCS Manager tracks the cumulative count of VLAN port instances and enforces the VLAN port limit during validation. Cisco UCS Manager reserves some pre-defined VLAN port resources for control traffic. These include management VLANs configured under HIF and NIF ports.

VLAN Port Limit Enforcement

Cisco UCS Manager validates VLAN port availability during the following operations:

- Configuring and unconfiguring border ports and border port channels
- Adding or removing VLANs from a cloud
- Configuring or unconfiguring SAN or NAS ports
- Associating or disassociating service profiles that contain configuration changes
- Configuring or unconfiguring VLANs under vNICs or vHBAs
- Receiving creation or deletion notifications from a VMWare vNIC and from an ESX hypervisor

Note This is outside the control of the Cisco UCS Manager.

- Fabric interconnect reboot
- Cisco UCS Manager upgrade or downgrade

Cisco UCS Manager strictly enforces the VLAN port limit on service profile operations. If Cisco UCS Manager detects that the VLAN port limit is exceeded, the service profile configuration fails during deployment.

Exceeding the VLAN port count in a border domain is less disruptive. When the VLAN port count is exceeded in a border domain Cisco UCS Manager changes the allocation status to Exceeded. To change the status back to **Available**, complete one of the following actions:

- Unconfigure one or more border ports
- Remove VLANs from the LAN cloud
- Unconfigure one or more vNICs or vHBAs

Configuring Named VLANs

Creating a Named VLAN

In a Cisco UCS domain that is configured for high availability, you can create a named VLAN that is accessible to both fabric interconnects or to only one fabric interconnect.

Important You cannot create VLANs with IDs from 4030 to 4047. This range of VLAN IDs is reserved.

The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch.

VLANs in the LAN cloud and FCoE VLANs in the SAN cloud must have different IDs. Using the same ID for a VLAN and an FCoE VLAN in a VSAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 On the **LAN** tab, click the **LAN** node.

Step 3 In the **Work** pane, click the **VLANs** tab.

Step 4 On the icon bar to the right of the table, click +.
If the + icon is disabled, click an entry in the table to enable it.

Step 5 In the **Create VLANs** dialog box, complete the required fields.

Step 6 If you clicked the **Check Overlap** button, do the following:

a) Click the **Overlapping VLANs** tab and review the fields to verify that the VLAN ID does not overlap with any IDs assigned to existing VLANs.
b) Click the **Overlapping VSANs** tab and review the fields to verify that the VLAN ID does not overlap with any FCoE VLAN IDs assigned to existing VSANs.
c) Click **OK**.
d) If Cisco UCS Manager identified any overlapping VLAN IDs or FCoE VLAN IDs, change the VLAN ID to one that does not overlap with an existing VLAN.

Step 7 Click **OK**.
Cisco UCS Manager adds the VLAN to one of the following **VLANs** nodes:

- The **LAN Cloud** > **VLANs** node for a VLAN accessible to both fabric interconnects.

- The *Fabric_Interconnect_Name* > **VLANs** node for a VLAN accessible to only one fabric interconnect.

Deleting a Named VLAN

If Cisco UCS Manager includes a named VLAN with the same VLAN ID as the one you delete, the VLAN is not removed from the fabric interconnect configuration until all named VLANs with that ID are deleted.

If you are deleting a private primary VLAN, ensure that you reassign the secondary VLANs to another working primary VLAN.

Before You Begin

Before you delete a VLAN from a fabric interconnect, ensure that the VLAN was removed from all vNICs and vNIC templates.

Note If you delete a VLAN that is assigned to a vNIC or vNIC template, the vNIC might allow that VLAN to flap.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 On the **LAN** tab, click the **LAN** node.

Step 3 In the **Work** pane, click the **VLANs** tab.

Step 4 Click one of the following subtabs, based on the VLAN that you want to delete:

Subtab	Description
All	Displays all VLANs in the Cisco UCS domain.
Dual Mode	Displays the VLANs that are accessible to both fabric interconnects.
Fabric A	Displays the VLANs that are accessible to only fabric interconnect A.
Fabric B	Displays the VLANs that are accessible to only fabric interconnect B.

Step 5 In the table, click the VLAN that you want to delete.
You can use the Shift key or Ctrl key to select multiple entries.

Step 6 Right-click the highlighted VLAN or VLANs and click **Delete**.

Step 7 If a confirmation dialog box displays, click **Yes**.

Configuring Private VLANs

Creating a Primary VLAN for a Private VLAN

In a Cisco UCS domain that is configured for high availability, you can create a primary VLAN that is accessible to both fabric interconnects or to only one fabric interconnect.

Important You cannot create VLANs with IDs from 4030 to 4047. This range of VLAN IDs is reserved.

The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch.

VLANs in the LAN cloud and FCoE VLANs in the SAN cloud must have different IDs. Using the same ID for a VLAN and an FCoE VLAN in a VSAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 On the **LAN** tab, click the **LAN** node.

Step 3 In the **Work** pane, click the **VLANs** tab.

Step 4 On the icon bar to the right of the table, click +.
If the + icon is disabled, click an entry in the table to enable it.

Step 5 In the **Create VLANs** dialog box, complete the required fields.

Step 6 If you clicked the **Check Overlap** button, do the following:

a) Click the **Overlapping VLANs** tab and review the fields to verify that the VLAN ID does not overlap with any IDs assigned to existing VLANs.
b) Click the **Overlapping VSANs** tab and review the fields to verify that the VLAN ID does not overlap with any FCoE VLAN IDs assigned to existing VSANs.
c) Click **OK**.
d) If Cisco UCS Manager identified any overlapping VLAN IDs or FCoE VLAN IDs, change the VLAN ID to one that does not overlap with an existing VLAN.

Step 7 Click **OK**.
Cisco UCS Manager adds the primary VLAN to one of the following **VLANs** nodes:

- The **LAN Cloud** > **VLANs** node for a primary VLAN accessible to both fabric interconnects.
- The ***Fabric_Interconnect_Name*** > **VLANs** node for a primary VLAN accessible to only one fabric interconnect.

Creating a Secondary VLAN for a Private VLAN

In a Cisco UCS domain that is configured for high availability, you can create a secondary VLAN that is accessible to both fabric interconnects or to only one fabric interconnect.

Important You cannot create VLANs with IDs from 4030 to 4047. This range of VLAN IDs is reserved.

The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch.

VLANs in the LAN cloud and FCoE VLANs in the SAN cloud must have different IDs. Using the same ID for a VLAN and an FCoE VLAN in a VSAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Before You Begin

Create the primary VLAN.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 On the **LAN** tab, click the **LAN** node.

Step 3 In the **Work** pane, click the **VLANs** tab.

Step 4 On the icon bar to the right of the table, click +.
If the + icon is disabled, click an entry in the table to enable it.

Step 5 In the **Create VLANs** dialog box, specify the required fields.
Note The multicast policy is associated to the primary VLAN, not the secondary VLAN.

Step 6 If you clicked the **Check Overlap** button, do the following:

a) Click the **Overlapping VLANs** tab and review the fields to verify that the VLAN ID does not overlap with any IDs assigned to existing VLANs.
b) Click the **Overlapping VSANs** tab and review the fields to verify that the VLAN ID does not overlap with any FCoE VLAN IDs assigned to existing VSANs.
c) Click **OK**.
d) If Cisco UCS Manager identified any overlapping VLAN IDs or FCoE VLAN IDs, change the VLAN ID to one that does not overlap with an existing VLAN.

Step 7 Click **OK**.
Cisco UCS Manager adds the primary VLAN to one of the following **VLANs** nodes:

- The **LAN Cloud** > **VLANs** node for a primary VLAN accessible to both fabric interconnects.
- The ***Fabric_Interconnect_Name*** > **VLANs** node for a primary VLAN accessible to only one fabric interconnect.

Community VLANs

Cisco UCS Manager supports Community VLANs in UCS Fabric Interconnects. Community ports communicate with each other and with promiscuous ports. Community ports have Layer 2 isolation from all other ports in other communities, or isolated ports within the PVLAN. Broadcasts are transmitted between the community ports associated with the PVLAN only and the other promiscuous ports. A promiscuous port can communicate with all interfaces, including the isolated and community ports within a PVLAN.

Creating a Community VLAN

In a Cisco UCS domain configured for high availability, you can create a Community VLAN accessible to both fabric interconnects or to only one fabric interconnect.

Important You cannot create VLANs with IDs from 4030 to 4047. This range of VLAN IDs is reserved.

The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch.

VLANs in the LAN cloud and FCoE VLANs in the SAN cloud must have different IDs. Using the same ID for a VLAN and an FCoE VLAN in a VSAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 On the **LAN** tab, click the **LAN** node.

Step 3 In the **Work** pane, click the **VLANs** tab.

Step 4 On the icon bar to the right of the table, click +.
If the + icon is disabled, click an entry in the table to enable it.

Step 5 In the **Create VLANs** dialog box, complete the following fields:

Name	Description
VLAN Name/Prefix field	For a single VLAN, this is the VLAN name. For a range of VLANs, this is the prefix that the system uses for each VLAN name. The VLAN name is case sensitive. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Multicast Policy drop-down list	The multicast policy associated with this VLAN.

Name	Description
Create Multicast Policy link	Click this link to create a new multicast policy that will be available to all VLANs.
Configuration options	You can choose one of the following: • **Common/Global—**The VLANs apply to both fabrics and use the same configuration parameters in both cases. • **Fabric A—**The VLANs only apply to fabric A. • **Fabric B—**The VLAN only apply to fabric B. • **Both Fabrics Configured Differently—**The VLANs apply to both fabrics, but you can specify different VLAN IDs for each fabric. For upstream disjoint L2 networks, Cisco recommends that you choose **Common/Global** to create VLANs that apply to both fabrics.
VLAN IDs field	To create one VLAN, enter a single numeric ID. To create multiple VLANs, enter individual IDs or ranges of IDs separated by commas. A VLAN ID can: • Be between 1 and 3967 • Be between 4048 and 4093 • Overlap with other VLAN IDs already defined on the system For example, to create six VLANs with the IDs 4, 22, 40, 41, 42, and 43, enter 4, 22, 40-43. **Important** You cannot create VLANs with IDs from 4030 to 4047. This range of VLAN IDs is reserved. The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch. VLANs in the LAN cloud and FCoE VLANs in the SAN cloud must have different IDs. Using the same ID for a VLAN and an FCoE VLAN in a VSAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Name	Description
Sharing Type field	Whether this VLAN is subdivided into private or secondary VLANs. This can be one of the following: • **None—**This VLAN does not have any secondary or private VLANs. • **Primary—**This VLAN can have one or more secondary VLANs, as shown in the **Secondary VLANs** area. • **Isolated—**This is a private VLAN. The primary VLAN with which it is associated is shown in the **Primary VLAN** drop-down list. • **Community—**This VLAN can communicate with other ports on the same PVLAN as well as the promiscuous port. To create a community VLAN, the sharing type should be set to **Community**.
Primary VLAN drop-down list	If the **Sharing Type** field is set to **Isolated** or **Community**, this is the primary VLAN associated with the Isolated or CommunityVLAN.
Permitted Orgs for VLAN(s)	Select the organization from the list for the VLAN. This VLAN will be available for the organizations that you select.
Check Overlap button	Click this button to determine whether the VLAN ID overlaps with any other IDs on the system.

Step 6 If you clicked the **Check Overlap** button, do the following:

a) Click the **Overlapping VLANs** tab and review the following fields to verify that the VLAN ID does not overlap with any IDs assigned to existing VLANs.

Name	Description
Fabric ID column	This can be one of the following: • **A** • **B** • **Dual—**The component is accessible to either fabric interconnect. This setting applies to virtual LAN and SAN networks created at the system level as opposed to the fabric-interconnect level.
Name column	The name of the VLAN.
VLAN column	The numeric id for the VLAN.
DN column	The full path to the VLAN. Click the link in this column to view the properties for the VLAN.

b) Click the **Overlapping VSANs** tab and review the following fields to verify that the VLAN ID does not overlap with any FCoE VLAN IDs assigned to existing VSANs:

Name	Description
Fabric ID column	This can be one of the following: • **A** • **B** • **Dual—**The component is accessible to either fabric interconnect. This setting applies to virtual LAN and SAN networks created at the system level as opposed to the fabric-interconnect level.
Name column	The name of the VSAN.
ID column	The numeric id for the VSAN.
FCoE VLAN ID column	The unique identifier assigned to the VLAN used for Fibre Channel connections.
DN column	The full path to the VSAN. Click the link in this column to view the properties for the VSAN.

c) Click **OK**.

d) If Cisco UCS Manager identified any overlapping VLAN IDs or FCoE VLAN IDs, change the VLAN ID to one that does not overlap with an existing VLAN.

Step 7 Click **OK**.

Cisco UCS Manager adds the Community VLAN to one of the following **VLANs** nodes:

- The **LAN Cloud** > **VLANs** node for a VLAN accessible to both fabric interconnects.
- The ***Fabric_Interconnect_Name*** > **VLANs** node for a VLAN accessible to only one fabric interconnect.

Creating Promiscuous Access on Appliance Port

Cisco UCS Manager supports Promiscuous access on appliance ports. The following procedure details the configurations steps.

Before You Begin

Create the PVLANs in Appliance Cloud.

Procedure

Step 1 In the **Navigation** pane, click the **LAN** tab.

Step 2 On the **LAN** tab, expand **LAN** > **Appliances** > ***Fabric*** > **Interfaces**.
The Interfaces pane displays.

Step 3 In the **Interfaces pane**on the icon bar to the right of the table, click + .
The **Appliance Links** pane displays.

Step 4 In the **Appliance Links** pane, click the **Unconfigured Ethernet Ports** to expand the **Unconfigured Ethernet Ports**.
All available Unconfigured Ethernet Ports display.

Step 5 Click the **Unconfigured Ethernet Ports** that you want to make an Appliance Port.

Step 6 Click **Make Appliance Port**.
The **Configure as Appliance Port** confirmation box displays.

Step 7 Click **Yes** to configure the appliance port.
The **Configure Appliance Port** dialog box opens.

Step 8 On the **LAN** tab, expand **LAN** > **Appliances** > ***Fabric*** > **Interfaces**.

Step 9 Expand **Appliance Ports**.

Step 10 Click the appliance port for which you want to modify the properties.

Step 11 In the **Interfaces pane**on the icon bar to the right of the table, click the **Modify** icon.
The **Properties for Appliance Interface** dialog box displays.

Step 12 In the **VLANs** pane, click the **Access** radio button.

Step 13 Select a Primary VLAN from the **Select VLAN** drop-down list to assign to the appliance port.
A list of secondary VLANs associated with the primary VLAN displays.

Step 14 Select a set of secondary VLANs allowed on the port.
Selecting an **Isolated** or **Community** VLAN turns the **VLAN** into a **Promiscuous Port**. If you select the Primary VLAN from the **Select VLAN** drop-down list, you must select the required secondary VLAN.

Step 15 Click **Apply** to configure **Promiscuous Access on Appliance Port**.

Creating a Promiscuous Trunk on Appliance Port

Cisco UCS Manager supports Promiscuous Trunks on appliance ports. The following procedure details the configurations steps.

Before You Begin

Create the Private VLANs in the Appliance Cloud.

Procedure

Step 1 In the **Navigation** pane, click the **LAN** tab.

Step 2 On the **LAN** tab, expand **LAN** > **Appliances** > ***Fabric*** > **Interfaces**.

The Interfaces pane displays.

Step 3 In the **Interfaces pane**on the icon bar to the right of the table, click + .
The **Appliance Links** pane displays.

Step 4 In the **Appliance Links** pane, click the **Unconfigured Ethernet Ports** to expand the **Unconfigured Ethernet Ports**.
All available Unconfigured Ethernet Ports display.

Step 5 Click the **Unconfigured Ethernet Ports** that you to want make an Appliance Port.

Step 6 Click **Make Appliance Port**.
The Configure as Appliance Port confirmation box displays.

Step 7 Click **Yes** to configure the appliance port.

Step 8 On the **LAN** tab, expand **LAN** > **Appliances** > ***Fabric*** > **Interfaces**.

Step 9 Expand **Appliance Ports**.

Step 10 Click the appliance port for which you want to modify the properties.

Step 11 In the **Interfaces pane** on the icon bar to the right of the table, click the **Modify** icon.
The **Properties for Appliance Interface** dialog box displays.

Step 12 In the **VLANs** pane, click the **Trunk** radio button.

Step 13 Select a **VLAN** from the available VLANs.
From the list of VLANs, you can select multiple**Isolated**,**Community**, **Primary** and **Regular** VLANs to apply on the port to make it a promiscuous trunk port.

Step 14 Click **Apply** to configure **Promiscuous on Trunk on Appliance Port**.

Allowing Private VLANs on vNICs - Community Access Mode

In a Cisco UCS domain that is configured for high availability, you can create a primary VLAN that is accessible to both fabric interconnects or to only one fabric interconnect.

Important You cannot create VLANs with IDs from 4030 to 4047. This range of VLAN IDs is reserved.

The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch.

VLANs in the LAN cloud and FCoE VLANs in the SAN cloud must have different IDs. Using the same ID for a VLAN and an FCoE VLAN in a VSAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 On the **LAN** tab, click the **LAN** node.

Step 3 In the **Work** pane, click the **VLANs** tab.

Step 4 On the icon bar to the right of the table, click +.
If the + icon is disabled, click an entry in the table to enable it.

Step 5 In the **Create VLANs** dialog box, complete the following fields:

Name	Description
VLAN Name/Prefix field	For a single VLAN, this is the VLAN name. For a range of VLANs, this is the prefix that the system uses for each VLAN name. The VLAN name is case sensitive. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Multicast Policy drop-down list	The multicast policy associated with this VLAN.
Create Multicast Policy link	Click this link to create a new multicast policy that will be available to all VLANs.
Configuration options	You can choose one of the following: • **Common/Global—**The VLANs apply to both fabrics and use the same configuration parameters in both cases. • **Fabric A—**The VLANs only apply to fabric A. • **Fabric B—**The VLAN only apply to fabric B. • **Both Fabrics Configured Differently—**The VLANs apply to both fabrics but you can specify different VLAN IDs for each fabric. For upstream disjoint L2 networks, Cisco recommends that you choose **Common/Global** to create VLANs that apply to both fabrics.

Name	Description
VLAN IDs field	To create one VLAN, enter a single numeric ID. To create multiple VLANs, enter individual IDs or ranges of IDs separated by commas. A VLAN ID can: • Be between 1 and 3967 • Be between 4048 and 4093 • Overlap with other VLAN IDs already defined on the system For example, to create six VLANs with the IDs 4, 22, 40, 41, 42, and 43, you would enter 4, 22, 40-43. **Important** You cannot create VLANs with IDs from 4030 to 4047. This range of VLAN IDs is reserved. The VLAN IDs you specify must also be supported on the switch that you are using. For example, on Cisco Nexus 5000 Series switches, the VLAN ID range from 3968 to 4029 is reserved. Before you specify the VLAN IDs in Cisco UCS Manager, make sure that the same VLAN IDs are available on your switch. VLANs in the LAN cloud and FCoE VLANs in the SAN cloud must have different IDs. Using the same ID for a VLAN and an FCoE VLAN in a VSAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that VLAN. Ethernet traffic is dropped on any VLAN which has an ID that overlaps with an FCoE VLAN ID.
Sharing Type field	Whether this VLAN is subdivided into private or secondary VLANs. This can be one of the following: • **None—**This VLAN does not have any secondary or private VLANs. This is a regular VLAN and not a PVLAN . • **Primary—**This VLAN can have one or more secondary VLANs, as shown in the **Secondary VLANs** area. • **Isolated—**This is a private VLAN. The primary VLAN with which it is associated is shown in the **Primary VLAN** drop-down list. • **Community—**This VLAN can communicate with other ports on the same PVLAN as well as the promiscuous port. Select this sharing type for to configure a **Community VLAN**.
Primary VLAN drop-down list	If the **Sharing Type** field is set to **Isolated** or **Community**, this is the primary VLAN associated with the Isolated or CommunityVLAN.
Permitted Orgs for VLAN(s)	Select the organization from the list for the VLAN. This VLAN will be available for the organizations you select here.

Name	Description
Check Overlap button	Click this button to determine whether the VLAN ID overlaps with any other IDs on the system.

Step 6 If you clicked the **Check Overlap** button, do the following:

a) Click the **Overlapping VLANs** tab and review the following fields to verify that the VLAN ID does not overlap with any IDs assigned to existing VLANs.

Name	Description
Fabric ID column	This can be one of the following: • **A** • **B** • **Dual—**The component is accessible to either fabric interconnect. This setting applies to virtual LAN and SAN networks created at the system level as opposed to the fabric-interconnect level.
Name column	The name of the VLAN.
VLAN column	The numeric id for the VLAN.
DN column	The full path to the VLAN. Click the link in this column to view the properties for the VLAN.

b) Click the **Overlapping VSANs** tab and review the following fields to verify that the VLAN ID does not overlap with any FCoE VLAN IDs assigned to existing VSANs:

Name	Description
Fabric ID column	This can be one of the following: • **A** • **B** • **Dual—**The component is accessible to either fabric interconnect. This setting applies to virtual LAN and SAN networks created at the system level as opposed to the fabric-interconnect level.
Name column	The name of the VSAN.
ID column	The numeric id for the VSAN.
FCoE VLAN ID column	The unique identifier assigned to the VLAN used for Fibre Channel connections.

Name	Description
DN column	The full path to the VSAN. Click the link in this column to view the properties for the VSAN.

c) Click **OK**.
d) If Cisco UCS Manager identified any overlapping VLAN IDs or FCoE VLAN IDs, change the VLAN ID to one that does not overlap with an existing VLAN.

Step 7 Click **OK**.
Cisco UCS Manager adds the Community VLAN to one of the following **VLANs** nodes:

- The **LAN Cloud** > **VLANs** node for a VLAN accessible to both fabric interconnects.
- The ***Fabric_Interconnect_Name*** > **VLANs** node for a VLAN accessible to only one fabric interconnect.

Configuring a Access Mode for Community Server

The Cisco UCS domain provides support for Private VLANs on vNICs. Configuring a VLAN on vNICs with access mode allows the server to operate as an Community Access Server.

Procedure

Step 1 In the **Navigation** pane, click the **Servers** tab.
Step 2 On the **Servers** tab, expand **Servers** > **Service Profiles** > ***Service_Profile_Name***.
Step 3 On the **Service Profile** , select a**vNIC** you want to control.
vNICs Property page displays where you can modify **VLANS**.
Step 4 Click the **Modify VLANs link**.
Displays the list of available **VLANS**.
Step 5 Click one of the Community **VLANs** you previously created.
Step 6 Click**OK**.
The Community **VLAN** is now associated with the **vNIC**. When you apply a Community VLAN on a vNIC the server operates as a *Community Access Server*.

Viewing the VLAN Port Count

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Fabric Interconnects**.

Step 3 Click the fabric interconnect for which you want to view the VLAN port count.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **General** tab, click the down arrows on the **VLAN Port Count** bar to expand that area.
Cisco UCS Manager GUI displays the following details:

Name	Description
VLAN Port Limit field	The maximum number of VLAN ports allowed on this fabric interconnect.
Access VLAN Port Count field	The number of available VLAN access ports.
Border VLAN Port Count field	The number of available VLAN border ports.
Allocation Status field	The VLAN port allocation status.

VLAN Port Count Optimization

VLAN port count optimization enables mapping the state of multiple VLANs into a single internal state. When you enable the VLAN port count optimization, Cisco UCS Manager logically groups VLANs based on the port VLAN membership. This grouping increases the port VLAN count limit. VLAN port count optimization also compresses the VLAN state and reduces the CPU load on the fabric interconnect. This reduction in the CPU load enables you to deploy more VLANs over more vNICs. Optimizing VLAN port count does not change any of the existing VLAN configuration on the vNICs.

VLAN port count optimization is disabled by default. You can enable or disable the option based on your requirements.

Important

- Enabling VLAN port count optimization increases the number of available VLAN ports for use. If the port VLAN count exceeds the maximum number of VLANs in a non-optimized state, you cannot disable the VLAN port count optimization.
- VLAN port count optimization is not supported in Cisco UCS 6100 Series fabric interconnect.

Enabling Port VLAN Count Optimization

By default, the port VLAN count optimization is disabled. You can enable the port VLAN count optimization to optimize the CPU usage and to increase the port VLAN count.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **LAN Cloud**.

Step 3 In the **Work** pane, click the **Global Policies** tab.

Step 4 In the **Port, VLAN Count Optimization** section, choose **Enabled**.

Step 5 Click **Save Changes**.

Step 6 If the **Port, VLAN Count Optimization** option is successfully enabled, a confirmation message displays. Click **OK** to close the dialog box.

Disabling Port VLAN Count Optimization

By default, the port VLAN count optimization is disabled. You can disable the port VLAN count optimization option if you enabled it to increase the port VLAN count and to optimize the CPU usage.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **LAN Cloud**.

Step 3 In the **Work** pane, click the **Global Policies** tab.

Step 4 In the **Port, VLAN Count Optimization** section, choose **Disabled**.

Step 5 Click **Save Changes**.

Step 6 If the **Port, VLAN Count Optimization** option is successfully disabled, a confirmation message displays. Click **OK** to close the dialog box.

Viewing VLAN Optimization Sets

Cisco UCS Manager automatically creates VLAN port count optimization groups based on the VLAN IDs in the system. All of the VLANs in the group share the same IGMP policy. The following VLANs are not included in the VLAN port count optimization group:

- FCoE VLANs
- Primary PVLANs and secondary PVLANs
- VLANs that are specified as a SPAN source

- VLANs configured as a single allowed VLAN on an interface and port profiles with a single VLAN

Cisco UCS Manager GUI automatically groups the optimized VLANs.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **LAN Cloud**.

Step 3 In the **Navigation** pane, click **Fabric A** or **Fabric B** to expand the list.

Step 4 Click **VLAN Optimization Sets**.
The **Work** pane displays the list of VLAN optimization groups with **Name** and **Size**.

VLAN Groups

VLAN groups allow you to group VLANs on Ethernet uplink ports, by function or by VLANs that belong to a specific network. You can define VLAN membership and apply the membership to multiple Ethernet uplink ports on the fabric interconnect.

Note Cisco UCS Manager supports a maximum of 200 VLAN Groups. If Cisco UCS Manager determines that you create more than 200 VLAN groups, the system disables VLAN compression.

You can configure inband and out-of-band (OOB) VLAN groups to use to access the Cisco Integrated Management Interface (CIMC) on blade and rack servers. Cisco UCS Manager supports OOB IPv4 and inband IPv4 and IPv6 VLAN groups for use with the uplink interfaces or uplink port channels.

After you assign a VLAN to a VLAN group, any changes to the VLAN group are applied to all Ethernet uplink ports that are configured with the VLAN group. The VLAN group also enables you to identify VLAN overlaps between disjoint VLANs.

You can configure uplink ports under a VLAN group. When you configure an uplink port for a VLAN group, that uplink port will support all the VLANs that are part of the associated VLAN groups and individual VLANs that are associated with the uplink using LAN Uplinks Manager, if any. Further, any uplink that is not selected for association with that VLAN group will stop supporting the VLANs that are part of that VLAN group.

You can create VLAN groups from the **LAN Cloud** or from the **LAN Uplinks Manager**.

Creating a VLAN Group

You can create a **VLAN Group** from **LAN Cloud** or the **LAN Uplinks Manager**. This procedure explains creating a VLAN group from the **LAN Cloud**. You can create separate VLAN groups to use for inband and out-of-band access using service profiles.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **LAN Cloud**.

Step 3 Right-click **LAN Cloud** and choose **Create VLAN Group** from the drop-down list.
The **Create VLAN Group** wizard launches.

Step 4 In the **Select VLANs** dialog box, specify the name and VLANs, then click **Next**.

Step 5 (Optional) In **Add Uplink Ports** dialog box, select the **Uplink Ports** from the list and add the ports to the **Selected Uplink Ports**, then click **Next**.

Step 6 (Optional) In **Add Port Channels** dialog box, select the **Port Channels**, and add the port channels to the **Selected Port Channels**, then click **Next**.

Step 7 (Optional) In the **Org Permissions** dialog box, select the appropriate groups from the list, then click **Next**.
The VLANs that belong to the group that you are creating can only access the groups that you select.

Step 8 Click **Finish**.
This VLAN group is added to the list of **VLAN Groups** under **LAN** > **LAN Cloud** > **VLAN Groups**.

Editing the Members of a VLAN Group

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **LAN Cloud**.

Step 3 In the **Navigation** pane, click **VLAN Groups** to expand the VLAN group list.

Step 4 From the list of VLAN groups, choose the VLAN group name to edit the group member VLANs.
You can use the Shift key or Ctrl key to select multiple entries.

Step 5 Right-click the highlighted VLAN group or VLAN groups and choose **Edit VLAN Group Members**.
The **Modify VLAN Group** *VLAN Group Name* dialog box opens.

Step 6 In the **Modify VLAN Group** *VLAN Group Name* dialog box, select the VLANs that you want to remove or add from the list and click **Next**.

Step 7 (Optional) In **Add Port Channels** pane, choose the **Port Channels**, and add them to the **Selected Port Channels**.

Step 8 (Optional) In the **Org Permissions** pane, choose the appropriate groups from the list.
The VLANs that belong to the group that you are creating can only access the groups that you select.

Step 9 Click **Finish**.

Step 10 This VLAN group is modified based on your selections.

Modifying the Organization Access Permissions for a VLAN Group

When you modify the organization access permissions for a VLAN group, the change in permissions applies to all VLANs that are in that VLAN group.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 On the **LAN** tab, expand **LAN** > **LAN Cloud** > **VLAN Group**, select *VLAN group name*.

Step 3 In the **Work** pane, click the **General** tab.

Step 4 In **Actions**, click **Modify VLAN Groups Org Permissions**.
The **Modify VLAN Groups Org Permissions** dialog box opens.

Step 5 In **Org Permissions**, do the following:

- To add organizations, select the organizations.
- To remove access permission from an organization, click to remove the selection.

Step 6 Click **OK**.

Deleting a VLAN Group

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **LAN Cloud**.

Step 3 In the **Navigation** pane, click **VLAN Groups** to expand the VLAN group list.

Step 4 From the displayed list of VLAN groups, choose the VLAN group name you want to delete.
You can use the Shift key or Ctrl key to select multiple entries.

Step 5 Right-click the highlighted VLAN group or VLAN groups and choose **Delete**.

Step 6 If a confirmation dialog box displays, click **Yes**.

VLAN Permissions

VLAN permissions restrict access to VLANs based on specified organizations and on the service profile organizations to which the VLANs belong. VLAN permissions also restrict the set of VLANs that you can assign to service profile vNICs. VLAN permissions is an optional feature and is disabled by default. You can enable or disable the feature based on your requirements. If you disable the feature, all of the VLANs are globally accessible to all organizations.

Note If you enable the org permission in **LAN** > **LAN Cloud** > **Global Policies** > **Org Permissions**, when you create a VLAN, the **Permitted Orgs for VLAN(s)** option displays in the **Create VLANs** dialog box. If you do not enable the **Org Permissions**, the **Permitted Orgs for VLAN(s)** option does not display.

Enabling the org permission allows you to specify the organizations for the VLAN. When you specify the organizations, the VLAN becomes available to that specific organization and all of the sub organizations below the structure. Users from other organizations cannot access this VLAN. You can also modify the VLAN permission anytime based on changes to your VLAN access requirements.

Caution When you assign the VLAN org permission to an organization at the root level, all sub organizations can access the VLANs. After assigning the org permission at the root level, and you change the permission for a VLAN that belongs to a sub organization, that VLAN becomes unavailable to the root level organization.

Enabling VLAN Permissions

By default, VLAN permissions is disabled. If you want to restrict VLAN access by creating permissions for different organizations, you must enable the org permission option.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **LAN Cloud**.

Step 3 In the **Work** pane, click the **Global Policies** tab.

Step 4 In the **Org Permissions** section, choose **Enabled**.

Step 5 Click **Save Changes**.

Step 6 If the **Org Permissions** option is successfully enabled, a confirmation message displays. Click **OK** to close the dialog box.

Disabling VLAN Permissions

By default, VLAN permissions is disabled. You can enable VLAN permissions and assign a VLAN to a different network group or organization. You can also disable the VLAN permission globally; however, the permissions assigned to the VLANs continue to exist in the system, but are not enforced. If you want to use the org permissions later, you can enable the feature to use the assigned permissions.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **LAN Cloud**.

Step 3 In the **Work** pane, click the **Global Policies** tab.

Step 4 In the **Org Permissions** section, choose **Disabled**.

Step 5 Click **Save Changes**.

Step 6 If the **Org Permissions** option is successfully disabled, a confirmation message displays. Click **OK** to close the dialog box.

Adding or Modifying VLAN Permissions

You can add or delete the permitted organization for a VLAN.

Note When you add an organization as a permitted organization for a VLAN, all of the descendant organizations can access the VLAN. When you remove the permission to access a VLAN from an organization, the descendant organizations no longer have access to the VLAN.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 On the **LAN** tab, expand **LAN** > **LAN Cloud** > **VLANs**, select *VLAN name*.

Step 3 In the **Work** pane, click the **General** tab.

Step 4 In **Actions**, click **Modify VLAN Org Permissions**.
The **Modify VLAN Org Permissions** dialog box opens.

Step 5 In **Permitted Orgs for VLAN(s)**,

- To add organizations, select the organizations.
- To remove access permission from an organization, click to remove the selection.

Step 6 Click **OK**.

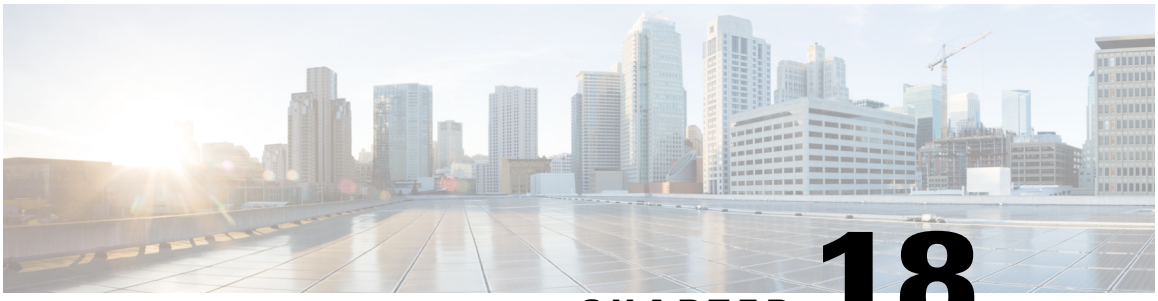

CHAPTER 18

Configuring LAN Pin Groups

This chapter includes the following sections:

LAN Pin Groups

Cisco UCS uses LAN pin groups to pin Ethernet traffic from a vNIC on a server to an uplink Ethernet port or port channel on the fabric interconnect. You can use this pinning to manage the distribution of traffic from the servers.

To configure pinning for a server, you must include the LAN pin group in a vNIC policy. The vNIC policy is then included in the service profile assigned to that server. All traffic from the vNIC travels through the I/O module to the specified uplink Ethernet port.

Note If you do not assign a pin group to a server interface through a vNIC policy, Cisco UCS Manager chooses an uplink Ethernet port or port channel for traffic from that server interface dynamically. This choice is not permanent. A different uplink Ethernet port or port channel may be used for traffic from that server interface after an interface flap or a server reboot.

If an uplink is part of a LAN pin group, the uplink is not necessarily reserved for only that LAN pin group. Other vNIC's policies that do not specify a LAN pin group can use the uplink as a dynamic uplink.

Creating a LAN Pin Group

In a system with two fabric interconnects, you can associate the pin group with only one fabric interconnect or with both fabric interconnects.

Before You Begin

Configure the ports and port channels with which you want to configure the pin group. You can only include ports and port channels configured as uplink ports in a LAN pin group.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **LAN Cloud**.

Step 3 Right-click **LAN Pin Groups** and select **Create LAN Pin Group**.

Step 4 In the **Create LAN Pin Group** dialog box, enter a unique name and description for the pin group.

Step 5 To pin traffic for fabric interconnect A, do the following in the **Targets** area:

a) Check the **Fabric Interconnect A** check box.

b) Click the drop-down arrow on the **Interface** field and navigate through the tree-style browser to select the port or port channel you want to associate with the pin group.

Step 6 To pin traffic for fabric interconnect B, do the following in the **Targets** area:

a) Check the **Fabric Interconnect B** check box.

b) Click the drop-down arrow on the **Interface** field and navigate through the tree-style browser to select the port or port channel you want to associate with the pin group.

Step 7 Click **OK**.

What to Do Next

Include the pin group in a vNIC template.

Deleting a LAN Pin Group

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 In the **LAN** tab, expand **LAN** > **LAN Cloud** > **LAN Pin Groups**.

Step 3 Right-click the LAN pin group you want to delete and select **Delete**.

Step 4 If a confirmation dialog box displays, click **Yes**.

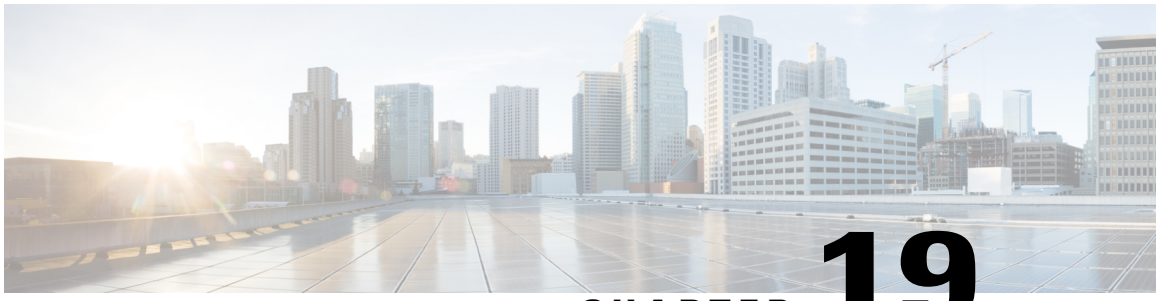

CHAPTER 19

Configuring MAC Pools

This chapter includes the following sections:

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 multitenancy, you can use the organizational hierarchy to ensure that MAC pools can be used only by specific applications or business services. Cisco UCS 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.

Creating a MAC Pool

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Pools**.

Step 3 Expand the node for the organization where you want to create the pool.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **MAC Pools** and select **Create MAC Pool**.

Step 5 In the **Define Name and Description** page of the **Create MAC Pool** wizard, complete the following fields:

Name	Description
Name field	The name of the MAC pool. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the MAC pool. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Assignment Order field	This can be one of the following: • **Default—**Cisco UCS Manager selects a random identity from the pool. • **Sequential—**Cisco UCS Manager selects the lowest available identity from the pool.

Step 6 Click **Next**.

Step 7 In the **Add MAC Addresses** page of the **Create MAC Pool** wizard, click **Add**.

Step 8 In the **Create a Block of MAC Addresses** dialog box, complete the following fields:

Name	Description
First MAC Address field	The first MAC address in the block.
Size field	The number of MAC addresses in the block.

Step 9 Click **OK**.

Step 10 Click **Finish**.

What to Do Next

Include the MAC pool in a vNIC template.

Deleting a MAC Pool

If you delete a pool, Cisco UCS Manager does not reallocate any addresses from that pool that were assigned to vNICs or vHBAs. All assigned addresses from a deleted pool remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 In the **LAN** tab, expand **LAN** > **Pools** > *Organization_Name* .

Step 3 Expand the **MAC Pools** node.

Step 4 Right-click the MAC pool you want to delete and select **Delete**.

Step 5 If a confirmation dialog box displays, click **Yes**.

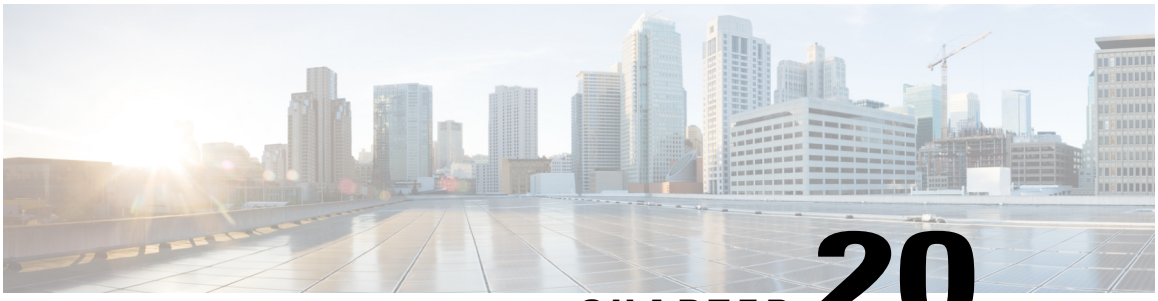

CHAPTER 20

Configuring Quality of Service

This chapter includes the following sections:

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.

Configuring System Classes

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 14: 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.

Configuring QoS System Classes

The type of adapter in a server might limit the maximum MTU supported. For example, network MTU above the maximums might cause the packet to be dropped for the following adapters:

- The Cisco UCS M71KR CNA adapter, which supports a maximum MTU of 9216.
- The Cisco UCS 82598KR-CI adapter, which supports a maximum MTU of 14000.

Important Use the same CoS (Class of Service) values on UCS and N5K for all the no-drop policies. To insure that end-to-end PFC works correctly, have the same QoS policy configured on all intermediate switches.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 In the **LAN** tab, expand **LAN** > **LAN Cloud**.

Step 3 Select the **QoS System Class** node.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 Update the properties for the system class that you want to configure to meet the traffic management needs of the system.

Note Some properties may not be configurable for all system classes:

- The maximum value for MTU is 9216.
- Changes saved to the drop displays the following warning message: Warning: The operation will cause momentary disruption to traffic forwarding.
- Changes saved to the MTU displays the following warning message: Warning: The operation will cause momentary disruption to traffic forwarding.

Step 6 Click **Save Changes**.

Enabling a QoS System Class

The Best Effort or Fibre Channel system classes are enabled by default.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 In the **LAN** tab, expand **LAN** > **LAN Cloud**.

Step 3 Select the **QoS System Class** node.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 Check the **Enabled** check box for the QoS system that you want to enable.

Step 6 Click **Save Changes**.

Disabling a QoS System Class

You cannot disable the Best Effort or Fibre Channel system classes.

All QoS policies that are associated with a disabled system class default to Best Effort or, if the disabled system class is configured with a Cos of 0, to the Cos 0 system class.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 In the **LAN** tab, expand **LAN** > **LAN Cloud**.
Step 3 Select the **QoS System Class** node.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 Uncheck the **Enabled** check box for the QoS system that you want to disable.
Step 6 Click **Save Changes**.

Configuring Quality of Service Policies

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.

Creating a QoS Policy

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 In the **LAN** tab, expand **LAN** > **Policies**.
Step 3 Expand the node for the organization where you want to create the pool.
If the system does not include multitenancy, expand the **root** node.
Step 4 Right-click **QoS Policy** and select **Create QoS Policy**.
Step 5 In the **Create QoS Policy** dialog box, complete the required fields.
Step 6 Click **OK**.

What to Do Next

Include the QoS policy in a vNIC or vHBA template.

Deleting a QoS Policy

If you delete a QoS policy that is in use or you disable a system class that is used in a QoS policy, any vNIC or vHBA that uses that QoS policy is assigned to the Best Effort system class or to the system class with a CoS of 0. In a system that implements multitenancy, Cisco UCS Manager first attempts to find a matching QoS policy in the organization hierarchy.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 Expand **Servers** > **Policies** > ***Organization_Name***.
Step 3 Expand the **QoS Policies** node.
Step 4 Right-click the QoS policy you want to delete and select **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Configuring Flow Control Policies

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.

Creating a Flow Control Policy

Before You Begin

Configure the network port with the corresponding setting for the flow control that you need. For example, if you enable the send setting for flow-control pause frames in the policy, ensure that the receive parameter in the network port is set to on or to desired. If you want the Cisco UCS port to receive flow-control frames, ensure that the send parameter is set to on or to desire on the network port. If you do not want to use flow control, you can set the send and receive parameters on the network port to off.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Policies**.

Step 3 Expand the **root** node.
You can only create a flow control policy in the root organization. You cannot create a flow control policy in a sub-organization.

Step 4 Right-click the **Flow Control Policies** node and select **Create Flow Control Policy**.

Step 5 In the **Create Flow Control Policy** wizard, complete the required fields.

Step 6 Click **OK**.

What to Do Next

Associate the flow control policy with an uplink Ethernet port or port channel.

Deleting a Flow Control Policy

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Policies** > *Organization_Name*.

Step 3 Expand the **Flow Control Policies** node.

Step 4 Right-click the policy you want to delete and select **Delete**.

Step 5 If a confirmation dialog box displays, click **Yes**.

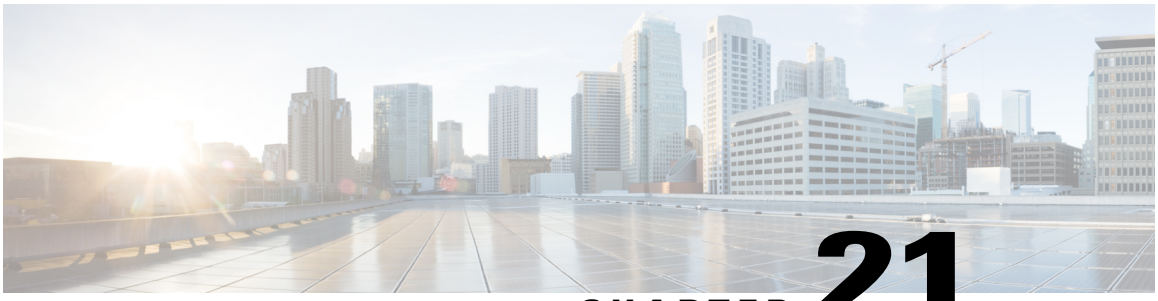

CHAPTER 21

Configuring Network-Related Policies

This chapter includes the following sections:

Configuring vNIC Templates

vNIC Template

The vNIC LAN connectivity policy defines how a vNIC on a server connects to the LAN.

Cisco UCS Manager does not automatically create a VM-FEX port profile with the correct settings when you create a vNIC template. If you want to create a VM-FEX port profile, you must configure the target of the vNIC template as a VM. You must include this policy in a service profile for it to take effect.

You can select VLAN groups in addition to any individual VLAN while creating a vNIC template.

Note If your server has two Emulex or QLogic NICs (Cisco UCS CNA M71KR-E or Cisco UCS CNA M71KR-Q), you must configure vNIC policies for both adapters in your service profile to get a user-defined MAC address for both NICs. If you do not configure policies for both NICs, Windows still detects both of them in the PCI bus. Then because the second eth is not part of your service profile, Windows assigns it a hardware MAC address. If you then move the service profile to a different server, Windows sees additional NICs because one NIC did not have a user-defined MAC address.

Creating a vNIC Template

Before You Begin

This policy requires that one or more of the following resources already exist in the system:

- Named VLAN
- MAC pool
- QoS policy
- LAN pin group
- Statistics threshold policy

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the **vNIC Templates** node and choose **Create vNIC Template**.

Step 5 In the **Create vNIC Template** dialog box:

a) In the **General** area, complete the following fields:

Name	Description
Name field	The name of the vNIC template. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A user-defined description of the template. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Name	Description
Fabric ID field	The fabric interconnect associated with the component. If you want vNICs created from this template to be able to access the second fabric interconnect if the default one is unavailable, check the **Enable Failover** check box. **Note** Do not enable vNIC fabric failover under the following circumstances: • If the Cisco UCS domain is running in Ethernet switch mode. vNIC fabric failover is not supported in that mode. If all Ethernet uplinks on one fabric interconnect fail, the vNICs do not fail over to the other. • If you plan to associate one or more vNICs created from this template to a server with an adapter that does not support fabric failover, such as the Cisco UCS 82598KR-CI 10-Gigabit Ethernet Adapter. If so, Cisco UCS Manager generates a configuration fault when you associate the service profile with the server.
Redundancy Type	The Redundancy type that you choose initiates a fabric failover using vNIC/HBA redundancy pairs. • **Primary Template—** Creates configurations that can be shared with the Secondary template. Any other shared changes on the Primary template are automatically synchronized to the Secondary template. • **Secondary Template—** All shared configurations are inherited from the Primary template. • **No Redundancy—** Legacy vNIC/vNHBA template behavior. Select this option if you do not want to use redundancy.
Target list box	A list of the possible targets for vNICs created from this template. The target you choose determines whether or not Cisco UCS Manager automatically creates a VM-FEX port profile with the appropriate settings for the vNIC template. This can be one of the following: • **Adapter—**The vNICs apply to all adapters. No VM-FEX port profile is created if you choose this option. • **VM—**The vNICs apply to all virtual machines. A VM-FEX port profile is created if you choose this option.

Name	Description
Template Type field	• **Initial Template**: vNICs created from this template are not updated if the template changes. • **Updating Template**: vNICs created from this template are updated if the template changes.

b) In the **VLANs** area, use the table to select the VLAN to assign to vNICs created from this template. The table contains the following columns:

Name	Description
Select column	Check the check box in this column for each VLAN that you want to use. **Note** VLANs and PVLANs can not be assigned to the same vNIC.
Name column	The name of the VLAN.
Native VLAN column	To designate one of the VLANs as the native VLAN, click the radio button in this column.

c) In the **Policies** area, complete the following fields:

Name	Description
CDN Source field	This can be one of the following options: • **vNIC Name** —Uses the vNIC template name of the vNIC instance as the CDN name. This is the default option. • **User Defined** — Displays the CDN Name field for you to enter a user-defined CDN name for the vNIC template.
MTU field	The maximum transmission unit, or packet size, that vNICs created from this vNIC template should use. Enter an integer between 1500 and 9000. **Note** If the vNIC template has an associated QoS policy, the MTU specified here must be equal to or less than the MTU specified in the associated QoS system class. If this MTU value exceeds the MTU value in the QoS system class, packets may be dropped during data transmission.
MAC Pool drop-down list	The MAC address pool that vNICs created from this vNIC template should use.

Name	Description
QoS Policy drop-down list	The quality of service policy that vNICs created from this vNIC template should use.
Network Control Policy drop-down list	The network control policy that vNICs created from this vNIC template should use.
Pin Group drop-down list	The LAN pin group that vNICs created from this vNIC template should use.
Stats Threshold Policy drop-down list	The statistics collection policy that vNICs created from this vNIC template should use.

Step 6 Click **OK**.

What to Do Next

Include the vNIC template in a service profile.

Creating vNIC Template Pairs

Procedure

Step 1 In the Navigation pane, click the **LAN** tab. On the **LAN** tab, expand **LAN > Policies.**

Step 2 Expand the node for the organization where you want to create the policy. If the system does not include multi-tenancy, expand the root node.

Step 3 Right-click the **vNIC Templates** node and choose **Create vNIC Template**. In the **Create vNIC Template** dialog box, assign a **Name**, **Description**, and select the **Fabric ID** for the template.

Step 4 Select the **Redundancy Type** as **Primary** or **Secondary** or **No Redundancy**. See the redundancy type descriptions below.

Step 5 Select the **Peer Redundancy Template—**to choose the name of the corresponding **Primary** or **Secondary** redundancy template to perform the template pairing from the **Primary** or **Secondary** redundancy template.

- **Primary—**Creates configurations that can be shared with the Secondary template. Any other shared changes on the Primary template are automatically synchronized to the Secondary template.
  - **VLANS**
  - **Template Type**
  - **MTU**
  - **Network Control Policies**
  - **Connection Policies**

- **QoS Policy**
- **Stats Threshold Policy**

Note

Following is a list of non-shared configurations:

- **Fabric ID**

  Note The Fabric ID must be mutually exclusive. If you assign the Primary template to Fabric A, then Fabric B is automatically assigned to the Secondary template as part of the synchronization from the Primary template.

- **CDN Source**
- **MAC Pool**
- **Description**
- **Pin Group Policy**

- **Secondary—**

  All shared configurations are inherited from the Primary template.

- **No Redundancy—**

  Legacy vNIC template behavior.

Step 6 Click **OK**.

What to Do Next

After you create the vNIC redundancy template pair, you can use the redundancy template pair to create redundancy vNIC pairs for any service profile in the same organization or sub-organization.

Undo vNIC Template Pairs

You can undo the vNIC template pair by changing the Peer Redundancy Template so that there is no peer template for the Primary or the Secondary template. When you undo a vNIC template pair, the corresponding vNIC pairs also becomes undone.

Procedure

Select **not set** from the **Peer Redundancy Template** drop-down list to undo the paring between the peer Primary or Secondary redundancy template used to perform the template pairing. You can also select **None** as the **Redundancy Type** to undo the pairing.

Note If you delete one template in a pair, you are prompt to delete the other template in the pair. If you do not delete the other template in the pair, that template resets its peer reference and retains its redundancy type.

Binding a vNIC to a vNIC Template

You can bind a vNIC associated with a service profile to a vNIC template. When you bind the vNIC to a vNIC template, Cisco UCS Manager configures the vNIC with the values defined in the vNIC template. If the existing vNIC configuration does not match the vNIC template, Cisco UCS Manager reconfigures the vNIC. You can only change the configuration of a bound vNIC through the associated vNIC template. You cannot bind a vNIC to a vNIC template if the service profile that includes the vNIC is already bound to a service profile template.

Important If the vNIC is reconfigured when you bind it to a template, Cisco UCS Manager reboots the server associated with the service profile.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that includes the service profile with the vNIC you want to bind. If the system does not include multi-tenancy, expand the **root** node.

Step 4 Expand *Service_Profile_Name* > **vNICs**.

Step 5 Click the vNIC you want to bind to a template.

Step 6 In the **Work** pane, click the **General** tab.

Step 7 In the **Actions** area, click **Bind to a Template**.

Step 8 In the **Bind to a vNIC Template** dialog box, do the following:

a) From the **vNIC Template** drop-down list, choose the template to which you want to bind the vNIC.

b) Click **OK**.

Step 9 In the warning dialog box, click **Yes** to acknowledge that Cisco UCS Manager may need to reboot the server if the binding causes the vNIC to be reconfigured.

Unbinding a vNIC from a vNIC Template

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that includes the service profile with the vNIC you want to unbind. If the system does not include multi-tenancy, expand the **root** node.

Step 4 Expand *Service_Profile_Name* > **vNICs**.
Step 5 Click the vNIC you want to unbind from a template.
Step 6 In the **Work** pane, click the **General** tab.
Step 7 In the **Actions** area, click **Unbind from a Template**.
Step 8 If a confirmation dialog box displays, click **Yes**.

Deleting a vNIC Template

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 Expand **LAN** > **Policies** > *Organization_Name*.
Step 3 Expand the **vNIC Templates** node.
Step 4 Right-click the policy you want to delete and choose **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Configuring Ethernet Adapter Policies

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 a cluster configuration with two fabric interconnects

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.
- LUN Queue Depth—The LUN queue depth setting is available for Windows system FC adapter policies. Queue depth is the number of commands that the HBA can send and receive in a single transmission per LUN. Windows Storport driver sets this to a default value of 20 for physical miniports and to 250 for virtual miniports. This setting adjusts the initial queue depth for all LUNs on the adapter. Valid range for this value is 1 to 254. The default LUN queue depth is 20. This feature only works with Cisco UCS Manager version 3.1(2) and higher.
- IO TimeOut Retry—When the target device is not responding to an IO request within the specified timeout, the FC adapter will abort the pending command then resend the same IO after the timer expires. The FC adapter valid range for this value is 1 to 59 seconds. The default IO retry timeout is 5 seconds. This feature only works with Cisco UCS Manager version 3.1(2) and higher.

Operating System Specific Adapter Policies

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.

Important We recommend that you use the values in these policies for the applicable operating system. Do not modify any of the values in the default policies unless directed to do so by Cisco Technical Support.

However, if you are creating an Ethernet adapter policy for a Windows OS (instead of using the default Windows adapter policy), you must use the following formulas to calculate values that work with Windows:

Completion Queues = Transmit Queues + Receive Queues
Interrupt Count = (Completion Queues + 2) rounded up to nearest power of 2

For example, if Transmit Queues = 1 and Receive Queues = 8 then:

Completion Queues = 1 + 8 = 9
Interrupt Count = (9 + 2) rounded up to the nearest power of 2 = 16

Accelerated Receive Flow Steering

Accelerated Receive Flow Steering (ARFS) is hardware-assisted receive flow steering that can increase CPU data cache hit rate by steering kernel level processing of packets to the CPU where the application thread consuming the packet is running.

Using ARFS can improve CPU efficiency and reduce traffic latency. Each receive queue of a CPU has an interrupt associated with it. You can configure the Interrupt Service Routine (ISR) to run on a CPU. The ISR moves the packet from the receive queue to the backlog of one of the current CPUs, which processes the packet later. If the application is not running on this CPU, the CPU must copy the packet to non-local memory, which adds to latency. ARFS can reduce this latency by moving that particular stream to the receive queue of the CPU on which the application is running.

ARFS is disabled by default and can be enabled through Cisco UCS Manager. To configure ARFS, do the following:

1. Create an adapter policy with ARFS enabled.
2. Associate the adapter policy with a service profile.
3. Enable ARFS on a host.
    1. Turn off Interrupt Request Queue (IRQ) balance.
    2. Associate IRQ with different CPUs.
    3. Enable ntuple by using ethtool.

Guidelines and Limitations for Accelerated Receive Flow Steering

- ARFS supports 64 filters per vNIC
- ARFS is supported on the following adapters:
    - Cisco UCS VIC 1280, 1240, 1340, and 1380
    - Cisco UCS VIC 1225, 1225T, 1285, 1223, 1227, 1227T, 1385, 1387
- ARFS is supported on the following Operating Systems:
    - Red Hat Enterprise Linux 6.5, and 6.6
    - Red Hat Enterprise Linux 7.0 and higher versions
    - SUSE Linux Enterprise Server 11 SP2 and SP3
    - SUSE Linux Enterprise Server 12 and higher versions
    - Ubuntu 14.04.2

Interrupt Coalescing

Adapters typically generate a large number of interrupts that a host CPU must service. Interrupt coalescing reduces the number of interrupts serviced by the host CPU. This is done by interrupting the host only once for multiple occurrences of the same event over a configurable coalescing interval.

When interrupt coalescing is enabled for receive operations, the adapter continues to receive packets, but the host CPU does not immediately receive an interrupt for each packet. A coalescing timer starts when the first packet is received by the adapter. When the configured coalescing interval times out, the adapter generates one interrupt with the packets received during that interval. The NIC driver on the host then services the multiple packets that are received. Reduction in the number of interrupts generated reduces the time spent by the host CPU on context switches. This means that the CPU has more time to process packets, which results in better throughput and latency.

Adaptive Interrupt Coalescing

Due to the coalescing interval, the handling of received packets adds to latency. For small packets with a low packet rate, this latency increases. To avoid this increase in latency, the driver can adapt to the pattern of traffic flowing through it and adjust the interrupt coalescing interval for a better response from the server.

Adaptive interrupt coalescing (AIC) is most effective in connection-oriented low link utilization scenarios including email server, databases server, and LDAP server. It is not suited for line-rate traffic.

Guidelines and Limitations for Adaptive Interrupt Coalescing

- Adaptive Interrupt Coalescing (AIC) does not provide any reduction in latency when the link utilization is more than 80 percent.
- Enabling AIC disables static coalescing.
- AIC is supported on the following Operating Systems:
  - Red Hat Enterprise Linux 6.4 and higher versions
  - Red Hat Enterprise Linux 7.0 and higher versions
  - SUSE Linux Enterprise Server 11 SP2 and SP3
  - SUSE Linux Enterprise Server 12
  - XenServer 6.5
  - Ubuntu 14.04.2

RDMA Over Converged Ethernet for SMB Direct

RDMA over Converged Ethernet (RoCE) allows direct memory access over an Ethernet network. RoCE is a link layer protocol, and hence, it allows communication between any two hosts in the same Ethernet broadcast domain. RoCE delivers superior performance compared to traditional network socket implementations because of lower latency, lower CPU utilization and higher utilization of network bandwidth. Windows 2012 R2 and later versions use RDMA for accelerating and improving the performance of SMB file sharing and Live Migration.

Cisco UCS Manager Release 2.2(4) supports RoCE for Microsoft SMB Direct. It sends additional configuration information to the adapter while creating or modifying an Ethernet adapter policy.

Guidelines and Limitations for SMB Direct with RoCE

- Microsoft SMB Direct with RoCE is supported:

◦ on Windows 2012 R2 for Cisco UCS Manager release 2.2(4) and later releases.

.

◦ on Windows 2016 for Cisco UCS Manager release 2.2(8) and later releases.

.

- Microsoft SMB Direct with RoCE is supported only with third generation Cisco UCS VIC 1340, 1380, 1385, 1387 adapters. Second generation UCS VIC 1225 and 1227 adapters are not supported.
- RoCE configuration is supported between Cisco adapters. Interoperability between Cisco adapters and third party adapters is not supported.
- Cisco UCS Manager does not support more than 4 RoCE-enabled vNICs per adapter.
- Cisco UCS Manager does not support RoCE with NVGRE, VXLAN, NetFlow, VMQ, or usNIC.
- Maximum number of queue pairs per adapter is 8192.
- Maximum number of memory regions per adapter is 524288.
- If you do not disable RoCE before downgrading Cisco UCS Manager from Release 2.2(4), downgrade will fail.
- Cisco UCS Manager does not support fabric failover for vNICs with RoCE enabled.

Creating an Ethernet Adapter Policy

Tip If the fields in an area do not display, click the **Expand** icon to the right of the heading.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **Adapter Policies** and choose **Create Ethernet Adapter Policy**.

Step 5 Enter a **Name** and optional **Description** for the policy.
This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Step 6 (Optional) In the **Resources** area, adjust the following values:

Name	Description
Transmit Queues field	The number of transmit queue resources to allocate. Enter an integer between 1 and 256.

Name	Description
Ring Size field	The number of descriptors in each transmit queue. Enter an integer between 64 and 4096.
Receive Queues field	The number of receive queue resources to allocate. Enter an integer between 1 and 256.
Ring Size field	The number of descriptors in each receive queue. Enter an integer between 64 and 4096.
Completion Queues field	The number of completion queue resources to allocate. In general, the number of completion queue resources you should allocate is equal to the number of transmit queue resources plus the number of receive queue resources. Enter an integer between 1 and 512.
Interrupts field	The number of interrupt resources to allocate. In general, this value should be equal to the number of completion queue resources. Enter an integer between 1 and 514.

Step 7 (Optional) In the **Options** area, adjust the following values:

Name	Description
Transmit Checksum Offload field	This can be one of the following: • **Disabled—**The CPU calculates all packet checksums. • **Enabled—**The CPU sends all packets to the hardware so that the checksum can be calculated. This option may reduce CPU overhead. **Note** This option affects only packets sent from the interface.
Receive Checksum Offload field	This can be one of the following: • **Disabled—**The CPU validates all packet checksums. • **Enabled—**The CPU sends all packet checksums to the hardware for validation. This option may reduce CPU overhead. **Note** This option affects only packets received by the interface.

Name	Description
TCP Segmentation Offload field	This can be one of the following: • **Disabled—**The CPU segments large TCP packets. • **Enabled—**The CPU sends large TCP packets to the hardware to be segmented. This option may reduce CPU overhead and increase throughput rate. **Note** This option is also known as Large Send Offload (LSO) and affects only packets sent from the interface.
TCP Large Receive Offload field	This can be one of the following: • **Disabled—**The CPU processes all large packets. • **Enabled—**The hardware reassembles all segmented packets before sending them to the CPU. This option may reduce CPU utilization and increase inbound throughput. **Note** This option affects only packets received by the interface.
Receive Side Scaling field	RSS distributes network receive processing across multiple CPUs in multiprocessor systems. This can be one of the following: • **Disabled—**Network receive processing is always handled by a single processor even if additional processors are available. • **Enabled—**Network receive processing is shared across processors whenever possible.
Accelerated Receive Flow Steering field	Packet processing for a flow must be performed on the local CPU. This is supported for Linux operating systems only. This can be one of the following: • **Disabled—**The CPU is not specified. • **Enabled—**Packet processing is performed on the local CPU.
Network Virtualization using Generic Routing Encapsulation field	Whether NVGRE overlay hardware offloads for TSO and checksum are enabled. This can be one of the following: • **Disabled—**NVGRE overlay hardware offloads are not enabled. • **Enabled—**NVGRE overlay hardware offloads are enabled.
Virtual Extensible LAN field	Whether VXLAN overlay hardware offloads for TSO and checksum are enabled. This can be one of the following: • **Disabled—**VXLAN overlay hardware offloads are not enabled. • **Enabled—**VXLAN overlay hardware offloads are enabled.

Name	Description
Failback Timeout field	After a vNIC has started using its secondary interface, this setting controls how long the primary interface must be available before the system resumes using the primary interface for the vNIC. Enter a number of seconds between 0 and 600.
Interrupt Mode field	The preferred driver interrupt mode. This can be one of the following: • **MSI X—**Message Signaled Interrupts (MSI) with the optional extension. This is the recommended option. • **MSI—**MSI only. • **INTx—**PCI INTx interrupts.
Interrupt Coalescing Type field	This can be one of the following: • **Min—**The system waits for the time specified in the **Interrupt Timer** field before sending another interrupt event. • **Idle—**The system does not send an interrupt until there is a period of no activity lasting as least as long as the time specified in the **Interrupt Timer** field.
Interrupt Timer field	The time to wait between interrupts or the idle period that must be encountered before an interrupt is sent. Enter a value between 1 and 65535. To turn off interrupt coalescing, enter 0 (zero) in this field.
RoCE field	Whether Remote Direct Memory Access over an Ethernet network is enabled. This can be one of the following: • **Disabled—**RoCE is disabled on the Ethernet adapter. • **Enabled—**RoCE is enabled on the Ethernet adapter.
RoCE Properties area	Lists the RoCE properties. This area is enabled only if you enable RoCE.
Queue Pairs	The number of queue pairs per adapter. Enter an integer between 1 and 8192. It is recommended that this number be an integer power of 2.
Memory Regions	The number of memory regions per adapter. Enter an integer between 1 and 524288. It is recommended that this number be an integer power of 2.

Name	Description
Resource Groups	The number of resource groups per adapter. Enter an integer between 1 and 128. It is recommended that this number be an integer power of 2 greater than or equal to the number of CPU cores on the system for optimum performance.

Step 8 Click **OK**.

Step 9 If a confirmation dialog box displays, click **Yes**.

Configuring an Ethernet Adapter Policy to Enable eNIC Support for MRQS on Linux Operating Systems

Cisco UCS Manager includes eNIC support for the Multiple Receive Queue Support (MRQS) feature on Red Hat Enterprise Linux Version 6.x and SUSE Linux Enterprise Server Version 11.x.

Procedure

Step 1 Create an Ethernet adapter policy.
Use the following parameters when creating the Ethernet adapter policy:

- Transmit Queues = 1
- Receive Queues = n (up to 8)
- Completion Queues = # of Transmit Queues + # of Receive Queues
- Interrupts = # Completion Queues + 2
- Receive Side Scaling (RSS) = Enabled
- Interrupt Mode = Msi-X

See Creating an Ethernet Adapter Policy, on page 294.

Step 2 Install an eNIC driver Version 2.1.1.35 or later.
See Cisco UCS Virtual Interface Card Drivers for Linux Installation Guide.

Step 3 Reboot the server

Configuring an Ethernet Adapter Policy to Enable Stateless Offloads with NVGRE

Cisco UCS Manager supports stateless offloads with NVGRE only with Cisco UCS VIC 1340 and/or Cisco UCS VIC 1380 adapters that are installed on servers running Windows Server 2012 R2 operating systems. Stateless offloads with NVGRE cannot be used with NetFlow, usNIC, or VM-FEX.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **Adapter Policies** and choose **Create Ethernet Adapter Policy**.

a) In the **Resources** area, set the following options:

- Transmit Queues = 1
- Receive Queues = n (up to 8)
- Completion Queues = # of Transmit Queues + # of Receive Queues
- Interrupts = # Completion Queues + 2

b) In the **Options** area, set the following options:

- Network Virtualization using Generic Routing Encapsulation = Enabled
- Interrupt Mode = Msi-X

For more information on creating an Ethernet adapter policy, see Creating an Ethernet Adapter Policy, on page 294.

Step 5 Click **OK** to create the Ethernet adapter policy.

Step 6 Install an eNIC driver Version 3.0.0.8 or later.
For more information, see http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/vic_drivers/install/Windows/b_Cisco_VIC_Drivers_for_Windows_Installation_Guide.html.

Step 7 Reboot the server.

Configuring an Ethernet Adapter Policy to Enable Stateless Offloads with VXLAN

Cisco UCS Manager supports stateless offloads with VXLAN only with Cisco UCS VIC 1340 and/or Cisco UCS VIC 1380 adapters that are installed on servers running VMWare ESXi Release 5.5 and later releases of the operating system. Stateless offloads with VXLAN cannot be used with NetFlow, usNIC, or VM-FEX.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **Adapter Policies** and choose **Create Ethernet Adapter Policy**.

a) In the **Resources** area, set the following options:

- Transmit Queues = 1
- Receive Queues = n (up to 8)
- Completion Queues = # of Transmit Queues + # of Receive Queues
- Interrupts = # Completion Queues + 2

b) In the **Options** area, set the following options:

- Virtual Extensible LAN = Enabled
- Interrupt Mode = Msi-X

For more information on creating an ethernet adapter policy, see Creating an Ethernet Adapter Policy, on page 294.

Step 5 Click **OK** to create the Ethernet adapter policy.

Step 6 Install an eNIC driver Version 2.1.2.59 or later.
For more information, see http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/vic_drivers/install/ESX/2-0/b_Cisco_VIC_Drivers_for_ESX_Installation_Guide.html.

Step 7 Reboot the server.

Deleting an Ethernet Adapter Policy

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Policies** > ***Organization_Name***.

Step 3 Expand the **Adapter Policies** node.

Step 4 Right-click the Ethernet adapter policy that you want to delete and choose **Delete**.

Step 5 If a confirmation dialog box displays, click **Yes**.

Configuring the Default vNIC Behavior Policy

Default vNIC Behavior Policy

Default vNIC behavior policy allows you to configure how vNICs are created for a service profile. You can choose to create vNICS manually, or you can create them automatically.

You can configure the default vNIC behavior policy to define how vNICs are created. This can be one of the following:

- **None—**Cisco UCS Manager does not create default vNICs for a service profile. All vNICs must be explicitly created.
- **HW Inherit—**If a service profile requires vNICs and none have been explicitly defined, Cisco UCS Manager creates the required vNICs based on the adapter installed in the server associated with the service profile.

Note If you do not specify a default behavior policy for vNICs, **HW Inherit** is used by default.

Configuring a Default vNIC Behavior Policy

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Policies**.

Step 3 Expand the **root** node.
You can configure only the default vNIC behavior policy in the root organization. You cannot configure the default vNIC behavior policy in a sub-organization.

Step 4 Click **Default vNIC Behavior**.

Step 5 On the **General Tab**, in the **Properties** area, click one of the following radio buttons in the **Action** field:

- **None—**Cisco UCS Manager does not create default vNICs for a service profile. All vNICs must be explicitly created.
- **HW Inherit—**If a service profile requires vNICs and none have been explicitly defined, Cisco UCS Manager creates the required vNICs based on the adapter installed in the server associated with the service profile.

Step 6 Click **Save Changes**.

Configuring LAN Connectivity Policies

About the LAN and SAN Connectivity Policies

Connectivity policies determine the connections and the network communication resources between the server and the LAN or SAN on the network. These policies use pools to assign MAC addresses, WWNs, and WWPNs to servers and to identify the vNICs and vHBAs that the servers use to communicate with the network.

Note We do not recommend that you use static IDs in connectivity policies, because these policies are included in service profiles and service profile templates and can be used to configure multiple servers.

Privileges Required for LAN and SAN Connectivity Policies

Connectivity policies enable users without network or storage privileges to create and modify service profiles and service profile templates with network and storage connections. However, users must have the appropriate network and storage privileges to create connectivity policies.

Privileges Required to Create Connectivity Policies

Connectivity policies require the same privileges as other network and storage configurations. For example, you must have at least one of the following privileges to create connectivity policies:

- admin—Can create LAN and SAN connectivity policies
- ls-server—Can create LAN and SAN connectivity policies
- ls-network—Can create LAN connectivity policies
- ls-storage—Can create SAN connectivity policies

Privileges Required to Add Connectivity Policies to Service Profiles

After the connectivity policies have been created, a user with ls-compute privileges can include them in a service profile or service profile template. However, a user with only ls-compute privileges cannot create connectivity policies.

Interactions between Service Profiles and Connectivity Policies

You can configure the LAN and SAN connectivity for a service profile through either of the following methods:

- LAN and SAN connectivity policies that are referenced in the service profile
- Local vNICs and vHBAs that are created in the service profile
- Local vNICs and a SAN connectivity policy
- Local vHBAs and a LAN connectivity policy

Cisco UCS maintains mutual exclusivity between connectivity policies and local vNIC and vHBA configuration in the service profile. You cannot have a combination of connectivity policies and locally created vNICs or vHBAs. When you include a LAN connectivity policy in a service profile, all existing vNIC configuration is erased, and when you include a SAN connectivity policy, all existing vHBA configuration in that service profile is erased.

Creating a LAN Connectivity Policy

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **LAN Connectivity Policies** and choose **Create LAN Connectivity Policy**.

Step 5 In the **Create LAN Connectivity Policy** dialog box, enter a name and optional description.

Step 6 Do one of the following:

- To add vNICs to the LAN connectivity policy, continue with Step 7.
- To add iSCSI vNICs to the LAN connectivity policy and use iSCSI boot with the server, continue with Step 8.

Step 7 To add vNICs, click **Add** next to the plus sign and complete the following fields in the **Create vNIC** dialog box:

a) In the **Create vNIC** dialog box, enter the name, select a **MAC Address Assignment**, and check the **Use vNIC Template** check box to use an existing vNIC template.
You can also create a MAC pool from this area.

b) Choose the **Fabric ID**, select the **VLANs** that you want to use, enter the **MTU**, and choose a **Pin Group**.
You can also create a VLAN and a LAN pin group from this area.

Note Cisco recommends using the native VLAN 1 setting to prevent traffic interruptions if using the Cisco Nexus 1000V Series Switches because changing the native VLAN 1 setting on a vNIC causes the port to turn on and off. You can only change the native VLAN setting on a Virtual Private Cloud (VPC) secondary port, and then change the primary port on the VPC.

c) In the **Operational Parameters** area, choose a **Stats Threshold Policy**.

d) In the Adapter Performance Profile area, choose an **Adapter Policy**, **QoS Policy**, and a **Network Control Policy**.
You can also create an Ethernet adapter policy, QoS policy, and network control policy from this area.

e) In the Connection Policies area, choose the **Dynamic vNIC**, **usNIC** or **VMQ** radio button, then choose the corresponding policy.
You can also create a dynamic vNIC, usNIC, or VMQ connection policy from this area.

f) Click **OK**.

Step 8 If you want to use iSCSI boot with the server, click the down arrows to expand the **Add iSCSI vNICs** bar and do the following:

a) Click **Add** on the table icon bar.

b) In the **Create iSCSI vNIC** dialog box, enter the **Name** and choose the **Overlay vNIC**, **iSCSI Adapter Policy**, and **VLAN**.
You can also create an iSCSI adapter policy from this area.

Note For the Cisco UCS M81KR Virtual Interface Card and the Cisco UCS VIC-1240 Virtual Interface Card, the VLAN that you specify must be the same as the native VLAN on the overlay vNIC.

For the Cisco UCS M51KR-B Broadcom BCM57711 Adapter, the VLAN that you specify can be any VLAN assigned to the overlay vNIC.

c) In the **MAC Address Assignment** drop-down list in the **iSCSI MAC Address** area, choose one of the following:

- Leave the MAC address unassigned, select **Select (None used by default)**. Select this option if the server that will be associated with this service profile contains a Cisco UCS M81KR Virtual Interface Card adapter or a Cisco UCS VIC-1240 Virtual Interface Card.

  Important If the server that will be associated with this service profile contains a Cisco UCS NIC M51KR-B adapter, you must specify a MAC address.

- A specific MAC address, select **00:25:B5:XX:XX:XX** and enter the address in the **MAC Address** field. To verify that this address is available, click the corresponding link.

- A MAC address from a pool, select the pool name from the list. Each pool name is followed by a pair of numbers in parentheses. The first number is the number of available MAC addresses in the pool and the second is the total number of MAC addresses in the pool.

  If this Cisco UCS domain is registered with Cisco UCS Central, there might be two pool categories. **Domain Pools** are defined locally in the Cisco UCS domain and **Global Pools** are defined in Cisco UCS Central.

d) (Optional) If you want to create a MAC pool that will be available to all service profiles, click **Create MAC Pool** and complete the fields in the **Create MAC Pool** wizard.
For more information, see Creating a MAC Pool, on page 271.

e) Click **OK**.

Step 9 After you have created all the vNICs or iSCSI vNICs you need for the policy, click **OK**.

What to Do Next

Include the policy in a service profile or service profile template.

Creating a vNIC for a LAN Connectivity Policy

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Policies** > *Organization_Name*.

Step 3 Expand the **LAN Connectivity Policies** node.

Step 4 Choose the policy to which you want to add a vNIC.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 On the icon bar of the **vNICs** table, click **Add**.

Step 7 In the **Create vNIC** dialog box, enter the name, select a **MAC Address Assignment**, and check the **Use vNIC Template** check box if you want to use an existing vNIC template.
You can also create a MAC pool from this area.

Step 8 Choose the **Fabric ID**, select the **VLANs** that you want to use, enter the **MTU**, and choose a **Pin Group**.
You can also create a VLAN and a LAN pin group from this area.

Step 9 In the **Operational Parameters** area, choose a **Stats Threshold Policy**.

Step 10 In the Adapter Performance Profile area, choose an **Adapter Policy**, **QoS Policy**, and a **Network Control Policy**.
You can also create an Ethernet adapter policy, QoS policy, and network control policy from this area.

Step 11 In the Connection Policies area, choose the **Dynamic vNIC**, **usNIC** or **VMQ** radio button, then choose the corresponding policy.
You can also create a dynamic vNIC, usNIC, or VMQ connection policy from this area.

Step 12 Click **OK**.

Step 13 Click **Save Changes**.

Deleting a vNIC from a LAN Connectivity Policy

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Policies** > *Organization_Name*.

Step 3 Expand the **LAN Connectivity Policies** node.

Step 4 Select the policy from which you want to delete the vNIC.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 In the **vNICs** table, do the following:

a) Click the vNIC you want to delete.

b) On the icon bar, click **Delete**.

Step 7 If a confirmation dialog box displays, click **Yes**.

Step 8 Click **Save Changes**.

Creating an iSCSI vNIC for a LAN Connectivity Policy

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Policies** > *Organization_Name*.

Step 3 Expand the **LAN Connectivity Policies** node.

Step 4 Choose the policy to which you want to add an iSCSI vNIC.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 On the icon bar of the **Add iSCSI vNICs** table, click **Add**.

Step 7 In the **Create iSCSI vNIC** dialog box, complete the following fields:

Name	Description
Name field	The name of the iSCSI vNIC. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Overlay vNIC drop-down list	The LAN vNIC associated with this iSCSI vNIC, if any.
iSCSI Adapter Policy drop-down list	The iSCSI adapter policy associated with this iSCSI vNIC, if any.
Create iSCSI Adapter Policy link	Click this link to create a new iSCSI adapter policy that will be available to all iSCSI vNICs.
VLAN drop-down list	The virtual LAN associated with this iSCSI vNIC. The default VLAN is **default**. **Note** For the Cisco UCS M81KR Virtual Interface Card and the Cisco UCS VIC-1240 Virtual Interface Card, the VLAN that you specify must be the same as the native VLAN on the overlay vNIC. For the Cisco UCS M51KR-B Broadcom BCM57711 Adapter, the VLAN that you specify can be any VLAN assigned to the overlay vNIC.

Step 8 In the **MAC Address Assignment** drop-down list in the **iSCSI MAC Address** area, choose one of the following:

- Leave the MAC address unassigned, select **Select (None used by default)**. Select this option if the server that will be associated with this service profile contains a Cisco UCS M81KR Virtual Interface Card adapter or a Cisco UCS VIC-1240 Virtual Interface Card.

  Important If the server that will be associated with this service profile contains a Cisco UCS NIC M51KR-B adapter, you must specify a MAC address.

- A specific MAC address, select **00:25:B5:XX:XX:XX** and enter the address in the **MAC Address** field. To verify that this address is available, click the corresponding link.
- A MAC address from a pool, select the pool name from the list. Each pool name is followed by a pair of numbers in parentheses. The first number is the number of available MAC addresses in the pool and the second is the total number of MAC addresses in the pool.

  If this Cisco UCS domain is registered with Cisco UCS Central, there might be two pool categories. **Domain Pools** are defined locally in the Cisco UCS domain and **Global Pools** are defined in Cisco UCS Central.

Step 9 (Optional) If you want to create a MAC pool that will be available to all service profiles, click **Create MAC Pool** and complete the fields in the **Create MAC Pool** wizard.
For more information, see Creating a MAC Pool, on page 271.

Step 10 Click **OK**.

Step 11 Click **Save Changes**.

Deleting an iSCSI vNIC from a LAN Connectivity Policy

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Policies** > ***Organization_Name***.

Step 3 Expand the **LAN Connectivity Policies** node.

Step 4 Chose the policy from which you want to delete the iSCSI vNIC.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 In the **Add iSCSI vNICs** table, do the following:

a) Click the iSCSI vNIC that you want to delete.
b) On the icon bar, click **Delete**.

Step 7 If a confirmation dialog box displays, click **Yes**.

Step 8 Click **Save Changes**.

Deleting a LAN Connectivity Policy

If you delete a LAN connectivity policy that is included in a service profile, it also deletes all vNICs and iSCSI vNICs from that service profile, and disrupt LAN data traffic for the server associated with the service profile.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Policies** > *Organization_Name*.

Step 3 Expand the **LAN Connectivity Policies** node.

Step 4 Right-click the policy that you want to delete and choose **Delete**.

Step 5 If a confirmation dialog box displays, click **Yes**.

Configuring Network Control Policies

Network Control Policy

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

- Whether the Cisco Discovery Protocol (CDP) is enabled or disabled
- How the virtual interface (VIF) behaves if no uplink port is available in end-host mode
- The action that Cisco UCS Manager takes on the remote Ethernet interface, vEthernet interface , or vFibre Channel interface when the associated border port fails
- Whether the server can use different MAC addresses when sending packets to the fabric interconnect
- Whether MAC registration occurs on a per-VNIC basis or for all VLANs

Action on Uplink Fail

By default, the **Action on Uplink Fail** property in the network control policy is configured with a value of link-down. For adapters such as the Cisco UCS M81KR Virtual Interface Card, this default behavior directs Cisco UCS Manager to bring the vEthernet or vFibre Channel interface down if the associated border port fails. For Cisco UCS systems using a non-VM-FEX capable converged network adapter that supports both Ethernet and FCoE traffic, such as Cisco UCS CNA M72KR-Q and the Cisco UCS CNA M72KR-E, this default behavior directs Cisco UCS Manager to bring the remote Ethernet interface down if the associated border port fails. In this scenario, any vFibre Channel interfaces that are bound to the remote Ethernet interface are brought down as well.

Note If your implementation includes those types of non-VM-FEX capable converged network adapters mentioned in this section and the adapter is expected to handle both Ethernet and FCoE traffic, we recommend that you configure the **Action on Uplink Fail** property with a value of warning. Note that this configuration might result in an Ethernet teaming driver not being able to detect a link failure when the border port goes down.

MAC Registration Mode

MAC addresses are installed only on the native VLAN by default, which maximizes the VLAN port count in most implementations.

Note If a trunking driver is being run on the host and the interface is in promiscuous mode, we recommend that you set the MAC Registration Mode to All VLANs.

NIC Teaming and Port Security

NIC teaming is a grouping together of network adapters to build in redundancy, and is enabled on the host. This teaming or bonding facilitates various functionalities, including load balancing across links and failover. When NIC teaming is enabled and events such as failover or reconfiguration take place, MAC address conflicts and movement may happen.

Port security, which is enabled on the fabric interconnect side, prevents MAC address movement and deletion. Therefore, you must not enable port security and NIC teaming together.

Configuring Link Layer Discovery Protocol for Fabric Interconnect vEthernet Interfaces

Cisco UCS Manager Release 2.2.4 allows you to enable and disable LLDP on a vEthernet interface. You can also retrieve information about these LAN uplink neighbors. This information is useful while learning the topology of the LAN connected to the UCS system and while diagnosing any network connectivity issues from the Fabric Interconnect (FI). The FI of a UCS system is connected to LAN uplink switches for LAN connectivity and to SAN uplink switches for storage connectivity. When using Cisco UCS with Cisco Application Centric Infrastructure (ACI), LAN uplinks of the FI are connected to ACI leaf nodes. Enabling LLDP on a vEthernet interface will help the Application Policy Infrastructure Controller (APIC) to identify the servers connected to the FI by using vCenter.

To permit the discovery of devices in a network, support for Link Layer Discovery Protocol (LLDP), a vendor-neutral device discovery protocol that is defined in the IEEE 802.1ab standard, is introduced. LLDP is a one-way protocol that allows network devices to advertise information about themselves to other devices on the network. LLDP transmits information about the capabilities and current status of a device and its interfaces. LLDP devices use the protocol to solicit information only from other LLDP devices.

You can enable or disable LLDP on a vEthernet interface based on the Network Control Policy (NCP) that is applied on the vNIC in the service profile.

Creating a Network Control Policy

MAC address-based port security for Emulex converged Network Adapters (N20-AE0102) is not supported. When MAC address-based port security is enabled, the fabric interconnect restricts traffic to packets that contain the MAC address that it first learns. This is either the source MAC address used in the FCoE Initialization Protocol packet, or the MAC address in an ethernet packet, whichever is sent first by the adaptor. This configuration can result in either FCoE or Ethernet packets being dropped.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the **Network Control Policies** node and select **Create Network Control Policy**.

Step 5 In the **Create Network Control Policy** dialog box, complete the required fields.

Step 6 In the LLDP area, do the following:

a) To enable the transmission of LLDP packets on an interface, click **Enabled** in the **Transmit** field.

b) To enable the reception of LLDP packets on an interface, click **Enabled** in the **Receive** field.

Step 7 In the **MAC Security** area, do the following to determine whether the server can use different MAC addresses when sending packets to the fabric interconnect:

a) Click the **Expand** icon to expand the area and display the radio buttons.

b) Click one of the following radio buttons to determine whether forged MAC addresses are allowed or denied when packets are sent from the server to the fabric interconnect:

- **Allow—** All server packets are accepted by the fabric interconnect, regardless of the MAC address associated with the packets.
- **Deny—** After the first packet has been sent to the fabric interconnect, all other packets must use the same MAC address or they will be silently rejected by the fabric interconnect. In effect, this option enables port security for the associated vNIC.

If you plan to install VMware ESX on the associated server, you must configure the **MAC Security** to **allow** for the network control policy applied to the default vNIC. If you do not configure **MAC Security** for **allow**, the ESX installation may fail because the MAC security permits only one MAC address while the installation process requires more than one MAC address.

Step 8 Click **OK**.

Deleting a Network Control Policy

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 Expand **LAN** > **Policies** > *Organization_Name*.
Step 3 Expand the **Network Control Policies** node.
Step 4 Right-click the policy you want to delete and select **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Configuring Multicast Policies

Multicast Policy

This policy is used to configure Internet Group Management Protocol (IGMP) snooping and IGMP querier. IGMP Snooping dynamically determines hosts in a VLAN that should be included in particular multicast transmissions. You can create, modify, and delete a multicast policy that can be associated to one or more VLANs. When a multicast policy is modified, all VLANs associated with that multicast policy are re-processed to apply the changes. For private VLANs, you can set a multicast policy for primary VLANs but not for their associated isolated VLANs due to a Cisco NX-OS forwarding implementation.

By default, IGMP snooping is enabled and IGMP querier is disabled. When IGMP snooping is enabled, the fabric interconnects send the IGMP queries only to the hosts. They do not send IGMP queries to the upstream network. To send IGMP queries to the upstream, do one of the following:

- Configure IGMP querier on the upstream fabric interconnect with IGMP snooping enabled
- Disable IGMP snooping on the upstream fabric interconnect
- Change the fabric interconnects to switch mode

The following limitations and guidelines apply to multicast policies:

- On a 6200 series fabric interconnect, user-defined multicast policies can also be assigned along with the default multicast policy.
- Only the default multicast policy is allowed for a global VLAN.
- If a Cisco UCS domain includes 6300 and 6200 series fabric interconnects, any multicast policy can be assigned.
- We highly recommend you use the same IGMP snooping state on the fabric interconnects and the associated LAN switches. For example, if IGMP snooping is disabled on the fabric interconnects, it should be disabled on any associated LAN switches as well.

Creating a Multicast Policy

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 Expand **LAN** > **Policies**.
Step 3 Expand the **root** node.
Step 4 Right-click the **Multicast Policies** node and select **Create Multicast Policy**.
Step 5 In the **Create Multicast Policy** dialog box, specify the name and IGMP snooping information.
Step 6 Click **OK**.

Modifying a Multicast Policy

This procedure describes how to change the IGMP snooping state and the IGMP snooping querier state of an existing multicast policy.

Note You cannot change the name of the multicast policy once it has been created.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 Expand **LAN** > **Policies**.
Step 3 Expand the **root** node.
Step 4 Click the policy that you want to modify.
Step 5 In the work pane, edit the fields as needed.
Step 6 Click **Save Changes**.

Deleting a Multicast Policy

Note If you assigned a non-default (user-defined) multicast policy to a VLAN and then delete that multicast policy, the associated VLAN inherits the multicast policy settings from the default multicast policy until the deleted policy is re-created.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 Expand **LAN** > **Policies**.
Step 3 Expand the **root** node.
Step 4 Right-click the **Multicast Policies** node and select **Delete Multicast Policy**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Configuring LACP policies

LACP Policy

Link Aggregation combines multiple network connections in parallel to increase throughput and to provide redundancy. Link aggregation control protocol (LACP) provides additional benefits for these link aggregation groups. Cisco UCS Manager enables you to configure LACP properties using LACP policy.

You can configure the following for a lacp policy:

- **Suspended-individual**: If you do not configure the ports on an upstream switch for lacp, the fabric interconnects treat all ports as uplink Ethernet ports to forward packets. You can place the lacp port in suspended state to avoid loops. When you set suspend-individual on a port-channel with LACP, if a port that is part of the port-channel does not receive PDUs from the peer port, it will go into suspended state.
- **Timer values**: You can configure rate-fast or rate-normal. In rate-fast configuration, the port is expected to receive 1 PDU every 1 second from the peer port. The time out for this is 3 seconds. In rate-normal configuration, the port is expected to receive 1 PDU every 30 seconds. The timeout for this is 90 seconds.

System creates a default LACP policy at system start up. You can modify this policy or create a new policy. You can also apply one LACP policy to multiple port-channels.

Creating a LACP Policy

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 Expand **LAN** > **Policies**.
Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 In the **Work Pane**, click **LACP Policies** tab, and click the + sign.

Step 5 In the **Create LACP Policy** dialog box, fill in the required fields.

Step 6 Click **OK**.

Modifying a LACP Policy

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 In the **Work Pane**, **LACP Policies** tab, and click on the policy you want to edit.

Step 5 Click the **Properties** icon on the right.

Step 6 In the **Properties** dialog box, make the required changes and click **Apply**.

Step 7 Click **OK**.

Configuring UDLD Link Policies

Understanding UDLD

UniDirectional Link Detection (UDLD) is a Layer 2 protocol that enables devices connected through fiber-optic or twisted-pair Ethernet cables to monitor the physical configuration of the cables and detect when a unidirectional link exists. All connected devices must support UDLD for the protocol to successfully identify and disable unidirectional links. When UDLD detects a unidirectional link, it marks the link as unidirectional. Unidirectional links can cause a variety of problems, including spanning-tree topology loops.

UDLD works with the Layer 1 mechanisms to determine the physical status of a link. At Layer 1, autonegotiation takes care of physical signaling and fault detection. UDLD performs tasks that autonegotiation cannot perform, such as detecting the identities of neighbors and shutting down misconnected interfaces. When you enable both autonegotiation and UDLD, the Layer 1 and Layer 2 detections work together to prevent physical and logical unidirectional connections and the malfunctioning of other protocols.

A unidirectional link occurs whenever traffic sent by a local device is received by its neighbor but traffic from the neighbor is not received by the local device.

Modes of Operation

UDLD supports two modes of operation: normal (the default) and aggressive. In normal mode, UDLD can detect unidirectional links due to misconnected interfaces on fiber-optic connections. In aggressive mode,

UDLD can also detect unidirectional links due to one-way traffic on fiber-optic and twisted-pair links and to misconnected interfaces on fiber-optic links.

In normal mode, UDLD detects a unidirectional link when fiber strands in a fiber-optic interface are misconnected and the Layer 1 mechanisms do not detect this misconnection. If the interfaces are connected correctly but the traffic is one way, UDLD does not detect the unidirectional link because the Layer 1 mechanism, which is supposed to detect this condition, does not do so. In case, the logical link is considered undetermined, and UDLD does not disable the interface. When UDLD is in normal mode, if one of the fiber strands in a pair is disconnected and autonegotiation is active, the link does not stay up because the Layer 1 mechanisms did not detect a physical problem with the link. In this case, UDLD does not take any action, and the logical link is considered undetermined.

UDLD aggressive mode is disabled by default. Configure UDLD aggressive mode only on point-to-point links between network devices that support UDLD aggressive mode. With UDLD aggressive mode enabled, when a port on a bidirectional link that has a UDLD neighbor relationship established stops receiving UDLD packets, UDLD tries to reestablish the connection with the neighbor and administratively shuts down the affected port. UDLD in aggressive mode can also detect a unidirectional link on a point-to-point link on which no failure between the two devices is allowed. It can also detect a unidirectional link when one of the following problems exists:

- On fiber-optic or twisted-pair links, one of the interfaces cannot send or receive traffic.
- On fiber-optic or twisted-pair links, one of the interfaces is down while the other is up.
- One of the fiber strands in the cable is disconnected.

Methods to Detect Unidirectional Links

UDLD operates by using two mechanisms:

- Neighbor database maintenance

  UDLD learns about other UDLD-capable neighbors by periodically sending a hello packet (also called an advertisement or probe) on every active interface to keep each device informed about its neighbors. When the switch receives a hello message, it caches the information until the age time (hold time or time-to-live) expires. If the switch receives a new hello message before an older cache entry ages, the switch replaces the older entry with the new one.

  UDLD clears all existing cache entries for the interfaces affected by the configuration change whenever an interface is disabled and UDLD is running, whenever UDLD is disabled on an interface, or whenever the switch is reset. UDLD sends at least one message to inform the neighbors to flush the part of their caches affected by the status change. The message is intended to keep the caches synchronized.

- Event-driven detection and echoing

  UDLD relies on echoing as its detection mechanism. Whenever a UDLD device learns about a new neighbor or receives a resynchronization request from an out-of-sync neighbor, it restarts the detection window on its side of the connection and sends echo messages in reply. Because this behavior is the same on all UDLD neighbors, the sender of the echoes expects to receive an echo in reply.

  If the detection window ends and no valid reply message is received, the link might shut down, depending on the UDLD mode. When UDLD is in normal mode, the link might be considered undetermined and might not be shut down. When UDLD is in aggressive mode, the link is considered unidirectional, and the interface is shut down.

If UDLD in normal mode is in the advertisement or in the detection phase and all the neighbor cache entries are aged out, UDLD restarts the link-up sequence to resynchronize with any potentially out-of-sync neighbors.

If you enable aggressive mode when all the neighbors of a port have aged out either in the advertisement or in the detection phase, UDLD restarts the link-up sequence to resynchronize with any potentially out-of-sync neighbor. UDLD shuts down the port if, after the fast train of messages, the link state is still undetermined.

UDLD Configuration Guidelines

The following guidelines and recommendations apply when you configure UDLD:

- A UDLD-capable interface also cannot detect a unidirectional link if it is connected to a UDLD-incapable port of another switch.
- When configuring the mode (normal or aggressive), make sure that the same mode is configured on both sides of the link.
- UDLD should be enabled only on interfaces that are connected to UDLD capable devices. The following interface types are supported:
  - Ethernet uplink
  - FCoE uplink
  - Ethernet uplink port channel member
  - FCoE uplink port channel member

Creating a Link Profile

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 Expand **LAN** > **Policies** > **LAN Cloud**.
Step 3 Right-click the **Link Profile** node and choose **Create Link Profile**.
Step 4 In the **Create Link Profile** dialog box, specify the name and the UDLD link policy.
Step 5 Click **OK**.

Creating a UDLD Link Policy

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 Expand **LAN** > **Policies** > **LAN Cloud**.
Step 3 Right-click the **UDLD Link Policy** node and choose **Create UDLD Link Policy**.
Step 4 In the **Create UDLD Link Policy** dialog box, specify the name, admin state, and mode.
Step 5 Click **OK**.

Modifying the UDLD System Settings

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 Expand **LAN** > **Policies** > **LAN Cloud**.
Step 3 On the LAN tab, expand **LAN** > **Policies** > **root**.
Step 4 Expand the **Link Protocol Policy** node and click **UDLD System Settings**.
Step 5 In the **Work** pane, click the **General** tab.
Step 6 In the **Properties** area, modify the fields as needed.
Step 7 Click **Save Changes**.

Assigning a Link Profile to a Port Channel Ethernet Interface

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 Expand **LAN** > **LAN Cloud** > *Fabric* > **Port Channels**.
Step 3 Expand the port channel node and click the Eth Interface where you want to assign a link profile.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Properties** area, choose the link profile that you want to assign.
Step 6 Click **Save Changes**.

Assigning a Link Profile to an Uplink Ethernet Interface

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 On the **LAN** tab, expand **LAN** > **LAN Cloud** > ***Fabric*** > **Uplink Eth Interface**
Step 3 Click the Eth Interface where you want to assign a link profile.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Properties** area, choose the link profile that you want to assign.
Step 6 Click **Save Changes**.

Assigning a Link Profile to a Port Channel FCoE Interface

Procedure

Step 1 In the **Navigation** pane, click **SAN**.
Step 2 On the **SAN** tab, expand **SAN** > **SAN Cloud** > ***Fabric*** > **FCoE Port Channels**
Step 3 Expand the FCoE port channel node and click the FCoE Interface where you want to assign a link profile.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Properties** area, choose the link profile that you want to assign.
Step 6 Click **Save Changes**.

Assigning a Link Profile to an Uplink FCoE Interface

Procedure

Step 1 In the **Navigation** pane, click **SAN**.
Step 2 On the **SAN** tab, expand **SAN** > **SAN Cloud** > ***Fabric*** > **Uplink FC Interfaces**
Step 3 Click the FC0E interface where you want to assign a link profile.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Properties** area, choose the link profile that you want to assign.
Step 6 Click **Save Changes**.

Configuring VMQ Connection Policies

VMQ Connection Policy

Cisco UCS Manager enables you to configure VMQ connection policy for a vNIC. VMQ provides improved network performance to the entire management operating system. Configuring a VMQ vNIC connection policy involves the following:

- Create a VMQ connection policy
- Create a static vNIC in a service profile
- Apply the VMQ connection policy to the vNIC

If you want to configure the VMQ vNIC on a service profile for a server, at least one adapter in the server must support VMQ. Make sure the servers have at least one the following adapters installed:

- UCS-VIC-M82-8P
- UCSB-MLOM-40G-01
- UCSC-PCIE-CSC-02

The following are the supported Operating Systems for VMQ:

- Windows 2012
- Windows 2012R2

You can apply only any one of the vNIC connection policies on a service profile at any one time. Make sure to select one of the three options such as Dynamic, usNIC or VMQ connection policy for the vNIC. When a VMQ vNIC is configured on service profile, make sure you have the following settings:

- Select SRIOV in the BIOS policy.
- Select Windows in the Adapter policy.

Creating a VMQ Connection Policy

Before you create a VMQ connection policy, consider the following:

- VMQ Tuning on the Windows Server — When an adapter is placed on a virtual switch, running the **Get-NetAdapterVmq** cmdlet displays True for VMQ. For more information on NIC teaming see Performance Tuning for Hyper-V Servers .
- Virtual machine level — By default, VMQ is enabled on all newly deployed VMs. VMQ can be enabled or disabled on existing VMs.
- Microsoft SCVMM — VMQ must be enabled on the port profile. If not, you will not be able to successfully create the virtual switch in SCVMM.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the **VMQ Connection Policies** node and select **Create VMQ Connection Policy**.

Step 5 In the **Create VMQ Connection Policy** dialog box, complete the following fields:

Name	Description
Name field	The VMQ connection policy name.
Description field	The description of the VMQ connection policy.
Number of VMQs field	The number of VMQs per adapter must be one more than the maximum number of VM NICs. **Note** Make sure that the total number of synthetic NICs present on the VMs is either equal to or greater than the number of VMs.
Number of Interrupts field	The number of CPU threads or logical processors available in the server. **Note** You cannot set this value to be more than the maximum number of available CPUs.

Step 6 Click **OK**.

Assigning Virtualization Preference to a vNIC

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 On the **Servers** tab, expand **Servers** > **target service profile** > **root** > **vNICs**.

Step 3 Click on the vNIC name to display properties on the work pane.

Step 4 In the **Connection Policies** section, select the radio button for **VMQ** and select the **VMQ Connection Policy** from the drop down.
In the **Properties** area **Virtualization Preference** for this vNIC changes to **VMQ**.

Enabling VMQ and NVGRE Offloading on the same vNIC

Perform the tasks in the table below to enable VMQ and NVGRE offloading on the same vNIC.

Note Currently, VMQ is not supported along with VXLAN on the same vNIC.

Task	Description	See
Enable normal NVGRE offloading	Perform this task by setting the corresponding flags in the adapter profile which is associated with the given vNIC. **Note** The Transmit checksum offload and TSO must be enabled for the NVGRE offloading to be effective.	Configuring an Ethernet Adapter Policy to Enable Stateless Offloads with NVGRE, on page 299 Applying an NVGRE Adapter Policy to a vNIC, on page 322
Enable VMQ	Perform this task by setting the appropriate connection policy when you add a vNIC to the service profile.	Creating a VMQ Connection Policy, on page 319 Assigning Virtualization Preference to a vNIC, on page 320

Configuring an Ethernet Adapter Policy to Enable Stateless Offloads with NVGRE

Cisco UCS Manager supports stateless offloads with NVGRE only with Cisco UCS VIC 1340 and/or Cisco UCS VIC 1380 adapters that are installed on servers running Windows Server 2012 R2 operating systems. Stateless offloads with NVGRE cannot be used with NetFlow, usNIC, or VM-FEX.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **Adapter Policies** and choose **Create Ethernet Adapter Policy**.

a) In the **Resources** area, set the following options:

- Transmit Queues = 1
- Receive Queues = n (up to 8)
- Completion Queues = # of Transmit Queues + # of Receive Queues

- Interrupts = # Completion Queues + 2

b) In the **Options** area, set the following options:

- Network Virtualization using Generic Routing Encapsulation = Enabled
- Interrupt Mode = Msi-X

For more information on creating an Ethernet adapter policy, see Creating an Ethernet Adapter Policy, on page 294.

Step 5 Click **OK** to create the Ethernet adapter policy.

Step 6 Install an eNIC driver Version 3.0.0.8 or later.
For more information, see http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/vic_drivers/install/Windows/b_Cisco_VIC_Drivers_for_Windows_Installation_Guide.html.

Step 7 Reboot the server.

Applying an NVGRE Adapter Policy to a vNIC

Procedure

Step 1 In the **Navigation** pane, click the **Servers** tab.

Step 2 On the **Servers** tab, expand **Servers** > **Target Service Profile** > **root** > **vNICS**

Step 3 Click on the vNIC name to display properties in the work pane.

Step 4 In the **Policies** section, select the NVGRE policy from **Adapter Policy** drop-down list.

Step 5 Click **Save Changes** to apply the policy to the vNIC.

Creating a VMQ Connection Policy

Before you create a VMQ connection policy, consider the following:

- VMQ Tuning on the Windows Server — When an adapter is placed on a virtual switch, running the **Get-NetAdapterVmq** cmdlet displays True for VMQ. For more information on NIC teaming see Performance Tuning for Hyper-V Servers .
- Virtual machine level — By default, VMQ is enabled on all newly deployed VMs. VMQ can be enabled or disabled on existing VMs.
- Microsoft SCVMM — VMQ must be enabled on the port profile. If not, you will not be able to successfully create the virtual switch in SCVMM.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the **VMQ Connection Policies** node and select **Create VMQ Connection Policy**.

Step 5 In the **Create VMQ Connection Policy** dialog box, complete the following fields:

Name	Description
Name field	The VMQ connection policy name.
Description field	The description of the VMQ connection policy.
Number of VMQs field	The number of VMQs per adapter must be one more than the maximum number of VM NICs. **Note** Make sure that the total number of synthetic NICs present on the VMs is either equal to or greater than the number of VMs.
Number of Interrupts field	The number of CPU threads or logical processors available in the server. **Note** You cannot set this value to be more than the maximum number of available CPUs.

Step 6 Click **OK**.

Assigning Virtualization Preference to a vNIC

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 On the **Servers** tab, expand **Servers** > **target service profile** > **root** > **vNICs**.

Step 3 Click on the vNIC name to display properties on the work pane.

Step 4 In the **Connection Policies** section, select the radio button for **VMQ** and select the **VMQ Connection Policy** from the drop down.
In the **Properties** area **Virtualization Preference** for this vNIC changes to **VMQ**.

NetQueue

Information About NetQueue

NetQueue improves traffic performance by providing a network adapter with multiple receive queues. These queues allow the data interrupt processing that is associated with individual virtual machines to be grouped.

Note NetQueue is supported on servers running VMware ESXi operating systems.

Configuring NetQueue

Procedure

Step 1 Create a Virtual Machine Queue (VMQ) connection policy.

Step 2 Configure NetQueues in a service profile by selecting the VMQ connection policy.
Use the following when you are configuring NetQueue:

- The default ring size is rx512, tx256
- The interrupt count on each VNIC is VMQ count x 2 +2

  Note The number of interrupts depends on the number of NetQueues enabled.

- The driver supports up to 16 NetQueues per port for standard frame configurations.

  Note VMware recommends that you use up to eight NetQueues per port for standard frame configurations.

- NetQueue should be enabled only on MSIX systems.
- You should disable NetQueue on 1 GB NICs.

Step 3 Enable the MSIX mode in the adapter policy for NetQueue.

Step 4 Associate the service profile with the server.

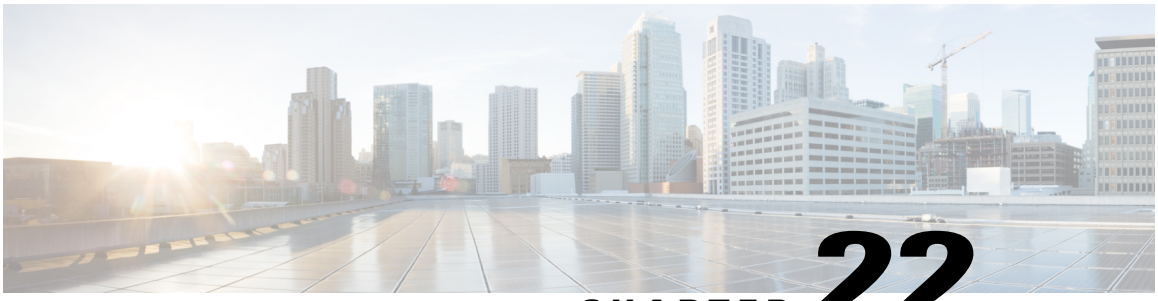

CHAPTER 22

Configuring Upstream Disjoint Layer-2 Networks

This chapter includes the following sections:

Upstream Disjoint Layer-2 Networks

Upstream disjoint layer-2 networks (disjoint L2 networks) are required if you have two or more Ethernet clouds that never connect, but must be accessed by servers or virtual machines located in the same Cisco UCS domain. For example, you could configure disjoint L2 networks if you require one of the following:

- Servers or virtual machines to access a public network and a backup network
- Servers or virtual machines for more than one customer are located in the same Cisco UCS domain, and that need to access the L2 networks for both customers in a multi-tenant system

Note By default, data traffic in Cisco UCS works on a principle of mutual inclusion. All traffic for all VLANs and upstream networks travels along all uplink ports and port channels. If you have upgraded from a release that does not support upstream disjoint layer-2 networks, you must assign the appropriate uplink interfaces to your VLANs, or traffic for those VLANs continues to flow along all uplink ports and port channels.

The configuration for disjoint L2 networks works on a principle of selective exclusion. Traffic for a VLAN that is designated as part of a disjoint network can only travel along an uplink Ethernet port or port channel

that is specifically assigned to that VLAN, and is selectively excluded from all other uplink ports and port channels. However, traffic for VLANs that are not specifically assigned to an uplink Ethernet port or port channel can still travel on all uplink ports or port channels, including those that carry traffic for the disjoint L2 networks.

In Cisco UCS, the VLAN represents the upstream disjoint L2 network. When you design your network topology for disjoint L2 networks, you must assign uplink interfaces to VLANs not the reverse.

For information about the maximum number of supported upstream disjoint L2 networks, see the appropriate *Cisco UCS Configuration Limits for Cisco UCS Manager Guide*.

Guidelines for Configuring Upstream Disjoint L2 Networks

When you plan your configuration for upstream disjoint L2 networks, consider the following:

Ethernet Switching Mode Must Be End-Host Mode

Cisco UCS only supports disjoint L2 networks when the Ethernet switching mode of the fabric interconnects is configured for end-host mode. You cannot connect to disjoint L2 networks if the Ethernet switching mode of the fabric interconnects is switch mode.

Symmetrical Configuration Is Recommended for High Availability

If a Cisco UCS domain is configured for high availability with two fabric interconnects, we recommend that both fabric interconnects are configured with the same set of VLANs.

VLAN Validity Criteria Are the Same for Uplink Ethernet Ports and Port Channels

The VLAN used for the disjoint L2 networks must be configured and assigned to an uplink Ethernet port or uplink Ethernet port channel. If the port or port channel does not include the VLAN, Cisco UCS Manager considers the VLAN invalid and does the following:

- Displays a configuration warning in the **Status Details** area for the server.
- Ignores the configuration for the port or port channel and drops all traffic for that VLAN.

Note The validity criteria are the same for uplink Ethernet ports and uplink Ethernet port channels. Cisco UCS Manager does not differentiate between the two.

Overlapping VLANs Are Not Supported

Cisco UCS does not support overlapping VLANs in disjoint L2 networks. You must ensure that each VLAN only connects to one upstream disjoint L2 domain.

Each vNIC Can Only Communicate with One Disjoint L2 Network

A vNIC can only communicate with one disjoint L2 network. If a server needs to communicate with multiple disjoint L2 networks, you must configure a vNIC for each of those networks.

To communicate with more than two disjoint L2 networks, a server must have a Cisco VIC adapter that supports more than two vNICs.

Appliance Port Must Be Configured with the Same VLAN as Uplink Ethernet Port or Port Channel

For an appliance port to communicate with a disjoint L2 network, you must ensure that at least one uplink Ethernet port or port channel is in the same network and is therefore assigned to the same VLANs that are used by the appliance port. If Cisco UCS Manager cannot identify an uplink Ethernet port or port channel that includes all VLANs that carry traffic for an appliance port, the appliance port experiences a pinning failure and goes down.

For example, a Cisco UCS domain includes a global VLAN named vlan500 with an ID of 500. vlan500 is created as a global VLAN on the uplink Ethernet port. However, Cisco UCS Manager does not propagate this VLAN to appliance ports. To configure an appliance port with vlan500, you must create another VLAN named vlan500 with an ID of 500 for the appliance port. You can create this duplicate VLAN in the **Appliances** node on the **LAN** tab of the Cisco UCS Manager GUI or the **eth-storage** scope in the Cisco UCS Manager CLI. If you are prompted to check for VLAN Overlap, accept the overlap and Cisco UCS Manager creates the duplicate VLAN for the appliance port.

Default VLAN 1 Cannot Be Configured Explicitly on an Uplink Ethernet Port or Port Channel

Cisco UCS Manager implicitly assigns default VLAN 1 to all uplink ports and port channels. Even if you do not configure any other VLANs, Cisco UCS uses default VLAN 1 to handle data traffic for all uplink ports and port channels.

Note After you configure VLANs in a Cisco UCS domain, default VLAN 1 remains implicitly on all uplink ports and port channels. You cannot explicitly assign default VLAN 1 to an uplink port or port channel, nor can you remove it from an uplink port or port channel.

If you attempt to assign default VLAN 1 to a specific port or port channel, Cisco UCS Manager raises an Update Failed fault.

Therefore, if you configure a Cisco UCS domain for disjoint L2 networks, do not configure any vNICs with default VLAN 1 unless you want all data traffic for that server to be carried on all uplink Ethernet ports and port channels and sent to all upstream networks.

VLANs for Both FIs Must be Concurrently Assigned

When you assign a port to a global VLAN, the VLAN is removed from all of the ports that are not explicitly assigned to the VLAN on both fabric interconnects. The ports on both FIs must be configured at the same time. If the ports are only configured on the first FI, traffic on the second FI will be disrupted.

Pinning Considerations for Upstream Disjoint L2 Networks

Communication with an upstream disjoint L2 network requires that you ensure that the pinning is properly configured. Whether you implement soft pinning or hard pinning, a VLAN membership mismatch causes traffic for one or more VLANs to be dropped.

Soft Pinning

Soft pinning is the default behavior in Cisco UCS. If you plan to implement soft pinning, you do not need to create LAN pin groups to specify a pin target for a vNIC. Instead, Cisco UCS Manager pins the vNIC to an uplink Ethernet port or port channel according to VLAN membership criteria.

With soft pinning, Cisco UCS Manager validates data traffic from a vNIC against the VLAN membership of all uplink Ethernet ports and port channels. If you have configured disjoint L2 networks, Cisco UCS Manager must be able to find an uplink Ethernet port or port channel that is assigned to all VLANS on the vNIC. If no uplink Ethernet port or port channel is configured with all VLANs on the vNIC, Cisco UCS Manager does the following:

- Brings the link down.
- Drops the traffic for all of the VLANs on the vNIC.
- Raises the following faults:
  - Link Down
  - VIF Down

Cisco UCS Manager does not raise a fault or warning about the VLAN configuration.

For example, a vNIC on a server is configured with VLANs 101, 102, and 103. Interface 1/3 is assigned only to VLAN 102. Interfaces 1/1 and 1/2 are not explicitly assigned to a VLAN, which makes them available for traffic on VLANs 101 and 103. As a result of this configuration, the Cisco UCS domain does not include a border port interface that can carry traffic for all three VLANS for which the vNIC is configured. As a result, Cisco UCS Manager brings down the vNIC, drops traffic for all three VLANs on the vNIC, and raises the Link Down and VIF Down faults.

Hard Pinning

Hard pinning occurs when you use LAN pin groups to specify the pinning target for the traffic intended for the disjoint L2 networks. In turn, the uplink Ethernet port or port channel that is the pinning target must be configured to communicate with the appropriate disjoint L2 network.

With hard pinning, Cisco UCS Manager validates data traffic from a vNIC against the VLAN membership of all uplink Ethernet ports and port channels, and validates the LAN pin group configuration to ensure it includes the VLAN and the uplink Ethernet port or port channel. If the validation fails at any point, Cisco UCS Manager does the following:

- Raises a Pinning VLAN Mismatch fault with a severity of Warning.
- Drops traffic for the VLAN.
- Does not bring the link down, so that traffic for other VLANs can continue to flow along it.

For example, if you want to configure hard pinning for an upstream disjoint L2 network that uses VLAN 177, do the following:

- Create a LAN pin group with the uplink Ethernet port or port channel that carries the traffic for the disjoint L2 network.
- Configure at least one vNIC in the service profile with VLAN 177 and the LAN pin group.
- Assign VLAN 177 to an uplink Ethernet port or port channel included in the LAN pin group

If the configuration fails at any of these three points, then Cisco UCS Manager warns for a VLAN mismatch for VLAN 177 and drops the traffic for that VLAN only.

Note If changes are made to soft pinning configurations resulting in vNIC VLANs not resolving with disjoint L2 uplink, a warning dialog box is displayed. The warning dialog box allows you to proceed with your configuration or cancel it. If you decide to proceed with the mis- configuration, you will experience a reduction is server traffic performance.

Configuring Cisco UCS for Upstream Disjoint L2 Networks

When you configure a Cisco UCS domain to connect with upstream disjoint L2 networks, you need to ensure that you complete all of the following steps.

Before You Begin

Before you begin this configuration, ensure that the ports on the fabric interconnects are properly cabled to support your disjoint L2 networks configuration.

Procedure

	Command or Action	Purpose
Step 1	Configure Ethernet switching mode for both fabric interconnects in Ethernet End-Host Mode.	The Ethernet switching mode must be in End-Host Mode for Cisco UCS to be able to communicate with upstream disjoint L2 networks. See Configuring Ethernet Switching Mode.
Step 2	Configure the ports and port channels that you require to carry traffic for the disjoint L2 networks.	See Configuring Ports and Port Channels, on page 79.
Step 3	Configure the LAN pin groups required to pin the traffic for the appropriate uplink Ethernet ports or port channels.	(Optional) See Configuring LAN Pin Groups, on page 269.
Step 4	Create one or more VLANs.	These can be named VLANs or private VLANs. For a cluster configuration, we recommend that you create the VLANs in the VLAN Manager and use the Common/Global configuration to ensure they are accessible to both fabric interconnects. See Creating a VLAN for an Upstream Disjoint L2 Network, on page 330.
Step 5	Assign the desired ports or port channels to the VLANs for the disjoint L2 networks.	When this step is completed, traffic for those VLANs can only be sent through the trunks for the assigned ports and/or port channels. Assigning Ports and Port Channels to VLANs, on page 330

	Command or Action	Purpose
Step 6	Ensure that the service profiles for all servers that need to communicate with the disjoint L2 networks include the correct LAN connectivity configuration to ensure the vNICs send the traffic to the appropriate VLAN.	You can complete this configuration through one or more vNIC templates or when you configure the networking options for the service profile. See Service Profiles, on page 589.

Creating a VLAN for an Upstream Disjoint L2 Network

For upstream disjoint L2 networks, we recommend that you create VLANs in the VLAN Manager.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 On the **LAN** tab, click the **LAN** node.

Step 3 In the **Work** pane, click the **LAN Uplinks Manager** link on the **LAN Uplinks** tab.
The LAN Uplinks Manager opens in a separate window.

Step 4 In the LAN Uplinks Manager, click **VLANs** > **VLAN Manager**.
You can create the VLAN on any of the subtabs. However, if you use the **All** subtab, you can view all of the configured VLANs in the table.

Step 5 On the icon bar to the right of the table, click +.
If the + icon is disabled, click an entry in the table to enable it.

Step 6 In the **Create VLANs** dialog box, specify the required fields and then click **OK**.
You cannot create VLANs with IDs from 3968 to 4047. This range of VLAN IDs is reserved.

Step 7 Repeat Steps 6 and 7 to create additional VLANs.

What to Do Next

Assign ports and port channels to the VLANs.

Assigning Ports and Port Channels to VLANs

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 On the **LAN** tab, click the **LAN** node.

Step 3 In the **Work** pane, click the **LAN Uplinks Manager** link on the **LAN Uplinks** tab.

The LAN Uplinks Manager opens in a separate window.

Step 4 In the LAN Uplinks Manager, click **VLANs** > **VLAN Manager**.
You can create the VLAN on any of the subtabs. However, if you use the **All** subtab, you can view all of the configured VLANs in the table.

Step 5 Click one of the following subtabs to configure ports and port channels on that fabric interconnect:

Subtab	Description
Fabric A	Displays the ports, port channels, and VLANs that are accessible to fabric interconnect A.
Fabric B	Displays the ports, port channels, and VLANs that are accessible to fabric interconnect B.

Step 6 In the **Ports and Port Channels** table, do the following:

- To assign an Uplink Ethernet port channel to a VLAN, expand the **Port Channels** node and click the port channel you want to assign to the VLAN.
- To assign an Uplink Ethernet port to the VLAN, expand the **Uplink Interfaces** node and click the port you want to assign to the VLAN

You can hold down the Ctrl key and click multiple ports or port channels to assign to them to the same VLAN or set of VLANs .

Step 7 In the **VLANs** table, expand the appropriate node if necessary and click the VLAN to which you want to assign the port or port channel.
You can hold down the Ctrl key and click multiple VLANs if you want to assign the same set of ports and/or port channels to them.

Step 8 Click the **Add to VLAN/VLAN Group** button.

Step 9 If a confirmation dialog box displays, click **Yes**.

Step 10 To assign additional ports or port channels to VLANs on the same fabric, repeat Steps 6, 7, and 8.

Step 11 To assign additional ports or port channels to VLANs on a different fabric, repeat Steps 5 through 8.
If the Cisco UCS domain is configured for high availability with two fabric interconnects, we recommend that you create the same set of VLANs on both fabric interconnects.

Step 12 If a confirmation dialog box displays, click **Yes**.

Step 13 Click **Apply** if you want to continue to work in the VLAN Manager, or click **OK** to close the window.
After a port or port channel is assigned to one or more VLANs, it is removed from all other VLANs.

Removing Ports and Port Channels from VLANs

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 On the **LAN** tab, click the **LAN** node.

Step 3 In the **Work** pane, click the **LAN Uplinks Manager** link on the **LAN Uplinks** tab.
The LAN Uplinks Manager opens in a separate window.

Step 4 In the LAN Uplinks Manager, click **VLANs** > **VLAN Manager**.
You can create the VLAN on any of the subtabs. However, if you use the **All** subtab, you can view all of the configured VLANs in the table.

Step 5 Click one of the following subtabs to configure ports and port channels on that fabric interconnect:

Subtab	Description
Fabric A	Displays the ports, port channels, and VLANs that are accessible to fabric interconnect A.
Fabric B	Displays the ports, port channels, and VLANs that are accessible to fabric interconnect B.

Step 6 In the **VLANs** table, expand the appropriate node and the VLAN from which you want to remove a port or port channel.

Step 7 Click the port or port channel that you want to remove from the VLAN.
Hold down the Ctrl key to click multiple ports or port channels.

Step 8 Click the **Remove from VLAN/VLAN Group** button.

Step 9 If a confirmation dialog box displays, click **Yes**.

Step 10 Click **Apply** if you want to continue to work in the VLAN Manager, or click **OK** to close the window.

Important If you remove all port or port channel interfaces from a VLAN, the VLAN returns to the default behavior and data traffic on that VLAN flows on all uplink ports and port channels. Based on the configuration in the Cisco UCS domain, this default behavior can cause Cisco UCS Manager to drop traffic for that VLAN. To avoid this occurrence, Cisco recommends that you assign at least one interface to the VLAN or delete the VLAN.

Viewing Ports and Port Channels Assigned to VLANs

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 On the **LAN** tab, click the **LAN** node.

Step 3 In the **Work** pane, click the **LAN Uplinks Manager** link on the **LAN Uplinks** tab.
The LAN Uplinks Manager opens in a separate window.

Step 4 In the LAN Uplinks Manager, click **VLANs** > **VLAN Manager**.
You can create the VLAN on any of the subtabs. However, if you use the **All** subtab, you can view all of the configured VLANs in the table.

Step 5 Click one of the following subtabs to configure ports and port channels on that fabric interconnect:

Subtab	Description
Fabric A	Displays the ports, port channels, and VLANs that are accessible to fabric interconnect A.
Fabric B	Displays the ports, port channels, and VLANs that are accessible to fabric interconnect B.

Step 6 In the **VLANs** table, expand the appropriate node and the VLAN for which you want to view the assigned ports or port channels.

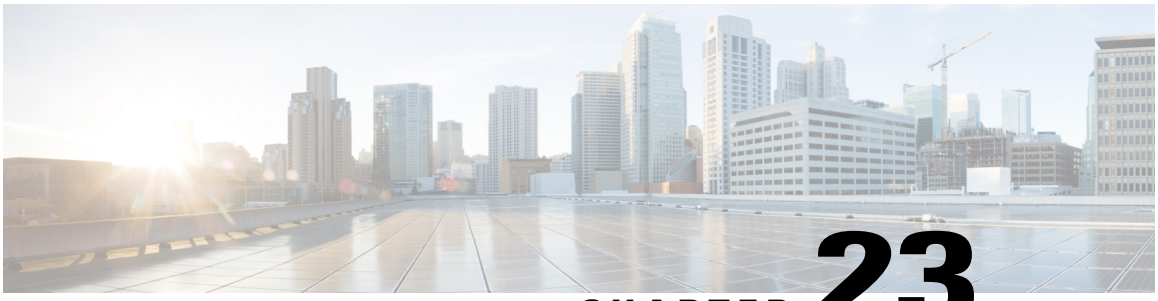

CHAPTER 23

Configuring Named VSANs

This chapter includes the following sections:

Named VSANs

A named VSAN creates a connection to a specific external SAN. The VSAN isolates traffic to that external SAN, including broadcast traffic. The traffic on one named VSAN knows that the traffic on another named VSAN exists, but cannot read or access that traffic.

Like a named VLAN, the name that you assign to a VSAN ID adds a layer of abstraction that allows you to globally update all servers associated with service profiles that use the named VSAN. You do not need to reconfigure the servers individually to maintain communication with the external SAN. You can create more than one named VSAN with the same VSAN ID.

Named VSANs in Cluster Configurations

In a cluster configuration, a named VSAN can be configured to be accessible only to the Fibre Channel uplink ports on one fabric interconnect or to the Fibre Channel uplink ports on both fabric interconnects.

Named VSANs and the FCoE VLAN ID

You must configure each named VSAN with an FCoE VLAN ID. This property determines which VLAN is used for transporting the VSAN and its Fibre Channel packets.

For FIP-capable, converged network adapters, such as the Cisco UCS CNA M72KR-Q and the Cisco UCS CNA M72KR-E, the named VSAN must be configured with a named VLAN that is not the native VLAN for the FCoE VLAN ID. This configuration ensures that FCoE traffic can pass through these adapters.

In the following sample configuration, a service profile with a vNIC and vHBA mapped to fabric A is associated with a server that has FIP capable, converged network adapters:

- The vNIC is configured to use VLAN 10.
- VLAN 10 is also designated as the native VLAN for the vNIC.
- The vHBA is configured to use VSAN 2.
- Therefore, VSAN 2 cannot be configured with VLAN 10 as the FCoE VLAN ID. VSAN 2 can be mapped to any other VLAN configured on fabric A.

Fibre Channel Uplink Trunking for Named VSANs

You can configure Fibre Channel uplink trunking for the named VSANs on each fabric interconnect. If you enable trunking on a fabric interconnect, all named VSANs in a Cisco UCS domain are allowed on all Fibre Channel uplink ports on that fabric interconnect.

Guidelines and Recommendations for VSANs

The following guidelines and recommendations apply to all named VSANs, including storage VSANs.

VSAN 4079 is a Reserved VSAN ID

Do not configure a VSAN as 4079. This VSAN is reserved and cannot be used in either FC switch mode or FC end-host mode.

If you create a named VSAN with ID 4079, Cisco UCS Manager marks that VSAN with an error and raises a fault.

Reserved VSAN Range for Named VSANs in FC Switch Mode

If you plan to use FC switch mode in a Cisco UCS domain, do not configure VSANs with an ID in the range from 3040 to 4078.

VSANs in that range are not operational if the fabric interconnects are configured to operate in FC switch mode. Cisco UCS Manager marks that VSAN with an error and raises a fault.

Reserved VSAN Range for Named VSANs in FC End-Host Mode

If you plan to use FC end-host mode in a Cisco UCS domain, do not configure VSANs with an ID in the range from 3840 to 4079.

VSANs in that range are not operational if the following conditions exist in a Cisco UCS domain:

- The fabric interconnects are configured to operate in FC end-host mode.
- The Cisco UCS domain is configured with Fibre Channel trunking or SAN port channels.

If these configurations exist, Cisco UCS Manager does the following:

1 Renders all VSANs with an ID in the range from 3840 to 4079 non-operational.

2 Raises a fault against the non-operational VSANs.

3 Transfers all non-operational VSANs to the default VSAN.

4 Transfers all vHBAs associated with the non-operational VSANs to the default VSAN.

If you disable Fibre Channel trunking and delete any existing SAN port channels, Cisco UCS Manager returns all VSANs in the range from 3840 to 4078 to an operational state and restores any associated vHBAs back to those VSANs.

Range Restrictions for Named VSAN IDs in FC Switch Mode

If you plan to use FC switch mode in a Cisco UCS domain, do not configure VSANs in the range from 3040 to 4078.

When a fabric interconnect operating in FC switch mode is connected to MDS as the upstream switch, VSANs configured in Cisco UCS Manager in the range from 3040 to 4078 and assigned as port VSANs cannot be created in MDS. This configuration results in a possible port VSAN mismatch.

Guidelines for FCoE VLAN IDs

Note FCoE VLANs in the SAN cloud and VLANs in the LAN cloud must have different IDs. Using the same ID for an FCoE VLAN in a VSAN and a VLAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that FCoE VLAN. Ethernet traffic is dropped on any VLAN with an ID that overlaps with an FCoE VLAN ID.

VLAN 4048 is user configurable. However, Cisco UCS Manager uses VLAN 4048 for the following default values. If you want to assign 4048 to a VLAN, you must reconfigure these values:

- After an upgrade to Cisco UCS, Release 2.0—The FCoE storage port native VLAN uses VLAN 4048 by default. If the default FCoE VSAN was set to use VLAN 1 before the upgrade, you must change it to a VLAN ID that is not used or reserved. For example, consider changing the default to 4049 if that VLAN ID is not in use.
- After a fresh install of Cisco UCS, Release 2.0—The FCoE VLAN for the default VSAN uses VLAN 4048 by default. The FCoE storage port native VLAN uses VLAN 4049.

Creating a Named VSAN

Note FCoE VLANs in the SAN cloud and VLANs in the LAN cloud must have different IDs. Using the same ID for an FCoE VLAN in a VSAN and a VLAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that FCoE VLAN. Ethernet traffic is dropped on any VLAN with an ID that overlaps with an FCoE VLAN ID.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **SAN Cloud**.

Step 3 In the **Work** pane, click the **VSANs** tab.

Step 4 On the icon bar to the right of the table, click +.
If the + icon is disabled, click an entry in the table to enable it.

Step 5 In the **Create VSAN** dialog box, complete the required.

Step 6 Click **OK**.
Cisco UCS Manager GUI adds the VSAN to one of the following **VSANs** nodes:

- The **SAN Cloud** > **VSANs** node for a storage VSAN accessible to both fabric interconnects.
- The **SAN Cloud** > ***Fabric_Name*** > **VSANs** node for a VSAN accessible to only one fabric interconnect.

Creating a Storage VSAN

Note FCoE VLANs in the SAN cloud and VLANs in the LAN cloud must have different IDs. Using the same ID for an FCoE VLAN in a VSAN and a VLAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that FCoE VLAN. Ethernet traffic is dropped on any VLAN with an ID that overlaps with an FCoE VLAN ID.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 On the **SAN** tab, expand **SAN** > **Storage Cloud**.

Step 3 In the **Work** pane, click the **VSANs** tab.

Step 4 On the icon bar to the right of the table, click +.
If the + icon is disabled, click an entry in the table to enable it.

Step 5 In the **Create VSAN** dialog box, complete the required fields.

Step 6 Click **OK**.
Cisco UCS Manager GUI adds the VSAN to one of the following **VSANs** nodes:

- The **Storage Cloud** > **VSANs** node for a storage VSAN accessible to both fabric interconnects.
- The **Storage Cloud** > ***Fabric_Name*** > **VSANs** node for a VSAN accessible to only one fabric interconnect.

Deleting a VSAN

If Cisco UCS Manager includes a named VSAN with the same VSAN ID as the one you delete, the VSAN is not removed from the fabric interconnect configuration until all named VSANs with that ID are deleted.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 In the **SAN** tab, click the **SAN** node.

Step 3 In the **Work** pane, click the **VSANs** tab.

Step 4 Click one of the following subtabs, depending upon what type of VSAN you want to delete:

Subtab	Description
All	Displays all VSANs in the Cisco UCS domain.
Dual Mode	Displays the VSANs that are accessible to both fabric interconnects.
Switch A	Displays the VSANs that are accessible to only fabric interconnect A.
Switch B	Displays the VSANs that are accessible to only fabric interconnect B.

Step 5 In the table, click the VSAN you want to delete.
You can use the Shift key or Ctrl key to select multiple entries.

Step 6 Right-click the highlighted VSAN or VSANs and choose **Delete**.

Step 7 If a confirmation dialog box displays, click **Yes**.

Changing the VLAN ID for the FCoE VLAN for a Storage VSAN

Caution Changing the VLAN ID of the FCoE VLAN for a storage VSAN causes an brief traffic outage.

Changing the FCoE VLAN for the default VSAN or any configured VSAN under a global policy may result in storage disconnect or complete shut down.

Note FCoE VLANs in the SAN cloud and VLANs in the LAN cloud must have different IDs. Using the same ID for an FCoE VLAN in a VSAN and a VLAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that FCoE VLAN. Ethernet traffic is dropped on any VLAN with an ID that overlaps with an FCoE VLAN ID.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.
Step 2 On the **SAN** tab, expand **SAN** > **Storage Cloud** > **VSANs**.
Step 3 Choose the VSAN for which you want to modify the FCoE VLAN ID.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **FCoE VLAN** field, enter the desired VLAN ID.
Step 6 Click **Save Changes**.

Enabling Fibre Channel Uplink Trunking

Note If the fabric interconnects are configured for Fibre Channel end-host mode, enabling Fibre Channel uplink trunking renders all VSANs with an ID in the range from 3840 to 4079 non-operational.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.
Step 2 Expand **SAN** > **SAN Cloud**.
Step 3 Click the node for the fabric where you want to enable FC uplink trunking.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Enable FC Uplink Trunking**.
Step 6 If a confirmation dialog box displays, click **Yes**.

Disabling Fibre Channel Uplink Trunking

Procedure

Step 1 In the **Navigation** pane, click **SAN**.
Step 2 Expand **SAN** > **SAN Cloud**.
Step 3 Click the node for the fabric where you want to disable Fibre Channel uplink trunking.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Disable FC Uplink Trunking**.
Step 6 If a confirmation dialog box displays, click **Yes**.

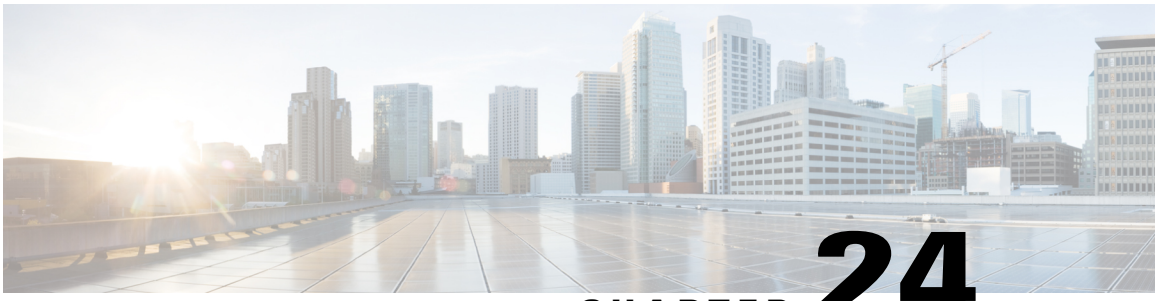

CHAPTER 24

Configuring SAN Pin Groups

This chapter includes the following sections:

SAN Pin Groups

Cisco UCS uses SAN pin groups to pin Fibre Channel traffic from a vHBA on a server to an uplink Fibre Channel port on the fabric interconnect. You can use this pinning to manage the distribution of traffic from the servers.

Note In Fibre Channel switch mode, SAN pin groups are irrelevant. Any existing SAN pin groups will be ignored.

To configure pinning for a server, you must include the SAN pin group in a vHBA policy. The vHBA policy is then included in the service profile assigned to that server. All traffic from the vHBA will travel through the I/O module to the specified uplink Fibre Channel port.

You can assign the same pin group to multiple vHBA policies. As a result, you do not need to manually pin the traffic for each vHBA.

Important Changing the target interface for an existing SAN pin group disrupts traffic for all vHBAs which use that pin group. The fabric interconnect performs a log in and log out for the Fibre Channel protocols to re-pin the traffic.

Creating a SAN Pin Group

In a system with two fabric interconnects, you can associate the pin group with only one fabric interconnect or with both fabric interconnects.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 In the **SAN** tab, expand **SAN** > **SAN Cloud**.

Step 3 Right-click **SAN Pin Groups** and select **Create SAN Pin Group**.

Step 4 Enter a unique name and description for the pin group.

Step 5 To pin traffic for fabric interconnect A, do the following in the **Targets** area:

a) Check the **Fabric A** check box.

b) Click the drop-down arrow on the **Interface** field and navigate through the tree-style browser to select the uplink Fibre Channel port you want to associate with the pin group.

Step 6 To pin traffic for fabric interconnect B, do the following in the **Targets** area:

a) Check the **Fabric B** check box.

b) Click the drop-down arrow on the **Interface** field and navigate through the tree-style browser to select the uplink Fibre Channel port you want to associate with the pin group.

Step 7 Click **OK**.

What to Do Next

Include the pin group in a vHBA template.

Deleting a SAN Pin Group

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 In the **SAN** tab, expand **SAN** > **SAN Cloud** > **SAN Pin Groups**.

Step 3 Right-click the SAN pin group you want to delete and select **Delete**.

Step 4 If a confirmation dialog box displays, click **Yes**.

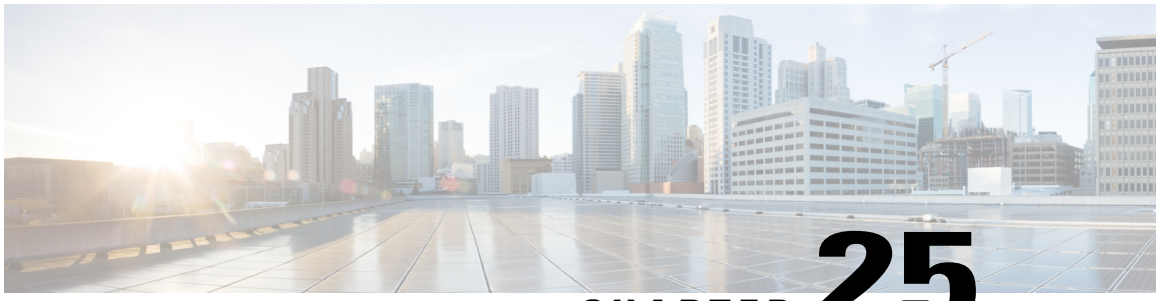

CHAPTER 25

Configuring WWN Pools

This chapter includes the following sections:

WWN Pools

A World Wide Name (WWN) pool is a collection of WWNs for use by the Fibre Channel vHBAs in a Cisco UCS domain. You create separate pools for the following:

- WW node names assigned to the vHBA
- WW port names assigned to the vHBA
- Both WW node names and WW port names

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. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNs in the SAN fabric, Cisco recommends using the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:XX:XX:XX

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.

WWNN Pools

A WWNN pool is a WWN pool that 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 that 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.

WWxN Pools

A WWxN pool is a WWN pool that contains both WW node names and WW port names. You can specify how many ports per node are created with WWxN pools. The pool size must be a multiple of *ports-per-node* + 1. For example, if you specify 7 ports per node, the pool size must be a multiple of 8. If you specify 63 ports per node, the pool size must be a multiple of 64.

You can use a WWxN pool whenever you select a WWNN or WWPN pool. The WWxN pool must be created before it can be assigned.

- For WWNN pools, the WWxN pool is displayed as an option in the **WWNN Assignment** drop-down list.
- For WWPN pools, choose **Derived** in the **WWPN Assignment** drop-down list.

Configuring WWNN Pools

Creating a WWNN Pool

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. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNs in the SAN fabric, Cisco recommends using the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:XX:XX:XX

A WWNN pool with the last four digits ending in 00:01 causes the vHBA to not initialize, no output from the lunlist command, and displays the Waiting for Flogi error. This error occurs if the WWPN is in the same block as the WWNN ending in 00:01. To ensure that the WWNN and WWPN addresses do not overlap, we recommend using a unique WWN address.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 In the **SAN** tab, expand **SAN** > **Pools**.

Step 3 Expand the node for the organization where you want to create the pool.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **WWNN Pools** and select **Create WWNN Pool**.

Step 5 In the **Define Name and Description** dialog box of the **Create WWNN Pool** wizard, complete the following fields:

Name	Description
Name field	The name of the World Wide Node Name pool. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the pool. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Assignment Order field	This can be one of the following: • **Default—**Cisco UCS Manager selects a random identity from the pool. • **Sequential—**Cisco UCS Manager selects the lowest available identity from the pool.

Step 6 Click **Next**.

Step 7 In the **Add WWN Blocks** page of the **Create WWNN Pool** wizard, click **Add**.

Step 8 In the **Create WWN Block** dialog box, complete the following fields:

Name	Description
From field	The first WWN in the block.
Size field	The number of WWNs in the block. For WWxN pools, the pool size must be a multiple of *ports-per-node* + 1. For example, if there are 7 ports per node, the pool size must be a multiple of 8. If there are 63 ports per node, the pool size must be a multiple of 64.

Step 9 Click **OK**.

Step 10 Click **Finish**.

What to Do Next

Include the WWNN pool in a service profile and template.

Adding a WWN Block to a WWNN Pool

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. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNs in the SAN fabric, Cisco recommends using the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:XX:XX:XX

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 In the **SAN** tab, expand **SAN** > **Pools** > ***Organization_Name*** .

Step 3 Expand the **WWNN Pools** node.

Step 4 Right-click the WWNN pool to which you want to add a WWN block and select **Create WWN Block**.

Step 5 In the **Create WWN Block** dialog box, complete the following fields:

Name	Description
From field	The first WWN in the block.
Size field	The number of WWNs in the block. For WWxN pools, the pool size must be a multiple of *ports-per-node* + 1. For example, if there are 7 ports per node, the pool size must be a multiple of 8. If there are 63 ports per node, the pool size must be a multiple of 64.

Step 6 Click **OK**.

Deleting a WWN Block from a WWNN Pool

If you delete an address block from a pool, Cisco UCS Manager does not reallocate any addresses in that block that were assigned to vNICs or vHBAs. All assigned addresses from a deleted block remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 In the **SAN** tab, expand **SAN** > **Pools** > *Organization_Name* > **WWNN Pools** > *WWNN_Pool_Name* .

Step 3 Right-click the WWN block that you want to delete and select **Delete**.

Step 4 If a confirmation dialog box displays, click **Yes**.

Adding a WWNN Initiator to a WWNN Pool

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. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNs in the SAN fabric, Cisco recommends using the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:XX:XX:XX

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 In the **SAN** tab, expand **SAN** > **Pools** > *Organization_Name* .

Step 3 Expand the **WWNN Pools** node.

Step 4 Right-click the WWNN pool to which you want to add a WWNN initiator and select **Create WWNN Initiiator**.

Step 5 In the **Create WWNN Initiator** dialog box, complete the following fields:

Name	Description
World Wide Name field	The WWN.
Name field	The name of the WWNN initiator. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A user-defined description of the WWNN initiator. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Step 6 Click **OK**.

Deleting a WWNN Initiator from a WWNN Pool

Procedure

Step 1 In the **Navigation** pane, click **SAN**.
Step 2 In the **SAN** tab, expand **SAN** > **Pools** > *Organization_Name* .
Step 3 Expand the **WWPN Pools** node.
Step 4 Choose the WWNN pool from which you want to delete a WWNN initiator.
Step 5 In the **Work** pane, click the **Initiators** tab.
Step 6 Right-click the initiator that you want to delete and choose **Delete**.
Step 7 If a confirmation dialog box displays, click **Yes**.

Deleting a WWNN Pool

If you delete a pool, Cisco UCS Manager does not reallocate any addresses from that pool that were assigned to vNICs or vHBAs. All assigned addresses from a deleted pool remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.
Step 2 In the **SAN** tab, expand **SAN** > **Pools** > *Organization_Name* .
Step 3 Expand the **WWNN Pools** node.
Step 4 Right-click the WWNN pool you want to delete and select **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Configuring WWPN Pools

Creating a WWPN Pool

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. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNs in the SAN fabric, Cisco recommends using the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:XX:XX:XX

A WWNN pool with the last four digits ending in 00:01 causes the vHBA to not initialize, no output from the lunlist command, and displays the Waiting for Flogi error. This error occurs if the WWPN is in the same block as the WWNN ending in 00:01. To ensure that the WWNN and WWPN addresses do not overlap, we recommend using a unique WWN address.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 In the **SAN** tab, expand **SAN** > **Pools**.

Step 3 Expand the node for the organization where you want to create the pool.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **WWPN Pools** and select **Create WWPN Pool**.

Step 5 In the **Define Name and Description** page of the **Create WWPN Pool** wizard, complete the following fields:

Name	Description
Name field	The name of the World Wide Port Name pool. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the pool. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Assignment Order field	This can be one of the following: • **Default**—Cisco UCS Manager selects a random identity from the pool. • **Sequential**—Cisco UCS Manager selects the lowest available identity from the pool.

Step 6 Click **Next**.

Step 7 In the **Add WWN Blocks** page of the **Create WWPN Pool** wizard, click **Add**.

Step 8 In the **Create WWN Block** dialog box, complete the following fields:

Name	Description
From field	The first WWN in the block.
Size field	The number of WWNs in the block. For WWxN pools, the pool size must be a multiple of *ports-per-node* + 1. For example, if there are 7 ports per node, the pool size must be a multiple of 8. If there are 63 ports per node, the pool size must be a multiple of 64.

Step 9 Click **OK**.

Step 10 Click **Finish**.

What to Do Next

Include the WWPN pool in a vHBA template.

Adding a WWN Block to a WWPN Pool

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. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNs in the SAN fabric, Cisco recommends using the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:XX:XX:XX

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 In the **SAN** tab, expand **SAN** > **Pools** > ***Organization_Name*** .

Step 3 Expand the **WWPN Pools** node.

Step 4 Right-click the WWPN pool to which you want to add a WWN block and select **Create WWN Block**.

Step 5 In the **Create WWN Block** dialog box, complete the following fields:

Name	Description
From field	The first WWN in the block.

Name	Description
Size field	The number of WWNs in the block. For WWxN pools, the pool size must be a multiple of *ports-per-node* + 1. For example, if there are 7 ports per node, the pool size must be a multiple of 8. If there are 63 ports per node, the pool size must be a multiple of 64.

Step 6 Click **OK**.

Deleting a WWN Block from a WWPN Pool

If you delete an address block from a pool, Cisco UCS Manager does not reallocate any addresses in that block that were assigned to vNICs or vHBAs. All assigned addresses from a deleted block remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 In the **SAN** tab, expand **SAN** > **Pools** > *Organization_Name* > **WWPN Pools** > *WWPN_Pool_Name* .

Step 3 Right-click the WWN block that you want to delete and select **Delete**.

Step 4 If a confirmation dialog box displays, click **Yes**.

Adding a WWPN Initiator to a WWPN Pool

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. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNs in the SAN fabric, Cisco recommends using the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:XX:XX:XX

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 In the **SAN** tab, expand **SAN** > **Pools** > *Organization_Name* .

Step 3 Expand the **WWPN Pools** node.

Step 4 Right-click the WWPN pool to which you want to add a WWPN initiator and select **Create WWPN Initiator**.

Step 5 In the **Create WWPN Initiator** dialog box, complete the following fields:

Name	Description
World Wide Name field	The WWN.
Name field	The name of the WWPN initiator. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A user-defined description of the WWPN initiator. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Step 6 If you want to add a SAN boot target, expand the **Boot Target** area and complete the following fields:

Name	Description
Boot Target WWPN field	The WWPN that corresponds to the location of the boot image.
Boot Target LUN field	The LUN that corresponds to the location of the boot image.

Step 7 Click **OK**.

Deleting a WWPN Initiator from a WWPN Pool

Procedure

Step 1 In the **Navigation** pane, click **SAN**.
Step 2 In the **SAN** tab, expand **SAN** > **Pools** > *Organization_Name* .
Step 3 Expand the **WWPN Pools** node.
Step 4 Choose the WWPN pool from which you want to delete a WWPN initiator.
Step 5 In the **Work** pane, click the **Initiators** tab.
Step 6 Right-click the initiator that you want to delete and choose **Delete**.
Step 7 If a confirmation dialog box displays, click **Yes**.

Deleting a WWPN Pool

If you delete a pool, Cisco UCS Manager does not reallocate any addresses from that pool that were assigned to vNICs or vHBAs. All assigned addresses from a deleted pool remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.
Step 2 In the **SAN** tab, expand **SAN** > **Pools** > *Organization_Name* .
Step 3 Expand the **WWPN Pools** node.
Step 4 Right-click the WWPN pool you want to delete and select **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Configuring WWxN Pools

Creating a WWxN Pool

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. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNs in the SAN fabric, Cisco recommends using the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:XX:XX:XX

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 In the **SAN** tab, expand **SAN** > **Pools**.

Step 3 Expand the node for the organization where you want to create the pool.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **WWxN Pools** and select **Create WWxN Pool**.

Step 5 In the **Define Name and Description** page of the **Create WWxN Pool** wizard, complete the following fields:

Name	Description
Name field	The name of the World Wide Port Name pool. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the pool. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Max Ports per Node field	The maximum number of ports that can be assigned to each node name in this pool. You cannot change this value once the object has been saved.
Assignment Order field	This can be one of the following: • **Default**—Cisco UCS Manager selects a random identity from the pool. • **Sequential**—Cisco UCS Manager selects the lowest available identity from the pool.

Step 6 Click **Next**.

Step 7 In the **Add WWN Blocks** page of the **Create WWxN Pool** wizard, click **Add**.

Step 8 In the **Create WWN Block** dialog box, complete the following fields:

Name	Description
From field	The first WWN in the block.
Size field	The number of WWNs in the block. For WWxN pools, the pool size must be a multiple of *ports-per-node* + 1. For example, if there are 7 ports per node, the pool size must be a multiple of 8. If there are 63 ports per node, the pool size must be a multiple of 64.

Step 9 Click **OK**.

Step 10 Click **Finish**.

What to Do Next

Include the WWxN pool in a service profile and template.

Adding a WWN Block to a WWxN Pool

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. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNs in the SAN fabric, Cisco recommends using the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:XX:XX:XX

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 In the **SAN** tab, expand **SAN** > **Pools** > ***Organization_Name*** .

Step 3 Expand the **WWxN Pools** node.

Step 4 Right-click the WWxN pool to which you want to add a WWN block and select **Create WWN Block**.

Step 5 In the **Create WWN Block** dialog box, complete the following fields:

Name	Description
From field	The first WWN in the block.

Name	Description
Size field	The number of WWNs in the block. For WWxN pools, the pool size must be a multiple of *ports-per-node* + 1. For example, if there are 7 ports per node, the pool size must be a multiple of 8. If there are 63 ports per node, the pool size must be a multiple of 64.

Step 6 Click **OK**.

Deleting a WWN Block from a WWxN Pool

If you delete an address block from a pool, Cisco UCS Manager does not reallocate any addresses in that block that were assigned to vNICs or vHBAs. All assigned addresses from a deleted block remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 In the **SAN** tab, expand **SAN** > **Pools** > *Organization_Name* > **WWxN Pools** > *WWxN_Pool_Name* .

Step 3 Right-click the WWN block that you want to delete and select **Delete**.

Step 4 If a confirmation dialog box displays, click **Yes**.

Deleting a WWxN Pool

If you delete a pool, Cisco UCS Manager does not reallocate any addresses from that pool that were assigned to vNICs or vHBAs. All assigned addresses from a deleted pool remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 In the **SAN** tab, expand **SAN** > **Pools** > ***Organization_Name*** .

Step 3 Expand the **WWxN Pools** node.

Step 4 Right-click the WWxN pool you want to delete and select **Delete**.

Step 5 If a confirmation dialog box displays, click **Yes**.

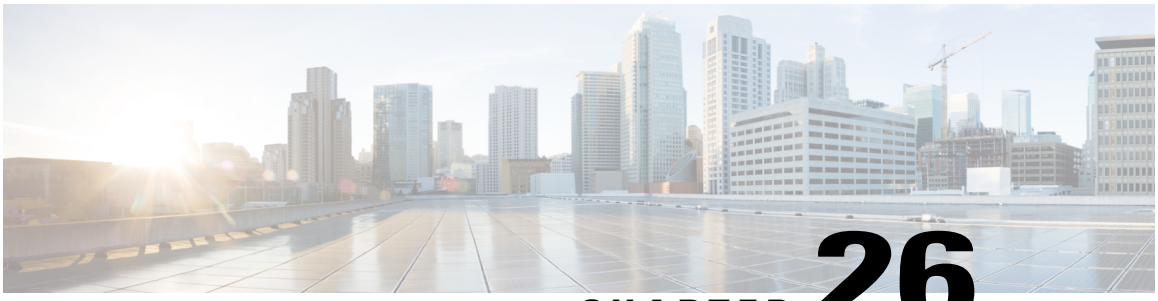

CHAPTER 26

Configuring Storage-Related Policies

This chapter includes the following sections:

Configuring vHBA Templates

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

Creating a vHBA Template

Before You Begin

This policy requires that one or more of the following resources already exist in the system:

- Named VSAN
- WWNN pool or WWPN pool
- SAN pin group
- Statistics threshold policy

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the **vHBA Templates** node and choose **Create vHBA Template**.

Step 5 In the **Create vHBA Template** dialog box, complete the following fields:

Name	Description
Name field	The name of the virtual HBA template. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A user-defined description of the template. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Fabric ID field	The name of the fabric interconnect that vHBAs created with this template are associated with.
Select VSAN drop-down list	The VSAN to associate with vHBAs created from this template.
Create VSAN link	Click this link if you want to create a VSAN.
Template Type field	This can be one of the following: • **Initial Template—**vHBAs created from this template are not updated if the template changes. • **Updating Template—**vHBAs created from this template are updated if the template changes.
Max Data Field Size field	The maximum size of the Fibre Channel frame payload bytes that the vHBA supports. Enter an integer between 256 and 2112. The default is 2048.
WWPN Pool drop-down list	The WWPN pool that a vHBA created from this template uses to derive its WWPN address.
QoS Policy drop-down list	The QoS policy that is associated with vHBAs created from this template.

Name	Description
Pin Group drop-down list	The SAN pin group that is associated with vHBAs created from this template.
Stats Threshold Policy drop-down list	The statistics collection policy that is associated with vHBAs created from this template.

Step 6 Click **OK**.

What to Do Next

Include the vHBA template in a service profile.

Creating vHBA Template Pairs

Procedure

Step 1 In the **Navigation** pane, click the **SAN** tab. On the **SAN** tab, expand **SAN** > **Policies**.

Step 2 Expand the node for the organization where you want to create the policy. If the system does not include multi-tenancy, expand the root node.

Step 3 Right-click the vHBA Templates node and choose **Create vHBA Template**. In the **Create vHBA Template** dialog box, assign a **Name**, **Description**, and select the **Fabric ID** for the template.

Step 4 Select the **Redundancy Type** as **Primary**, **Secondary** or **No Redundancy**. See the redundancy type descriptions below.

Step 5 Select the **Peer Redundancy Template—**to choose the name of the corresponding **Primary** or **Secondary** redundancy template to perform the template pairing using the **Primary** or **Secondary** redundancy template.

- **Primary—**Creates configurations that can be shared with the Secondary template. Any other shared changes on the Primary template are automatically synchronized to the Secondary template.

Note Following is a list of shared configurations:

- **VSANS**
- **Template Type**
- **Maximum Data Field Size**
- **QoS Policy**
- **Stats Threshold Policy**

Following is a list of non-shared configurations:

- **Fabric ID**

  Note The Fabric ID must be mutually exclusive. If you assign the Primary template to Fabric A, then Fabric B is automatically assigned to the Secondary template as part of the synchronization from the Primary template.

- **Description**
- **WWPN Pool**
- **Pin Group Policy**

- **Secondary—**

  All shared configurations are inherited from the Primary template.

- **No Redundancy—**

  Legacy vNIC template behavior.

Step 6 Click **OK**.

What to Do Next

After you create the vHBA redundancy template pair, you can use the redundancy template pair to create redundancy vHBA pairs for any service profile in the same organization or sub-organization.

Undo vHBA Template Pairs

You can undo the vHBA template pair by changing the Peer Redundancy Template so that there is no peer template for the Primary or the Secondary template. When you undo a vHBA template pair, the corresponding vHBA pairs also becomes undone.

Procedure

Select **not set** from the **Peer Redundancy Template** drop-down list to undo the paring between the peer Primary or Secondary redundancy template used to perform the template pairing. You can also select **None** as the **Redundancy Type** to undo the pairing.

Note If you delete one template in a pair, you are prompt to delete the other template in the pair. If you do not delete the other template in the pair, that template resets its peer reference and retains its redundancy type.

Binding a vHBA to a vHBA Template

You can bind a vHBA associated with a service profile to a vHBA template. When you bind the vHBA to a vHBA template, Cisco UCS Manager configures the vHBA with the values defined in the vHBA template. If the existing vHBA configuration does not match the vHBA template, Cisco UCS Manager reconfigures the vHBA. You can only change the configuration of a bound vHBA through the associated vHBA template. You cannot bind a vHBA to a vHBA template if the service profile that includes the vHBA is already bound to a service profile template.

Important If the vHBA is reconfigured when you bind it to a template, Cisco UCS Manager reboots the server associated with the service profile.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that includes the service profile with the vHBA you want to bind. If the system does not include multi-tenancy, expand the **root** node.

Step 4 Expand *Service_Profile_Name* > **vHBAs**.

Step 5 Click the vHBA you want to bind to a template.

Step 6 In the **Work** pane, click the **General** tab.

Step 7 In the **Actions** area, click **Bind to a Template**.

Step 8 In the **Bind to a vHBA Template** dialog box, do the following:

a) From the **vHBA Template** drop-down list, choose the template to which you want to bind the vHBA.

b) Click **OK**.

Step 9 In the warning dialog box, click **Yes** to acknowledge that Cisco UCS Manager may need to reboot the server if the binding causes the vHBA to be reconfigured.

Unbinding a vHBA from a vHBA Template

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that includes the service profile with the vHBA you want to unbind.

If the system does not include multi-tenancy, expand the **root** node.

Step 4 Expand ***Service_Profile_Name*** > **vHBAs**.

Step 5 Click the vHBA you want to unbind from a template.

Step 6 In the **Work** pane, click the **General** tab.

Step 7 In the **Actions** area, click **Unbind from a Template**.

Step 8 If a confirmation dialog box displays, click **Yes**.

Deleting a vHBA Template

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **Policies** > ***Organization_Name***.

Step 3 Expand the **vHBA Templates** node.

Step 4 Right-click the vHBA template that you want to delete and choose **Delete**.

Step 5 If a confirmation dialog box displays, click **Yes**.

Configuring Fibre Channel Adapter Policies

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 a cluster configuration with two fabric interconnects

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.
- LUN Queue Depth—The LUN queue depth setting is available for Windows system FC adapter policies. Queue depth is the number of commands that the HBA can send and receive in a single transmission per LUN. Windows Storport driver sets this to a default value of 20 for physical miniports and to 250 for virtual miniports. This setting adjusts the initial queue depth for all LUNs on the adapter. Valid range for this value is 1 to 254. The default LUN queue depth is 20. This feature only works with Cisco UCS Manager version 3.1(2) and higher.
- IO TimeOut Retry—When the target device is not responding to an IO request within the specified timeout, the FC adapter will abort the pending command then resend the same IO after the timer expires. The FC adapter valid range for this value is 1 to 59 seconds. The default IO retry timeout is 5 seconds. This feature only works with Cisco UCS Manager version 3.1(2) and higher.

Operating System Specific Adapter Policies

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.

Important We recommend that you use the values in these policies for the applicable operating system. Do not modify any of the values in the default policies unless directed to do so by Cisco Technical Support.

However, if you are creating an Ethernet adapter policy for a Windows OS (instead of using the default Windows adapter policy), you must use the following formulas to calculate values that work with Windows:

Completion Queues = Transmit Queues + Receive Queues
Interrupt Count = (Completion Queues + 2) rounded up to nearest power of 2

For example, if Transmit Queues = 1 and Receive Queues = 8 then:

Completion Queues = 1 + 8 = 9
Interrupt Count = (9 + 2) rounded up to the nearest power of 2 = 16

Creating a Fibre Channel Adapter Policy

Tip If the fields in an area do not display, click the **Expand** icon to the right of the heading.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **Fibre Channel Policies** and choose **Create Fibre Channel Adapter Policy**.

Step 5 Enter a name and description for the policy in the following fields:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Owner field	This can be one of the following: • **Local—**This policy is available only to service profiles and service profile templates in this Cisco UCS domain. • **Pending Global—**Control of this policy is being transferred to Cisco UCS Central. Once the transfer is complete, this policy will be available to all Cisco UCS domains registered with Cisco UCS Central. • **Global—**This policy is managed by Cisco UCS Central. Any changes to this policy must be made through Cisco UCS Central.

Step 6 (Optional) In the **Resources** area, adjust the following values:

Name	Description
Transmit Queues field	The number of transmit queue resources to allocate. This value cannot be changed.
Ring Size field	The number of descriptors in each transmit queue. This parameter applies to Extended Link Services (ELS) and Common Transport (CT) fibre channel frames for generic services. It does not affect adapter performance. Enter an integer between 64 and 128. The default is 64.
Receive Queues field	The number of receive queue resources to allocate. This value cannot be changed.
Ring Size field	The number of descriptors in each receive queue. This parameter applies to Extended Link Services (ELS) and Common Transport (CT) fibre channel frames for generic services. It does not affect adapter performance. Enter an integer between 64 and 128. The default is 64.
SCSI I/O Queues field	The number of SCSI IO queue resources the system should allocate. Enter an integer between 1 and 8. The default is 1. **Note** At this time, the Cisco UCS M81KR Virtual Interface Card adapter supports only one SCSI I/O queue.
Ring Size field	The number of descriptors in each SCSI I/O queue. Enter an integer between 64 and 512. The default is 512. **Note** The number of descriptors can affect the performance of the adapter, so we recommend that you do not change the default value.

Step 7 (Optional) In the **Options** area, adjust the following values:

Name	Description
FCP Error Recovery field	Whether the system uses FCP Sequence Level Error Recovery (FC-TAPE) protocol for sequence level error recovery with tape devices. This enables or disables the Read Exchange Concise (REC) and Sequence Retransmission Request (SRR) functions on the VIC firmware. This can be one of the following: • **Disabled—**This is the default. • **Enabled—**You should select this option if your system is connected to one or more tape drive libraries. **Note** This parameter only applies to a server with a Virtual Interface Card (VIC) adapter, such as the Cisco UCS M81KR Virtual Interface Card.
Flogi Retries field	The number of times that the system tries to log in to the fabric after the first failure. Enter any integer. To specify that the system continue to try indefinitely, enter infinite in this field. We recommend you consult your storage array documentation for the optimal value for this parameter. **Note** This parameter only applies to a server with a VIC adapter, or a converged network adapter such as the Cisco UCS M71KR-E Emulex Converged Network Adapter.
Flogi Timeout field	The number of milliseconds that the system waits before it tries to log in again. Enter an integer between 1000 and 255000. The default is 4,000. We recommend you consult your storage array documentation for the optimal value for this parameter. **Note** This parameter only applies to a server with a VIC adapter or a converged network adapter. When a Flogi timeout value of 20 seconds or more is configured for a boot vHBA, it could lead to a SAN boot failure if the adapter does not receive an accept to the initial Flogi. For a boot-enabled vHBA, the recommended timeout values is 5 seconds or less.
Plogi Retries field	The number of times that the system tries to log into a port after the first failure. Enter an integer between 0 and 255. The default is 8. We recommend you consult your storage array documentation for the optimal value for this parameter. **Note** This parameter only applies to a server with a VIC adapter.

Name	Description
Plogi Timeout field	The number of milliseconds that the system waits before it tries to log in again. Enter an integer between 1000 and 255000. The default is 20,000. We recommend you consult your storage array documentation for the optimal value for this parameter. For an HBA that is going to be used to boot a Windows OS from SAN, the recommended value for this field is 4,000 ms. **Note** This parameter only applies to a server with a VIC adapter. When a Plogi timeout value of 20 seconds or more is configured for a boot vHBA, it could lead to a SAN boot failure if the adapter does not receive an accept to the initial Plogi. For a boot-enabled vHBA, the recommended timeout values is 5 seconds or less.
Error Detect Timeout field	The number of milliseconds to wait before the system assumes that an error has occurred. This value cannot be changed.
Port Down Timeout field	The number of milliseconds a remote Fibre Channel port should be offline before informing the SCSI upper layer that the port is unavailable. This parameter is important for host multi-pathing drivers and it is one of the key indicators used for error processing. Enter an integer between 0 and 240000. The default is 30,000. For a server with a VIC adapter running ESX, the recommended value is 10,000. For a server with a port that is going to be used to boot a Windows OS from SAN, the recommended value for this field is 5000 milliseconds. We recommend you consult your storage array documentation for the optimal value for this parameter. **Note** This parameter only applies to a server with a VIC adapter.
IO Retry Timeout (seconds)	The number of seconds that the FC adapter waits before aborting the pending command and resending the same IO. This happens when the network device does not responding to an IO request within the specified time. Enter an integer between 0 and 59 seconds. The default IO retry timeout is 5 seconds.

Name	Description
Port Down IO Retry field	The number of times an IO request to a port is returned because the port is busy before the system decides the port is unavailable. Enter an integer between 0 and 255. The default is 8. We recommend you consult your storage array documentation for the optimal value for this parameter. **Note** This parameter only applies to a server with a VIC adapter running Windows.
Link Down Timeout field	The number of milliseconds the uplink port should be offline before it informs the system that the uplink port is down and fabric connectivity has been lost. Enter an integer between 0 and 240000. The default is 30,000. We recommend you consult your storage array documentation for the optimal value for this parameter. **Note** This parameter only applies to a server with a VIC adapter running Windows.
Resource Allocation Timeout field	The number of milliseconds to wait before the system assumes that a resource cannot be properly allocated. This value cannot be changed.
IO Throttle Count field	The maximum number of data or control I/O operations that can be pending in the vHBA at one time. If this value is exceeded, the additional I/O operations wait in the queue until the number of pending I/O operations decreases and the additional operations can be processed. **Note** This parameter is not the same as the LUN queue depth, which is controlled by Cisco UCS Manager based on the operating system installed on the server. Enter an integer between 256 and 1024. The default is 256. We recommend you consult your storage array documentation for the optimal value for this parameter.
Max LUNs Per Target field	The maximum number of LUNs that the Fibre Channel driver will export or show. The maximum number of LUNs is usually controlled by the operating system running on the server. Enter an integer between 1 and 1024. The default value is 256. For servers running ESX or Linux, the recommended value is 1024. We recommend you consult your operating system documentation for the optimal value for this parameter. **Note** This parameter only applies to a server with a VIC adapter or a network adapter.
LUN Queue Depth	The number of commands that the HBA can send and receive in a single transmission per LUN. Enter an integer between 1 and 254. The default LUN queue depth is 20.

Name	Description
Interrupt Mode field	The method used to send interrupts to the operating system from the driver. This can be one of the following: • **MSI-X—**Message Signaled Interrupts (MSI) with the optional extension. We recommend that you select this option if the operating system on the server supports it. • **MSI—**MSI only. • **INTx—**PCI INTx interrupts. **Note** This parameter only applies to a server with a VIC adapter or a network adapter running an operating system other than Windows. The Windows operating system ignores this parameter.

Step 8 Click **OK**.

Step 9 If a confirmation dialog box displays, click **Yes**.

Deleting a Fibre Channel Adapter Policy

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **Policies** > *Organization_Name*.

Step 3 Expand the **Fibre Channel Policies** node.

Step 4 Right-click the policy you want to delete and choose **Delete**.

Step 5 If a confirmation dialog box displays, click **Yes**.

Configuring the Default vHBA Behavior Policy

Default vHBA Behavior Policy

Default vHBA behavior policy allow you to configure how vHBAs are created for a service profile. You can choose to create vHBAs manually, or you can allow them to be created automatically.

You can configure the default vHBA behavior policy to define how vHBAs are created. This can be one of the following:

- **None—**Cisco UCS Manager does not create default vHBAs for a service profile. All vHBAs must be explicitly created.
- **HW Inherit—**If a service profile requires vHBAs and none have been explicitly defined, Cisco UCS Manager creates the required vHBAs based on the adapter installed in the server associated with the service profile.

Note If you do not specify a default behavior policy for vHBAs, **none** is used by default.

Configuring a Default vHBA Behavior Policy

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **Policies**.

Step 3 Expand the **root** node.
You can configure only the default vHBA behavior policy in the root organization. You cannot configure the default vHBA behavior policy in a sub-organization.

Step 4 Click **Default vHBA Behavior**.

Step 5 On the **General Tab**, in the **Properties** area, click one of the following radio buttons in the **Action** field:

- **None—**Cisco UCS Manager does not create default vHBAs for a service profile. All vHBAs must be explicitly created.
- **HW Inherit—**If a service profile requires vHBAs and none have been explicitly defined, Cisco UCS Manager creates the required vHBAs based on the adapter installed in the server associated with the service profile.

Step 6 Click **Save Changes**.

Configuring SAN Connectivity Policies

About the LAN and SAN Connectivity Policies

Connectivity policies determine the connections and the network communication resources between the server and the LAN or SAN on the network. These policies use pools to assign MAC addresses, WWNs, and WWPNs to servers and to identify the vNICs and vHBAs that the servers use to communicate with the network.

Note We do not recommend that you use static IDs in connectivity policies, because these policies are included in service profiles and service profile templates and can be used to configure multiple servers.

Privileges Required for LAN and SAN Connectivity Policies

Connectivity policies enable users without network or storage privileges to create and modify service profiles and service profile templates with network and storage connections. However, users must have the appropriate network and storage privileges to create connectivity policies.

Privileges Required to Create Connectivity Policies

Connectivity policies require the same privileges as other network and storage configurations. For example, you must have at least one of the following privileges to create connectivity policies:

- admin—Can create LAN and SAN connectivity policies
- ls-server—Can create LAN and SAN connectivity policies
- ls-network—Can create LAN connectivity policies
- ls-storage—Can create SAN connectivity policies

Privileges Required to Add Connectivity Policies to Service Profiles

After the connectivity policies have been created, a user with ls-compute privileges can include them in a service profile or service profile template. However, a user with only ls-compute privileges cannot create connectivity policies.

Interactions between Service Profiles and Connectivity Policies

You can configure the LAN and SAN connectivity for a service profile through either of the following methods:

- LAN and SAN connectivity policies that are referenced in the service profile
- Local vNICs and vHBAs that are created in the service profile
- Local vNICs and a SAN connectivity policy
- Local vHBAs and a LAN connectivity policy

Cisco UCS maintains mutual exclusivity between connectivity policies and local vNIC and vHBA configuration in the service profile. You cannot have a combination of connectivity policies and locally created vNICs or vHBAs. When you include a LAN connectivity policy in a service profile, all existing vNIC configuration is erased, and when you include a SAN connectivity policy, all existing vHBA configuration in that service profile is erased.

Creating a SAN Connectivity Policy

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **SAN Connectivity Policies** and choose **Create SAN Connectivity Policy**.

Step 5 In the **Create SAN Connectivity Policy** dialog box, enter a name and optional description.

Step 6 From the **WWNN Assignment** drop-down list in the **World Wide Node Name** area, choose one of the following:

- Choose **Select (pool default used by default)** to use the default WWN pool.
- Choose one of the options listed under **Manual Using OUI** and then enter the WWN in the **World Wide Node Name** field.

  You can specify a WWNN in the range 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. You can click the **here** link to verify that the WWNN you specified is available.
- Choose a WWN pool name from the list to have a WWN assigned from the specified pool. Each pool name is followed by two numbers in parentheses that show the number of WWNs still available in the pool and the total number of WWNs in the pool.

Step 7 In the **vHBAs** table, click **Add**.

Step 8 In the **Create vHBAs** dialog box, enter the name and optional description.

Step 9 Choose the **Fabric ID**, **Select VSAN**, **Pin Group**, **Persistent Binding**, and **Max Data Field Size**.
You can also create a VSAN or SAN pin group from this area.

Step 10 In the **Operational Parameters** area, choose the **Stats Threshold Policy**.

Step 11 In the **Adapter Performance Profile** area, choose the **Adapter Policy** and **QoS Policy**.
You can also create a fibre channel adapter policy or QoS policy from this area.

Step 12 After you have created all the vHBAs you need for the policy, click **OK**.

What to Do Next

Include the policy in a service profile or service profile template.

Creating a vHBA for a SAN Connectivity Policy

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 On the **SAN** tab, expand **SAN** > **Policies** > *Organization_Name* > **San Connectivity Policies**.

Step 3 Choose the policy for which you want to create a vHBA.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the table icon bar, click the + button.

Step 6 In the **Create vHBAs** dialog box, enter the name and optional description.

Step 7 Choose the **Fabric ID**, **Select VSAN**, **Pin Group**, **Persistent Binding**, and **Max Data Field Size**.
You can also create a VSAN or SAN pin group from this area.

Step 8 In the **Operational Parameters** area, choose the **Stats Threshold Policy**.

Step 9 In the **Adapter Performance Profile** area, choose the **Adapter Policy** and **QoS Policy**.
You can also create a fibre channel adapter policy or QoS policy from this area.

Step 10 Click **Save Changes**.

Deleting a vHBA from a SAN Connectivity Policy

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **Policies** > *Organization_Name*.

Step 3 Choose the policy from which you want to delete the vHBA.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **vHBAs** table, do the following:

a) Click the vHBA that you want to delete.

b) On the icon bar, click **Delete**.

Step 6 If a confirmation dialog box displays, click **Yes**.

Creating an Initiator Group for a SAN Connectivity Policy

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **Policies** > *Organization_Name*.

Step 3 Choose the policy for which you want to create an initiator group.

Step 4 In the **Work** pane, click the **vHBA Initiator Groups** tab.

Step 5 In the table icon bar, click the + button.

Step 6 In the **Create vHBA Initiator Group** dialog box, complete the following fields:

Name	Description
Name field	The name of the vHBA initiator group. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the group. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Select vHBA Initiators table	Check the check box in the **Select** column for each vHBA that you want to use.
Storage Connection Policy drop-down list	The storage connection policy associated with this vHBA initiator group. If you want to: • Use an existing storage connection policy, then choose that policy from the drop-down list. The Cisco UCS Manager GUI displays information about the policy and its FC target endpoints in the **Global Storage Connection Policy** area. Create a new storage connection policy that will be globally available, then click the **Create Storage Connection Policy** link. • Create a local storage connection policy that is available only to this vHBA initiator group, then choose the **Specific Storage Connection Policy** option. The Cisco UCS Manager GUI displays the **Specific Storage Connection Policy** area that allows you to configure the local storage connection policy.
Create Storage Connection Policy link	Click this link to create a new storage connection policy that will be available to all service profiles and service profile templates.

Step 7 Click **OK**.

Deleting an Initiator Group from a SAN Connectivity Policy

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **Policies** > ***Organization_Name***.

Step 3 Choose the policy from which you want to delete the initiator group

Step 4 In the **Work** pane, click the **vHBA Initiator Groups** tab.

Step 5 In the table, do the following:

a) Click the initiator group that you want to delete.

b) On the icon bar, click **Delete**.

Step 6 If a confirmation dialog box displays, click **Yes**.

Deleting a SAN Connectivity Policy

If you delete a SAN connectivity policy that is included in a service profile, it also deletes all vHBAs from that service profile and disrupts SAN data traffic for the server associated with the service profile.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **Policies** > ***Organization_Name***.

Step 3 Expand the **SAN Connectivity Policies** node.

Step 4 Right-click the policy that you want to delete and choose **Delete**.

Step 5 If a confirmation dialog box displays, click **Yes**.

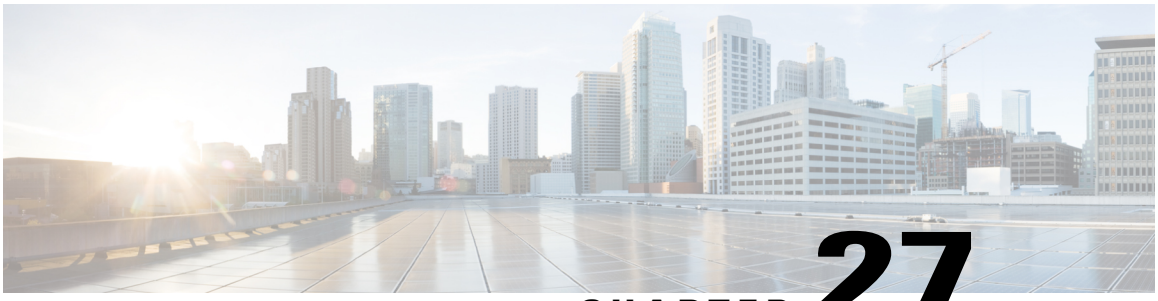

CHAPTER 27

Configuring Fibre Channel Zoning

This chapter includes the following sections:

Information About Fibre Channel Zoning

Fibre Channel zoning allows you to partition the Fibre Channel fabric into one or more zones. Each zone defines the set of Fibre Channel initiators and Fibre Channel targets that can communicate with each other in a VSAN. Zoning also enables you to set up access control between hosts and storage devices or user groups.

The access and data traffic control provided by zoning does the following:

- Enhances SAN network security
- Helps prevent data loss or corruption
- Reduces performance issues

Information About Zones

A zone consists of multiple zone members and has the following characteristics:

- Members in a zone can access each other; members in different zones cannot access each other.
- Zones can vary in size.
- Devices can belong to more than one zone.
- A physical fabric can have a maximum of 8,000 zones.

Information About Zone Sets

Each zone set consists of one or more zones. You can use zone sets to enforce access control within the Fibre Channel fabric. In addition, zone sets provide you with the following advantages:

- Only one zone set can be active at any time.
- All zones in a zone set can be activated or deactivated as a single entity across all switches in the fabric.
- A zone can be a member of more than one zone set.
- A switch in a zone can have a maximum of 500 zone sets.

Support for Fibre Channel Zoning in Cisco UCS Manager

Cisco UCS Manager supports switch-based Fibre Channel zoning and Cisco UCS Manager-based Fibre Channel zoning. You cannot configure a combination of zoning types in the same Cisco UCS domain. You can configure a Cisco UCS domain with one of the following types of zoning:

- Cisco UCS Manager-based Fibre Channel zoning—This configuration combines direct attach storage with local zoning. Fibre Channel or FCoE storage is directly connected to the fabric interconnects and zoning is performed in Cisco UCS Manager, using Cisco UCS local zoning. Any existing Fibre Channel or FCoE uplink connections need to be disabled. Cisco UCS does not currently support active Fibre Channel or FCoE uplink connections coexisting with the utilization of the UCS Local Zoning feature.
- Switch-based Fibre Channel zoning—This configuration combines direct attach storage with uplink zoning. The Fibre Channel or FCoE storage is directly connected to the fabric interconnects and zoning is performed externally to the Cisco UCS domain through an MDS or Nexus 5000 switch. This configuration does not support local zoning in the Cisco UCS domain.

Note Zoning is configured on a per-VSAN basis. You cannot enable zoning at the fabric level.

Cisco UCS Manager-Based Fibre Channel Zoning

With Cisco UCS Manager-based zoning, Cisco UCS Manager controls the Fibre Channel zoning configuration for the Cisco UCS domain, including creating and activating zones for all VSANs that you set up with this type of zoning. This type of zoning is also know as local zoning or direct attach storage with local zoning.

Note You cannot implement Cisco UCS Manager-based zoning if the VSAN is also configured to communicate with a VSAN on an upstream switch and includes Fibre Channel or FCoE uplink ports.

Supported Fibre Channel Zoning Modes

Cisco UCS Manager-based zoning supports the following types of zoning:

- Single initiator single target—Cisco UCS Manager automatically creates one zone for each vHBA and storage port pair. Each zone has two members. We recommend that you configure this type of zoning unless you expect the number of zones to exceed the maximum supported.
- Single initiator multiple targets—Cisco UCS Manager automatically creates one zone for each vHBA. We recommend that you configure this type of zoning if you expect the number of zones to reach or exceed the maximum supported.

vHBA Initiator Groups

vHBA initiator groups determine the Fibre Channel zoning configuration for all vHBAs in a service profile. Cisco UCS Manager does not include any default vHBA initiator groups. You must create vHBA initiator groups in any service profile that is to be assigned to servers included in a zone.

The configuration in a vHBA initiator group determines the following:

- The vHBAs included in the initiator group, which are sometimes referred to as vHBA initiators.
- A Fibre Channel storage connection policy, which includes the associated VSAN and the Fibre Channel target ports on the storage array.
- The type of Fibre Channel zoning to be configured for the vHBAs included in the group.

Fibre Channel Storage Connection Policy

The Fibre Channel storage connection policy contains a collection of target storage ports on storage arrays that you use to configure Cisco UCS Manager-based Fibre Channel zoning. You can create this policy under an organization or an initiator group.

The storage arrays in these zones must be directly connected to the fabric interconnects. The target storage ports on these arrays that you include in the Fibre Channel storage connection policy can be either Fibre Channel storage ports or FCoE storage ports. You use the WWN of a port to add it to the policy and to identify the port for the Fibre Channel zone.

Note Cisco UCS Manager does not create default Fibre Channel storage.

Fibre Channel Active Zone Set Configuration

In each VSAN that has been enabled for Fibre Channel zoning, Cisco UCS Manager automatically configures one zone set and multiple zones. The zone membership specifies the set of initiators and targets that are allowed to communicate with each other. Cisco UCS Manager automatically activates that zone set.

Cisco UCS Manager processes the user-configured vHBA initiator groups and their associated Fibre Channel storage connection policy to determine the desired connectivity between Fibre Channel initiators and targets. Cisco UCS Manager uses the following information to build pair-wise zone membership between initiators and targets:

- The port WWNs of the vHBA initiators derived from the vHBA initiator groups.
- The port WWNs of the storage array derived from the storage connection policy.

Switch-Based Fibre Channel Zoning

With switch-based zoning, a Cisco UCS domain inherits the zoning configuration from the upstream switch. You cannot configure or view information about your zoning configuration in Cisco UCS Manager. You have to disable zoning on a VSAN in Cisco UCS Manager to use switch-based zoning for that VSAN.

Guidelinesandrecommendationsfor Cisco UCS Manager-Based Fibre Channel Zoning

When you plan your configuration for Fibre Channel zoning, consider the following guidelines and recommendations:

Fibre Channel Switching Mode Must Be Switch Mode for Cisco UCS Manager Configurations

If you want Cisco UCS Manager to handle Fibre Channel zoning, the fabric interconnects must be in Fibre Channel Switch mode. You cannot configure Fibre Channel zoning in End-Host mode.

Symmetrical Configuration Is Recommended for High Availability

If a Cisco UCS domain is configured for high availability with two fabric interconnects, we recommend that both fabric interconnects are configured with the same set of VSANs.

Configuring Cisco UCS Manager Fibre Channel Zoning

Note This procedure provides a high level overview of the steps required to configure a Cisco UCS domain for Fibre Channel zoning that is controlled by Cisco UCS Manager. You must ensure that you complete all of the following steps.

Procedure

	Command or Action	Purpose
Step 1	If you have not already done so, disconnect the fabric interconnects in the Cisco UCS domain from any external Fibre Channel switches, such as an MDS.	
Step 2	If the Cisco UCS domain still includes zones that were managed by the external Fibre Channel switch, run the **clear-unmanaged-fc-zone-all** command on every affected VSAN to remove those zones.	This functionality is not currently available in the Cisco UCS Manager GUI. You must perform this step in the Cisco UCS Manager CLI.

	Command or Action	Purpose
Step 3	Configure the Fibre Channel switching mode for both fabric interconnects in Fibre Channel Switch mode.	You cannot configure Fibre Channel zoning in End-Host mode. See Configuring Fibre Channel Switching Mode, on page 76.
Step 4	Configure the Fibre Channel and FCoE storage ports that you require to carry traffic for the Fibre Channel zones.	See Configuring Ports and Port Channels, on page 79.
Step 5	Create one or more VSANs and enable Fibre Channel zoning on all VSANs that you require to carry traffic for the Fibre Channel zones.	For a cluster configuration, we recommend that you create the VSANs that you intend to include in a Fibre Channel zone in the SAN Uplinks Manager and use the common/global configuration to ensure they are accessible to both fabric interconnects. See Creating a VSAN for Fibre Channel Zoning, on page 385.
Step 6	Create one or more Fibre Channel storage connection policies.	You can perform this step when you configure Fibre Channel zoning in the service profiles, if you prefer. See Creating a Fibre Channel Storage Connection Policy, on page 388.
Step 7	Configure zoning in service profiles or service profile templates for servers that need to communicate through Fibre Channel zones.	Complete the following steps to complete this configuration: • Enable zoning in the VSAN or VSANs assigned to the VHBAs. • Configure one or more vHBA initiator groups.

Creating a VSAN for Fibre Channel Zoning

Note FCoE VLANs in the SAN cloud and VLANs in the LAN cloud must have different IDs. Using the same ID for an FCoE VLAN in a VSAN and a VLAN results in a critical fault and traffic disruption for all vNICs and uplink ports using that FCoE VLAN. Ethernet traffic is dropped on any VLAN with an ID that overlaps with an FCoE VLAN ID.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 On the **SAN** tab, click the **SAN** node.

Step 3 In the **Work** pane, click the **SAN Uplinks Manager** link on the **SAN Uplinks** tab.
The SAN Uplinks Manager opens in a separate window.

Step 4 In the SAN Uplinks Manager, click the **VSAN** tab.
You can create the VSAN on any of the subtabs. However, if you use the **All** subtab, you can view all of the configured VSANs in the table.

Step 5 On the icon bar to the right of the table, click +.
If the + icon is disabled, click an entry in the table to enable it.

Step 6 In the **Create VSAN** dialog box, complete the following fields:

Name	Description
Name field	The name assigned to the network. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
FC Zoning field	Click the radio button to determine whether Cisco UCS Manager configures Fibre Channel zoning for the Cisco UCS domain. This can be one of the following: • **Disabled—**The upstream switch handles Fibre Channel zoning, or Fibre Channel zoning is not implemented for the Cisco UCS domain. Cisco UCS Manager does not configure Fibre Channel zoning. • **Enabled—**Cisco UCS Manager configures and controls Fibre Channel zoning for the Cisco UCS domain. **Note** If you enable Fibre Channel zoning through Cisco UCS Manager, do not configure the upstream switch with any VSANs that are being used for Fibre Channel zoning.

Name	Description
Type radio button	Click the radio button to determine how the VSAN should be configured. This can be one of the following: • **Common/Global—**The VSAN maps to the same VSAN ID in all available fabrics. • **Fabric A—**The VSAN maps to the a VSAN ID that exists only in fabric A. • **Fabric B—**The VSAN maps to the a VSAN ID that exists only in fabric B. • **Both Fabrics Configured Differently—**The VSAN maps to a different VSAN ID in each available fabric. If you choose this option, Cisco UCS Manager GUI displays a **VSAN ID** field and a **FCoE VLAN** field for each fabric.
VSAN ID field	The unique identifier assigned to the network. The ID can be between 1 and 4078, or between 4080 and 4093. 4079 is a reserved VSAN ID. In addition, if you plan to use FC end-host mode, the range between 3840 to 4079 is also a reserved VSAN ID range.
FCoE VLAN field	The unique identifier assigned to the VLAN used for Fibre Channel connections. VLAN 4048 is user configurable. However, Cisco UCS Manager uses VLAN 4048 for the following default values. If you want to assign 4048 to a VLAN, you must reconfigure these values: • After an upgrade to Cisco UCS, Release 2.0—The FCoE storage port native VLAN uses VLAN 4048 by default. If the default FCoE VSAN was set to use VLAN 1 before the upgrade, you must change it to a VLAN ID that is not used or reserved. For example, consider changing the default to 4049 if that VLAN ID is not in use. • After a fresh install of Cisco UCS, Release 2.0—The FCoE VLAN for the default VSAN uses VLAN 4048 by default. The FCoE storage port native VLAN uses VLAN 4049. For FIP-capable, converged network adapters, such as the Cisco UCS CNA M72KR-Q and the Cisco UCS CNA M72KR-E, the named VSAN must be configured with a named VLAN that is not the native VLAN for the FCoE VLAN ID. This configuration ensures that FCoE traffic can pass through these adapters.

Step 7 Click **OK**.

Configuring Fibre Channel Storage Connection Policies

Creating a Fibre Channel Storage Connection Policy

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the **Storage Connection Policies** node and choose **Create Storage Connection Policy**.

Step 5 In the **Create Storage Connection Policy** dialog box, complete the following fields:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Step 6 In the **Zoning Type** field, click one of the following radio buttons:

- **None—**Cisco UCS Manager does not configure Fibre Channel zoning.
- **Single Initiator Single Target—**Cisco UCS Manager automatically creates one zone for each vHBA and storage port pair. Each zone has two members. We recommend that you configure this type of zoning unless you expect the number of zones to exceed the maximum supported.
- **Single Initiator Multiple Targets—**Cisco UCS Manager automatically creates one zone for each vHBA. We recommend that you configure this type of zoning if you expect the number of zones to reach or exceed the maximum supported.

Step 7 In the **FC Target Endpoints** table, click + on the icon bar to the right of the table.

If the + icon is disabled, click an entry in the table to enable it.

Step 8 In the **Create FC Target Endpoint** dialog box, complete the following fields and then click **OK**:

Name	Description
WWPN field	The WWPN (WWN) assigned to the physical target port on the Fibre Channel or FCoE storage array that the server uses to access the LUNs configured on the storage array.
Description field	A description of the target endpoint. We recommend that you include information about the port, LUNs, or storage array to which the target endpoint connects. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Path field	The fabric interconnect used for communications with the target endpoint.
Select VSAN drop-down list	The VSAN used for communications with the target endpoint.
Create VSAN link	Click this link if you want to create a VSAN.

Repeat this step until you have created all desired target endpoints for the policy.

Step 9 After you have created all desired target endpoints for the policy, click **OK**.

Deleting a Fibre Channel Storage Connection Policy

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **Policies** > *Organization_Name*.

Step 3 Expand the **Storage Connection Policies** node.

Step 4 Right-click the policy you want to delete and select **Delete**.

Step 5 If a confirmation dialog box displays, click **Yes**.

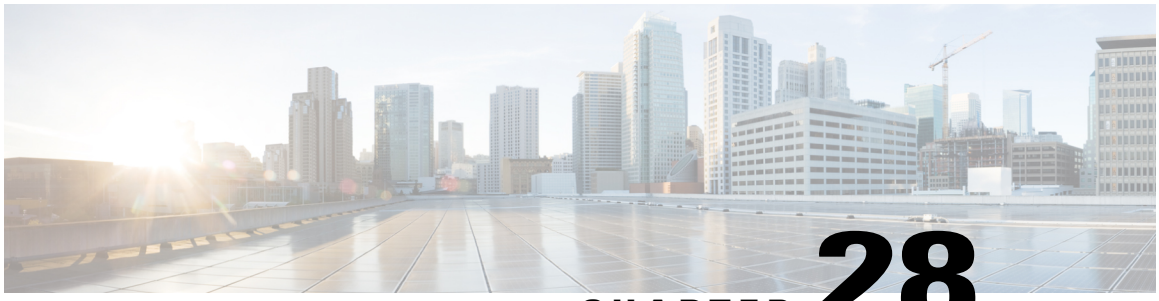

CHAPTER 28

Configuring Server-Related Pools

This chapter includes the following sections:

Configuring Server Pools

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 multitenancy 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.

Creating a Server Pool

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Pools**.

Step 3 Expand the node for the organization where you want to create the pool.

If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the **Server Pools** node and select **Create Server Pool**.

Step 5 On the **Set Name and Description** page of the **Create Server Pool** wizard, complete the following fields:

Name	Description
Name field	The name of the server pool. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A user-defined description of the server pool. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Step 6 Click **Next**.

Step 7 On the **Add Servers** page of the **Create Server Pool** wizard:

a) Select one or more servers from the **Available Servers** table.

b) Click the **>>** button to add the servers to the server pool.

c) When you have added all desired servers to the pool, click **Finish**.

Deleting a Server Pool

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Pools** > *Organization_Name*.

Step 3 Expand the **Server Pools** node.

Step 4 Right-click the pool you want to delete and select **Delete**.

Step 5 If a confirmation dialog box displays, click **Yes**.

Adding Servers to a Server Pool

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Pools** > *Organization_Name*.

Step 3 Right-click the pool to which you want to add one or more servers and select **Add Servers to Server Pool**.

Step 4 In the **Add Servers to Server Pool** dialog box, do the following:

a) In the **Servers** table, select the servers that you want to add to the server pool.
You can use the Shift key or Ctrl key to select multiple entries.

b) Click the **>>** button to move those servers to the **Pooled Servers** table and add them to the server pool.

c) Click **OK**.

Removing Servers from a Server Pool

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Pools** > *Organization_Name*.

Step 3 Right-click the pool from which you want to remove one or more servers and select **Add Servers to Server Pool**.

Step 4 In the **Add Servers to Server Pool** dialog box, do the following:

a) In the **Pooled Servers** table, select the servers that you want to remove from the server pool.
You can use the Shift key or Ctrl key to select multiple entries.

b) Click the **<<** button to move those servers to the **Servers** table and remove them from the server pool.

c) Click **OK**.

Configuring UUID Suffix Pools

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, are 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.

Creating a UUID Suffix Pool

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Pools**.

Step 3 Expand the node for the organization where you want to create the pool.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **UUID Suffix Pools** and select **Create UUID Suffix Pool**.

Step 5 In the **Define Name and Description** page of the **Create UUID Suffix Pool** wizard, complete the following fields:

Name	Description
Name field	The name of the UUID pool. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	The user-defined description of the pool. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Prefix field	This can be one of the following: • **Derived—**The system creates the suffix. • **other—**You specify the desired suffix. If you select this option, Cisco UCS Manager GUI displays a text field where you can enter the desired suffix, in the format *XXXXXXXX-XXXX-XXXX*.
Assignment Order field	This can be one of the following: • **Default—**Cisco UCS Manager selects a random identity from the pool. • **Sequential—**Cisco UCS Manager selects the lowest available identity from the pool.

Step 6 Click **Next**.

Step 7 In the **Add UUID Blocks** page of the **Create UUID Suffix Pool** wizard, click **Add**.

Step 8 In the **Create a Block of UUID Suffixes** dialog box, complete the following fields:

Name	Description
From field	The first UUID in the block.
Size field	The number of UUIDs in the block.

Step 9 Click **OK**.

Step 10 Click **Finish** to complete the wizard.

What to Do Next

Include the UUID suffix pool in a service profile and/or template.

Deleting a UUID Suffix Pool

If you delete a pool, Cisco UCS Manager does not reallocate any addresses from that pool that were assigned to vNICs or vHBAs. All assigned addresses from a deleted pool remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Pools** > *Organization_Name*.

Step 3 Expand the **UUID Suffix Pools** node.

Step 4 Right-click the pool you want to delete and select **Delete**.

Step 5 If a confirmation dialog box displays, click **Yes**.

Configuring IP Pools

IP Pools

IP pools are collections of IP addresses that do not have a default purpose. You can create IPv4 or IPv6 address pools in Cisco UCS Manager to do the following:

-
- Replace the default management IP pool **ext-mgmt** for servers that have an associated service profile. Cisco UCS Manager reserves each block of IP addresses in the IP pool for external access that terminates in the Cisco Integrated Management Controller (CIMC) on a server. If there is no associated service profile, you must use the **ext-mgmt** IP pool for the CIMC to get an IP address.
- Replace the management inband or out-of-band IP addresses for the CIMC.

Note You cannot create iSCSI boot IPv6 pools in Cisco UCS Manager.

You can create IPv4 address pools in Cisco UCS Manager to do the following:

- Replace the default iSCSI boot IP pool **iscsi-initiator-pool**. Cisco UCS Manager reserves each block of IP addresses in the IP pool that you specify.
- Replace both the management IP address and iSCSI boot IP addresses.

Note The IP pool must not contain any IP addresses that were assigned as static IP addresses for a server or service profile.

Creating an IP Pool

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Pools** > *Organization_Name* .

Step 3 Right-click **IP Pools** and select **Create IP Pool**.

Step 4 In the **Define Name and Description** page of the **Create IP Pool** wizard, complete the following fields:

Name	Description
Name field	The name of the IP address pool. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	The user-defined description of the IP address pool. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Assignment Order field	This can be one of the following: • **Default—**Cisco UCS Manager selects a random identity from the pool. • **Sequential—**Cisco UCS Manager selects the lowest available identity from the pool.

Step 5 Click **Next**.

Step 6 In the **Add IPv4 Blocks** page of the **Create IP Pool** wizard, click **Add**.

Step 7 In the **Create a Block of IPv4 Addresses** dialog box, complete the following fields:

Name	Description
From field	The first IPv4 address in the block.
Size field	The number of IP addresses in the pool.
Subnet Mask field	The subnet mask associated with the IPv4 addresses in the block.
Default Gateway field	The default gateway associated with the IPv4 addresses in the block.
Primary DNS field	The primary DNS server that this block of IPv4 addresses should access.
Secondary DNS field	The secondary DNS server that this block of IPv4 addresses should access.

Step 8

Step 9 Click **Next**.

Step 10 In the **Add IPv6 Blocks** page of the **Create IP Pool** wizard, click **Add**.

Step 11 In the **Create a Block of IPv6 Addresses** dialog box, complete the following fields:

Name	Description
From field	The first IPv6 address in the block.
Size field	The last IPv6 address in the block.
Default Gateway field	The default gateway associated with the IPv6 addresses in the block.
Prefix	The network address prefix associated with the IPv6 addresses in the block.
Primary DNS field	The primary DNS server that this block of IPv6 addresses should access.
Secondary DNS field	The secondary DNS server that this block of IPv6 addresses should access.

Step 12 Click **OK**.

Step 13 Click **Finish** to complete the wizard.

What to Do Next

Include the IP pool in a service profile and template.

Adding a Block to an IP Pool

You can add blocks of IPv4 or IPv6 addresses to IP pools.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Pools** > *Organization_Name* .

Step 3 Expand the **IP Pools** node.

Step 4 Right-click the desired IP pool and select one of:

- **Create Block of IPv4 Addresses**
- **Create Block of IPv6 Addresses**

Step 5 Complete the fields in the appropriate dialog box.

a) In the **Create a Block of IPv4 Addresses** dialog box, complete the following fields:

Name	Description
Name column	The range of IPv4 addresses assigned to the block.

Name	Description
From column	The first IPv4 address in the block.
To column	The last IPv4 address in the block.
Subnet column	The subnet mask associated with the IPv4 addresses in the block.
Default Gateway column	The default gateway associated with the IPv4 addresses in the block.
Primary DNS column	The primary DNS server that this block of IPv4 addresses should access.
Secondary DNS column	The secondary DNS server that this block of IPv4 addresses should access.

b) In the **Create a Block of IPv6 Addresses** dialog box, complete the following fields:

Name	Description
Name column	The range of IPv4 addresses assigned to the block.
From column	The first IPv4 address in the block.
To column	The last IPv4 address in the block.
Subnet column	The subnet mask associated with the IPv4 addresses in the block.
Default Gateway column	The default gateway associated with the IPv4 addresses in the block.
Primary DNS column	The primary DNS server that this block of IPv4 addresses should access.
Secondary DNS column	The secondary DNS server that this block of IPv4 addresses should access.

Step 6 Click **OK**.

Deleting a Block from an IP Pool

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 Expand **LAN** > **Pools** > **Root** .
Step 3 Expand the **IP Pools** node.
Step 4 Expand the pool for which you want to delete a block of IP addresses.
Step 5 Right-click the IP address block that you want to delete and select **Delete**.
Step 6 If a confirmation dialog box displays, click **Yes**.

Deleting an IP Pool

If you delete a pool, Cisco UCS Manager does not reallocate any addresses from that pool that were assigned to vNICs or vHBAs. All assigned addresses from a deleted pool remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 Expand **LAN** > **Pools** > ***Organization_Name*** .
Step 3 Expand the **IP Pools** node.
Step 4 Right-click the IP pool you want to delete and select **Delete**.
Note You cannot delete the default pools **ext-mgmt** and **iscsi-initiator-pool**.
Step 5 If a confirmation dialog box displays, click **Yes**.

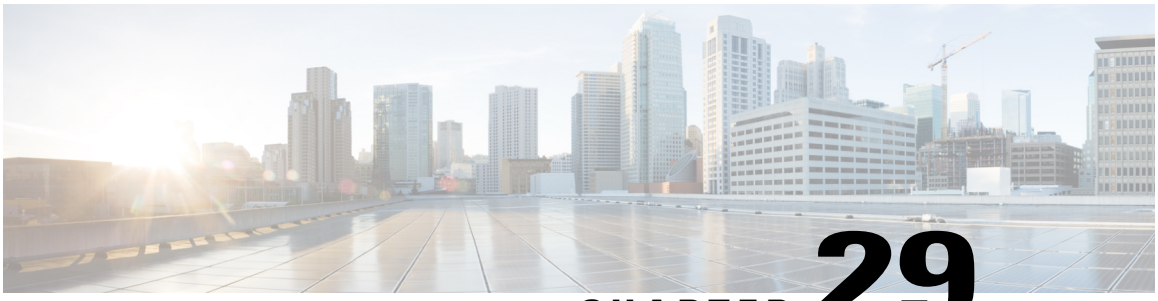

CHAPTER 29

Setting the Management IP Address

This chapter includes the following sections:

Management IP Address

Each server in a Cisco UCS domain must have a one or more management IP addresses assigned to its Cisco Integrated Management Controller (CIMC) or to the service profile associated with the server. Cisco UCS Manager uses these IP addresses for external access that terminates in the CIMC. This external access can be through one of the following services:

- KVM console
- Serial over LAN
- An IPMI tool

The management IP addresses used to access the CIMC on a server can be out-of-band (OOB) addresses, through which traffic traverses the fabric interconnect via the management port, or inband addresses, through which traffic traverses the fabric interconnect via the fabric uplink port. Up to six IP addresses can be configured to access the CIMC on a server, two out-of-band (OOB) and four inband.

You can configure the following management IP addresses:

- A static OOB IPv4 address assigned directly to the server
- An OOB IPv4 address assigned to the server from a global ext-mgmt pool
- An inband IPv4 address derived from a service profile associated with the server

- An inband IPv4 address drawn from a management IP pool and assigned to a service profile or service profile template
- An static inband IPv6 address assigned directly to the server
- An inband IPv6 address derived from a service profile associated with the server

You can assign multiple management IP addresses to each CIMC on the server and to the service profile associated with the server. If you do so, you must use different IP addresses for each of them.

A management IP address that is assigned to a service profile moves with that service profile. If KVM or SoL sessions are active when you migrate the service profile to another server, Cisco UCS Manager terminates the sessions and does not restart them after the migration is completed. You configure the IP address when you create or modify a service profile.

Note You cannot assign a static IP address to a server or service profile if that IP address has already been assigned to a server or service profile in the Cisco UCS domain. If you attempt to do so, Cisco UCS Manager warns you that the IP address is already in use and rejects the configuration.

An ARP request will be sent to the gateway IP address every second from each server that is configured with an Inband IP address. This is to check if connectivity for the Inband traffic through the current Fabric Interconnect is up, and to initiate a failover to the other Fabric Interconnect if it is down. The path selected for Inband and the failover operations are completely independent of the server data traffic.

Configuring the Management IP Address on a Blade Server

Configuring a Blade Server to Use a Static IP Address

If this action is greyed out, the server has already been assigned a static IP address.

You can configure a total of three static management addresses per server:

- Outband IPv4
- Inband IPv4
- Inband IPv6

Note You are not required to configure all three.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.

Step 3 Click the server for which you want to configure IP addresses.

Step 4 In the **Work** pane, click the **Inventory** tab.

Step 5 Click the **CIMC** subtab.
In the **Actions** area, two choices are available for management IP addresses:

- **Modify Outband Static Management IP**
- **Change Inband Management IP**

Step 6 To modify the outband static management IP address, in the **Actions** area, click **Modify Outband Static Management IP**:

Step 7 In the **Modify Outband Static Management IP** dialog box, complete the following fields:

Field	Description
IP Address	The static IPv4 address to be assigned to the server.
Subnet Mask	The subnet mask for the IP address.
Default Gateway	The default gateway that the IP address should use.

Step 8 Click **OK**.

Step 9 To modify the inband management IP address, click **Change Inband Management IP**.
In the **Change Management IP Address** dialog box, there are two tabs:

- **Inband IPv4**
- **Inband IPv6**

a) To change the static inband IPv4 management address, click the **Inband IPv4** subtab.
b) In the **Change Management IP Address** dialog box, complete the following fields:

Field	Description
Management IP Address Policy drop-down	Click **Static**.
IP Address	The static IPv4 address to be assigned to the server.
Subnet Mask	The subnet mask for the IP address.
Default Gateway	The default gateway that the IP address should use.

c) Click **OK**.
d) To change the static inband management IPv6 address, click the **Inband IPv6** subtab.

e) In the **Change Management IP Address** dialog box, complete the following fields:

Field	Description
Management IP Address Policy drop-down	Click **Static**.
IP Address	The static IPv6 address to be assigned to the server.
Prefix	The network prefix for the IP address.
Default Gateway	The default gateway that the IP address should use.

Step 10 Click **OK**.

Step 11 If a confirmation dialog box displays, click **Yes**.

Configuring a Blade Server to Use a Management IP Pool

If any action is specified in this procedure is greyed out, it means that the configuration has already been completed.

You can configure a total of three management IP pools per server:

- Outband IPv4
- Inband IPv4
- Inband IPv6

Note You are not required to configure all three.

Before You Begin

Before configuring servers to use management IP pools, configure management IP pools. For additional information on inband management and vLAN groups, see Inband Management Support, on page 27.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > ***Chassis Number*** > **Servers**.

Step 3 Click the server that you want to configure to use the management IP pool.

Step 4 In the **Work** pane, click the **Inventory** tab.

Step 5 Click the **CIMC** subtab.

- To configure an outband IP pooled management IP address policy, proceed with Step 6.

- To configure inband IPv4 and/or IPv6 management IP address policies, proceed to Step 8.

Step 6 In the **Actions** area, click **Use Outband Pooled Management IP**.

Step 7 Click **Yes** in the **Use Outband Pooled Management IP** confirmation dialog box, then click **OK**.
The management IP address policy is now switched to using an OOB IP address from the outband management IP pool.

Step 8 In the **Actions** area, click **Change Inband Management IP**.

Step 9 In the **Change Management IP Dialog** box, there are two tabs:

- **Inband IPv4**
- **Inband IPv6**

a) To change the inband IPv4 management IP pool, click the **Inband IPv4** tab, and complete the following fields:

Field	Description
Network drop-down list	A VLAN selected from the associated VLAN group. For more information on how to configure the VLAN group, see Configure a VLAN and VLAN Group.
Management IP Address Policy drop-down list	The management IP pool you want to assign to the server. There are two types of pools available: • **Domain Pools** • **Global Pools** Select one of the pools available from either the **Domain Pools** entries or the **Global Pools** entries.

b) To change the inband IPv6 management IP pool, click the **Inband IPv6** tab, and complete the following fields:

Field	Description
Network drop-down list	A VLAN selected from the associated VLAN group. For more information on how to configure the VLAN group, see Configure a VLAN and VLAN Groupt
Management IP Address Policy drop-down list	The management IP pool you want to assign to the server. There are two types of pools available: • **Domain Pools** • **Global Pools** Select one of the pools available from either the **Domain Pools** entries or the **Global Pools** entries.

Step 10 Click **OK**.

Step 11 If a confirmation dialog box displays, click **Yes**.

Deleting the Inband Configuration from a Blade Server

This procedure removes the inband management IP address configuration from a blade server. If this action is greyed out, no inband configuration was completed.

Procedure

Step 1 In the **Navigation** pane, click the **Servers** tab.

Step 2 On the **Equipment** tab, expand **Equipment** > **Chassis** > *Chassis Number* > **Servers** > *Server Name*.

Step 3 In the **Work** area, click the **Inventory** tab.

Step 4 Click the **CIMC** subtab.

Step 5 In the **Actions** area, click **Delete Inband Configuration**.

Step 6 Click **Yes** in the **Delete** confirmation dialog box.
The inband configuration for the server is deleted.

Note If an inband service profile is configured in Cisco UCS Manager with a default VLAN and pool name, the server CIMC will automatically get an inband configuration from the inband profile approximate one minute after deleting the inband configuration here.

Configuring the Management IP Address on a Rack Server

Configuring a Rack Server to Use a Static IP Address

If this action is greyed out, the server has already been assigned a static IP address.

You can configure a total of three static management addresses per server:

- Outband IPv4
- Inband IPv4
- Inband IPv6

Note You are not required to configure all three.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Rack Mounts** > **Servers**.

Step 3 Click the server for which you want to configure IP addresses.

Step 4 In the **Work** pane, click the **Inventory** tab.

Step 5 Click the **CIMC** subtab.
In the **Actions** area, two choices are available for management IP addresses:

- **Modify Outband Static Management IP**
- **Change Inband Management IP**

Step 6 To modify the outband static management IP address, in the **Actions** area, click **Modify Outband Static Management IP**.

Step 7 In the **Modify Outband Static Management IP** dialog box, complete the following fields:

Field	Description
IP Address	The static IPv4 address to be assigned to the server.
Subnet Mask	The subnet mask for the IP address.
Default Gateway	The default gateway that the IP address should use.

Step 8 Click **OK**.

Step 9 To modify the inband management IP address, click **Change Inband Management IP**.
In the **Change Management IP Address** dialog box, there are two tabs:

- **Inband IPv4**
- **Inband IPv6**

Step 10 To change the static inband IPv4 management address, click the **Inband IPv4** tab.

Step 11 In the **Change Management IP Address** dialog box, complete the following fields:

Field	Description
Management IP Address Policy drop-down	Click **Static**.
IP Address	The static IPv4 address to be assigned to the server.
Subnet Mask	The subnet mask for the IP address.
Default Gateway	The default gateway that the IP address should use.

Step 12 Click **OK**.

Step 13 To change the static inband IPv6 management address, click the **Inband IPv6** tab.

Step 14 In the **Change Management IP Address** dialog box, complete the following fields:

Field	Description
Management IP Address Policy drop-down	Click **Static**.
IP Address	The static IPv6 address to be assigned to the server.
Prefix	The prefix for the IP address.
Default Gateway	The default gateway that the IP address should use.

Step 15 Click **OK**.

Step 16 If a confirmation dialog box displays, click **Yes**.

Configuring a Rack Server to Use a Management IP Pool

If any action is specified in this procedure is greyed out, it means that the configuration has already been completed.

You can configure a total of three management IP pools per server:

- Outband IPv4
- Inband IPv4
- Inband IPv6

Note You are not required to configure all three.

Before You Begin

Configure management IP pools before configuring servers to use management IP pools.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Rack Mounts** > **Servers**.

Step 3 Click the server that you want to configure to use the management IP pool.

Step 4 In the **Work** pane, click the **Inventory** tab.

Step 5 Click the **CIMC** subtab.

- To configure an outband IP pooled management IP address policy, proceed with Step 6.

- To configure inband IPv4 and/or IPv6 management IP address policies, proceed to Step 8.

Step 6 In the **Actions** area, click **Use Outband Pooled Management IP**.

Step 7 Click **Yes** in the **Use Outband Pooled Management IP** confirmation dialog box, then click **OK**.
The management IP address policy is now switched to using an OOB IP address from the outband management IP pool.

Step 8 In the **Actions** area, click **Change Inband Management IP**.

Step 9 In the **Change Management IP Dialog** box, there are two tabs:

- **Inband IPv4**
- **Inband IPv6**

a) To change the inband IPv4 management IP pool, click the **Inband IPv4** tab, and complete the following fields:

Field	Description
Network drop-down list	The network (VLAN) that you want the server to use. For more information on how to configure the VLAN group, see Configure a VLAN and VLAN Group.
Management IP Address Policy drop-down list	The management IP pool you want to assign to the server. There are two types of pools available: • **Domain Pools** • **Global Pools** Select one of the pools available from either the **Domain Pools** entries or the **Global Pools** entries.

b) To change the inband IPv6 management IP pool, click the **Inband IPv6** tab, and complete the following fields:

Field	Description
Network drop-down list	The network (VLAN) that you want the server to use. For more information on how to configure the VLAN group, see Configure a VLAN and VLAN Group.
Management IP Address Policy drop-down list	The management IP pool you want to assign to the server. There are two types of pools available: • **Domain Pools** • **Global Pools** Select one of the pools available from either the **Domain Pools** entries or the **Global Pools** entries.

Step 10 Click **OK**.

Step 11 If a confirmation dialog box displays, click **Yes**.

Deleting the Inband Configuration from a Rack Server

This procedure removes the inband management IP address configuration from a rack server. If this action is greyed out, no inband configuration was configured.

Procedure

Step 1 In the **Navigation** pane, click the **Servers** tab.

Step 2 On the **Equipment** tab, expand **Equipment** > **Rack-Mounts** > **Servers** > ***Server Name***.

Step 3 In the **Work** area, click the **Inventory** tab.

Step 4 Click the **CIMC** subtab.

Step 5 In the **Actions** area, click **Delete Inband Configuration**.

Step 6 Click **Yes** in the **Delete** confirmation dialog box.
The inband configuration for the server is deleted.

Note If an inband service profile is configured in Cisco UCS Manager with a default VLAN and pool name, the server CIMC automatically gets an inband configuration from the inband profile approximate one minute after deleting the inband configuration here.

Setting the Management IP Addresses on a Service Profile

You can set the following management IP addresses on a service profile:

- Out-of-band (OOB) IPv4 address
- Inband IPv4 address
- Inband IPv6 address

This procedure guides you through the process of setting the management IP addresses on a service profile.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that contains the service profile for which you want to set the management IP addresses.

If the system does not include multitenancy, expand the **root** node.

Step 4 Click the service profile for which you want to set the management IP addresses.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 In the **Actions** area, click **Change Management IP Address**.

Step 7 To change the out-of-band IPv4 address, click the **Outband IPv4** tab, and complete the following fields:

Name	Description
Management IP Address Policy drop-down	How the OOB management IPv4 address is derived for the server associated with this service profile. This can be one of the following: • **None—**No management IP address is assigned to the server through the service profile. The management IP address is based on the CIMC management IP address defined on the server. • **Static—**The service profile assigns a static management IP address to the associated server. If the service profile is migrated to a new server, the management IP address moves with the service profile. • **Domain Pools—**The service profile assigns an IP management address from the pool selected from the list of pools to the associated server. If the service profile is migrated to a new server, the management IP address moves with the service profile.
IP Address field	The OOB management IPv4 address assigned to the server through the service profile. **Note** The IP address fields are only displayed if **Management IP Address Policy** is set to **Static** or one of the **Domain Pools** IP address pools is selected from the **Management IP Address Policy** field drop-down list.
Subnet Mask field	The subnet mask for the OOB management IPv4 address.
Default Gateway field	The default gateway for the OOB management IPv4 address.

a) If you selected **Static** from the **Management IP Address Policy** drop-down list, complete the following fields:

Name	Description
IP Address field	The management IP address assigned to the server through the service profile.
Subnet Mask field	The subnet mask for the management IP address.
Default Gateway field	The default gateway for the management IP address.

Step 8 To change the inband IPv4 address, click the **Inband** tab, then click the **Inband IPv4** subtab, and complete the following fields:

Name	Description
Network drop-down	A VLAN selected from the associated VLAN group. **Note** The network drop-down does not appear in the service profile unless it is first configured under the LAN tab. For more information on how to configure the VLAN group, see Configure a VLAN and VLAN Group .
Management IP Address Policy drop-down	How the inband management IPv4 address is derived for the server associated with this service profile. This can be one of the following: • **None—**No management IP address is assigned to the server through the service profile. The management IP address is based on the CIMC management IP address defined on the server. • **Static—**The service profile assigns a static management IP address to the associated server. If the service profile is migrated to a new server, the management IP address moves with the service profile. • **Domain Pools—**The service profile assigns an IP management address from the pool selected from the list of pools to the associated server. If the service profile is migrated to a new server, the management IP address moves with the service profile.
IP Address field	The inband management IPv4 address assigned to the server through the service profile. **Note** The IP address fields are only displayed if **Management IP Address Policy** is set to **Static** or one of the **Domain Pools** IP address pools is selected from the **Management IP Address Policy** field drop-down list.
Subnet Mask field	The subnet mask for the inband management IPv4 address.
Default Gateway field	The default gateway for the inband management IPv4 address.

a) If you selected **Static** from the **Management IP Address Policy** drop-down list, complete the following fields:

Name	Description
IP Address field	The inband management IPv4 address for the server associated with this service profile.
Subnet Mask field	The subnet mask for the OOB management IPv4 address.
Default Gateway field	The default gateway for the OOB management IPv4 address.

Step 9 To change the inband IPv6 address, click the **Inband** tab, then click the **Inband IPv6** subtab, and complete the following fields:

Name	Description
Network drop-down	A VLAN selected from the associated VLAN group. **Note** The network drop-down does not appear in the service profile unless it is first configured under the LAN tab. For more information on how to configure the VLAN group, see Configure a VLAN and VLAN Group .
Management IP Address Policy drop-down	How the inband management IPv6 address is derived for the server associated with this service profile. This can be one of the following: • **None—**No management IP address is assigned to the server through the service profile. The management IP address is based on the CIMC management IP address defined on the server. • **Static—**The service profile assigns a static management IP address to the associated server. If the service profile is migrated to a new server, the management IP address moves with the service profile. • **Domain Pools—**The service profile assigns an IP management address from the pool selected from the list of pools to the associated server. If the service profile is migrated to a new server, the management IP address moves with the service profile.
IP Address field	The inband management IPv6 address assigned to the server through the service profile. **Note** The IP address fields are only displayed if **Management IP Address Policy** is set to **Static** or one of the **Domain Pools** IP address pools is selected from the **Management IP Address Policy** field drop-down list.
Prefix field	The prefix for the inband management IPv6 address.
Default Gateway field	The default gateway for the inband management IPv6 address.

a) If you selected **Static** from the **Management IP Address Policy** drop-down list, complete the following fields:

Name	Description
IP Address field	The inband management IPv6 address assigned to the server through the service profile.
Prefix field	The prefix for the inband management IPv6 address.
Default Gateway field	The default gateway for the inband management IPv6 address.

Step 10 Click **Save Changes**.

Setting the Management IP Address on a Service Profile Template

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profile Templates**.

Step 3 Expand the node for the organization that contains the service profile template for which you want to set the management IP address.
If the system does not include multitenancy, expand the **root** node.

Step 4 Click the service profile template for which you want to set the management IP address.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 Expand the **Management IP Address** area.

Step 7 In the **Actions** area, click **Change Management IP Address**.

Step 8 Complete the fields in the **Change Management IP Address** dialog box.

Step 9 Click **Save Changes**.

Configuring the Management IP Pool

Management IP Pools

The default management IP pool, **IP Pool ext-mgmt** is a collection of external IPv4 and IPv6 addresses. Cisco UCS Manager reserves each block of IP addresses in the management IP pool for external access that terminates in the CIMC on a server.

By default, the **IP Pool ext-mgmt** is used to configure the CIMC outbound management IP address. You cannot change this IP pool if already a static IP address is assigned to the server from this pool. If you want to configure the outbound management IP address for CIMC from a static IP address, then you can delete the IP addresses from the default management IP pool.

You can configure separate out-of-band IPv4 address pools, and in-band IPv4 or IPv6 address pools. You can configure in-band pools that contain both IPv4 and IPv6 address blocks.

Tip To avoid assigning an IP pool that contains only IPv4 addresses as the in-band IPv6 policy, or assigning an IP pool that contains only IPv6 addresses as the in-band IPv4 policy to a server CIMC, it is suggested that you configure separate in-band address pools, each with only IPv4 or IPv6 addresses.

You can configure service profiles and service profile templates to use IP addresses from the management IP pools. You cannot configure servers to use the management IP pool.

All IP addresses in the management IP pool must be in the same IPv4 subnet, or have the same IPv6 network prefix as the IP address of the fabric interconnect.

Note The management IP pool must not contain any IP addresses that were assigned as static IP addresses for a server or service profile.

Creating an IPv4 Address Block in the Management IP Pool

The management IP pool must not contain any IP addresses that were assigned as static IP addresses for a server or service profile.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 In the **LAN** tab, expand **LAN** > **Pools** > *Organization_Name* .

Step 3 Expand the **IP Pools** node.

Step 4 Right-click **IP Pool ext-mgmt** and select **Create Block of IP Addresses**.

Step 5 In the **Create a Block of IPv4 Addresses** dialog box, complete the following fields:

Name	Description
Name column	The range of IPv4 addresses assigned to the block.
From column	The first IPv4 address in the block.
To column	The last IPv4 address in the block.
Subnet column	The subnet mask associated with the IPv4 addresses in the block.
Default Gateway column	The default gateway associated with the IPv4 addresses in the block.
Primary DNS column	The primary DNS server that this block of IPv4 addresses should access.
Secondary DNS column	The secondary DNS server that this block of IPv4 addresses should access.

Step 6 Click **OK**.

What to Do Next

Configure one or more service profiles or service profile templates to obtain the CIMC IP address from the management IP pool.

Creating an IPv6 Address Block in the Management IP Pool

The management IP pool must not contain any IP addresses that were assigned as static IP addresses for a server or service profile.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 In the **LAN** tab, expand **LAN** > **Pools** > *Organization_Name* .
Step 3 Expand the **IP Pools** node.
Step 4 Right-click **IP Pool ext-mgmt** and select **Create Block of IP Addresses**.
Step 5 In the **Create a Block of IPv6 Addresses** dialog box, specify the required information.
Step 6 Click **OK**.

What to Do Next

Configure one or more service profiles or service profile templates to obtain the CIMC IP address from the management IP pool.

Deleting an IP Address Block from the Management IP Pool

Procedure

Step 1 In the **Navigation** pane, click **LAN**.
Step 2 In the **LAN** tab, expand **LAN** > **Pools** > *Organization_Name* .
Step 3 Expand the **IP Pools** node.
Step 4 Select **IP Pool ext-mgmt**.
Step 5 Right-click the IP address block that you want to delete and select **Delete**.
Step 6 If a confirmation dialog box displays, click **Yes**.

CHAPTER 30

Configuring Server-Related Policies

This chapter includes the following sections:

Configuring BIOS Settings

Server BIOS Settings

Cisco UCS provides two methods for making global modifications to the BIOS settings on servers in an Cisco UCS domain. You can create one or more BIOS policies that include a specific grouping of BIOS settings that match the needs of a server or set of servers, or you can use the default BIOS settings for a specific server platform.

Both the BIOS policy and the default BIOS settings for a server platform enable you to fine tune the BIOS settings for a server managed by Cisco UCS Manager.

Depending upon the needs of the data center, you can configure BIOS policies for some service profiles and use the BIOS defaults in other service profiles in the same Cisco UCS domain, or you can use only one of them. You can also use Cisco UCS Manager to view the actual BIOS settings on a server and determine whether they are meeting current needs.

Note Cisco UCS Manager pushes BIOS configuration changes through a BIOS policy or default BIOS settings to the Cisco Integrated Management Controller (CIMC) buffer. These changes remain in the buffer and do not take effect until the server is rebooted.

We recommend that you verify the support for BIOS settings in the server that you want to configure. Some settings, such as Mirroring Mode for RAS Memory, are not supported by all Cisco UCS servers.

Main BIOS Settings

The following table lists the main server BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Reboot on BIOS Settings Change	When the server is rebooted after you change one or more BIOS settings. If you enable this setting, the server is rebooted according to the maintenance policy in the server's service profile. For example, if the maintenance policy requires user acknowledgment, the server is not rebooted and the BIOS changes are not applied until a user acknowledges the pending activity. If you do not enable this setting, the BIOS changes are not applied until the next time the server is rebooted, whether as a result of another server configuration change or a manual reboot.
Quiet Boot	What the BIOS displays during Power On Self-Test (POST). This can be one of the following: • **disabled—**The BIOS displays all messages and Option ROM information during boot. • **enabled—**The BIOS displays the logo screen, but does not display any messages or Option ROM information during boot. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Post Error Pause	What happens when the server encounters a critical error during POST. This can be one of the following: • **disabled—**The BIOS continues to attempt to boot the server. • **enabled—**The BIOS pauses the attempt to boot the server and opens the Error Manager when a critical error occurs during POST. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Resume Ac On Power Loss	How the server behaves when power is restored after an unexpected power loss. This can be one of the following: • **stay-off—**The server remains off until manually powered on. • **last-state—**The server is powered on and the system attempts to restore its last state. • **reset—**The server is powered on and automatically reset. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Front Panel Lockout	Whether the power and reset buttons on the front panel are ignored by the server. This can be one of the following: • **disabled—**The power and reset buttons on the front panel are active and can be used to affect the server. • **enabled—**The power and reset buttons are locked out. The server can only be reset or powered on or off from the CIMC GUI. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Consistent Device Naming	Consistent Device Naming allows Ethernet interfaces to be named in a consistent manner. This makes Ethernet interface names more uniform, easy to identify, and persistent when adapter or other configuration changes are made. Whether consistent device naming is enabled or not. This can be one of the following: • **disabled—**Consistent device naming is disabled for the BIOS policy. • **enabled—**Consistent device naming is enabled for the BIOS policy. This enables Ethernet interfaces to be named consistently. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Processor BIOS Settings

The following table lists the processor BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Turbo Boost	Whether the processor uses Intel Turbo Boost Technology, which allows the processor to automatically increase its frequency if it is running below power, temperature, or voltage specifications. This can be one of the following: • **disabled—**The processor does not increase its frequency automatically. • **enabled—**The processor uses Turbo Boost Technology if required. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Enhanced Intel Speedstep	Whether the processor uses Enhanced Intel SpeedStep Technology, which allows the system to dynamically adjust processor voltage and core frequency. This technology can result in decreased average power consumption and decreased average heat production. This can be one of the following: • **disabled—**The processor never dynamically adjusts its voltage or frequency. • **enabled—**The processor utilizes Enhanced Intel SpeedStep Technology and enables all supported processor sleep states to further conserve power. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. We recommend that you contact your operating system vendor to make sure your operating system supports this feature.
Hyper Threading	Whether the processor uses Intel Hyper-Threading Technology, which allows multithreaded software applications to execute threads in parallel within each processor. This can be one of the following: • **disabled—**The processor does not permit hyperthreading. • **enabled—**The processor allows for the parallel execution of multiple threads. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. We recommend that you contact your operating system vendor to make sure the operating system supports this feature.

Name	Description
Core Multi Processing	Sets the state of logical processor cores per CPU in a package. If you disable this setting, Intel Hyper Threading technology is also disabled. This can be one of the following: • **all—**Enables multiprocessing on all logical processor cores. • **1** through ***n*—**Specifies the number of logical processor cores per CPU that can run on the server. To disable multiprocessing and have only one logical processor core per CPU running on the server, choose 1. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. We recommend that you contact your operating system vendor to make sure your operating system supports this feature.
Execute Disabled Bit	Classifies memory areas on the server to specify where the application code can execute. As a result of this classification, the processor disables code execution if a malicious worm attempts to insert code in the buffer. This setting helps to prevent damage, worm propagation, and certain classes of malicious buffer overflow attacks. This can be one of the following: • **disabled—**The processor does not classify memory areas. • **enabled—**The processor classifies memory areas. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. We recommend that you contact your operating system vendor to make sure your operating system supports this feature.
Virtualization Technology (VT)	Whether the processor uses Intel Virtualization Technology, which allows a platform to run multiple operating systems and applications in independent partitions. This can be one of the following: • **disabled—**The processor does not permit virtualization. • **enabled—**The processor allows multiple operating systems in independent partitions. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** If you change this option, you must power cycle the server before the setting takes effect.

Name	Description
Hardware Pre-fetcher	Whether the processor allows the Intel hardware prefetcher to fetch streams of data and instruction from memory into the unified second-level cache when necessary. This can be one of the following: • **disabled—**The hardware prefetcher is not used. • **enabled—**The processor uses the hardware prefetcher when cache issues are detected. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** **CPU Performance** must be set to **Custom** in order to specify this value. For any value other than **Custom**, this option is overridden by the setting in the selected CPU performance profile.
Adjacent Cache Line Pre-fetcher	Whether the processor fetches cache lines in even/odd pairs instead of fetching just the required line. This can be one of the following: • **disabled—**The processor only fetches the required line. • **enabled—**The processor fetches both the required line and its paired line. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** **CPU Performance** must be set to **Custom** in order to specify this value. For any value other than **Custom**, this option is overridden by the setting in the selected CPU performance profile.
DCU Streamer Pre-fetch	Whether the processor uses the DCU IP Prefetch mechanism to analyze historical cache access patterns and preload the most relevant lines in the L1 cache. This can be one of the following: • **disabled—**The processor does not try to anticipate cache read requirements and only fetches explicitly requested lines. • **enabled—**The DCU prefetcher analyzes the cache read pattern and prefetches the next line in the cache if it determines that it may be needed. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
DCU IP Pre-fetcher	Whether the processor uses the DCU IP Prefetch mechanism to analyze historical cache access patterns and preload the most relevant lines in the L1 cache. This can be one of the following: • **disabled—**The processor does not preload any cache data. • **enabled—**The DCU IP prefetcher preloads the L1 cache with the data it determines to be the most relevant. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Direct Cache Access	Allows processors to increase I/O performance by placing data from I/O devices directly into the processor cache. This setting helps to reduce cache misses. This can be one of the following: • **disabled—**Data from I/O devices is not placed directly into the processor cache. • **enabled—**Data from I/O devices is placed directly into the processor cache. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Processor C State	Whether the system can enter a power savings mode during idle periods. This can be one of the following: • **disabled—**The system remains in a high-performance state even when idle. • **enabled—**The system can reduce power to system components such as the DIMMs and CPUs. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. We recommend that you contact your operating system vendor to make sure your operating system supports this feature.

Name	Description
Processor C1E	Allows the processor to transition to its minimum frequency upon entering C1. This setting does not take effect until after you have rebooted the server. This can be one of the following: • **disabled—**The CPU continues to run at its maximum frequency in the C1 state. • **enabled—**The CPU transitions to its minimum frequency. This option saves the maximum amount of power in the C1 state. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Processor C3 Report	Whether the processor sends the C3 report to the operating system. This can be one of the following: • **disabled—**The processor does not send the C3 report. • **acpi-c2—**The processor sends the C3 report using the advanced configuration and power interface (ACPI) C2 format. • **acpi-c3—**The processor sends the C3 report using the ACPI C3 format. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. On the Cisco UCS B440 Server, the BIOS Setup menu uses enabled and disabled for these options. If you specify acpi-c2 or acpi-c2, the server sets the BIOS value for that option to enabled.
Processor C6 Report	Whether the processor sends the C6 report to the operating system. This can be one of the following: • **disabled—**The processor does not send the C6 report. • **enabled—**The processor sends the C6 report. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Processor C7 Report	Whether the processor sends the C7 report to the operating system. This can be one of the following: • **disabled—**The processor does not send the C7 report. • **enabled—**The processor sends the C7 report. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Processor CMCI field	Enables CMCI generation.
CPU Performance	Sets the CPU performance profile for the server. This can be one of the following: • **enterprise—**For M3 servers, all prefetchers and data reuse are enabled. For M1 and M2 servers, data reuse and the DCU IP prefetcher are enabled, and all other prefetchers are disabled. • **high-throughput—**Data reuse and the DCU IP prefetcher are enabled, and all other prefetchers are disabled. • **hpc—**All prefetchers are enabled and data reuse is disabled. This setting is also known as high-performance computing.
Max Variable MTRR Setting	Allows you to select the number of mean time to repair (MTRR) variables. This can be one of the following: • **auto-max—**BIOS uses the default value for the processor. • **8—**BIOS uses the number specified for the variable MTRR. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Local X2 APIC	Allows you to set the type of Application Policy Infrastructure Controller (APIC) architecture. This can be one of the following: • **xapic—**Uses the standard xAPIC architecture. • **x2apic—**Uses the enhanced x2APIC architecture to support 32 bit addressability of processors. • **auto—**Automatically uses the xAPIC architecture that is detected. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Power Technology	Enables you to configure the CPU power management settings for the following options: • Enhanced Intel Speedstep Technology • Intel Turbo Boost Technology • Processor Power State C6 Power Technology can be one of the following: • **disabled—**The server does not perform any CPU power management and any settings for the BIOS parameters mentioned above are ignored. • **Energy Efficient—**The server determines the best settings for the BIOS parameters mentioned above and ignores the individual settings for these parameters. • **performance—**The server automatically optimizes the performance for the BIOS parameters mentioned above. • **custom—**The server uses the individual settings for the BIOS parameters mentioned above. You must select this option if you want to change any of these BIOS parameters. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Energy Performance	Allows you to determine whether system performance or energy efficiency is more important on this server. This can be one of the following: • **performance** • **balanced-performance** • **balanced-energy** • **energy-efficient** • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** **Power Technology** must be set to **Custom** or the server ignores the setting for this parameter.
Frequency Floor Override	Whether the CPU is allowed to drop below the maximum non-turbo frequency when idle. This can be one of the following: • **disabled—** The CPU can drop below the maximum non-turbo frequency when idle. This option decreases power consumption but may reduce system performance. • **enabled—** The CPU cannot drop below the maximum non-turbo frequency when idle. This option improves system performance but may increase power consumption. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
P-STATE Coordination	Allows you to define how BIOS communicates the P-state support model to the operating system. There are 3 models as defined by the Advanced Configuration and Power Interface (ACPI) specification. • **hw-all—**The processor hardware is responsible for coordinating the P-state among logical processors with dependencies (all logical processors in a package). • **sw-all—**The OS Power Manager (OSPM) is responsible for coordinating the P-state among logical processors with dependencies (all logical processors in a physical package), and must initiate the transition on all of the logical processors. • **sw-all—**The OS Power Manager (OSPM) is responsible for coordinating the P-state among logical processors with dependencies (all logical processors in a package), and may initiate the transition on any of the logical processors in the domain. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** **Power Technology** must be set to **Custom** or the server ignores the setting for this parameter.
DRAM Clock Throttling	Allows you to tune the system settings between the memory bandwidth and power consumption. This can be one of the following: • **balanced—** DRAM clock throttling is reduced, providing a balance between performance and power. • **performance—**DRAM clock throttling is disabled, providing increased memory bandwidth at the cost of additional power. • **Energy Efficient—**DRAM clock throttling is increased to improve energy efficiency. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Channel Interleaving	Whether the CPU divides memory blocks and spreads contiguous portions of data across interleaved channels to enable simultaneous read operations. This can be one of the following: • **Auto—**The CPU determines what interleaving is done. • **1-way—**Some channel interleaving is used. • **2-way** • **3-way** • **4-way—**The maximum amount of channel interleaving is used. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Rank Interleaving	Whether the CPU interleaves physical ranks of memory so that one rank can be accessed while another is being refreshed. This can be one of the following: • **Auto—**The CPU determines what interleaving is done. • **1-way—**Some rank interleaving is used. • **2-way** • **4-way** • **8-way—**The maximum amount of rank interleaving is used. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Memory Interleaving	Whether the CPU interleaves the physical memory so that the memory can be accessed while another is being refreshed. This can be one of the following: • **Auto—**The CPU determines what interleaving is done. • **1-way—**Some memory interleaving is used. • **2-way** • **4-way** • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Demand Scrub	Whether the system corrects single bit memory errors encountered when the CPU or I/O makes a demand read. This can be one of the following: • **disabled—** Single bit memory errors are not corrected. • **enabled—** Single bit memory errors are corrected in memory and the corrected data is set in response to the demand read. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Patrol Scrub	Whether the system actively searches for, and corrects, single bit memory errors even in unused portions of the memory on the server. This can be one of the following: • **disabled—**The system checks for memory ECC errors only when the CPU reads or writes a memory address. • **enabled—**The system periodically reads and writes memory searching for ECC errors. If any errors are found, the system attempts to fix them. This option may correct single bit errors before they become multi-bit errors, but it may adversely affect performance when the patrol scrub is running. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Altitude	The approximate number of meters above sea level at which the physical server is installed. This can be one of the following: • **auto—**The CPU determines the physical elevation. • **—**The server is approximately 300 meters above sea level. • **—**The server is approximately 900 meters above sea level. • **—**The server is approximately 1500 meters above sea level. • **—**The server is approximately 3000 meters above sea level. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Package C State Limit	The amount of power available to the server components when they are idle. This can be one of the following: • • **—**The server may enter any available C state. • **—**The server provides all server components with full power at all times. This option maintains the highest level of performance and requires the greatest amount of power. • **—**When the CPU is idle, the system slightly reduces the power consumption. This option requires less power than C0 and allows the server to return quickly to high performance mode. • **—**When the CPU is idle, the system reduces the power consumption further than with the C1 option. This requires less power than C1 or C0, but it takes the server slightly longer to return to high performance mode. • **—**When the CPU is idle, the system reduces the power consumption further than with the C3 option. This option saves more power than C0, C1, or C3, but there may be performance issues until the server returns to full power. • **—**When the CPU is idle, the system reduces the power consumption further than with the C1 option. This requires less power than C1 or C0, but it takes the server slightly longer to return to high performance mode. • **—**When the CPU is idle, the server makes a minimal amount of power available to the components. This option saves the maximum amount of power but it also requires the longest time for the server to return to high performance mode. • **—**When the CPU is idle, the server makes a minimal amount of power available to the components. This option saves more power than C7, but it also requires the longest time for the server to return to high performance mode. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
CPU Hardware Power Management	Enables processor Hardware Power Management (HWPM). This can be one of the following: • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. • **—**HWPM is disabled. • **—**HWPM native mode is enabled. • **—**HWPM Out-Of-Box mode is enabled.
Energy Performance Tuning	Determines if the BIOS or Operating System can turn on the energy performance bias tuning. The options are BIOS and OS.
Workload Configuration	This feature allows for workload optimization. The options are Balanced and I/O Sensitive. Cisco recommends using Balanced.

Intel Directed I/O BIOS Settings

The following table lists the Intel Directed I/O BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
VT for Directed IO	Whether the processor uses Intel Virtualization Technology for Directed I/O (VT-d). This can be one of the following: • **disabled—**The processor does not use virtualization technology. • **enabled—**The processor uses virtualization technology. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** This option must be enabled if you want to change any of the other Intel Directed I/O BIOS settings.

Name	Description
Interrupt Remap	Whether the processor supports Intel VT-d Interrupt Remapping. This can be one of the following: • **disabled—**The processor does not support remapping. • **enabled—**The processor uses VT-d Interrupt Remapping as required. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Coherency Support	Whether the processor supports Intel VT-d Coherency. This can be one of the following: • **disabled—**The processor does not support coherency. • **enabled—**The processor uses VT-d Coherency as required. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
ATS Support	Whether the processor supports Intel VT-d Address Translation Services (ATS). This can be one of the following: • **disabled—**The processor does not support ATS. • **enabled—**The processor uses VT-d ATS as required. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Pass Through DMA Support	Whether the processor supports Intel VT-d Pass-through DMA. This can be one of the following: • **disabled—**The processor does not support pass-through DMA. • **enabled—**The processor uses VT-d Pass-through DMA as required. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

RAS Memory BIOS Settings

The following table lists the RAS memory BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Memory RAS Config	How the memory reliability, availability, and serviceability (RAS) is configured for the server. This can be one of the following: • **maximum-performance—**System performance is optimized. • **mirroring—**System reliability is optimized by using half the system memory as backup. • **lockstep—**If the DIMM pairs in the server have an identical type, size, and organization and are populated across the SMI channels, you can enable lockstep mode to minimize memory access latency and provide better performance. Lockstep is enabled by default for B440 servers. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
NUMA	Whether the BIOS supports NUMA. This can be one of the following: • **disabled—**The BIOS does not support NUMA. • **enabled—**The BIOS includes the ACPI tables that are required for NUMA-aware operating systems. If you enable this option, the system must disable Inter-Socket Memory interleaving on some platforms. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Mirroring Mode	Memory mirroring enhances system reliability by keeping two identical data images in memory. This option is only available if you choose the **mirroring** option for **Memory RAS Config**. It can be one of the following: • **inter-socket—**Memory is mirrored between two Integrated Memory Controllers (IMCs) across CPU sockets. • **intra-socket—**One IMC is mirrored with another IMC in the same socket. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Sparing Mode	Sparing optimizes reliability by holding memory in reserve so that it can be used in case other DIMMs fail. This option provides some memory redundancy, but does not provide as much redundancy as mirroring. The available sparing modes depend on the current memory population. This option is only available if you choose **sparing** option for **Memory RAS Config**. It can be one of the following: • **dimm-sparing—**One DIMM is held in reserve. If a DIMM fails, the contents of a failing DIMM are transferred to the spare DIMM. • **rank-sparing—**A spare rank of DIMMs is held in reserve. If a rank of DIMMs fails, the contents of the failing rank are transferred to the spare rank. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
LV DDR Mode	Whether the system prioritizes low voltage or high frequency memory operations. This can be one of the following: • **power-saving-mode—**The system prioritizes low voltage memory operations over high frequency memory operations. This mode may lower memory frequency in order to keep the voltage low. • **performance-mode—**The system prioritizes high frequency operations over low voltage operations. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
DRAM Refresh Rate **set dram-refresh-rate-config dram-refresh**	The refresh interval rate for internal memory. This can be one of the following: • **1x** • **2x** • **3x** • **4x** • **auto** • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
DDR3 Voltage Selection	The voltage to be used by the dual-voltage RAM. This can be one of the following: • **ddr3-1500mv** • **ddr3-1350mv** • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Serial Port BIOS Settings

The following table lists the serial port BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Serial Port A	Whether serial port A is enabled or disabled. This can be one of the following: • **disabled—**The serial port is disabled. • **enabled—**The serial port is enabled. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

USB BIOS Settings

The following table lists the USB BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Make Device Non Bootable	Whether the server can boot from a USB device. This can be one of the following: • **disabled—**The server can boot from a USB device. • **enabled—**The server cannot boot from a USB device. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Legacy USB Support	Whether the system supports legacy USB devices. This can be one of the following: • **disabled—**USB devices are only available to EFI applications. • **enabled—**Legacy USB support is always available. • **auto—**Disables legacy USB support if no USB devices are connected. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
USB System Idle Power Optimizing Setting	Whether the USB System Idle Power Optimizing setting is used to reduce USB EHCI idle power consumption. Depending upon the value you choose, this setting can have an impact on performance. This can be one of the following: • **high-performance—**The USB System Idle Power Optimizing setting is disabled, because optimal performance is preferred over power savings. Selecting this option can significantly improve performance. We recommend you select this option unless your site has server power restrictions. • **lower-idle-power—**The USB System Idle Power Optimizing setting is enabled, because power savings are preferred over optimal performance. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
USB Front Panel Access Lock	USB front panel lock is configured to enable or disable the front panel access to USB ports. This can be one of the following: • **disabled** • **enabled** • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Port 60/64 Emulation	Whether the system supports 60h/64h emulation for complete USB keyboard legacy support. This can be one of the following: • **disabled—**60h/64 emulation is not supported. • **enabled—**60h/64 emulation is supported. You should select this option if you are using a non-USB aware operating system on the server. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
USB Port:Front	Whether the front panel USB devices are enabled or disabled. This can be one of the following: • **disabled—**Disables the front panel USB ports. Devices connected to these ports are not detected by the BIOS and operating system. • **enabled—**Enables the front panel USB ports. Devices connected to these ports are detected by the BIOS and operating system. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
USB Port:Internal	Whether the internal USB devices are enabled or disabled. This can be one of the following: • **disabled—**Disables the internal USB ports. Devices connected to these ports are not detected by the BIOS and operating system. • **enabled—**Enables the internal USB ports. Devices connected to these ports are detected by the BIOS and operating system. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
USB Port:KVM	Whether the KVM ports are enabled or disabled. This can be one of the following: • **disabled—**Disables the KVM keyboard and/or mouse devices. Keyboard and/or mouse will not work in the KVM window. • **enabled—**Enables the KVM keyboard and/or mouse devices. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
USB Port:Rear	Whether the rear panel USB devices are enabled or disabled. This can be one of the following: • **disabled—**Disables the rear panel USB ports. Devices connected to these ports are not detected by the BIOS and operating system. • **enabled—**Enables the rear panel USB ports. Devices connected to these ports are detected by the BIOS and operating system. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
USB Port:SD Card	Whether the SD card drives are enabled or disabled. This can be one of the following: • **disabled—**Disables the SD card drives. The SD card drives are not detected by the BIOS and operating system. • **enabled—**Enables the SD card drives. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
USB Port:VMedia	Whether the virtual media devices are enabled or disabled. This can be one of the following: • **disabled—**Disables the vMedia devices. • **enabled—**Enables the vMedia devices. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
All USB Devices	Whether all physical and virtual USB devices are enabled or disabled. This can be one of the following: • **disabled—**All USB devices are disabled. • **enabled—**All USB devices are enabled. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
xHCI Mode Support	Whether xHCI mode support is enabled or disabled. This can be one of the following: • **disabled—**xHCI mode support is disabled. • **enabled—**xHCI mode support is enabled. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

PCI Configuration BIOS Settings

The following table lists the PCI configuration BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Max Memory Below 4G	Whether the BIOS maximizes memory usage below 4GB for an operating system without PAE support, depending on the system configuration. This can be one of the following: • **disabled—**Does not maximize memory usage. Choose this option for all operating systems with PAE support. • **enabled—**Maximizes memory usage below 4GB for an operating system without PAE support. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Memory Mapped IO Above 4Gb Config	Whether to enable or disable memory mapped I/O of 64-bit PCI devices to 4GB or greater address space. Legacy option ROMs are not able to access addresses above 4GB. PCI devices that are 64-bit compliant but use a legacy option ROM may not function correctly with this setting enabled. This can be one of the following: • **disabled—**Does not map I/O of 64-bit PCI devices to 4GB or greater address space. • **enabled—**Maps I/O of 64-bit PCI devices to 4GB or greater address space. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
VGA Priority	Allows you to set the priority for VGA graphics devices if multiple VGA devices are found in the system. This can be one of the following: • **onboard—**Priority is given to the onboard VGA device. BIOS post screen and OS boot are driven through the onboard VGA port. • **offboard—**Priority is given to the PCIE Graphics adapter. BIOS post screen and OS boot are driven through the external graphics adapter port. • **onboard-vga-disabled—**Priority is given to the PCIE Graphics adapter, and the onboard VGA device is disabled. **Note** The vKVM does not function when the onboard VGA is disabled. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** Only onboard VGA devices are supported with Cisco UCS B-Series servers.
ASPM Support	Allows you to set the level of ASPM (Active Power State Management) support in the BIOS. This can be one of the following: • **disabled—**ASPM support is disabled in the BIOS. • **auto—**The CPU determines the power state. • **forcel0—**Force all links to L0 standby (L0s) state. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

QPI BIOS Settings

The following table lists the QPI BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
QPI Link Frequency	The Intel QuickPath Interconnect (QPI) link frequency, in megatransfers per second (MT/s). This can be one of the following: • **6400** • **7200** • **8000** • **9600** • **Auto—**The CPU determines the QPI link frequency. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
QPI Snoop Mode	This can be one of the following: • **home-snoop—**The snoop is always spawned by the home agent (centralized ring stop) for the memory controller. This mode has a higher local latency than early snoop, but it provides extra resources for a larger number of outstanding transactions. • **cluster-on-die—**This mode is available only for processors that have 10 or more cores. It is the best mode for highly NUMA optimized workloads. • **home-directory-snoop-with-osb** • **early-snoop—**The distributed cache ring stops can send a snoop probe or a request to another caching agent directly. This mode has lower latency and it is best for workloads that have shared data sets across threads and can benefit from a cache-to-cache transfer, or for workloads that are not NUMA optimized. • **auto —**The CPU determines the QPI Snoop mode. • **platform-default —**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

LOM and PCIe Slots BIOS Settings

The following table lists the USB BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
PCIe Slot:SAS OptionROM	Whether Option ROM is available on the SAS port. This can be one of the following: • **disabled—**The expansion slot is not available. • **enabled—**The expansion slot is available. • **uefi-only —**The expansion slot is available for UEFI only. • **legacy-only—**The expansion slot is available for legacy only. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe Slot:*n* Link Speed	This option allows you to restrict the maximum speed of an adapter card installed in PCIe slot ***n***. This can be one of the following: • **gen1—**2.5GT/s (gigatransfers per second) is the maximum speed allowed. • **gen2—**5GT/s is the maximum speed allowed. • **gen3—**8GT/s is the maximum speed allowed. • **auto—**The maximum speed is set automatically. • **disabled—**The maximum speed is not restricted. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
PCIe Slot:*n* OptionROM	Whether Option ROM is available on the port. This can be one of the following: • **disabled—**The expansion slot is not available. • **enabled—**The expansion slot is available. • **uefi-only —**The expansion slot is available for UEFI only. • **legacy-only—**The expansion slot is available for legacy only. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe Slot:HBA OptionROM	Whether Option ROM is available on the HBA port. This can be one of the following: • **disabled—**The expansion slot is not available. • **enabled—**The expansion slot is available. • **uefi-only —**The expansion slot is available for UEFI only. • **legacy-only—**The expansion slot is available for legacy only. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe Slot:MLOM OptionROM	Whether Option ROM is available on the MLOM port. This can be one of the following: • **disabled—**The expansion slot is not available. • **enabled—**The expansion slot is available. • **uefi-only —**The expansion slot is available for UEFI only. • **legacy-only—**The expansion slot is available for legacy only. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
PCIe Slot:N1 OptionROM	Whether Option ROM is available on the port. This can be one of the following: • **disabled—**The expansion slot is not available. • **enabled—**The expansion slot is available. • **uefi-only —**The expansion slot is available for UEFI only. • **legacy-only—**The expansion slot is available for legacy only. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe Slot:N2 OptionROM	Whether Option ROM is available on the port. This can be one of the following: • **Disabled—**The expansion slot is not available. • **Enabled—**The expansion slot is available. • **uefi-only —**The expansion slot is available for UEFI only. • **legacy-only—**The expansion slot is available for legacy only. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe 10G LOM 2 Link	Whether Option ROM is available on the 10G LOM port. This can be one of the following: • **disabled—**The expansion slot is not available. • **enabled—**The expansion slot is available. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
PCI ROM CLP **set pci-rom-clp-support pci-rom-clp-config**	• **disabled—**The expansion slot is not available. • **enabled—**The expansion slot is available. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SIOC1 Option ROM **set sioc1-optionrom-config sioc1-optionrom**	• **disabled—**The expansion slot is not available. • **enabled—**The expansion slot is available. • **uefi-only —**The expansion slot is available for UEFI only. • **legacy-only—**The expansion slot is available for legacy only. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SIOC2 Option ROM **set sioc2-optionrom-config sioc2-optionrom**	• **disabled—**The expansion slot is not available. • **enabled—**The expansion slot is available. • **uefi-only —**The expansion slot is available for UEFI only. • **legacy-only—**The expansion slot is available for legacy only. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
SB MEZZ1 Option ROM **set sbmezz1-optionrom-config sbmezz1-optionrom**	• **disabled—**The expansion slot is not available. • **enabled—**The expansion slot is available. • **uefi-only —**The expansion slot is available for UEFI only. • **legacy-only—**The expansion slot is available for legacy only. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
IOE Slot1 Option ROM **set ioeslot1-optionrom-config ioeslot1-optionrom**	• **disabled—**The expansion slot is not available. • **enabled—**The expansion slot is available. • **uefi-only —**The expansion slot is available for UEFI only. • **legacy-only—**The expansion slot is available for legacy only. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
IOE MEZZ 1 Option ROM **set ioemezz1-optionrom-config ioemezz1-optionrom**	• **disabled—**The expansion slot is not available. • **enabled—**The expansion slot is available. • **uefi-only —**The expansion slot is available for UEFI only. • **legacy-only—**The expansion slot is available for legacy only. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
IOE Slot2 Option ROM **set ioeslot2-optionrom-config ioeslot2-optionrom**	• **disabled—**The expansion slot is not available. • **enabled—**The expansion slot is available. • **uefi-only —**The expansion slot is available for UEFI only. • **legacy-only—**The expansion slot is available for legacy only. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
IO ENVME1 Option ROM **set ioenvme1-optionrom-config ioenvme1-optionrom**	• **disabled—**The expansion slot is not available. • **enabled—**The expansion slot is available. • **uefi-only —**The expansion slot is available for UEFI only. • **legacy-only—**The expansion slot is available for legacy only. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
IO ENVME2 Option ROM **set ioenvme2-optionrom-config ioenvme2-optionrom**	• **disabled—**The expansion slot is not available. • **enabled—**The expansion slot is available. • **uefi-only —**The expansion slot is available for UEFI only. • **legacy-only—**The expansion slot is available for legacy only. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
SBNVME1 Option ROM **set sbnvme1-optionrom-config sbnvme1-optionrom**	• **disabled—**The expansion slot is not available. • **enabled—**The expansion slot is available. • **uefi-only —**The expansion slot is available for UEFI only. • **legacy-only—**The expansion slot is available for legacy only. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Graphics Configuration BIOS Settings

The following tables list the graphics configuration BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Integrated Graphics	Enables integrated graphics. This can be one of the following: • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. • **—**Integrated graphic is enabled. • **—**Integrated graphics is disabled.
Aperture Size	Allows you to set the size of mapped memory for the integrated graphics controller. This can be one of the following: • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. • • • • • •

Name	Description
Onboard Graphics	Enables onboard graphics (KVM). This can be one of the following: • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. • **—**Onboard graphics is enabled. • **—**Onboard graphics is disabled.

Boot Options BIOS Settings

The following table lists the boot options BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Boot Option Retry	Whether the BIOS retries NON-EFI based boot options without waiting for user input. This can be one of the following: • **disabled—**Waits for user input before retrying NON-EFI based boot options. • **enabled—**Continually retries NON-EFI based boot options without waiting for user input. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Intel Entry SAS RAID	Whether the Intel SAS Entry RAID Module is enabled. This can be one of the following: • **disabled—**The Intel SAS Entry RAID Module is disabled. • **enabled—**The Intel SAS Entry RAID Module is enabled. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Intel Entry SAS RAID Module	How the Intel SAS Entry RAID Module is configured. This can be one of the following: • **it-ir-raid**—Configures the RAID module to use Intel IT/IR RAID. • **intel-esrtii**—Configures the RAID module to use Intel Embedded Server RAID Technology II. • **Platform Default**—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Onboard SCU Storage Support	Whether the onboard software RAID controller is available to the server. This can be one of the following: • **disabled**—The software RAID controller is not available. • **enabled**—The software RAID controller is available. • **Platform Default**—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Note BIOS parameter virtualization capability in Cisco UCS Manager maps a unified set of BIOS settings in a service profile to the actual BIOS supporting parameters. However, not all BIOS setting items are applicable to every server model/platform. When you create a custom BIOS policy and have the **Boot Option Retry** selected, and when there is no bootable option available, the reboot fails on the Cisco UCS B420 M3 or Cisco UCS B420 M4 servers and Cisco UCS Manager displays this message : *Reboot and Select proper Boot device or Insert Boot Media in selected Boot device and press a key*. You must manually set a boot option after the boot path is corrected, in order to enable the servers to reboot after a power outage. For more information about BIOS default server policies and the BIOS options and their default settings, see BIOS Policy, on page 458 and Server BIOS Settings, on page 417.

Server Management BIOS Settings

The following tables list the server management BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

General Settings

Name	Description
Assert Nmi on Serr	Whether the BIOS generates a non-maskable interrupt (NMI) and logs an error when a system error (SERR) occurs. This can be one of the following: • **disabled—**The BIOS does not generate an NMI or log an error when a SERR occurs. • **enabled—**The BIOS generates an NMI and logs an error when a SERR occurs. You must enable this setting if you want to enable **Assert Nmi on Perr**. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Assert Nmi on Perr	Whether the BIOS generates a non-maskable interrupt (NMI) and logs an error when a processor bus parity error (PERR) occurs. This can be one of the following: • **disabled—**The BIOS does not generate an NMI or log an error when a PERR occurs. • **enabled—**The BIOS generates an NMI and logs an error when a PERR occurs. You must enable **Assert Nmi on Serr** to use this setting. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
OS Boot Watchdog Timer	Whether the BIOS programs the watchdog timer with a predefined timeout value. If the operating system does not complete booting before the timer expires, the CIMC resets the system and an error is logged. This can be one of the following: • **disabled—**The watchdog timer is not used to track how long the server takes to boot. • **enabled—**The watchdog timer tracks how long the server takes to boot. If the server does not boot within the predefined length of time, the CIMC resets the system and logs an error. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. This feature requires either operating system support or Intel Management software.

Name	Description
OS Boot Watchdog Timer Timeout Policy	What action the system takes if the watchdog timer expires. This can be one of the following: • **power-off—**The server is powered off if the watchdog timer expires during OS boot. • **reset—**The server is reset if the watchdog timer expires during OS boot. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. This option is only available if you enable the OS Boot Watchdog Timer.
OS Boot Watchdog Timer Timeout	What timeout value the BIOS uses to configure the watchdog timer. This can be one of the following: • **5-minutes—**The watchdog timer expires 5 minutes after the OS begins to boot. • **10-minutes—**The watchdog timer expires 10 minutes after the OS begins to boot. • **15-minutes—**The watchdog timer expires 15 minutes after the OS begins to boot. • **20-minutes—**The watchdog timer expires 20 minutes after the OS begins to boot. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. This option is only available if you enable the OS Boot Watchdog Timer.
FRB-2 Timer	Whether the FRB-2 timer is used to recover the system if it hangs during POST. This can be one of the following: • **Disabled—**The FRB-2 timer is not used. • **Enabled—**The FRB-2 timer is started during POST and used to recover the system if necessary. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Console Redirection Settings

Name	Description
Console Redirection	Allows a serial port to be used for console redirection during POST and BIOS booting. After the BIOS has booted and the operating system is responsible for the server, console redirection is irrelevant and has no effect. This can be one of the following: • **disabled—**No console redirection occurs during POST. • **serial-port-a—**Enables serial port A for console redirection during POST. This option is valid for blade servers and rack-mount servers. • **serial-port-b—**Enables serial port B for console redirection and allows it to perform server management tasks. This option is only valid for rack-mount servers. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** If you enable this option, you also disable the display of the Quiet Boot logo screen during POST.
Flow Control	Whether a handshake protocol is used for flow control. Request to Send / Clear to Send (RTS/CTS) helps to reduce frame collisions that can be introduced by a hidden terminal problem. This can be one of the following: • **none—**No flow control is used. • **rts-cts—**RTS/CTS is used for flow control. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** This setting must match the setting on the remote terminal application.

Name	Description
BAUD Rate	What BAUD rate is used for the serial port transmission speed. If you disable Console Redirection, this option is not available. This can be one of the following: • **9600—**A 9600 BAUD rate is used. • **19200—**A 19200 BAUD rate is used. • **38400—**A 38400 BAUD rate is used. • **57600—**A 57600 BAUD rate is used. • **115200—**A 115200 BAUD rate is used. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** This setting must match the setting on the remote terminal application.
Terminal Type	What type of character formatting is used for console redirection. This can be one of the following: • **pc-ansi—**The PC-ANSI terminal font is used. • **vt100—**A supported vt100 video terminal and its character set are used. • **vt100-plus—**A supported vt100-plus video terminal and its character set are used. • **vt-utf8—**A video terminal with the UTF-8 character set is used. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** This setting must match the setting on the remote terminal application.
Legacy OS Redirect	Whether redirection from a legacy operating system, such as DOS, is enabled on the serial port. This can be one of the following: • **disabled—**The serial port enabled for console redirection is hidden from the legacy operating system. • **enabled—** The serial port enabled for console redirection is visible to the legacy operating system. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Putty KeyPad **set console-redir-config putty-function-keypad**	Allows you to change the action of the PuTTY function keys and the top row of the numeric keypad. This can be one of the following: • **vt100—**The function keys generate ESC OP through ESC O[. • **linux—**Mimics the Linux virtual console. Function keys F6 to F12 behave like the default mode, but F1 to F5 generate ESC [[A through ESC [[E. • **xtermr6—**Function keys F5 to F12 behave like the default mode. Function keys F1 to F4 generate **ESC OP** through ESC OS, which are the sequences produced by the top row of the keypad on Digital terminals. • **sco—**The function keys F1 to F12 generate ESC [M through ESC [X. The function and shift keys generate ESC [Y through ESC [j. The control and function keys generate ESC [k through ESC [v. The shift, control and function keys generate ESC [w through ESC [{. • **escn—**The default mode. The function keys match the general behavior of Digital terminals. The function keys generate sequences such as ESC [11~ and ESC [12~. • **vt400—**The function keys behave like the default mode. The top row of the numeric keypad generates ESC OP through ESC OS. • **Platform Default—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Out of Band Management	Used for Windows Special Administration Control (SAC).
Redirection After BIOS POST	

BIOS Policy

The BIOS policy is a policy that automates the configuration of BIOS settings for a server or group of servers. You can create global BIOS policies available to all servers in the root organization, or you can create BIOS policies in sub-organizations that are only available to that hierarchy.

To use a BIOS policy, do the following:

1. Create the BIOS policy in Cisco UCS Manager.
2. Assign the BIOS policy to one or more service profiles.
3. Associate the service profile with a server.

During service profile association, Cisco UCS Manager modifies the BIOS settings on the server to match the configuration in the BIOS policy. If you do not create and assign a BIOS policy to a service profile, the server uses the default BIOS settings for that server platform.

Default BIOS Settings

Cisco UCS Manager includes a set of default BIOS settings for each type of server supported by Cisco UCS. The default BIOS settings are available only in the root organization and are global. Only one set of default BIOS settings can exist for each server platform supported by Cisco UCS. You can modify the default BIOS settings, but you cannot create an additional set of default BIOS settings.

Each set of default BIOS settings are designed for a particular type of supported server and are applied to all servers of that specific type which do not have a BIOS policy included in their service profiles.

Unless a Cisco UCS implementation has specific needs that are not met by the server-specific settings, we recommend that you use the default BIOS settings that are designed for each type of server in the Cisco UCS domain.

Cisco UCS Manager applies these server platform-specific BIOS settings as follows:

- The service profile associated with a server does not include a BIOS policy.
- The BIOS policy is configured with the platform-default option for a specific setting.

You can modify the default BIOS settings provided by Cisco UCS Manager. However, any changes to the default BIOS settings apply to all servers of that particular type or platform. If you want to modify the BIOS settings for only certain servers, we recommend that you use a BIOS policy.

Creating a BIOS Policy

Note Cisco UCS Manager pushes BIOS configuration changes through a BIOS policy or default BIOS settings to the Cisco Integrated Management Controller (CIMC) buffer. These changes remain in the buffer and do not take effect until the server is rebooted.

We recommend that you verify the support for BIOS settings in the server that you want to configure. Some settings, such as Mirroring Mode for RAS Memory, are not supported by all Cisco UCS servers.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **BIOS Policies** and select **Create BIOS Policy**.

Step 5 On the **Main** page of the **Create BIOS Policy** wizard, enter a name for the BIOS policy in the **Name** field.

This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Step 6 In the **Create BIOS Policy** wizard, do the following to configure the BIOS settings:

a) If you want to change a BIOS setting, click the desired radio button or make the appropriate choice from the drop-down list.
For descriptions and information about the options for each BIOS setting, see the following topics:

- **Main** page: Main BIOS Settings, on page 418
- **Processor** page: Processor BIOS Settings, on page 420
- **Intel Directed IO** page: Intel Directed I/O BIOS Settings, on page 433
- **RAS Memory** page: RAS Memory BIOS Settings, on page 435
- **Serial Port** page: Serial Port BIOS Settings, on page 437
- **USB** page: USB BIOS Settings, on page 438
- **PCI Configuration** page: PCI Configuration BIOS Settings, on page 441
- **Boot Options** page: Boot Options BIOS Settings, on page 452
- **Server Management** page: Server Management BIOS Settings, on page 453

b) Click **Next** after each page.

Step 7 After you have configured all of the BIOS settings for the policy, click **Finish**.

Modifying the BIOS Defaults

We recommend that you verify the support for BIOS settings in the server that you want to configure. Some settings, such as Mirroring Mode for RAS Memory, are not supported by all Cisco UCS servers.

Unless a Cisco UCS implementation has specific needs that are not met by the server-specific settings, we recommend that you use the default BIOS settings that are designed for each type of server in the Cisco UCS domain.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Expand **BIOS Defaults** and select the server model number for which you want to modify the default BIOS settings.

Step 5 In the **Work** pane, click the appropriate tab and then click the desired radio button or make a choice from the drop-down list to modify the default BIOS settings:

For descriptions and information about the options for each BIOS setting, see the following topics. Not all BIOS settings are available for each type of server.

- **Main** tab: Main BIOS Settings, on page 418
- **Advanced** tab:
  - **Processor** subtab: Processor BIOS Settings, on page 420
  - **Intel Directed IO** subtab: Intel Directed I/O BIOS Settings, on page 433
  - **RAS Memory** subtab: RAS Memory BIOS Settings, on page 435
  - **Serial Port** subtab: Serial Port BIOS Settings, on page 437
  - **USB** subtab: USB BIOS Settings, on page 438
  - **PCI Configuration** subtab: PCI Configuration BIOS Settings, on page 441
- **Boot Options** tab: Boot Options BIOS Settings, on page 452
- **Server Management** tab: Server Management BIOS Settings, on page 453

Step 6 Click **Save Changes**.

Viewing the Actual BIOS Settings for a Server

Follow this procedure to see the actual BIOS settings on a server.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.

Step 3 Choose the server for which you want to view the actual BIOS settings.

Step 4 On the **Work** pane, click the **Inventory** tab.

Step 5 Click the **Motherboard** subtab.

Step 6 In the **BIOS Settings** area, click the **Expand** icon to the right of the heading to open that area.
Each tab in the **BIOS Settings** area displays the settings for that server platform. Some of the tabs contain subtabs with additional information.

Configuring Trusted Platform Module

Trusted Platform Module

The Trusted Platform Module (TPM) is a component that can securely store artifacts that are used to authenticate the server. These artifacts can include passwords, certificates, or encryption keys. A TPM can also be used to store platform measurements that help ensure that the platform remains trustworthy. Authentication (ensuring that the platform can prove that it is what it claims to be) and attestation (a process helping to prove that a platform is trustworthy and has not been breached) are necessary steps to ensure safer computing in all environments. It is a requirement for the Intel Trusted Execution Technology (TXT) security feature, which must be enabled in the BIOS settings for a server equipped with a TPM. Cisco UCS M4 blade and rack-mount servers include support for TPM. TPM is enabled by default on these servers.

Important

- If you upgrade Cisco UCS Manager to Release 2.2(4), TPM is enabled.
- When TPM is enabled and you downgrade Cisco UCS Manager from Release 2.2(4), TPM is disabled.

Intel Trusted Execution Technology

Intel Trusted Execution Technology (TXT) provides greater protection for information that is used and stored on the business server. A key aspect of that protection is the provision of an isolated execution environment and associated sections of memory where operations can be conducted on sensitive data, invisible to the rest of the system. Intel TXT provides for a sealed portion of storage where sensitive data such as encryption keys can be kept, helping to shield them from being compromised during an attack by malicious code. Cisco UCS M4 blade and rack-mount servers include support for TXT. TXT is disabled by default on these servers.

TXT can be enabled only after TPM, Intel Virtualization technology (VT) and Intel Virtualization Technology for Directed I/O (VT-d) are enabled. When you only enable TXT, it also implicitly enables TPM, VT, and VT-d.

Configuring Trusted Platform

Cisco UCS M4 blade and rack-mount servers include support for TPM and TXT. UCS Manager Release 2.2(4) allows you to perform the following operations on TPM and TXT:

- Configuring Trusted Platform, on page 463
- Clearing TPM for a Blade Server, on page 699 or Clearing TPM for a Rack-Mount Server, on page 713

Note For Cisco UCS M3 blade servers, press **F2** to enter the BIOS setup menu and change the settings.

Configuring Trusted Platform

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to configure TPM.

Step 4 Expand **BIOS Policies** and select the BIOS policy for which you want to configure TPM.

Step 5 In the **Work** pane, click the **Advanced** tab.

Step 6 Click the **Trusted Platform** subtab.

Step 7 To configure TPM, click one of the following:

Option	Description
disabled	Disables TPM
enable	Enables TPM
Platform Default	Enables TPM

Step 8 To configure TXT, click one of the following:

Option	Description
disabled	Disables TXT
enable	Enables TXT
Platform Default	Disables TXT

Step 9 Click **Save Changes**.

Consistent Device Naming

When there is no mechanism for the Operating System to label Ethernet interfaces in a consistent manner, it becomes difficult to manage network connections with server configuration changes. Consistent Device Naming (CDN), introduced in Cisco UCS Manager Release 2.2(4), allows Ethernet interfaces to be named in a consistent manner. This makes Ethernet interface names more persistent when adapter or other configuration changes are made.

To configure CDN for a vNIC, do the following:

- Enable consistent device naming in the BIOS policy.
- Associate the BIOS policy with a service profile.
- Configure consistent naming for a vNIC.

Guidelines and Limitations for Consistent Device Naming

- CDN is supported only on Windows 2012 R2. It is not supported on any other Operating System.
- Consistent device naming (CDN) is supported on all M3 and higher blade and rack-mount servers.
- BIOS and adapter firmware must be part of the Release 2.2(4) bundle to support CDN.
- In Cisco UCS Manager Release 2.2(4), CDN is supported only on the following adapters:
  - Cisco UCS VIC 1225 (UCSC-PCIE-CSC-02)
  - Cisco UCS MLOM 1227 (UCSC-MLOM-CSC-02)
  - Cisco UCS VIC 1225T (UCSC-PCIE-C10T-02)
  - Cisco UCS MLOM 1227T (UCSC-MLOM-C10T-02)
  - Cisco UCS VIC 1240 (UCSB-MLOM-40G-01)
  - Cisco UCS VIC 1280 (UCS-VIC-M82-8P)
  - Cisco UCS VIC 1340 (UCSB-MLOM-40G-03)
  - Cisco UCS VIC 1380 (UCSB-VIC-M83-8P)
- CDN is not supported for vNIC template and dynamic vNIC.
- Multiple vNICs within the same service profile cannot have the same CDN name.
- When a CDN name is not specified for a vNIC, the vNIC name is used as the CDN name.
- The CDN name that you configure for a vNIC appears as **Admin CDN Name**. The CDN name that is finally applied to the vNIC appears as **Oper CDN Name**. For example, if the **Admin CDN Name** for a vNIC called "vnic0" is cdn0, then the **Oper CDN Name** for this vNIC will be cdn0, but if the **Admin CDN Name** for the same vNIC is not specified, the **Oper CDN Name** will be vnic0.
- In Cisco UCS Manager Release 2.2(4), downgrade of Cisco UCS Manager is prevented if CDN is enabled in a BIOS policy that is assigned to an associated server.
- In Cisco UCS Manager Release 2.2(4), downgrade of the BIOS firmware is prevented if a CDN-enabled BIOS policy is assigned to a server.
- In Cisco UCS Manager Release 2.2(4), downgrade of the adapter firmware is prevented if a CDN-enabled BIOS policy is assigned to a server.
- When the applied BIOS policy is changed from CDN-disabled to CDN-enabled or from CDN-enabled to CDN-disabled, the host reboots with a warning, irrespective of whether reboot on BIOS update is enabled or not.
- It is recommended that you enable CDN in the BIOS policy and add CDN names to the vNICS before the Windows Operating System is installed.
- If the Windows Operating System is already installed on the server and CDN is then enabled in the BIOS policy, do the following:
  1. Uninstall the network drivers.
  2. Scan the system for hidden devices and uninstall them.

3 Rescan the system for new hardware and install the network drivers again.

If this is not done, the vNICs will not come up with the configured CDN names.

- When the applied BIOS policy is changed from CDN-disabled to CDN-enabled or from CDN-enabled to CDN-disabled on a service profile, do the following:

1 Uninstall the network drivers.

2 Scan the system for hidden devices and delete them.

3 Rescan the system for new hardware and install the network drivers again.

Note When the BIOS policy is changed from CDN-enabled to CDN-disabled, ensure that the CDN names are removed from all the vNICs on the system.

- If any change is made to the vNICs, the BDF of all the devices on the system also changes. Following are some of the scenarios that trigger a change in the BDF of all the vNICs present on the system:
  - When a vNIC is added or deleted
  - When a vNIC is moved from one adapter on the system to another adapter on the system

When these changes are made to the system, do the following:

1 Uninstall the network driver from all the present network interfaces.

2 Scan the system for hidden devices and uninstall them.

3 Rescan the system for new hardware and install the network driver on the network controllers again.

If the hidden devices are not deleted, the CDN names of the network adapters will not appear as configured on Cisco UCS Manager.

CDN with a Mixed Set of Adapters

When a CDN name is configured for a vNIC in a system with a mixed set of CDN-supported adapters and CDN-unsupported adapters, then system placement may not place CDN-configured vNICs on adapters that support CDN.

If CDN is enabled in the BIOS policy, and system placement places a CDN-configured vNIC (Admin CDN configured) on an adapter that does not support CDN, an info fault will be raised, but the configuration issue for the service profile will be ignored.

If CDN is enabled in the BIOS policy, and system placement places a vNIC (Admin CDN not configured) on an adapter that does not support CDN, an info fault will be raised, but the configuration issue for the service profile will be ignored. The **Oper CDN Name** in this case will be empty and will not be derived from the vNIC name.

If you want to deploy the CDN name as the host network interface name for a server, you must manually place a vNIC on a supported adapter.

Configuring Consistent Device Naming in a BIOS Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to configure Consistent Device Naming (CDN).

Step 4 Expand **BIOS Policies** and select the BIOS policy for which you want to configure CDN.

Step 5 In the **Work** pane, click the **Main** tab.

Step 6 In the **Properties** area, click one of the following in the **Consistent Device Naming** field to configure CDN:

Option	Description
disabled	Disables CDN in the BIOS policy
enable	Enables CDN in the BIOS policy
Platform Default	The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Step 7 Click **Save Changes**.

Configuring a CDN Name for a vNIC

When a CDN name is not specified for a vNIC, the vNIC name is used as the CDN name.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that contains the vNIC for which you want to configure a CDN name.

Step 4 Expand the service profile and **vNICs** node that contain the vNIC for which you want to configure a CDN name.

Step 5 Select the vNIC.

Step 6 In the **Work** pane, click the **General** tab.

Step 7 In the **Properties** area, enter the CDN name for the vNIC in the**Admin CDN Name** field.

Note The CDN name that you configure for a vNIC appears as **Admin CDN Name**. The CDN name that is finally applied to the vNIC appears as **Oper CDN Name**. For example, if the **Admin CDN Name** for a vNIC called "vnic0" is cdn0, then the **Oper CDN Name** for this vNIC will be cdn0, but if the **Admin CDN Name** for the same vNIC is not specified, the **Oper CDN Name** will be vnic0.

Step 8 Click **Save Changes**.

CIMC Security Policies

Cisco UCS Manager provides the following policies to increase security:

- KVM Management Policy
- IPMI Access Profile

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 CIMC. 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 can also restrict remote connectivity by disabling or enabling IPMI over LAN in the IPMI access profile. IPMI over LAN is disabled by default on all unassociated servers, and on all servers without an IPMI access policy. When an IPMI access policy is created, the IPMI over LAN is set to enabled by default. If you do not change the value to disabled, IPMI over LAN will be enabled on all associated servers.

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

Creating an IPMI Access Profile

Before You Begin

An IPMI profile requires that one or more of the following resources already exist in the system:

- Username with appropriate permissions that can be authenticated by the operating system of the server
- Password for the username
- Permissions associated with the username

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **IPMI Access Profiles** and select **Create IPMI Access Profile**.

Step 5 In the **Create IPMI Access Profile** dialog box:

a) Enter a unique name and description for the profile.

b) In the **IPMI Over LAN** field, choose whether to allow or restrict remote connectivity.

c) Click **OK**.

Step 6 In the **IPMI Users** area of the navigator, click +.

Step 7 In the **Create IPMI User** dialog box:

a) Complete the following fields:

Name	Description
Name field	The username to associate with this IPMI profile. Enter 1 to 16 alphanumeric characters. You can also use @ (at sign), _ (underscore), and - (hyphen). You cannot change this name once the profile has been saved.
Password field	The password associated with this username. Enter 1 to 20 standard ASCII characters, except for = (equal sign), $ (dollar sign), and | (vertical bar).
Confirm Password field	The password a second time for confirmation purposes.
Role field	The user role. This can be one of the following: • **Admin** • **Read Only**

b) Click **OK**.

Step 8 Repeat Steps 6 and 7 to add another user.

Step 9 Click **OK** to return to the IPMI profiles in the **Work** pane.

What to Do Next

Include the IPMI profile in a service profile and/or template.

Deleting an IPMI Access Profile

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 In the **Servers** tab, expand **Servers** > **Policies** > *Organization_Name*

Step 3 Expand the **IPMI Profiles** node.

Step 4 Right-click the profile you want to delete and select **Delete**.

Step 5 If a confirmation dialog box displays, click **Yes**.

KVM Management Policy

The KVM Management policy allows you to determine whether vMedia encryption is enabled when you access a server via KVM.

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

Note After a KVM vMedia session is mapped, if you change the KVM management policy, it will result in a loss of the vMedia session. You must re-map the KVM vMedia session again.

Creating a KVM Management Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **KVM Management Policies** and select **Create KVM Management Policy**.

Step 5 In the **Create KVM Management Policy** dialog box:

a) Enter a unique name and description for the policy.

b) In the **vMedia Encryption** field, choose whether to enable vMedia encryption.

c) Click **OK**.

Note After a KVM vMedia session is mapped, if you change the KVM management policy, it will result in a loss of the vMedia session. You must re-map the KVM vMedia session again.

Configuring Local Disk Configuration Policies

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:

- **No Local Storage—**For a diskless server 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.

- **RAID 0 Striped—**Data is striped across all disks in the array, providing fast throughput. There is no data redundancy, and all data is lost if any disk fails.
- **RAID 1 Mirrored—**Data is written to two disks, providing complete data redundancy if one disk fails. The maximum array size is equal to the available space on the smaller of the two drives.
- **Any Configuration—**For a server configuration that carries forward the local disk configuration without any changes.
- **No RAID—**For a server configuration that removes the RAID and leaves the disk MBR and payload unaltered.

  If you choose **No RAID** and you apply this policy to a server that already has an operating system with RAID storage configured, the system does not remove the disk contents. Therefore, there may be no visible differences on the server after you apply the **No RAID** mode. This can lead to a mismatch between the RAID configuration in the policy and the actual disk configuration shown in the **Inventory** > **Storage** tab for the server.

  To make sure that any previous RAID configuration information is removed from a disk, apply a scrub policy that removes all disk information after you apply the **No RAID** configuration mode.
- **RAID 5 Striped Parity—**Data is striped across all disks in the array. Part of the capacity of each disk stores parity information that can be used to reconstruct data if a disk fails. RAID 5 provides good data throughput for applications with high read request rates.
- **RAID 6 Striped Dual Parity—**Data is striped across all disks in the array and two parity disks are used to provide protection against the failure of up to two physical disks. In each row of data blocks, two sets of parity data are stored.
- **RAID 10 Mirrored and Striped—**RAID 10 uses mirrored pairs of disks to provide complete data redundancy and high throughput rates.
- **RAID 50 Striped Parity and Striped —**Data is striped across multiple striped parity disk sets to provide high throughput and multiple disk failure tolerance.
- **RAID 60 Striped Dual Parity and Striped —**Data is striped across multiple striped dual parity disk sets to provide high throughput and greater disk failure tolerance.

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

Note For a Cisco UCS C-Series server integrated with Cisco UCS Manager, with an embedded on-board RAID controller, the local disk mode should always be **Any Configuration**, and the RAID must be configured directly on the controller.

Guidelines for all Local Disk Configuration Policies

Before you create a local disk configuration policy, consider the following guidelines:

No Mixed HDDs and SSDs

Do not include HDDs and SSDs in a single server or RAID configuration.

Do Not Assign a Service Profile with the Default Local Disk Configuration Policy from a B200 M1 or M2 to a B200 M3

Due to the differences in the RAID/JBOD support provided by the storage controllers of B200 M1 and M2 servers and those of the B200 M3 server, you cannot assign or re-assign a service profile that includes the default local disk configuration policy from a B200M1 or M2 server to a B200 M3 server. The default local disk configuration policy includes those with Any Configuration or JBOD configuration.

JBOD Mode Support

The B200 M3 server supports JBOD mode for local disks.

Note Only B200 M1, B200 M2, B200 M3, B250 M1, B250 M2 and B22 M3 blade servers support the JBOD mode for local disks.

Guidelines for Local Disk Configuration Policies Configured for RAID

Configure RAID Settings in Local Disk Configuration Policy for Servers with MegaRAID Storage Controllers

If a blade server or integrated rack-mount server has a MegaRAID controller, you must configure RAID settings for the drives in the Local Disk Configuration policy included in the service profile for that server. You can do this either by configuring the local disk configuration policy in the service profile using one of the defined RAID modes for that server, or you can use the **Any Configuration** mode with the LSI Utilities toolset to create the RAID volumes.

If you do not configure your RAID LUNs before installing the OS, disk discovery failures might occur during the installation and you might see error messages such as “No Device Found.”

Server May Not Boot After RAID1 Cluster Migration if Any Configuration Mode Specified in Service Profile

After RAID1 clusters are migrated, you need to associate a service profile with the server. If the local disk configuration policy in the service profile is configured with **Any Configuration** mode rather than **RAID1**, the RAID LUN remains in "inactive" state during and after association. As a result, the server cannot boot.

To avoid this issue, ensure that the service profile you associate with the server contains the identical local disk configuration policy as the original service profile before the migration and does not include the **Any Configuration** mode.

Do Not Use JBOD Mode on Servers with MegaRAID Storage Controllers

Do not configure or use JBOD mode or JBOD operations on any blade server or integrated rack-mount server with a MegaRAID storage controllers. JBOD mode and operations are not intended for nor are they fully functional on these servers.

Maximum of One RAID Volume and One RAID Controller in Integrated Rack-Mount Servers

A rack-mount server that has been integrated with Cisco UCS Manager can have a maximum of one RAID volume irrespective of how many hard drives are present on the server.

All the local hard drives in an integrated rack-mount server must be connected to only one RAID Controller. Integration with Cisco UCS Manager does not support the connection of local hard drives to multiple RAID

Controllers in a single rack-mount server. We therefore recommend that you request a single RAID Controller configuration when you order rack-mount servers to be integrated with Cisco UCS Manager.

In addition, do not use third party tools to create multiple RAID LUNs on rack-mount servers. Cisco UCS Manager does not support that configuration.

Maximum of One RAID Volume and One RAID Controller in Blade Servers

A blade server can have a maximum of one RAID volume irrespective of how many drives are present in the server. All the local hard drives must be connected to only one RAID controller. For example, a B200 M3 server has an LSI controller and an Intel Patsburg controller, but only the LSI controller can be used as a RAID controller.

In addition, do not use third party tools to create multiple RAID LUNs on blade servers. Cisco UCS Manager does not support that configuration.

Number of Disks Selected in Mirrored RAID Should Not Exceed Two

If the number of disks selected in the Mirrored RAID exceed two, RAID 1 is created as a RAID 10 LUN. This issue can occur with the Cisco UCS B440 M1 and B440 M2 servers.

License Required for Certain RAID Configuration Options on Some Servers

Some Cisco UCS servers require a license for certain RAID configuration options. When Cisco UCS Manager associates a service profile containing this local disk policy with a server, Cisco UCS Manager verifies that the selected RAID option is properly licensed. If there are issues, Cisco UCS Manager displays a configuration error during the service profile association.

For RAID license information for a specific Cisco UCS server, see the *Hardware Installation Guide* for that server.

B420 M3 Server Does Not Support All Configuration Modes

The B420 M3 server does not support the following configuration modes in a local disk configuration policy:

- No RAID
- RAID 6 Striped Dual Parity

In addition, the B420 M3 does not support JBOD modes or operations.

Single-Disk RAID 0 Configurations Not Supported on Some Blade Servers

A single-disk RAID 0 configuration is not supported in the following blade servers:

- Cisco UCS B200 M1
- Cisco UCS B200 M2
- Cisco UCS B250 M1
- Cisco UCS B250 M2

Creating a Local Disk Configuration Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **Local Disk Config Policies** and choose **Create Local Disk Configuration Policy**.

Step 5 In the **Create Local Disk Configuration Policy** dialog box, complete the following fields:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Name	Description
Mode drop-down list	This can be one of the following local disk policy modes: • **No Local Storage** • **RAID 0 Striped** • **RAID 1 Mirrored** • **Any Configuration** • **No RAID** If you choose **No RAID** and you apply this policy to a server that already has an operating system with RAID storage configured, the system does not remove the disk contents. Therefore, there may be no visible differences on the server after you apply the **No RAID** mode. This can lead to a mismatch between the RAID configuration in the policy and the actual disk configuration shown in the **Inventory** > **Storage** tab for the server. To make sure that any previous RAID configuration information is removed from a disk, apply a scrub policy that removes all disk information after you apply the **No RAID** configuration mode. • **RAID 5 Striped Parity** • **RAID 6 Striped Dual Parity** • **RAID 10 Mirrored and Striped** • **RAID 50 Striped Parity and Striped** • **RAID 60 Striped Dual Parity and Striped** **Note** Some Cisco UCS servers require a license for certain RAID configuration options. When Cisco UCS Manager associates a service profile containing this local disk policy with a server, Cisco UCS Manager verifies that the selected RAID option is properly licensed. If there are issues, Cisco UCS Manager displays a configuration error during the service profile association. For RAID license information for a specific Cisco UCS server, see the *Hardware Installation Guide* for that server.

Name	Description
Protect Configuration check box	If checked, the server retains the configuration in the local disk configuration policy even if the server is disassociated from the service profile. **Caution** Protect Configuration becomes non-functional if one or more disks in the server are defective or faulty. This property is checked by default. When a service profile is disassociated from a server and a new service profile associated, the setting for the Protect Configuration property in the new service profile takes precedence and overwrites the setting in the previous service profile. With this option enabled, the data on the disk is protected even after the server is decommissioned and then recommissioned. Hence, reassociation of the server with a service profile fails. **Note** If you disassociate the server from a service profile with this option enabled and then associate it with a new service profile that includes a local disk configuration policy with different properties, the server returns a configuration mismatch error and the association fails.
FlexFlash State radio button	To enable or disable the FlexFlash controller on the SD card click the appropriate button. **Note** This parameter only applies to a server with an SD card module.
FlexFlash RAID Reporting State radio button	To enable or disable RAID reporting click the appropriate button. When RAID reporting is enabled, the RAID status is monitored and faults are enabled. **Note** If only one SD card is installed, the RAID state will be displayed as Disabled and the RAID health as NA even if RAID reporting is enabled.

Step 6 Click **OK**.

Changing a Local Disk Configuration Policy

This procedure describes how to change a local disk configuration policy from an associated service profile. You can also change a local disk configuration policy from the **Policies** node of the **Servers** tab.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the organization that includes the service profile with the local disk configuration policy you want to change.
If the system does not include multitenancy, expand the **root** node.

Step 4 Click the service profile that contains the local disk configuration policy you want to change.

Step 5 In the **Work** pane, click the **Storage** tab.

Step 6 In the **Actions** area, click **Change Local Disk Configuration Policy**.

Step 7 In the **Change Local Disk Configuration Policy** dialog box, choose one of the following options from the **Select the Local Disk Configuration Policy** drop-down list.

Option	Description
Use a Disk Policy	Select an existing local disk configuration policy from the list below this option. Cisco UCS Manager assigns this policy to the service profile.
Create a Local Disk Policy	Enables you to create a local disk configuration policy that can only be accessed by the selected service profile.
No Disk Policy	Selects the default local disk policy. **Note** If a UCS server is attached to the Cisco UCS Manager, selecting the No Disk Policy can erase and replace the RAID with individual RAID 0 disks if the default RAID configuration is not supported on the attached server.

Step 8 Click **OK**.

Step 9 (Optional) Expand the **Local Disk Configuration Policy** area to confirm that the change has been made.

Deleting a Local Disk Configuration Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies** > *Organization_Name*.

Step 3 Expand the **Local Disk Config Policies** node.

Step 4 Right-click the policy you want to delete and select **Delete**.

Step 5 If a confirmation dialog box displays, click **Yes**.

FlexFlash Support

Overview

Cisco UCS B-Series, C-Series M3 and higher, and S-Series M4 servers support internal Secure Digital (SD) memory cards. The SD cards are hosted by the Cisco Flexible Flash storage controller, a PCI-based controller which has two slots for SD cards. The cards contain a single partition called HV. When FlexFlash is enabled, Cisco UCS Manager displays the HV partition as a USB drive to both the BIOS and the host operating system.

You can populate one or both the SD card slots that are provided. If two SD cards are populated, you can use them in a mirrored mode.

Note Do not mix different capacity cards in the same server.

The SD cards can be used to store operating system boot images or other information. The following figure illustrates the SD card slots.

Figure 2: SD Card Slots

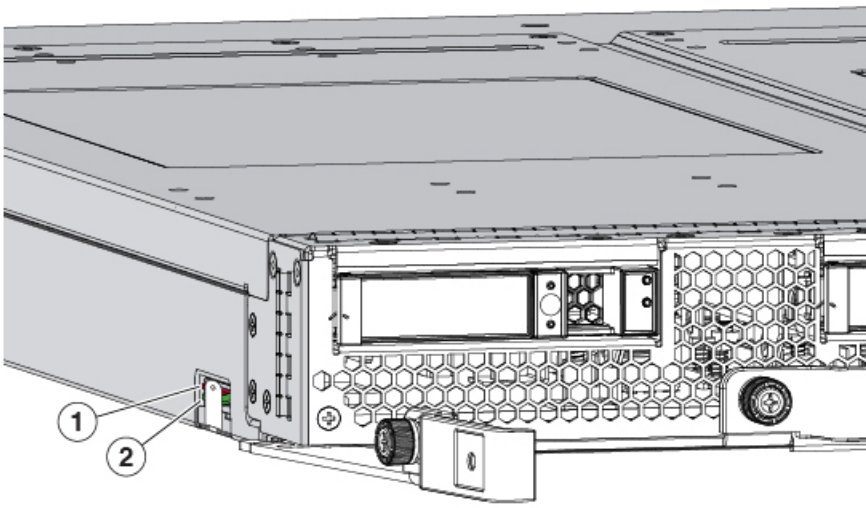

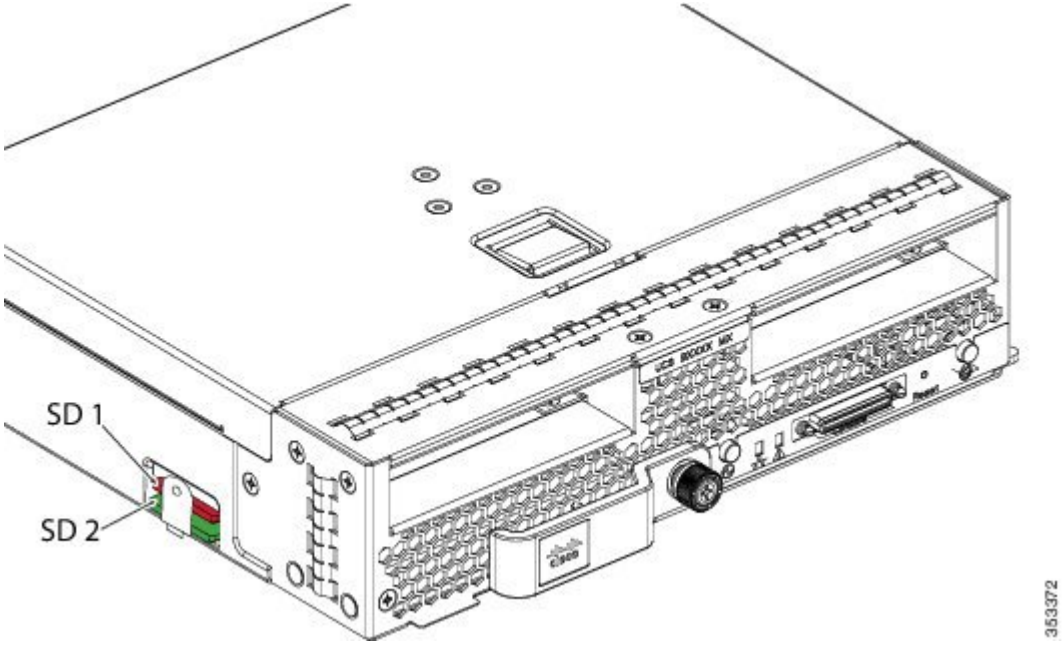

FlexFlash is disabled by default. You can enable FlexFlash in a local disk policy used in a service profile. When FlexFlash is enabled in a local disk policy, and the server is capable of supporting SD cards, the FlexFlash controller is enabled during service profile association. If a server is not capable of supporting SD cards or has an older CIMC version, a config failure message is displayed.

If you disable FlexFlash in a supported server, the Hypervisor or HV partition is immediately disconnected from the host. The FlexFlash controller will also be disabled as part of a related service profile disassociation.

The FlexFlash controller supports RAID-1 for dual SD cards. The FlexFlash scrub policy erases the HV partition in both cards, and brings the cards to a healthy RAID state.

You can configure new SD cards in a RAID pair and format them using one of the following methods:

- Format the SD cards. Formatting the SD Cards, on page 481provides detailed information.
- For an associated server, create a FlexFlash scrub policy and disassociate the service profile from the server. For an unassociated server, create a FlexFlash scrub policy and reacknowledge the server after modifying the default scrub policy.

  The *Scrub Policy Settings* section in the *Cisco UCS Manager Server Management Guide* provides more details about the usage of the scrub policy.

Note Disable the scrub policy as soon as the pairing is complete.

To boot from the HV partition, the SD card must be present in the boot policy used in the service profile.

FlexFlash Firmware Management

The FlexFlash controller firmware is bundled as part of the CIMC image. When you upgrade the CIMC, if a newer firmware version is available for the FlexFlash controller, the controller can no longer be managed, and the FlexFlash inventory displays the **Controller State** as **Waiting For User Action** and the **Controller Health** as **Old Firmware Running**. To upgrade the FlexFlash controller firmware, you need to perform a

board controller update. For more information, see the appropriate *Cisco UCS B-Series Firmware Management Guide*, available at the following URL: http://www.cisco.com/en/US/products/ps10281/products_installation_and_configuration_guides_list.html.

Limitations for the Cisco Flexible Flash Storage Controller:

- The Cisco Flexible Flash storage controller only supports 16 GB, 32 GB, and 64 GB SD cards.

Note 16 GB and 32 GB cards are supported only on the B200-M3 blade servers, and the 64 GB SD cards are supported only on the B200-M4 blade servers.

- We do not recommend using an SD card from a rack server in a blade server, or using an SD card from a blade server in a rack server. Switching SD cards between server types might result in data loss from the SD card.
- Some Cisco UCS C-Series rack-mount servers have SD cards with four partitions: HV, HUU, SCU, and Drivers. Only the HV partition is visible in Cisco UCS Manager. You can migrate a four-partition SD card to a single HV partition card with a FlexFlash scrub policy.
- The FlexFlash controller does not support RAID-1 sync (mirror rebuild). If the SD cards are in a degraded RAID state, or if any metadata errors are reported by the controller, you must run the FlexFlash scrub policy to pair the cards for RAID. For more information about the FlexFlash scrub policy, see Server-Related Policies. The following conditions might result in degraded RAID or metadata errors:
  - Inserting a new or used SD card in one slot, when the server already has an SD card populated in the second slot.
  - Inserting two SD cards from different servers.
- The server firmware version must be at 2.2(1a) or higher.

FlexFlash FX3S Support

Beginning with Release 2.2(3), Cisco UCS Manager allows additional FlexFlash support with the FX3S controller. The FX3S controller is present on the following servers:

- Cisco UCS B200 M4 blade server
- Cisco UCS C220 M4 rack server
- Cisco UCS C240 M4 rack server

FlexFlash operations with the FX3S control are similar to those with the Cisco Flexible Flash storage controller. FlexFlash is disabled by default, and is enabled using a local disk policy. You can also reset the controller, format the SD cards, and enable automatic synchronization of your paired SD cards.

The SD cards for the FX3S controller contain a single partition called Hypervisor.

Limitations for the Cisco FX3S Controller:

- The FX3S controller supports only 32 GB and 64 GB SD cards. 16 GB cards are not supported.

- We do not recommend using an SD card from a rack server in a blade server, or using an SD card from a blade server in a rack server. Switching SD cards between server types might result in data loss from the SD card.
- The server firmware version must be at 2.2(3a) or higher.

Enabling FlexFlash SD Card Support

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Expand **Local Disk Config Policies** and choose the local disk config policy for which you want to enable FlexFlash support.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 In the **FlexFlash State** field, click the **Enable** radio button.

Step 7 In the **FlexFlash RAID Reporting State** field, click the **Enable** radio button.

Step 8 Click **Save Changes**.

Enabling Auto-Sync

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > ***Chassis Number*** > **Servers**.

Step 3 Click the server for which you want to enable auto-sync.

Step 4 In the **Work** pane, click the **Inventory** tab.

Step 5 Click the **Storage** subtab.

Step 6 In the **Actions** area, click **Enable Auto-sync**.

Step 7 In the **Enable Auto-sync** dialog box, choose the **Admin Slot Number** for the SD card that you want to use as the primary.

Step 8 Click **OK**.

Formatting the SD Cards

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
Step 3 Click the server for which you want to format the SD cards.
Step 4 In the **Work** pane, click the **Inventory** tab.
Step 5 Click the **Storage** subtab.
Step 6 In the **Actions** area, click **Format SD Cards**.
Step 7 Click **Yes** to format the SD cards.

Resetting the FlexFlash Controller

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
Step 3 Click the server for which you want to reset the FlexFlash controller.
Step 4 In the **Work** pane, click the **Inventory** tab.
Step 5 Click the **Storage** subtab.
Step 6 In the **Actions** area, click **Reset FlexFlash Controller**.
Step 7 Click **Yes** to reset the FlexFlash controller.

Configuring Scrub Policies

Scrub Policy Settings

This policy determines what happens to local data and to the BIOS settings on a server during the discovery process, when the server is re-acknowledged, or when the server is disassociated from a service profile.

Note Local disk scrub policies only apply to hard drives that are managed by Cisco UCS Manager and do not apply to other devices such as USB drives.

Depending upon how you configure a scrub policy, the following can occur at those times:

Disk scrub

One of the following occurs to the data on any local drives on disassociation:

- If enabled, destroys all data on any local drives.
- If disabled, preserves all data on any local drives, including local storage configuration.

BIOS Settings Scrub

One of the following occurs to the BIOS settings when a service profile containing the scrub policy is disassociated from a server:

- If enabled, erases all BIOS settings for the server and resets them to the BIOS defaults for that server type and vendor.
- If disabled, preserves the existing BIOS settings on the server.

FlexFlash Scrub

FlexFlash Scrub enables you to pair new or degraded SD cards, resolve FlexFlash metadata configuration failures, and migrate older SD cards with 4 partitions to single partition SD cards. One of the following occurs to the SD card when a service profile containing the scrub policy is disassociated from a server, or when the server is reacknowledged:

- If enabled, the HV partition on the SD card is formatted using the PNUOS formatting utility. If two SD cards are present, the cards are RAID-1 paired, and the HV partitions in both cards are marked as valid. The card in slot 1 is marked as primary, and the card in slot 2 is marked as secondary.
- If disabled, preserves the existing SD card settings.

Note

- Because the FlexFlash scrub erases the HV partition on the SD cards, we recommend that you take a full backup of the SD card(s) using your preferred host operating system utilities before performing the FlexFlash Scrub.
- To resolve metadata config failures in a service profile, you need to disable FlexFlash in the local disk config policy before you run the FlexFlash scrub, then enable FlexFlash after the server is reacknowledged.
- Disable the scrub policy as soon as the pairing is complete or the metadata failures are resolved.
- FlexFlash scrub is not supported for Cisco UCS S3260 Storage Server.

Creating a Scrub Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **Scrub Policies** and select **Create Scrub Policy**.

Step 5 In the **Create Scrub Policy** wizard, complete the following fields:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Disk Scrub field	If this field is set to **Yes**, when a service profile containing this scrub policy is disassociated from a server, all data on the server local drives is completely erased. If this field is set to **No**, the data on the local drives is preserved, including all local storage configuration.
BIOS Settings Scrub field	If the field is set to **Yes**, when a service profile containing this scrub policy is disassociated from a server, the BIOS settings for that server are erased and reset to the defaults for that server type and vendor. If this field is set to **No**, the BIOS settings are preserved.
FlexFlash Scrub field	If the field is set to **Yes**, the HV partition on the SD card is formatted using the PNUOS formatting utility when the server is reacknowledged. If this field is set to **No**, the SD card is preserved.

Step 6 Click **OK**.

Deleting a Scrub Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Policies** > *Organization_Name*.
Step 3 Expand the **Scrub Policies** node.
Step 4 Right-click the policy you want to delete and select **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Configuring DIMM Error Management

DIMM Correctable Error Handling

In Cisco UCS Manager, when a DIMM encounters a significant correctable error in a given predefined window, it is stated as degraded and considered as a non-functional device.

The DIMM correctable error handling feature enables you to reset all the correctable and uncorrectable memory errors on all the DIMMs in a server. When you reset the error configuration, the error count of a given DIMM is cleared, the status changes to operable, and it resets the sensor state of the given DIMM.

Resetting Memory Errors

Use this procedure to reset all correctable and uncorrectable memory errors encountered by Cisco UCS Manager and the baseboard management controller (BMC).

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
Step 3 Right-click on the server for which you want to reset the error configuration, and select **Reset All Memory Errors**. You can also select **Reset All Memory Errors** from the **Actions** area.
Step 4 If the Cisco UCS Manager GUI displays a confirmation dialog box, click **Yes**.

DIMM Blacklisting

In Cisco UCS Manager, the state of the Dual In-line Memory Module (DIMM) is based on SEL event records. When the BIOS encounters a noncorrectable memory error during memory test execution, the DIMM is marked as faulty. A faulty DIMM is a considered a nonfunctional device.

If you enable DIMM blacklisting, Cisco UCS Manager monitors the memory test execution messages and blacklists any DIMMs that encounter memory errors in the DIMM SPD data. To allow the host to map out any DIMMs that encounter uncorrectable ECC errors.

Enabling DIMM Blacklisting

The memory policy is a global policy that you can apply to existing servers on a Cisco UCS domain and also to the servers that are added after you set the memory policy.

Note

- This feature is supported both on the Cisco UCS B-Series blade servers and UCS C-Series rack servers.

  

  Note Cisco UCS C-Series 420 M3 rack server do not support this feature.

- This global policy cannot be added to a service profile.

Before You Begin

- For Cisco B-Series blade server, the server firmware must be at Release 2.2(1) or a later release.
- For Cisco C-Series rack server, the server firmware must be at Release 2.2(3).
- You must be logged in with one of the following privileges:
  - Admin
  - Server policy
  - Server profile server policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to enable the blacklisting.
If the system does not include multitenancy, expand the **root** node.

Step 4 Expand **Memory Policy** and choose **default**.

Step 5 In the **Blacklisting** area, click the **Enabled** radio button.

The DIMM blacklisting is enabled for the domain level policy and these changes apply to all the servers on that particular domain.

Note If the Cisco IMC of a server does not support DIMM blacklisting, an information level fault is generated.

Configuring Serial over LAN Policies

Serial over LAN Policy Overview

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.

Creating a Serial over LAN Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **Serial over LAN Policies** and select **Create Serial over LAN Policy**.

Step 5 In the **Create Serial over LAN Policy** wizard, complete the following fields:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Name	Description
Owner field	This can be one of the following: • **Local—**This policy is available only to service profiles and service profile templates in this Cisco UCS domain. • **Pending Global—**Control of this policy is being transferred to Cisco UCS Central. Once the transfer is complete, this policy will be available to all Cisco UCS domains registered with Cisco UCS Central. • **Global—**This policy is managed by Cisco UCS Central. Any changes to this policy must be made through Cisco UCS Central.
Serial over LAN State field	This can be one of the following: • **Disable—**Serial over LAN access is blocked. • **Enable—**Serial over LAN access is permitted.
Speed drop-down list	This can be one of the following: • **9600** • **19200** • **38400** • **57600** • **115200**

Step 6 Click **OK**.

Deleting a Serial over LAN Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Policies** > *Organization_Name*.
Step 3 Expand the **Serial over LAN Policies** node.
Step 4 Right-click the policy you want to delete and select **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Configuring Server Autoconfiguration Policies

Server Autoconfiguration Policy Overview

Cisco UCS Manager uses this policy to determine how to configure a new server. If you create a server autoconfiguration policy, the following occurs when a new server starts:

1. The qualification in the server autoconfiguration policy is executed against the server.
2. If the server meets the required qualifications, the server is associated with a service profile created from the service profile template configured in the server autoconfiguration policy. The name of that service profile is based on the name given to the server by Cisco UCS Manager.
3. The service profile is assigned to the organization configured in the server autoconfiguration policy.

Creating an Autoconfiguration Policy

Before You Begin

This policy requires that one or more of the following resources already exist in the system:

- Server pool policy qualifications
- Service profile template
- Organizations, if a system implements multitenancy

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Click the **Equipment** node.

Step 3 In the **Work** pane, click the **Policies** tab.

Step 4 Click the **Autoconfig Policies** subtab.

Step 5 On the icon bar to the right of the table, click +.
If the + icon is disabled, click an entry in the table to enable it.

Step 6 In the **Create Autoconfiguration Policy** dialog box, complete the following fields:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Name	Description
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Qualification drop-down list	The server pool policy qualification associated with this auto-configuration policy. If a new server is discovered that matches the criteria specified in the server pool policy qualification, Cisco UCS automatically creates a service profile based on the service profile template selected in the **Service Profile Template Name** drop-down list and associates the newly created service profile with the server.
Org drop-down list	The organization associated with this autoconfiguration policy. If Cisco UCS automatically creates a service profile to associate with a server, it places the service profile under the organization selected in this field.
Service Profile Template Name drop-down list	The service profile template associated with this policy.

Step 7 Click **OK**.

Deleting an Autoconfiguration Policy

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Click the **Equipment** node.
Step 3 In the **Work** pane, click the **Policies** tab.
Step 4 Click the **Autoconfig Policies** subtab.
Step 5 Right-click the autoconfiguration policy that you want to delete and choose **Delete**.
Step 6 If a confirmation dialog box displays, click **Yes**.

Configuring Server Discovery Policies

Server Discovery Policy Overview

The server discovery policy determines how the UCS Manager reacts when you add a new UCS Blade Server and UCS Mini. 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.

If you create a server discovery policy, the following occurs when a new server starts:

1. The server discovery policy qualification is executed against the server.
2. If the server meets the required qualifications, Cisco UCS Manager applies the following to the server:
   - Depending on the option that you select for the action, UCS Manager discovers the new server immediately, or waits for a user acknowledgment of the new server
   - Applies the scrub policy to the server

If automatic deep discovery is triggered by any hardware insertion, removal, or replacement, the following occurs:

1. The server is moved to a “pending activities” list.
2. A critical hardware mismatch fault is raised on the server, indicating that UCSM has detected a hardware mismatch.
3. User must explicitly acknowledge the server to trigger the deep discovery.

Important In Cisco UCS Manager Release 2.2 (4), blade servers do not support drives with a block size of 4K, but rack-mount servers support such drives. If a drive with a block size of 4K is inserted into a blade server, discovery fails and the following error message appears:

```
Unable to get Scsi Device Information from the system
```

If this error occurs, do the following:

1. Remove the 4K drive.
2. Reacknowledge the server.

Reacknowledging the server causes the server to reboot and results in loss of service.

Creating a Server Discovery Policy

Before You Begin

If you plan to associate this policy with a server pool, create server pool policy qualifications.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Click the **Equipment** node.
Step 3 In the **Work** pane, click the **Policies** tab.
Step 4 Click the **Server Discovery Policies** subtab.
Step 5 Click the + icon on the table icon bar to open the **Create Server Discovery Policy** dialog box.
Step 6 In the **Description** field, enter a description for the discovery policy.
Step 7 In the **Action** field, select one of the following options:

- **Immediate—**Cisco UCS Manager attempts to discover new servers automatically
- **User Acknowledged—**Cisco UCS Manager waits until the user tells it to search for new servers

Step 8 (Optional) To associate this policy with a server pool, select server pool policy qualifications from the **Qualification** drop-down list.
Step 9 (Optional) To include a scrub policy, select a policy from the **Scrub Policy** drop-down list.
Step 10 Click **OK**.

What to Do Next

Include the server discovery policy in a service profile and/or template.

Deleting a Server Discovery Policy

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Click the **Equipment** node.
Step 3 In the **Work** pane, click the **Policies** tab.
Step 4 Click the **Server Discovery Policies** subtab.
Step 5 Right-click the server discover policy that you want to delete and choose **Delete**.
Step 6 If a confirmation dialog box displays, click **Yes**.

Configuring Server Inheritance Policies

Server Inheritance Policy Overview

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.

Creating a Server Inheritance Policy

A blade server or rack-mount server with a VIC adapter, such as the Cisco UCS M81KR Virtual Interface Card, does not have server identity values burned into the server hardware at manufacture. As a result, the identity of the adapter must be derived from default pools. If the default pools do not include sufficient entries for one to be assigned to the server, service profile association fails with a configuration error.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Click the **Equipment** node.

Step 3 In the **Work** pane, click the **Policies** tab.

Step 4 Click the **Server Inheritance Policies** subtab.

Step 5 On the icon bar at the bottom of the table, click **+ Add**.
If **+ Add** is disabled, click an entry in the table to enable it.

Step 6 In the **Create Server Inheritance Policy** dialog box, complete the following fields:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Name	Description
Qualification drop-down list	To associate this policy with one or more specific server pools, choose the server pool qualification policy that identifies these pools.
Org drop-down list	If you want to associate an organization with this policy, or if you want to change the current association, choose the organization from the drop-down list.

Step 7 Click **OK**.

Deleting a Server Inheritance Policy

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Click the **Equipment** node.
Step 3 In the **Work** pane, click the **Policies** tab.
Step 4 Click the **Server Inheritance Policies** subtab.
Step 5 Right-click the server inheritance policy that you want to delete and choose **Delete**.
Step 6 If a confirmation dialog box displays, click **Yes**.

Configuring Server Pool Policies

Server Pool Policy Overview

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.

Creating a Server Pool Policy

Before You Begin

This policy requires that one or more of the following resources already exist in the system:

- A minimum of one server pool

- Server pool policy qualifications, if you choose to have servers automatically added to pools

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **Server Pool Policies** and select **Create Server Pool Policy**.

Step 5 In the **Create Server Pool Policy** dialog box, complete the following fields:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Target Pool drop-down list	If you want to associate this policy with a server pool, select that pool from the drop-down list.
Qualification drop-down list	To associate this policy with one or more specific server pools, choose the server pool qualification policy that identifies these pools.

Step 6 Click **OK**.

Deleting a Server Pool Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Policies** > *Organization_Name*.
Step 3 Expand the **Server Pool Policies** node.
Step 4 Right-click the policy you want to delete and select **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Configuring Server Pool Policy Qualifications

Server Pool Policy Qualification Overview

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.

You can use the server pool policy qualifications to qualify servers according to the following criteria:

- Adapter type
- Chassis location
- Memory type and configuration
- Power group
- CPU cores, type, and configuration
- Storage configuration and capacity
- Server model

Depending upon the implementation, you might need to configure several policies with server pool policy qualifications including the following:

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

- Server pool policy

Creating Server Pool Policy Qualifications

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the **Server Pool Policy Qualifications** node and select **Create Server Pool Policy Qualification**.

Step 5 In the **Create Server Pool Policy Qualification** dialog box, enter a unique name and description for the policy.

Step 6 (Optional) To use this policy to qualify servers according to their adapter configuration, do the following:

a) Click **Create Adapter Qualifications**.

b) In the **Create Adapter Qualifications** dialog box, complete the following fields:

Name	Description
Type drop-down list	The adapter type. Once you save the adapter qualification, this type cannot be changed.
PID field	A regular expression that the adapter PID must match.
Maximum Capacity field	The maximum capacity for the selected type. To specify a capacity, choose **select** and enter the desired maximum capacity. You can enter an integer between 1 and 65535.

c) Click **OK**.

Step 7 (Optional) To use this policy to qualify servers according to the chassis in which they physically reside, do the following:

a) Click **Create Chassis/Server Qualifications**.

b) In the **Chassis Qualifications** area of the **Create Chassis and Server Qualifications** dialog box, complete the following fields to specify the range of chassis you want to use:

- **First Chassis ID** field—The first chassis ID from which server pools associated with this policy can draw.
- **Number of Chassis** field—The total number of chassis to include in the pool, starting with the chassis identified in the **First Chassis ID** field.

Example:

For example, if you want to use chassis 5, 6, 7, and 8, enter 5 in the **First Chassis ID** field and 4 in the **Number of Chassis** field. If you want to use only chassis 3, enter 3 in the **First Chassis ID** field and 1 in the **Number of Chassis** field.

Tip If you want to use chassis 5, 6, and 9, create a chassis/server qualification for the range 5-6 and another qualification for chassis 9. You can add as many chassis/server qualifications as needed.

c) Click **Finish**.

Step 8 (Optional) To use this policy to qualify servers according to both the chassis and slot in which they physically reside, do the following:

a) Click **Create Chassis/Server Qualifications**.

b) In the **Chassis Qualifications** area of the **Create Chassis and Server Qualifications** dialog box, complete the following fields to specify the range of chassis you want to use:

- **First Chassis ID** field—The first chassis ID from which server pools associated with this policy can draw.
- **Number of Chassis** field—The total number of chassis to include in the pool, starting with the chassis identified in the **First Chassis ID** field.

c) In the **Server Qualifications** table, click **Add**.

d) In the **Create Server Qualifications** dialog box, complete the following fields to specify the range of server locations you want to use:

- **First Slot ID** field—The first slot ID from which server pools associated with this policy can draw.
- **Number of Slots** field—The total number of slots from which server pools associated with this policy can draw.

e) Click **Finish Stage**.

f) To add another range of slots, click **Add** and repeat steps d and e.

g) When you have finished specifying the slot ranges, click **Finish**.

Step 9 (Optional) To use this policy to qualify servers according to their memory configuration, do the following:

a) Click **Create Memory Qualifications**.

b) In the **Create Memory Qualifications** dialog box, complete the following fields:

Name	Description
Clock field	The minimum clock speed required, in megahertz.
Latency field	The maximum latency allowed, in nanoseconds.
Min Cap field	The minimum memory capacity required, in megabytes.
Max Cap field	The maximum memory capacity allowed, in megabytes.
Width field	The minimum width of the data bus.
Units field	The unit of measure to associate with the value in the **Width** field.

c) Click **OK**.

Step 10 (Optional) To use this policy to qualify servers according to their CPU/Cores configuration, do the following:

a) Click **Create CPU/Cores Qualifications**.

b) In the **Create CPU/Cores Qualifications** dialog box, complete the following fields:

Name	Description
Processor Architecture drop-down list	The CPU architecture to which this policy applies.
PID field	A regular expression that the processor PID must match.
Min Number of Cores field	The minimum number of CPU cores required. To specify a capacity, choose **select** and enter an integer between 1 and 65535 in the associated text field.
Max Number of Cores field	The maximum number of CPU cores allowed. To specify a capacity, choose **select** and enter an integer between 1 and 65535 in the associated text field.
Min Number of Threads field	The minimum number of CPU threads required. To specify a capacity, choose **select** and enter an integer between 1 and 65535 in the associated text field.
Max Number of Threads field	The maximum number of CPU threads allowed. To specify a capacity, choose **select** and enter an integer between 1 and 65535 in the associated text field.
CPU Speed field	The minimum CPU speed required. To specify a capacity, choose **select** and enter the minimum CPU speed.
CPU Stepping field	The minimum CPU version required. To specify a capacity, choose **select** and enter the maximum CPU speed.

c) Click **OK**.

Step 11 (Optional) To use this policy to qualify servers according to their storage configuration and capacity, do the following:

a) Click **Create Storage Qualifications**.

b) In the **Create Storage Qualifications** dialog box, complete the following fields:

Name	Description
Diskless field	Whether the available storage must be diskless. This can be one of the following: • **Unspecified—**Either storage type is acceptable. • **Yes—**The storage must be diskless. • **No—**The storage cannot be diskless.
Number of Blocks field	The minimum number of blocks required. To specify a capacity, choose **select** and enter the number of blocks.
Block Size field	The minimum block size required, in bytes. To specify a capacity, choose **select** and enter the block size.
Min Cap field	The minimum storage capacity across all disks in the server, in megabytes. To specify a capacity, choose **select** and enter the minimum storage capacity.
Max Cap field	The maximum storage capacity allowed, in megabytes. To specify a capacity, choose **select** and enter the maximum storage capacity.
Per Disk Cap field	The minimum storage capacity per disk required, in gigabytes. To specify a capacity, choose **select** and enter the minimum capacity on each disk.
Units field	The number of units. To specify a capacity, choose **select** and enter the desired units.
Number of Flex Flash Cards field	The number of FlexFlash Cards. To specify a capacity, choose **select** and enter the desired units.
Disk Type field	The disk type. This can be one of the following: • **Unspecified—**Either disk type is acceptable. • **HDD—**The disk must be HDD. • **SSD—**The disk must be SSD (SATA or SAS).

c) Click **OK**.

Step 12 To use this policy to qualify servers according to the model of the server, do the following:

a) Click **Create Server PID Qualifications**.

b) In the **Create Server PID Qualifications** dialog box, select the PID of the server model from the **PID** drop-down list.
c) Click **OK**.

Step 13 (Optional) To use this policy to qualify servers according to power group, do the following:
a) Click **Create Power Group Qualifications**.
b) In the **Create Power Group Qualifications** dialog box, choose a power group from the **Power Group** drop-down list.
c) Click **OK**.

Step 14 (Optional) To use this policy to qualify the rack-mount servers that can be added to the associated server pool, do the following:
a) Click **Create Rack Qualifications**.
b) In the **Create Rack Qualifications** dialog box, complete the following fields:

Name	Description
First Slot ID field	The first rack-mount server slot ID from which server pools associated with this policy can draw.
Number of Slots field	The total number of rack-mount server slots from which server pools associated with this policy can draw.

Step 15 Verify the qualifications in the table and correct if necessary.
Step 16 Click **OK**.

Deleting Server Pool Policy Qualifications

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Policies** > *Organization_Name*.
Step 3 Expand the **Server Pool Policy Qualifications** node.
Step 4 Right-click the policy qualifications you want to delete and select **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Deleting Qualifications from Server Pool Policy Qualifications

Use this procedure to modify Server Pool Policy Qualifications by deleting one or more sets of qualifications.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies** > *Organization_Name*.

Step 3 Expand the **Server Pool Policy Qualifications** node.

Step 4 Choose the policy you want to modify.

Step 5 In the **Work** pane, choose the **Qualifications** tab.

Step 6 To delete a set of qualifications:

a) In the table, choose the row that represents the set of qualifications.

b) Right-click the row and select **Delete**.

Step 7 Click **Save Changes**.

Configuring vNIC/vHBA Placement Policies

vNIC/vHBA Placement Policies

vNIC/vHBA placement policies are used to determine the following:

- How the virtual network interface connections (vCons) are mapped to the physical adapters on a server.
- What types of vNICs or vHBAs can be assigned to each vCon.

Each vNIC/vHBA placement policy contains four vCons that are virtual representations of the physical adapters. When a vNIC/vHBA placement policy is assigned to a service profile, and the service profile is associated with a server, the vCons in the vNIC/vHBA placement policy are assigned to the physical adapters and the vNICs and vHBAs are assigned to those vCons.

For blade or rack servers that contain one adapter, Cisco UCS assigns all vCons to that adapter. For servers that contain four adapters, Cisco UCS assigns vCon1 to Adapter1, vCon2 to Adapter2, vCon3 to Adapter3, and vCon4 to Adapter4.

For blade or rack servers that contain two or three adapters, Cisco UCS assigns the vCons based on the type of server and the selected virtual slot mapping scheme, which can be **Round Robin** or **Linear Ordered**. For details about the available mapping schemes, see vCon to Adapter Placement, on page 502.

After Cisco UCS assigns the vCons, it assigns the vNICs and vHBAs based on the **Selection Preference** for each vCon. This can be one of the following:

Note You can specify the PCI order for the vHBA; however, the desired order works within a class of devices, such as vNICs or vHBAs and not across them. Within an adapter, vNICs are always placed ahead of the vHBAs.

- **All—**All configured vNICs and vHBAs can be assigned to the vCon, whether they are explicitly assigned to it, unassigned, or dynamic. This is the default.

- **Assigned Only—**vNICs and vHBAs must be explicitly assigned to the vCon. You can assign them explicitly through the service profile or the properties of the vNIC or vHBA.
- **Exclude Dynamic—**Dynamic vNICs and vHBAs cannot be assigned to the vCon. The vCon can be used for all static vNICs and vHBAs, whether they are unassigned or explicitly assigned to it.
- **Exclude Unassigned—**Unassigned vNICs and vHBAs cannot be assigned to the vCon. The vCon can be used for dynamic vNICs and vHBAs and for static vNICs and vHBAs that are explicitly assigned to it.
- **Exclude usNIC—**Cisco usNICs cannot be assigned to the vCon. The vCon can be used for all other configured vNICs and vHBAs, whether they are explicitly assigned to it, unassigned, or dynamic.

Note An SRIOV usNIC that is explicitly assigned to a vCon set to **Exclude usNIC** will remain assigned to that vCon.

If you do not include a vNIC/vHBA placement policy in the service profile, Cisco UCS Manager defaults to the **Round Robin** vCon mapping scheme and the **All** vNIC/vHBA selection preference, distributing the vNICs and vHBAs between the adapters based on the capabilities and relative capacities of each adapter.

vCon to Adapter Placement

Cisco UCS maps every vCon in a service profile to a physical adapter on the server. How that mapping occurs and how the vCons are assigned to a specific adapter in a server depends on the following:

- The type of server. N20-B6620-2 and N20-B6625-2 blade servers with two adapter cards use a different mapping scheme than other supported rack or blade servers.
- The number of adapters in the server.
- The setting of the virtual slot mapping scheme in the vNIC/vHBA placement policy, if applicable.

You must consider this placement when you configure the vNIC/vHBA selection preference to assign vNICs and vHBAs to vCons.

Note vCon to adapter placement is not dependent upon the PCIE slot number of the adapter. The adapter numbers used for the purpose of vCon placement are not the PCIE slot numbers of the adapters, but the ID assigned to them during server discovery.

vCon to Adapter Placement for N20-B6620-2 and N20-B6625-2 Blade Servers

In N20-B6620-2 and N20-B6625-2 blade servers, the two adapters are numbered left to right while vCons are numbered right to left. If one of these blade servers has a single adapter, Cisco UCS assigns all vCons to that adapter. If the server has two adapters, the vCon assignment depends upon the virtual slot mapping scheme:

- **Round Robin—**Cisco UCS assigns vCon2 and vCon4 to Adapter1 and vCon1 and vCon3 to Adapter2. This is the default.

- **Linear Ordered—**Cisco UCS assigns vCon3 and vCon4 to Adapter1 and vCon1 and vCon2 to Adapter2.

vCon to Adapter Placement for All Other Supported Servers

For all other servers supported by Cisco UCS in addition to the N20-B6620-2 and N20-B6625-2 blade servers, the vCon assignment depends on the number of adapters in the server and the virtual slot mapping scheme.

For blade or rack servers that contain one adapter, Cisco UCS assigns all vCons to that adapter. For servers that contain four adapters, Cisco UCS assigns vCon1 to Adapter1, vCon2 to Adapter2, vCon3 to Adapter3, and vCon4 to Adapter4.

For blade or rack servers that contain two or three adapters, Cisco UCS assigns the vCons based on the selected virtual slot mapping scheme: Round Robin or Linear Ordered.

Table 15: vCon to Adapter Placement Using the Round - Robin Mapping Scheme

Number of Adapters	vCon1 Assignment	vCon2 Assignment	vCon3 Assignment	vCon4 Assignment
1	Adapter1	Adapter1	Adapter1	Adapter1
2	Adapter1	Adapter2	Adapter1	Adapter2
3	Adapter1	Adapter2	Adapter3	Adapter2
4	Adapter1	Adapter2	Adapter3	Adapter4

Round Robin is the default mapping scheme.

Table 16: vCon to Adapter Placement Using the Linear Ordered Mapping Scheme

Number of Adapters	vCon1 Assignment	vCon2 Assignment	vCon3 Assignment	vCon4 Assignment
1	Adapter1	Adapter1	Adapter1	Adapter1
2	Adapter1	Adapter1	Adapter2	Adapter2
3	Adapter1	Adapter2	Adapter3	Adapter3
4	Adapter1	Adapter2	Adapter3	Adapter4

Note If you are using a vCon policy with two adapters in the Cisco UCS B440 M2 Blade Server, be aware of the following mapping.

- vCon 2 to adapter 1 maps first
- vCon 1 to adapter 2 maps second ZXA Q

vNIC/vHBA to vCon Assignment

Cisco UCS Manager provides two options for assigning vNICs and vHBAs to vCons through the vNIC/vHBA placement policy: explicit assignment and implicit assignment.

Explicit Assignment of vNICs and vHBAs

With explicit assignment, you specify the vCon and, therefore, the adapter to which a vNIC or vHBA is assigned. Use this assignment option when you need to determine how the vNICs and vHBAs are distributed between the adapters on a server.

To configure a vCon and the associated vNICs and vHBAs for explicit assignment, do the following:

- Set the vCon configuration to any of the available options. You can configure the vCons through a vNIC/vHBA placement policy or in the service profile associated with the server. If a vCon is configured for **All**, you can still explicitly assign a vNIC or vHBA to that vCon.
- Assign the vNICs and vHBAs to a vCon. You can make this assignment through the virtual host interface placement properties of the vNIC or vHBA or in the service profile associated with the server.

If you attempt to assign a vNIC or vHBA to a vCon that is not configured for that type of vNIC or vHBA, Cisco UCS Manager displays a message advising you of the configuration error.

During service profile association, Cisco UCS Manager validates the configured placement of the vNICs and vHBAs against the number and capabilities of the physical adapters in the server before assigning the vNICs and vHBAs according to the configuration in the policy. Load distribution is based upon the explicit assignments to the vCons and adapters configured in this policy.

If the adapters do not support the assignment of one or more vNICs or vHBAs, Cisco UCS Manager raises a fault against the service profile.

Note You can specify the PCI order for the vHBA; however, the desired order works within a class of devices, such as vNICs or vHBAs and not across them. Within an adapter, vNICs are always placed ahead of the vHBAs.

Implicit Assignment of vNICs and vHBAs

With implicit assignment, Cisco UCS Manager determines the vCon and, therefore, the adapter to which a vNIC or vHBA is assigned according to the capability of the adapters and their relative capacity. Use this assignment option if the adapter to which a vNIC or vHBA is assigned is not important to your system configuration.

To configure a vCon for implicit assignment, do the following:

- Set the vCon configuration to **All**, **Exclude Dynamic**, or **Exclude Unassigned**. You can configure the vCons through a vNIC/vHBA placement policy or in the service profile associated with the server.
- Do not set the vCon configuration to **Assigned Only**. Implicit assignment cannot be performed with this setting.
- Do not assign any vNICs or vHBAs to a vCon.

During service profile association, Cisco UCS Manager verifies the number and capabilities of the physical adapters in the server and assigns the vNICs and vHBAs accordingly. Load distribution is based upon the capabilities of the adapters, and placement of the vNICs and vHBAs is performed according to the actual order determined by the system. For example, if one adapter can accommodate more vNICs than another, that adapter is assigned more vNICs.

If the adapters cannot support the number of vNICs and vHBAs configured for that server, Cisco UCS Manager raises a fault against the service profile.

Implicit Assignment of vNICs in a Dual Adapter Environment

When you use implicit vNIC assignment for a dual slot server with an adapter card in each slot, Cisco UCS Manager typically assigns the vNICs/vHBAs as follows:

- If the server has the same adapter in both slots, Cisco UCS Manager assigns half the vNICs and half the vHBAs to each adapter.
- If the server has one non-VIC adapter and one VIC adapter, Cisco UCS Manager assigns two vNICs and two vHBAs to the non-VIC adapter and the remaining vNICs and vHBAs to the VIC adapter.
- If the server has two different VIC adapters, Cisco UCS Manager assigns the vNICs and vHBAs proportionally, based on the relative capabilities of the two adapters.

The following examples show how Cisco UCS Manager would typically assign the vNICs and vHBAs with different combinations of supported adapter cards:

- If you want to configure four vNICs and the server contains two Cisco UCS M51KR-B Broadcom BCM57711 adapters (with two vNICs each), Cisco UCS Manager assigns two vNICs to each adapter.
- If you want to configure 50 vNICs and the server contains a Cisco UCS CNA M72KR-E adapter (2 vNICs) and a Cisco UCS M81KR Virtual Interface Card adapter (128 vNICs), Cisco UCS Manager assigns two vNICs to the Cisco UCS CNA M72KR-E adapter and 48 vNICs to the Cisco UCS M81KR Virtual Interface Card adapter.
- If you want to configure 150 vNICs and the server contains a Cisco UCS M81KR Virtual Interface Card adapter (128 vNICs) and a Cisco UCS VIC-1240 Virtual Interface Card adapter (256 vNICs), Cisco UCS Manager assigns 50 vNICs to the Cisco UCS M81KR Virtual Interface Card adapter and 100 vNICs to the Cisco UCS VIC-1240 Virtual Interface Card adapter.

Note Exceptions to this implicit assignment occur if you configure the vNICs for fabric failover and if you configure dynamic vNICs for the server.

For a configuration that includes vNIC fabric failover where one adapter does not support vNIC failover, Cisco UCS Manager implicitly assigns all vNICs that have fabric failover enabled to the adapter that supports them. If the configuration includes only vNICs that are configured for fabric failover, no vNICs are implicitly assigned to the adapter that does not support them. If some vNICs are configured for fabric failover and some are not, Cisco UCS Manager assigns all failover vNICs to the adapter that supports them and a minimum of one nonfailover vNIC to the adapter that does not support them, according to the ratio above.

For a configuration that includes dynamic vNICs, the same implicit assignment would occur. Cisco UCS Manager assigns all dynamic vNICs to the adapter that supports them. However, with a combination of dynamic vNICs and static vNICs, at least one static vNIC is assigned to the adapter that does not support dynamic vNICs.

Creating a vNIC/vHBA Placement Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **vNIC/vHBA Placement Policies** and choose **Create Placement Policy**.

Step 5 In the **Create Placement Policy** dialog box, do the following:

a) Complete the following fields:

Name	Description
Name field	The name for this placement policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Name	Description
Virtual Slot Mapping Scheme field	Cisco UCS assigns virtual network interface connections (vCons) to the PCIe adapter cards in the server. Each vCon is a virtual representation of a physical adapter that can be assigned vNICs and vHBAs. For blade or rack servers that contain one adapter, Cisco UCS assigns all vCons to that adapter. For servers that contain four adapters, Cisco UCS assigns vCon1 to Adapter1, vCon2 to Adapter2, vCon3 to Adapter3, and vCon4 to Adapter4. For blade or rack servers that contain two or three adapters, Cisco UCS assigns the vCons based on the selected virtual slot mapping scheme. This can be one of the following: • **Round Robin—** In a server with two adapter cards, Cisco UCS assigns vCon1 and vCon3 to Adapter1, then assigns vCon2 and vCon4 to Adapter2. In a server with three adapter cards, Cisco UCS assigns vCon1 to Adapter1, vCon2 and vCon4 to Adapter2, and vCon3 to Adapter3. This is the default scheme. • **Linear Ordered—** In a server with two adapter cards, Cisco UCS assigns vCon1 and vCon2 to Adapter1, then assigns vCon3 and vCon4 to Adapter2. In a server with three adapter cards, Cisco UCS assigns vCon1 to Adapter1 and vCon2 to Adapter2, then assigns vCon3 and vCon4 to Adapter3. **Note** In N20-B6620-2 and N20-B6625-2 blade servers, the two adapters are numbered left to right while vCons are numbered right to left. If one of these blade servers has a single adapter, Cisco UCS assigns all vCons to that adapter. If the server has two adapters, the vCon assignment depends upon the virtual slot mapping scheme: • **Round Robin—**Cisco UCS assigns vCon2 and vCon4 to Adapter1 and vCon1 and vCon3 to Adapter2. This is the default. • **Linear Ordered—**Cisco UCS assigns vCon3 and vCon4 to Adapter1 and vCon1 and vCon2 to Adapter2. After Cisco UCS assigns the vCons, it assigns the vNICs and vHBAs based on the **Selection Preference** for each vCon.

b) In the **Selection Preference** column for each **Virtual Slot**, choose one of the following from the drop-down list:

- **All—**All configured vNICs and vHBAs can be assigned to the vCon, whether they are explicitly assigned to it, unassigned, or dynamic. This is the default.
- **Assigned Only—**vNICs and vHBAs must be explicitly assigned to the vCon. You can assign them explicitly through the service profile or the properties of the vNIC or vHBA.
- **Exclude Dynamic—**Dynamic vNICs and vHBAs cannot be assigned to the vCon. The vCon can be used for all static vNICs and vHBAs, whether they are unassigned or explicitly assigned to it.
- **Exclude Unassigned—**Unassigned vNICs and vHBAs cannot be assigned to the vCon. The vCon can be used for dynamic vNICs and vHBAs and for static vNICs and vHBAs that are explicitly assigned to it.
- **Exclude usNIC—**Cisco usNICs cannot be assigned to the vCon. The vCon can be used for all other configured vNICs and vHBAs, whether they are explicitly assigned to it, unassigned, or dynamic.

  Note An SRIOV usNIC that is explicitly assigned to a vCon set to **Exclude usNIC** will remain assigned to that vCon.

c) Click **OK**.

Deleting a vNIC/vHBA Placement Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Policies** > *Organization_Name*.
Step 3 Expand the **vNIC/vHBA Placement Policies** node.
Step 4 Right-click the policy you want to delete and choose **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Explicitly Assigning a vNIC to a vCon

Before You Begin

Configure the vCons through a vNIC/vHBA placement policy or in the service profile with one of the following values:

- **Assigned Only**
- **Exclude Dynamic**
- **Exclude Unassigned**

If a vCon is configured for **All**, you can explicitly assign a vNIC or vHBA to that vCon. However, there is less control with this configuration.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization which contains the service profile whose vNICs you want to explicitly assign to a vCon.
If the system does not include multitenancy, expand the **root** node.

Step 4 Expand ***Service_Profile_Name*** > **vNICs**.

Step 5 Click on the vNIC that you want to explicitly assign to a vCon.

Step 6 In the **Work** pane, click the **General** tab.

Step 7 In the **Virtual Host Interface Placement** section, complete the following fields:

Name	Description
Desired Placement drop-down list	The user-specified virtual network interface connection (vCon) placement for the vNIC. This can be one of the following: • **Any—**Allows Cisco UCS Manager to determine the vCon to which the vNIC is assigned. • **1—**Explicitly assigns the vNIC to vCon1. • **2—**Explicitly assigns the vNIC to vCon2. • **3—**Explicitly assigns the vNIC to vCon3. • **4—**Explicitly assigns the vNIC to vCon4.
Actual Assignment field	The actual vCon assignment of the vNIC on the server.

If you attempt to assign a vNIC to a vCon that is not configured for that type of vNIC, Cisco UCS Manager displays a message box to advise you of the configuration error. You must either assign the vNIC to another vCon or change the vCon configuration in the service profile.

Step 8 In the **Order** section, complete the following fields:

Name	Description
Desired Order field	The user-specified PCI order for the vNIC. Enter an integer between 0 and 128. You cannot create more than 128 vNICs for a server.
Actual Order field	The actual PCI order of the vNIC on the server.

Step 9 Click **Save Changes**.

Explicitly Assigning a vHBA to a vCon

Before You Begin

Configure the vCons through a vNIC/vHBA placement policy or in the service profile with one of the following values:

- **Assigned Only**
- **Exclude Dynamic**
- **Exclude Unassigned**

If a vCon is configured for **All**, you can explicitly assign a vNIC or vHBA to that vCon. However, there is less control with this configuration.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization which contains the service profile whose vHBAs you want to explicitly assign to a vCon.
If the system does not include multitenancy, expand the **root** node.

Step 4 Expand ***Service_Profile_Name*** > **vHBAs**.

Step 5 Click on the vHBA that you want to explicitly assign to a vCon.

Step 6 In the **Work** pane, click the **General** tab.

Step 7 In the **Virtual Host Interface Placement** section, complete the following fields:

Name	Description
Desired Placement field	The user-specified virtual network interface connection (vCon) placement for the vHBA. This can be one of the following: • **Any—**Allows Cisco UCS Manager to determine the vCon to which the vHBA is assigned. • **1—**Explicitly assigns the vHBA to vCon1. • **2—**Explicitly assigns the vHBA to vCon2. • **3—**Explicitly assigns the vHBA to vCon3. • **4—**Explicitly assigns the vHBA to vCon4.
Actual Assignment field	The actual vCon assignment of the vHBA on the server.

If you attempt to assign a vHBA to a vCon that is not configured for that type of vHBA, Cisco UCS Manager displays a message box to advise you of the configuration error. You must either assign the vHBA to another vCon or change the vCon configuration in the service profile.

Step 8 In the **Order** section, complete the following fields:

Name	Description
Desired Order field	The user-specified PCI order for the vHBA. Enter an integer between 0 and 128. You cannot create more than 128 vHBAs for a server.
Actual Order field	The actual PCI order of the vHBA on the server.

Step 9 Click **Save Changes**.

Placing Static vNICs Before Dynamic vNICs

For optimal performance, static vNICs and vHBAs should be placed before dynamic vNICs on the PCIe bus. Static vNICs refer to both static vNICs and vHBAs. Cisco UCS Manager Release 2.1 provides the following functionality regarding the order of static and dynamic vNICs:

- After upgrading to Cisco UCS Manager Release 2.1, if no change is made to existing service profiles (profiles that are defined in releases prior to Cisco UCS Manager Release 2.1), the vNIC order does not change.
- After an upgrade to Cisco UCS Manager Release 2.1, any vNIC-related change would reorder the vNIC map. As a result, all dynamic vNICs would be placed after the static vNICs.
- For newly created service profiles in Cisco UCS Manager Release 2.1, static vNICs are always ordered before dynamic vNICs.
- The above behavior is independent of the sequence of creating or deleting static or dynamic vNICs.
- For SRIOV-enabled service profiles, UCSM places the vNIC Physical Function(PF) before the corresponding Virtual Functions (VFs). This scheme guarantees that the VFs are placed close to the parent PF vNIC on the PCIe bus and BDFs are in successive incremental order for the VFs.

Example

Beginning Device Order in Cisco UCS Manager Release 2.0:

```
dyn-vNIC-1 1
dyn-vNIC-2 2
```

New Device Order in Cisco UCS Manager Release 2.0 (Add 2 static vNICs):

```
dyn-vNIC-1 1
dyn-vNIC-2 2
eth-vNIC-1 3
eth-vNIC-2 4
```

After upgrading to Cisco UCS Manager Release 2.1, (Before any vNIC-related change is made to the service profile.)

```
dyn-vNIC-1 1
dyn-vNIC-2 2
eth-vNIC-1 3
eth-vNIC-2 4
```

New Device Order in Cisco UCS Manager Release 2.1 (Add 2 dynamic vNICs by changing the policy count from 2 to 4.)

```
dyn-vNIC-1 3
dyn-vNIC-2 4
eth-vNIC-1 1
eth-vNIC-2 2
dyn-vNIC-3 5
dyn-vNIC-4 6
```

Dynamic vNICs as Multifunction PCIe Devices

Cisco UCS Manager Version 2.1 provisions static vNICs as 0-function devices (new BUS for every static vNIC). Multifunction dynamic vNICs are placed from the new Bus-slot after the last static vNIC/vHBA.

Note Cisco UCS Manager Version 2.1 supports the new StaticZero mode.

Table 17: Version Compatibility

Cisco UCS Manager		
Version 1.4 **Scheme: ZeroFunction**	**Version 2.0** **Scheme: ZeroFunction / MultiFunction**	**Version 2.1** **Scheme: ZeroFunction / MultiFunction / StaticZero**
Static and Dynamic vNICs are all on Bus [0-57], Function [0] < ZeroFunction Mode >	Static vNICs and Dynamic vNICs are on Bus [0-57], Function [0-7]. Bus 0, Function 0 Bus 0, Function 7 Bus 1, Function 0 < MultiFunction Mode >	Static vNICs or PFs will be on Bus [0-57], Function [0]. SRIOV: Corresponding VFs will be on the same Bus and Functions [1-255] No-SRIOV: Dynamic vNICs are on Bus [0-57], Function [0-7] < StaticZero Mode >
	Upgrade from Balboa will not renumber BDFs (remain in ZeroFunction mode) until Bus <= 57. Once devices exceed 58, switch to MultiFunction mode.	Upgrade from Balboa will not renumber BDFs (remain in ZeroFunction mode) until Bus <=57. Once devices exceed 58 or Platform specific maximum PCIe Bus number or change to SRIOV configuration, switch to StaticZero mode.

Cisco UCS Manager		
Version 1.4 **Scheme: ZeroFunction**	**Version 2.0** **Scheme: ZeroFunction / MultiFunction**	**Version 2.1** **Scheme: ZeroFunction / MultiFunction / StaticZero**
		Upgrade from Cisco UCS Manager Version 2.0 will not renumber BDFs (remain in ZeroFunction / MultiFunction mode). Once devices exceed 58 or Platfor specific maximum PCIe Bus number OR Change to SRIOV configuration, switch to StaticZero mode.

vNIC/vHBA Host Port Placement

After a vNIC/vHBA is assigned to a vCON, it can be placed on one of the host ports of specific adapters. You can either explicitly specify the host port for placement, or allow Cisco UCS Manager to automatically assign vNICs/vHBAs to host ports.

Note You can perform vNIC/vHBA host port placement on servers that support Cisco UCS VIC 1340 and VIC 1380 adapters.

The host port placement of the vNIC/vHBA determines the order of the vNIC/vHBA on the adapter. The vNICs/vHBAs placed on the first host port will be enumerated first, followed by the vNICs/vHBAs on the second host port.

Configuring Host Port Placement

You can configure host port placement for vNICs on servers that support Cisco UCS VIC 1340 and VIC 1380 adapters.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Select the service profile which is associated with the vNIC that you want to place on a host port.

Step 4 Expand ***Service_Profile_Name*** > **vNICs** > ***vNIC_Name***

Step 5 In the **Work** pane, click the **General** tab.

Step 6 In the **Host Port** section in the **Properties** area, select the one of the following as **Admin Host Port**:

- **Any—**Allows Cisco UCS Manager to determine the host port to which the vNIC is assigned.
- **1—**Explicitly assigns the vNIC to host port 1.

- **2—**Explicitly assigns the vNIC to host port 2.

Actual Host Port displays the actual assignment of the vNIC on a host port. When this feature is not supported, this will appear as **None**.

Step 7 Click **Save Changes**.

CIMC Mounted vMedia

Using Scriptable vMedia

Cisco UCS Manager allows provisioning of vMedia devices iso images for remote UCS servers. Using Scriptable vMedia, you can programmatically mount an IMG or an ISO image on a remote server. CIMC mounted vMedia provide communications between other mounted media inside your datacenter with no additional requirements media connection. Scriptable vMedia allows you to control virtual media devices without using a browser to manually map each UCS server individually.

Scriptable vMedia supports multiple share types including NFS, CIFS, HTTP, and HTTPS shares. **Scriptable vMedia** is enabled through BIOS configuration and configured through a Web GUI and CLI interface.

Cisco UCS Manager Scriptable vMedia supports the following functionality:

- Booting from a specific vMedia device
- Copying files from a mounted share to a local disk
- Installation and updating OS drivers

Note Cisco UCS Manager support for Scriptable vMedia is applicable for CIMC mapped devices only. Existing KVM based vMedia devices are not supported.

vMedia mount fails when the following conditions are met:

1. The remote vMedia image filename in the vMedia policy is set to **Service-Profile-Name**.
2. The service profile is renamed.

This is because the change in the name of the service profile does not change the remote vMedia image filename in the vMedia policy. The image filename still points to the older image on the remote device, which cannot be found.

Note Cisco UCS B200M2 Blade Server and Cisco UCS B230M2 Blade Server cannot use a vMedia policy as the policy is not supported on these blade servers.

Creating a vMedia Policy

A vMedia policy is used to configure the mapping information for remote vMedia devices. Two vMedia devices and mappings for CD and HDD are allowed in a vMedia policy. You can configure one ISO and one IMG at a time. ISO configurations maps to a CD drive and IMG configurations maps to a HDD device.

Note If you want to map a device to a remote folder, you must create an IMG and map it as a HDD device.

Before You Begin

Make sure that you have access to the following:

- Remote vMedia Server
- vMedia Devices

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **vMedia Policies** and select **Create vMedia Policy**.

Step 5 In the **Create vMedia Policy** dialog box, complete the following fields:

Name	Description
Name	The name of the vMedia policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description	A description of the policy. We recommend including information about where and when the policy should be used. Maximum 115 characters.
Retry on Mount Failure	Designates if the vMedia will continue mounting when a mount failure occurs. This can be: • **Yes** • **No** **Note** The default setting is **Yes**. When **Yes** is selected the remote server will continue to try to mount the vMedia mount process until it is successful or you disable this option. If you select **No**, a warning message will appear indicating retry on mount failure will not work in case of mount failure.

Step 6 On the icon bar to the right of the table, click +.

Step 7 In the **Create vMedia Mount** dialog box, complete the following fields:

Name	Description
Name	Name of the vMedia Mount policy.
Device Type	The type of remote vMedia you plan to mount. This can be: • **CDD—**Scriptable vMedia CD. • **HDD—**Scriptable vMedia HDD.
Protocol	The protocol to use when communicating with the remote server. Click one of the following radio buttons to indicate the protocol you want to use to communicate with the mounted remote server. This can be: • **NFS** - Network Files System. • **CIFS** - Common Internet File System. • **HTTP** - Hypertext Transfer Protocol. • **HTTPS** - Hypertext Transfer Protocol over Secure.

Name	Description
Authentication Protocol	The protocol to use for authentication when you use CIFS as the protocol for communicating with the remote server. When you use any protocol other than CIFS, this field is not available. Select one of the following from the drop-down list to specify the authentication protocol. • **Default—**NT LAN Manager Security Support Provider (NTLMSSP) protocol. Use this option only with Windows 2008 R2 and Windows 2012 R2. • **None—**No authentication is used • **Ntlm—**NT LAN Manager (NTLM) security protocol. Use this option only with Windows 2008 R2 and Windows 2012 R2. • **Ntlmi—**NTLMi security protocol. Use this option only when you enable Digital Signing in the CIFS Windows server. • **Ntlmssp—**NT LAN Manager Security Support Provider (NTLMSSP) protocol. Use this option only with Windows 2008 R2 and Windows 2012 R2. • **Ntlmsspi—**NTLMSSPi protocol. Use this option only when you enable Digital Signing in the CIFS Windows server. • **Ntlmv2—**NTLMv2 security protocol. Use this option only with Samba Linux. • **Ntlmv2i—**NTLMv2i security protocol. Use this option only with Samba Linux. **Note** The authentication protocol options are available only when you select **CIFS** as the protocol. For all other protocols, the **Authentication Protocol** field is disabled.
Hostname/IPAddress	Enter the IP address or hostname of the location where the backup file is to be stored. This can be a server, storage array, local drive, or any read/write media that the fabric interconnect can access through the network. If you use a hostname, you must configure Cisco UCS Manager to use a DNS server. The hostname (DNS) can be used when **Inband** network is configured for that server.

Name	Description
Image Name Variable	The name to be used for the image. This can be: • **None—**Filename must be entered in the **Remote File** field. • **Service Profile Name—**Filename automatically becomes the name of the service profile that the vMedia Policy is associated with. **Note** • If you select **Service Profile Name** as the **Image Name** variable, the **Remote File** field is disabled. • If you select **Service Profile Name** as the **Image Name** variable, do not rename the service profile. Renaming the service profile can result in vMedia mount failure.
Remote File	Enter the full path to the ISO or other image file. **Note** Ensure that the full path to the file begins with "/" after the share name. This field can contain the filename [with the file extension] only.
Remote Path	Enter the share name on the remote server, for example "Share".
Username	Enter the username that Cisco UCS Manager should use to log in to the remote server. This field does not apply if the protocol is NFS. This field is optional if the protocol is HTTP.
Password	Enter the password associated with the username. This field does not apply if the protocol is NFS. This field is optional if the protocol is HTTP.

Step 8 Click **OK**.
The remote server details are listed in the **vMedia Mounts** area of the **Create vMedia Mount** dialog box.

What to Do Next

Create a vMedia boot policy.

Adding a vMedia Policy to a Service Profile

Before you can use Scriptable vMedia, you must add the vMedia and Boot Policies to a Service Profile. After the vMedia and Boot Policies are added to a service profile you can associate the service profile with a Cisco UCS server. The following procedure describes how to add a vMedia policy to a Service Profile.

Before You Begin

Configure the vMedia Policy you want to add to a service profile.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization where you want to create the service profile.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the organization and select **Create Service Profile (expert)**.
The Unified Computing System Manager pane displays.

Step 5 In the **Name** field, enter a unique name that you can use to identify the service profile.
This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.

This name must be unique within the organization or sub-organization in which you are creating the service profile.

Step 6 From the **UUID Assignment** drop-down list, do one of the following:

Option	Description
Select (pool default used by default)	Assigns a UUID from the default UUID Suffix pool. Continue with Step 8.
Hardware Default	Uses the UUID assigned to the server by the manufacturer. If you choose this option, the UUID remains unassigned until the service profile is associated with a server. At that point, the UUID is set to the UUID value assigned to the server by the manufacturer. If the service profile is later moved to a different server, the UUID is changed to match the new server. Continue with Step 8.
XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX	Uses the UUID that you manually assign. Continue with Step 7.

Option	Description
Pools *Pool_Name*	Assigns a UUID from the UUID Suffix pool that you select from the list at the bottom of the drop-down list. Each pool name is followed by two numbers in parentheses that show the number of UUIDs still available in the pool and the total number of UUIDs in the pool. If you do not want use any of the existing pools, but instead want to create a pool that all service profiles can access, continue with Step 4. Otherwise, continue with Step 8.

Step 7 (Optional) If you selected the **XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX** option, do the following:

a) In the **UUID** field, enter the valid UUID that you want to assign to the server which uses this service profile.

b) To verify that the selected UUID is available, click the **here** link.

Step 8 (Optional) If you want to create a new UUID Suffix pool to use to use in this service profile, click **Create UUID Suffix Pool** and complete the fields in the **Create UUID Suffix Pool** wizard.
For more information, see Creating a UUID Suffix Pool, on page 394.

Step 9 (Optional) In the text box, enter a description of this service profile.
The user-defined description for this service profile.

Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Step 10 Click **Next**.

Step 11 From the **vMedia** drop down list, choose one of the following:

Option	Description
Select vMedia Policy to use	Enables you to assign a vMedia policy to this service profile. Continue with Step 12.
Create a Specific vMedia Policy	Enables you to create a local vMedia policy that can only be accessed by this service profile.

Option	Description
vMedia Policies *Policy_Name*	Assigns an existing vMedia policy to the service profile. If you choose this option, Cisco UCS Manager displays the details of the policy. If you do not want use any of the existing policies but instead want to create a policy that all service profiles can access, click **Create vMedia Policy** . Otherwise, choose a policy from the list and continue with Step 13.

Step 12 If you created a new vmedia policy accessible to all service profiles and template, choose that policy from the **vMedia** drop down list .

Step 13 Click **Next**.

Viewing CIMC vMedia Policy

Before You Begin

vMedia Policies are configured.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 On the **Servers** tab, expand **Policies** > **vMedia Policies**.

Step 3 Expand the **vMedia Policies** node to view the list of **vMedia Policies**.

Step 4 Double-click the name of a vMedia policy to view the properties for the selected **vMedia Mount**.
On the **Properties** page, you can modify the properties used for the **vMedia Mounts**.

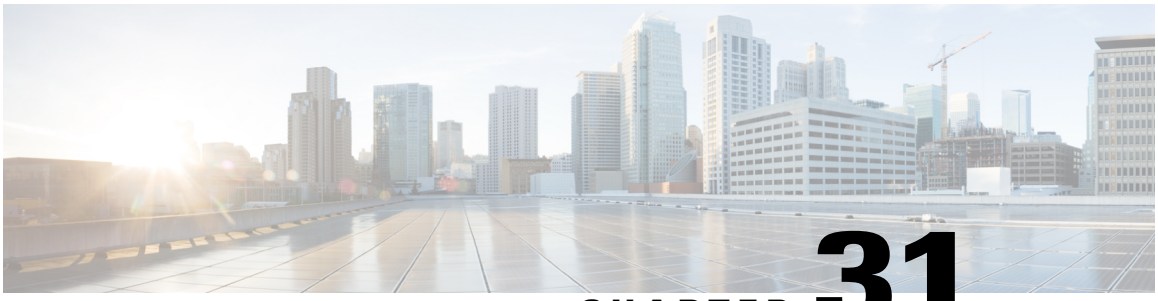

CHAPTER 31

Configuring Server Boot

This chapter includes the following sections:

Boot Policy

The Cisco UCS Manager enables you to create a boot policy for blade servers, rack servers, and modular servers.

The Cisco UCS Manager boot policy overrides the boot order in the BIOS setup menu and determines the following:

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

For example, you can 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 can either create a named boot policy to associate with one or more service profiles, or create a boot policy for a specific service profile. A boot policy must be included 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, Cisco UCS Manager applies the default boot policy.

Note Changes to a boot policy might be propagated to all servers created with an updating service profile template that includes that boot policy. Re-association of the service profile with the server to rewrite the boot order information in the BIOS is automatically triggered.

You can also specify the following for the boot policy:

- Local LUN name. The name specified is the logical name in the storage profile, not the deployed name. For modular servers, you can specify both a primary and secondary name. For other servers, specify only a primary name. Specifying a secondary name results in a configuration error.
- Specific JBOD disk number for booting from JBOD disks. This is not supported for the Modular servers.
- Any LUN for backward compatibility; however, we do not recommend this. Other devices must not have bootable images to ensure a successful boot.

UEFI Boot Mode

Unified Extensible Firmware Interface (UEFI) is a specification that defines a software interface between an operating system and platform firmware. Cisco UCS Manager uses UEFI to replace the BIOS firmware interfaces. This allows the BIOS to run in UEFI mode while still providing legacy support.

You can choose either legacy or UEFI boot mode when you create a boot policy. Legacy boot mode is supported for all Cisco UCS servers. UEFI boot mode is supported only on M3 and higher servers, and allows you to enable UEFI secure boot mode.

UEFI PXE boot is supported with all Cisco VIC adapters on Cisco UCS rack servers integrated with Cisco UCS Manager Release 2.2(4) and later releases. Beginning with Cisco UCS Manager Release 2.2(1), UEFI PXE boot is supported on all Cisco blade servers.

The following limitations apply to the UEFI boot mode:

- UEFI boot mode is not supported with the following combinations:
  - Gen-3 Emulex and QLogic adapters on Cisco UCS blade and rack servers integrated with Cisco UCS Manager.
  - iSCSI boot for all adapters on Cisco UCS rack servers integrated with Cisco UCS Manager.
- If you want to use UEFI boot mode with two iSCSI LUNs, you must manually specify a common iSCSI initiator name in the service profile that is applied to both underlying iSCSI eNICs rather than allowing Cisco UCS Manager to select the name from an IQN suffix pool. If you do not supply a common name, Cisco UCS Manager will not be able to detect the second iSCSI LUN.
- You cannot mix UEFI and legacy boot mode on the same server.

- The server will boot correctly in UEFI mode only if the boot devices configured in the boot policy have UEFI-aware operating systems installed. If a compatible OS is not present, the boot device is not displayed on the **Actual Boot Order** tab in the **Boot Order Details** area.
- In some corner cases, the UEFI boot may not succeed because the UEFI boot manager entry was not saved correctly in the BIOS NVRAM. You can use the UEFI shell to enter the UEFI boot manager entry manually. This situation could occur in the following situations:
  - If a blade server with UEFI boot mode enabled is disassociated from the service profile, and the blade is manually powered on using the **Equipment** tab or the front panel.
  - If a blade server with UEFI boot mode enabled is disassociated from the service profile, and a direct VIC firmware upgrade is attempted.
  - If a blade or rack server with UEFI boot mode enabled is booted off SAN LUN, and the service profile is migrated.

You can create UEFI boot parameters in Cisco UCS Manager. UEFI Boot Parameters, on page 563 provides more information.

UEFI Secure Boot

Cisco UCS Manager supports UEFI secure boot on Cisco UCS B-Series M3 and M4 Blade servers, Cisco UCS C-Series M3 and M4 Rack servers, and Cisco UCS S-Series M4 Rack servers. When UEFI secure boot is enabled, all executables, such as boot loaders and adapter drivers, are authenticated by the BIOS before they can be loaded. To be authenticated, the images must be signed by either the Cisco Certificate Authority (CA) or a Microsoft CA.

The following limitations apply to UEFI secure boot:

- UEFI boot mode must be enabled in the boot policy.
- The Cisco UCS Manager software and the BIOS firmware must be at Release 2.2 or greater.

Note UEFI boot mode is supported on Cisco UCS C-Series and S-Series rack servers beginning with Release 2.2(3a).

- User-generated encryption keys are not supported.
- UEFI secure boot can only be controlled by Cisco UCS Manager.
- If you want to downgrade to an earlier version of Cisco UCS Manager, and you have a server in secure boot mode, you must disassociate, then re-associate the server before downgrading. Otherwise, server discovery is not successful.

CIMC Secure Boot

With CIMC secure boot, only Cisco signed firmware images can be installed and run on the servers. When the CIMC is updated, the image is certified before the firmware is flashed. If certification fails, the firmware is not flashed. This prevents unauthorized access to the CIMC firmware.

Guidelines and Limitations for CIMC Secure Boot

- CIMC secure boot is supported on Cisco UCS M3 rack servers.

Note CIMC secure boot is enabled by default on the Cisco UCS C220 M4, C240 M4 rack servers, and is automatically enabled on the Cisco UCS C460 M4 rack server after upgrading to CIMC firmware release 2.2(3) or higher.

- After CIMC secure boot is enabled, you cannot disable it.
- After CIMC secure boot is enabled on a server, you cannot downgrade to a CIMC firmware image prior to 2.1(3).

Determining the CIMC Secure Boot Status

Procedure

Step 1 In the **Navigation** pane, click the **Equipment** tab.
Step 2 Expand **Equipment** > **Rack-Mounts** > **Servers** > ***Server Name***.
Step 3 In the **Work** area, click the **Inventory** tab.
Step 4 Click the **CIMC** subtab.
Step 5 In the **CIMC** area, note the **Secure Boot Operational State** field.
This can be one of the following:

- **Unsupported—**CIMC secure boot is not supported on the server.
- **Disabled—**CIMC secure boot is supported, but is disabled on the server.
- **Enabling—**CIMC secure boot was enabled, and the operation is in process.
- **Enabled—**CIMC secure boot is enabled on the server.

Enabling CIMC Secure Boot on a Rack Server

Procedure

Step 1 In the **Navigation** pane, click the **Equipment** tab.
Step 2 Expand **Equipment** > **Rack-Mounts** > **Servers** > ***Server Name***.
Step 3 In the **Work** area, click the **Inventory** tab.
Step 4 Click the **CIMC** subtab.
Step 5 In the **Actions** area, click **Enable Secure Boot**.

CIMC secure boot is only supported on Cisco UCS M3 rack servers. If CIMC secure boot is not supported or is already enabled, this action is greyed.

Step 6 Click **Yes** in the **Enable Secure Boot** confirmation dialog box.

Note After enabled, you cannot disable CIMC secure boot.

Creating a Boot Policy

You can also create a local boot policy that is restricted to a service profile or service profile template. However, Cisco recommends that you create a global boot policy that can be included in multiple service profiles or service profile templates.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **Boot Policies** and select **Create Boot Policy**.
The **Create Boot Policy** wizard displays.

Step 5 Enter a unique name and description for the policy.
This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Step 6 (Optional) After you make changes to the boot order, check the **Reboot on Boot Order Change** check box to reboot all servers that use this boot policy.
For boot policies applied to a server with a non-Cisco VIC adapter, even if the **Reboot on Boot Order Change** check box is not checked, when SAN devices are added, deleted, or their order is changed, the server always reboots when boot policy changes are saved.

Step 7 (Optional) If desired, check the **Enforce vNIC/vHBA/iSCSI Name** check box.

- If checked, Cisco UCS Manager displays a configuration error and reports whether one or more of the vNICs, vHBAs, or iSCSI vNICs listed in the **Boot Order** table match the server configuration in the service profile.
- If not checked, Cisco UCS Manager uses the vNICs or vHBAs (as appropriate for the boot option) from the service profile.

Step 8 In the Boot Mode field, choose the **Legacy** or **UEFI** radio button.

Step 9 If you selected UEFI, check the **Boot Security** checkbox if you want to enable UEFI boot security.

Step 10 Configure one or more of the following boot options for the boot policy and set their boot order:

- Local Devices boot—To boot from local devices, such as local disks on the server, virtual media, or remote virtual disks, continue with Configuring a Local Disk Boot for a Boot Policy, on page 556.
- SAN boot—To boot from an operating system image on the SAN, continue with Configuring a SAN Boot for a Boot Policy, on page 528.

  You can specify a primary and a secondary SAN boot. If the primary boot fails, the server attempts to boot from the secondary.
- LAN boot—To boot from a centralized provisioning server, continue with Configuring a LAN Boot for a Boot Policy, on page 555.
- iSCSI boot—To boot from an iSCSI LUN, continue with Creating an iSCSI Boot Policy, on page 539.

What to Do Next

Include the boot policy in a service profile and template.

After a server is associated with a service profile that includes this boot policy, you can verify the boot order in the **Boot Order Details** area on the **General** tab for the server.

SAN Boot

You can configure a boot policy to boot one or more servers from an operating system image on the SAN. The boot policy can include a primary and a secondary SAN boot. If the primary boot fails, the server attempts to boot from the secondary.

Cisco recommends using 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 same operating system image. Therefore, the new server appears as the same server to the network.

To use a SAN boot, ensure that the following is configured:

- The Cisco UCS domain must be able to communicate with the SAN storage device that hosts the operating system image.
- A boot target LUN (Logical Unit Number) on the device where the operating system image is located.

Note SAN boot is not supported on Gen-3 Emulex adapters on Cisco UCS blade and rack servers.

Configuring a SAN Boot for a Boot Policy

You can also create a local boot policy that is restricted to a service profile or service profile template. However, Cisco recommends that you create a global boot policy that can be included in multiple service profiles or service profile templates.

Tip If you configure a local disk and a SAN LUN for the boot order storage type and the operating system or logical volume manager (LVM) is configured incorrectly, the server might boot from the local disk rather than the SAN LUN.

For example, on a server with Red Hat Linux installed, where the LVM is configured with default LV names and the boot order is configured with a SAN LUN and a local disk, Linux reports that there are two LVs with the same name and boots from the LV with the lowest SCSI ID, which could be the local disk.

This procedure continues directly from Creating a Boot Policy, on page 527.

Before You Begin

Note If you are creating a boot policy that boots the server from a SAN LUN and you require reliable SAN boot operations, Cisco recommends that you first remove all local disks and other SAN LUNs from the boot policy in the server service profile.

This does not apply to the UCS Mini Series.

Procedure

Step 1 Click the down arrows to expand the **vHBAs** area.

Step 2 Click the **Add SAN Boot** link.

Step 3 In the **Add San Boot** dialog box, specify the vHBA and type, then click **OK**.

Step 4 If this vHBA points to a bootable SAN image, click the **Add SAN Boot Target** link, and in the **Add SAN Boot Target** dialog box specify the boot target LUN, boot target WWPN, and type, then click **OK**:

Step 5 Do one of the following:

- Add another boot device to the **Boot Order** table.
- Click **OK** to finish.

What to Do Next

Include the boot policy in a service profile and template.

After a server is associated with a service profile that includes this boot policy, you can verify the boot order in the **Boot Order Details** area on the **General** tab for the server.

iSCSI Boot

iSCSI boot enables a server to boot its operating system from an iSCSI target machine located remotely over a network.

iSCSI boot is supported on the following Cisco UCS hardware:

- Cisco UCS blade servers that have the Cisco UCS M51KR-B Broadcom BCM57711 network adapter and use the default MAC address provided by Broadcom.
- Cisco UCS M81KR Virtual Interface Card
- Cisco UCS VIC-1240 Virtual Interface Card
- Cisco UCS VIC-1280 Virtual Interface Card
- Cisco UCS rack servers that have the Cisco UCS M61KR-B Broadcom BCM57712 network adapter.
- Cisco UCS P81E Virtual Interface Card
- Cisco UCS VIC 1225 Virtual Interface Cardon Cisco UCS rack servers

There are prerequisites that must be met before you configure iSCSI boot. For a list of these prerequisites, see iSCSI Boot Guidelines and Prerequisites, on page 531.

For a high-level procedure for implementing iSCSI boot, see Configuring iSCSI Boot, on page 533.

iSCSI Boot Process

Cisco UCS Manager uses the iSCSI vNIC and iSCSI boot information created for the service profile in the association process to program the adapter, located on the server. After the adapter is programmed, the server reboots with the latest service profile values. After the power on self-test (POST), the adapter attempts to initialize using these service profile values. If the adapter can use the values and log in to its specified target, the adapter initializes and posts an iSCSI Boot Firmware Table (iBFT) to the host memory and a valid bootable LUN to the system BIOS. The iBFT that is posted to the host memory contains the initiator and target configuration that is programmed on the primary iSCSI VNIC.

Note Previously, the host could see only one of the boot paths configured, depending on which path completed the LUN discovery first, and would boot from that path. Now, when there are two iSCSI boot vNICs configured, the host sees both of the boot paths. So for multipath configurations, a single IQN must be configured on both the boot vNICs. If there are different IQNs configured on the boot vNICs on a host, the host boots with the IQN that is configured on the boot vNIC with the lower PCI order.

The next step, which is the installation of the operating system (OS), requires an OS that is iBFT capable. During installation of the OS, the OS installer scans the host memory for the iBFT table and uses the information in the iBFT to discover the boot device and create an iSCSI path to the target LUN. Some OSs requires a NIC driver to complete this path. If this step is successful, the OS installer finds the iSCSI target LUN on which to install the OS.

Note The iBFT works at the OS installation software level and might not work with HBA mode (also known as TCP offload). Whether iBFT works with HBA mode depends on the OS capabilities during installation. Also, for a server that includes a Cisco UCS M51KR-B Broadcom BCM57711 adapter, the iBFT normally works at a maximum transmission unit (MTU) size of 1500, regardless of the MTU jumbo configuration. If the OS supports HBA mode, you might need to set HBA mode, dual-fabric support, and jumbo MTU size after the iSCSI installation process.

iSCSI Boot Guidelines and Prerequisites

These guidelines and prerequisites must be met before configuring iSCSI boot:

- After the iSCSI boot policies are created, a user with ls-compute privileges can include them in a service profile or service profile template. However, a user with only ls-compute privileges cannot create iSCSI boot policies.
- To set up iSCSI boot from a Windows 2008 server where the second vNIC (failover vNIC) must boot from an iSCSI LUN, consult Microsoft Knowledge Base Article 976042. Microsoft has a known issue where Windows might fail to boot from an iSCSI drive or cause a bugcheck error if the networking hardware is changed. To work around this issue, follow the resolution recommended by Microsoft.
- The storage array must be licensed for iSCSI boot and the array side LUN masking must be properly configured.
- Two IP addresses must be determined, one for each iSCSI initiator. If possible, the IP addresses should be on the same subnet as the storage array. The IP addresses are assigned statically or dynamically using the Dynamic Host Configuration Protocol (DHCP).
- You cannot configure boot parameters in the Global boot policy. Instead, after configuring boot parameters, include the boot policy in the appropriate service profile.
- The operating system (OS) must be iSCSI Boot Firmware Table (iBFT) compatible.
- For Cisco UCS M51KR-B Broadcom BCM57711 network adapters:
  - Servers that use iSCSI boot must contain the Cisco UCS M51KR-B Broadcom BCM57711 network adapter. For information on installing or replacing an adapter card, see the *Cisco UCS B250 Extended Memory Blade Server Installation and Service Note*. The service note is accessible from the *Cisco UCS B-Series Servers Documentation Roadmap* at http://www.cisco.com/go/unifiedcomputing/b-series-doc.
  - Set the MAC addresses on the iSCSI device.
  - If you are using the DHCP Vendor ID (Option 43), configure the MAC address of an iSCSI device in /etc/dhcpd.conf.
  - HBA mode (also known as TCP offload) and the boot to target setting are supported. However, only Windows OS supports HBA mode during installation.
  - Before installing the OS, disable the boot to target setting in the iSCSI adapter policy, then after installing the OS, re-enable the boot to target setting.

    Note Each time you change an adapter policy setting, the adapter reboots to apply the new setting.

  - When installing the OS on the iSCSI target, the iSCSI target must be ordered *before* the device where the OS image resides. For example, if you are installing the OS on the iSCSI target from a CD, the boot order should be the iSCSI target and then the CD.
  - After the server is iSCSI booted, do not modify the Initiator Name, Target name, LUN, iSCSI device IP, or Netmask/gateway using the Broadcom tool.

- Do not interrupt the POST (power on self-test) process or the Cisco UCS M51KR-B Broadcom BCM57711 network adapter will fail to initialize.

- For Cisco UCS M81KR Virtual Interface Card and Cisco UCS VIC-1240 Virtual Interface Card:

  For Cisco UCS VIC-1240 Virtual Interface Card:

  - Do not set MAC addresses on the iSCSI device.
  - HBA mode and the boot to target setting are *not* supported.
  - When installing the OS on the iSCSI target, the iSCSI target must be ordered *after* the device where the OS image resides. For example, if you are installing the OS on the iSCSI target from a CD, the boot order should be the CD and then the iSCSI target.
  - If you are using the DHCP Vendor ID (Option 43), the MAC address of the overlay vNIC must be configured in /etc/dhcpd.conf.
  - After the server is iSCSI booted, do not modify the IP details of the overlay vNIC.

- The VMware ESX/ESXi operating system does not support storing a core dump file to an iSCSI boot target LUN. Dump files must be written to a local disk.

Initiator IQN Configuration

Cisco UCS uses the following rules to determine the initiator IQN for an adapter iSCSI vNIC at the time a service profile is associated with a physical server:

- An initiator IQN at the service profile level *and* at the iSCSI vNIC level cannot be used together in a service profile.
- If an initiator IQN is specified at the service profile level, all of the adaptor iSCSI vNICs are configured to use the same initiator IQN, except in the case of DHCP Option 43, where the initiator IQN is set to empty on the adapter iSCSI vNIC.
- When an initiator IQN is set at the iSCSI vNIC level, the initiator IQN at the service profile level is removed, if one is present.
- If there are two iSCSI vNIC in a service profile and only one of them has the initiator IQN set, the second one is configured with the default IQN pool. You can change this configuration later. The only exception is if DHCP Option 43 is configured. In this case, the initiator IQN on the second iSCSI vNIC is removed during service profile association.

Note If you change an iSCSI vNIC to use the DHCP Option 43 by setting the vendor ID, it does not remove the initiator IQN configured at the service profile level. The initiator IQN at the service profile level can still be used by another iSCSI vNIC which does not use the DHCP Option 43.

Enabling MPIO on Windows

You can enable (MPIO) to optimize connectivity with storage arrays.

Note If you change the networking hardware, Windows might fail to boot from an iSCSI drive. For more information, see Microsoft support Article ID: 976042.

Before You Begin

The server on which you enable the Microsoft Multipath I/O (MPIO) must have a Cisco VIC driver.

If there are multiple paths configured to the boot LUN, only one path should be enabled when the LUN is installed.

Procedure

Step 1 In the service profile associated with the server, configure the primary iSCSI vNIC.
For more information, see Creating an iSCSI vNIC for a Service Profile, on page 540.

Step 2 Using the primary iSCSI vNIC, install the Windows operating system on the iSCSI target LUN.

Step 3 After Windows installation completes, enable MPIO on the host.

Step 4 In the service profile associated with the server, add the secondary iSCSI vNIC to the boot policy.
For more information, see Creating an iSCSI Boot Policy, on page 539.

Configuring iSCSI Boot

When you configure an adapter or blade in Cisco UCS to iSCSI boot from a LUN target, complete all of the following steps.

Procedure

	Command or Action	Purpose
Step 1	Configure the iSCSI boot adapter policy.	(Optional) For more information, see Creating an iSCSI Boot Policy, on page 539
Step 2	Configure the authentication profiles for the initiator and target.	(Optional) For more information, see Creating an iSCSI Authentication Profile, on page 536
Step 3	To configure the iSCSI initiator to use an IP address from a pool of IP addresses, add a block of IP addresses to the iSCSI initiator pool.	(Optional) For more information, see Creating an iSCSI Initiator IP Pool, on page 538

	Command or Action	Purpose
Step 4	Create a boot policy that can be used in any service profile. Alternatively, you can create a local boot policy only for the specific service policy. However, Cisco recommends that you create a boot policy that can be shared with multiple service profiles.	For more information about creating a boot policy that can be used in any service profile, see Creating an iSCSI Boot Policy, on page 539.
Step 5	If you created a boot policy that can be used in any service profile, assign it to the service profile. Otherwise, proceed to the next step.	You can assign the boot policy to the service profile while configuring the iSCSI boot and vNIC parameters in the service profile in step 7.
Step 6	Create an iSCSI vNIC in a service profile.	For more information, see Creating an iSCSI vNIC for a Service Profile, on page 540
Step 7	Configure the iSCSI boot parameters, including the iSCSI qualifier name (IQN), initiator, target interfaces, and iSCSI vNIC parameters in a service profile in expert mode or service profile template.	For more information, see Creating a Service Profile with the Expert Wizard, on page 592 or Creating a Service Profile Template, on page 595, respectively.
Step 8	Verify the iSCSI boot operation.	For more information, see *Verifying iSCSI Boot*.
Step 9	Install the OS on the server.	For more information, see one of the following guides: • *Cisco UCS B-Series Blade Servers VMware Installation Guide* • *Cisco UCS B-Series Blade Servers Linux Installation Guide* • *Cisco UCS B-Series Blade Servers Windows Installation Guide*
Step 10	Boot the server.	

Creating an iSCSI Adapter Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **Adapter Policies** and choose **Create iSCSI Adapter Policy**.

Step 5 In the **Create iSCSI Adapter Policy** dialog box, complete the following fields:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Connection Timeout field	The number of seconds to wait until Cisco UCS assumes that the initial login has failed and the iSCSI adapter is unavailable. Enter an integer between 0 and 255. If you enter 0, Cisco UCS uses the value set in the adapter firmware (default: 15 seconds).
LUN Busy Retry Count field	The number of times to retry the connection in case of a failure during iSCSI LUN discovery. Enter an integer between 0 and 60. If you enter 0, Cisco UCS uses the value set in the adapter firmware (default: 15 seconds).
DHCP Timeout field	The number of seconds to wait before the initiator assumes that the DHCP server is unavailable. Enter an integer between 60 and 300 (default: 60 seconds).
Enable TCP Timestamp check box	Check this box if you want to use a TCP Timestamp. With this setting, transmitted packets are given a time stamp of when the packet was sent so that the packet's round-trip time can be calculated, when needed. **Note** This option only applies to servers with the Cisco UCS NIC M51KR-B adapter.
HBA Mode check box	Check this box to enable HBA mode (also known as TCP offload). **Important** This option should only be enabled for servers with the Cisco UCS NIC M51KR-B adapter running the Windows operating system.
Boot to Target check box	Check this box to boot from the iSCSI target. **Note** This option only applies to servers with the Cisco UCS NIC M51KR-B adapter. It should be disabled until you have installed an operating system on the server.

Name	Description
Owner field	This can be one of the following: • **Local—**This policy is available only to service profiles and service profile templates in this Cisco UCS domain. • **Pending Global—**Control of this policy is being transferred to Cisco UCS Central. Once the transfer is complete, this policy will be available to all Cisco UCS domains registered with Cisco UCS Central. • **Global—**This policy is managed by Cisco UCS Central. Any changes to this policy must be made through Cisco UCS Central.

Step 6 Click **OK**.

What to Do Next

Include the adapter policy in a service profile and template.

Deleting an iSCSI Adapter Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Policies**.
Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.
Step 4 Expand the **Adapter Policies** node.
Step 5 Right-click the adapter policy and choose **Delete.**
Step 6 If a confirmation dialog box displays, click **Yes**.

Creating an iSCSI Authentication Profile

For iSCSI boot, you need to create both an initiator and a target iSCSI authentication profile.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **iSCSI Authentication Profiles** and choose **Create iSCSI Authentication Profile**.

Step 5 In the **Create Authentication Profile** dialog box, complete the following fields:

Name	Description
Name field	The name of the authentication profile. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
User Id field	The user Id associated with this profile. Enter between 1 and 128 characters, spaces, or special characters.
Password field	The password associated with this profile. Enter between 12 and 16 characters, including special characters.
Confirm Password field	The password again for confirmation purposes.

Step 6 Click **OK**.

What to Do Next

Include the authentication profile in a service profile and template.

Deleting an iSCSI Authentication Profile

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Expand the **iSCSI Authentication Profiles** node.

Step 5 Right-click the IP pool you want to delete and choose **Delete**.

Step 6 If a confirmation dialog box displays, click **Yes**.

Creating an iSCSI Initiator IP Pool

You can create a group of IP addresses to be used for iSCSI boot. Cisco UCS Manager reserves the block of IPv4 addresses you specify.

The IP pool must not contain any IP addresses that were assigned as static IP addresses for a server or service profile.

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN** > **Pools**.

Step 3 Expand the node for the organization where you want to create the pool.
If the system does not include multitenancy, expand the **root** node.

Step 4 Expand the **IP Pools** node.

Step 5 Right-click **IP Pool iscsi-initiator-pool** and choose **Create Block of IPv4 Addresses**.

Step 6 In the **Create a Block of IPv4 Addresses** dialog box, complete the following fields:

Name	Description
Name column	The range of IPv4 addresses assigned to the block.
From column	The first IPv4 address in the block.
To column	The last IPv4 address in the block.
Subnet column	The subnet mask associated with the IPv4 addresses in the block.
Default Gateway column	The default gateway associated with the IPv4 addresses in the block.
Primary DNS column	The primary DNS server that this block of IPv4 addresses should access.
Secondary DNS column	The secondary DNS server that this block of IPv4 addresses should access.

Step 7 Click **OK**.

What to Do Next

Configure one or more service profiles or service profile templates to obtain the iSCSI initiator IP address from the iSCSI initiator IP pool.

Creating an iSCSI Boot Policy

You can add up to two iSCSI vNICs per boot policy. One vNIC acts as the primary iSCSI boot source, and the other acts as the secondary iSCSI boot source.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **Boot Policies** and choose **Create Boot Policy**.
The **Create Boot Policy** wizard displays.

Step 5 Enter a unique name and description for the policy.
This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Step 6 (Optional) To reboot a server that uses this boot policy after you make changes to the boot order, check the **Reboot on Boot Order Change** check box.
In the Cisco UCS Manager GUI, if the **Reboot on Boot Order Change** check box is checked for a boot policy, and if CD-ROM or Floppy is the last device in the boot order, deleting or adding the device does not directly affect the boot order and the server does not reboot.

Note This applies only to servers using the standard boot order.

Step 7 (Optional) If desired, check the **Enforce vNIC/vHBA/iSCSI Name** check box.

- If checked, Cisco UCS Manager displays a configuration error and reports whether one or more of the vNICs, vHBAs, or iSCSI vNICs listed in the **Boot Order** table match the server configuration in the service profile.
- If not checked, Cisco UCS Manager uses the vNICs or vHBAs (as appropriate for the boot option) from the service profile.

Step 8 To add a iSCSI boot to the boot policy, do the following:

a) Click the down arrows to expand the iSCSI vNICs area.
b) Click the **Add iSCSI Boot** link.
c) In the **Add iSCSI Boot** dialog box, enter a name for the iSCSI vNIC, and click **OK**.
d) Repeat steps b and c to create another iSCSI vNIC.

What to Do Next

Include the boot policy in a service profile and template.

After a server is associated with a service profile that includes this boot policy, you can verify the actual boot order in the **Boot Order Details** area on the **General** tab for the server.

Creating an iSCSI vNIC for a Service Profile

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that contains the service profile for which you want to create an iSCSI vNIC.

Step 4 Expand the service profile for which you want to create a iSCSI vNIC.

Step 5 Right-click the **iSCSI vNICs** node and choose **Create vNICs**.

Step 6 In the **Create iSCSI vNIC** dialog box, complete the following fields:

Name	Description
Name field	The name of the iSCSI vNIC. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Overlay vNIC drop-down list	The LAN vNIC associated with this iSCSI vNIC, if any.
iSCSI Adapter Policy drop-down list	The iSCSI adapter policy associated with this iSCSI vNIC, if any.
Create iSCSI Adapter Policy link	Click this link to create a new iSCSI adapter policy that will be available to all iSCSI vNICs.
MAC Address field	The MAC address associated with this iSCSI vNIC, if any. If the MAC address is not set, the Cisco UCS Manager GUI displays **Derived**.
MAC Pool field	The MAC pool associated with this iSCSI vNIC, if any.

Name	Description
VLAN drop-down list	The virtual LAN associated with this iSCSI vNIC. The default VLAN is **default**. **Note** For the Cisco UCS M81KR Virtual Interface Card and the Cisco UCS VIC-1240 Virtual Interface Card, the VLAN that you specify must be the same as the native VLAN on the overlay vNIC. For the Cisco UCS M51KR-B Broadcom BCM57711 Adapter, the VLAN that you specify can be any VLAN assigned to the overlay vNIC.

Step 7 In the **MAC Address Assignment** drop-down list in the **iSCSI MAC Address** area, choose one of the following:

- Leave the MAC address unassigned, select **Select (None used by default)**. Select this option if the server that will be associated with this service profile contains a Cisco UCS M81KR Virtual Interface Card adapter or a Cisco UCS VIC-1240 Virtual Interface Card.

  Important If the server that will be associated with this service profile contains a Cisco UCS NIC M51KR-B adapter, you must specify a MAC address.

- A specific MAC address, select **00:25:B5:XX:XX:XX** and enter the address in the **MAC Address** field. To verify that this address is available, click the corresponding link.
- A MAC address from a pool, select the pool name from the list. Each pool name is followed by a pair of numbers in parentheses. The first number is the number of available MAC addresses in the pool and the second is the total number of MAC addresses in the pool.

  If this Cisco UCS domain is registered with Cisco UCS Central, there might be two pool categories. **Domain Pools** are defined locally in the Cisco UCS domain and **Global Pools** are defined in Cisco UCS Central.

Step 8 (Optional) If you want to create a MAC pool that will be available to all service profiles, click **Create MAC Pool** and complete the fields in the **Create MAC Pool** wizard.
For more information, see Creating a MAC Pool, on page 271.

Step 9 Click **OK**.

Step 10 (Optional) If you want to set or change the initiator name, from the **iSCSI vNICs** tab, click **Reset Initiator Name** or **Change Initiator Name** and complete the fields in the **Change Initiator Name** dialog box or click . For more information, see either Setting the Initiator IQN at the Service Profile Level, on page 542 or Setting the Initiator IQN at the Service Profile Level, on page 542.

Deleting an iSCSI vNIC from a Service Profile

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Service Profiles**.
Step 3 Expand the node for the organization that contains the service profile from which you want to delete an iSCSI vNIC.
Step 4 Expand the service profile from which you want to delete an iSCSI vNIC.
Step 5 Expand the **iSCSI vNICs** node.
Step 6 Right-click the iSCSI vNIC you want to delete and choose **Delete**.
Step 7 If a confirmation dialog box displays, click **Yes**.

Setting the Initiator IQN at the Service Profile Level

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Service Profiles**.
Step 3 Expand the desired node for the organization.
Step 4 Click the service profile that has the iSCSI vNIC that you want to change.
Step 5 In the **Work** pane, click the **iSCSI vNICs** tab.
Step 6 Click **Reset Initiator Name**.
Step 7 If a confirmation dialog box displays, click **Yes**.

Changing the Initiator IQN at the Service Profile Level

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the desired node for the organization.

Step 4 Click the service profile that has the iSCSI vNIC that you want to change.

Step 5 In the **Work** pane, click the **iSCSI vNICs** tab.

Step 6 In the **Actions** area, click **Change Initiator Name**.

Step 7 In the **Change Initiator Name** dialog box, change the values in the following fields

Name	Description
Initiator Name Assignment drop-down list	Choose the IQN initiator name that you want to use from the drop-down list.
Initiator Name field	If you selected a manual initiator name assignment, enter the initiator name.
Create IQN Suffix Pool link	Click to create a new IQN suffix pool.

Step 8 Click **OK**.

Setting iSCSI Boot Parameters

You can set iSCSI boot parameters, including the boot order, boot policy, iSCSI authentication profile, initiator interface, and target interface for an iSCSI vNIC.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that contains the service profile for which you want to create iSCSI boot parameters. If the system does not include multi-tenancy, expand the root node.

Step 4 Click the service profile for which you want to create iSCSI boot parameters.

Step 5 Click the **Boot Order** tab.

Step 6 In the **Specific Boot Policy** area, click the down arrows to expand the **iSCSI vNICs** area.

Step 7 In the **iSCSI vNICs** area, double-click the iSCSI vNICs from which you want to boot the server to add them to the **Boot Order** table.

Step 8 In the **iSCSI vNICs** area, click the **Set Boot Parameters** link.
If there are two iSCSI vNICs, choose the one for which you want to set boot parameters.

Step 9 In the **Set iSCSI Boot Parameters** dialog box, complete the following fields:

Name	Description
Name field	The name of the iSCSI vNIC for which you are setting the boot parameters.
Authentication Profile drop-down list	The name of the associated iSCSI authentication profile.
Create Authentication Profile link	Click this link to create a new iSCSI authentication profile that will be available to all iSCSI vNICs.

Step 10 In the **Initiator Name** area, complete the following fields:

Name	Description
Initiator Name Assignment drop-down list	Select how the iSCSI boot initiator name is assigned. Choose one of the following methods: • **Manual—**You will enter a name in the **Initiator Name** field. The initiator name can contain up to 223 characters. • **Pools—**Choose an IQN suffix pool from which the name will be assigned. **Note** Setting the Initiator Name from the Set iSCSI Boot Parameters dialog box sets the initiator IQN at the iSCSI vNIC level and not at the service profile level. If more than one path is configured, you must set the initiator IQN from the **iSCSI vNICs** tab or when creating a service profile. If you need to, you can change or reset the initiator name. For more information, see Changing the Initiator IQN at the Service Profile Level, on page 543.

Name	Description
Create IQN Suffix Pool link	Click this link to create a new IQN suffix pool that will be available to all iSCSI vNICs.
Initiator Name field	A regular expression that defines the name of the iSCSI initiator. You can enter any alphanumeric string as well as the following special characters: • . (period) • : (colon) • - (dash)

Step 11 From the **Initiator IP Address Policy** drop-down list, choose of the following:

Option	Description
Select (DHCP used by default)	The system selects an interface automatically using DHCP. Proceed to Step 13.
Static	A static IPv4 address is assigned to the iSCSI boot vNIC based on the information entered in this area. Proceed to Step 12.
Pool	An IPv4 address is assigned to the iSCSI boot vNIC from the management IP address pool. Proceed to Step 13.

Step 12 If you chose **Static** from the **Initiator IP Address Policy** drop-down list, complete the following fields:

Name	Description
IPv4 Address field	The IPv4 address assigned to the iSCSI boot vNIC. If you want to specify this address, you must select **Static** in the **Initiator IP Address Policy** drop-down list.
Subnet Mask field	The subnet mask associated with the IPv4 address.
Default Gateway field	The default gateway associated with the IPv4 address.
Primary DNS field	The primary DNS server address.
Secondary DNS field	The secondary DNS server address.

Step 13 For the iSCSI target interface, choose one of the following radio buttons:

Option	Description
iSCSI Static Target Interface	The system creates a static target interface that you need to configure. Proceed to Step 14.
iSCSI Auto Target Interface	The system creates an auto target interface. You need to specify whether the auto target uses an initiator or a DCHP vendor ID. Proceed to Step 16.

Step 14 If you chose **iSCSI Static Target Interface**, in the **Static Target Interface** table, click **Add**.

Step 15 In the **Create iSCSI Static Target** dialog box, complete the following fields:

Name	Description
iSCSI Target Name field	A regular expression that defines the iSCSI Qualified Name (IQN) or Extended Unique Identifier (EUI) name of the iSCSI target. You can enter any alphanumeric characters as well as the following special characters: • . (period) • : (colon) • - (dash) **Important** This name must be properly formatted using standard IQN or EUI guidelines. The following examples show properly formatted iSCSI target names: • iqn.2001-04.com.example • iqn.2001-04.com.example:storage:diskarrays-sn-a8675309 • iqn.2001-04.com.example:storage.tape1.sys1.xyz • iqn.2001-04.com.example:storage.disk2.sys1.xyz • eui.02004567A425678D
Priority field	The system-assigned priority for the iSCSI target.
Port field	The port associated with the iSCSI target. Enter an integer between 1 and 65535. The default is 3260.
Authentication Profile drop-down list	The name of the associated iSCSI authentication profile.
Create iSCSI Authentication Profile link	Click this link to create a new iSCSI authentication profile that will be available to all iSCSI vNICs.
IPv4 Address field	The IPv4 address assigned to the iSCSI target.

Name	Description
LUN Id field	The LUN identifier in the iSCSI target.

Step 16 If you chose **iSCSI Auto Target Interface**, enter either the initiator name or the DHCP vendor ID in the **DHCP Vendor Id** field. The initiator must have already been configured. The vendor ID can be up to 32 alphanumeric characters.

Step 17 Click **OK**.

Modifying iSCSI Boot Parameters

You can modify iSCSI boot parameters, including the boot order, boot policy, iSCSI authentication profile, initiator interface, and target interface for an iSCSI vNIC.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that contains the service profile for which you want to modify iSCSI boot parameters. If the system does not include multi-tenancy, expand the root node.

Step 4 Click the service profile for which you want to modify iSCSI boot parameters.

Step 5 Click the **Boot Order** tab.

Step 6 In the **Specific Boot Policy** area, click the down arrows to expand the **iSCSI vNICs** area.

Step 7 To add or delete an iSCSI vNIC from the boot order or to change the boot order, do one of the following:

- To add an iSCSI vNIC, in the **iSCSI vNICs** area, double-click an iSCSI vNICs to add it to the **Boot Order** table.
- To delete an iSCSI vNIC from the boot order, in the **Boot Order** table, select the iSCSI vNIC and click **Delete**.
- To change the iSCSI vNIC boot order, in the **Boot Order** table, select the iSCSI vNIC and click either **Move Up** or **Move Down**.

Step 8 To change the boot parameters, in the **iSCSI vNICs** area, click the **Set Boot Parameters** link.
If there are two iSCSI vNICs, choose the one for which you want to change boot parameters.

Step 9 In the **Set iSCSI Boot Parameters** dialog box, change the values in any of the following fields:

Name	Description
Name field	The name of the iSCSI vNIC for which you are setting the boot parameters.

Name	Description
Authentication Profile drop-down list	The name of the associated iSCSI authentication profile.
Create Authentication Profile link	Click this link to create a new iSCSI authentication profile that will be available to all iSCSI vNICs.

Step 10 In the **Initiator Name** area, complete the following fields:

Name	Description
Initiator Name Assignment drop-down list	Select how the iSCSI boot initiator name is assigned. Choose one of the following methods: • **Manual—**You will enter a name in the **Initiator Name** field. The initiator name can contain up to 223 characters. • **Pools—**Choose an IQN suffix pool from which the name will be assigned. **Note** Setting the Initiator Name from the Set iSCSI Boot Parameters dialog box sets the initiator IQN at the iSCSI vNIC level and not at the service profile level. If more than one path is configured, you must set the initiator IQN from the **iSCSI vNICs** tab or when creating a service profile. If you need to, you can change or reset the initiator name. For more information, see Changing the Initiator IQN at the Service Profile Level, on page 543.
Create IQN Suffix Pool link	Click this link to create a new IQN suffix pool that will be available to all iSCSI vNICs.
Initiator Name field	A regular expression that defines the name of the iSCSI initiator. You can enter any alphanumeric string as well as the following special characters: • . (period) • : (colon) • - (dash)

Step 11 From the **Initiator IP Address Policy** drop-down list, change the selection to one of the following:

Option	Description
Select (DHCP used by default)	The system selects an interface automatically using DHCP. Proceed to Step 13.

Option	Description
Static	A static IPv4 address is assigned to the iSCSI boot vNIC based on the information entered in this area. Proceed to Step 12.
Pool	An IPv4 address is assigned to the iSCSI boot vNIC from the management IP address pool. Proceed to Step 13.

Step 12 If you chose **Static** from the **Initiator IP Address Policy** drop-down list, complete or change the following fields:

Name	Description
IPv4 Address field	The IPv4 address assigned to the iSCSI boot vNIC. If you want to specify this address, you must select **Static** in the **Initiator IP Address Policy** drop-down list.
Subnet Mask field	The subnet mask associated with the IPv4 address.
Default Gateway field	The default gateway associated with the IPv4 address.
Primary DNS field	The primary DNS server address.
Secondary DNS field	The secondary DNS server address.

Step 13 For the iSCSI target interface, choose one of the following radio buttons:

Option	Description
iSCSI Static Target Interface	The system creates a static target interface that you need to configure. Proceed to Step 14.
iSCSI Auto Target Interface	The system creates an auto target interface. You need to specify whether the auto target uses an initiator or a DCHP vendor ID. Proceed to Step 15.

Step 14 If you chose **iSCSI Static Target Interface**, do one of the following in the **Static Target Interface** table:

- To add an iSCSI static target interface, click **Add** or to modify an iSCSI target interface, select the iSCSI target interface that you want to change and click **Modify**. Then and complete or change the following fields in the **Create iSCSI Static Target** dialog box:

Name	Description
iSCSI Target Name field	A regular expression that defines the iSCSI Qualified Name (IQN) or Extended Unique Identifier (EUI) name of the iSCSI target. You can enter any alphanumeric characters as well as the following special characters: • . (period) • : (colon) • - (dash) **Important** This name must be properly formatted using standard IQN or EUI guidelines. The following examples show properly formatted iSCSI target names: • iqn.2001-04.com.example • iqn.2001-04.com.example:storage:diskarrays-sn-a8675309 • iqn.2001-04.com.example:storage.tape1.sys1.xyz • iqn.2001-04.com.example:storage.disk2.sys1.xyz • eui.02004567A425678D
Priority field	The system-assigned priority for the iSCSI target.
Port field	The port associated with the iSCSI target. Enter an integer between 1 and 65535. The default is 3260.
Authentication Profile drop-down list	The name of the associated iSCSI authentication profile.
Create iSCSI Authentication Profile link	Click this link to create a new iSCSI authentication profile that will be available to all iSCSI vNICs.
IPv4 Address field	The IPv4 address assigned to the iSCSI target.
LUN Id field	The LUN identifier in the iSCSI target.

- To delete an iSCSI target interface, select the iSCSI target interface that you want to delete and click **Delete**.

Note If you have two iSCSI static targets and you delete the first priority target, the second priority target becomes the first priority target, although Cisco UCS Manager still shows it as the second priority target.

Step 15 If you chose **iSCSI Auto Target Interface**, change the entry to either the initiator name or the DHCP vendor ID in the **DHCP Vendor Id** field. The initiator must have already been configured. The vendor ID can be up to 32 alphanumeric characters.

Step 16 Click **OK**.

IQN Pools

An IQN pool is a collection of iSCSI Qualified Names (IQNs) for use as initiator identifiers by iSCSI vNICs in a Cisco UCS domain.

IQN pool members are of the form *prefix*:*suffix*:*number*, where you can specify the prefix, suffix, and a block (range) of numbers.

An IQN pool can contain more than one IQN block, with different number ranges and different suffixes, but sharing the same prefix.

Creating an IQN Pool

Note In most cases, the maximum IQN size (prefix + suffix + additional characters) is 223 characters. When using the Cisco UCS NIC M51KR-B adapter, you must limit the IQN size to 128 characters.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **Pools**.

Step 3 Expand the node for the organization where you want to create the pool.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **IQN Pools** and select **Create IQN Suffix Pool**.

Step 5 In the **Define Name and Description** page of the **Create IQN Suffix Pool** wizard, fill in the following fields:

Field	Description
Name	The name of the iSCSI Qualified Name (IQN) pool. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description	The user-defined description of the pool. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Field	Description
Prefix	The prefix for any IQN blocks created for this pool. Enter from 1 to 150 characters. You can use any letter or number, as well as the special characters . (period), : (colon), and - (hyphen). For example, you could use iqn1.alpha.com.
Assignment Order field	This can be one of the following: • **Default—**Cisco UCS Manager selects a random identity from the pool. • **Sequential—**Cisco UCS Manager selects the lowest available identity from the pool.

Step 6 Click **Next**.

Step 7 In the **Add IQN Blocks** page of the **Create IQN Suffix Pool** wizard, click **Add**.

Step 8 In the **Create a Block of IQN Suffixes** dialog box, fill in the following fields:

Name	Description
Suffix field	The suffix for this bock of iSCSI Qualified Names (IQNs). Enter from 1 to 64 characters. You can use any letter or number, as well as the special characters . (period), : (colon), and - (hyphen). For example, you could use alphadc-1.
From field	The first suffix number in the block.
Size field	The number of suffixes in the block.

Step 9 Click **OK**.

Step 10 Click **Finish** to complete the wizard.

What to Do Next

Include the IQN suffix pool in a service profile and template.

Adding a Block to an IQN Pool

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **Pools**.

Step 3 Expand the node for the organization containing the pool.
If the system does not include multitenancy, expand the **root** node.

Step 4 Expand the **IQN Pools** node.

Step 5 Right-click the desired IQN pool and select **Create a Block of IQN Suffixes**.

Step 6 In the **Create a Block of IQN Suffixes** dialog box, fill in the following fields:

Name	Description
Suffix field	The suffix for this bock of iSCSI Qualified Names (IQNs). Enter from 1 to 64 characters. You can use any letter or number, as well as the special characters . (period), : (colon), and - (hyphen). For example, you could use alphadc-1.
From field	The first suffix number in the block.
Size field	The number of suffixes in the block.

Step 7 Click **OK**.

Deleting a Block from an IQN Pool

If you delete an address block from a pool, Cisco UCS Manager does not reallocate any addresses in that block that were assigned to vNICs or vHBAs. All assigned addresses from a deleted block remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **Pools**.

Step 3 Expand the node for the organization containing the pool.

If the system does not include multitenancy, expand the **root** node.

Step 4 Expand the **IQN Pools** node.

Step 5 Choose the IQN pool for which you want to delete a block of IQN suffixes.

Step 6 In the **Work pane**, click the **IQN Blocks** tab.

Step 7 Right-click the block to be deleted and select **Delete**.

Step 8 Click **Yes** to confirm the deletion.

Step 9 Click **Save Changes**.

Deleting an IQN Pool

If you delete a pool, Cisco UCS Manager does not reallocate any addresses from that pool that were assigned to vNICs or vHBAs. All assigned addresses from a deleted pool remain with the vNIC or vHBA to which they are assigned until one of the following occurs:

- The associated service profiles are deleted.
- The vNIC or vHBA to which the address is assigned is deleted.
- The vNIC or vHBA is assigned to a different pool.

Procedure

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **Pools**.

Step 3 Expand the node for the organization containing the pool.
If the system does not include multitenancy, expand the **root** node.

Step 4 Expand the **IQN Pools** node.

Step 5 Right-click the pool that you want to delete and select **Delete**.

Step 6 If a confirmation dialog box displays, click **Yes**.

LAN Boot

You can configure a boot policy to boot one or more servers from a centralized provisioning server on the LAN. A LAN (or PXE) boot is frequently used to install operating systems on a server from that LAN server.

You can add more than one type of boot device to a LAN boot policy. For example, you could add a local disk or virtual media boot as a secondary boot device.

Configuring a LAN Boot for a Boot Policy

You can also create a local boot policy that is restricted to a service profile or service profile template. However, Cisco recommends that you create a global boot policy that can be included in multiple service profiles or service profile templates.

You can add more than one type of boot device to a boot policy. For example, you can add a local disk or virtual media boot as a secondary boot device.

This procedure continues directly from Creating a Boot Policy, on page 527.

Procedure

Step 1 Click the down arrows to expand the **vNICs** area.

Step 2 Click the **Add LAN Boot** link.

Step 3 In the **Add LAN Boot** dialog box, enter the name of the vNIC that you want to use for the LAN boot in the **vNIC** field, then click **OK**.

Step 4 Do one of the following:

- Add another boot device to the **Boot Order** table.
- Click **OK** to finish.

What to Do Next

Include the boot policy in the service profile template.

After a server is associated with a service profile that includes this boot policy, you can verify the actual boot order in the **Boot Order Details** area on the **General** tab for the server.

Local Devices Boot

Cisco UCS Manager allows you to boot from different local devices.

Note For Cisco UCS M3 and M4 blade and rack servers using enhanced boot order, you can select both top-level and second-level boot devices. For Cisco UCS M1 and M2 blade and rack servers using standard boot order, you can only select a top-level device.

Local Disk Boot

If a server has a local drive, you can configure a boot policy to boot the server from the top-level local disk device or from any of the following second-level devices:

- Local LUN—Enables boot from local disk or local LUN.
- Local JBOD—Enables boot from a bootable JBOD.

- SD card—Enables boot from SD card.
- Internal USB—Enables boot for internal USB.
- External USB—Enables boot from external USB.
- Embedded Local LUN—Enables boot from the embedded local LUN on the Cisco UCS 240 M4 server.
- Embedded Local Disk—Enables boot from the embedded local disk on the Cisco UCS C240 M4SX and the M4L servers.

Note Second-level devices are only available for Cisco UCS M3 and M4 blade and rack servers using enhanced boot order. For Cisco UCS M1 and M2 blade and rack servers using standard boot order, you can choose only the top-level **Add Local Disk**.

Virtual Media Boot

You can configure a boot policy to boot one or more servers from a virtual media device that is accessible from the server. A virtual media device mimics the insertion of a physical CD/DVD disk (read-only) or floppy disk (read-write) into a server. This type of server boot is typically used to manually install operating systems on a server.

Note Second-level devices are only available for Cisco UCS M3 and M4 blade and rack servers using enhanced boot order. For Cisco UCS M1 and M2 blade and rack servers using standard boot order, you can choose only the top-level **Add CD/DVD** or **Add Floppy**.

Remote Virtual Drive Boot

You can configure a boot policy to boot one or more servers from a remote virtual drive that is accessible from the server.

Configuring a Local Disk Boot for a Boot Policy

You can also create a local boot policy that is restricted to a service profile or service profile template. However, Cisco recommends that you create a global boot policy that can be included in multiple service profiles or service profile templates.

You can add more than one type of boot device to a boot policy. For example, you can add an SD card boot as a secondary boot device.

This procedure continues directly from Creating a Boot Policy, on page 527.

Procedure

Step 1 Click the down arrows to expand the **Local Devices** area.

Step 2 Click any of the following links to add the device to the **Boot Order** table:

- **Add Local Disk** or

- **Add Local LUN**
- **Add SD Card**
- **Add Internal USB**
- **Add External USB**

Note For Cisco UCS M3 and M4 blade and rack servers using enhanced boot order, you can select both top-level and second-level boot devices. For Cisco UCS M1 and M2 blade and rack servers using standard boot order, you can only select a top-level device.

Step 3 Do one of the following:

- Add another boot device to the **Boot Order** table.
- Click **OK** to finish.

What to Do Next

Include the boot policy in a service profile and template.

After a server is associated with a service profile that includes this boot policy, you can verify the actual boot order in the **Boot Order Details** area on the **General** tab for the server.

Configuring a Virtual Media Boot for a Boot Policy

You can also create a local boot policy that is restricted to a service profile or service profile template. However, Cisco recommends that you create a global boot policy that can be included in multiple service profiles or service profile templates.

You can add more than one type of boot device to a boot policy. For example, you can add a local disk boot as a second boot device.

Note Virtual Media requires the USB to be enabled. If you modify the BIOS settings that affect the USB functionality, you also affect the Virtual Media. Therefore, Cisco recommends that you leave the following USB BIOS defaults for best performance:

- Make Device Non Bootable—set to **disabled**
- USB Idle Power Optimizing Setting—set to **high-performance**

This procedure continues directly from Creating a Boot Policy, on page 527.

Procedure

Step 1 Click the down arrows to expand the **Local Devices** area.

Step 2 Click any of the following links to add the device to the **Boot Order** table:

- **Add CD/DVD** or

- **Add Local CD/DVD**
- **Add Remote CD/DVD** (For KVM CD/DVD in rack servers)

- **Add Floppy** or
  - **Add Local Floppy**
  - **Add Remote Floppy**
- **Add Remote Virtual Drive**

Note For Cisco UCS M3 and M4 blade and rack servers using enhanced boot order, you can select both top-level and second-level boot devices. For Cisco UCS M1 and M2 blade and rack servers using standard boot order, you can only select a top-level device.

Step 3 Do one of the following:

- Add another boot device to the **Boot Order** table.
- Click **OK** to finish.

What to Do Next

Include the boot policy in a service profile and template.

After a server is associated with a service profile that includes this boot policy, you can verify the actual boot order in the **Boot Order Details** area on the **General** tab for the server.

Creating a vMedia Boot Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Right-click **Boot Policies** and select **Create Boot Policy**.

Step 4 In the **Create Boot Policy** dialog box, complete the following fields:

Name	Description
Name field	The name of the policy.
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Name	Description
Reboot on Boot Order Change check box	If checked, the system reboots the server if you make changes to the boot order. If this option is selected, the following action occurs: • If any **CDD/HDD** device is present in vMedia policy and same device is present in boot policy, any change in vmedia policy for that particular device causes the server to reboot. The purpose of this reboot is to boot from the latest mounted image. • If any **CDD/HDD** device is present in vMedia policy and same device is not present in boot policy, any change in vmedia policy for that particular device does not cause a server reboot. **Note** If this option is not selected, any modification in vMedia policy does not cause the host to reboot. If you add or change boot devices for a non-virtualized adapter, the system always reboots the server.
Boot Mode field	The type of boot mode that is enabled. This can be one of the following: • **Legacy—**Select if the system is not UEFI-enabled. • **Uefi—**Select if the system is UEFI-enabled. **Note** UEFI boot mode requires that the adapter firmware is UEFI-enabled, the boot device has a UEFI-aware operating system, and the service profile be associated with a Cisco UCS M3 blade or rack server.

Step 5 Expand the **CIMC Mounted vMedia** to add a virtual Remote vMedia Device to the boot order.

Step 6 Click the **Add Remote vMedia HDD or Add Remote vMedia CD/DVD** link.
Depending on the **vMedia** device you select, the **Add Remote vMedia** dialog box displays.

Step 7 Add the **CIMC Mounted CD/DVD** or **CIMC Mounted HDD** from the drop-down list.

Step 8 Click **OK**.
The vMedia Boot Policy is added to the **Boot Order** pane.

What to Do Next

Associate the vMedia and Boot Policies with a **Service Profile**.

Adding a Boot Policy to a vMedia Service Profile

This procedure describes how to set the boot policy options for vMedia on the **Server Boot Order** page of the **Create Service Profile (expert)** wizard.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization where you want to create the service profile.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the organization and select **Create Service Profile (expert)**.
The Unified Computing System Manager pane displays.

Step 5 In the **Name** field, enter a unique name that you can use to identify the service profile.
This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.

This name must be unique within the organization or sub-organization in which you are creating the service profile.

Step 6 From the **UUID Assignment** drop-down list, do one of the following:

Option	Description
Select (pool default used by default)	Assigns a UUID from the default UUID Suffix pool. Continue with Step 8.
Hardware Default	Uses the UUID assigned to the server by the manufacturer. If you choose this option, the UUID remains unassigned until the service profile is associated with a server. At that point, the UUID is set to the UUID value assigned to the server by the manufacturer. If the service profile is later moved to a different server, the UUID is changed to match the new server. Continue with Step 8.
XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX	Uses the UUID that you manually assign. Continue with Step 7.

Option	Description
Pools *Pool_Name*	Assigns a UUID from the UUID Suffix pool that you select from the list at the bottom of the drop-down list. Each pool name is followed by two numbers in parentheses that show the number of UUIDs still available in the pool and the total number of UUIDs in the pool. If you do not want use any of the existing pools, but instead want to create a pool that all service profiles can access, continue with Step 4; otherwise, continue with Step 8.

Step 7 (Optional) If you selected the **XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX** option, do the following:

a) In the **UUID** field, enter the valid UUID that you want to assign to the server which uses this service profile.

b) To verify that the selected UUID is available, click the **here** link.

Step 8 (Optional) If you want to create a new UUID Suffix pool to use to use in this service profile, click **Create UUID Suffix Pool** and complete the fields in the **Create UUID Suffix Pool** wizard.
For more information, see Creating a UUID Suffix Pool, on page 394.

Step 9 (Optional) In the text box, enter a description of this service profile.
The user-defined description for this service profile.

Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Step 10 Click **Next**.

Step 11 **Navigate** to **Create Service Profile (expert)** and click **Server Boot Order**.
The **Boot Policy** pane displays.

Step 12 From the **Boot Policy** drop-down list, choose one of the following:

Option	Description
Select Boot Policy to use	Assigns the default boot policy to this service profile. Continue with Step 13.
Create a Specific Boot Policy	Enables you to create a local boot policy that can only be accessed by this service profile.

Option	Description
Boot Policies *Policy_Name*	Assigns an existing boot policy to the service profile. If you choose this option, Cisco UCS Manager displays the details of the policy. If you do not want use any of the existing policies but instead want to create a policy that all service profiles can access, click **Create Boot Policy**; otherwise, choose a policy from the list and continue with Step 13.

Step 13 If you created a new boot policy accessible to all service profiles and template, choose that policy from the **Boot Policy** drop down list .

Step 14 Click **Next**.

What to Do Next

Associate your service profile with a Cisco UCS server.

Configuring an EFI Shell Boot for a Boot Policy

You can create a boot policy with an EFI Shell as the boot device. Booting from an EFI Shell prevents loss of data and provides more options to script, debug, and control various booting scenarios. EFI Shell is supported as a boot device only in the **Uefi** boot mode.

This procedure continues directly from Creating a Boot Policy, on page 527.

Before You Begin

To configure EFI Shell as a boot device, ensure that the boot mode is set to **Uefi**.

Important In an EFI Shell boot policy, If you edit the boot mode to **Legacy**, Cisco UCS Manager removes the EFI Shell boot device and sets the boot policy to default.

Procedure

Step 1 Select **Uefi** as the **Boot Mode**.

Step 2 Click the down arrows to expand the **EFI Shell** area.

Step 3 Click the **Add EFI Shell** link.
EFI Shell appears as a boot device in the **Boot Order** table

Step 4 Click **OK** to finish.

What to Do Next

Include the boot policy in a service profile and template.

Deleting a Boot Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies** > *Organization_Name*.

Step 3 Expand the **Boot Policies** node.

Step 4 Right-click the policy you want to delete and choose **Delete**.

Step 5 If a confirmation dialog box displays, click **Yes**.

UEFI Boot Parameters

UEFI boot mode for servers is dependent on information that is stored on the platform hardware. The boot entry, which contains information about the UEFI OS boot loader, is stored in the BIOS flash of the server. In Cisco UCS Manager releases earlier than Release 2.2(4), when a service profile is migrated from one server to another server, the boot loader information is not available on the destination server. Hence, the BIOS cannot load the boot loader information for the server to boot in UEFI boot mode.

Cisco UCSM Release 2.2(4) introduces UEFI boot parameters to provide the BIOS with information about the location of the UEFI OS boot loader on the destination server from where the BIOS loads it. Now, the server can use the boot loader information and boot in UEFI boot mode.

Guidelines and Limitations for UEFI Boot Parameters

- You can configure UEFI boot parameters only if the boot mode is UEFI.
- When you upgrade Cisco UCS Manager to Release 2.2(4), UEFI boot failure during service profile migration is not handled automatically. You must explicitly create the UEFI boot parameters in the target device to successfully boot to the UEFI-capable OS.
- UEFI boot parameters are supported on all M3 and higher servers that support second-level boot order.
- You can specify UEFI boot parameters for the following device types:
  - SAN LUN
  - ISCSI LUN
  - Local LUN
- UEFI boot parameters are specific to each operating system. You can specify UEFI boot parameters for the following operating systems:
  - VMware ESX
  - SuSE Linux

- Microsoft Windows
- Red Hat Enterprise Linux 7

Setting UEFI Boot Parameters

Before You Begin

Ensure that the **Boot Mode** of the boot policy is **Uefi**.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand **Boot Policies** and select the boot policy for which you want to configure UEFI boot parameters.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 To set UEFI boot parameters for a LUN, select the LUN in the **Boot Order** area and click **Set Uefi Boot Parameters**.

Important You can configure UEFI boot parameters only for local LUNs, SAN LUNs, and iSCSI LUNs.

Step 6 In the **Set Uefi Boot Parameters** dialog box, enter the following information:

Field	Description
Boot Loader Name	Specifies the name of the boot loader. This is a mandatory field.
Boot Loader Path	Specifies the path where the boot loader is located. This is a mandatory field.
Boot Loader Description	Describes the boot loader.

Step 7 Click **OK**

Step 8 Click **Save Changes**.

Modifying UEFI Boot Parameters

Before You Begin

Ensure that the **Boot Mode** of the boot policy is **Uefi**.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand **Boot Policies**, and select the boot policy for which you want to modify UEFI boot parameters.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 To modify UEFI boot parameters for a LUN with UEFI boot parameters, select the LUN in the **Boot Order** area and click **Modify Uefi Boot Parameters**.

Important You can configure UEFI boot parameters only for local LUNs, SAN LUNs, and iSCSI LUNs.

Step 6 In the **Modify Uefi Boot Parameters** dialog box, enter the following information:

Field	Description
Boot Loader Name	Specifies the name of the boot loader. This is a mandatory field.
Boot Loader Path	Specifies the path where the boot loader is located. This is a mandatory field.
Boot Loader Description	Describes the boot loader.

Step 7 Click **OK**

Step 8 Click **Save Changes**.

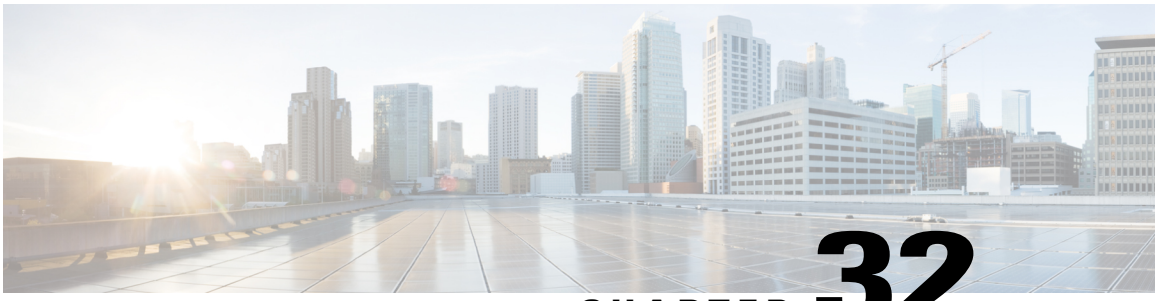

CHAPTER 32

Deferring Deployment of Service Profile Updates

This chapter includes the following sections:

Service Profile Deferred Deployments

Some modifications to a service profile or to an updating service profile template can be disruptive and require a reboot of the server. You can, however, configure deferred deployment to control when those disruptive configuration changes are implemented. For example, you can choose to deploy the service profile changes immediately or have them deployed during a specified maintenance window. You can also choose whether or not a service profile deployment requires explicit user acknowledgment.

Deferred deployment is available for all configuration changes that occur through the association of a service profile with a server. These configuration changes can be prompted by a change to a service profile, to a policy that is included in a service profile, or to an updating service profile template. For example, you can defer the upgrade and activation of firmware through host firmware packages and management firmware packages, such as server BIOS, RAID controller, host HBA, and network adapters. However, you cannot defer the direct deployment of firmware images for components that do not use either of the firmware packages, such as Cisco UCS Manager, fabric interconnects, and I/O modules.

Deferred deployment is not available for the following actions which require the reboot of a server:

- Initial association of a service profile with a server
- Final disassociation of a service profile from a server, without associating the service profile with a different server
- Decommissioning a server
- Re-acknowledging a server
- Resetting a server

If you want to defer the deployment of service profile changes, you must configure one or more maintenance policies and configure each service profile with a maintenance policy. If you want to define the time period when the deployment should occur, you also need to create at least one schedule with one or more recurring occurrences or one time occurrences, and include that schedule in a maintenance policy.

Schedules for Deferred Deployments

A schedule contains a set of occurrences. These occurrences can be one time only or can recur at a specified time and day each week. The options defined in the occurrence, such as the duration of the occurrence or the maximum number of tasks to be run, determine whether a service profile change is deployed. For example, if a change cannot be deployed during a given maintenance window because the maximum duration or number of tasks was reached, that deployment is carried over to the next maintenance window.

Each schedule checks periodically to see whether the Cisco UCS domain entered one or more maintenance windows. If so, the schedule executes the deployments that are eligible according to the constraints specified in the maintenance policy.

A schedule contains one or more occurrences, which determine the maintenance windows associated with that schedule. An occurrence can be one of the following:

One Time Occurrence

One time occurrences define a single maintenance window. These windows continue until the maximum duration of the window or the maximum number of tasks that can be run in the window is reached.

Recurring Occurrence

Recurring occurrences define a series of maintenance windows. These windows continue until the maximum number of tasks or the end of the day specified in the occurrence was reached.

Maintenance Policy

A maintenance policy determines how Cisco UCS Manager reacts when a change that requires a server reboot is made to a service profile associated with a server or to an updating service profile bound to one or more service profiles.

The maintenance policy specifies how Cisco UCS Manager deploys the service profile changes. The deployment can occur in one of the following ways:

- Immediately
- When acknowledged by a user with administrator privileges
- Automatically at the time specified in a schedule
- On the next reboot or shutdown without waiting for the user acknowledgment or the timer scheduling option

Note If the **On Next Boot** option is enabled in a maintenance policy, and you downgrade from Cisco UCS Manager Release 3.1(1) or later releases to any release earlier than Cisco UCS Manager Release 2.2(8), firmware downgrade will fail. Disable **On Next Boot** from the maintenance policy to continue with the downgrade.

You can use the soft shutdown timer in the maintenance policy to configure the wait time for performing a hard shutdown. The soft shutdown timer is applicable when you reboot the server for the following:

- Reset the server using the **Gracefully Restart OS** option.
- Shut down the server with the **In case of graceful shutdown failure, a hard shutdown will be issued after X seconds** option.
- Modify a service profile that requires a server reboot.

If the maintenance policy is configured to deploy the change during a scheduled maintenance window, the policy must include a valid schedule. The schedule deploys the changes in the first available maintenance window.

Note A maintenance policy only prevents an immediate server reboot when a configuration change is made to an associated service profile. However, a maintenance policy does not prevent the following actions from taking place right away:

- Deleting an associated service profile from the system
- Disassociating a server profile from a server
- Directly installing a firmware upgrade without using a service policy
- Resetting the server

Pending Activities for Deferred Deployments

If you configure a deferred deployment in a Cisco UCS domain, Cisco UCS Manager enables you to view all pending activities. You can see activities that are waiting for user acknowledgement and those that are scheduled.

If a Cisco UCS domain has pending activities, Cisco UCS Manager GUI notifies users with admin privileges when they log in.

Cisco UCS Manager displays information about all pending activities, including the following:

- Name of the service profile to deploy and associate with a server
- Server affected by the deployment
- Disruption caused by the deployment
- Change performed by the deployment

Note You cannot specify the maintenance window in which a specific pending activity is applied to the server. The maintenance window depends upon how many activities are pending and which maintenance policy is assigned to the service profile. However, any user with admin privileges can manually initiate a pending activity and reboot the server immediately, whether it is waiting for user acknowledgment or for a maintenance window.

Guidelines and Limitations for Deferred Deployments

Service Profile Association Changes and Maintenance Policy Options

When changing service profile association, the following maintenance policy options can affect how the changes are applied:

- If the **On Next Boot** and **User Ack** options are enabled in a maintenance policy, the service profile association change displays a warning that an acknowledgement is required. However, association will happen immediately.
- If the **On Next Boot** and **User Ack** options are not enabled in a maintenance policy, the service profile association change displays a warning that an acknowledgement is required, and will remain pending until acknowledged.

Cannot Undo All Changes to Service Profiles or Service Profile Templates

If you cancel a pending change, Cisco UCS Manager attempts to roll back the change without rebooting the server. However, for complex changes, Cisco UCS Manager may have to reboot the server a second time to roll back the change. For example, if you delete a vNIC, Cisco UCS Manager reboots the server according to the maintenance policy included in the service profile. You cannot cancel this reboot and change, even if you restore the original vNIC in the service profile. Instead, Cisco UCS Manager schedules a second deployment and reboot of the server.

Association of Service Profile Can Exceed Boundaries of Maintenance Window

After Cisco UCS Manager begins the association of the service profile, the scheduler and maintenance policy do not have any control over the procedure. If the service profile association does not complete within the allotted maintenance window, the process continues until it is completed. For example, this can occur if the association does not complete in time because of retried stages or other issues.

Cannot Specify Order of Pending Activities

Scheduled deployments run in parallel and independently. You cannot specify the order in which the deployments occur. You also cannot make the deployment of one service profile change dependent upon the completion of another.

Cannot Perform Partial Deployment of Pending Activity

Cisco UCS Manager applies all changes made to a service profile in the scheduled maintenance window. You cannot make several changes to a service profile at the same time and then have those changes be spread across several maintenance windows. When Cisco UCS Manager deploys the service profile changes, it updates the service profile to match the most recent configuration in the database.

Configuring Schedules

Creating a Schedule

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 On the **Servers** tab, right-click **Schedules** and choose **Create Schedule**.

Step 3 In the **Identify Schedule** page of the **Create Schedule** wizard, complete the following fields:

Name	Description
Name field	The name of the schedule. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the schedule. We recommend including information about where and when the schedule should be used. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Owner field	The owner of the schedule. This can be one of the following: • **Local—**Cisco UCS Manager owns the schedule, which is configured in this Cisco UCS domain. • **Pending Global—**Cisco UCS Manager is in the process of transferring this schedule to Cisco UCS Central. • **Global—**Cisco UCS Central owns the schedule, which is configured on a remote server.

Step 4 Click **Next**.

Step 5 On the **One Time Occurrences** page, click one of the following:

Option	Description
Next	Moves to the next page. Choose this option if you do not want to create a one time occurrence for this schedule. If you choose this option, continue with Step 8.

Option	Description
Add	Opens the **Create a One Time Occurrence** dialog box, where you can specify a single time when this schedule should be run. If you choose this option, continue with Step 6.

Step 6 (Optional) In the **Create a One Time Occurrence** dialog box, do the following:

a) Complete the following fields:

Name	Description
Name field	The name of the one time occurrence of this schedule. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Start Time field	The date and time that the occurrence will run. Click the down arrow at the end of the field to select the date from a calendar.

b) Click the down arrows to expand the **Options** area.

c) In the **Options** area, complete the following fields:

Name	Description
Max Duration field	The maximum length of time that the scheduled occurrence can run. This can be one of the following: • **None—**The occurrence runs until all tasks are completed. • **other—**Cisco UCS Manager GUI displays the **dd:hh:mm:ss** field allowing you to specify the maximum amount of time that the occurrence can run. Cisco UCS completes as many scheduled tasks as possible within the specified time. By default, the maximum duration is set to **none**. If you do not change this setting and you do not set a maximum number of tasks, the maintenance window continues until all pending activities are completed.

Name	Description
Max Number of Tasks field	The maximum number of scheduled tasks that can be run during this occurrence. This can be one of the following: • **Unlimited—**Cisco UCS runs all scheduled tasks unless those tasks exceed the maximum time specified in the **Max Duration** field. If **Max Duration** is set to **none** and you select this option, the maintenance window continues until all pending activities are completed. • **other—**Cisco UCS Manager GUI displays a text field allowing you to specify the maximum number of tasks that can be run during this occurrence. Enter an integer between 1 and 65535. **Note** This option does not apply if this schedule is associated with a fault suppression task.
Max Number of Concurrent Tasks field	The maximum number of tasks that can run concurrently during this occurrence. This can be one of the following: • **Unlimited—**Cisco UCS runs as many concurrent tasks as the system can handle. • **other—**Cisco UCS Manager GUI displays a text field allowing you to specify the maximum number of concurrent tasks that can be run during this occurrence. Enter an integer between 1 and 65535. **Note** This option does not apply if this schedule is associated with a fault suppression task.
Minimum Interval Between Tasks field	The minimum length of time that the system should wait before starting a new task. This setting is meaningful only if the maximum number of concurrent tasks is set to a value other than None. This can be one of the following: • **None—**Cisco UCS runs the next task as soon as possible. • **other—**Cisco UCS Manager GUI displays the **dd:hh:mm:ss** field allowing you to specify the minimum amount of time that Cisco UCS will wait between tasks. **Note** This option does not apply if this schedule is associated with a fault suppression task.

d) Click **OK**.

Step 7 To add another one time occurrence, click **Add** and repeat step 6. Otherwise, click **Next**.

Step 8 (Optional) If you want to define a recurring occurrence for this schedule, on the **Recurring Occurrences** page, click **Add**.

a) In the **Create a Recurring Occurrence** dialog box, complete the following fields:

Name	Description
Name field	The name of the recurring occurrence of this schedule. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Day field	The day on which Cisco UCS runs an occurrence of this schedule. This can be one of the following: • **every day** • **Monday** • **Tuesday** • **Wednesday** • **Thursday** • **Friday** • **Saturday** • **Sunday** • **odd days** • **even days**
Hour field	The hour of the specified day at which this occurrence of the schedule starts. This can be an integer between 0 and 24, where 0 and 24 are both equivalent to midnight. **Note** Cisco UCS ends all recurring occurrences on the same day in which they start, even if the maximum duration has not been reached. For example, if you specify a start time of 11 p.m. and a maximum duration of 3 hours, Cisco UCS starts the occurrence at 11 p.m. but ends it at 11:59 p.m. after only 59 minutes. Ensure that the start time you specify is early enough so that the recurring occurrence finishes before 11:59 p.m.
Minute field	The minute of the hour at which the schedule occurrence starts. This can be an integer between 0 and 60.

b) Click the down arrows to expand the **Options** area.
c) In the **Options** area, complete the following fields:

Name	Description
Max Duration field	The maximum length of time that each occurrence of this schedule can run. This can be one of the following: • **None—**The occurrence runs until all tasks are completed. • **other—**Cisco UCS Manager GUI displays the **dd:hh:mm:ss** field allowing you to specify the maximum amount of time that the occurrence can run. Cisco UCS completes as many scheduled tasks as possible within the specified time.
Max Number of Tasks field	The maximum number of scheduled tasks that can be run during each occurrence. This can be one of the following: • **Unlimited—**Cisco UCS runs all scheduled tasks unless those tasks exceed the maximum time specified in the **Max Duration** field. If **Max Duration** is set to **none** and you select this option, the maintenance window continues until all pending activities are completed. • **other—**Cisco UCS Manager GUI displays a text field allowing you to specify the maximum number of tasks that can be run during this occurrence. Enter an integer between 1 and 65535. **Note** This option does not apply if this schedule is associated with a fault suppression task.
Max Number of Concurrent Tasks field	The maximum number of tasks that can run concurrently during each occurrence. This can be one of the following: • **Unlimited—**Cisco UCS runs as many concurrent tasks as the system can handle. • **other—**Cisco UCS Manager GUI displays a text field allowing you to specify the maximum number of concurrent tasks that can be run during this occurrence. Enter an integer between 1 and 65535. **Note** This option does not apply if this schedule is associated with a fault suppression task.

Name	Description
Minimum Interval Between Tasks field	The minimum length of time that the system should wait before starting a new task. This setting is meaningful only if the maximum number of concurrent tasks is set to a value other than None. This can be one of the following: • **None—**Cisco UCS runs the next task as soon as possible. • **other—**Cisco UCS Manager GUI displays the **dd:hh:mm:ss** field allowing you to specify the minimum amount of time that Cisco UCS will wait between tasks. **Note** This option does not apply if this schedule is associated with a fault suppression task.

d) Click **OK**.
e) To add another recurring occurrence, click **Add** and repeat this step.

Step 9 Click **Finish**.

Creating a One Time Occurrence for a Schedule

Note By default, the maximum duration and the maximum number of tasks are set to **none**. If you do not change either of these defaults, Cisco UCS Manager does not impose any limit to the length of time that the maintenance window lasts. All pending activities are applied as soon as the scheduled maintenance window begins, and Cisco UCS Manager continues to reboot the servers impacted by the pending activities until all of those tasks are complete.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Schedules**.
Step 3 Right-click the schedule to which you want to add an occurrence and choose **Create a One Time Occurrence**.
Step 4 In the **Create a One Time Occurrence** dialog box, complete the following fields:

Name	Description
Name field	The name of the one time occurrence of this schedule. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Name	Description
Start Time field	The date and time that the occurrence will run. Click the down arrow at the end of the field to select the date from a calendar.

Step 5 Click the down arrows to expand the **Options** area.

Step 6 In the **Options** area, complete the following fields:

Name	Description
Max Duration field	The maximum length of time that the scheduled occurrence can run. This can be one of the following: • **None—**The occurrence runs until all tasks are completed. • **other—**Cisco UCS Manager GUI displays the **dd:hh:mm:ss** field allowing you to specify the maximum amount of time that the occurrence can run. Cisco UCS completes as many scheduled tasks as possible within the specified time. By default, the maximum duration is set to **none**. If you do not change this setting and you do not set a maximum number of tasks, the maintenance window continues until all pending activities are completed.
Max Number of Tasks field	The maximum number of scheduled tasks that can be run during this occurrence. This can be one of the following: • **Unlimited—**Cisco UCS runs all scheduled tasks unless those tasks exceed the maximum time specified in the **Max Duration** field. If **Max Duration** is set to **none** and you select this option, the maintenance window continues until all pending activities are completed. • **other—**Cisco UCS Manager GUI displays a text field allowing you to specify the maximum number of tasks that can be run during this occurrence. Enter an integer between 1 and 65535. **Note** This option does not apply if this schedule is associated with a fault suppression task.

Name	Description
Max Number of Concurrent Tasks field	The maximum number of tasks that can run concurrently during this occurrence. This can be one of the following: • **Unlimited—**Cisco UCS runs as many concurrent tasks as the system can handle. • **other—**Cisco UCS Manager GUI displays a text field allowing you to specify the maximum number of concurrent tasks that can be run during this occurrence. Enter an integer between 1 and 65535. **Note** This option does not apply if this schedule is associated with a fault suppression task.
Minimum Interval Between Tasks field	The minimum length of time that the system should wait before starting a new task. This setting is meaningful only if the maximum number of concurrent tasks is set to a value other than None. This can be one of the following: • **None—**Cisco UCS runs the next task as soon as possible. • **other—**Cisco UCS Manager GUI displays the **dd:hh:mm:ss** field allowing you to specify the minimum amount of time that Cisco UCS will wait between tasks. **Note** This option does not apply if this schedule is associated with a fault suppression task.

Step 7 Click **OK**.

Creating a Recurring Occurrence for a Schedule

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Schedules**.

Step 3 Right-click the schedule to which you want to add an occurrence and choose **Create a Recurring Occurrence**.

Step 4 In the **Create a Recurring Occurrence** dialog box, complete the following fields:

Name	Description
Name field	The name of the recurring occurrence of this schedule. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Day field	The day on which Cisco UCS runs an occurrence of this schedule. This can be one of the following: • **every day** • **Monday** • **Tuesday** • **Wednesday** • **Thursday** • **Friday** • **Saturday** • **Sunday** • **odd days** • **even days**
Hour field	The hour of the specified day at which this occurrence of the schedule starts. This can be an integer between 0 and 24, where 0 and 24 are both equivalent to midnight. **Note** Cisco UCS ends all recurring occurrences on the same day in which they start, even if the maximum duration has not been reached. For example, if you specify a start time of 11 p.m. and a maximum duration of 3 hours, Cisco UCS starts the occurrence at 11 p.m. but ends it at 11:59 p.m. after only 59 minutes. Ensure that the start time you specify is early enough so that the recurring occurrence finishes before 11:59 p.m.
Minute field	The minute of the hour at which the schedule occurrence starts. This can be an integer between 0 and 60.

Step 5 Click the down arrows to expand the **Options** area.

Step 6 In the **Options** area, complete the following fields:

Name	Description
Max Duration field	The maximum length of time that each occurrence of this schedule can run. This can be one of the following: • **None—**The occurrence runs until all tasks are completed. • **other—**Cisco UCS Manager GUI displays the **dd:hh:mm:ss** field allowing you to specify the maximum amount of time that the occurrence can run. Cisco UCS completes as many scheduled tasks as possible within the specified time.
Max Number of Tasks field	The maximum number of scheduled tasks that can be run during each occurrence. This can be one of the following: • **Unlimited—**Cisco UCS runs all scheduled tasks unless those tasks exceed the maximum time specified in the **Max Duration** field. If **Max Duration** is set to **none** and you select this option, the maintenance window continues until all pending activities are completed. • **other—**Cisco UCS Manager GUI displays a text field allowing you to specify the maximum number of tasks that can be run during this occurrence. Enter an integer between 1 and 65535. **Note** This option does not apply if this schedule is associated with a fault suppression task.
Max Number of Concurrent Tasks field	The maximum number of tasks that can run concurrently during each occurrence. This can be one of the following: • **Unlimited—**Cisco UCS runs as many concurrent tasks as the system can handle. • **other—**Cisco UCS Manager GUI displays a text field allowing you to specify the maximum number of concurrent tasks that can be run during this occurrence. Enter an integer between 1 and 65535. **Note** This option does not apply if this schedule is associated with a fault suppression task.

Name	Description
Minimum Interval Between Tasks field	The minimum length of time that the system should wait before starting a new task. This setting is meaningful only if the maximum number of concurrent tasks is set to a value other than None. This can be one of the following: • **None**—Cisco UCS runs the next task as soon as possible. • **other**—Cisco UCS Manager GUI displays the **dd:hh:mm:ss** field allowing you to specify the minimum amount of time that Cisco UCS will wait between tasks. **Note** This option does not apply if this schedule is associated with a fault suppression task.

Step 7 Click **OK**.

Deleting a One Time Occurrence from a Schedule

If this is the only occurrence in a schedule, that schedule is reconfigured with no occurrences. If the schedule is included in a maintenance policy and that policy is assigned to a service profile, any pending activities related to the server associated with the service profile cannot be deployed. You must add a one time occurrence or a recurring occurrence to the schedule to deploy the pending activity.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Schedules** > ***Schedule_Name***.
Step 3 Expand **One Time Occurrences**.
Step 4 Right-click the occurrence you want to delete and choose **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Deleting a Recurring Occurrence from a Schedule

If this is the only occurrence in a schedule, that schedule is reconfigured with no occurrences. If the schedule is included in a maintenance policy and that policy is assigned to a service profile, any pending activities related to the server associated with the service profile cannot be deployed. You must add a one time occurrence or a recurring occurrence to the schedule to deploy the pending activity.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Schedules** > ***Schedule_Name***.
Step 3 Expand **Recurring Occurrences**.
Step 4 Right-click the occurrence you want to delete and choose **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Deleting a Schedule

If this schedule is included in a maintenance policy, the policy is reconfigured with no schedule. If that policy is assigned to a service profile, any pending activities related to the server associated with the service profile cannot be deployed. You must add a schedule to the maintenance policy to deploy the pending activity.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Schedules**.
Step 3 Right-click the schedule you want to delete and choose **Delete**.
Step 4 If a confirmation dialog box displays, click **Yes**.

Configuring Maintenance Policies

Creating a Maintenance Policy

Before You Begin

If you plan to configure this maintenance policy for automatic deferred deployment, create a schedule.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Policies**.
Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.
Step 4 Right-click **Maintenance Policies** and choose **Create Maintenance Policy**.
Step 5 In the **Create Maintenance Policy** dialog box, complete the following fields:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Soft Shutdown Timer drop-down list	This timer allows you to specify the time in seconds when Cisco UCS Manager performs a server shut down and reboot. Cisco UCS Manager waits until the specified time in the maintenance policy before performing a hard shut down. This can be one of the following: • **150 Secs—**Cisco UCS Manager waits until 150 seconds before performing a hard shut down and reboot of the server. • **300 Secs—**Cisco UCS Manager waits until 300 seconds before performing a hard shut down and reboot of the server. • **600 Secs—**Cisco UCS Manager waits for 600 seconds before performing a hard shut down and reboot of the server. • **Never—**Cisco UCS Manager never performs a server shut down.

Name	Description
Reboot Policy field	When a service profile is associated with a server, or when changes are made to a service profile that is already associated with a server, you must reboot the server to complete the process. The **Reboot Policy** field determines when the reboot occurs for servers associated with any service profiles that include this maintenance policy. This can be one of the following: • **Immediate—**The server reboots automatically as soon as the service profile association is complete or when you save service profile changes. • **User Ack—**You must explicitly acknowledge the pending activities for the changes made to the service profile to be applied to the associated server. • **Timer Automatic—**Cisco UCS defers all service profile associations and changes until the maintenance window defined by the schedule shown in the **Schedule** field. • **On Next Boot—**This option is used in combination with either **User Ack** or **Timer Automatic**. When the **On Next Boot** option is enabled, the host OS reboot, shutdown, and reset, or server reset and shutdown also triggers the associated FSM to apply the changes waiting for the**User Ack**, or the **Timer Automatic** maintenance window. **Note** De-selecting the On Next Boot option disables the Maintenance Policy on the BMC.
Schedule drop-down list	If the **Reboot Policy** is set to **Timer Automatic**, the schedule specifies when maintenance operations can be applied to the server. Cisco UCS reboots the server and completes the service profile changes at the scheduled time.
Create Schedule link	Creates a new schedule that is available to all objects in this Cisco UCS domain.

Step 6 Click **OK**.

What to Do Next

Include the policy in a service profile or service profile template.

Deleting a Maintenance Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Policies** > *Organization_Name*.
Step 3 Expand **Maintenance Policies**.
Step 4 Right-click the maintenance policy you want to delete and choose **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Managing Pending Activities

Viewing Pending Activities

Procedure

Step 1 On the toolbar, click **Pending Activities**.
Step 2 Click one of the following tabs:

- **User Acknowledged Activities—**Contains the **Service Profiles** and **Fabric Interconnects** tabs that display the tasks requiring user acknowledgment before they can complete.
- **Scheduled Activities—**Displays the tasks that will be performed based on the associated maintenance schedule.

Step 3 Click a row in the table to view the details of that pending activity.
If you click the link in the **Server** column, Cisco UCS Manager displays the properties of that server.

Deploying a Service Profile Change Waiting for User Acknowledgement

Important You cannot stop Cisco UCS Manager from rebooting the affected server after you acknowledge a pending activity.

Procedure

Step 1 On the toolbar, click **Pending Activities**.

Step 2 In the **Pending Activities** dialog box, click the **User Acknowledged Activities** tab, then the **Service Profiles** tab.

Step 3 Check the check box in the **Reboot Now** column for each pending activity that you want to deploy immediately.

Step 4 Click **OK**.
Cisco UCS Manager immediately reboots the server affected by the pending activity.

Deploying All Service Profile Changes Waiting for User Acknowledgement

Important You cannot stop Cisco UCS Manager from rebooting the affected server after you acknowledge a pending activity.

Procedure

Step 1 On the toolbar, click **Pending Activities**.

Step 2 In the **Pending Activities** dialog box, click the **User Acknowledged Activities** tab and then the **Service Profiles** tab.

Step 3 In the toolbar, check the **Acknowledge All** check box.
Cisco UCS Manager GUI checks the **Reboot Now** check boxes for all pending activities listed in the table.

Step 4 Click **OK**.
Cisco UCS Manager immediately reboots all servers affected by the pending activities listed in the table.

Deploying a Scheduled Service Profile Change Immediately

Important You cannot stop Cisco UCS Manager from rebooting the affected server after you acknowledge a pending activity.

Procedure

Step 1 On the toolbar, click **Pending Activities**.

Step 2 In the **Pending Activities** dialog box, click the **Scheduled Activities** tab.

Step 3 Check the check box in the **Reboot Now** column for each pending activity you want to deploy immediately.

Step 4 Click **OK**.
Cisco UCS Manager immediately reboots the server affected by the pending activity.

Deploying All Scheduled Service Profile Changes Immediately

Important You cannot stop Cisco UCS Manager from rebooting the affected server after you acknowledge a pending activity.

Procedure

Step 1 On the toolbar, click **Pending Activities**.

Step 2 In the **Pending Activities** dialog box, click the **Scheduled Activities** tab.

Step 3 In the toolbar, check the **Acknowledge All** check box.
Cisco UCS Manager GUI checks the **Reboot Now** check boxes for all pending activities listed in the table.

Step 4 Click **OK**.
Cisco UCS Manager immediately reboots all servers affected by the pending activities listed in the table.

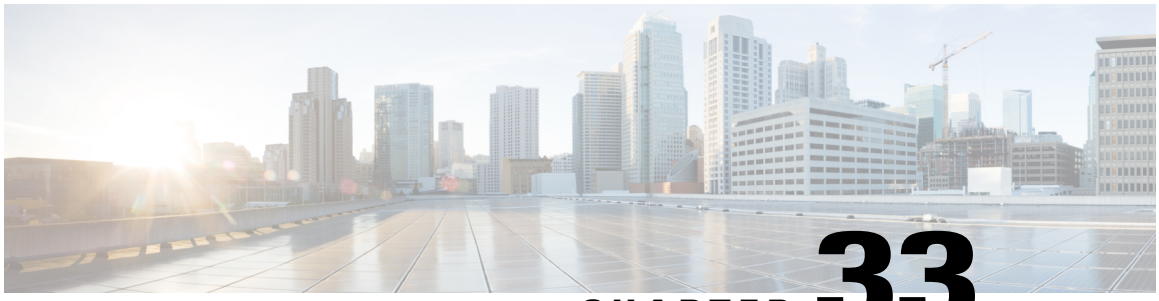

CHAPTER 33

Service Profiles

This chapter includes the following sections:

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 pres**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.

Guidelines and Recommendations for Service Profiles

In addition to any guidelines or recommendations that are specific to policies and pools included in service profiles and service profile templates, such as the local disk configuration policy, adhere to the following guidelines and recommendations that impact the ability to associate a service profile with a server:

Limit to the Number of vNICs that Can Be Configured on a Rack-Mount Server

You can configure up to 56 vNICs per supported adapter, such as the Cisco UCS P81E Virtual Interface Card (N2XX-ACPCI01), on any rack-mount server that is integrated with Cisco UCS Manager.

No Power Capping Support for Rack-Mount Servers

Power capping is not supported for rack servers. If you include a power control policy in a service profile that is associated with a rack-mount server, the policy is not implemented.

QoS Policy Guidelines for vNICs

You can only assign a QoS policy to a vNIC if the priority setting for that policy is not set to **fc**, which represents the Fibre Channel system class. You can configure the priority for the QoS policy with any other system class.

QoS Policy Guidelines for vHBAs

You can only assign a QoS policy to a vHBA if the priority setting for that policy is set to **fc**, which represents the Fibre Channel system class.

The Host Control setting for a QoS policy applies to vNICs only. It has no effect on a vHBA.

Creating Service Profiles

Creating a Service Profile with the Expert Wizard

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization where you want to create the service profile.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the organization and select **Create Service Profile (expert)**.

Step 5 In the **Identify Service Profile** panel, specify the service profile **Name**, UUID assignment and click **Next**.
You can provide an optional description for this service profile. If the UUID is not available, you can also create a UUID Suffix Pool from this panel.

Note To create a service profile quickly, you can click **Finish** after specifying the name. Cisco UCS Manager creates a new service profile with the specified name and all system default values.

Step 6 (Optional) In the **Networking** panel, specify the required information for the **Dynamic vNIC Connection Policy** and **LAN Connectivity** sections, then click **Next**.
You can create a dynamic vNIC connection policy and LAN connectivity policy from this panel.

Step 7 (Optional) In the **Storage** panel, specify the SAN configuration information such as, **Local Storage Policy**, **SAN Connectivity**, **WWNN** and **VSAN**, then click **Next**.
You can create a local disk configuration policy and SAN connectivity policy from this panel.

Step 8 (Optional) In the **Zoning** panel, specify the required zoning information, then click**Next**.
You can create the vHBA initiator groups from this panel.

Step 9 (Optional) In the **vNIC/vHBA Placement** panel, specify the placement method and PCI order, then click**Next**.
You can create a placement policy from this panel.

Step 10 (Optional) In the **Server Boot Order** panel, specify the **Boot Policy** from the drop-down list, then click **Next**.
You can create a boot policy from this panel.

Step 11 (Optional) In the **Maintenance Policy** panel, specify the maintenance policy, then click **Next**.
You can create a new maintenance policy and specify a maintenance schedule from this panel.

Step 12 (Optional) In the **Server Assignment** panel, specify the **Server Assignment** from the drop down list and the power state to apply on assignment, then click **Next**.
You can create a server pool or a host firmware package from this panel.

Step 13 (Optional) In the **Operational Policies** panel, specify the system operational information such as, **BIOS Configuration**, **External IPMI Management Configuration**, **Management IP Address**, **Monitoring Configuration (Thresholds)**, **Power Control Policy Configuration**, and **Scrub Policy**, then click **Finish**.

Note To set up an Outband IPv4 address or an Inband IPv4 or IPv6 address, click the respective tabs and complete the required fields.

If you do not find the policies you need for each of these configurations, you can create them from this panel.

Creating a Service Profile that Inherits Server Identity

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization where you want to create the service profile.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the organization and select **Create Service Profile**.

Step 5 In the **Naming** area of the **Create Service Profile** dialog box, complete the following fields:

a) In the **Name** field, enter a unique name that you can use to identify the service profile.
This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.

b) In the **Description** field, enter a description of this service profile.

Step 6 In the **vNICs** area of the **Create Service Profile** dialog box, choose the primary and secondary vNICs.

Step 7 In the **vHBAs** area of the **Create Service Profile** dialog box, choose the primary and secondary vHBAs.

Step 8 In the **Boot Order** area of the **Create Service Profile** dialog box, choose the primary and secondary boot devices.

Step 9 (Optional) In the **Select** column of the **Server Association (optional)** area, click the radio button for a server to associate this service profile with that server.

Step 10 Click **OK**.

Creating a Hardware Based Service Profile for a Blade Server

You cannot move a hardware based service profile to another server.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.

Step 3 Choose the server for which you want to create a hardware based service profile.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Create Service Profile**.

Step 6 In the **Create Service Profile for Server** dialog box, do the following:

a) From the **Create Service Profile in Organization** drop-down list, select the organization in which you want to create the service profile.

b) Click the **Hardware Based Service Profile** radio button.
c) In the **Name** field, enter a unique name for the service profile.
This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.
d) If you want Cisco UCS Manager to create vNICs for the service profile, check the **Create Default vNICs** check box.
e) If you want Cisco UCS Manager to create vHBAs for the service profile, check the **Create Default vHBAs** check box.
f) Click **OK**.

Cisco UCS Manager inherits and automatically applies the identity and configuration information in the server, creates the service profile, and associates it with the server.

Creating a Hardware Based Service Profile for a Rack-Mount Server

You cannot move a hardware based service profile to another server.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Rack Mounts** > **Servers**.
Step 3 Choose the server for which you want to create a hardware based service profile.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Create Service Profile**.
Step 6 In the **Create Service Profile for Server** dialog box, do the following:
a) From the **Create Service Profile in Organization** drop-down list, select the organization in which you want to create the service profile.
b) Click the **Hardware Based Service Profile** radio button.
c) In the **Name** field, enter a unique name for the service profile.
This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.
d) If you want Cisco UCS Manager to create vNICs for the service profile, check the **Create Default vNICs** check box.
e) If you want Cisco UCS Manager to create vHBAs for the service profile, check the **Create Default vHBAs** check box.
f) Click **OK**.

Cisco UCS Manager inherits and automatically applies the identity and configuration information in the server, creates the service profile, and associates it with the server.

Working with Service Profile Templates

Creating a Service Profile Template

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profile Templates**.

Step 3 Expand the node for the organization where you want to create the service profile template.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the organization and choose **Create Service Profile Template**.

Step 5 In the **Identify Service Profile Template** panel, specify the service profile **Name**, **Type**, and **UUID Assignment**, then click **Next**.
You can provide an optional description for this service profile template.

Note To create a service profile template quickly, you can click **Finish** after specifying the name. Cisco UCS Manager creates a new service profile template with the specified name and all system default values.

Step 6 (Optional) In the **Networking** panel, specify the required information for the **Dynamic vNIC Connection Policy** and **LAN Connectivity** sections, then click **Next**
You can create a dynamic vNIC connection policy and LAN connectivity policy from this panel.

Step 7 (Optional) In the **Storage** panel, specify the SAN configuration information such as, **Local Storage Policy**, **SAN Connectivity**, **WWNN**, and **vHBAs**, then click **Next**.
You can create a local disk configuration policy and SAN connectivity policy from this panel.

Step 8 (Optional) In the **Zoning** panel, specify the required zoning information, then click**Next**.
You can create the vHBA initiator groups from this panel.

Step 9 (Optional) In the **vNIC/vHBA Placement** panel, specify the placement method and PCI order, then click**Next**.
You can create a placement policy from this panel.

Step 10 (Optional) In the **Server Boot Order** panel, specify the **Boot Policy** from the drop-down list, then click **Next**.
You can create a boot policy from this panel.

Step 11 (Optional) In the **Maintenance Policy** panel, specify the maintenance policy, then click **Next**.
You can create a new maintenance policy and specify a maintenance schedule from this panel.

Step 12 (Optional) In the **Server Assignment** panel, specify the **Pool Assignment** from the drop down list and the power state to apply on assignment, then click **Next**.
You can create a server pool or a host firmware package from this panel.

Step 13 (Optional) In the **Operational Policies** panel, specify the system operational information such as, **BIOS Configuration**, **External IPMI Management Configuration**, **Management IP Address**, **Monitoring Configuration (Thresholds)**, **Power Control Policy Configuration**, and **Scrub Policy**, then click **Finish**.

Note To set up an Outband IPv4 address or an Inband IPv4 or IPv6 address, click the respective tabs and complete the required fields.

If you do not find the policies you need for each of these configurations, you can create them from this panel.

Creating One or More Service Profiles from a Service Profile Template

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profile Templates**.

Step 3 Expand the node for the organization that contains the service profile template that you want to use as the basis for your service profiles.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the service profile template from which you want to create the profiles and select **Create Service Profiles From Template**.

Step 5 In the **Create Service Profiles From Template** dialog box, complete the required fields.

Step 6 Click **OK**.

Creating a Template Based Service Profile for a Blade Server

Before You Begin

A qualified service profile template with the desired values must exist in Cisco UCS Manager.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.

Step 3 Choose the server for which you want to create a template based service profile.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Create Service Profile**.

Step 6 In the **Create Service Profile for Server** dialog box, do the following:

a) Click the **Template Based Service Profile** radio button.

b) In the **Name** field, enter a unique name for the service profile.
This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.

c) From the **Service Profile Template** drop-down list, select the template from which you want to create the service profile associated with this server.
Note The drop-down list only lists service profile templates compatible with the selected blade server.

d) Click **OK**.

Creating a Template Based Service Profile for a Rack-Mount Server

Before You Begin

A qualified service profile template with the desired values must exist in Cisco UCS Manager.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Rack Mounts** > **Servers**.

Step 3 Choose the server for which you want to create a template based service profile.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Create Service Profile**.

Step 6 In the **Create Service Profile for Server** dialog box, do the following:

a) Click the **Template Based Service Profile** radio button.

b) In the **Name** field, enter a unique name for the service profile.
This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.

c) From the **Service Profile Template** drop-down list, select the template from which you want to create the service profile associated with this server.

d) Click **OK**.

Creating a Service Profile Template from a Service Profile

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that contains the service profile that you want to use as the basis for your template.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the service profile from which you want to create the template and select **Create a Service Profile Template**.

Step 5 In the **Create Template From Service Profile** dialog box, complete the required fields.

Step 6 Click **OK**.

Managing Service Profiles

Cloning a Service Profile

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization where you want to create the service profile.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the service profile you want to clone and select **Create a Clone**.

Step 5 In the **Create Clone From Service Profile** dialog box:

a) Enter the name you want to use for the new profile in the **Clone Name** field.
This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.

This name must be unique within the organization or sub-organization in which you are creating the service profile.

b) Click **OK**.

Step 6 Navigate to the service profile you just created and make sure that all options are correct.

Associating a Service Profile with a Server or Server Pool

Follow this procedure if you did not associate the service profile with a blade server or server pool when you created it, or to change the blade server or server pool with which a service profile is associated.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that contains the service profile that you want to associate with a new server or server pool.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the service profile you want to associate with a server and select **Associate Service Profile**.

Step 5 In the **Associate Service Profile** dialog box, select one of the following options:

Option	Description
Server Pool	Select a server pool from the drop-down list. Cisco UCS Manager assigns a server from this pool to the service profile. Continue with Step 7.
Server	Navigate to the desired available server in the navigation tree and select the server which will be assigned to the service profile. Continue with Step 7.
Custom Server	Specifies the chassis and slot that contains the server that will be assigned to the service profile. If the server is not in the slot or is otherwise unavailable, the service profile will be associated with the server when it becomes available. Continue with Step 6.

Step 6 If you chose **Custom Server**, do the following:

a) In the **Chassis Id** field, enter the number of the chassis where the selected server is located.

b) In the **Server Id** field, enter the number of the slot where the selected server is located.

Step 7 If you want to restrict the migration of the service profile after it is associated with a server, check the **Restrict Migration.** check box.
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.

Step 8 Click **OK**.

Disassociating a Service Profile from a Server or Server Pool

When you disassociate a service profile, Cisco UCS Manager attempts to shutdown the operating system on the server. If the operating system does not shutdown within a reasonable length of time, Cisco UCS Manager forces the server to shutdown.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that contains the service profile that you want to disassociate from a server or server pool.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the service profile you want to disassociate from a server and select **Disassociate Service Profile**.

Step 5 In the **Disassociate Service Profile** dialog box, click **Yes** to confirm that you want to disassociate the service profile.

Step 6 (Optional) Monitor the status and FSM for the server to confirm that the disassociation completed.

Deleting the Inband Configuration from a Service Profile

This procedure removes the inband management IP address configuration from a service profile. If this action is greyed out, no inband configuration was configured.

Procedure

Step 1 In the **Navigation** pane, click the **Servers** tab.

Step 2 On the **Servers** tab, expand **Servers** > **Service Profiles** > ***Service_Profile_Name***.

Step 3 In the **Work** pane, click the **General** tab.

Step 4 In the **Actions** area, click **Delete Inband Configuration**.

Step 5 Click **Yes** in the **Delete** confirmation dialog box.
The inband management IP address configuration for the service profile is deleted.

Renaming a Service Profile

When you rename a service profile, the following occurs:

- Event logs and audit logs that reference the previous name for the service profile are retained under that name.
- A new audit record is created to log the rename operation.
- All records of faults against the service profile under its previous name are transferred to the new service profile name.

Note You cannot rename a service profile with pending changes.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that includes the service profile you want to rename.
If the system does not include multitenancy, expand the **root** node.

Step 4 Click the service profile you want to rename.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 In the **Actions** area, click **Rename Service Profile**.

Step 7 In the **Rename Service Profile** dialog box, enter the new name for the service profile in the **New Name** field. This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.

Step 8 Click **OK**.

Changing the UUID in a Service Profile

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that contains the service profile for which you want to change the UUID.
If the system does not include multitenancy, expand the **root** node.

Step 4 Choose the service profile that requires the UUID for the associated server to be changed.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 In the **Actions** area, click **Change UUID**.

Step 7 From the **UUID Assignment** drop-down list, do one of the following:

Option	Description
Select (pool default used by default)	Assigns a UUID from the default UUID Suffix pool. Continue with Step 9.

Option	Description
Hardware Default	Uses the UUID assigned to the server by the manufacturer. If you choose this option, the UUID remains unassigned until the service profile is associated with a server. At that point, the UUID is set to the UUID value assigned to the server by the manufacturer. If the service profile is later moved to a different server, the UUID is changed to match the new server. Continue with Step 9.
XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX	Uses the UUID that you manually assign. Continue with Step 8.
Pools ***Pool_Name***	Assigns a UUID from the UUID Suffix pool that you select from the list at the bottom of the drop-down list. Each pool name is followed by two numbers in parentheses that show the number of UUIDs still available in the pool and the total number of UUIDs in the pool. Continue with Step 9.

Step 8 (Optional) If you selected the **XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX** option, do the following:

a) In the **UUID** field, enter the valid UUID that you want to assign to the server which uses this service profile.

b) To verify that the selected UUID is available, click the **here** link.

Step 9 Click **OK**.

Modifying the Boot Order in a Service Profile

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that includes the service profile for which you want to change the boot order.
If the system does not include multi-tenancy, expand the **root** node.

Step 4 Click the service profile for which you want to change the boot order.

Step 5 In the **Work** pane, click the **Boot Order** tab.

Step 6 Click **Modify Boot Policy** to change the existing boot policy.

Step 7 In the **Modify Boot Policy** dialog box, choose one of the following from the **Boot Policy** drop-down list:

Option	Description
Select Boot Policy to use	Assigns the default boot policy to this service profile. Continue with Step 14.
Create a Specific Boot Policy	Enables you to create a local boot policy that can only be accessed by this service profile. Continue with Step 8.
Boot Policies *Policy_Name*	Assigns an existing boot policy to the service profile. If you choose this option, Cisco UCS Manager displays the details of the policy. If you do not want use any of the existing policies, but instead want to create a policy that all service profiles can access, click **Create Boot Policy** and continue with Step 2. Otherwise, continue with Step 14.

Step 8 If you chose to create a boot policy, in the **Create Boot Policy** dialog box, enter a unique name and description for the policy.
This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Step 9 (Optional) To reboot all servers that use this boot policy after you make changes to the boot order, check the **Reboot on Boot Order Change** check box.
In the Cisco UCS Manager GUI, if the **Reboot on Boot Order Change** check box is checked for a boot policy, and if CD-ROM or Floppy is the last device in the boot order, deleting or adding the device does not directly affect the boot order and the server does not reboot.

Step 10 (Optional) If desired, check the **Enforce vNIC/vHBA/iSCSI Name** check box.

- If checked, Cisco UCS Manager displays a configuration error and reports whether one or more of the vNICs, vHBAs, or iSCSI vNICs listed in the **Boot Order** table match the server configuration in the service profile.
- If not checked, Cisco UCS Manager uses the vNICs or vHBAs (as appropriate for the boot option) from the service profile.

Step 11 To add a local disk, virtual CD-ROM, or virtual floppy to the boot order, do the following:

a) Click the down arrows to expand the **Local Devices** area.

b) Click one of the following links to add the device to the **Boot Order** table:

- **Add Local Disk** or
  - **Add Local LUN**
  - **Add SD Card**
  - **Add Internal USB**

- **Add External USB**

- **Add CD/DVD** or
  - **Add Local CD/DVD**
  - **Add Local Remote CD/DVD**

c) Add another boot device to the **Boot Order** table, or click **OK** to finish.

Step 12 To add a LAN boot to the boot order, do the following:

a) Click the down arrows to expand the **vNICs** area.
b) Click the **Add LAN Boot** link.
c) In the **Add LAN Boot** dialog box, enter the name of the vNIC that you want to use for the LAN boot in the **vNIC** field, then click **OK**.
d) Add another device to the **Boot Order** table, or click **OK** to finish.

Step 13 To add a SAN boot to the boot order, do the following:

a) Click the down arrows to expand the **vHBAs** area.
b) Click the **Add SAN Boot** link.
c) In the **Add San Boot** dialog box, specify the vHBA and type, then click **OK**.
d) If this vHBA points to a bootable SAN image, click the **Add SAN Boot Target** link and, in the **Add SAN Boot Target** dialog box, specify the boot target LUN, boot target WWPN, and type, then click **OK**.
e) Add another boot device to the **Boot Order** table, or click **OK** to finish.

Step 14 Click **OK**.

Creating a vNIC for a Service Profile

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that contains the service profile for which you want to create a vNIC.

Step 4 Expand the service profile for which you want to create a vNIC.

Step 5 Right-click the **vNICs** node and choose **Create vNICs**.

Step 6 In the **Create vNIC** dialog box, enter the name, select a **MAC Address Assignment**, and check the **Use vNIC Template** check box only if one or more vNIC templates exist in the system.
You can also create a MAC pool from this area.

Note You can also create vNIC pairs for any service profile in the same organization or sub-organization. This eliminates the need to configure vNIC pairs independently using one or more templates. For information on creating vNIC pairs, see Creating vNIC Pairs for a Service Profile, on page 605.

Step 7 Select the **vNIC Template** from the drop-down list.

You can also click **Create vNIC Template** to display the**Create vNIC Template** dialog box to create a different vNIC template.

Step 8 Choose the **Fabric ID**, select the **VLANs** that you want to use, enter the **MTU**, and choose a **Pin Group**.
You can also create a VLAN and a LAN pin group from this area.

Step 9 In the **Operational Parameters** area, choose a **Stats Threshold Policy**.

Step 10 In the Adapter Performance Profile area, choose an **Adapter Policy**, **QoS Policy**, and a **Network Control Policy**.
You can also create an Ethernet adapter policy, QoS policy, and network control policy from this area.

Step 11 In the Connection Policies area, choose the **Dynamic vNIC**, **usNIC** or **VMQ** radio button, then choose the corresponding policy.
You can also create a dynamic vNIC, usNIC, or VMQ connection policy from this area.

Step 12 Click **OK**.

Creating vNIC Pairs for a Service Profile

You can use any vNIC template pair to create vNIC pairs for any service profile in the same organization or sub-organization. This eliminates the need to configure vNIC pairs independently using one or more templates.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that contains the service profile for which you want to create a vNIC.

Step 4 Expand the service profile for which you want to create a vNIC.

Step 5 Right-click the **vNICs** node and choose **Create vNICs**.

Step 6 Enter a name for the vNIC. You can also select a **MAC Address Assignment**.

Step 7 Check the **Use vNIC Template** check box only if one or more vNIC templates exist in the system.
You can also create a MAC pool from this area.

Step 8 Check the **Redundancy Pair** check box to use a vNIC Template Pair to group vNICs to belong to a specific server.

Step 9 Enter the name of the Primary vNIC and Secondary vNIC, which is the peer vNIC in the vNIC pair.

Step 10 Select the **vNIC Template** from the drop-down list.
You can also click **Create vNIC Template** to display the**Create vNIC Template** dialog box to create a different vNIC template to use for the vNIC pair.

Note If you plan to use global vNIC templates for your redundancy pair when using a local service profile, you can assign the vNIC template for the primary vNIC and set the peer name for the second vNIC. However, you need to modify the secondary vNIC and manually assign the secondary vNIC template.

Step 11 Click **OK**.

Deleting a vNIC from a Service Profile

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Service Profiles**.
Step 3 Expand the node for the organization that contains the service profile from which you want to delete a vNIC.
Step 4 Expand the service profile from which you want to delete a vNIC.
Step 5 Expand the **vNICs** node.
Step 6 Right-click the vNIC you want to delete and choose **Delete**.
Step 7 If a confirmation dialog box displays, click **Yes**.

Creating a vHBA for a Service Profile

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Service Profiles**.
Step 3 Expand the node for the organization that contains the service profile for which you want to create a vHBA.
Step 4 Expand the service profile for which you want to create a vHBA.
Step 5 Right-click the **vHBAs** node and choose **Create vHBAs**.
Step 6 In the **Create vHBAs** dialog box, enter the name and optional description.
Step 7 Choose the **Fabric ID**, **Select VSAN**, **Pin Group**, **Persistent Binding**, and **Max Data Field Size**.
You can also create a VSAN or SAN pin group from this area.
Step 8 In the **Operational Parameters** area, choose the **Stats Threshold Policy**.
Step 9 In the **Adapter Performance Profile** area, choose the **Adapter Policy** and **QoS Policy**.
You can also create a fibre channel adapter policy or QoS policy from this area.
Step 10 Click **OK**.

Creating a vHBA Pair for a Service Profile

You can use any vHBA template pair to create vHBA pairs for any service profile in the same organization or sub-organization. This eliminates the need to configure vHBA pairs independently using one or more templates

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that contains the service profile for which you want to create a vHBA.

Step 4 Expand the service profile for which you want to create a vHBA.

Step 5 Right-click the **vHBAs** node and choose **Create vHBAs**.

Step 6 Enter a name for the vHBA.

Step 7 Check the **Use vHBA Template** check box only if one or more vHBA templates exist in the system.
You can also select the **World Wide Port Name** from this area.

Step 8 Check the **Redundancy Pair** check box to use a vHBA Template Pair to group vHBAs to belong to a specific server.

Step 9 Enter the name of the **Primary** vHBA and **Secondary** vHBA, which will be peer vHBAs in the vHBA pair.

Step 10 Select one of the vHBA templates in the template pair (Primary or Secondary). UCS Manager automatically gets the peer template in the pair to create the peer vHBA. You can also click **Create vHBA Template** to create a new vHBA Template pair to use for this vHBA pair.

Note If you plan to use global vHBA templates for your redundancy pair when using a local service profile, you can assign the vHBA template for the primary vHBA and set the peer name for the second vHBA. However, you need to modify the secondary vHBA and manually assign the secondary vHBA template.

Step 11 Choose the **Fabric ID**, select the **VSANs** that you want to use, then enable or disable the **Persistent Binding**, and enter the **Max Data Field Size**.
You can also create a VSAN and a SAN Pin Group from this area.

Step 12 In the **Operational Parameters** area, choose a **Stats Threshold Policy**.

Step 13 In the Adapter Performance Profile area, choose an **Adapter Policy** and a **QoS Policy**.
You can also create a Fibre Channel Adapter Policy and a QoS policy.

Step 14 Click **OK**.

Changing the WWPN for a vHBA

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Service Profiles**.
Step 3 Expand the node for the organization that contains the service profile for which you want to change the WWPN.
Step 4 Expand ***Service_Profile_Name*** > **vHBAs**.
Step 5 Click the vHBA for which you want to change the WWPN.
Step 6 In the **Work** pane, click the **General** tab.
Step 7 In the **Actions** area, click **Change World Wide Name**.
Step 8 In the **Change World Wide Port Name** dialog box, complete the required fields.
Step 9 Click **OK**.

Clearing Persistent Binding for a vHBA

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Service Profiles**.
Step 3 Expand the node for the organization that contains the service profile for which you want to modify the vHBA.
Step 4 Expand ***Service_Profile_Name*** > **vHBAs**.
Step 5 Click the vHBA for which you want to clear the persistent binding.
Step 6 In the **Work** pane, click the **General** tab.
Step 7 In the **Actions** area, click **Clear Persistent Binding**.
Step 8 If a confirmation dialog box displays, click **Yes**.

Deleting a vHBA from a Service Profile

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Service Profiles**.
Step 3 Expand the node for the organization that contains the service profile from which you want to delete a vHBA.
Step 4 Expand the service profile from which you want to delete a vHBA.
Step 5 Expand the **vHBAs** node.
Step 6 Right-click the vHBA you want to delete and choose **Delete**.
Step 7 If a confirmation dialog box displays, click **Yes**.

Adding a vHBA Initiator Group to a Service Profile

Procedure

Step 1 Expand **Servers** > **Service Profiles**.
Step 2 Expand the node for the organization that contains the service profile to which you want to add a vHBA initiator group.
If the system does not include multitenancy, expand the **root** node.
Step 3 Choose the service profile to which you want to add a vHBA initiator group.
Step 4 In the **Work** pane, click the **Storage** > **vHBA Initiator Groups**.
Step 5 On the icon bar at the right of the **Select vHBA Initiator Groups** table, click +.
Step 6 In the **Create vHBA Initiator Group** dialog box, complete the following fields to set the name and description:

Name	Description
Name field	The name of the vHBA initiator group. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the group. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Step 7 In the **Select vHBA Initiators** table, check the check box in the **Select** column for each vHBA you want to include in the vHBA initiator group.

Step 8 To add a storage connection policy to the initiator group, choose one of the following options:

- Choose an existing storage connection policy from the **Storage Connection Policy** drop-down list. Continue with Step 10.
- Click the **Create Storage Connection Policy** link if you want to create a new storage connection policy that will be available for use by other vHBA initiator groups within the Cisco UCS domain. For more information, see Creating a Fibre Channel Storage Connection Policy, on page 388. After you create the storage connection policy, continue with Step 10.
- Choose the **Specific Storage Connection Policy** option to create a storage connection policy that is only available to this vHBA initiator group. Continue with Step 9.

Step 9 In the **Specific Storage Connection Policy** area, complete the following fields to create a storage connection policy that is only available to this vHBA initiator group:

Name	Description
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Zoning Type field	This can be one of the following: • **None—**Cisco UCS Manager does not configure Fibre Channel zoning. • **Single Initiator Single Target—**Cisco UCS Manager automatically creates one zone for each vHBA and storage port pair. Each zone has two members. We recommend that you configure this type of zoning unless you expect the number of zones to exceed the maximum supported. • **Single Initiator Multiple Targets—**Cisco UCS Manager automatically creates one zone for each vHBA. We recommend that you configure this type of zoning if you expect the number of zones to reach or exceed the maximum supported.

Name	Description
FC Target Endpoints table	The Fibre Channel target endpoints associated with this policy. This table contains the following columns and buttons: • **WWPN** column—The World Wide Port Name associated with the endpoint. • **Path** column—The path to the endpoint. • **VSAN** column—The VSAN associated with the endpoint. • **Add** button—Creates a new FC target endpoint. • **Delete** button—Deletes the selected endpoint. • **Properties** button—Displays all properties for the selected endpoint.

Step 10 Click **OK**.

Step 11 If a confirmation dialog box displays, click **Yes**.

Binding a Service Profile to a Service Profile Template

You can bind a service profile to a service profile template. When you bind the service profile to a template, Cisco UCS Manager configures the service profile with the values defined in the service profile template. If the existing service profile configuration does not match the template, Cisco UCS Manager reconfigures the service profile. You can only change the configuration of a bound service profile through the associated template.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that includes the service profile you want to bind.
If the system does not include multi-tenancy, expand the **root** node.

Step 4 Click the service profile you want to bind.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 In the **Actions** area, click **Bind to a Template**.

Step 7 In the **Bind to a Service Profile Template** dialog box, do the following:

a) From the **Service Profile Template** drop-down list, choose the template to which you want to bind the service profile.

b) Click **OK**.

Unbinding a Service Profile from a Service Profile Template

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Service Profiles**.
Step 3 Expand the node for the organization that includes the service profile you want to unbind.
If the system does not include multi-tenancy, expand the **root** node.
Step 4 Click the service profile you want to unbind.
Step 5 In the **Work** pane, click the **General** tab.
Step 6 In the **Actions** area, click **Unbind from the Template**.
Step 7 If a confirmation dialog box displays, click **Yes**.

Deleting a Service Profile

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 In the **Servers** tab, expand **Servers** > **Service Profiles** > *Organization_Name* .
Step 3 Right-click the service profile you want to delete and select **Delete**.
Step 4 If a confirmation dialog box displays, click **Yes**.
Step 5 Click **OK**.

Managing Service Profile Templates

Associating a Service Profile Template with a Server Pool

Follow this procedure if you did not associate the service profile template with a server pool when you created it, or to change the server pool with which a service profile created from this template is associated.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profile Templates**.

Step 3 Expand the node for the organization that contains the service profile that you want to associate with a server pool.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the service profile template you want to associate with a server pool and select **Associate with Server Pool**.
The **Associate with Server Pool** dialog box opens.

Step 5 From the **Server Pool** section of the **Pool Assignment** drop-down list, select a server pool.
If you select **Assign Later**, the service profile template is not associated with a server pool.

Step 6 (Optional) From the **Select Qualification** drop-down list, select the server pool policy qualifications you want to apply to a server that is associated with a service profile created from this template.

Step 7 Click **OK**.

Disassociating a Service Profile Template from its Server Pool

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profile Templates**.

Step 3 Expand the node for the organization that contains the service profile that you want to disassociate from its server pool.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the service profile template you want to disassociate from its server pool and select **Disassociate Template**.

Step 5 If a confirmation dialog box displays, click **Yes**.

Changing the UUID in a Service Profile Template

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profile Templates**.

Step 3 Expand the node for the organization that contains the service profile template for which you want to change the UUID.
If the system does not include multitenancy, expand the **root** node.

Step 4 Choose the service profile template whose UUID assignment you want to change.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 In the **Actions** area, click **Change UUID**.

Step 7 From the **UUID Assignment** drop-down list, choose one of the following:

Option	Description
Select (pool default used by default)	Assigns a UUID from the default UUID Suffix pool.
Hardware Default	Uses the UUID assigned to the server by the manufacturer. If you choose this option, the UUID remains unassigned until the service profile is associated with a server. At that point, the UUID is set to the UUID value assigned to the server by the manufacturer. If the service profile is later moved to a different server, the UUID is changed to match the new server.
Pools ***Pool_Name***	Assigns a UUID from the UUID Suffix pool that you select from the list at the bottom of the drop-down list. Each pool name is followed by two numbers in parentheses that show the number of UUIDs still available in the pool and the total number of UUIDs in the pool.

Step 8 Click **OK**.

Resetting the UUID Assigned to a Service Profile from a Pool in a Service Profile Template

If you change the UUID suffix pool assigned to an updating service profile template, Cisco UCS Manager does not change the UUID assigned to a service profile created with that template. If you want Cisco UCS Manager to assign a UUID from the newly assigned pool to the service profile, and therefore to the associated server, you must reset the UUID. You can only reset the UUID assigned to a service profile and its associated server under the following circumstances:

- The service profile was created from an updating service profile template and includes a UUID assigned from a UUID suffix pool.
- The UUID suffix pool name is specified in the service profile. For example, the pool name is not empty.
- The UUID value is not 0, and is therefore not derived from the server hardware.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that contains the service profile for which you want to reset the UUID. If the system does not include multitenancy, expand the **root** node.

Step 4 Choose the service profile that requires the UUID for the associated server to be reset to a different UUID suffix pool.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 In the **Actions** area, click **Reset UUID**.
If this action is not visible, then the UUID configuration in the service profile does not meet the requirements for resetting a UUID.

Step 7 If a confirmation dialog box displays, click **Yes**.

Step 8 Click **OK**

Resetting the MAC Address Assigned to a vNIC from a Pool in a Service Profile Template

If you change the MAC pool assigned to an updating service profile template, Cisco UCS Manager does not change the MAC address assigned to a service profile created with that template. If you want Cisco UCS Manager to assign a MAC address from the newly assigned pool to the service profile, and therefore to the associated server, you must reset the MAC address. You can only reset the MAC address assigned to a service profile and its associated server under the following circumstances:

- The service profile was created from an updating service profile template and includes a MAC address assigned from a MAC pool.
- The MAC pool name is specified in the service profile. For example, the pool name is not empty.
- The MAC address value is not 0, and is therefore not derived from the server hardware.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that contains the service profile for which you want to reset the MAC address.
If the system does not include multitenancy, expand the **root** node.

Step 4 Expand ***Service_Profile_Name*** > **vNICs**.

Step 5 Click the vNIC for which you want to reset the MAC address.

Step 6 In the **Work** pane, click the **General** tab.

Step 7 In the **Actions** area, click **Reset MAC Address**.

Step 8 If a confirmation dialog box displays, click **Yes**.

Step 9 Click **OK**.

Resetting the WWPN Assigned to a vHBA from a Pool in a Service Profile Template

If you change the WWPN pool assigned to an updating service profile template, Cisco UCS Manager does not change the WWPN assigned to a service profile created with that template. If you want Cisco UCS Manager to assign a WWPN from the newly assigned pool to the service profile, and therefore to the associated server, you must reset the WWPN. You can only reset the WWPN assigned to a service profile and its associated server under the following circumstances:

- The service profile was created from an updating service profile template and includes a WWPN assigned from a WWPN pool.
- The WWPN pool name is specified in the service profile. For example, the pool name is not empty.
- The WWPN value is not 0, and is therefore not derived from the server hardware.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that contains the service profile for which you want to reset the WWPN.
If the system does not include multitenancy, expand the **root** node.

Step 4 Expand *Service_Profile_Name* > **vHBAs**.
Step 5 Click the vHBA for which you want to reset the WWPN.
Step 6 In the **Work** pane, click the **General** tab.
Step 7 In the **Actions** area, click **Reset WWPN**.
Step 8 If a confirmation dialog box displays, click **Yes**.
Step 9 Click **OK**.

Deleting the Inband Configuration from a Service Profile Template

This procedure removes the inband management IP address configuration from a service profile template. If this action is greyed out, no inband configuration was configured.

Procedure

Step 1 In the **Navigation** pane, click the **Servers** tab.
Step 2 On the **Servers** tab, expand **Servers** > **Service Profile Template** > *Service_Profile_Template_Name*.
Step 3 In the **Work** pane, click the **General** tab.
Step 4 In the **Actions** area, click **Delete Inband Configuration**.
Step 5 Click **Yes** in the **Delete** confirmation dialog box.
The inband management IP address configuration for the service profile template is deleted.

What to Do Next

CHAPTER 34

Configuring Storage Profiles

This part contains the following chapters:

Storage Profiles

To allow flexibility in defining the number of storage disks, roles and usage of these disks, and other storage parameters, you can create and use storage profiles. A storage profile encapsulates the storage requirements for one or more service profiles. LUNs configured in a storage profile can be used as boot LUNs or data LUNs, and can be dedicated to a specific server. You can also specify a local LUN as a boot device. However, LUN resizing is not supported. The introduction of storage profiles allows you to do the following:

- Configure multiple virtual drives and select the physical drives that are used by a virtual drive. You can also configure the storage capacity of a virtual drive.
- Configure the number, type and role of disks in a disk group.
- Associate a storage profile with a service profile.

You can create a storage profile both at an org level and at a service-profile level. A service profile can have a dedicated storage profile as well as a storage profile at an org level.

Disk Groups and Disk Group Configuration Policies

You can select and configure the disks to be used for storage. A logical collection of these physical disks is called a disk group. Disk groups allow you to organize local disks. The storage controller controls the creation and configuration of disk groups.

A disk group configuration policy defines how a disk group is created and configured. The policy specifies the RAID level to be used for the disk group. It also specifies either a manual or an automatic selection of disks for the disk group, and roles for disks. You can use a disk group policy to manage multiple disk groups. However, a single disk group can be managed only by one disk group policy.

A hot spare is an unused extra disk that can be used by a disk group in the case of failure of a disk in the disk group. Hot spares can be used only in disk groups that support a fault-tolerant RAID level. In addition, a disk can be allocated as a global hot spare, which means that it can be used by any disk group.

Virtual Drives

A disk group can be partitioned into virtual drives. Each virtual drive appears as an individual physical device to the Operating System.

All virtual drives in a disk group must be managed by using a single disk group policy.

Configuration States

Indicates the configuration states of a virtual drive. Virtual drives can have the following configuration states:

- Applying—Creation of the virtual drive is in progress.
- Applied—Creation of the virtual drive is complete, or virtual disk policy changes are configured and applied successfully.
- Failed to apply—Creation, deletion, or renaming of a virtual drive has failed due to errors in the underlying storage subsystem.
- Orphaned—The service profile that contained this virtual drive is deleted or the service profile is no longer associated with a storage profile.

Deployment States

Indicates the actions that you are performing on virtual drives. Virtual drives can have the following deployment states:

- No action—No pending work items for the virtual drive.
- Creating—Creation of the virtual drive is in progress.
- Deleting—Deletion of the virtual drive is in progress.
- Modifying—Modification of the virtual drive is in progress.

Operability States

Indicates the operating condition of a virtual drive. Virtual drives can have the following operability states:

- Optimal—The virtual drive operating condition is good. All configured drives are online.
- Degraded—The virtual drive operating condition is not optimal. One of the configured drives has failed or is offline.
- Cache-degraded—The virtual drive has been created with a write policy of **write back** mode, but the BBU has failed, or there is no BBU.

Note This state does not occur if you select the **always write back** mode.

- Partially degraded—The operating condition in a RAID 6 virtual drive is not optimal. One of the configured drives has failed or is offline. RAID 6 can tolerate up to two drive failures.
- Offline—The virtual drive is not available to the RAID controller. This is essentially a failed state.
- Unknown—The state of the virtual drive is not known.

Presence States

Indicates the presence of virtual drive components. Virtual drives have the following presence states:

- Equipped—The virtual drive is available.
- Mismatched—A virtual drive deployed state is different from its configured state.
- Missing—Virtual drive is missing.

RAID Levels

The RAID level of a disk group describes how the data is organized on the disk group for the purpose of ensuring availability, redundancy of data, and I/O performance.

The following are features provided by RAID:

- Striping—Segmenting data across multiple physical devices. This improves performance by increasing throughput due to simultaneous device access.
- Mirroring—Writing the same data to multiple devices to accomplish data redundancy.
- Parity—Storing of redundant data on an additional device for the purpose of error correction in the event of device failure. Parity does not provide full redundancy, but it allows for error recovery in some scenarios.
- Spanning—Allows multiple drives to function like a larger one. For example, four 20 GB drives can be combined to appear as a single 80 GB drive.

The supported RAID levels include the following:

- RAID 0 Striped—Data is striped across all disks in the array, providing fast throughput. There is no data redundancy, and all data is lost if any disk fails.
- RAID 1 Mirrored—Data is written to two disks, providing complete data redundancy if one disk fails. The maximum array size is equal to the available space on the smaller of the two drives.

- RAID 5 Striped Parity—Data is striped across all disks in the array. Part of the capacity of each disk stores parity information that can be used to reconstruct data if a disk fails. RAID 5 provides good data throughput for applications with high read request rates.

  RAID 5 distributes parity data blocks among the disks that are part of a RAID-5 group and requires a minimum of three disks.

- RAID 6 Striped Dual Parity—Data is striped across all disks in the array and two sets of parity data are used to provide protection against failure of up to two physical disks. In each row of data blocks, two sets of parity data are stored.

  Other than addition of a second parity block, RAID 6 is identical to RAID 5 . A minimum of four disks are required for RAID 6.

- RAID 10 Mirrored and Striped—RAID 10 uses mirrored pairs of disks to provide complete data redundancy and high throughput rates through block-level striping. RAID 10 is mirroring without parity and block-level striping. A minimum of four disks are required for RAID 10.

- RAID 50 Striped Parity and Striped—Data is striped across multiple striped parity disk sets to provide high throughput and multiple disk failure tolerance.

- RAID 60 Striped Dual Parity and Striped—Data is striped across multiple striped dual parity disk sets to provide high throughput and greater disk failure tolerance.

Automatic Disk Selection

When you specify a disk group configuration, and do not specify the local disks in it, Cisco UCS Manager determines the disks to be used based on the criteria specified in the disk group configuration policy. Cisco UCS Manager can make this selection of disks in multiple ways.

When all qualifiers match for a set of disks, then disks are selected sequentially according to their slot number. Regular disks and dedicated hot spares are selected by using the lowest numbered slot.

The following is the disk selection process:

1. Iterate over all local LUNs that require the creation of a new virtual drive. Iteration is based on the following criteria, in order:
   a. Disk type
   b. Minimum disk size from highest to lowest
   c. Space required from highest to lowest
   d. Disk group qualifier name, in alphabetical order
   e. Local LUN name, in alphabetical order
2. Select regular disks depending on the minimum number of disks and minimum disk size. Disks are selected sequentially starting from the lowest numbered disk slot that satisfies the search criteria.

Note If you specify **Any** as the type of drive, the first available drive is selected. After this drive is selected, subsequent drives will be of a compatible type. For example, if the first drive was SATA, all subsequent drives would be SATA.

3 Select dedicated hot spares by using the same method as normal disks. Disks are only selected if they are in an **Unconfigured Good** state.

4 If a provisioned LUN has the same disk group policy as a deployed virtual drive, then try to deploy the new virtual drive in the same disk group. Otherwise, try to find new disks for deployment.

Supported LUN Modifications

Some modifications that are made to the LUN configuration when LUNs are already deployed on an associated server are supported.

The following are the types of modifications that can be performed:

- Creation of a new virtual drive.
- Deletion of an existing virtual drive, which is in the orphaned state.
- Non-disruptive changes to an existing virtual drive. These changes can be made on an existing virtual drive without loss of data, and without performance degradation:
  - Policy changes. For example, changing the write cache policy.
  - Modification of boot parameters

The removal of a LUN will cause a warning to be displayed. Ensure that you take action to avoid loss of data.

Unsupported LUN Modifications

Some modifications to existing LUNs are not possible without destroying the original virtual drive and creating a new one. All data is lost in these types of modification, and these modifications are not supported.

Disruptive modifications to an existing virtual drive are not supported. The following are unsupported disruptive changes:

- Any supported RAID level change that can be handled through reconstruction. For example, RAID0 to RAID1.
- Increasing the size of a virtual drive through reconstruction.
- Addition and removal of disks through reconstruction.

Destructive modifications are also not supported. The following are unsupported destructive modifications:

- RAID-level changes that do not support reconstruction. For example, RAID5 to RAID1.
- Shrinking the size of a virtual drive.
- RAID-level changes that support reconstruction, but where there are other virtual drives present on the same drive group.
- Disk removal when there is not enough space left on the disk group to accommodate the virtual drive.
- Explicit change in the set of disks used by the virtual drive.

Disk Insertion Handling

When the following sequence of events takes place:

1. The LUN is created in one of the following ways:
   1. You specify the slot specifically by using a local disk reference
   2. The system selects the slot based on criteria specified by you
2. The LUN is successfully deployed, which means that a virtual drive is created, which uses the slot.
3. You remove a disk from the slot, possibly because the disk failed.
4. You insert a new working disk into the same slot.

The following scenarios are possible:

- Non-Redundant Virtual Drives, on page 624
- Redundant Virtual Drives with No Hot Spare Drives, on page 624
- Redundant Virtual Drives with Hot Spare Drives, on page 624
- Replacing Hot Spare Drives, on page 625
- Inserting Physical Drives into Unused Slots, on page 625

Non-Redundant Virtual Drives

For non-redundant virtual drives (RAID 0), when a physical drive is removed, the state of the virtual drive is **Inoperable**. When a new working drive is inserted, the new physical drive goes to an **Unconfigured Good** state.

For non-redundant virtual drives, there is no way to recover the virtual drive. You must delete the virtual drive and re-create it.

Redundant Virtual Drives with No Hot Spare Drives

For redundant virtual drives (RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, RAID 60) with no hot spare drives assigned, virtual drive mismatch, virtual drive member missing, and local disk missing faults appear until you insert a working physical drive into the same slot from which the old physical drive was removed.

If the physical drive size is greater than or equal to that of the old drive, the storage controller automatically uses the new drive for the virtual drive. The new drive goes into the **Rebuilding** state. After rebuild is complete, the virtual drive goes back into the **Online** state.

Redundant Virtual Drives with Hot Spare Drives

For redundant virtual drives (RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, RAID 60) with hot spare drives assigned, when a drive fails, or when you remove a drive, the dedicated hot spare drive, if available, goes into

the **Rebuilding** state with the virtual drive in the **Degraded** state. After rebuilding is complete, that drive goes to the **Online** state.

Cisco UCSM raises a disk missing and virtual drive mismatch fault because although the virtual drive is operational, it does not match the physical configuration that Cisco UCSM expects.

if you insert a new disk in the slot with the disk missing, automatic copy back starts from the earlier hot spare disk to the newly inserted disk. After copy back, the hot spare disk is restored. In this state all faults are cleared.

If automatic copy back does not start, and the newly inserted disk remains in the **Unconfigured Good**, **JBOD**, or **Foreign Configuration** state, remove the new disk from the slot, reinsert the earlier hot spare disk into the slot, and import foreign configuration. This initiates the rebuilding process and the drive state becomes **Online**. Now, insert the new disk in the hot spare slot and mark it as hot spare to match it exactly with the information available in Cisco UCSM.

Replacing Hot Spare Drives

If a hot spare drive is replaced, the new hot spare drive will go to the **Unconfigured Good**, **Unconfigured Bad**, **JBOD**, or **Foreign Configuration** state.

Cisco UCSM will raise a virtual drive mismatch or virtual drive member mismatch fault because the hot spare drive is in a state different from the state configured in Cisco UCSM.

You must manually clear the fault. To do this, you must perform the following actions:

1. Clear the state on the newly inserted drive to **Unconfigured Good**.
2. Configure the newly inserted drive as a hot spare drive to match what is expected by Cisco UCSM.

Inserting Physical Drives into Unused Slots

If you insert new physical drives into unused slots, neither the storage controller nor Cisco UCSM will make use of the new drive even if the drive is in the **Unconfigured Good** state and there are virtual drives that are missing good physical drives.

The drive will simply go into the **Unconfigured Good** state. To make use of the new drive, you will need to modify or create LUNs to reference the newly inserted drive.

Virtual Drive Naming

When you use UCSM to create a virtual drive, UCSM assigns a unique ID that can be used to reliably identify the virtual drive for further operations. UCSM also provides the flexibility to provide a name to the virtual drive at the time of service profile association. Any virtual drive without a service profile or a server reference is marked as an orphan virtual drive.

In addition to a unique ID, a name is assigned to the drive. Names can be assigned in two ways:

- When configuring a virtual drive, you can explicitly assign a name that can be referenced in storage profiles.
- If you have not preprovisioned a name for the virtual drive, UCSM generates a unique name for the virtual drive.

You can rename virtual drives that are not referenced by any service profile or server.

LUN Dereferencing

A LUN is dereferenced when it is no longer used by any service profile. This can occur as part of the following scenarios:

- The LUN is no longer referenced from the storage profile
- The storage profile is no longer referenced from the service profile
- The server is disassociated from the service profile
- The server is decommissioned

When the LUN is no longer referenced, but the server is still associated, re-association occurs.

When the service profile that contained the LUN is deleted, the LUN state is changed to **Orphaned**.

Controller Constraints and Limitations

- For Cisco UCS C240, C220, C24, and C22 servers, the storage controller allows 24 virtual drives per server. For all other servers, the storage controller allows 16 virtual drives per server.
- In Cisco UCS Manager Release 2.2(4), blade servers do not support drives with a block size of 4K, but rack-mount servers support such drives. If a drive with a block size of 4K is inserted into a blade server, discovery fails and the following error message appears:`Unable to get Scsi Device Information from the system`.

Configuring Storage Profiles

Configuring a Disk Group Policy

Configuring a disk group involves the following:

1. Setting the RAID level
2. Automatically or manually configuring disks in a disk group policy
3. Configuring virtual drive properties

Configuring a Disk Group Policy

You can configure the disks in a disk group policy automatically or manually.

Procedure

Step 1 In the **Navigation** pane, click **Storage**.

Step 2 Expand **Storage** > **Storage Provisioning** > **Storage Policies**

Step 3 Expand the node for the organization where you want to create the disk group policy.

Step 4 Right-click **Disk Group Policies** in the organization and select **Create Disk Group Policy**.

Step 5 In the **Create Disk Group Policy** dialog box, specify the following:

Name	Description
Name field	The name of the policy This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object has been saved.
Description field	A description of the policy. We recommend that you include information about where and when the policy should be used. Enter up to 256 characters. You can use any characters or spaces except `(accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
RAID Level drop-down list	This can be one of the following: • **RAID 0 Striped** • **RAID 1 Mirrored** • **RAID 5 Striped Parity** • **RAID 6 Striped Dual Parity** • **RAID 10 Mirrored and Striped** • **RAID 50 Striped Parity and Striped** • **RAID 60 Striped Dual Parity and Striped**

Step 6 To automatically configure the disks in a disk group policy, select **Disk Group Configuration (Automatic)** and specify the following:

Name	Description
Number of drives field	Specifies the number of drives for the disk group. The range for drives is from 0 to 24 drivesfor Cisco UCS C240, C220, C24, and C22 servers. For all other servers, the limit is 16 drives per server.. **Unspecified** is the default number of drives. When you select the number of drives as **Unspecified**, the number of drives will be selected according to the disk selection process.

Name	Description
Drive Type field	Drive type for the disk group. You can select: • **Unspecified** • **HDD** • **SSD** **Unspecified** is the default drive type. If you select any of the drive types, the system automatically selects the first available unconfigured good drive o that type. After the system selects the drive type, all subsequent drives would be of that same type. For example, if the first unconfigured good drive selected is SATA, all subsequent drives would be SATA.
Number of Hot Spares field	Number of dedicated hot spares for the disk group. The range for dedicated hot spares is from 0 to 24 hot spares. **Unspecified** is the default number of dedicated hot spares. When you select the number of dedicated hot spares as **Unspecified**, the hot spares will be selected according to the disk selection process.
Number of Global Hot Spares field	Number of global hot spares for the disk group. The range for dedicated hot spares is from 0 to 24 hot spares. **Unspecified** is the default number of global hot spares. When you select the number of global hot spares as **Unspecified**, the hot spares will be selected according to the disk selection process.
Min Drive Size field	Minimum drive size for the disk group. Only disks that match this criteria are available for selection. The range for minimum drive size is from 0 to 10240 GB. **Unspecified** is the default minimum drive size. When you select the minimum drive size as **Unspecified**, drives of all sizes will be available for selection.
Use Remaining Disks checkbox	Indicates whether the remaining disks in the disk group should be used or not. By default, this check box is not checked.

Step 7 To manually configure the disks in a disk group policy, select **Disk Group Configuration (Manual)** and do the following:

a) On the icon bar to the right of the table, click +

b) In the **Create Local Disk Configuration Reference** dialog box, complete the following fields:

Name	Description
Slot field	Slot for which the local disk reference is configured.
Role field	Role of the local disk in the disk group. You can select: • **Dedicated Hot Spare** • **Normal**

Name	Description
Span ID field	Span ID of the span group to which the disk belongs. Disks belonging to a single span group can be treated as a single disk with a larger size. The values range from 0 to 8. You can also set the Span ID as **Unspecified** when spanning information is not required. **Unspecified** is the default Span ID of the local disk.

Step 8 In the **Virtual Drive Configuration** area, specify the following:

Name	Description
Strip Size (KB) field	Stripe size for a virtual drive. This can only be **Platform Default**.
Access Policy field	Access policy for a virtual drive. This can be one of the following: • **Platform Default** • **Read Write** • **Read Only** • **Blocked**
Read Policy field	Read policy for a virtual drive. This can be one of the following: • **Platform Default** • **Read Ahead** • **Normal**
Write Cache Policy field	Write-cache-policy for a virtual drive. This can be one of the following: • **Platform Default** • **Write Through** • **Write Back Good Bbu** • **Always Write Back**
IO Policy field	I/O policy for a virtual drive. This can be one of the following: • **Platform Default** • **Direct** • **Cached**

Name	Description
Drive Cache field	State of the drive cache. This can be one of the following: • **Platform Default** • **No Change** • **Enable** • **Disable**

All virtual drives in a disk group should be managed by using the same disk group policy.

If you try to associate to a server that does not support these properties, a configuration error will be generated.

Only the following storage controllers support these properties:

- LSI 6G MegaRAID SAS 9266-8i
- LSI 6G MegaRAID SAS 9271-8i
- LSI 6G MegaRAID 9265-8i
- LSI MegaRAID SAS 2208 ROMB
- LSI MegaRAID SAS 9361-8i

For the LSI MegaRAID SAS 2208 ROMB controller, these properties are supported only in the B420-M3 blade server. For the other controllers, these properties are supported in multiple rack servers.

Step 9 Click **OK**.

Creating a Storage Profile

You can create storage profile policies from the **Storage** tab in the **Navigation** pane. Additionally, you can also configure the default storage profile that is specific to a service profile from the **Servers** tab.

Procedure

Step 1 In the **Navigation** pane, click **Storage**.

Step 2 Expand **Storage** > **Storage Provisioning** > **Storage Profiles**

Step 3 Expand the node for the organization where you want to create the storage profile.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click the organization and select **Create Storage Profile**.

Step 5 In the **Create Storage Profile** dialog box, specify the storage profile **Name**. You can provide an optional **Description** for this storage profile.

Step 6 (Optional) In the **Storage Items** area, **Create Local LUNs** and add them to this storage profile.

Step 7 Click **OK**.

Creating a Specific Storage Profile

Procedure

Step 1 Expand **Servers** > **Service Profiles**.

Step 2 Expand the node for the organization that contains the service profile for which you want to create a specific storage profile.
If the system does not include multitenancy, expand the **root** node.

Step 3 Choose the service profile for which you want to create a specific storage profile.

Step 4 In the **Work** pane, click the **Storage** > **LUN Configuration** tab.

Step 5 In the **Actions** area, click **Modify Storage Profile**.

Step 6 In the **Modify Storage Profile** dialog box, click the **Specific Storage Profile** tab.

Step 7 Click **Create Specific Storage Profile**.

Step 8 (Optional) In the **Specific Storage Profile** area, complete the **Description** field to set the description of the storage profile.
Each service profile can have only one specific storage profile. Hence, the name of this storage profile is provided by default.

Step 9 In the **Storage Items** area, **Create Local LUNs** and add them to this storage profile.

Step 10 Click **OK**.

Step 11 If a confirmation dialog box displays, click **Yes**.

Deleting a Storage Profile

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click **Storage**.	
Step 2	Expand **Storage** > **Storage Provisioning** > **Storage Profiles**	
Step 3	Expand the node for the organization that contains the storage profile that you want to delete.	

	Command or Action	Purpose
Step 4	Right-click the storage profile that you want to delete and select **Delete**.	
Step 5	Click **Yes** in the confirmation box that appears.	

Configuring Local LUNs

You can create local LUNs within a storage profile policy from the **Storage** tab in the **Navigation** pane. Additionally, you can also create local LUNs within the default storage profile that is specific to a service profile from the **Servers** tab.

Procedure

Step 1 In the **Navigation** pane, click **Storage**.

Step 2 Expand **Storage** > **Storage Provisioning** > **Storage Profiles**

Step 3 Expand the node for the organization that contains the storage profile within which you want to create a local LUN.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Create Local LUN**.

Step 6 In the Create Local LUN dialog box, complete the following fields:

Name	Description
Name field	Name for the new local LUN.
Size (GB) field	Size of this LUN in GB. The size can range from 1 to 10240 GB. **Note** You do not need to specify a LUN size while claiming an orphaned LUN.
Expand To Available field	Specifies that this LUN can be expanded to use the entire available disk group. For each service profile, only one LUN can use this option.
Auto Deploy field	Whether the local LUN should be automatically deployed or not.
Select Disk Group Configuration field	The disk group configuration to be applied to this local LUN.

Step 7 (Optional) Click **Create Disk Group Policy** to create a new disk group policy for this local LUN.

Step 8 Click **OK**.

Deleting Local LUNs

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click **Storage**.	
Step 2	Expand **Storage** > **Storage Provisioning** > **Storage Profiles**	
Step 3	Expand the node for the organization that contains the storage profile from which you want to delete a local LUN.	
Step 4	Expand **Local LUNs** for the storage profile that you want and select the LUN that you want to delete.	
Step 5	Right-click the LUN that you want to delete and select **Delete**.	A confirmation dialog box appears.
Step 6	Click **Yes**.	

PCH SSD Controller Definition

Cisco UCS Manager Platform Controller Hub (PCH) Solid State Drive (SSD) Controller Definition provides a local storage configuration in storage profiles where you can configure all the disks in a single RAID or in a JBOD disk array.

The PCH Controller Definition configuration provides the following features:

- Ability to configure a single LUN RAID across two internal SSDs connected to the onboard PCH controller
- A way to configure the controller in two modes: AHCI (JBOD) and SWRAID (RAID).
- Ability to configure the PCH storage device in an Embedded Local LUN and Embedded Local Disk boot policy so precision control for boot order is achieved even with the presence of other bootable local storage devices in the server. Do not use the Local LUN or the Local JBOD options to boot from PCH disks
- Scrub policy support for the internal SSD drives. This is applicable only for the SWRAID mode. This does not apply for the AHCI and NORAID of PCH Controller modes.
- Firmware upgrade support for the internal SSD drives. Disk firmware upgrade is supported only when the PCH Controller is in SWRAID mode. It is not supported for AHCI mode.

You configure PCH controller SSDs in a storage profile policy. You enable or disable protect configuration which saves the LUN configuration even after a service profile disassociation. You choose a controller mode. The PCH controller configuration supports only these two RAID options: RAID0 and RAID1. Use No RAID configuration option for AHCI mode where all the disks connected to the controller configured as JBOD disks. The configuration deployment happens as part of the storage profile association to a service profile process.

Cisco UCS Manager supports the following the internal SSDs:

- UCSC-C240-M4L - 12 SFF disks
- UCSC-C240-M4SX - 24 SFF disks

Embedded RAID Hub Controllers are split into two controllers: SATA and sSATA (Secondary SATA. Cisco UCS Manager support for PCH Controller definition is limited only for the first SATA controller and two internal SSDs which are embedded into the riser and connected to the first SATA controller. The first SATA controller controls the two SSDs in internal riser and also the front panel drives in slot 1 to 4. The sSATA controller controls only the front panel drives in slot 5 to 8. The CPU sees these controllers as two independent devices. There are two different sets of PCI definition for SATA and sSATA controller. The Cisco UCS Manager support will be added only for the first SATA controller which manages the internal SSDs.

For the PCH Controller Definition configuration in a Cisco UCS Manager boot policy two new devices exist to select: PCH LUN and PCH Disk. EmbeddedLocalLun represents the boot device in SWRAID mode and EmbeddedLocalDisk represent the boot devices in AHCI mode.

The system uses the same scrub policy is used to scrub supported SSDs. If the scrub is Yes, configured LUNs are destroyed as part of disassociation or re-discovery. If the scrub is No, configured LUNs are saved during disassociation and re-discovery.

Cisco UCS Manager supports firmware upgrade for the internal SSDs only when the PCH Controller is in SWRAID mode. It is not supported in the AHCI mode.

Creating a Storage Profile PCH Controller Definition

The PCH Controller Definition provides a storage configuration in Storage Profiles where you can configure internal SSDs connected to a PCH controller. You create a name for the controller definition, specify whether you want the storage profile to retain the configuration even if the storage profile is disassociated from the service profile, and chose the RAID level to indicate the controller mode.

Procedure

Step 1 In the **Navigation** pane, click the **Storage** tab.

Step 2 Right-click **Storage Profiles**.

Step 3 Choose **Create Storage Profile** from the pop-up menu or click Storage Profile or click the Storage Profile link on the **Getting Started** tab.

Step 4 In the **Navigation** pane, right-click a specific **Storage Profile** and chose **Show Navigator** from the pop-up menu.

Step 5 In the **Create Storage Profile** dialog box, click the **Controller Definitions** tab and configure the following information:

Step 6 Type a storage profile **Name**.
The name can be no linger than 32 characters long.

Step 7 (Optional) Type a **Description** for this storage profile.

Step 8 Click [+] at the right of the dialog box to display the **Create PCH Controller Definition**.

Step 9 In **Create PCH Controller Definition** dialog box, configure the following information:

Name	Description
Name field	The name of the storage controller. **Note** Once you save a PCH Controller Definition, you can not modify the name from the General Tab Properties area. Enter up to 16 characters. You can use any alphanumeric characters. Special characters and spaces are not supported.
Protect Configuration check box	If checked, the storage profile retains the configuration even if the storage profile is disassociated from the service profile. **Note** If you disassociate the storage profile from a service profile with this option enabled, and then associate it with a new service profile that includes a local disk configuration policy with different properties, the server returns a configuration mismatch error and the association fails.

Name	Description
RAID Level drop-down list	

Name	Description
	This can be one of the following disk policy modes: • **No Local Storage—**(Supported for PCH SSD Controller Definition) For a diskless server 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. • **RAID 0 Striped—**(Supported for PCH SSD Controller Definition) Data is striped across all disks in the array, providing fast throughput. There is no data redundancy, and all data is lost if any disk fails. • **RAID 1 Mirrored—**(Supported for PCH SSD Controller Definition) Data is written to two disks, providing complete data redundancy if one disk fails. The maximum array size is equal to the available space on the smaller of the two drives. • **Any Configuration—**For a server configuration that carries forward the local disk configuration without any changes. • **No RAID—**For a server configuration that removes the RAID and leaves the disk MBR and payload unaltered. If you choose No RAID and you apply this policy to a server that already has an operating system with RAID storage configured, the system does not remove the disk contents. Therefore, there may be no visible differences on the server after you apply the No RAID mode. This can lead to a mismatch between the RAID configuration in the policy and the actual disk configuration shown in the Inventory > Storage tab for the server. To make sure that any previous RAID configuration information is removed from a disk, apply a scrub policy that removes all disk information after you apply the No RAID configuration mode. • **RAID 5 Striped Parity—**Data is striped across all disks in the array. Part of the capacity of each disk stores parity information that can be used to reconstruct data if a disk fails. RAID 5 provides good data throughput for applications with high read request rates. • **RAID 6 Striped Dual Parity—**Data is striped across all disks in the array and two parity disks are used to provide protection against the failure of up to two physical disks. In each row of data blocks, two sets of parity data are stored. • **RAID 10 Mirrored and Striped—**RAID 10 uses mirrored pairs of disks to provide complete data redundancy and high throughput rates. • **RAID 50 Striped Parity and Striped—**Data is striped across multiple striped parity disk sets to provide high throughput and multiple disk failure tolerance. • **RAID 60 Striped Dual Parity and Striped—**Data is striped

Name	Description
	across multiple striped dual parity disk sets to provide high throughput and greater disk failure tolerance. **Note** Some Cisco UCS servers require a license for certain RAID configuration options. When Cisco UCS Manager associates a service profile containing this local disk policy with a server, Cisco UCS Manager verifies that the selected RAID option is properly licensed. If there are issues, Cisco UCS Manager displays a configuration error during the service profile association. For RAID license information for a specific Cisco UCS server, see the Hardware Installation Guide for that server.

Step 10 Click OK.
The new PCH Controller Definition appears in the navigation pane.

Modifying a Service Profile PCH Controller Definition

Procedure

Step 1 In the **Navigation** pane, click the **Storage** tab.
Step 2 Expand **Storage Profiles** to the specific storage profile name that you want.
Step 3 Expand **Controller Definitions** and click the specific controller definition that you want.
Step 4 On the **General** tab, modify the following information:

Name	Description
Name field	The name of the storage controller. **Note** Once you save a PCH Controller Definition, you can not modify the name from the General Tab Properties area. Enter up to 16 characters. You can use any alphanumeric characters. Special characters and spaces are not supported.
Protect Configuration check box	If checked, the storage profile retains the configuration even if the storage profile is disassociated from the service profile. **Note** If you disassociate the storage profile from a service profile with this option enabled, and then associate it with a new service profile that includes a local disk configuration policy with different properties, the server returns a configuration mismatch error and the association fails.

Name	Description
RAID Level drop-down list	

Name	Description
	This can be one of the following disk policy modes: • **No Local Storage—**(Supported for PCH SSD Controller Definition) For a diskless server 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. • **RAID 0 Striped—**(Supported for PCH SSD Controller Definition) Data is striped across all disks in the array, providing fast throughput. There is no data redundancy, and all data is lost if any disk fails. • **RAID 1 Mirrored—**(Supported for PCH SSD Controller Definition) Data is written to two disks, providing complete data redundancy if one disk fails. The maximum array size is equal to the available space on the smaller of the two drives. • **Any Configuration—**For a server configuration that carries forward the local disk configuration without any changes. • **No RAID—**For a server configuration that removes the RAID and leaves the disk MBR and payload unaltered. If you choose No RAID and you apply this policy to a server that already has an operating system with RAID storage configured, the system does not remove the disk contents. Therefore, there may be no visible differences on the server after you apply the No RAID mode. This can lead to a mismatch between the RAID configuration in the policy and the actual disk configuration shown in the Inventory > Storage tab for the server. To make sure that any previous RAID configuration information is removed from a disk, apply a scrub policy that removes all disk information after you apply the No RAID configuration mode. • **RAID 5 Striped Parity—**Data is striped across all disks in the array. Part of the capacity of each disk stores parity information that can be used to reconstruct data if a disk fails. RAID 5 provides good data throughput for applications with high read request rates. • **RAID 6 Striped Dual Parity—**Data is striped across all disks in the array and two parity disks are used to provide protection against the failure of up to two physical disks. In each row of data blocks, two sets of parity data are stored. • **RAID 10 Mirrored and Striped—**RAID 10 uses mirrored pairs of disks to provide complete data redundancy and high throughput rates. • **RAID 50 Striped Parity and Striped—**Data is striped across multiple striped parity disk sets to provide high throughput and multiple disk failure tolerance. • **RAID 60 Striped Dual Parity and Striped—**Data is striped

Name	Description
	across multiple striped dual parity disk sets to provide high throughput and greater disk failure tolerance. **Note** Some Cisco UCS servers require a license for certain RAID configuration options. When Cisco UCS Manager associates a service profile containing this local disk policy with a server, Cisco UCS Manager verifies that the selected RAID option is properly licensed. If there are issues, Cisco UCS Manager displays a configuration error during the service profile association. For RAID license information for a specific Cisco UCS server, see the Hardware Installation Guide for that server.

Step 5 Click OK.
The system displays whether it saved the modified PCH Controller Definition successfully.

Deleting a Storage Profile PCH Controller Definition

Procedure

Step 1 In the **Navigation** pane, click the **Storage** tab.
Step 2 Expand **Storage Profiles**.
Step 3 Expand **PCH Controller Definitions**.
Step 4 In the **Navigation** pane, click the specific Controller Definition that you want.
Step 5 In the **General** tab **Actions** area, click **Delete**.
Step 6 Confirm whether you want to delete the definition.
The system displays whether it deleted the definition successfully. if not, see PCH Controller Definition Configuration Troubleshooting , on page 641
Step 7 If successfully deleted, click OK.

PCH Controller Definition Configuration Troubleshooting

PCH Controller Definition Creation

Unsuccessful PCH Controller Definition configuration exists under the following situations:

- You try to configure a Controller definition for an unsupported server model

- You try to use the legacy local disk configuration policy and also configures the PCH storage in storage profile
- You try to configure same controller using storage profile controller definition and also by using storage profile Local LUN configuration interface
- If the **Protect Configuration** checkbox is ON and you configured the RAID Type differently than the deployed configuration in SWRAID mode.
- If the **Protect Configuration** checkbox is ON and the RAID Type does not match the present controller mode.

Warning Any configuration change in the PCH storage configuration (like Controller mode change, RAID level change or controller qualifier change) for an already associated server triggers a PNUOS boot to happen causing a down time for the host OS.

Boot Policy

A configuration error occurs for any of the following cases:

- You select PCH Disk in boot policy but the primary or secondary target path slot number did not match with any of the inventoried internal SSD slot numbers.
- You select both PCH LUN and PCH Disk at the sam a time in the boot policy.

Firmware

For an incompatible software combination, there will not be any configuration error to at the time of association. However the storage configuration for the PCH SSD controller might fail or might not be deployed during association if you do not use the supported software combinations. Also, booting from the PCH SSD controller internal SSD might fail at the end of association for an incompatible software combination.

Associating a Storage Profile with an Existing Service Profile

You can associate a storage profile with an existing service profile or a new service profile. *Creating a Service Profile with the Expert Wizard* in the *Cisco UCS Manager GUI Configuration Guide, Release 2.2* provides more information about associating a storage profile with a new service profile.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that contains the service profile that you want to associate with a storage profile.

Step 4 Choose the service profile that you want to associate with a storage profile.

Step 5 In the **Work** pane, click the **Storage** tab.

Step 6 Click the **LUN Configuration** subtab.

Step 7 In the **Actions** area, click **Modify Storage Profile**. The **Modify Storage Profile** dialog box appears.

Step 8 Click the **Storage Profile Policy** tab.

Step 9 To associate an existing storage profile with this service profile, select the storage profile that you want to associate from the **Storage Profile** drop-down list, and click **OK**. The details of the storage profile appear in the **Storage Items** area.

Step 10 To create a new storage profile and associate it with this service profile, click **Create Storage Profile**, complete the required fields, and click **OK**. Creating a Storage Profile, on page 630 provides more information on creating a new storage profile.

Step 11 (Optional) To dissociate the service profile from a storage profile, select **No Storage Profile** from the **Storage Profile** drop-down list, and click **OK**.

Displaying Details of All Local LUNs Inherited By a Service Profile

Storage profiles can be defined under org and as a dedicated storage profile under service profile. Thus, a service profile inherits local LUNs from both possible storage profiles. It can have a maximum of 2 such local LUNs. You can display the details of all local LUNs inherited by a service profile by using the following command:

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization that contains the service profile that you want to display.

Step 4 Choose the service profile whose inherited local LUNs you want to display.

Step 5 In the **Work** pane, click the **Storage** tab.

Step 6 Click the **LUN Configuration** subtab, and then click the **Local LUNs** tab.
Displays the following detailed information about all the local LUNs inherited by the specified service profile:

- **Name—**LUN name in the storage profile.
- **Admin State—**Specifies whether a local LUN should be deployed or not. Admin state can be **Online** or **Undeployed**.

When the local LUN is being referenced by a service profile, if the auto-deploy status is **no-auto-deploy** then the admin state will be **Undeployed** , else it will be **Online**. After the local LUN is referenced by a service profile, any change made to this local LUN's auto-deploy status is not reflected in the admin state of the LUN inherited by the service profile.

- **RAID Level—**Summary of the RAID level of the disk group used.
- **Provisioned Size (GB)—**Size, in GB, of the LUN specified in the storage profile.
- **Assigned Size (MB)—**Size, in MB, assigned by UCSM.
- **Config State—**State of LUN configuration. The states can be one of the following:
  - **Applying—**Admin state is online, the LUN is associated with a server, and the virtual drive is being created.
  - **Applied—**Admin state is online, the LUN is associated with a server, and the virtual drive is created.
  - **Apply Failed—**Admin stage is online, the LUN is associated with a server, but the virtual drive creation failed.
  - **Not Applied—**The LUN is not associated with a server, or the LUN is associated with a service profile, but admin state is undeployed.
- **Referenced LUN Name—**The preprovisioned virtual drive name, or UCSM-generated virtual drive name.
- **Deploy Name—**The virtual drive name after deployment.
- **ID—**LUN ID.
- **Order—**Order of LUN visibility to the server.
- **Bootable—**Whether the LUN is bootable or not.
- **LUN New Name—**New name of the LUN.
- **Drive State—**State of the virtual drive. The states are:
  - **Unknown**
  - **Optimal**
  - **Degraded**
  - **Inoperable**
  - **Partially Degraded**

Importing Foreign Configurations for a RAID Controller on a Blade Server

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click **Equipment**.	
Step 2	Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.	
Step 3	Choose the server of the RAID controller for which you want to import foreign configurations.	
Step 4	In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.	
Step 5	Click the **Controller** subtab.	
Step 6	In the **Actions** area, click **Import Foreign Configuration**.	

Importing Foreign Configurations for a RAID Controller on a Rack Server

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click **Equipment**.	
Step 2	Expand **Equipment** > **Rack Mounts** > **Servers**.	
Step 3	Choose the server of the RAID controller for which you want to import foreign configurations.	
Step 4	In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.	
Step 5	Click the **Controller** subtab.	
Step 6	In the **Actions** area, click **Import Foreign Configuration**.	

Configuring Local Disk Operations on a Blade Server

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click **Equipment**.	
Step 2	Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.	

	Command or Action	Purpose
Step 3	Choose the server for which you want to configure local disk operations.	
Step 4	In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.	
Step 5	Click the **Disks** subtab.	
Step 6	Right-click the disk that you want and select one of the following operations: • **Clear Foreign Configuration State—**Clears any foreign configuration that exists in a local disk when it is introduced into a new configuration. • **Set Unconfigured Good—**Specifies that the local disk can be configured. • **Set Prepare For Removal—**Specifies that the local disk is marked for removal from the chassis. • **Set Undo Prepare For Removal—**Specifies that the local disk is no longer marked for removal from the chassis. • **Mark as Dedicated Hot Spare—**Specifies the local disk as a dedicated hot spare. You can select the virtual drive from the available drives. • **Remove Hot Spare—**Specifies that the local disk is no longer a hot spare. • **Set JBOD to Unconfigured Good—**Specifies that the new local disk can be configured after being marked as Unconfigured Good.	

Configuring Local Disk Operations on a Rack Server

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click **Equipment**.	
Step 2	Expand **Equipment** > **Rack Mounts** > **Servers**.	
Step 3	Choose the server for which you want to configure local disk operations.	
Step 4	In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.	
Step 5	Click the **Disks** subtab.	
Step 6	Right-click the disk that you want and select one of the following operations: • **Clear Foreign Configuration State—**Clears any foreign configuration that exists in a local disk when it is introduced into a new configuration. • **Set Unconfigured Good—**Specifies that the local disk can be configured. • **Set Prepare For Removal—**Specifies that the local disk is marked for removal.	

	Command or Action	Purpose
	• **Set Undo Prepare For Removal—**Specifies that the local disk is no longer marked for removal. • **Mark as Dedicated Hot Spare—**Specifies the local disk as a dedicated hot spare. You can select the virtual drive from the available drives. • **Remove Hot Spare—**Specifies that the local disk is no longer a hot spare. • **Set JBOD to Unconfigured Good—**Specifies that the new local disk can be configured after being marked as **Unconfigured Good**.	

Configuring Virtual Drive Operations

The following operations can be performed only on orphaned virtual drives:

- Delete an orphaned virtual drive
- Rename an orphaned virtual drive

Deleting an Orphan Virtual Drive on a Blade Server

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click **Equipment**.	
Step 2	Expand **Equipment** > **Chassis** > ***Chassis Number*** > **Servers**.	
Step 3	Choose the server for which you want to delete an orphan virtual drive.	
Step 4	In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.	
Step 5	Click the **LUNs** subtab.	
Step 6	Right-click the virtual drive that you want and select **Delete Orphaned LUN**.	A confirmation dialog box appears.
Step 7	Click **Yes**.	

Deleting an Orphan Virtual Drive on a Rack Server

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click **Equipment**.	
Step 2	Expand **Equipment** > **Rack Mounts** > **Servers**.	
Step 3	Choose the server for which you want to delete an orphan virtual drive.	
Step 4	In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.	
Step 5	Click the **LUNs** subtab.	
Step 6	Right-click the virtual drive that you want and select **Delete Orphaned LUN**.	A confirmation dialog box appears.
Step 7	Click **Yes**.	

Renaming an Orphan Virtual Drive on a Blade Server

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click **Equipment**.	
Step 2	Expand **Equipment** > **Chassis** > ***Chassis Number*** > **Servers**.	
Step 3	Choose the server for which you want to rename an orphan virtual drive.	
Step 4	In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.	
Step 5	Click the **LUNs** subtab.	
Step 6	Right-click the virtual drive that you want and select **Rename Referenced LUN**.	
Step 7	In the **Rename Referenced LUN** dialog box that appears, enter the new **LUN Name**.	
Step 8	Click **OK**.	

Renaming an Orphan Virtual Drive on a Rack Server

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click **Equipment**.	
Step 2	Expand **Equipment** > **Rack Mounts** > **Servers**.	
Step 3	Choose the server for which you want to rename an orphan virtual drive.	
Step 4	In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.	
Step 5	Click the **LUNs** subtab.	
Step 6	Right-click the virtual drive that you want and select **Rename Referenced LUN**.	
Step 7	In the **Rename Referenced LUN** dialog box that appears, enter the new **LUN Name**.	
Step 8	Click **OK**.	

Boot Policy for Local Storage

You can specify the primary boot device for a storage controller as a local LUN or a JBOD disk. Each storage controller can have one primary boot device. However, in a storage profile, you can set only one device as the primary boot LUN.

Configuring the Boot Policy for a Local Device

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click **Servers**.	
Step 2	Expand **Servers** > **Policies**.	
Step 3	Expand the node for the organization where you want to create the policy.	If the system does not include multitenancy, expand the root node.
Step 4	Select the boot policy that you want to configure.	
Step 5	In the **Work** pane, click the **General** tab.	
Step 6	Click the down arrows to expand the **Local Devices** area.	
Step 7	Click **Add Local LUN** to configure the boot order of the local LUN.	

	Command or Action	Purpose
Step 8	To configure the local LUN as the primary boot device, select **Primary**.	
Step 9	In the **LUN Name** field, enter the name of the LUN to be configured as the primary boot device.	
Step 10	Click **OK**.	

Configuring the Boot Policy for a Local JBod Device

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click **Servers**.	
Step 2	Expand **Servers** > **Policies**.	
Step 3	Expand the node for the organization where you want to create the policy.	If the system does not include multitenancy, expand the root node.
Step 4	Select the boot policy that you want to configure.	
Step 5	In the **Work** pane, click the **General** tab.	
Step 6	Click the down arrows to expand the **Local Devices** area.	
Step 7	Click **Add Local JBod** to configure the local JBod device as the primary boot device.	BOD is supported only on the following servers: • Cisco UCS B200 M3 blade server • Cisco UCS B260 M4 blade server • Cisco UCS B460 M4 blade server • Cisco UCS B200 M4 blade server • Cisco UCS C220 M4 rack-mount server • Cisco UCS C240 M4 rack-mount server • Cisco UCS C460 M4 rack-mount server
Step 8	In the **Disk Slot Number** field, enter the slot number of the JBod disk to be configured as the primary boot device.	
Step 9	Click **OK**.	

Local LUN Operations in a Service Profile

Preprovisioning a LUN Name

Preprovisioning a LUN name can be done only when the admin state of the LUN is **Undeployed**. If this LUN name exists and the LUN is orphaned, its is claimed by the service profile. If this LUN does not exist, a new LUN is created with the specified name.

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click **Servers**.	
Step 2	Expand **Servers** > **Service Profiles** > *Service_Profile_Name*.	
Step 3	In the **Work** pane, click the **Storage** tab.	
Step 4	Click the **LUN Configuration** tab.	
Step 5	In the **Local LUNs** subtab, right-click the LUN for which you want to preprovision a LUN name and select **Pre-Provision LUN Name**.	
Step 6	In the **Set Pre-Provision LUN Name** dialog box, enter the LUN name.	
Step 7	Click **OK**.	

Claiming an Orphan LUN

Claiming an orphan LUN can be done only when the admin state of the LUN is **Undeployed**. You can explicitly change the admin state of the LUN to **Undeployed** for claiming an orphan LUN.

If the LUN name is empty, set a LUN name before claiming it.

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click **Servers**.	
Step 2	Expand **Servers** > **Service Profiles** > *Service_Profile_Name*.	
Step 3	In the **Work** pane, click the **Storage** tab.	
Step 4	Click the **LUN Configuration** tab.	
Step 5	In the **Local LUNs** subtab, right-click the LUN that you want to claim and select **Claim Orphan LUN**.	

	Command or Action	Purpose
Step 6	In the **Claim Orphan LUN** dialog box that appears, select an orphaned LUN.	
Step 7	Right-click the LUN and select **Set Admin State**.	
Step 8	In the **Set Admin State** dialog box that appears, select **Undeployed** to undeploy a LUN and claim ownership.	
Step 9	Click **OK**.	

Deploying and Undeploying a LUN

You can deploy or undeploy a LUN. If the admin state of a local LUN is **Undeployed**, the reference of that LUN is removed and the LUN is not deployed.

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click **Servers**.	
Step 2	Expand **Servers** > **Service Profiles** > ***Service_Profile_Name***.	
Step 3	In the **Work** pane, click the **Storage** tab.	
Step 4	Click the **LUN Configuration** tab.	
Step 5	In the **Local LUNs** subtab, right-click the LUN that you want to deploy or undeploy and select **Set Admin State**.	
Step 6	In the **Set Admin State** dialog box that appears, select **Online** to deploy a LUN or **Undeployed** to undeploy a LUN.	
Step 7	Click **OK**.	

Renaming a Service Profile Referenced LUN

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click **Servers**.	
Step 2	Expand **Servers** > **Service Profiles** > ***Service_Profile_Name***.	
Step 3	In the **Work** pane, click the **Storage** tab.	
Step 4	Click the **LUN Configuration** tab.	

	Command or Action	Purpose
Step 5	In the **Local LUNs** subtab, right-click the LUN for which you want to rename the referenced LUN, and select **Rename Referenced LUN**.	
Step 6	In the **Rename Referenced LUN** dialog box that appears, enter the new name of the referenced LUN.	
Step 7	Click **OK**.	

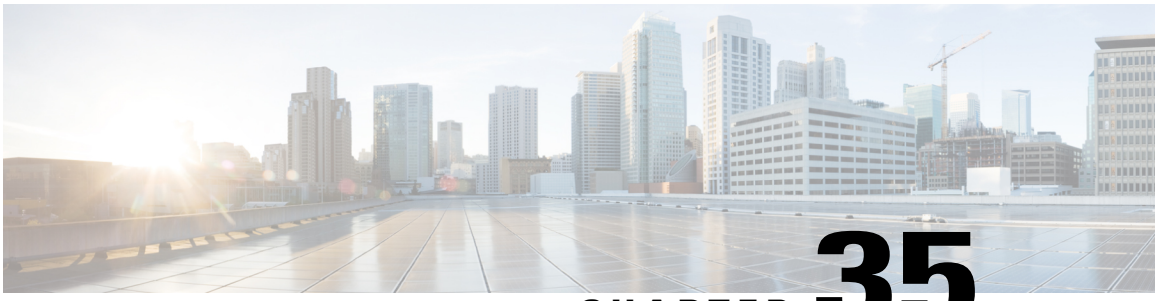

CHAPTER 35

Managing Power in Cisco UCS

This chapter includes the following sections:

Power Capping in Cisco UCS

You can control the maximum power consumption on a server through power capping, as well as manage the power allocation in the Cisco UCS Manager for the UCS B-Series Blade Servers, UCS Mini, and mixed UCS domains.

UCS Manager supports power capping on the following servers:

- UCS Mini 6324
- UCS 6300 Series Fabric Interconnects

You can use Policy Driven Chassis Group Power Cap, or Manual Blade Level Power Cap methods to allocate power that applies to all of the servers in a chassis.

Cisco UCS Manager provides the following power management policies to help you allocate power to your servers:

Power Management Policies	Description
Power Policy	Specifies the redundancy for power supplies in all chassis in a Cisco UCS domain.
Power Control Policies	Specifies the priority to calculate the initial power allocation for each blade in a chassis.
Global Power Allocation	Specifies the Policy Driven Chassis Group Power Cap or the Manual Blade Level Power Cap to apply to all servers in a chassis.
Global Power Profiling	Specifies how the power cap values of the servers are calculated. If it is enabled, the servers will be profiled during discovery through benchmarking. This policy applies when the Global Power Allocation Policy is set to Policy Driven Chassis Group Cap.

Rack Server Power Management

Power capping is not supported for rack servers.

Power Management Precautions

If the CIMC is reset, the power monitoring functions of Cisco UCS become briefly unavailable until the CIMC reboots. Typically, the reset only takes 20 seconds; however, it is possible that the peak power cap can exceed during that time. To avoid exceeding the configured power cap in a low power-capped environment, consider staggering the rebooting or activation of CIMCs.

UCS Power Policy

Power Policy for Cisco UCS Servers

The power policy is global and is inherited by all of the chassis' managed by the Cisco UCS Manager instance. You can add the power policy to a service profile to specify the redundancy for power supplies in all chassis' in the Cisco UCS domain. This policy is also known as the PSU policy.

For more information about power supply redundancy, see *Cisco UCS 5108 Server Chassis Hardware Installation Guide*.

Configuring the Power Policy

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Click the **Equipment** node.

Step 3 In the **Work** pane, click the **Policies** tab.

Step 4 Click the **Global Policies** subtab.

Step 5 In the **Power Policy** area, click one of the following radio buttons in the **Redundancy** field:

- **Non Redundant—**Cisco UCS Manager turns on the minimum number of power supplies (PSUs) needed and balances the load between them. If any additional PSUs are installed, Cisco UCS Manager sets them to a "turned-off" state. If the power to any PSU is disrupted, the system may experience an interruption in service until Cisco UCS Manager can activate a new PSU and rebalance the load.

  In general, a Cisco UCS chassis requires at least two PSUs for non-redundant operation. Only smaller configurations (requiring less than 2500W) can be powered by a single PSU.

- **N+1—**The total number of PSUs to satisfy non-redundancy, plus one additional PSU for redundancy, are turned on and equally share the power load for the chassis. If any additional PSUs are installed, Cisco UCS Manager sets them to a "turned-off" state. If the power to any PSU is disrupted, Cisco UCS Manager can recover without an interruption in service.

  In general, a Cisco UCS chassis requires at least three PSUs for N+1 operation.

- **Grid—**Two power sources are turned on, or the chassis requires greater than N+1 redundancy. If one source fails (which causes a loss of power to one or two PSUs), the surviving PSUs on the other power circuit continue to provide power to the chassis.

For more information about power supply redundancy, see *Cisco UCS 5108 Server Chassis Hardware Installation Guide*.

Step 6 Click **Save Changes**.

Global Power Allocation Policy Configuration

Global Power Allocation Policy

The Global Power Allocation Policy allows you to specify the Policy Driven Chassis Group Power Cap or Manual Blade-level Power Cap power allocation method applied to servers in a chassis.

Cisco recommends using the default Policy Driven Chassis Group Power Cap power allocation method.

Important Any change to the Manual Blade level Power Cap configuration results in the loss of any groups or configuration options set for the Policy Driven Chassis Group Power Cap.

Configuring the Global Power Allocation Policy

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Click the **Equipment** node.

Step 3 In the **Work** pane, click the **Policies** tab.

Step 4 Click the **Global Policies** subtab.

Step 5 In the **Global Power Allocation Policy** area, click one of the following radio buttons in the **Allocation Method** field to determine the power cap management mode used in the Cisco UCS domain:

- **Manual Blade Level Cap—**Power allocation is configured on each individual blade server in all chassis. If you select this option, you cannot create power groups.
- **Policy Driven Chassis Group Cap—**Power allocation is configured at the chassis level through power control policies included in the associated service profiles. If you select this option, you can also create power groups that contain one or more chassis in the Cisco UCS domain.

By default, power allocation is done for each chassis through a power control policy.

Step 6 Click **Save Changes**.

Policy Driven Power Capping

Policy Driven Chassis Group Power Capping

When you select the Policy Driven Chassis Group Power Cap in the Global Cap Policy, Cisco UCS can maintain the over-subscription of servers without risking power failures. You can achieve over-subscription through a two-tier process. For example, at the chassis level, Cisco UCS divides the amount of power available among members of the power group, and at the blade level, the amount of power allotted to a chassis is divided among blades based on priority.

Each time a service profile is associated or disassociated, Cisco UCS Manager recalculates the power allotment for each blade server within the chassis. If necessary, power from lower-priority service profiles is redistributed to higher-priority service profiles.

UCS power groups cap power in less than one second to safely protect data center circuit breakers. A blade must stay at its cap for 20 seconds before the chassis power distribution is optimized. This is intentionally carried out over a slower timescale to prevent reacting to transient spikes in demand.

Note The system reserves enough power to boot a server in each slot, even if that slot is empty. This reserved power cannot be leveraged by servers requiring more power. Blades that fail to comply with the power cap are penalized.

Power Groups

Power Groups in UCS Manager

A power group is a set of chassis that all draw power from the same power distribution unit (PDU). In Cisco UCS Manager, you can create power groups that include one or more chassis, then set a peak power cap in AC watts for that power grouping.

Implementing power capping at the chassis level requires the following:

- IOM, CIMC, and BIOS version 1.4 or higher
- Two Power Supply Units (PSUs)

The peak power cap is a static value that represents the maximum power available to all blade servers within a given power group. If you add or remove a blade from a power group, but do not manually modify the peak power value, the power group adjusts the peak power cap to accommodate the basic power-on requirements of all blades within that power group.

A minimum of 890 AC watts should be set for each chassis. This converts to 800 watts of DC power, which is the minimum amount of power required to power an empty chassis. To associate a half-width blade, the group cap needs to be set to 1475 AC watts. For a full-width blade, it needs to be set to 2060 AC watts.

After a chassis is added to a power group, all service profile associated with the blades in the chassis become part of that power group. Similarly, if you add a new blade to a chassis, that blade inherently becomes part of the chassis' power group.

Note Creating a power group is not the same as creating a server pool. However, you can populate a server pool with members of the same power group by creating a power qualifier and adding it to server pool policy.

When a chassis is removed or deleted, the chassis gets removed from the power group.

UCS Manager supports explicit and implicit power groups.

- **Explicit**: You can create a power group, add chassis' and racks, and assign a budget for the group.
- **Implicit**: Ensures that the chassis is always protected by limiting the power consumption within safe limits. By default, all chassis that are not part of an explicit power group are assigned to the default group and the appropriate caps are placed. New chassis that connect to UCS Manager are added to the default power group until you move them to a different power group.

The following table describes the error messages you might encounter while assigning power budget and working with power groups.

Error Message	Cause	Recommended Action

`Insufficient budget for power group POWERGROUP_NAME`
and/or
`Chassis N cannot be capped as group cap is low. Please consider raising the cap.`
and/or
`Admin committed insufficient for power group GROUP_NAME, using previous value N`
and/or
`Power cap application failed for chassis N`
	One of these messages displays if you did not meet the minimum limit when assigning the power cap for a chassis, or the power requirement increased because of the addition of blades or change of power policies.	Increase the power cap limit to the **Minimum Power Cap for Allowing Operations (W)** value displayed on the **Power Group** page for the specified power group.

`Chassis N cannot be capped as the available PSU power is not enough for the chassis and the blades. Please correct the problem by checking input power or replace the PSU`	Displays when the power budget requirement for the chassis is more than the PSU power that is available.	Check the PSU input power and redundancy policy to ensure that enough power is available for the chassis. If a PSU failed, replace the PSU.
`Power cap application failed for server N`	Displays when the server is consuming more power than allocated and cannot be capped, or the server is powered on when no power is allocated.	Do not power on un-associated servers.
`P-State lowered as consumption hit power cap for server`	Displays when the server is capped to reduce the power consumption below the allocated power.	This is an information message. If a server should not be capped, in the service profile set the value of the power control policy **Power Capping** field to **no-cap**.
`Chassis N has a mix of high-line and low-line PSU input power sources.`	This fault is raised when a chassis has a mix of high-line and low-line PSU input sources connected.	This is an unsupported configuration. All PSUs must be connected to similar power sources.

Creating a Power Group

Before You Begin

Make sure that the Global Power Allocation Policy is set to **Policy Driven Chassis Group Cap** on the **Global Policies** tab.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Click the **Equipment** node.

Step 3 In the **Work** pane, click the **Policies** tab.

Step 4 Click the **Power Groups** subtab.

Step 5 On the icon bar to the right of the table, click +.
If the + icon is disabled, click an entry in the table to enable it.

Step 6 On the first page of the **Create Power Group** wizard, complete the following fields:

a) Enter a unique name and description for the power group.
This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

b) Click **Next**.

Step 7 On the **Add Chassis Members** page of the **Create Power Group** wizard, do the following:

a) In the **Chassis** table, choose one or more chassis to include in the power group.

b) Click the >> button to add the chassis to the **Selected Chassis** table that displays all chassis included in the power group.
You can use the << button to remove one or more chassis from the power group.

c) Click **Next**.

Step 8 On the **Power Group Attributes** page of the **Create Power Group** wizard, do the following:

a) Complete the following fields:

Name	Description
Power Cap field	The maximum peak power (in watts) available to the power group. Enter an integer between 0 and 10000000.
Enable Dynamic Reallocation field	This can be one of the following: • **Chassis—**Cisco UCS monitors power usage and changes the blade allocations as required to maximize power utilization. • **None—**Blade allocations are not adjusted dynamically.

b) Click **Finish**.

Adding a Chassis to a Power Group

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Click the **Equipment** node.

Step 3 In the **Work** pane, click the **Power Groups** tab.

Step 4 Right-click the power group to which you want to add a chassis and choose **Add Chassis Members**.

Step 5 In the **Add Chassis Members** dialog box, do the following:

a) In the **Chassis** table, choose one or more chassis to include in the power group.

b) Click the **>>** button to add the chassis to the **Selected Chassis** table that displays all chassis included in the power group.
You can use the **<<** button to remove one or more chassis from the power group.

c) Click **OK**.

Removing a Chassis from a Power Group

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Click the **Equipment** node.

Step 3 In the **Work** pane, click the **Power Groups** tab.

Step 4 Expand the power group from which you want to remove a chassis.

Step 5 Right-click the chassis that you want to remove from the power group and choose **Delete**.

Step 6 If a confirmation dialog box displays, click **Yes**.

Deleting a Power Group

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Click the **Equipment** node.
Step 3 In the **Work** pane, click the **Power Groups** tab.
Step 4 Right-click the power group that you want to delete and choose **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Power Control Policy in UCS Manager

Power Control Policy

Cisco UCS uses the priority set in the power control policy along with the blade type and configuration to calculate the initial power allocation for each blade within a chassis. During normal operation, the active blades within a chassis can borrow power from idle blades within the same chassis. If all blades are active and reach the power cap, service profiles with higher priority power control policies take precedence over service profiles with lower priority power control policies.

Priority is ranked on a scale of 1-10, where 1 indicates the highest priority and 10 indicates lowest priority. The default priority is 5.

For mission-critical application a special priority called no-cap is also available. Setting the priority to no-cap prevents Cisco UCS from leveraging unused power from a particular server. With this setting, the server is allocated the maximum amount of power possible for that type of server.

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

Creating a Power Control Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Policies**.
Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **Power Control Policies** and choose **Create Power Control Policy**.

Step 5 In the **Create Power Control Policy** dialog box, complete the following fields:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Fan Speed Policy drop-down	Fan speed is for rack servers only. This drop-down option was introduced in Cisco UCS Manager Release 2.2(6). The fan speed can be one of the following: • **Low Power—**The fan runs at the minimum speed required to keep the server cool. • **Balanced—**The fan runs faster when needed based on the heat generated by the server. When possible, the fan returns to the minium required speed. • **Performance—**The fan is kept at the speed needed for better server performance. This draws more power but means the fan is already at speed if the server begins to heat up. • **High Power—**The fan is kept at an even higher speed that emphasizes performance over power consumption. • **Max Power—**The fan is kept at the maximum speed at all times. This option provides the most cooling and uses the most power. • **Any—**The server determines the optimal fan speed. **Note** The **Fan Speed Policy**

Name	Description
Owner field	This can be one of the following: • **Local—**This policy is available only to service profiles and service profile templates in this Cisco UCS domain. • **Pending Global—**Control of this policy is being transferred to Cisco UCS Central. Once the transfer is complete, this policy will be available to all Cisco UCS domains registered with Cisco UCS Central. • **Global—**This policy is managed by Cisco UCS Central. Any changes to this policy must be made through Cisco UCS Central.
Power Capping field	What happens to a server when the demand for power within a power group exceeds the power supply. This can be one of the following: • **No Cap—**The server runs at full capacity regardless of the power requirements of the other servers in its power group. • **cap—**The server is allocated a minimum amount of power capacity based on the server's priority relative to the other servers in its server group. If more power becomes available, Cisco UCS allows the capped servers to exceed their original allocations. It only lowers the allocations if there is a drop in the total power available to the power group. When you select **cap**, Cisco UCS Manager GUI displays the **Priority** field.
Priority field	The priority the server has within its power group when power capping is in effect. Enter an integer between 1 and 10, where 1 is the highest priority.

Step 6 Click **OK**.

What to Do Next

Include the policy in a service profile or service profile template.

Deleting a Power Control Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Policies** > *Organization_Name*.
Step 3 Expand the **Power Control Policies** node.
Step 4 Right-click the policy you want to delete and select **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

Blade Level Power Capping

Manual Blade Level Power Cap

When manual blade-level power cap is configured in the global cap policy, you can set a power cap for each blade server in a Cisco UCS domain.

The following configuration options are available:

- **Watts—**You can specify the maximum amount of power that the server can consume at one time. This maximum can be any amount between 0 watts and 1100 watts.
- **Unbounded—**No power usage limitations are imposed on the server. The server can use as much power as it requires.

If the server encounters a spike in power usage that meets or exceeds the maximum configured for the server, Cisco UCS Manager does not disconnect or shut down the server. Instead, Cisco UCS Manager reduces the power that is made available to the server. This reduction can slow down the server, including a reduction in CPU speed.

Note If you configure the manual blade-level power cap using **Equipment** > **Policies** > **Global Policies** > **Global Power Allocation Policy**, the priority set in the Power Control Policy is no longer relevant.

Setting the Blade-Level Power Cap for a Server

Before You Begin

Make sure the global power allocation policy is set to **Manual Blade Level Cap** on the **Global Policies** tab.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.

Step 3 Choose the server for which you want to set the power budget.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Power Budget** area, do the following:

a) Click the **Expand** icon to the right of the heading to display the fields.

b) Complete the following fields:

Name	Description
Admin Status field	Whether this server is power capped. This can be one of the following: • **Unbounded—**The server is not power capped under any circumstances. • **Enabled—**The Cisco UCS Manager GUI displays the **Watts** field. **Note** Power capping goes into effect only if there is insufficient power available to the chassis to meet the demand. If there is sufficient power, the server can use as many watts as it requires.
Watts field	The maximum number of watts that the server can use if there is not enough power to the chassis to meet the demand. The value range is from 0 and 10000000.

Step 6 Click **Save Changes**.

Viewing the Blade Level Power Cap

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis**.

Step 3 Choose the chassis for which you want to view the server power usage.

Step 4 Do one of the following:

• To view the power usage for all servers in the chassis, click the **Power** tab in the **Work** pane.

- To view the power usage for one server in the chassis, expand the chassis and click the server. Then click the **Power** tab in the **Work** pane.

Step 5 If necessary, expand the **Motherboards** node to view the power counters.

Power Sync Policy

Cisco UCS Manager includes a global (default) power sync policy to address power synchronization issues between the associated service profiles and the servers. You can use the power sync policy to synchronize the power state when the desired power state of the service profile differs from the actual power state of the server. The power sync policy allows you to control when to synchronize the desired power state on the associated service profiles for M-series modular servers, rack-mount servers, and blade servers. The power sync policy does not affect other power-related policies.

The power sync policy applies to all the service profiles by default. You cannot delete the default power sync policy, but you can edit the default policy. You can create your own power sync policies and apply them to the service profiles. You can also create a power sync policy that is specific to a service profile and it always takes precedence over the default policy.

Cisco UCS Manager creates a fault on the associated service profile when the power sync policy referenced in the service profile does not exist. Cisco UCS Manager automatically clears the fault once you create a power sync policy for the specified service profile or change the reference to an existing policy in the service profile.

Power Synchronization Behavior

Cisco UCS Manager synchronizes the power state only when the actual power state of the server is OFF. The current power synchronization behavior is based on the actual power state and the desired power state after shallow association occurs.
For example, the following events trigger shallow association:

- Fabric Interconnects(FI) and IOM disconnected.
- IOM reset
- FI power loss or reboot
- Chassis reacknowledgment
- Chassis power loss
- Service profile change

The following table describes the current power synchronization behavior:

Event	Desired Power State	Actual Power State Before Event	Actual Power State After Event
Shallow Association	ON	OFF	ON
Shallow Association	OFF	OFF	OFF

Event	Desired Power State	Actual Power State Before Event	Actual Power State After Event
Shallow Association	ON	ON	ON
Shallow Association	OFF	ON	ON

Creating a Power Sync Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Right-click **Power Sync Policies** and choose **Create Power Sync Policy**.

Step 5 In the **Create Power Sync Policy** dialog box, complete the following fields:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Owner field	The owner of the policy. This can be one of the following: • **Local—**This policy is available only to service profiles and service profile templates in this Cisco UCS domain. • **Pending Global—**Control of this policy is being transferred to Cisco UCS Central. Once the transfer is complete, this policy is available to all Cisco UCS domains registered with Cisco UCS Central. • **Global—**This policy is managed by Cisco UCS Central. Any changes to this policy must be made through Cisco UCS Central.

Name	Description
Sync-Option field	The options that allow you to synchronize the desired power state of the associated service profile to the physical server. This can be one of the following: • **Default Sync—**After the initial server association, any configuration change or management connectivity changes that you perform trigger a server reassociation. This option synchronizes the desired power state to the physical server if the physical server power state is off and the desired power state is on. This is the default behavior. • **Always Sync—**When the initial server association or the server reassociation occurs, this option always synchronize the desired power state to the physical power state even if the physical server power state is on and desired power state is off. • **Initial Only Sync—**This option only synchronizes the power to a server when a service profile is associated to the server for the first time or when the server is re-commissioned. When you set this option, resetting the power state from the physical server side does not affect the desired power state on the service profile.

Step 6 Click **OK**.

What to Do Next

Include the policy in a service profile or service profile template.

Changing a Power Sync Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Policies**.

Step 3 Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.

Step 4 Choose a service profile policy from the **root** node.

Step 5 In the **Work** pane, click the **Policies** tab.

Step 6 Click the **Change Power Sync Policy** from the **Actions** area.
The information displayed depends on what you choose in the **Select the Power Sync Policy** drop-down list. You can choose:

- **No Power Sync Policy—**If you choose this option, Cisco UCS Manager GUI does not display any other information. When you choose this option, Cisco UCS Manager implicitly uses the default power sync policy. Cisco UCS Manager searches for the default power sync policy under service profile organizations. If the policy is not found, then it uses the default power sync policy under root.
- **Use an Existing Power Sync Policy—**if you want to select a global policy. Cisco UCS Manager GUI displays the **Power Sync Policy** drop-down list that enables you to choose an existing policy.
- **Create a Local Power Sync Policy—**if you want to create a power sync policy that can only be accessed by this service profile. You can also create a power sync policy by using the **Create Power Sync Policy** link from the Power Sync Policy area.

Deleting a Power Sync Policy

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Policies** > *Organization_Name*.
Step 3 Expand the **Power Sync Policies** node.
Step 4 Right-click the policy you want to delete and select **Delete**.
Step 5 If a confirmation dialog box displays, click **Yes**.

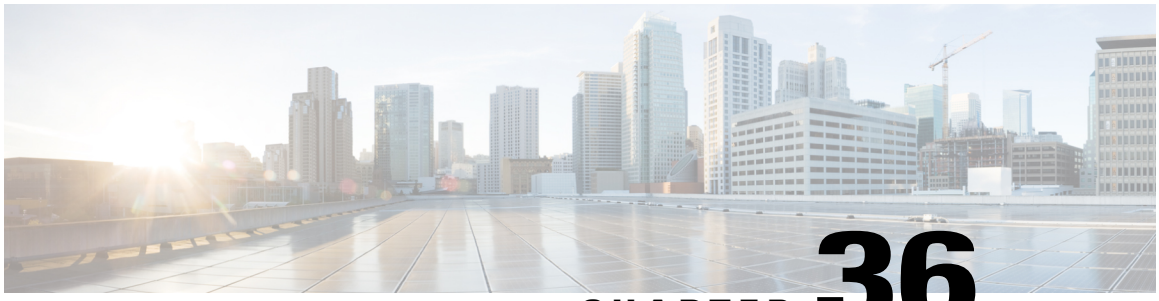

CHAPTER 36

Managing Time Zones

This chapter includes the following sections:

Time Zones

Cisco UCS requires a domain-specific time zone setting and an NTP server to ensure the correct time display in Cisco UCS Manager. If you do not configure both of these settings in a Cisco UCS domain, the time does not display correctly.

Setting the Time Zone

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 In the **Admin** tab, expand **All**.

Step 3 Click **Time Zone Management**.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 From the **Time Zone** drop-down list, select the time zone you want to use for the Cisco UCS domain.

Step 6 Click **Save Changes**.

Adding an NTP Server

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 In the **Admin** tab, expand **All**.

Step 3 Click **Time Zone Management**.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **NTP Servers** area, click the + button on the table icon bar.

Step 6 In the **Add NTP Server** dialog box, do the following:

a) In the **NTP Server** field, enter the IPv4 or IPv6 address or hostname of the NTP server you want to use for this Cisco UCS domain.

b) Click **OK**.

Deleting an NTP Server

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 In the **Admin** tab, expand **All**.

Step 3 Click **Time Zone Management**.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **NTP Servers** area, right-click the server you want to delete and select **Delete**.

Step 6 If a confirmation dialog box displays, click **Yes**.

Step 7 Click **Save Changes**.

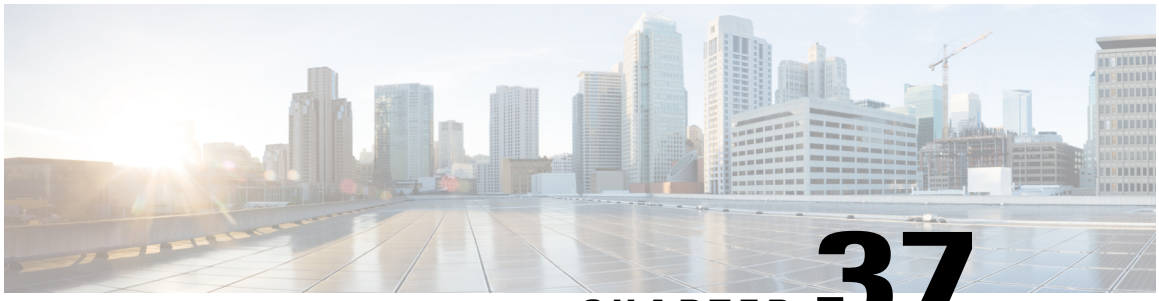

CHAPTER 37

Managing the Chassis

This chapter includes the following sections:

Chassis Management in Cisco UCS Manager GUI

You can manage and monitor all chassis in a Cisco UCS domain through Cisco UCS Manager GUI.

Guidelines for Removing and Decommissioning Chassis

Consider the following guidelines when deciding whether to remove or decommission a chassis using Cisco UCS Manager:

Decommissioning a Chassis

Decommissioning is performed when a chassis is physically present and connected but you want to temporarily remove it from the Cisco UCS Manager configuration. Because it is expected that a decommissioned chassis

will be eventually recommissioned, a portion of the chassis' information is retained by Cisco UCS Manager for future use.

Removing a Chassis

Removing is performed when you physically remove a chassis from the system. Once the physical removal of the chassis is completed, the configuration for that chassis can be removed in Cisco UCS Manager.

Note You cannot remove a chassis from Cisco UCS Manager if it is physically present and connected.

If you need to add a removed chassis back to the configuration, it must be reconnected and then rediscovered. During rediscovery Cisco UCS Manager will assign the chassis a new ID that may be different from ID that it held before.

Acknowledging a Chassis

Perform the following procedure if you increase or decrease the number of links that connect the chassis to the fabric interconnect. Acknowledging the chassis ensures that Cisco UCS Manager is aware of the change in the number of links and that traffics flows along all available links.

After you enable or disable a port on a fabric interconnect, wait for at least 1 minute before you re-acknowledge the chassis. If you re-acknowledge the chassis too soon, the pinning of server traffic from the chassis might not get updated with the changes to the port that you enabled or disabled.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis**.

Step 3 Choose the chassis that you want to acknowledge.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Acknowledge Chassis**.

Step 6 If Cisco UCS Manager displays a confirmation dialog box, click **Yes**.
Cisco UCS Manager disconnects the chassis and then rebuilds the connections between the chassis and the fabric interconnect or fabric interconnects in the system.

Decommissioning a Chassis

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis**.

Step 3 Choose the chassis that you want to decommission.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Decommission Chassis**.

Step 6 If Cisco UCS Manager GUI displays a confirmation dialog box, click **Yes**.
The decommission may take several minutes to complete. After the chassis has been removed from the configuration, Cisco UCS Manager adds the chassis to the **Decommissioned** tab.

Removing a Chassis

Before You Begin

Physically remove the chassis before performing the following procedure.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis**.

Step 3 Choose the chassis that you want to remove.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Remove Chassis**.

Step 6 If Cisco UCS Manager displays a confirmation dialog box, click **Yes**.
The removal may take several minutes to complete.

Recommissioning a Single Chassis

This procedure returns the chassis to the configuration and applies the chassis discovery policy to the chassis. After this procedure, you can access the chassis and any servers in it.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 In the **Equipment** tab, expand the **Equipment** node.

Step 3 Click the **Chassis** node.

Step 4 In the **Work** pane, click the **Decommissioned** tab.

Step 5 For the chassis that you want to recommission, do the following:

a) Right-click the chassis and choose **Re-commission Chassis**.

b) In the **Chassis ID** field of the **Re-commission Chassis** dialog box, type or use the arrows to choose the ID that you want to assign to the chassis

c) Click **OK**.

Step 6 If Cisco UCS Manager GUI displays a confirmation dialog box, click **Yes**.
This procedure may take several minutes to complete. After the chassis has been recommissioned, Cisco UCS Manager runs the chassis discovery policy and adds the chassis to the list in the **Navigation** pane.

Recommissioning Multiple Chassis

This procedure returns the chassis to the configuration and applies the chassis discovery policy to the chassis. After this procedure, you can access the chassis and any servers in it.

Note You cannot renumber the chassis when you recommission multiple chassis at the same time. Cisco UCS Manager assigns the same ID that the chassis had previously.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 In the **Equipment** tab, expand the **Equipment** node.

Step 3 Click the **Chassis** node.

Step 4 In the **Work** pane, click the **Decommissioned** tab.

Step 5 In the row for each chassis that you want to recommission, check the **Re-commission** check box.

Step 6 Click **Save Changes**.

Step 7 If Cisco UCS Manager GUI displays a confirmation dialog box, click **Yes**.
This procedure may take several minutes to complete. After the chassis has been recommissioned, Cisco UCS Manager runs the chassis discovery policy and adds the chassis to the list in the **Navigation** pane.

Renumbering a Chassis

Note You cannot renumber a blade server through Cisco UCS Manager. The ID assigned to a blade server is determined by its physical slot in the chassis. To renumber a blade server, you must physically move the server to a different slot in the chassis.

Before You Begin

If you are swapping IDs between chassis, you must first decommission both chassis, then wait for the chassis decommission FSM to complete before proceeding with the renumbering steps.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis**.

Step 3 Verify that the **Chassis** node does not include the following:

- The chassis you want to renumber
- A chassis with the number you want to use

If either of these chassis are listed in the **Chassis** node, decommission those chassis. You must wait until the decommission FSM is complete and the chassis are not listed in the **Chassis** node before continuing. This might take several minutes.

Step 4 On the **Equipment** tab, click the **Chassis** node.

Step 5 In the **Work** pane, click the **Decommissioned** tab.

Step 6 For the chassis that you want to renumber, do the following:

a) Right-click the chassis and choose **Re-commission Chassis**.

b) In the **Chassis ID** field of the **Re-commission Chassis** dialog box, type or use the arrows to choose the ID that you want to assign to the chassis

c) Click **OK**

Step 7 If a confirmation dialog box displays, click **Yes**.

Toggling the Locator LED

Local Disk Locator LED Status

The local disk locator LED, located on the slot where you insert the local disk, identifies where a specific disk is inserted in a blade or rack server. The locator LED is useful when you need to remove a disk from among many in a server for maintenance.

You can successfully turn on or off the local disk locator LED when:

- The server is powered on. When the server is powered off and you try to turn on or off, the locator LED, UCS Manager generates an error.
- CIMC version is 2.7 or higher.
- The RAID controller supports the out-of-band (OOB) storage interface.

Turning on the Locator LED for a Chassis

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis**.

Step 3 Click the chassis that you need to locate.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Turn on Locator LED**.
This action is not available if the locator LED is already turned on.

The LED on the chassis starts flashing.

Turning off the Locator LED for a Chassis

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis**.

Step 3 Choose the chassis for which you want to turn off the locator LED.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Turn off Locator LED**.
This action is not available if the locator LED is already turned off.

The LED on the chassis stops flashing.

Toggling the Local Disk Locator LED On and Off

Before You Begin

On and Off

- Ensure the server on which the disk is located is powered on. If the server is off, you are not able to turn on or off the local disk locator LED.

Procedure

Step 1 In the **Navigation** pane, click the **Equipment** tab.

Step 2 On the **Equipment** tab, expand **Equipment** > **Rack Mounts** > **Servers** > ***Server Number***.

Step 3 In the **Work** area, click the **Inventory > Storage > Disks** tabs.
The Storage Controller inventory appears.

Step 4 Click a disk.
The disk details appear.

Step 5 In the Actions area, click **Turn on Locator LED** or **Turn off Locator LED**.
The **Locator LED** state appears in the **Properties** area.

Step 6 Click **Save Changes.**

NVMe PCIe SSD Inventory

Cisco UCS Manager GUIdiscovers, identifies, and displays the inventory of Non-Volatile Memory Express (NVMe) Peripheral Component Interconnect Express (PCIe) SSD storage devices. You can view the health of the storage devices in the server. NVMe with PCIe SSD storage devices reduce latency, increased input/output operations per second (IOPS), and lower power consumption compared to SAS or SATA SSDs.

Viewing NVMe PCIe SSD Storage Inventory

Procedure

Step 1 In the **Navigation** pane, click the **Equipment** tab.

Step 2 On the **Equipment** tab, expand **Equipment** > **Rack Mounts** > **Servers**.

Step 3 Click the **Inventory** tab.

Step 4 Do one of the following:

a) Click the **Storage** tab.
You view the list of NVMe PCIe SSD storage devices named **Storage Controller NVME *ID number*.**
You view the name, size, serial number, operating status, state and other details.

b) Click the NVMe PCIe SSD storage device.
You see the following inventory details:

Name	Description
ID	The NVMe PCIe SSD storage device configured on the server.

Name	Description
Model	The NVMe PCIe SSD storage device model.
Revision	The NVMe PCIe SSD storage device revision.
RAID Support	Whether the NVMe PCIe SSD storage device is RAID enabled.
OOB Interface Support	Whether the NVMe PCIe SSD storage device support out-of-band management .
PCIe Address	The NVMe PCIe SSD storage device on the virtual interface card (VIC).
Number of Local Disks	The number of disks contained in the NVMe PCIe SSD storage device.
Rebuild Rate	Not applicable to NVMe PCIe SSD storage devices.
Vendor	The vendor that manufactured the NVMe PCIe SSD storage device.
PID	The NVMe PCIe SSD storage device product ID, also known as product name, model name, product number
Serial	The storage device serial number.

Health LED Alarms

The blade health LED is located on the front of each Cisco UCS B-Series blade server. Cisco UCS Manager allows you to view the sensor faults that cause the blade health LED to change color from green to amber or blinking amber.

The health LED alarms display the following information:

Name	Description
Severity column	The severity of the alarm. This can be one of the following: • Critical—The blade health LED is blinking amber. This is indicated with a red dot. • Minor—The blade health LED is amber. This is indicated with an orange dot.
Description column	A brief description of the alarm.
Sensor ID column	The ID of the sensor the triggered the alarm.
Sensor Name column	The name of the sensor that triggered the alarm.

Viewing Health LED Alarms

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > ***Chassis Number*** > **Servers**.

Step 3 Click the server for which you want to view health LED alarms.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **View Health LED Alarms**.
The **View Health LED Alarms** dialog box lists the health LED alarms for the selected server.

Step 6 Click **OK** to close the **View Health LED Alarms** dialog box.

Viewing Health LED Status

Procedure

	Command or Action	Purpose
Step 1	UCS-A# **scope server** *chassis-id* / *blade-id*	Enters chassis server mode for the specified server.
Step 2	UCS-A /chassis/server # **show health-led expand**	Displays the health LED and sensor alarms for the selected server.

The following example shows how to display the health LED status and sensor alarms for chassis 1 server 1:

```
UCS-A# scope server 1/1
UCS-A /chassis/server # show health-led
Health LED:
    Severity: Minor
    Reason:: P0V75_STBY:Voltage Threshold Crossed;TEMP_SENS_FRONT:Temperature Threshold
Crossed;
    Color: Amber
    Oper State:: On

    Sensor Alarm:
        Severity: Minor
        Sensor ID: 7
        Sensor Name: P0V75_STBY
        Alarm Desc: Voltage Threshold Crossed

        Severity: Minor
        Sensor ID: 76
        Sensor Name: TEMP_SENS_FRONT
        Alarm Desc: Temperature Threshold Crossed

        Severity: Minor
        Sensor ID: 91
        Sensor Name: DDR3_P1_D2_TMP
        Alarm Desc: Temperature Threshold Crossed

UCS-A /chassis/server #
```

Viewing the POST Results for a Chassis

You can view any errors collected during the Power On Self-Test process for all servers and adapters in a chassis.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis**.

Step 3 Choose the chassis for which you want to view the POST results.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **View POST Results**.
The **POST Results** dialog box lists the POST results for each server in the chassis and its adapters.

Step 6 (Optional) Click the link in the **Affected Object** column to view the properties of that adapter.

Step 7 Click **OK** to close the **POST Results** dialog box.

CHAPTER 38

Managing Blade Servers

This chapter includes the following sections:

- Health LED Alarms, page 701
- Viewing Health LED Alarms, page 702

Blade Server Management

You can manage and monitor all blade servers in a Cisco UCS domain through Cisco UCS Manager. You can perform some blade server management tasks, such as changes to the power state, from the server and service profile.

The remaining management tasks can only be performed on the server.

The power supply units go into power save mode when a chassis has two blades or less. When a third blade is added to the chassis and is fully discovered, the power supply units return to regular mode.

If a blade server slot in a chassis is empty, Cisco UCS Manager provides information, errors, and faults for that slot. You can also re-acknowledge the slot to resolve server mismatch errors and to have Cisco UCS Manager rediscover the blade server in the slot.

Cisco UCS B460 M4 Blade Server Management

The Cisco UCS B460 M4 blade server consists of two full-width Cisco UCS B260 blade servers that are connected by a Cisco UCS scalability connector. Each individual blade server is called a node and can be either the master or slave node.

Because each Cisco UCS B460 M4 blade server has two different nodes, you should note the following:

- The master node is always the node in the highest numbered slots.
- Whenever the Cisco UCS B460 blade server is referred to in Cisco UCS Manager, the reference is to the master slot number.
- If you remove the Cisco UCS scalability connector from the Cisco UCS B460 M4 blade server, the **Physical Display** area in the Cisco UCS Manager GUI displays **Needs Resolution** on both master node slots and both slave node slots.
- The health LED displays both the individual health of the master and slave node, and the combined health of both nodes together. The combined health LED always displays the status of the node with the worst health. Any health LED alarms are shown individually.
- In the Cisco UCS Manager GUI, you can turn on and off the locator LEDs for either the master or the slave node. In the Cisco UCS Manager CLI, you can turn on and off the locator LEDs individually, or both locator LEDs at the same time.
- Power capping on the Cisco UCS B460 M4 blade server is applied at the server level. Each node is capped at one half of the total value.
- Updating firmware updates both the master and slave node at the same time. You cannot update the firmware on an individual node.
- Local disk configuration is supported only on the master node.
- The Cisco UCS B460 blade server does not distinguish between the SEL logs that are generated by either the master or the slave node. The logs are displayed on the same page and are differentiated by the slot number.

- On the Cisco UCS Manager GUI **Storage** tab, the **Local Disk Configuration Policy** and **Actual Disk Configurations** areas display only the data for the Cisco UCS B460 blade server master node. No fields are displayed for the slave node.

Upgrading to a Cisco UCS B460 M4 Blade Server

If you have a Cisco UCS B260 M4 blade server, you can purchase an upgrade kit to convert to a Cisco UCS B460 M4 blade server. For more information, see the appropriate *Cisco UCS Hardware Installation Guide*.

Before You Begin

You must have two Cisco UCS B260 M4 blade servers and a Cisco UCS scalability connector.

Procedure

Step 1 Verify that the existing Cisco UCS B260 M4 blade server is not associated with a service profile.

Step 2 Insert the second Cisco UCS B260 M4 blade server into the chassis either above or below the first blade server.

Note If the second blade server does not have a Cisco UCS scalability terminator, use the terminator from the first blade server.

Step 3 Decommission both Cisco UCS B260 M4 blade servers.

Step 4 Synchronize the firmware.
Use the **Firmware Auto Sync Server** policy in Cisco UCS Manager to automatically update the new server. For more information, see the appropriate *Cisco UCS B-Series Firmware Management Guide*.

Step 5 Replace the Cisco UCS scalability terminators with the Cisco UCS scalability connector.
The presence of the slots changes to mismatch, but discovery is not triggered.

Step 6 Reacknowledge the new Cisco UCS B460 M4 blade server.

Guidelines for Removing and Decommissioning Blade Servers

Consider the following guidelines when deciding whether to remove or decommission a blade server using Cisco UCS Manager:

Decommissioning a Blade Server

If you want to temporarily decommission a physically present and connected blade server, you can temporarily remove it from the configuration. A portion of the server's information is retained by Cisco UCS Manager for future use, in case the blade server is recommissioned.

Removing a Blade Server

Removing is performed when you physically remove a blade server from the Cisco UCS Manager by disconnecting it from the chassis. You cannot remove a blade server from Cisco UCS Manager if it is physically present and connected to a chassis. After the physical removal of the blade server is completed, the configuration for that blade server can be removed in Cisco UCS Manager.

During removal, active links to the blade server are disabled, all entries from databases are removed, and the server is automatically removed from any server pools that it was assigned to during discovery.

Note Only servers added to a server pool automatically during discovery are removed automatically. Servers that were manually added to a server pool must be removed manually.

To add a removed blade server back to the configuration, it must be reconnected, then rediscovered. When a server is reintroduced to Cisco UCS Manager, it is treated as a new server and is subject to the deep discovery process. For this reason, it is possible for Cisco UCS Manager to assign the server a new ID that might be different from the ID that it held before.

Recommendations for Avoiding Unexpected Server Power Changes

If a server is not associated with a service profile, you can use any available means to change the server power state, including the physical Power or Reset buttons on the server.

If a server is associated with, or assigned to, a service profile, you should only use the following methods to change the server power state:

- In Cisco UCS Manager GUI, go to the **General** tab for the server or the service profile associated with the server and select **Boot Server** or **Shutdown Server** from the **Actions** area.
- In Cisco UCS Manager CLI, scope to the server or the service profile associated with the server and use the **power up** or **power down** commands.

Important Do *not* use any of the following options on an associated server that is currently powered off:

- **Reset** in the GUI
- **cycle cycle-immediate** or **reset hard-reset-immediate** in the CLI
- The physical Power or Reset buttons on the server

If you reset, cycle, or use the physical power buttons on a server that is currently powered off, the server's actual power state might become out of sync with the desired power state setting in the service profile. If the communication between the server and Cisco UCS Manager is disrupted or if the service profile configuration changes, Cisco UCS Manager might apply the desired power state from the service profile to the server, causing an unexpected power change.

Power synchronization issues can lead to an unexpected server restart, as shown below:

Desired Power State in Service Profile	Current Server Power State	Server Power State After Communication Is Disrupted
Up	Powered Off	Powered On

Desired Power State in Service Profile	Current Server Power State	Server Power State After Communication Is Disrupted
Down	Powered On	Powered On **Note** Running servers are not shut down regardless of the desired power state in the service profile.

Booting Blade Servers

Booting a Blade Server

If the **Boot Server** link is dimmed in the **Actions** area, you must shut down the server first.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
Step 3 Choose the server that you want to boot.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Boot Server**.
Step 6 If a confirmation dialog box displays, click **Yes**.

After the server boots, the **Overall Status** field on the **General** tab displays an OK status.

Booting a Server from the Service Profile

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Service Profiles**.
Step 3 Expand the node for the organization where you want to create the service profile.
If the system does not include multitenancy, expand the **root** node.
Step 4 Choose the service profile that requires the associated server to boot.
Step 5 In the **Work** pane, click the **General** tab.
Step 6 In the **Actions** area, click **Boot Server**.
Step 7 If a confirmation dialog box displays, click **Yes**.
Step 8 Click **OK** in the **Boot Server** dialog box.

After the server boots, the **Overall Status** field on the **General** tab displays an ok status or an up status.

Determining the Boot Order of a Blade Server

Tip You can also view the boot order tabs from the **General** tab of the service profile associated with a server.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.

Step 3 Click the server for which you want to determine the boot order.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 If the **Boot Order Details** area is not expanded, click the **Expand** icon to the right of the heading.

Step 6 To view the boot order assigned to the server, click the **Configured Boot Order** tab.

Step 7 To view what will boot from the various devices in the physical server configuration, click the **Actual Boot Order** tab.

Note The **Actual Boot Order** tab always shows "Internal EFI Shell" at the bottom of the boot order list.

Shutting Down Blade Servers

Shutting Down a Blade Server

When you use this procedure to shut down a server with an installed operating system, Cisco UCS Manager triggers the OS into a graceful shutdown sequence.

If the **Shutdown Server** link is dimmed in the **Actions** area, the server is not running.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
Step 3 Choose the server that you want to shut down.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Shutdown Server**.
Step 6 If a confirmation dialog box displays, click **Yes**.

After the server has been successfully shut down, the **Overall Status** field on the **General** tab displays a power-off status.

Shutting Down a Server from the Service Profile

When you use this procedure to shut down a server with an installed operating system, Cisco UCS Manager triggers the OS into a graceful shutdown sequence.

If the **Shutdown Server** link is dimmed in the **Actions** area, the server is not running.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Service Profiles**.
Step 3 Expand the node for the organization where you want to create the service profile.
If the system does not include multitenancy, expand the **root** node.
Step 4 Choose the service profile that requires the associated server to shut down.
Step 5 In the **Work** pane, click the **General** tab.
Step 6 In the **Actions** area, click **Shutdown Server**.
Step 7 If a confirmation dialog box displays, click **Yes**.

After the server successfully shuts down, the **Overall Status** field on the **General** tab displays a down status or a power-off status.

Resetting a Blade Server

When you reset a server, Cisco UCS Manager sends a pulse on the reset line. You can choose to gracefully shut down the operating system. If the operating system does not support a graceful shutdown, the server is power cycled. The option to have Cisco UCS Manager complete all management operations before it resets the server does not guarantee the completion of these operations before the server is reset.

Note If you are trying to boot a server from a power-down state, you should not use **Reset**.

If you continue the power-up with this process, the desired power state of the servers become out of sync with the actual power state and the servers might unexpectedly shut down at a later time. To safely reboot the selected servers from a power-down state, click **Cancel**, then select the **Boot Server** action.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.

Step 3 Choose the server that you want to reset.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Reset**.

Step 6 In the **Reset Server** dialog box, do the following:

a) Click the **Power Cycle** option.

b) (Optional) Check the check box if you want Cisco UCS Manager to complete all management operations that are pending on this server.

c) Click **OK**.

The reset may take several minutes to complete. After the server has been reset, the **Overall Status** field on the **General** tab displays an ok status.

Resetting a Blade Server to Factory Default Settings

You can now reset a blade server to its factory settings. By default, the factory reset operation does not affect storage drives and flexflash drives. This is to prevent any loss of data. However, you can choose to reset these devices to a known state as well.

Important Resetting storage devices will result in loss of data.

Perform the following procedure to reset the server to factory default settings.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.

Step 3 Choose the server that you want to reset to its factory default settings.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Server Maintenance**.

Step 6 In the **Maintenance** dialog box, do the following:

a) Click **Reset to Factory Default**.
b) Click **OK**.

Step 7 From the **Maintenance Server** dialog box that appears, select the appropriate options:

- To delete all storage, check the **Scrub Storage** checkbox.
- To place all disks into their initial state after deleting all storage, check the **Create Initial Volumes** checkbox.

  You can check this checkbox only if you check the **Scrub Storage** checkbox. For servers that support JBOD, the disks will be placed in a JBOD state. For servers that do not support JBOD, each disk will be initialized with a single R0 volume that occupies all the space in the disk.

  Important Do not check the **Create Initial Volumes** if you want to use storage profiles. Creating initial volumes when you are using storage profiles may result in configuration errors.

- To delete all flexflash storage, check the **Scrub FlexFlash** checkbox.

Cisco UCS Manager resets the server to its factory default settings.

Reacknowledging a Blade Server

Perform the following procedure to rediscover the server and all endpoints in the server. For example, you can use this procedure if a server is stuck in an unexpected state, such as the discovery state.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
Step 3 Choose the server that you want to acknowledge.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Server Maintenance**.
Step 6 In the **Maintenance** dialog box, click **Re-acknowledge**, then click **OK**.
Cisco UCS Manager disconnects the server and then builds the connections between the server and the fabric interconnect or fabric interconnects in the system. The acknowledgment may take several minutes to complete. After the server has been acknowledged, the **Overall Status** field on the **General** tab displays an OK status.

Removing a Server from a Chassis

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > ***Chassis Number*** > **Servers**.

Step 3 Choose the server that you want to remove from the chassis.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Server Maintenance**.

Step 6 In the **Maintenance** dialog box, click **Decommission**, then click **OK**.
The server is removed from the Cisco UCS configuration.

Step 7 Go to the physical location of the chassis and remove the server hardware from the slot.
For instructions on how to remove the server hardware, see the *Cisco UCS Hardware Installation Guide* for your chassis.

What to Do Next

If you physically re-install the blade server, you must re-acknowledge the slot for the Cisco UCS Manager to rediscover the server.

For more information, see Reacknowledging a Server Slot in a Chassis, on page 696.

Deleting the Inband Configuration from a Blade Server

This procedure removes the inband management IP address configuration from a blade server. If this action is greyed out, no inband configuration was completed.

Procedure

Step 1 In the **Navigation** pane, click the **Servers** tab.

Step 2 On the **Equipment** tab, expand **Equipment** > **Chassis** > ***Chassis Number*** > **Servers** > ***Server Name***.

Step 3 In the **Work** area, click the **Inventory** tab.

Step 4 Click the **CIMC** subtab.

Step 5 In the **Actions** area, click **Delete Inband Configuration**.

Step 6 Click **Yes** in the **Delete** confirmation dialog box.
The inband configuration for the server is deleted.

Note If an inband service profile is configured in Cisco UCS Manager with a default VLAN and pool name, the server CIMC will automatically get an inband configuration from the inband profile approximate one minute after deleting the inband configuration here.

Decommissioning a Blade Server

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > ***Chassis Number*** > **Servers**.

Step 3 Choose the server that you want to decommission.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Server Maintenance**.

Step 6 In the **Maintenance** dialog box, do the following:

a) Click **Decommission**.

b) Click **OK**.

The server is removed from the Cisco UCS configuration.

What to Do Next

If you physically re-install the blade server, you must re-acknowledge the slot for the Cisco UCS Manager to rediscover the server.

For more information, see Reacknowledging a Server Slot in a Chassis, on page 696.

Removing a Non-Existent Blade Server Entry

Perform the following procedure after decommissioning the server and physically removing the server hardware. This procedure removes the non existing stale entry of a blade server from the **Decommissioned** tab.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 In the **Work** pane, click the **Decommissioned** tab.

Step 3 On the row for each blade server that you want to remove from the list, check the check box in the **Recommission** column, then click **Save Changes**.

Step 4 If a confirmation dialog box displays, click **Yes**.

Recommissioning a Blade Server

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 On the **Equipment** tab, click the **Chassis** node.

Step 3 In the **Work** pane, click the **Decommissioned** tab.

Step 4 On the row for each blade server that you want to recommission, check the check box in the **Recommission** column, then click **Save Changes**.

Step 5 If a confirmation dialog box displays, click **Yes**.

Step 6 (Optional) Monitor the progress of the server recommission and discovery on the **FSM** tab for the server.

What to Do Next

-

Reacknowledging a Server Slot in a Chassis

Perform the following procedure if you decommissioned a blade server without removing the physical hardware from the chassis, and you want Cisco UCS Manager to rediscover and recommission the server.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.

Step 3 Choose the server whose slot you want to reacknowledge.

Step 4 If Cisco UCS Manager displays a **Resolve Slot Issue** dialog box, do one of the following:

Option	Description
The **here** link in the **Situation** area	Click this link and then click **Yes** in the confirmation dialog box. Cisco UCS Manager reacknowledges the slot and discovers the server in the slot.
OK	Click this button if you want to proceed to the **General** tab. You can use the **Reacknowledge Slot** link in the **Actions** area to have Cisco UCS Manager reacknowledge the slot and discover the server in the slot.

Removing a Non-Existent Blade Server from the Configuration Database

Perform the following procedure if you physically removed the server hardware without first decommissioning the server. You cannot perform this procedure if the server is physically present.

If you want to physically remove a server, see Removing a Server from a Chassis, on page 694.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > ***Chassis Number*** > **Servers**.

Step 3 Choose the server that you want to remove from the configuration database.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Server Maintenance**.

Step 6 In the **Maintenance** dialog box, click **Remove**, then click **OK**.
Cisco UCS Manager removes all data about the server from its configuration database. The server slot is now available for you to insert new server hardware.

Turning the Locator LED for a Blade Server On and Off

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > ***Chassis Number*** > **Servers**.

Step 3 Choose the server for which you want to turn the locator LED on or off.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click one of the following:

- **Turn on Locator LED—**Turns on the LED for the selected server.
- **Turn off Locator LED—**Turns off the LED for the selected server.
- **Turn on Master Locator LED—**For the Cisco UCS B460 M4 blade server, turns on the LED for the master node.
- **Turn off Master Locator LED—**For the Cisco UCS B460 M4 blade server, turns off the LED for the master node.
- **Turn on Slave Locator LED—**For the Cisco UCS B460 M4 blade server, turns on the LED for the slave node.

- **Turn off Locator LED—**For the Cisco UCS B460 M4 blade server, turns off the LED for the slave node.

Resetting the CMOS for a Blade Server

Sometimes, troubleshooting a server might require you to reset the CMOS. Resetting the CMOS is not part of the normal maintenance of a server.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Chassis** > ***Chassis Number*** > **Servers**.
Step 3 Choose the server for which you want to reset the CMOS.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Recover Server**.
Step 6 In the **Recover Server** dialog box, click **Reset CMOS**, then click **OK**.

Resetting the CIMC for a Blade Server

Sometimes, with the firmware, troubleshooting a server might require you to reset the CIMC. Resetting the CIMC is not part of the normal maintenance of a server. After you reset the CIMC, the CIMC reboots with the running version of the firmware for that server.

If the CIMC is reset, the power monitoring functions of Cisco UCS become briefly unavailable until the CIMC reboots. Typically, the reset only takes 20 seconds; however, it is possible that the peak power cap can exceed during that time. To avoid exceeding the configured power cap in a low power-capped environment, consider staggering the rebooting or activation of CIMCs.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Chassis** > ***Chassis Number*** > **Servers**.
Step 3 Choose the server for which you want to reset the CIMC.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Recover Server**.
Step 6 In the **Recover Server** dialog box, click **Reset CIMC (Server Controller)**, then click **OK**.

Clearing TPM for a Blade Server

You can clear TPM only on Cisco UCS M4 blade and rack-mount servers that include support for TPM.

Caution Clearing TPM is a potentially hazardous operation. The OS may stop booting. You may also see loss of data.

Before You Begin

TPM must be enabled.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
Step 3 Choose the server for which you want to clear TPM.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Recover Server**.
Step 6 In the **Recover Server** dialog box, click **Clear TPM**, then click **OK**.

Recovering the Corrupt BIOS on a Blade Server

Sometimes, an issue with a server might require you to recover the corrupted BIOS. This procedure is not part of the normal maintenance of a server. After you recover the BIOS, the server boots with the running version of the firmware for that server. This radio button might dim if the BIOS does not require recovery or the option is not available for a particular server.

Before You Begin

Important Remove all attached or mapped USB storage from a server before you attempt to recover the corrupt BIOS on that server. If an external USB drive is attached or mapped from vMedia to the server, BIOS recovery fails.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.

Step 3 Choose the server for which you want to recover the BIOS.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Recover Server**.

Step 6 In the **Recover Server** dialog box, do the following:

a) Click **Recover Corrupt BIOS**.

Note If this option is not available for a specific server, follow the instructions to update and activate the BIOS for a server.

b) Click **OK**.

Step 7 If a confirmation dialog box displays, click **Yes**.

Step 8 In the **Recover Corrupt BIOS** dialog box, do the following:

a) Complete the following fields:

Name	Description
Version To Be Activated drop-down list	Choose the firmware version from the drop-down list to activate.

b) Click **OK**.

Viewing the POST Results for a Blade Server

You can view any errors collected during the Power On Self-Test process for a server and its adapters.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.

Step 3 Choose the server for which you want to view the POST results.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **View POST Results**.
The **POST Results** dialog box lists the POST results for the server and its adapters.

Step 6 (Optional) Click the link in the **Affected Object** column to view the properties of that adapter.

Step 7 Click **OK** to close the **POST Results** dialog box.

Issuing an NMI from a Blade Server

Perform the following procedure if the system remains unresponsive and you need Cisco UCS Manager to issue a Non-Maskable Interrupt (NMI) to the BIOS or operating system from the CIMC. This action creates a core dump or stack trace, depending on the operating system installed on the server.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.

Step 3 Choose the server that you want to issue the NMI.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Server Maintenance**.

Step 6 In the **Maintenance** dialog box, do the following:

a) Click **Diagnostic Interrupt**.

b) Click **OK**.

Cisco UCS Manager sends an NMI to the BIOS or operating system.

Health LED Alarms

The blade health LED is located on the front of each Cisco UCS B-Series blade server. Cisco UCS Manager allows you to view the sensor faults that cause the blade health LED to change color from green to amber or blinking amber.

The health LED alarms display the following information:

Name	Description
Severity column	The severity of the alarm. This can be one of the following: • Critical—The blade health LED is blinking amber. This is indicated with a red dot. • Minor—The blade health LED is amber. This is indicated with an orange dot.
Description column	A brief description of the alarm.
Sensor ID column	The ID of the sensor the triggered the alarm.
Sensor Name column	The name of the sensor that triggered the alarm.

Viewing Health LED Alarms

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.

Step 3 Click the server for which you want to view health LED alarms.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **View Health LED Alarms**.
The **View Health LED Alarms** dialog box lists the health LED alarms for the selected server.

Step 6 Click **OK** to close the **View Health LED Alarms** dialog box.

CHAPTER 39

Managing Rack-Mount Servers

This chapter includes the following sections:

Rack-Mount Server Management

You can manage and monitor all rack-mount servers that are integrated with a Cisco UCS domain through Cisco UCS Manager. All management and monitoring features are supported for rack-mount servers except power capping. Some rack-mount server management tasks, such as changes to the power state, can be performed from both the server and service profile. The remaining management tasks can only be performed on the server.

Cisco UCS Manager provides information, errors, and faults for each rack-mount server that it has discovered.

Tip For information on how to integrate a supported Cisco UCS rack-mount server with Cisco UCS Manager, see the Cisco UCS C-series server integration guide or Cisco UCS S-series server integration guide for your Cisco UCS Manager release.

Guidelines for Removing and Decommissioning Rack-Mount Servers

Consider the following guidelines when deciding whether to remove or decommission a rack-mount server using Cisco UCS Manager:

Decommissioning a Rack-Mount server

Decommissioning is performed when a rack-mount server is physically present and connected but you want to temporarily remove it from the configuration. Because it is expected that a decommissioned rack-mount server will be eventually recommissioned, a portion of the server's information is retained by Cisco UCS Manager for future use.

Removing a Rack-Mount server

Removing is performed when you physically remove the server from the system by disconnecting the rack-mount server from the fabric extender. You cannot remove a rack-mount server from Cisco UCS Manager if it is physically present and connected to the fabric extender. Once the rack-mount server is disconnected, the configuration for that rack-mount server can be removed in Cisco UCS Manager.

During removal, management interfaces are disconnected, all entries from databases are removed, and the server is automatically removed from any server pools that it was assigned to during discovery.

Note Only those servers added to a server pool automatically during discovery will be removed automatically. Servers that have been manually added to a server pool have to be removed manually.

If you need to add a removed rack-mount server back to the configuration, it must be reconnected and then rediscovered. When a server is reintroduced to Cisco UCS Manager it is treated like a new server and is subject to the deep discovery process. For this reason, it's possible that Cisco UCS Manager will assign the server a new ID that may be different from the ID that it held before.

Recommendations for Avoiding Unexpected Server Power Changes

If a server is not associated with a service profile, you can use any available means to change the server power state, including the physical Power or Reset buttons on the server.

If a server is associated with, or assigned to, a service profile, you should only use the following methods to change the server power state:

- In Cisco UCS Manager GUI, go to the **General** tab for the server or the service profile associated with the server and select **Boot Server** or **Shutdown Server** from the **Actions** area.
- In Cisco UCS Manager CLI, scope to the server or the service profile associated with the server and use the **power up** or **power down** commands.

Important Do *not* use any of the following options on an associated server that is currently powered off:

- **Reset** in the GUI
- **cycle cycle-immediate** or **reset hard-reset-immediate** in the CLI
- The physical Power or Reset buttons on the server

If you reset, cycle, or use the physical power buttons on a server that is currently powered off, the server's actual power state might become out of sync with the desired power state setting in the service profile. If the communication between the server and Cisco UCS Manager is disrupted or if the service profile configuration changes, Cisco UCS Manager might apply the desired power state from the service profile to the server, causing an unexpected power change.

Power synchronization issues can lead to an unexpected server restart, as shown below:

Desired Power State in Service Profile	Current Server Power State	Server Power State After Communication Is Disrupted
Up	Powered Off	Powered On
Down	Powered On	Powered On **Note** Running servers are not shut down regardless of the desired power state in the service profile.

Booting Rack-Mount Servers

Booting a Rack-Mount Server

If the **Boot Server** link is dimmed in the **Actions** area, you must shut down the server first.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Rack Mounts** > **Servers**.
Step 3 Choose the server that you want to boot.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Boot Server**.
Step 6 If a confirmation dialog box displays, click **Yes**.

After the server boots, the **Overall Status** field on the **General** tab displays an OK status.

Booting a Server from the Service Profile

Procedure

Step 1 In the **Navigation** pane, click **Servers**.
Step 2 Expand **Servers** > **Service Profiles**.
Step 3 Expand the node for the organization where you want to create the service profile.
If the system does not include multitenancy, expand the **root** node.
Step 4 Choose the service profile that requires the associated server to boot.
Step 5 In the **Work** pane, click the **General** tab.
Step 6 In the **Actions** area, click **Boot Server**.
Step 7 If a confirmation dialog box displays, click **Yes**.
Step 8 Click **OK** in the **Boot Server** dialog box.
After the server boots, the **Overall Status** field on the **General** tab displays an ok status or an up status.

Determining the Boot Order of a Rack-Mount Server

Tip You can also view the boot order tabs from the **General** tab of the service profile associated with a server.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Rack Mounts** > **Servers**.

Step 3 Click the server for which you want to determine the boot order.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 If the **Boot Order Details** area is not expanded, click the **Expand** icon to the right of the heading.

Step 6 To view the boot order assigned to the server, click the **Configured Boot Order** tab.

Step 7 To view what will boot from the various devices in the physical server configuration, click the **Actual Boot Order** tab.

Note The **Actual Boot Order** tab always shows "Internal EFI Shell" at the bottom of the boot order list.

Shutting Down Rack-Mount Servers

Shutting Down a Rack-Mount Server

When you use this procedure to shut down a server with an installed operating system, Cisco UCS Manager triggers the OS into a graceful shutdown sequence.

If the **Shutdown server** link is dimmed in the **Actions** area, the server is not running.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Rack Mounts** > **Servers**.

Step 3 Choose the server that you want to shut down.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Shutdown Server**.

Step 6 If a confirmation dialog box displays, click **Yes**.

After the server has been successfully shut down, the **Overall Status** field on the **General** tab displays a power-off status.

Shutting Down a Server from the Service Profile

When you use this procedure to shut down a server with an installed operating system, Cisco UCS Manager triggers the OS into a graceful shutdown sequence.

If the **Shutdown Server** link is dimmed in the **Actions** area, the server is not running.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization where you want to create the service profile.
If the system does not include multitenancy, expand the **root** node.

Step 4 Choose the service profile that requires the associated server to shut down.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 In the **Actions** area, click **Shutdown Server**.

Step 7 If a confirmation dialog box displays, click **Yes**.

After the server successfully shuts down, the **Overall Status** field on the **General** tab displays a down status or a power-off status.

Resetting a Rack-Mount Server

When you reset a server, Cisco UCS Manager sends a pulse on the reset line. You can choose to gracefully shut down the operating system. If the operating system does not support a graceful shutdown, the server is power cycled. The option to have Cisco UCS Manager complete all management operations before it resets the server does not guarantee the completion of these operations before the server is reset.

Note If you are trying to boot a server from a power-down state, you should not use **Reset**.

If you continue the power-up with this process, the desired power state of the servers become out of sync with the actual power state and the servers might unexpectedly shut down at a later time. To safely reboot the selected servers from a power-down state, click **Cancel**, then select the **Boot Server** action.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Rack Mounts** > **Servers**.

Step 3 Choose the server that you want to reset.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Reset**.

Step 6 In the **Reset Server** dialog box, do the following:

a) Click the **Power Cycle** option.
b) (Optional) Check the check box if you want Cisco UCS Manager to complete all management operations that are pending on this server.
c) Click **OK**.

The reset may take several minutes to complete. After the server is reset, the **Overall Status** field on the **General** tab displays an ok status.

Reacknowledging a Rack-Mount Server

Perform the following procedure to rediscover the server and all endpoints in the server. For example, you can use this procedure if a server is stuck in an unexpected state, such as the discovery state.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Rack Mounts** > **Servers**.
Step 3 Choose the server that you want to acknowledge.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Server Maintenance**.
Step 6 In the **Maintenance** dialog box, do the following:
a) Click **Re-acknowledge**.
b) Click **OK**.

Cisco UCS Manager disconnects the server, then builds the connections between the server and the fabric interconnect or fabric interconnects in the system. The acknowledgment may take several minutes to complete. After the server is acknowledged, the **Overall Status** field on the **General** tab displays an OK status.

Deleting the Inband Configuration from a Rack Server

This procedure removes the inband management IP address configuration from a rack server. If this action is greyed out, no inband configuration was configured.

Procedure

Step 1 In the **Navigation** pane, click the **Servers** tab.

Step 2 On the **Equipment** tab, expand **Equipment** > **Rack-Mounts** > **Servers** > ***Server Name***.

Step 3 In the **Work** area, click the **Inventory** tab.

Step 4 Click the **CIMC** subtab.

Step 5 In the **Actions** area, click **Delete Inband Configuration**.

Step 6 Click **Yes** in the **Delete** confirmation dialog box.
The inband configuration for the server is deleted.

Note If an inband service profile is configured in Cisco UCS Manager with a default VLAN and pool name, the server CIMC automatically gets an inband configuration from the inband profile approximate one minute after deleting the inband configuration here.

Decommissioning a Rack-Mount Server

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Rack Mounts** > **Servers**.

Step 3 Choose the server that you want to decommission.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Server Maintenance**.

Step 6 In the **Maintenance** dialog box, click **Decommission**, then click **OK**.
The server is removed from the Cisco UCS configuration.

Recommissioning a Rack-Mount Server

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 On the **Equipment** tab, click the **Rack-Mounts** node.

Step 3 In the **Work** pane, click the **Decommissioned** tab.

Step 4 On the row for each rack-mount server that you want to recommission, do the following:

a) In the **Recommission** column, check the check box.

b) Click **Save Changes**

Step 5 If a confirmation dialog box displays, click **Yes**.

Step 6 (Optional) Monitor the progress of the server recommission and discovery on the **FSM** tab for the server.

Renumbering a Rack-Mount Server

Before You Begin

If you are swapping IDs between servers, you must first decommission both servers, then wait for the server decommission FSM to complete before proceeding with the renumbering steps.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Rack Mounts** > **Servers**.

Step 3 Expand the **Servers** node and verify that it does not include the following:

- The rack-mount server you want to renumber
- A rack-mount server with the number you want to use

If either of these servers are listed in the **Servers** node, decommission those servers. You must wait until the decommission FSM is complete and the servers are not listed in the node before continuing. This might take several minutes.

Step 4 Choose the rack-mount server that you want to renumber.

Step 5 On the **Equipment** tab, click the **Rack-Mounts** node.

Step 6 In the **Work** pane, click the **Decommissioned** tab.

Step 7 On the row for each rack-mount server that you want to renumber, do the following:

a) Double-click in the **ID** field, and enter the new number that you want to assign to the rack-mount server.

b) In the **Recommission** column, check the check box.

c) Click **Save Changes**

Step 8 If a confirmation dialog box displays, click **Yes**.

Step 9 (Optional) Monitor the progress of the server recommission and discovery on the **FSM** tab for the server.

Removing a Non-Existent Rack-Mount Server from the Configuration Database

Perform the following procedure if you physically removed the server hardware without first decommissioning the server. You cannot perform this procedure if the server is physically present.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Rack Mounts** > **Servers**.

Step 3 Choose the server that you want to remove from the configuration database.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click **Server Maintenance**.

Step 6 In the **Maintenance** dialog box, click **Remove**, then click **OK**.
Cisco UCS Manager removes all data about the server from its configuration database. The server slot is now available for you to insert new server hardware.

Turning the Locator LED for a Rack-Mount Server On and Off

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Rack Mounts** > **Servers**.

Step 3 Choose the server for which you want to turn the locator LED on or off.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click one of the following:

- **Turn on Locator LED**
- **Turn off Locator LED**

Resetting the CMOS for a Rack-Mount Server

Sometimes, troubleshooting a server might require you to reset the CMOS. Resetting the CMOS is not part of the normal maintenance of a server.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Rack Mounts** > **Servers**.
Step 3 Choose the server for which you want to reset the CMOS.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Recover Server**.
Step 6 In the **Recover Server** dialog box, click **Reset CMOS**, then click **OK**.

Resetting the CIMC for a Rack-Mount Server

Sometimes, with the firmware, troubleshooting a server might require you to reset the CIMC. Resetting the CIMC is not part of the normal maintenance of a server. After you reset the CIMC, the CIMC reboots with the running version of the firmware for that server.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Rack Mounts** > **Servers**.
Step 3 Choose the server for which you want to reset the CIMC.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Recover Server**.
Step 6 In the **Recover Server** dialog box, click **Reset CIMC (Server Controller)**, then click **OK**.

Clearing TPM for a Rack-Mount Server

You can clear TPM only on Cisco UCS M4 blade and rack-mount servers that include support for TPM.

Caution Clearing TPM is a potentially hazardous operation. The OS may stop booting. You may also see loss of data.

Before You Begin

TPM must be enabled.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Rack Mounts** > **Servers**.
Step 3 Choose the server for which you want to clear TPM.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Recover Server**.
Step 6 In the **Recover Server** dialog box, click **Clear TPM**, then click **OK**.

Recovering the Corrupt BIOS on a Rack-Mount Server

Sometimes, an issue with a server might require you to recover the corrupted BIOS. This procedure is not part of the normal maintenance of a server. After you recover the BIOS, the server boots with the running version of the firmware for that server. This radio button might dim if the BIOS does not require recovery or the option is not available for a particular server.

Before You Begin

Important Remove all attached or mapped USB storage from a server before you attempt to recover the corrupt BIOS on that server. If an external USB drive is attached or mapped from vMedia to the server, BIOS recovery fails.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Rack Mounts** > **Servers**.
Step 3 Choose the server for which you want to recover the BIOS.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Recover Server**.
Step 6 In the **Recover Server** dialog box, click **Recover Corrupt BIOS**, then click **OK**.
Step 7 If a confirmation dialog box displays, click **Yes**.
Step 8 In the **Recover Corrupt BIOS** dialog box, specify the version to be activated, then click **OK**.

Viewing the POST Results for a Rack-Mount Server

You can view any errors collected during the Power On Self-Test process for a server and its adapters.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Rack Mounts** > **Servers**.
Step 3 Choose the server for which you want to view the POST results.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **View POST Results**.
The **POST Results** dialog box lists the POST results for the server and its adapters.
Step 6 (Optional) Click the link in the **Affected Object** column to view the properties of that adapter.
Step 7 Click **OK** to close the **POST Results** dialog box.

Issuing an NMI from a Rack-Mount Server

Perform the following procedure if the system remains unresponsive and you need Cisco UCS Manager to issue a Non-Maskable Interrupt (NMI) to the BIOS or operating system from the CIMC. This action creates a core dump or stack trace, depending on the operating system installed on the server.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Rack Mounts** > **Servers**.
Step 3 Choose the server that you want to issue the NMI.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Server Maintenance**.
Step 6 In the **Maintenance** dialog box, click **Diagnostic Interrupt**, then click **OK**.
Cisco UCS Manager sends an NMI to the BIOS or operating system.

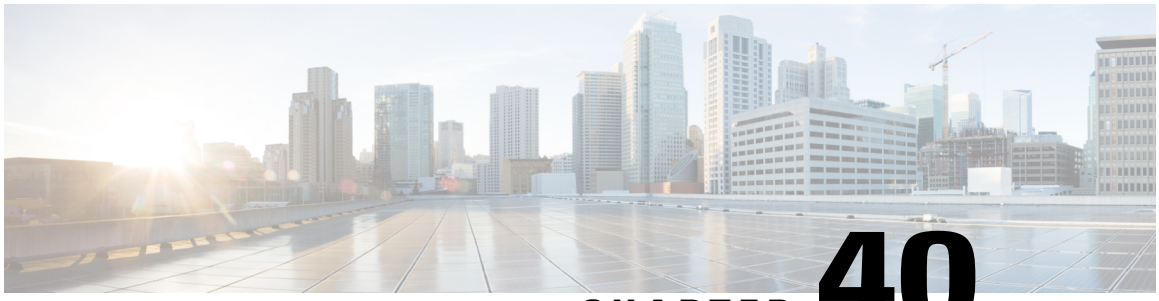

CHAPTER 40

Starting the KVM Console

This chapter includes the following sections:

KVM Console

The KVM console is an interface accessible from the Cisco UCS Manager GUI or the KVM Launch Manager that emulates a direct KVM connection. Unlike the KVM dongle, which requires you to be physically connected to the server, the KVM console allows you to connect to the server from a remote location across the network.

You must ensure that either the server or the service profile associated with the server is configured with a CIMC IP address if you want to use the KVM console to access the server. The KVM console uses the CIMC IP address assigned to a server or a service profile to identify and connect with the correct server in a Cisco UCS domain.

Instead of using CD/DVD or floppy drives directly connected to the server, the KVM console uses virtual media, which are actual disk drives or disk image files that are mapped to virtual CD/DVD or floppy drives. You can map any of the following to virtual drives:

- CD/DVD or floppy drives on your computer
- Disk image files on your computer
- CD/DVD or floppy drives on the network
- Disk image files on the network

Note When you launch the KVM console from the physical server, the system checks if the server is associated to a service profile. If the server is associated to a service profile with an associated management IP address, launches the KVM console using that management IP address. If no management IP address is associated in the service profile, then the system launches the KVM console using the physical server.

Recommendations for Using the KVM Console to Install a Server OS

To install an OS from a virtual CD/DVD or floppy drive, you must ensure that the virtual CD/DVD or floppy drive is set as the first boot device in the service profile.

Installing an OS using the KVM console may be slower than using the KVM dongle because the installation files must be downloaded across the network to the server. If you map a disk drive or disk image file from a network share to a virtual drive, the installation may be even slower because the installation files must be downloaded from the network to the KVM console (your computer) and then from the KVM console to the server. When using this installation method, we recommend that you have the installation media as close as possible to the system with the KVM console.

Virtual KVM Console

The KVM console is an interface accessible from CIMC that emulates a direct keyboard, video, and mouse (KVM) connection to the server. It allows you to connect to and control the server from a remote location and to map physical locations to virtual drives that can by accessed by the server during this KVM session.

Important The KVM console requires Java Runtime Environment (JRE) version 1.5.0 or higher.

KVM Console Tab

This tab provides command line access to the server. The menu options available in this tab are described below.

File Menu

Menu Item	Description
Capture to File button	Opens the **Save** dialog box that allows you to save the current screen as a JPG image.
Paste Text From Clipboard button	Allows you to copy text from a clipboard to the server using the KVM console.
Paste Text From File button	Allows you to copy text from a remote file to the server using the KVM console.
Exit button	Closes the KVM console.

View Menu

Menu Item	Description
Refresh button	Updates the console display with the server's current video output.
Fit button	Resizes the console window to the minimum size needed to display the video image from the server. This option is only available if the console is in **Windowed** mode.
Video Scaling button	Sizes the video image so that the complete image fits within the console window.
Full Screen button	Expands the KVM console so that it fills the entire screen.
Mini Mode button	Displays a thumbnail view of the host server display and provides no input for the keyboard or mouse.

Macros Menu

Choose the keyboard shortcut you want to execute on the remote system.

Menu Item	Description
Server Macros menu	Displays the server side macros downloaded from the Cisco IMC, if any. If no server side macros have been downloaded, then the menu item is disabled.
Static Macros menu	Displays a predefined set of macros.
User Defined Macros menu	Displays the user-defined macros that have been created.
Manage button	Opens the **Configure User Defined Macros** dialog box, which allows you to create and manage macros. System-defined macros cannot be deleted.

Tools Menu

Menu Item	Description
Session Options button, **General** tab	**General** tab in the **Session Options** dialog box allows you to specify whether all keystrokes are passed to the target system when the console is in Windowed mode. The default is no.

Menu Item	Description
Session Options button, **Mouse** tab	The **Mouse** tab in the **Session Options** dialog box allows you to specify the following: • **Termination Key—**The key to use that will terminate single cursor mode. The default is **F12**. • **Mouse Acceleration—**Allows you to change the positioning of a USB mouse. The default is **Absolute Positioning**. **Note** If you are experiencing mouse tracking issues between your system and the remote system, try changing the mouse acceleration mode. If none of the modes solve the problem, you can use the **Single Cursor** tool to ignore the remote mouse in favor of the local mouse. If that does not work, set this option to **No Acceleration** and then go into the operating system on the remote server and set mouse acceleration to **None** in the system settings.
Session Options button, **Security** tab	The **Security** tab in the **Session Options** dialog box allows you to specify the type of unsecured connections which will be accepted without user concurrence, and allows you to clear previously accepted connections.
Session Options button, **Scaling** tab	The **Scaling** tab in the **Session Options** dialog box allows you to specify whether the aspect ratio maintenance will be in effect when the console is in windowed mode or in full screen mode. The default is **Windowed Mode**.
Session Options button, **Mini-Mode** tab	The **Mini-Mode** tab in the **Session Options** dialog box allows you to specify the size of the KVM/vMedia client when in Mini-Mode.
Session Options button, **Certificate** tab	The **Certificate** tab in the **Session Options** dialog box displays the details of the certificate being used in the current session and allows you to set the name and path of the exported certificate file.
Single Cursor button	Turns on the single cursor feature, which offsets mouse alignment issues encountered on some remote operating systems. When you turn this feature on, the mouse pointer is trapped within the viewer window. To turn the feature off, press the termination key specified in the **Session Options** dialog box. The default is **F12**.

Menu Item	Description
Stats button	Opens the **Stats** dialog box, which displays the following: • The statistics of the KVM session • Frame rate measured in the number of frames per second • Bandwidth measured in the number of KBs per second • Compression measured in the percentage of compression being used • Packet rate measured in number of packets per second When vMedia is activated, the **Stats** dialog box displays the following: • Transfer rate of vMedia measured in data transported per second. • The type of local device or image file to which the host server device is mapped. • The elapsed time of the device to map. • The number of bytes sent or received by the server. • The **USB Reset** button to reset all the USB devices connected to the server.
Session User List button	Opens the **Session User List** dialog box that shows all the user IDs that have an active KVM session.
Chat button	Communicates with other Virtual Console users through a chatting interface available within the Virtual Console viewer. Chat is not supported for all Cisco UCS servers. **Note** The chat window cannot be minimized and the chat history is not stored once the window is closed. You can chat even if the server is powered off.

Virtual Media Menu

Name	Description
Create Image	Opens the **Create Image from Folder** dialog box that allows you to create an image of the folder you want to map on the server to the local folder that you want to map on the server. After the system has created the image, it saves the IMG file on your system.

Name	Description
Activate Virtual Devices	Activates a vMedia session that allows you to attach a drive or image file from your local computer or network. **Note** If you have not allowed unsecured connections, you will be prompted to accept the session. If you reject the session, the virtual media session is terminated.
Map CD/DVD	Choose the CD/DVD that you want to access, and click the **Map Device** button to map it to the host server device. **Note** If the **Read Only** checkbox is checked, the server cannot write to the vMedia device even if the device has write capability.
Map Removable Disk	Choose the removable disk that you want to access, and click the **Map Device** button to map it to the host server device. **Note** If the **Read Only** checkbox is checked, the server cannot write to the vMedia device even if the device has write capability.
Map Floppy	Choose the floppy that you want to access, and click the **Map Device** button to map it to the host server device. **Note** If the **Read Only** checkbox is checked, the server cannot write to the vMedia device even if the device has write capability.

KVM Direct Access

KVM direct access allows the administrators that manage the blade and rack servers in your Cisco UCS Manager domain access the KVM for their servers directly using a web browser. This feature allows you to restrict access to the IP addresses of the fabric interconnects, while still allowing your administrators to access the KVM console for the servers they manage.

Note Only out-of-band IPv4 management interface addresses are supported for KVM direct access.

KVM direct access also supports custom applications from which users can navigate to a server management IP address without using the Cisco UCS Manager GUI interface or the KVM Launch Manager.

KVM direct access is supported by providing a management IP address assigned directly to the server or associated to the server with a service profile, to the server's administrator. The server administrator enters the IP address in a browser, and navigates to the Cisco UCS KVM Direct log in page. In the log in page, the user enters their user name and password, and chooses an authentication domain. When they launch Cisco

UCS KVM Direct, the console for the server is displayed, the same way it would if they had accessed the server from the Cisco UCS Manager GUI.

KVM direct access employs self-signed certificates for authentication. When users access a server management IP address or service profile IP address for the first time, a dialog box will be displayed to alert them that they need to add a certificate exception to their browser's cache.

The default communications service that supports Cisco UCS KVM direct access is HTTPS. This cannot be disabled. When a user enters a management IP in a browser using HTTP as part of the address, they will be automatically redirected to the HTTPS service.

To accommodate KVM direct access, be sure that the CIMC Web Service communication service in Cisco UCS Manager is enabled.

Note The CIMC Web Service is enabled by default in Cisco UCS Manager.

Starting the KVM Console from a Server

You can start multiple KVM Console sessions using the addresses assigned to the server.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.

Step 3 Choose the server that you want to access through the **KVM Console**.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **Actions** area, click the **>>** button to the right of **KVM Console**.
The **KVM Console** opens in a separate window and displays a list of available out-of-band and inband addresses associated with the server.

Note If you click **KVM Console** and not the **>>** button, your session will be started using server addresses in the preferential order of inband IPv6 first, inband IPv4 second, and out-of-band IPv4 third.

Step 6 Choose an address from the **Select IP Address** list.
Addresses displayed as **(Inband)** access the server via the uplink ports and those displayed as **(Outband)** access the server via the management interface port.

Step 7 Click **OK**.
The KVM Console is launched using the address you selected.

Tip If the Caps Lock key on your keyboard is on when you open a KVM session, and you subsequently turn off your Caps Lock key, the **KVM Console** may continue to act as if Caps Lock is turned on. To synchronize the KVM Console and your keyboard, press Caps Lock once without the **KVM Console** in focus and then press Caps Lock again with the **KVM Console** in focus.

Step 8 To start another KVM session for the same server, repeat steps 5 through 7.
Another KVM session is started. You can start up to six sessions for a server, depending on the number of addresses that have been configured for it.

Starting the KVM Console from a Service Profile

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand **Servers** > **Service Profiles**.

Step 3 Expand the node for the organization which contains the service profile for which you want to launch the KVM console.
If the system does not include multitenancy, expand the **root** node.

Step 4 Choose the service profile for which you need KVM access to the associated server.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 In the **Actions** area, click the **>>** button to the right of **KVM Console**.
The **KVM Console** opens in a separate window and displays a list of available out-of-band and inband addresses associated with the server.

Note If you click **KVM Console** and not the**>>** button, your session will be started using server addresses in the preferential order of inband IPv6 first, inband IPv4 second, and out-of-band IPv4 third.

Step 7 Choose an address from the **Select IP Address** list.
Addresses displayed as **(Inband)** access the server via the uplink ports and those displayed as **(Outband)** access the server via the management interface port.

Step 8 Click **OK**.
The KVM Console is launched using the address you selected.

Tip If the Caps Lock key on your keyboard is on when you open a KVM session, and you subsequently turn off your Caps Lock key, the **KVM Console** may continue to act as if Caps Lock is turned on. To synchronize the KVM Console and your keyboard, press Caps Lock once without the **KVM Console** in focus and then press Caps Lock again with the **KVM Console** in focus.

Step 9 To start another session for the same server, repeat steps 6 through 8.
Another KVM session is started. You can start up to six sessions for a server, depending on the number of addresses that have been configured for it.

Starting the KVM Console from the KVM Launch Manager

The KVM Launch Manager enables you to access a server through the KVM console without logging in to Cisco UCS Manager.

Before You Begin

To access the KVM console for a server through the KVM Launch Manager, you need the following:

- Cisco UCS username and password.
- Name of the service profile associated with the server for which you want KVM access.

Procedure

Step 1 In your web browser, type or select the web link for Cisco UCS Manager GUI.

Example:
The default web link for HTTP access is `http://UCSManager_IP` for an IPv4 address, or `http://UCSManager_IP6` for an IPv6 address. The default web link for HTTPS access is `https://UCSManager_IP` for an IPv4 address, or `https://UCSManager_IP6` for an IPv6 address. In a standalone configuration, *UCSManager_IP* or *UCSManager_IP6* are the IPv4 or IPv6 addresses, respectively, for the management port on the fabric interconnect. In a cluster configuration, *UCSManager_IP* or *UCSManager_IP6* are the IPv4 or IPv6 addresses, respectively, assigned to Cisco UCS Manager.

Step 2 On the Cisco UCS Manager launch page, click **Launch KVM Manager**.

Step 3 If a **Security Alert** dialog box appears, click **Yes** to accept the security certificate and continue.

Step 4 On the **UCS - KVM Launch Manager Login** page, do the following:

a) Enter your Cisco UCS username and password.
b) (Optional) If your Cisco UCS implementation includes multiple domains, select the appropriate domain from the **Domain** drop-down list.
c) Click **OK**.

Step 5 In the **Service Profiles** table of the KVM Launch Manager, do the following:

a) Locate the row containing the service profile and associated server for which you need KVM access.
b) In the **Launch KVM** column for that server, click **Launch**.
The KVM console opens in a separate window.

Tip If the Caps Lock key on your keyboard is on when you open a KVM session, and you subsequently turn off your Caps Lock key, the **KVM Console** may continue to act as if Caps Lock is turned on. To synchronize the KVM Console and your keyboard, press Caps Lock once without the **KVM Console** in focus and then press Caps Lock again with the **KVM Console** in focus.

Starting the KVM Console from the Cisco UCS KVM Direct Web Page

The Cisco UCS KVM Direct login page enables you to access a server directly from a web browser without logging in to Cisco UCS Manager.

Before You Begin

To access the KVM console for a server using the Cisco UCS KVM Direct login page, you need the following:

- A Cisco UCS username and password.
- The server CIMC or service profile IPv4 management address for the server you want to access.

Procedure

Step 1 In your web browser, type or select the web link for the management IP address of the server you want to access.

Step 2 If a **Security Alert** dialog box appears, click **Yes** to create a security exception.
The security exception is permanently stored in your browser's cache.

Step 3 In the Cisco UCS**KVM Direct** dialog box, specify the name, password, and domain.

Step 4 Click **Lauch KVM.**
The KVM console is launched.

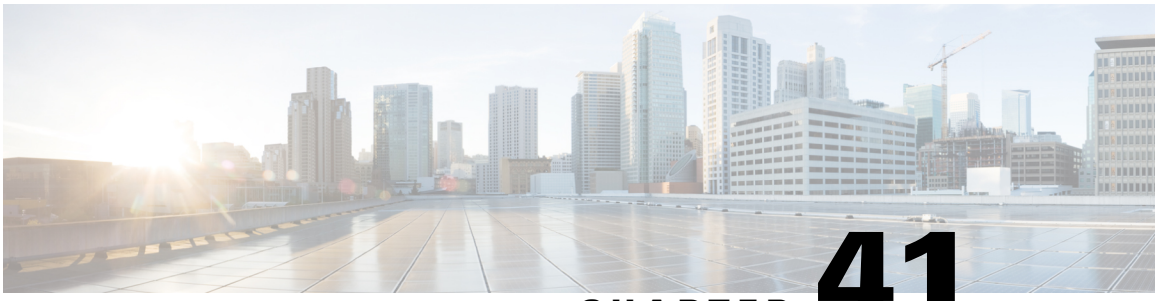

CHAPTER 41

CIMC Session Management

This chapter includes the following sections:

CIMC Session Management

You can view and close any KVM, vMedia, and SOL sessions in Cisco UCS Manager. If you have administrator privileges, you can discontinue the KVM, vMedia, and SoL sessions of any user. . Cisco Integrated Management Controller (CIMC) provides session information to Cisco UCS Manager. When Cisco UCS Manager gets an event from CIMC, it updates its session table and displays the information to all users.

The session information consists of the following information:

- Name—The name of the user who launched the session.
- Session ID—The ID associated with the session. The format of the session ID for blades is [unique identifier] _ [chassis id] _ [Blade id]. The format of the session ID for racks is [unique identifier] _ 0 _ [Rack id].
- Type of session—KVM, vMedia, or SoL.
- Privilege level of the user—Read-Write, Read Only, or Granted.
- Administrative state—Active or Inactive. The value is active if the session is active. The value is inactive if the session terminate command has been issued but the session has not been terminated. This situation occurs when FSM of the server is in progress with another operation or when the connectivity to CIMC is lost.
- Source Address—The IP address of the computer from which the session was opened.
- Service Profile—The service profile associated with the session. The service profile attribute value for a CIMC session is displayed only if the session is opened on an IP address that is provided from the service profile.
- Server—The name of the server associated with the session.
- Login time—The date and time the session started.
- Last Update Time—The last time the session information was updated by CIMC.

A new session is generally added when a user connects to KVM, vMedia, or SOL. A Pnuos vMedia session will be displayed in the session table during the server discovery with the user name __vmediausr__.

The CIMC session data is available under the **CIMC Sessions** tab in Cisco UCS Manager GUI. Any CIMC session terminated by the user is audit logged with proper details.

Note To perform the GUI and CLI tasks that are described in this guide, a CIMC image version of 2.1(2a) or above is required for the session management support for the blade servers. The latest CIMC image version of 1.5(1l) and above is required for the rack-servers.

Viewing All Open CIMC Sessions

This task describes one way to view all CIMC sessions opened globally on Cisco UCS Manager. You can view CIMC sessions of all servers opened by local, remote, or IPMI users in a single page.

Procedure

Step 1 In the **Navigation** pane, click the **Admin** tab.

Step 2 On the **Admin** tab, click **User Management** > **User Services**.

Step 3 In the **Work** pane, click the **CIMC Sessions** tab.

Viewing the CIMC Sessions of a Server

This task describes how to view the CIMC sessions of a specific server. You can view the CIMC sessions opened on the server and the service profile.

Procedure

Step 1 In the **Navigation** pane, click the **Equipment** tab.

Step 2 On the **Equipment** tab, expand **Chassis** > **Chassis Number** > **Servers** > ***Server Number***.

Step 3 In the **Work** pane, click the **CIMC Sessions** tab.

Viewing the CIMC Sessions of a Service Profile

This task describes how to view the CIMC sessions of a specific service profile.

Note A CIMC session will only be displayed under a service profile if the session was opened on an IP address provided from that service profile.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 On the **Servers** tab, expand **Servers** > **Service Profiles** > **Root** > ***Service Profile Name***.

Step 3 In the **Work** pane, click the **CIMC Sessions** tab.

Viewing the CIMC Sessions Opened by a Local User

This task describes how to view CIMC sessions opened by a local user.

Procedure

Step 1 In the **Navigation** pane, click the **Admin** tab.

Step 2 On the **Admin** tab, expand **User Management** > **User Services** > **Locally Authenticated Users** > ***User Name***.

Step 3 In the **Work** pane, click the **CIMC Sessions** tab.

Viewing the CIMC Sessions Opened by a Remote User

This task describes how to view CIMC sessions opened by a remote user.

Procedure

Step 1 In the **Navigation** pane, click the **Admin** tab.

Step 2 On the **Admin** tab, expand **User Management** > **User Services** > **Remotely Authenticated Users** > ***User Name***.

Step 3 In the **Work** pane, click the **CIMC Sessions** tab.

Clearing All Open CIMC Sessions

This task describes how to clear all open CIMC sessions. You can clear the CIMC sessions of all servers and service-profiles opened by the local, remote, or IPMI users.

Procedure

Step 1 In the **Navigation** pane, click the **Admin** tab.
Step 2 On the **Admin** tab, click **User Management**.
Step 3 In the **Work** pane, click the **CIMC Sessions** tab.
Step 4 Select the CIMC sessions you want to clear, right-click, and select **Clear CIMC Session**.
Step 5 If the Cisco UCS Manager GUI displays a confirmation dialog box, click **Yes**.

Clearing the CIMC Sessions of a Server

This task describes how to clear the CIMC session of a server. You can clear one or more CIMC sessions that are opened on a server.

Procedure

Step 1 In the **Navigation** pane, click the **Equipment** tab.
Step 2 On the **Equipment** tab, expand **Servers** > ***Server Name***.
Step 3 In the **Work** pane, click the **CIMC Sessions** tab.
Step 4 Select the CIMC sessions that you want to clear, right-click, and select **Clear CIMC Session**.
Step 5 If the Cisco UCS Manager GUI displays a confirmation dialog box, click **Yes**.

Clearing the CIMC Sessions of a Service Profile

This task describes how to clear the CIMC sessions of a service profile. You can clear one or more CIMC sessions opened with an IP address provided on the service-profile.

Procedure

Step 1 In the **Navigation** pane, click the **Servers** tab.
Step 2 On the **Servers** tab, expand **Servers** > **Service Profiles** > **root** > ***Service Profile Name***.
Step 3 In the **Work** pane, click the **CIMC Sessions** tab.
Step 4 Select the CIMC sessions that you want to clear, right-click, and select **Clear CIMC Session**.
Step 5 If the Cisco UCS Manager GUI displays a confirmation dialog box, click **Yes**.

Clearing the CIMC Sessions of a Local User

This task describes how to clear the CIMC sessions of a local user. You can clear one or more CIMC sessions opened by a local user.

Procedure

Step 1 In the **Navigation** pane, click the **Admin** tab.
Step 2 On the **Admin** tab, expand **User Services** > **Locally Authenticated Users** > ***User Name***.
Step 3 In the **Work** pane, click the **General** tab.
Step 4 Under the **General** tab, expand the **CIMC Sessions** section.
Step 5 Select the CIMC sessions you want to clear, right-click, and select **Clear CIMC Session**.
Step 6 If the Cisco UCS Manager GUI displays a confirmation dialog box, click **Yes**.

Clearing the CIMC Sessions of a Remote User

This task describes how to clear the CIMC sessions of a remote user. You can clear one or more CIMC sessions opened by a remote user.

Procedure

Step 1 In the **Navigation** pane, click the **Admin** tab.
Step 2 On the **Admin** tab, expand **User Services** > **Remotely Authenticated Users** > ***User Name***.
Step 3 In the **Work** pane, click the **General** tab.
Step 4 Under the **General** tab, expand the **CIMC Sessions** section.
Step 5 Select the CIMC sessions that you want to clear, right-click, and select **Clear CIMC Session**.
Step 6 If the Cisco UCS Manager GUI displays a confirmation dialog box, click **Yes**.

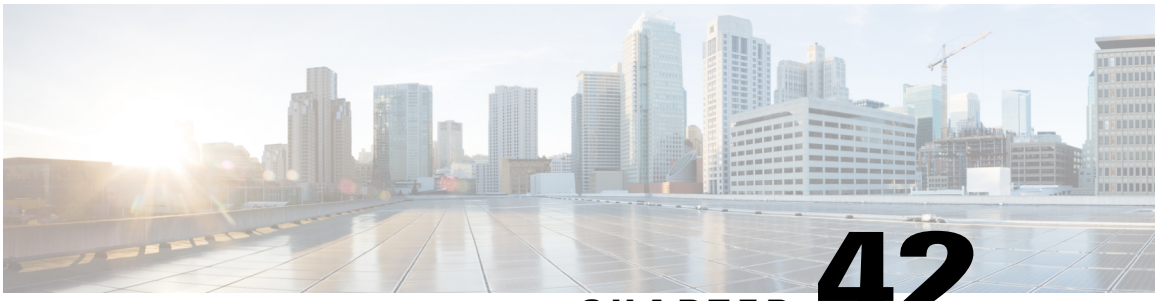

CHAPTER 42

Managing the I/O Modules

This chapter includes the following sections:

I/O Module Management in Cisco UCS Manager GUI

You can manage and monitor all I/O modules in a Cisco UCS domain through Cisco UCS Manager GUI.

Acknowledging an IO Module

Cisco UCS Manager Release 2.2(4) introduces the ability to acknowledge a specific IO module in a chassis.

Note This operation rebuilds the network connectivity between the IO module and the Fabrics to which it is connected.

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click **Equipment**.	
Step 2	Expand **Equipment** > **Chassis** > *Chassis Number* > **IO Modules**.	
Step 3	Choose the I/O module that you want to acknowledge.	
Step 4	In the **Work** pane, click the **General** tab.	
Step 5	In the **Actions** area, click **Acknowledge IO Module**.	

	Command or Action	Purpose
Step 6	In the **Acknowledge IO Module** confirmation box, click **Yes**.	

Resetting an I/O Module

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Chassis** > ***Chassis Number*** > **IO Modules**.
Step 3 Choose the I/O module that you want to reset.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Reset IO Module**.
Step 6 If a confirmation dialog box displays, click **Yes**.

Viewing the POST Results for an I/O Module

You can view any errors collected during the Power On Self-Test process for an I/O module.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Chassis** > ***Chassis Number*** > **IO Modules**.
Step 3 Choose the I/O module for which you want to view the POST results.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **View POST Results**.
The **POST Results** dialog box lists the POST results for the I/O module.
Step 6 Click **OK** to close the **POST Results** dialog box.

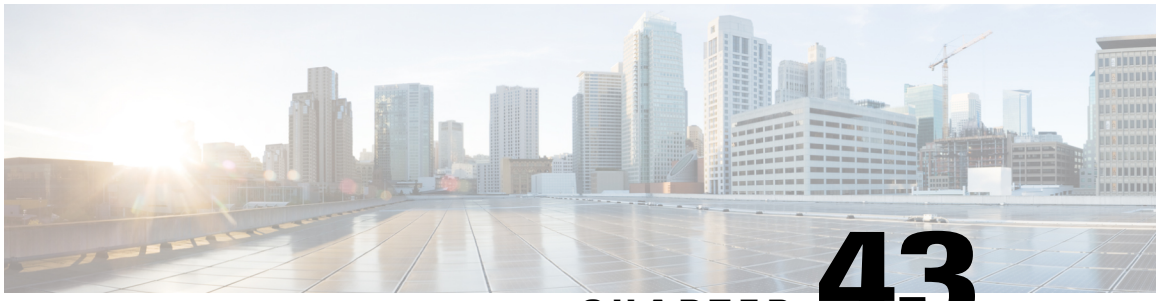

CHAPTER 43

Backing Up and Restoring the Configuration

This chapter includes the following sections:

Backup Operations in UCS

When you perform a backup through Cisco UCS Manager, you take a snapshot of all or part of the system configuration and export the file to a location on your network. You cannot use Cisco UCS Manager to back up data on the servers.

You can perform a backup while the system is up and running. The backup operation only saves information from the management plane. It does not have any impact on the server or network traffic.

Backup Types

You can perform one or more of the following types of backups in Cisco UCS Manager and Cisco UCS Central:

- **Full state—**A binary file that includes a snapshot of the entire system. You can use the file generated from this backup to restore the system during disaster recovery. This file can restore or rebuild the configuration on the original fabric interconnect, or recreate the configuration on a different fabric interconnect. You cannot use this file for an import.

Note You can only use a full state backup file to restore a system that is running the same version as the system from which the backup file was exported.

- **All configuration—**An XML file that includes all system and logical configuration settings. You can use the file generated from this backup to import these configuration settings to the original fabric interconnect or to a different fabric interconnect. You cannot use this file for a system restore. This file does not include passwords for locally authenticated users.
- **System configuration—**An XML file that includes all system configuration settings such as usernames, roles, and locales. You can use the file generated from this backup to import these configuration settings to the original fabric interconnect or to a different fabric interconnect. You cannot use this file for a system restore.
- **Logical configuration—**An XML file that includes all logical configuration settings such as service profiles, VLANs, VSANs, pools, and policies. You can use the file generated from this backup to import these configuration settings to the original fabric interconnect or to a different fabric interconnect. You cannot use this file for a system restore.

Considerations and Recommendations for Backup Operations

Before you create a backup operation, consider the following:

Backup Locations

The backup location is the destination or folder on the network where you want Cisco UCS Manager to export the backup file. You can maintain only one backup operation for each location where you plan to save a backup file.

Potential to Overwrite Backup Files

If you rerun a backup operation without changing the filename, Cisco UCS Manager overwrites the existing file on the server. To avoid overwriting existing backup files, change the filename in the backup operation or copy the existing file to another location.

Multiple Types of Backups

You can run and export more than one type of backup to the same location. Change the backup type before you rerun the backup operation. We recommend that you change the filename for easier identification and to avoid overwriting the existing backup file.

Scheduled Backups

You can create a backup operation in advance and leave the admin state disabled, until you are ready to run the backup. Cisco UCS Manager does not run the backup operation, save, or export the configuration file until you set the admin state of the backup operation to enabled.

Incremental Backups

You cannot perform incremental backups.

Encryption of Full State Backups

Full state backups are encrypted so that passwords and other sensitive information are not exported as clear text.

FSM Tasks for Backup Policy and Configuration Export Policy

When configuring both **Backup Policy** and **Config Export Policy** on the **Policy Backup & Export** tab and using the same hostname for both policies, Cisco UCS Manager will create only one **Backup Operation** in the **Backup Configuration** page to run both tasks. Each policy run will not have a separate FSM task.

To see a separate FSM task for each policy, you can create a hostname alias in your DNS server to point to the same FTP/TFTP/SCP/SFTP server. Then you can use one hostname for the **Backup Policy** and another hostname for the **Config Export Policy**.

Scheduled Backups

You can configure policies in Cisco UCS to schedule the following types of backups:

- Full state
- All configuration

You cannot schedule any other type of backup.

Full State Backup Policy

The full state backup policy allows you to schedule regular full state backups of a snapshot of the entire system. You can choose whether to configure the full state backup to occur on a daily, weekly, or biweekly basis.

Cisco UCS Manager maintains a maximum number of backup files on the remote server. The maxfiles parameter is used when Cisco UCS Manager is registered with Cisco UCS Central. The maxfiles parameter is user configurable on Cisco UCS Central and controls the number of backup files stored on Cisco UCS Central.

If Cisco UCS Manager is not registered with Cisco UCS Central, and the user is storing backup files on a remote backup server, the backup files are not managed by Cisco UCS Manager. The remote machine server administrator must monitor the disk usage and rotate the backup files to create space for new backup files.

All Configuration Export Policy

The all configuration backup policy allows you to schedule a regular backup and export of all system and logical configuration settings. This backup does not include passwords for locally authenticated users. You can choose whether to configure the all configuration backup to occur on a daily, weekly, or bi-weekly basis.

Cisco UCS maintains a maximum number of backup files on the remote server. When that number is exceeded, Cisco UCS overwrites the oldest backup file.

Import Configuration

You can import any configuration file that was exported from Cisco UCS. The file does not need to have been exported from the same Cisco UCS.

Note You cannot import configuration from a higher release to a lower release.

The import function is available for all configuration, system configuration, and logical configuration files. You can perform an import while the system is up and running. An import operation modifies information on the management plane only. Some modifications caused by an import operation, such as a change to a vNIC assigned to a server, can cause a server reboot or other operations that disrupt traffic.

You cannot schedule an import operation. You can, however, create an import operation in advance and leave the admin state disabled until you are ready to run the import.Cisco UCS will not run the import operation on the configuration file until you set the admin state to enabled.

You can maintain only one import operation for each location where you saved a configuration backup file.

Import Methods

You can use one of the following methods to import and update a system configuration through Cisco UCS:

- **Merge—**The information in the imported configuration file is compared with the existing configuration information. If there are conflicts, the import operation overwrites the information on the Cisco UCS domain with the information in the import configuration file.
- **Replace—**The current configuration information is replaced with the information in the imported configuration file one object at a time.

System Restore

You can use the restore function for disaster recovery.

You can restore a system configuration from any full state backup file that was exported from Cisco UCS. The file does not need to have been exported from Cisco UCS on the system that you are restoring. When restoring using a backup file that was exported from a different system, we recommend that you use a system with the same or similar system configuration and hardware, including fabric interconnects, servers, adapters, and I/O module or FEX connectivity. Mismatched hardware and system configuration can lead to the restored system not fully functioning. If there is a mismatch between the I/O module links or servers on the two systems, acknowledge the chassis and servers after the restore operation.

The restore function is only available for a full state backup file. You cannot import a full state backup file. You perform a restore through the initial system setup. For more information, see the appropriate *Cisco UCS Central Installation and Upgrade Guide*.

Note You can only use a full state backup file to restore a system that is running the same version as the system from which the backup file was exported.

Required User Role for Backup and Import Operations

You must have a user account that includes the admin role to create and run backup and import operations.

Configuring Backup Operations

Creating a Backup Operation

Before You Begin

Obtain the backup server IPv4 or IPv6 address and authentication credentials.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Click the **All** node.

Step 3 In the **Work** pane, click the **General** tab.

Step 4 In the **Actions** area, click **Backup Configuration**.

Step 5 In the **Backup Configuration** dialog box, click **Create Backup Operation**.

Step 6 In the **Create Backup Operation** dialog box, complete the following fields:

Name	Description
Admin State field	This can be one of the following: • **Enabled—**Cisco UCS Manager runs the backup operation as soon as you click **OK**. • **Disabled—**Cisco UCS Manager does not run the backup operation when you click **OK**. If you select this option, all fields in the dialog box remain visible. However, you must manually run the backup from the **Backup Configuration** dialog box.

Name	Description
Type field	The information saved in the backup configuration file. This can be one of the following: • **Full state—**A binary file that includes a snapshot of the entire system. You can use the file generated from this backup to restore the system during disaster recovery. This file can restore or rebuild the configuration on the original fabric interconnect, or recreate the configuration on a different fabric interconnect. You cannot use this file for an import. **Note** You can only use a full state backup file to restore a system that is running the same version as the system from which the backup file was exported. • **All configuration—**An XML file that includes all system and logical configuration settings. You can use the file generated from this backup to import these configuration settings to the original fabric interconnect or to a different fabric interconnect. You cannot use this file for a system restore. This file does not include passwords for locally authenticated users. • **System configuration—**An XML file that includes all system configuration settings such as usernames, roles, and locales. You can use the file generated from this backup to import these configuration settings to the original fabric interconnect or to a different fabric interconnect. You cannot use this file for a system restore. • **Logical configuration—**An XML file that includes all logical configuration settings such as service profiles, VLANs, VSANs, pools, and policies. You can use the file generated from this backup to import these configuration settings to the original fabric interconnect or to a different fabric interconnect. You cannot use this file for a system restore.

Name	Description
Preserve Identities check box	This checkbox remains selected for **All Configuration** and **System Configuration** type of backup operation, and provides the following functionality: • **All Configuration—**The backup file preserves all identities derived from pools, including vHBAs, WWPNs, WWNN, vNICs, MACs and UUIDs. Also, the identities for Chassis, FEX, Rack Servers, and user labels for Chassis, FEX, Rack Servers, IOMs and Blade Servers are preserved. **Note** If this check box is not selected the identities will be reassigned and user labels will be lost after a restore. • **System Configuration—**The backup file preserves identities for Chassis, FEX, Rack Servers, and user labels for Chassis, FEX, Rack Servers, IOMs and Blade Servers. **Note** If this check box is not selected the identities will be reassigned and user labels will be lost after a restore. If this checkbox is selected for **Logical Configuration** type of backup operation, the backup file preserves all identities derived from pools, including vHBAs, WWPNs, WWNN, vNICs, MACs and UUIDs. **Note** If this check box is not selected the identities will be reassigned and user labels will be lost after a restore.
Location of the Backup File field	Where the backup file should be saved. This can be one of the following: • **Remote File System—**The backup XML file is saved to a remote server. Cisco UCS Manager GUI displays the fields described below that allow you to specify the protocol, host, filename, username, and password for the remote system. • **Local File System—**The backup XML file is saved locally. Cisco UCS Manager GUI displays the **Filename** field with an associated **Browse** button that let you specify the name and location for the backup file. **Note** Once you click **OK**, the location cannot be changed.

Name	Description
Protocol field	The protocol to use when communicating with the remote server. This can be one of the following: • **FTP** • **TFTP** • **SCP** • **SFTP** • **USB A—**The USB drive inserted into fabric interconnect A. This option is only available for certain system configurations. • **USB B—**The USB drive inserted into fabric interconnect B. This option is only available for certain system configurations.
Hostname field	The hostname, IPv4 or IPv6 address of the location where the backup file is stored. This can be a server, storage array, local drive, or any read/write media that the fabric interconnect can access through the network. **Note** If you use a hostname rather than an IPv4 or IPv6 address, you must configure a DNS server. If the Cisco UCS domain is not registered with Cisco UCS Central or DNS management is set to **local**, configure a DNS server in Cisco UCS Manager. If the Cisco UCS domain is registered with Cisco UCS Central and DNS management is set to **global**, configure a DNS server in Cisco UCS Central.
Remote File field	The full path to the backup configuration file. This field can contain the filename as well as the path. If you omit the filename, the backup procedure assigns a name to the file.
User field	The username the system should use to log in to the remote server. This field does not apply if the protocol is TFTP or USB.
Password field	The password for the remote server username. This field does not apply if the protocol is TFTP or USB. Cisco UCS Manager does not store this password. Therefore, you do not need to enter this password unless you intend to enable and run the backup operation immediately.

Step 7 Click **OK**.

Step 8 If Cisco UCS Manager displays a confirmation dialog box, click **OK**.

If you set the **Admin State** field to enabled, Cisco UCS Manager takes a snapshot of the configuration type that you selected and exports the file to the network location. The backup operation displays in the **Backup Operations** table in the **Backup Configuration** dialog box.

Step 9 (Optional) To view the progress of the backup operation, do the following:

a) If the operation does not display in the **Properties** area, click the operation in the **Backup Operations** table.

b) In the **Properties** area, click the down arrows on the **FSM Details** bar.

The **FSM Details** area expands and displays the operation status.

Step 10 Click **OK** to close the **Backup Configuration** dialog box.
The backup operation continues to run until it is completed. To view the progress, re-open the **Backup Configuration** dialog box.

Running a Backup Operation

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Click the **All** node.

Step 3 In the **Work** pane, click the **General** tab.

Step 4 In the **Actions** area, click **Backup Configuration**.

Step 5 In the **Backup Operations** table of the **Backup Configuration** dialog box, click the backup operation that you want to run.
The details of the selected backup operation display in the **Properties** area.

Step 6 In the **Properties** area, complete the following fields:

a) In the **Admin State** field, click the **Enabled** radio button.

b) For all protocols except TFTP, enter the password for the username in the **Password** field.

c) (Optional) Change the content of the other available fields.

Note If you change other fields -- such as resetting a scheduled backup from weekly to daily -- you must re-enter your user name and password. Otherwise, an FI backup will fail.

Step 7 Click **Apply**.
Cisco UCS Manager takes a snapshot of the configuration type that you selected and exports the file to the network location. The backup operation displays in the **Backup Operations** table in the **Backup Configuration** dialog box.

Step 8 (Optional) To view the progress of the backup operation, click the down arrows on the **FSM Details** bar.
The **FSM Details** area expands and displays the operation status.

Step 9 Click **OK** to close the **Backup Configuration** dialog box.
The backup operation continues to run until it is completed. To view the progress, re-open the **Backup Configuration** dialog box.

Modifying a Backup Operation

You can modify a backup operation to save a file of another backup type to that location or to change the filename and avoid overwriting previous backup files.

Note You can only use a full state backup file to restore a system that is running the same version as the system from which the backup file was exported.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Click the **All** node.

Step 3 In the **Work** pane, click the **General** tab.

Step 4 In the **Actions** area, click **Backup Configuration**.

Step 5 In the **Backup Operations** area of the **Backup Configuration** dialog box, click the backup operation that you want to modify.
The details of the selected backup operation display in the **Properties** area. If the backup operation is in a disabled state, the fields are dimmed.

Step 6 In the **Admin State** field, click the **enabled** radio button.

Step 7 Modify the appropriate fields.
You do not have to enter the password unless you want to run the backup operation immediately.

Step 8 (Optional) If you do not want to run the backup operation immediately, click the **disabled** radio button in the **Admin State** field.

Step 9 Click **OK**.

Deleting One or More Backup Operations

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Click the **All** node.

Step 3 In the **Work** pane, click the **General** tab.

Step 4 In the **Actions** area, click **Backup Configuration**.

Step 5 In the **Backup Operations** table of the **Backup Configuration** dialog box, click the backup operations that you want to delete.

Tip You cannot click a backup operation in the table if the admin state of the operation is set to **Enabled.**

Step 6 Click the **Delete** icon in the icon bar of the **Backup Operations** table.

Step 7 If a confirmation dialog box displays, click **Yes**.

Step 8 In the **Backup Configuration** dialog box, click one of the following:

Option	Description
Apply	Deletes the selected backup operations without closing the dialog box.
OK	Deletes the selected backup operations and closes the dialog box.

Configuring Scheduled Backups

Configuring the Full State Backup Policy

Before You Begin

Obtain the backup server IPv4 or IPv6 address and authentication credentials.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Click the **All** node.

Step 3 In the **Work** pane, click the **Backup and Export Policy** tab.

Step 4 In the **Full State Backup Policy** area, complete the following fields:

Name	Description
Hostname field	The hostname, IPv4 or IPv6 address of the location where the policy backup file is stored. This can be a server, storage array, local drive, or any read/write media that the fabric interconnect can access through the network. **Note** If you use a hostname rather than an IPv4 or IPv6 address, you must configure a DNS server. If the Cisco UCS domain is not registered with Cisco UCS Central or DNS management is set to **local**, configure a DNS server in Cisco UCS Manager. If the Cisco UCS domain is registered with Cisco UCS Central and DNS management is set to **global**, configure a DNS server in Cisco UCS Central.

Name	Description
Protocol field	The protocol to use when communicating with the remote server. This can be one of the following: • **FTP** • **TFTP** • **SCP** • **SFTP** • **USB A—**The USB drive inserted into fabric interconnect A. This option is only available for certain system configurations. • **USB B—**The USB drive inserted into fabric interconnect B. This option is only available for certain system configurations.
User field	The username the system should use to log in to the remote server. This field does not apply if the protocol is TFTP or USB.
Password field	The password for the remote server username. This field does not apply if the protocol is TFTP or USB.
Remote File field	The full path to the policy backup file. This field can contain the filename as well as the path. If you omit the filename, the backup procedure assigns a name to the file.
Admin State field	This can be one of the following: • **Enabled—**Cisco UCS Manager backs up all policy information using the schedule specified in the **Schedule** field. • **Disabled—**Cisco UCS Manager does not back up policy information.
Schedule field	The frequency with which Cisco UCS Manager backs up policy information.
Max Files field	The maximum number of backup files that Cisco UCS Manager maintains. This value cannot be changed.
Description field	The description of the backup policy. The default description is **Database Backup Policy**. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Step 5 (Optional) In the **Backup/Export Config Reminder** area, complete the following fields:

Name	Description
Admin State column	This can be one of the following: • **Enable—**Cisco UCS Manager raises a fault if a backup is not taken during the specified time period. • **Disable—**Cisco UCS Manager does not raise a fault if a backup is not taken during the specified time period.
Remind Me After (days) column	The number of days before you are reminded to take a backup. Enter an integer between 1 and 365. The default value is 30 days.

Step 6 Click **Save Changes**.

Configuring the All Configuration Export Policy

Before You Begin

Obtain the backup server IPv4 or IPv6 address and authentication credentials.

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Click the **All** node.

Step 3 In the **Work** pane, click the **Policy Backup & Export** tab.

Step 4 In the **Config Export Policy** area, complete the following fields:

Name	Description
Hostname field	The hostname, IPv4 or IPv6 address of the location where the configuration backup file is stored. This can be a server, storage array, local drive, or any read/write media that the fabric interconnect can access through the network. **Note** If you use a hostname rather than an IPv4 or IPv6 address, you must configure a DNS server. If the Cisco UCS domain is not registered with Cisco UCS Central or DNS management is set to **local**, configure a DNS server in Cisco UCS Manager. If the Cisco UCS domain is registered with Cisco UCS Central and DNS management is set to **global**, configure a DNS server in Cisco UCS Central.

Name	Description
Protocol field	The protocol to use when communicating with the remote server. This can be one of the following: • **FTP** • **TFTP** • **SCP** • **SFTP** • **USB A—**The USB drive inserted into fabric interconnect A. This option is only available for certain system configurations. • **USB B—**The USB drive inserted into fabric interconnect B. This option is only available for certain system configurations.
User field	The username the system should use to log in to the remote server. This field does not apply if the protocol is TFTP or USB.
Password field	The password for the remote server username. This field does not apply if the protocol is TFTP or USB.
Remote File field	The full path to the backup configuration file. This field can contain the filename as well as the path. If you omit the filename, the backup procedure assigns a name to the file.
Admin State field	This can be one of the following: • **Enabled—**Cisco UCS Manager backs up all policy information using the schedule specified in the **Schedule** field. • **Disabled—**Cisco UCS Manager does not back up policy information.
Schedule field	The frequency with which Cisco UCS Manager backs up policy information.
Max Files field	The maximum number of configuration backup files that Cisco UCS Manager maintains. This value cannot be changed.
Description field	The description of the configuration export policy. The default description is **Configuration Export Policy**. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Step 5 (Optional) In the **Backup/Export Config Reminder** area, complete the following fields:

Name	Description
Admin State column	This can be one of the following: • **Enable—**Cisco UCS Manager raises a fault if a backup is not taken during the specified time period. • **Disable—**Cisco UCS Manager does not raise a fault if a backup is not taken during the specified time period.
Remind Me After (days) column	The number of days before you are reminded to take a backup. Enter an integer between 1 and 365. The default value is 30 days.

Step 6 Click **Save Changes**.

Configuring Import Operations

Creating an Import Operation

You cannot import a Full State backup file. You can import any of the following configuration files:

- All configuration
- System configuration
- Logical configuration

Before You Begin

Collect the following information to import a configuration file:

- Backup server IP address and authentication credentials
- Fully-qualified name of a backup file

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Click the **All** node.

Step 3 In the **Work** pane, click the **General** tab.

Step 4 In the **Actions** area, click **Import Configuration**.

Step 5 In the **Import Configuration** dialog box, click **Create Import Operation**.

Step 6 In the **Create Import Operation** dialog box, complete the following fields:

Name	Description
Admin State field	This can be one of the following: • **Enabled—**Cisco UCS Manager runs the import operation as soon as you click **OK**. • **Disabled—**Cisco UCS Manager does not run the import operation when you click **OK**. If you select this option, all fields in the dialog box remain visible. However, you must manually run the import from the **Import Configuration** dialog box.
Action field	This can be one of the following: • **Merge—**The configuration information is merged with the existing information. If there are conflicts, the system replaces the information on the current system with the information in the import configuration file. • **Replace—**The system takes each object in the import configuration file and overwrites the corresponding object in the current configuration.
Location of the Import File field	Where the backup file that you want to import is located. This can be one of the following: • **Remote File System—**The backup XML file is stored on a remote server. Cisco UCS Manager GUI displays the fields described below that allow you to specify the protocol, host, filename, username, and password for the remote system. • **Local File System—**The backup XML file is stored locally. Cisco UCS Manager GUI displays the **Filename** field with an associated **Browse** button that let you specify the name and location for the backup file to be imported.

Name	Description
Protocol field	The protocol to use when communicating with the remote server. This can be one of the following: • **FTP** • **TFTP** • **SCP** • **SFTP** • **USB A—**The USB drive inserted into fabric interconnect A. This option is only available for certain system configurations. • **USB B—**The USB drive inserted into fabric interconnect B. This option is only available for certain system configurations.
Hostname field	The hostname, IPv4 or IPv6 address from which the configuration file should be imported. **Note** If you use a hostname rather than an IPv4 or IPv6 address, you must configure a DNS server. If the Cisco UCS domain is not registered with Cisco UCS Central or DNS management is set to **local**, configure a DNS server in Cisco UCS Manager. If the Cisco UCS domain is registered with Cisco UCS Central and DNS management is set to **global**, configure a DNS server in Cisco UCS Central.
Remote File field	The name of the XML configuration file.
User field	The username the system should use to log in to the remote server. This field does not apply if the protocol is TFTP or USB.
Password field	The password for the remote server username. This field does not apply if the protocol is TFTP or USB. Cisco UCS Manager does not store this password. Therefore, you do not need to enter this password unless you intend to enable and run the import operation immediately.

Step 7 Click **OK**.

Step 8 In the confirmation dialog box, click **OK**.
If you set the **Admin State** to enabled, Cisco UCS Manager imports the configuration file from the network location. Depending upon which action you selected, the information in the file is either merged with the existing configuration or replaces the existing configuration. The import operation displays in the **Import Operations** table of the **Import Configuration** dialog box.

Step 9 (Optional) To view the progress of the import operation, do the following:

a) If the operation does not automatically display in the **Properties** area, click the operation in the **Import Operations** table.

b) In the **Properties** area, click the down arrows on the **FSM Details** bar.

The **FSM Details** area expands and displays the operation status.

Step 10 Click **OK** to close the **Import Configuration** dialog box.

The import operation continues to run until it is completed. To view the progress, re-open the **Import Configuration** dialog box.

Running an Import Operation

You cannot import a Full State backup file. You can import any of the following configuration files:

- All configuration
- System configuration
- Logical configuration

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Click the **All** node.

Step 3 In the **Work** pane, click the **General** tab.

Step 4 In the **Actions** area, click **Import Configuration**.

Step 5 In the **Import Operations** table of the **Import Configuration** dialog box, click the operation that you want to run.

The details of the selected import operation display in the **Properties** area.

Step 6 In the **Properties** area, complete the following fields:

a) In the **Admin State** field, click the **Enabled** radio button.

b) For all protocols except TFTP, enter the password for the username In the **Password** field.

c) (Optional) Change the content of the other available fields.

Step 7 Click **Apply**.

Cisco UCS Manager imports the configuration file from the network location. Depending upon which action you selected, the information in the file is either merged with the existing configuration or replaces the existing configuration. The import operation displays in the **Import Operations** table of the **Import Configuration** dialog box.

Step 8 (Optional) To view the progress of the import operation, click the down arrows on the **FSM Details** bar.

The **FSM Details** area expands and displays the operation status.

Step 9 Click **OK** to close the **Import Configuration** dialog box.

The import operation continues to run until it is completed. To view the progress, re-open the **Import Configuration** dialog box.

Modifying an Import Operation

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Click the **All** node.

Step 3 In the **Work** pane, click the **General** tab.

Step 4 In the **Actions** area, click **Import Configuration**.

Step 5 In the **Import Operations** area of the **Import Configuration** dialog box, click the import operation that you want to modify.
The details of the selected import operation display in the **Properties** area. If the import operation is in a disabled state, the fields are dimmed.

Step 6 In the **Admin State** field, click the **enabled** radio button.

Step 7 Modify the appropriate fields.
You do not have to enter the password unless you want to run the import operation immediately.

Step 8 (Optional) If you do not want to run the import operation immediately, click the **disabled** radio button in the **Admin State** field.

Step 9 Click **OK**.

Deleting One or More Import Operations

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Click the **All** node.

Step 3 In the **Work** pane, click the **General** tab.

Step 4 In the **Actions** area, click **Import Configuration**.

Step 5 In the **Import Operations** table of the **Backup Configuration** dialog box, click the import operations that you want to delete.

Tip You cannot click an import operation in the table if the admin state of the operation is set to **Enabled.**

Step 6 Click the **Delete** icon in the icon bar of the **Import Operations** table.

Step 7 If a confirmation dialog box displays, click **Yes**.

Step 8 In the **Import Configuration** dialog box, click one of the following:

Option	Description
Apply	Deletes the selected import operations without closing the dialog box.

Option	Description
OK	Deletes the selected import operations and closes the dialog box.

Restoring the Configuration for a Fabric Interconnect

It is recommended that you use a full state backup file to restore a system that is running the same version as the system from which the backup file was exported. You can also use a full state backup to restore a system if they have the same release train. For example, you can use a full state backup taken from a system running Release 2.1(3a) to restore a system running Release 2.1(3f).

To avoid issues with VSAN or VLAN configuration, a backup should be restored on the fabric interconnect that was the primary fabric interconnect at the time of backup.

Before You Begin

Collect the following information to restore the system configuration:

- Fabric interconnect management port IPv4 address and subnet mask, or IPv6 address and prefix
- Default gateway IPv4 or IPv6 address
- Backup server IPv4 or IPv6 address and authentication credentials
- Fully-qualified name of a Full State backup file

Note You must have access to a Full State configuration file to perform a system restore. You cannot perform a system restore with any other type of configuration or backup file.

Procedure

Step 1 Connect to the console port.

Step 2 If the fabric interconnect is off, power on the fabric interconnect.
You will see the power on self-test message as the fabric interconnect boots.

Step 3 At the installation method prompt, enter gui.

Step 4 If the system cannot access a DHCP server, you may be prompted to enter the following information:

- IPv4 or IPv6 address for the management port on the fabric interconnect
- Subnet mask or prefix for the management port on the fabric interconnect
- IPv4 or IPv6 address for the default gateway assigned to the fabric interconnect

Step 5 Copy the web link from the prompt into a web browser and go to the Cisco UCS Manager GUI launch page.

Step 6 On the launch page, select **Express Setup**.

Step 7 On the **Express Setup** page, select **Restore From Backup** and click **Submit**.

Step 8 In the **Protocol** area of the **Cisco UCS Manager Initial Setup** page, select the protocol you want to use to upload the full state backup file:

- **SCP**
- **TFTP**
- **FTP**
- **SFTP**

Step 9 In the **Server Information** area, complete the following fields:

Name	Description
Server IP	The IPv4 or IPv6 address of the computer where the full state backup file is located. This can be a server, storage array, local drive, or any read/write media that the fabric interconnect can access through the network.
Backup File Path	The file path where the full state backup file is located, including the folder names and filename. **Note** You can only use a full state backup file to restore a system that is running the same version as the system from which the backup file was exported.
User ID	The username the system should use to log in to the remote server. This field does not apply if the protocol is TFTP or USB.
Password	The password for the remote server username. This field does not apply if the protocol is TFTP or USB.

Step 10 Click **Submit**.
You can return to the console to watch the progress of the system restore.

The fabric interconnect logs in to the backup server, retrieves a copy of the specified full-state backup file, and restores the system configuration.

For a cluster configuration, you do not need to restore the secondary fabric interconnect. As soon as the secondary fabric interconnect reboots, Cisco UCS Manager synchronizes the configuration with the primary fabric interconnect.

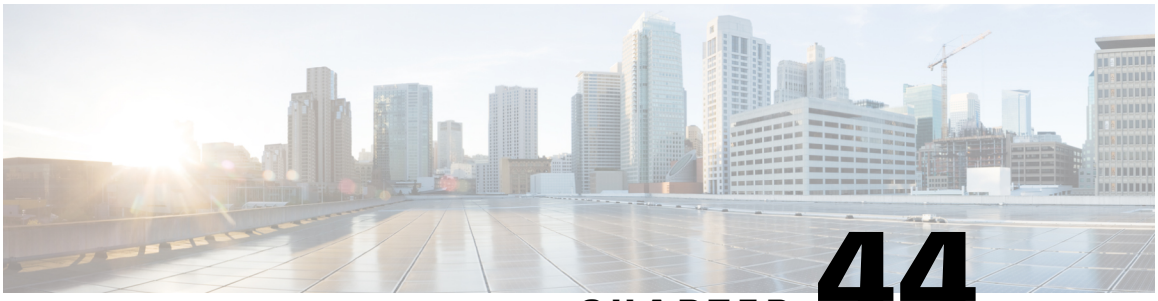

CHAPTER 44

Recovering a Lost Password

This chapter includes the following sections:

Password Recovery for the Admin Account

The admin account is the system administrator or superuser account. If an administrator loses the password to this account, you can have a serious security issue. The procedure to recover the password for the admin account requires you to power cycle all fabric interconnects and will lead to a temporary data transmission outage.

When you recover the password for the admin account, you actually change the password for that account. You cannot retrieve the original password for that account.

You can reset the password for all other local accounts through Cisco UCS Manager. However, you must log in to Cisco UCS Manager with an account that includes aaa or admin privileges.

Caution For Cisco UCS Mini, this procedure requires you to pull all the fabric interconnects in a Cisco UCS domain out of their chassis slots. As a result, all data transmission in the Cisco UCS domain is stopped until you slide the fabric interconnects back into their chassis slots.

For other Cisco UCS configurations, this procedure requires you to power down all fabric interconnects. As a result, all data transmission in the Cisco UCS domain is stopped until you restart the fabric interconnects.

Determining the Leadership Role of a Fabric Interconnect

Important To determine the role of the fabric interconnects in a cluster when the admin password is lost, open the Cisco UCS Manager GUI from the IP addresses of both fabric interconnects. The subordinate fabric interconnect fails with the following message:

```
UCSM GUI is not available on secondary node.
```

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 In the **Equipment** tab, expand **Equipment** > **Fabric Interconnects**.

Step 3 Click the fabric interconnect for which you want to identify the role.

Step 4 In the **Work** pane, click the **General** tab.

Step 5 In the **General** tab, click the down arrows on the **High Availability Details** bar to expand that area.

Step 6 View the **Leadership** field to determine whether the fabric interconnect is the primary or subordinate.

Verifying the Firmware Versions on a Fabric Interconnect

You can use the following procedure to verify the firmware versions on all fabric interconnects in a Cisco UCS domain. You can verify the firmware for a single fabric interconnect through the **Installed Firmware** tab for that fabric interconnect.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 In the **Equipment** tab, click the **Equipment** node.

Step 3 In the **Work** pane, click the **Firmware Management** tab.

Step 4 In the **Installed Firmware** tab, verify that the following firmware versions for each fabric interconnect match the version to which you updated the firmware:

- Kernel version
- System version

Recovering the Admin Account Password in a Standalone Configuration

This procedure will help you to recover the password that you set for the admin account when you performed an initial system setup on the fabric interconnect. The admin account is the system administrator or superuser account.

Before You Begin

1. Physically connect the console port on the fabric interconnect to a computer terminal or console server
2. Determine the running versions of the following firmware:
   - The firmware kernel version on the fabric interconnect
   - The firmware system version

Tip To find this information, you can log in with any user account on the Cisco UCS domain.

Procedure

Step 1 Connect to the console port.

Step 2 Power cycle the fabric interconnect:

a) For Cisco UCS Mini, pull the fabric interconnect out of its chassis slot. For all other configurations, turn off the power to the fabric interconnect.

b) For Cisco UCS Mini, slide the fabric interconnect back into its chassis slot. For all other configurations, turn on the power to the fabric interconnect.

Step 3 In the console, press one of the following key combinations as it boots to get the `loader` prompt:

- Ctrl+l
- Ctrl+Shift+r

You may need to press the selected key combination multiple times before your screen displays the `loader` prompt.

Step 4 Boot the kernel firmware version on the fabric interconnect.

```
loader >
```
boot /installables/switch/
kernel_firmware_version

Example:
```
loader > boot /installables/switch/ucs-6100-k9-kickstart.4.1.3.N2.1.0.11.gbin

loader > boot /installables/switch/ucs-mini-k9-kickstart.5.0.3.N2.3.01a.bin
```

Step 5 Enter config terminal mode.
```
Fabric(boot)#
config terminal
```

Step 6 Reset the admin password.
```
Fabric(boot)(config)#
admin-password
password
```

Choose a strong password that includes at least one capital letter and one number. The password cannot be blank.

The new password displays in clear text mode.

Step 7 Exit config terminal mode and return to the boot prompt.

Step 8 Boot the system firmware version on the fabric interconnect.
```
Fabric(boot)#
load /installables/switch/
system_firmware_version
```

Example:
```
Fabric(boot)# load /installables/switch/ucs-6100-k9-system.4.1.3.N2.1.0.211.bin

Fabric(boot)# load /installables/switch/ucs-mini-k9-system.5.0.3.N2.3.01a.bin
```

Step 9 After the system image loads, log in to Cisco UCS Manager.

Recovering the Admin Account Password in a Cluster Configuration

This procedure will help you to recover the password that you set for the admin account when you performed an initial system setup on the fabric interconnects. The admin account is the system administrator or superuser account.

Before You Begin

1. Physically connect a console port on one of the fabric interconnects to a computer terminal or console server
2. Obtain the following information:
   - The firmware kernel version on the fabric interconnect
   - The firmware system version
   - Which fabric interconnect has the primary leadership role and which is the subordinate

Tip To find this information, you can log in with any user account on the Cisco UCS domain.

Procedure

Step 1 Connect to the console port.

Step 2 For the subordinate fabric interconnect:

a) For Cisco UCS Mini, pull the fabric interconnect out of its chassis slot. For all other configurations, turn off the power to the fabric interconnect.

b) For Cisco UCS Mini, slide the fabric interconnect back into its chassis slot. For all other configurations, turn on the power to the fabric interconnect.

c) In the console, press one of the following key combinations as it boots to get the `loader` prompt:

- Ctrl+l
- Ctrl+Shift+r

You may need to press the selected key combination multiple times before your screen displays the `loader` prompt.

Step 3 Power cycle the primary fabric interconnect:

a) For Cisco UCS Mini, pull the fabric interconnect out of its chassis slot. For all other configurations, turn off the power to the fabric interconnect.

b) For Cisco UCS Mini, slide the fabric interconnect back into its chassis slot. For all other configurations, turn on the power to the fabric interconnect.

Step 4 In the console, press one of the following key combinations as it boots to get the `loader` prompt:

- Ctrl+l
- Ctrl+Shift+r

You may need to press the selected key combination multiple times before your screen displays the `loader` prompt.

Step 5 Boot the kernel firmware version on the primary fabric interconnect.

```
loader > boot /installables/switch/
kernel_firmware_version
```

Example:

```
loader > boot /installables/switch/ucs-6100-k9-kickstart.4.1.3.N2.1.0.11.gbin

loader > boot /installables/switch/ucs-mini-k9-kickstart.5.0.3.N2.3.01a.bin
```

Step 6 Enter config terminal mode.

```
Fabric(boot)# config terminal
```

Step 7 Reset the admin password.

```
Fabric(boot)(config)# admin-password password
```

Choose a strong password that includes at least one capital letter and one number. The password cannot be blank.

The new password displays in clear text mode.

Step 8 Exit config terminal mode and return to the boot prompt.

Step 9 Boot the system firmware version on the primary fabric interconnect.

```
Fabric(boot)# load /installables/switch/
system_firmware_version
```

Example:

```
Fabric(boot)# load /installables/switch/ucs-6100-k9-system.4.1.3.N2.1.0.211.bin

Fabric(boot)# load /installables/switch/ucs-mini-k9-system.5.0.3.N2.3.01a.bin
```

Step 10 After the system image loads, log in to Cisco UCS Manager.

Step 11 In the console for the subordinate fabric interconnect, do the following to bring it up:

a) Boot the kernel firmware version on the subordinate fabric interconnect.

```
loader > boot /installables/switch/
kernel_firmware_version
```

b) Boot the system firmware version on the subordinate fabric interconnect.

```
Fabric(boot)# load /installables/switch/
system_firmware_version
```