Cisco UCS Manager Infrastructure Management Guide, Release 3.2

First Published: 2017-08-18

Last Modified: 2018-03-21

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

© 2017 Cisco Systems, Inc. All rights reserved.

CONTENTS

PREFACE

Preface xi

Audience xi

Conventions xi

Related Cisco UCS Documentation xiii

Documentation Feedback xiii

CHAPTER 1

New and Changed Information 1

New and Changed Information for This Release 1

CHAPTER 2

Overview 3

Cisco UCS Manager User Documentation 3

Infrastructure Management Guide Overview 4

Cisco Unified Computing System Overview 5

Cisco UCS Building Blocks and Connectivity 7

Cisco UCS Fabric Infrastructure Portfolio 8

Cisco UCS I/O Modules and Cisco UCS Fabric Extenders 8

Cisco UCS Chassis 9

Cisco UCS Mini Infrastructure 9

Cisco UCS Infrastructure Virtualization 10

CHAPTER 3

Equipment Policies 13

Chassis/FEX Discovery Policy 13

Pinning 16

Port-Channeling 17

Configuring the Chassis/FEX Discovery Policy 17

Chassis Connectivity Policy 18

Configuring a Chassis Connectivity Policy 18
Rack Server Discovery Policy 19
Configuring the Rack Server Discovery Policy 19
Aging Time for the MAC Address Table 20
Configuring the Aging Time for the MAC Address Table 20

CHAPTER 4 **Chassis Management** 21
Chassis Management in Cisco UCS Manager GUI 21
The Cisco UCS S3260 Chassis 21
Cisco UCS 5108 Blade Server Chassis 22
Extended Chassis for UCS Mini 22
Guidelines for Removing and Decommissioning Chassis 23
Acknowledging a Chassis 23
Decommissioning a Chassis 24
Removing a Chassis 24
Recommissioning a Single Chassis 25
Recommissioning Multiple Chassis 25
Renumbering a Chassis 26
Turning on the Locator LED for a Chassis 27
Turning off the Locator LED for a Chassis 27
Creating a Zoning Policy from Inventory 28
Viewing the POST Results for a Chassis 28

CHAPTER 5 **I/O Module Management** 29
I/O Module Management in Cisco UCS Manager GUI 29
Acknowledging an IO Module 29
Resetting an I/O Module 30
Resetting an I/O Module from a Peer I/O Module 30
Viewing Health Events for an I/O Module 31
Viewing the POST Results for an I/O Module 32

CHAPTER 6 **SIOC Management** 33
SIOC Management in Cisco UCS Manager 33
SIOC Removal or Replacement 33

Acknowledging an SIOC 34
Resetting the CMC 35
CMC Secure Boot 35
Guidelines and Limitations for CMC Secure Boot 35
Enabling CMC Secure Boot 36

CHAPTER 7 **Power Management in Cisco UCS 37**
Power Capping in Cisco UCS 38
Power Policy Configuration 38
Power Policy for Cisco UCS Servers 38
Configuring the Power Policy 39
Power Supply for Redundancy Method 39
Configuring Policy Driven Chassis Group Power Capping 40
Policy Driven Chassis Group Power Capping 40
Power Control Policy 40
Creating a Power Control Policy 41
Deleting a Power Control Policy 43
Power Groups in UCS Manager 43
Creating a Power Group 45
Adding a Chassis to a Power Group 46
Removing a Chassis from a Power Group 47
Deleting a Power Group 47
Blade Level Power Capping 47
Manual Blade Level Power Cap 47
Setting the Blade-Level Power Cap for a Server 48
Viewing the Blade-Level Power Cap 49
Global Power Profiling Policy Configuration 49
Global Power Profiling Policy 49
Configuring the Global Power Profile Policy 50
Global Power Allocation Policy Configuration 50
Global Power Allocation Policy 50
Configuring the Global Power Allocation Policy 51
Power Management During Power-on Operations 51
Power Sync Policy Configuration 52

Power Sync Policy 52
Power Synchronization Behavior 52
Creating a Power Sync Policy 53
Changing a Power Sync Policy 54
Deleting a Power Sync Policy 55
Rack Server Power Management 55
UCS Mini Power Management 56

CHAPTER 8 **Blade Server Hardware Management** 57
Blade Server Management 57
Guidelines for Removing and Decommissioning Blade Servers 58
Recommendations for Avoiding Unexpected Server Power Changes 58
Booting a Blade Server 59
Booting a Rack-Mount Server from the Service Profile 60
Determining the Boot Order of a Blade Server 60
Shutting Down a Blade Server 61
Shutting Down a Server from the Service Profile 61
Resetting a Blade Server 62
Resetting a Blade Server to Factory Default Settings 62
Reacknowledging a Blade Server 63
Removing a Server from a Chassis 64
Deleting the Inband Configuration from a Blade Server 64
Decommissioning a Blade Server 65
Removing a Non-Existent Blade Server Entry 65
Recommissioning a Blade Server 66
Reacknowledging a Server Slot in a Chassis 66
Removing a Non-Existent Blade Server from the Configuration Database 67
Turning the Locator LED for a Blade Server On and Off 67
Turning the Local Disk Locator LED on a Blade Server On and Off 68
Resetting the CMOS for a Blade Server 68
Resetting the CIMC for a Blade Server 69
Clearing TPM for a Blade Server 69
Viewing the POST Results for a Blade Server 70
Issuing an NMI from a Blade Server 70

Viewing Health Events for a Blade Server 71
Health LED Alarms 72
Viewing Health LED Alarms 73
Smart SSD 73
Monitoring SSD Health 74

CHAPTER 9 **Rack-Mount Server Hardware Management** 75
Rack-Mount Server Management 75
Guidelines for Removing and Decommissioning Rack-Mount Servers 76
Recommendations for Avoiding Unexpected Server Power Changes 77
Booting a Rack-Mount Server 77
Booting a Rack-Mount Server from the Service Profile 78
Determining the Boot Order of a Rack-Mount Server 78
Shutting Down a Rack-Mount Server 79
Shutting Down a Server from the Service Profile 79
Resetting a Rack-Mount Server 80
Resetting a Rack-Mount Server to Factory Default Settings 81
Persistent Memory Scrub 81
Reacknowledging a Rack-Mount Server 82
Deleting the Inband Configuration from a Rack-Mount Server 83
Decommissioning a Rack-Mount Server 83
Recommissioning a Rack-Mount Server 84
Renumbering a Rack-Mount Server 84
Removing a Non-Existent Rack-Mount Server from the Configuration Database 85
Turning the Locator LED for a Rack-Mount Server On and Off 85
Turning the Local Disk Locator LED on a Rack-Mount Server On and Off 86
Resetting the CMOS for a Rack-Mount Server 86
Resetting the CIMC for a Rack-Mount Server 87
Clearing TPM for a Rack-Mount Server 87
Issuing an NMI from a Rack-Mount Server 88
Viewing Health Events for a Rack-Mount Server 88
Viewing the POST Results for a Rack-Mount Server 90
Viewing the Power Transition Log 90

CHAPTER 10 S3X60 Server Node Hardware Management 91

Cisco UCS S3260 Server Node Management 91

Booting a Cisco UCS S3260 Server Node 92

Booting a Cisco UCS S3260 Server Node from the Service Profile 92

Determining the Boot Order of a Cisco UCS S3260 Server Node 92

Shutting Down a Cisco UCS S3260 Server Node 93

Shutting Down a Cisco UCS S3260 Server Node from the Service Profile 93

Resetting a Cisco UCS S3260 Server Node 94

Resetting a Cisco UCS S3260 Server Node to Factory Default Settings 95

Reacknowledging a Cisco UCS S3260 Server Node 96

Removing a Cisco UCS S3260 Server Node from a Chassis 96

Deleting the Inband Configuration from a Cisco UCS S3260 Server Node 97

Decommissioning a Cisco UCS S3260 Server Node 97

Recommissioning a Cisco UCS S3260 Server Node 98

Reacknowledging a Server Slot in a S3260 Chassis 98

Removing a Non-Existent Cisco UCS S3260 Server Node from the Configuration Database 99

Turning the Locator LED for a Cisco UCS S3260 Server Node On and Off 99

Turning the Local Disk Locator LED on a Cisco UCS S3260 Server Node On and Off 100

Resetting the CIMC for a Cisco UCS S3260 Server Node 100

Resetting the CMOS for a Cisco UCS S3260 Server Node 101

Issuing an NMI from a Cisco UCS S3260 Server Node 101

Viewing the POST Results for a Cisco UCS S3260 Server Node 102

Viewing Health Events for a Cisco UCS S3260 Server Node 102

Health LED Alarms 104

Viewing Health LED Alarms 104

CHAPTER 11 Virtual Interface Management 105

Virtual Circuits 105

Virtual Interfaces 105

Virtual Interface Subscription Management and Error Handling 106

Virtualization in Cisco UCS 106

Overview of Virtualization 106

Overview of Cisco Virtual Machine Fabric Extender 107

Virtualization with Network Interface Cards and Converged Network Adapters 107
Virtualization with a Virtual Interface Card Adapter 107

CHAPTER 12 Troubleshoot Infrastructure 109
Recovering the Corrupt BIOS on a Blade Server 109
Recovering the Corrupt BIOS on a Rack-Mount Server 110

Preface

- Audience, on page xi
- Conventions, on page xi
- Related Cisco UCS Documentation, on page xiii
- Documentation Feedback, on page xiii

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`.
CLI commands	CLI command keywords appear in **this font**. Variables in a CLI command appear in *this font*.
[ ]	Elements in square brackets are optional.

Text Type	Indication
{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: https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/overview/guide/UCS_roadmap.html

For a complete list of all C-Series documentation, see the *Cisco UCS C-Series Servers Documentation Roadmap* available at the following URL: https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/overview/guide/ucs_rack_roadmap.html.

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@external.cisco.com. We appreciate your feedback.

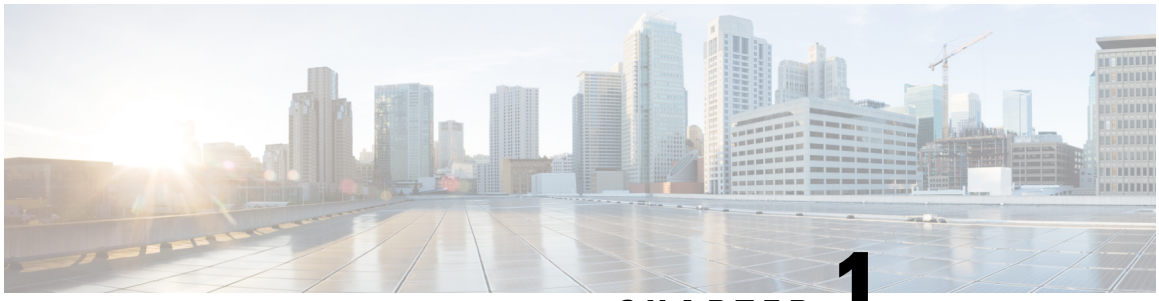

CHAPTER 1

New and Changed Information

• New and Changed Information for This Release, on page 1

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 this guide or of all new features in this release.

Table 1: New Features and Changed Behavior in Cisco UCS Manager, Release 3.2(2)

Feature	Description	Where Documented
Chassis dynamic power rebalance mechanism	Chassis dynamic power rebalance mechanism is enabled by default.	Power Control Policy, on page 40

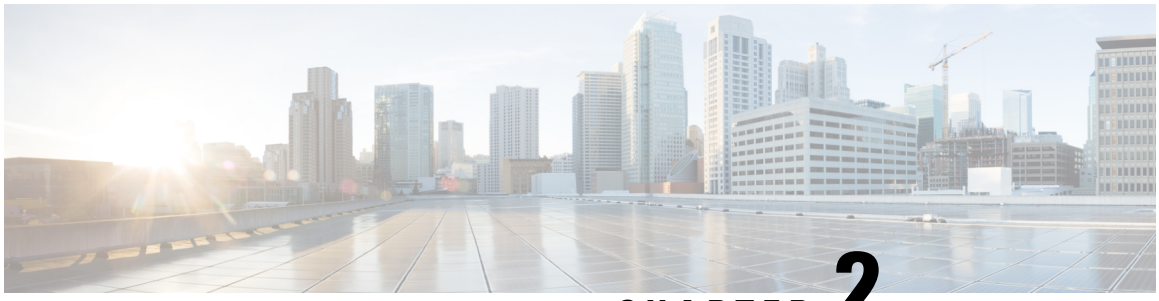

CHAPTER 2

Overview

- Cisco UCS Manager User Documentation, on page 3
- Infrastructure Management Guide Overview, on page 4
- Cisco Unified Computing System Overview, on page 5
- Cisco UCS Building Blocks and Connectivity, on page 7

Cisco UCS Manager User Documentation

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

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

Guide	Description
Cisco UCS Manager Server Management Guide	Discusses the new licenses, registering Cisco UCS domain with Cisco UCS Central, power capping, server boot, server profiles, and server-related policies.
Cisco UCS Manager Storage Management Guide	Discusses all aspects of storage management, such as SAN and VSAN in Cisco UCS Manager.
Cisco UCS Manager Network Management Guide	Discusses all aspects of network management, such as LAN and VLAN connectivity in Cisco UCS Manager.
Cisco UCS Manager System Monitoring Guide	Discusses all aspects of system and health monitoring, including system statistics in Cisco UCS Manager.
Cisco UCS S3260 Server Integration with Cisco UCS Manager	Discusses all aspects of management of UCS S-Series servers that are managed through Cisco UCS Manager.

Infrastructure Management Guide Overview

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

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

Topic	Description
Virtual Interface Management	Overview of Virtualization in Cisco UCS, Virtual Interfaces and procedures to manage them.
Server Troubleshooting	Common troubleshooting scenarios for Servers.

Cisco Unified Computing System Overview

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

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

Figure 1: Cisco Unified Computing System Architecture

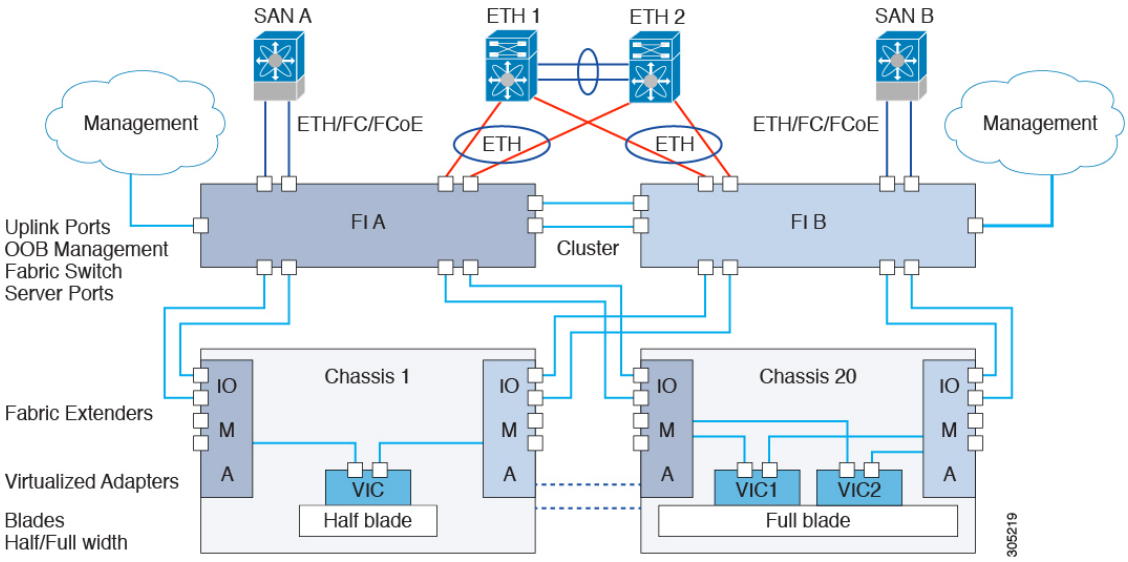

Architectural Simplification

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

High Availability

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

Scalability

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

Flexibility

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

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

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

Optimized for Server Virtualization

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

Cisco UCS Building Blocks and Connectivity

Figure 2: Cisco UCS Building Blocks and Connectivity

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

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

  For more information on Cisco UCS Fabric Interconnects, see Cisco UCS Fabric Infrastructure Portfolio, on page 8.

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

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

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

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

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

Cisco UCS Fabric Infrastructure Portfolio

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

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

Cisco UCS I/O Modules and Cisco UCS Fabric Extenders

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

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

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

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

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

I/O Module Management, on page 29 provides more details about managing I/O Modules.

Cisco UCS Chassis

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

Chassis Management, on page 21 provides details on managing the chassis through Cisco UCS Manager.

Cisco UCS 5108 Blade Server Chassis

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

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

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

Cisco UCS Mini Infrastructure

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

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

Figure 3: Cisco UCS Mini

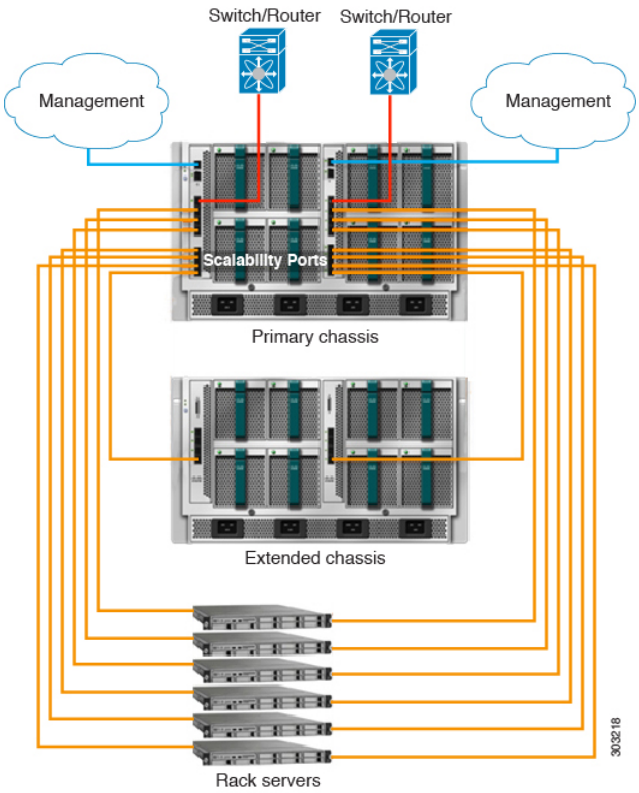

In the Cisco UCS Mini solution, the Cisco UCS 6324 fabric interconnect is collapsed into the IO Module form factor, and is inserted into the IOM slot of the blade server chassis. The Cisco UCS 6324 fabric interconnect has 24 10G ports available on it. Sixteen of these ports are server facing, two 10G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

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

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

Cisco UCS Infrastructure Virtualization

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

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

Figure 4: Virtualization Capabilities of Cisco UCS

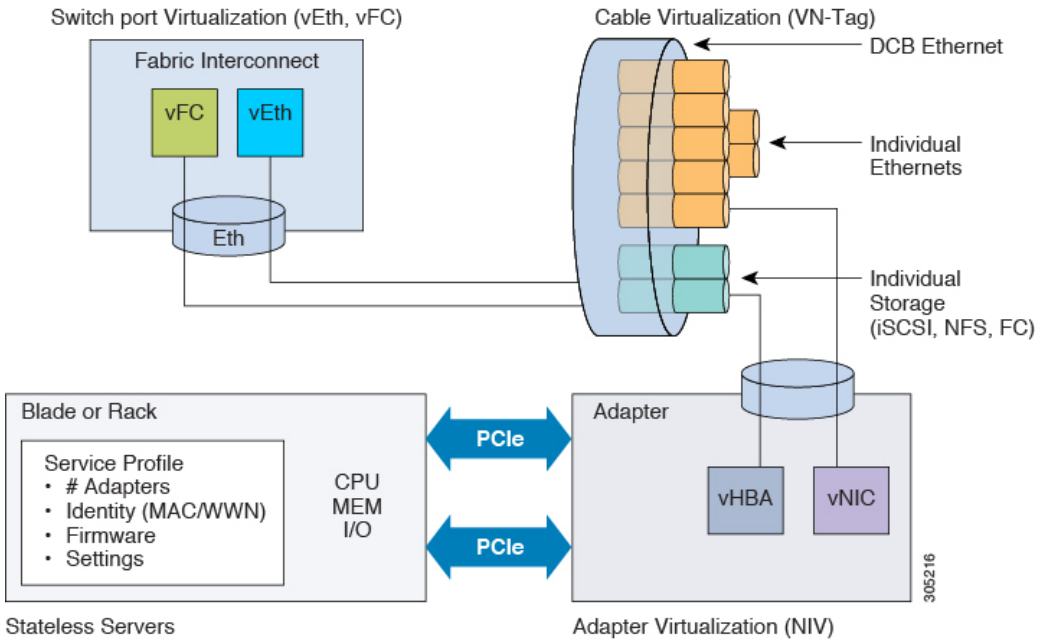

Switch Port Virtualization

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

Cable Virtualization

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

Adapter Virtualization

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

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

Server Virtualization

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

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

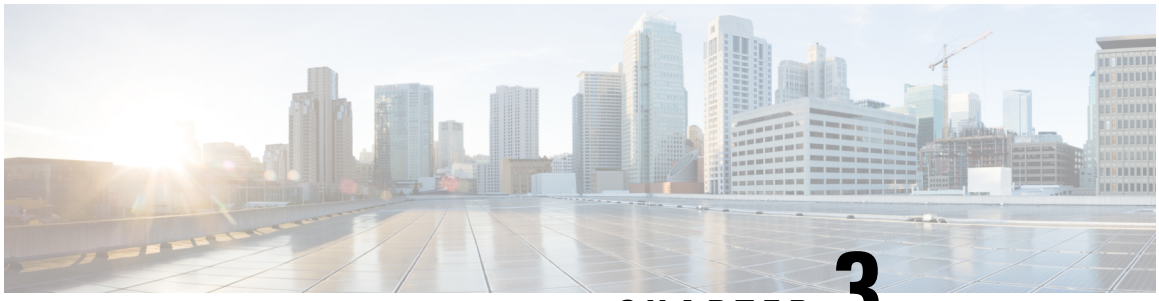

CHAPTER 3

Equipment Policies

- Chassis/FEX Discovery Policy, on page 13
- Chassis Connectivity Policy, on page 18
- Rack Server Discovery Policy, on page 19
- Aging Time for the MAC Address Table, on page 20

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.

In a Cisco UCS Mini setup, chassis discovery policy is supported only on the extended chassis.

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 of a chassis, if chassis/FEX discovery policy changes are done, acknowledge IO Modules rather than the entire Chassis to avoid disruption. The discovery policy changes can include increasing the number of links between Fabric Interconnect and IO Module, or changes to the Link Grouping preference.

Make sure that you monitor for faults before and after the IO Module acknowledgement to ensure that the connectivity is restored before proceeding to the other IO Module for the chassis.

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 2: 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 **None**, links from the IOM are pinned to the fabric interconnect.

After a fabric port channel is created through Cisco UCS Manager, 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.

Multicast Hardware Hash

In a port channel, by default, ingress multicast traffic on any port in the fabric interconnect (FI) selects a particular link between the IOM and the fabric interconnect to egress the traffic. To reduce potential issues with the bandwidth, and to provide effective load balancing of the ingress multicast traffic, hardware hashing is used for multicast traffic. When multicast hardware hashing is enabled, all links between the IOM and the fabric interconnect in a port channel can be used for multicast traffic.

Pinning

Pinning in Cisco UCS is only relevant to uplink ports. If you configure **Link Grouping Preference** as **None** during chassis discovery, the IOM forwards traffic from a specific server to the fabric interconnect through its uplink ports by using static route pinning.

The following table showcases how pinning is done between an IOM and the fabric interconnect based on the number of active fabric links between the IOM and the fabric interconnect.

Table 3: Pinning on an IOM

Number of Active Fabric Links	Server slot pinned to fabric link
1-Link	All the HIF ports are pinned to the active link
2-Link	1,3,5,7 to link-1 2,4,6,8 to link-2
4-Link	1,5 to link-1 2,6 to link-2 3,7 to link-3 4,8 to link-4

Number of Active Fabric Links	Server slot pinned to fabric link
8-Link (Applies only to 2208XP)	1 to link-1 2 to link-2 3 to link-3 4 to link-4 5 to link-5 6 to link-6 7 to link-7 8 to link-8

Only 1,2,4 and 8 links are supported. 3,5,6, and 7 links are not valid configurations.

Port-Channeling

While pinning traffic from a specific server to an uplink port provides you with greater control over the unified fabric and ensures optimal utilization of uplink port bandwidth, it could also mean excessive traffic over certain circuits. This issue can be overcome by using port channeling. Port channeling groups all links between the IOM and the fabric interconnect into one port channel. The port channel uses a load balancing algorithm to decide the link over which to send traffic. This results in optimal traffic management.

Cisco UCS supports port-channeling only through the Link Aggregation Control Protocol (LACP). 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 into 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 this parameter is set to **None**, links from the IOM to the fabric interconnect are not grouped in a fabric port channel.

Once a fabric port channel is created, links can be added or removed by changing the link group preference and reacknowledging the chassis, or by enabling or disabling the chassis from the 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.

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.

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

The chassis connectivity policy is created by Cisco UCS Manager only when the hardware configuration supports fabric port channels.

In a Cisco UCS Mini setup, the creation of a chassis connectivity policy is supported only on the extended chassis.

Configuring a Chassis Connectivity Policy

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

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**.

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**.

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**.

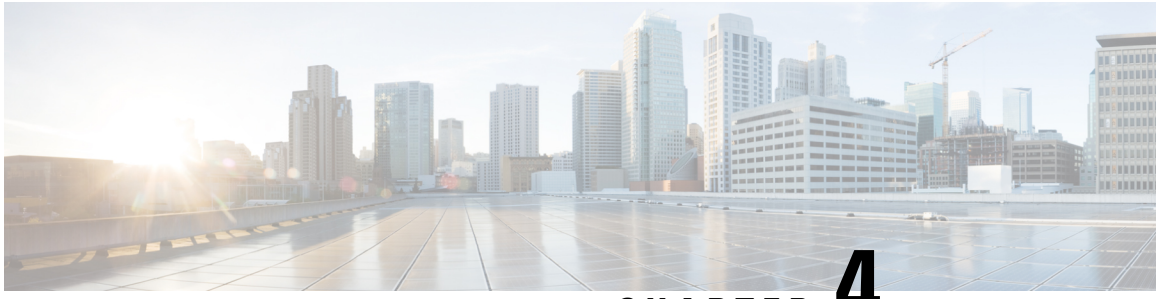

CHAPTER 4

Chassis Management

- Chassis Management in Cisco UCS Manager GUI , on page 21
- Guidelines for Removing and Decommissioning Chassis, on page 23
- Acknowledging a Chassis, on page 23
- Decommissioning a Chassis, on page 24
- Removing a Chassis, on page 24
- Recommissioning a Single Chassis, on page 25
- Recommissioning Multiple Chassis, on page 25
- Renumbering a Chassis, on page 26
- Turning on the Locator LED for a Chassis, on page 27
- Turning off the Locator LED for a Chassis, on page 27
- Creating a Zoning Policy from Inventory, on page 28
- Viewing the POST Results for a Chassis, on page 28

Chassis Management in Cisco UCS Manager GUI

You can manage and monitor all chassis in a Cisco UCS domain through Cisco UCS Manager GUI.

The Cisco UCS S3260 Chassis

Cisco UCS Manager Release 3.1(2) introduces support for the Cisco UCS S3260 chassis on Cisco UCS 6300 Series, and 6200 Series fabric interconnect setups.

The Cisco UCS S3260 chassis is a 4U chassis that is designed to operate in a standalone environment and also as part of the Cisco Unified Computing System. It has the following main components:

- Four 1050 Watt AC modular power supplies (2 + 2 shared and redundant mode of operation)
- Two System IO Controller (SIOC) slots
- Two storage server slots out of which one can be used for storage expansion

Note The second server slot in the chassis can be utilized by an HDD expansion tray module for an additional four 3.5” drives.

- 56 3.5” drive bays with an optional 4 x 3.5” HDD expansion tray module instead of the second server
- Up to 360 TB storage capacity by using 6 TB HDDs
- Serial Attached SCSI (SAS) expanders that can be configured to assign the 3.5” drives to individual server modules
- The two servers in the chassis can be replaced by a single, dual-height server with an IO expander

Cisco UCS 5108 Blade Server Chassis

The Cisco UCS 5100 Series Blade Server Chassis is logically part of the fabric interconnects, thus creating a single, coherent management domain and decreasing management complexity. In the management domain, server management is handled by the fabric interconnect, while I/O and network management is extended to every chassis and blade server. Basing the I/O infrastructure on a unified fabric allows the Cisco Unified Computing System to have a simple and streamlined chassis yet offer a comprehensive set of I/O options. This results in the chassis having only five basic components:

- The physical chassis with passive midplane and active environmental monitoring circuitry
- Four power-supply bays with power entry in the rear, and redundant-capable, hot-swappable power supply units accessible from the front panel
- Eight hot-swappable fan trays, each with two fans
- Two fabric extender slots accessible from the back panel
- Eight blade server slots accessible from the front panel

The blade server chassis has flexible partitioning with removable dividers to handle two blade server form factors:

- Half-width blade servers have access to power and two 10GBASE-KR connections, one to each fabric extender slot.
- Full-width blade servers connect to power and two connections to each fabric extender.

Extended Chassis for UCS Mini

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

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

To use a extended chassis, do the following:

- Connect the second Cisco UCS 5108 chassis to the existing single-chassis Cisco UCS 6324 Series fabric interconnect configuration through the scalability port.

- Configure the chassis discovery policy.
- Configure the server ports and wait for the second chassis to be discovered.

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

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.

Note Chassis acknowledgement causes complete loss of network and storage connectivity to the chassis.

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.

Note This procedure is not applicable for Cisco UCSC S3260 Chassis.

Procedure

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

Step 2 Expand **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 This procedure is not applicable for Cisco UCSC S3260 Chassis.

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 Expand **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.

Note This procedure is not applicable for Cisco UCSC S3260 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 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**.

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.

Creating a Zoning Policy from Inventory

You can create a disk zoning policy from the existing inventory and disk ownership.

Note Creating a disk zoning policy from the existing inventory is supported only on Cisco UCS S3260 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 create a zoning policy.

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

Step 5 In the **Actions** area, click **Create Zoning Policy from Inventory**.

Step 6 In the **Create Zoning Policy from Inventory** dialog box that appears, do the following:

a) Enter the **Disk Zoning Policy Name**.
b) Select the organization where you want to create the policy.
c) Click **OK**

Step 7 In the confirmation dialog box that appears, click **OK**.

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.

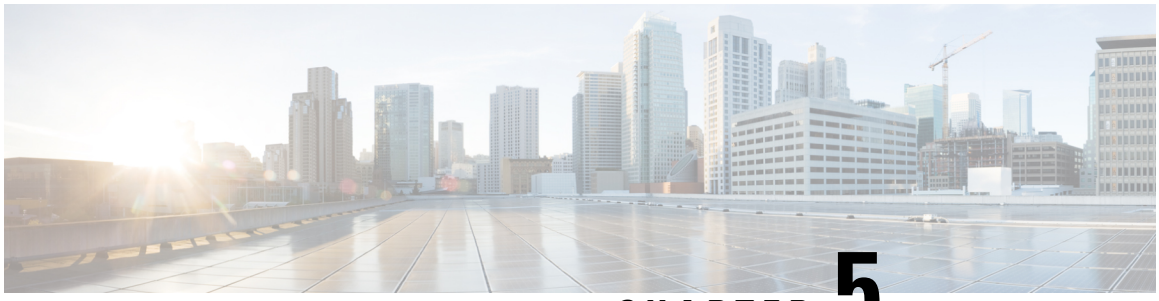

CHAPTER 5

I/O Module Management

- I/O Module Management in Cisco UCS Manager GUI , on page 29
- Acknowledging an IO Module, on page 29
- Resetting an I/O Module, on page 30
- Resetting an I/O Module from a Peer I/O Module, on page 30
- Viewing Health Events for an I/O Module, on page 31
- Viewing the POST Results for an I/O Module, on page 32

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.

Cisco UCS Manager Release 3.1(1) introduces the Cisco UCS-IOM-2304 I/O module with 40 GbE connectivity to the Cisco UCS 6300 Series Fabric Interconnect. The *Cisco UCS Manager Getting Started Guide* provides more information about this functionality.

Acknowledging an IO Module

Cisco UCS Manager Release 2.2(4) introduces the ability to acknowledge a specific IO module in a chassis.

Note

- After adding or removing physical links between Fabric Interconnect and IO Module, an acknowledgement of the IO Module is required to properly configure the connection.
- The ability to re-acknowledge each IO Module individually allows to rebuild the network connectivity between a single IO Module and its parent Fabric Interconnect without disrupting production traffic in the other Fabric Interconnect.

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 acknowledge.
Step 4 In the **Work** pane, click the **General** tab.
Step 5 In the **Actions** area, click **Acknowledge IO Module**.
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**.

Resetting an I/O Module from a Peer I/O Module

Sometimes, I/O module upgrades can result in failures or I/O modules can become unreachable from Cisco UCS Manager due to memory leaks. You can reboot an I/O module that is unreachable through its peer I/O module.

Resetting the I/O module restores the I/O module to factory default settings, deletes all cache files and temporary files, but retains the size-limited OBFL file.

Procedure

Step 1 In the **Navigation** pane, click **Equipment**.
Step 2 Expand **Equipment** > **Chassis** > ***Chassis Number*** > **IO Modules**.
Step 3 Choose the peer I/O module of 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 Peer IO Module**.

Viewing Health Events 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 health events.

Step 4 In the **Work** pane, click the **Health** tab

The health events triggered for this I/O module appear. The fields in this tab are:

Name	Description
Health Summary area	
Health Qualifier field	Comma-separated names of all the heath events that are triggered for the component.
Health Severity field	Highest severity of all the health events that are triggered for the component. This can be one of the following: • **critical** • **major** • **minor** • **warning** • **info** • **cleared** **Note** The severity levels listed here are from highest to lowest severity.
Health Details area	

Name	Description
Severity column	Severity of the health event. This can be one of the following: • **critical** • **major** • **minor** • **warning** • **info** • **cleared** **Note** The severity levels listed here are from highest to lowest severity.
Name column	Name of the health event.
Description column	Detailed description of the health event.
Value column	Current value of the health event.
Details area	The **Details** area displays the **Name**, **Description**, **Severity**, and **Value** details of any health event that you select in the **Health Details** area.

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.

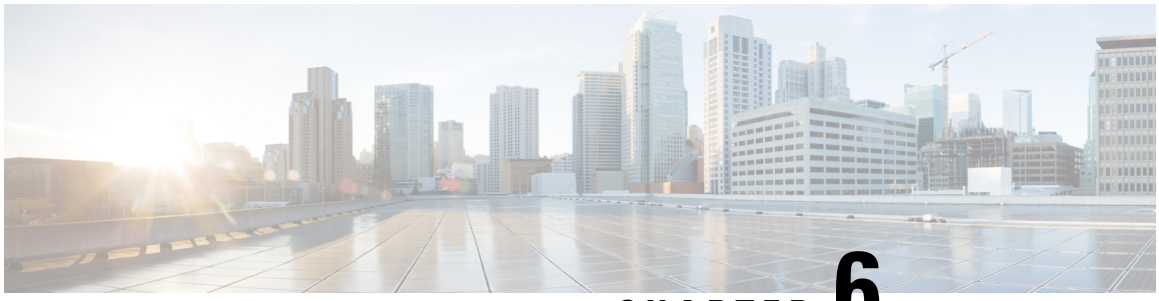

CHAPTER 6

SIOC Management

- SIOC Management in Cisco UCS Manager , on page 33
- Acknowledging an SIOC, on page 34
- Resetting the CMC, on page 35
- CMC Secure Boot, on page 35

SIOC Management in Cisco UCS Manager

You can manage and monitor all System Input/Output Controllers (SIOC) in a Cisco UCS domain through Cisco UCS Manager.

SIOC Removal or Replacement

You can remove or replace an SIOC from a chassis. Removal or replacement of an SIOC is a service-affecting operation, which requires you to power down the entire chassis.

Guidelines for SIOC Removal

- To remove the active SIOC, or both SIOCs, shut down and remove power from the entire chassis. You must disconnect all power cords to completely remove power.
- Removal of SIOCs from a chassis results in the entire chassis being disconnected from Cisco UCS Manager.

SIOC Removal

Do the following to remove an SIOC from the system:

1. Shut down and remove power from the entire chassis. You must disconnect all power cords to completely remove power.
2. Disconnect the cables connecting the SIOC to the system.
3. Remove the SIOC from the system.

SIOC Replacement

Do the following to remove an SIOC from the system and replace it with another SIOC:

1. Shut down and remove power from the entire chassis. You must disconnect all power cords to completely remove power.
2. Disconnect the cables connecting the SIOC to the system.
3. Remove the SIOC from the system.
4. Connect the new SIOC to the system.
5. Connect the cables to the SIOC.
6. Connect power cords and then power on the system.
7. Acknowledge the new SIOC.

   The server connected to the replaced SIOC is rediscovered.

Note If the firmware of the replaced SIOC is not the same version as the peer SIOC, then it is recommended to update the firmware of the replaced SIOC by re-triggering chassis profile association.

Acknowledging an SIOC

Cisco UCS Manager has the ability to acknowledge a specific SIOC in a chassis. Perform the following procedure when you replace an SIOC in a chassis.

Caution This operation rebuilds the network connectivity between the SIOC and the fabric interconnects to which it is connected. The server corresponding to this SIOC becomes unreachable, and traffic is disrupted.

Procedure

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

Step 2 Expand **Equipment** > **Chassis** > ***Chassis Number*** > **SIOC**

Step 3 Choose the SIOC that you want to acknowledge.

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

Step 5 In the **Actions** area, click **Acknowledge SIOC**.

Step 6 In the **Acknowledge SIOC** confirmation box, click **Yes**.

Resetting the CMC

Procedure

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

Step 2 Expand **Equipment** > **Chassis** > ***Chassis Number*** > **SIOC** > ***SIOC Number***

Step 3 In the **Work** pane, click the **Chassis Management Controller** tab.

Step 4 In the **Actions** area, click **Reset CMC**.

Step 5 In the **Reset CMC** confirmation box, click **Yes**.

CMC Secure Boot

With Chassis Management Controller (CMC) secure boot, only Cisco-signed firmware images can be installed and run on the CMC. When the CMC 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 CMC firmware.

Guidelines and Limitations for CMC Secure Boot

- CMC secure boot is supported only on the Cisco UCS S3260 chassis.
- When chassis association is in progress, enabling secure boot on one of the SIOCs will result in a failed operation.
- After CMC secure boot is enabled, it cannot be disabled.
- CMC secure boot is specific to the SIOC on which it is enabled. If you replace the SIOC on which CMC secure boot is enabled, the **Secure boot operational state** field will now display the secure boot status of the new SIOC.
- After CMC secure boot is enabled on a chassis, you cannot move the chassis back to standalone mode and downgrade the firmware to a CMC firmware image earlier than Cisco IMC Release 2.0(13).
- The **Secure boot operational state** field shows the secure boot status. This can be one of the following:
  - **Disabled**—When CMC secure boot is not enabled. This is the default state.
  - **Enabling**—When CMC secure boot is being enabled.
  - **Enabled**—When CMC secure boot is enabled.

Enabling CMC Secure Boot

Procedure

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

Step 2 Expand **Equipment** > **Chassis** > ***Chassis Number*** > **SIOC**

Step 3 Choose the SIOC on which you want to enable CMC secure boot.

Step 4 In the **Work** pane, click the **Chassis Management Controller** tab.

Step 5 In the **Actions** area, click **Enable Secure Boot**.

The **Enable Secure Boot** confirmation box appears with the following warning:

When committed, CMC secure boot and installation will be enabled. This is an irreversible operation. Are you sure you want to enable secure boot.

Step 6 Click **Yes**.

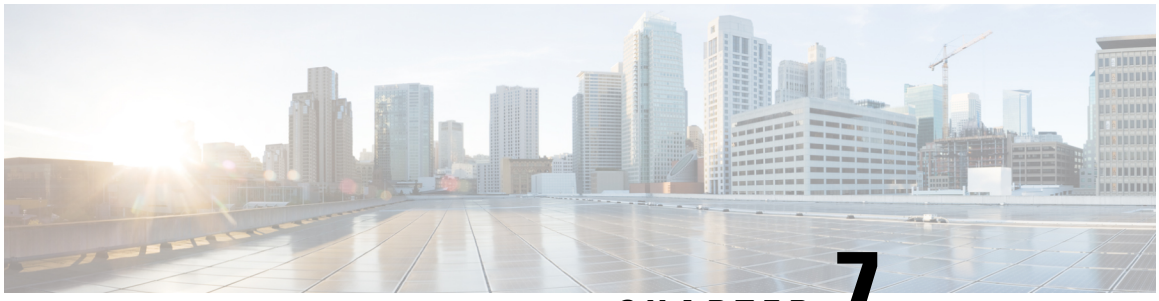

CHAPTER 7

Power Management in Cisco UCS

- Power Capping in Cisco UCS, on page 38
- Power Policy Configuration, on page 38
- Power Policy for Cisco UCS Servers, on page 38
- Configuring the Power Policy, on page 39
- Power Supply for Redundancy Method, on page 39
- Configuring Policy Driven Chassis Group Power Capping, on page 40
- Policy Driven Chassis Group Power Capping, on page 40
- Power Control Policy, on page 40
- Power Groups in UCS Manager, on page 43
- Blade Level Power Capping, on page 47
- Manual Blade Level Power Cap, on page 47
- Setting the Blade-Level Power Cap for a Server, on page 48
- Viewing the Blade-Level Power Cap, on page 49
- Global Power Profiling Policy Configuration, on page 49
- Global Power Profiling Policy, on page 49
- Configuring the Global Power Profile Policy, on page 50
- Global Power Allocation Policy Configuration, on page 50
- Global Power Allocation Policy, on page 50
- Configuring the Global Power Allocation Policy, on page 51
- Power Management During Power-on Operations, on page 51
- Power Sync Policy Configuration, on page 52
- Power Sync Policy, on page 52
- Power Synchronization Behavior, on page 52
- Creating a Power Sync Policy, on page 53
- Changing a Power Sync Policy, on page 54
- Deleting a Power Sync Policy, on page 55
- Rack Server Power Management, on page 55
- UCS Mini Power Management , on page 56

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 blade servers, UCS C220 and C240 M4/M5, and C480 M5 rack servers, UCS Mini, and mixed UCS domains.

Cisco UCS Manager supports power capping on the following:

- UCS 6200 Series Fabric Interconnects
- UCS 6300 Series Fabric Interconnects
- UCS 6324 Series Fabric Interconnects (Cisco UCS Mini)
- UCS 6400 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.
Power Save Policy	Globally manages the chassis to maximize energy efficiency or availability.
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.

Power Policy Configuration

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**.

Power Supply for Redundancy Method

PSU Redundancy	Max Power @ 220 V (Watts)	Max Power @ 110 V (Watts)
1+1 (N+1) OR 1 (N)	2500	1300
2+1 (N+1) OR 2 (N) or 2+2 (Grid)	5000	2600
3+1 (N+1) OR 3 (N)	5472	3900

PSU Redundancy	Max Power @ 220 V (Watts)	Max Power @ 110 V (Watts)
4 (N)	5472	5200

Configuring Policy Driven Chassis Group 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 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.

Starting with Cisco UCS Manager 3.2(2), chassis dynamic power rebalance mechanism is enabled by default. The mechanism continuously monitors the power usage of the blade servers and adjusts the power allocation accordingly. Chassis dynamic power rebalance mechanism operates within the overall chassis power budget set by Cisco UCS Manager, which is calculated from the available PSU power and Group power.

For mission-critical application a special priority called **no-cap** is also available. Setting the priority to **no-cap** does not guarantee that a blade server gets maximum power all the time, however, it prioritizes the blade server over other servers during the chassis dynamic power rebalance budget allocations.

Note If all the blade servers are set with no-cap priority and all of them run high power consuming loads, then there is a chance that some of the blade servers get capped under high power usage, based on the power distribution done through dynamic balance.

Global Power Control Policy options are inherited by all the chassis managed by the Cisco UCS Manager.

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 multi tenancy, 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).

Name	Description
Fan Speed Policy drop-down	Fan speed is for rack servers only. This 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 minimum 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. **Note** The **performance** option is not supported on Cisco UCS C-Series M5 and M6 servers. • **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.
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 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.

Name	Description
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**.

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.

Error Message	Cause	Recommended Action
`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 **Add Rack Members** page of the **Create Power Group** wizard, do the following:

a) In the **Rack Unit** table, choose one or more rack units to include in the power group.

b) Click the >> button to add the rack to the **Selected Rack Unit** table that displays all racks included in the power group.

You can use the << button to remove one or more rack units from the power group.

c) Click **Next**.

Step 9 On the **Add FEX Members** page of the **Create Power Group** wizard, do the following:

a) In the **FEX** table, choose one or more FEX to include in the power group.

b) Click the >> button to add the chassis to the **Selected FEX** table that displays all FEX included in the power group.

You can use the << button to remove one or more FEX from the power group.

c) Click **Next**.

Step 10 On the **Add FI Members** page of the **Create Power Group** wizard, do the following:

a) In the **FI** table, choose one or more FI to include in the power group.

b) Click the >> button to add the FI to the **Selected FI** table that displays all chassis included in the power group.

You can use the << button to remove one or more FI from the power group.

c) Click **Next**.

Step 11 On the **Power Group Attributes** page of the **Create Power Group** wizard, do the following:

a) Complete the following fields:

Name	Description
Input Power(W) field	The maximum peak power (in watts) available to the power group. Enter an integer between 0 and 10000000.
Recommended value for Input Power field	The recommended range of input power values for all the members of the power group.

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**.

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 1300 watts.

  Note B480 M5 systems using 256GB DIMMs must have a manual blade level cap at 1300W.

- **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** Manual blade level power capping will limit the power consumption of a single system, regardless of available power in the chassis.
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.

Global Power Profiling Policy Configuration

Global Power Profiling Policy

The Global Power Profiling Policy specifies how power allocation is applied to all of the servers in a chassis. The policy applies when you set the Global Power Allocation Policy to **Policy Driven Chassis Group Cap**. You can set the Global Power Profiling Policy to one of the following:

- **Disabled**—The minimum and maximum power cap values of the blades are calculated based on the static power consumption values of each of the components.
- **Enabled**—The minimum and maximum power cap values of the blades are measured as part of the server discovery. These values are similar to the actual power consumption of the blades.

Note After enabling the Global Power Profiling Policy, you must re-acknowledge the blades to obtain the minimum and maximum power cap.

Configuring the Global Power Profile 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 Profiling Policy** area, check the **Profile Power** checkbox to enable the Global Power Profiling Policy.

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**.

Power Management During Power-on Operations

Boot Staggering during Power on

Cisco UCS Manager attempts to boot as many blades as possible based on the amount of available power. If the power required to boot a blade is not available, Cisco UCS Manager staggers the boot in the Finite State Machine (FSM) CheckPowerAvailability stage, and raises the following fault on the blade: Insufficient power available to power-on server x/y.

When the required power becomes available, the FSM proceeds with blade power on. After a blade powers off, the allocated power budget is reclaimed.

Note When the power budget that was allocated to the blade is reclaimed, the allocated power displays as 0 Watts.

Limitation

If you power on a blade outside of the Cisco UCS Manager and if there is not enough power available for allocation, the following fault is raised:

```
Power cap application failed for server x/y
```

Power Allocation during Service Profile Association

The power allocated to a blade during service profile association depends on the Power Control Policy used, and the power that is available from the power group. After the power is allocated to a server during a successful service profile association, the blade is guaranteed the minimum power cap. If the Power Control Policy priority is set to no-cap, a blade is allocated a potential maximum power cap, which might exceed the measured maximum power cap that displays.

Note If the priority of an associated blade is changed to no-cap, and is not able to allocate the maximum power cap, you might see one of the following faults:

- `PSU-insufficient`—There is not enough available power for the PSU.
- `Group-cap-insufficient`—The group cap value is not sufficient for the blade.

Power Sync Policy Configuration

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 power state of the service profile differs from the actual power state of the server. The policy allows you to control when to synchronize the power state on the associated service profiles for the servers. The power sync policy does not affect other power-related policies.

The power synchronization 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 preferred 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	Preferred Power State	Actual Power State Before Event	Actual Power State After Event
Shallow Association	ON	OFF	ON
Shallow Association	OFF	OFF	OFF
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 multi tenancy, 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).

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 synchronizes 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 multi tenancy, 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**.

Rack Server Power Management

Power capping is supported for following rack servers:

- Cisco UCS C220 M4 Server
- Cisco UCS C240 M4 Server
- Cisco UCS C220 M5 Server
- Cisco UCS C240 M5 Server
- Cisco UCS C480 M5 Server

UCS Mini Power Management

You can manage power of the blade servers in the Cisco UCS 6324 Fabric Interconnect (FI), which is used for remote offices and branch sites, and for limited server deployments. UCS Manager supports Dual Line Power Supply Unit and 110V when used with the Cisco UCS 6324 Fabric Interconnect. You can manage how you want to allocate power when using 110V power supplies, because they might not provide enough power for a fully loaded chassis. Dual power supplies is standard for both AC and DC-48V on the Cisco UCS Mini 6324.

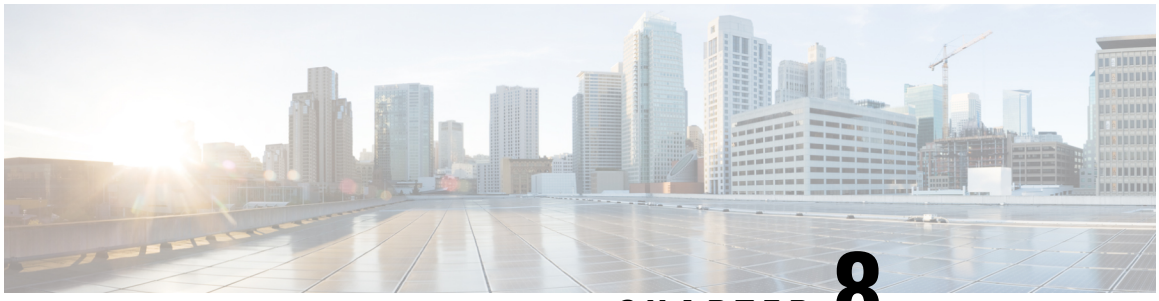

CHAPTER 8

Blade Server Hardware Management

- Blade Server Management, on page 57
- Booting a Blade Server, on page 59
- Booting a Rack-Mount Server from the Service Profile , on page 60
- Determining the Boot Order of a Blade Server, on page 60
- Shutting Down a Blade Server, on page 61
- Shutting Down a Server from the Service Profile , on page 61
- Resetting a Blade Server, on page 62
- Resetting a Blade Server to Factory Default Settings, on page 62
- Reacknowledging a Blade Server, on page 63
- Removing a Server from a Chassis, on page 64
- Deleting the Inband Configuration from a Blade Server, on page 64
- Decommissioning a Blade Server, on page 65
- Removing a Non-Existent Blade Server Entry, on page 65
- Recommissioning a Blade Server, on page 66
- Reacknowledging a Server Slot in a Chassis, on page 66
- Removing a Non-Existent Blade Server from the Configuration Database, on page 67
- Turning the Locator LED for a Blade Server On and Off, on page 67
- Turning the Local Disk Locator LED on a Blade Server On and Off, on page 68
- Resetting the CMOS for a Blade Server, on page 68
- Resetting the CIMC for a Blade Server, on page 69
- Clearing TPM for a Blade Server, on page 69
- Viewing the POST Results for a Blade Server, on page 70
- Issuing an NMI from a Blade Server, on page 70
- Viewing Health Events for a Blade Server, on page 71
- Health LED Alarms, on page 72
- Smart SSD, on page 73

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.

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
Down	Powered On	Powered On **Note** Running servers are not shut down regardless of the desired power state in the service profile.

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 Rack-Mount 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 multi tenancy, 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 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.

Note When a blade server that is associated with a service profile is shut down, the VIF down alerts F0283 and F0479 are automatically suppressed.

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 multi tenancy, 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** box 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 66.

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 **Equipment**.

Step 2 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 66.

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 Expand **Equipment** node.

Step 3 Click the **Chassis** node.

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

Step 5 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 6 If a confirmation dialog box displays, click **Yes**.

Step 7 (Optional) Monitor the progress of the server recommission and discovery on the **FSM** tab for the server.

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 64.

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.

Turning the Local Disk Locator LED on a Blade Server On and Off

Before you begin

- Ensure the server, on which the disk is located, is powered on. If the server is off, you are unable to turn on or off the local disk locator LED.

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 local disk locator LED on or off.

Step 4 In the **Work** pane, click the **Inventory** > **Storage** > **Disks** tabs.

The Storage Controller inventory appears.

Step 5 Click a disk.
The disk details appear.

Step 6 In the **Details** area, click **Toggle Locator LED**.
If the **Locator LED** state is **On**, it will turn **Off**. If the **Locator LED** state is **Off**, it will turn **On**.

Step 7 Click **Save Changes.**

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 the management controller of the blade 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 and higher 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**.

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.

Viewing Health Events for 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 for which you want to view health events.

Step 4 In the **Work** pane, click the **Health** tab

The health events triggered for this server appear. The fields in this tab are:

Name	Description
Health Summary area	
Health Qualifier field	Comma-separated names of all the heath events that are triggered for the component.
Health Severity field	Highest severity of all the health events that are triggered for the component. This can be one of the following: • **critical** • **major** • **minor** • **warning** • **info** • **cleared** **Note** The severity levels listed here are from highest to lowest severity.
Health Details area	

Name	Description
Severity column	Severity of the health event. This can be one of the following: • **critical** • **major** • **minor** • **warning** • **info** • **cleared** **Note** The severity levels listed here are from highest to lowest severity.
Name column	Name of the health event.
Description column	Detailed description of the health event.
Value column	Current value of the health event.
Details area	The **Details** area displays the **Name**, **Description**, **Severity**, and **Value** details of any health event that you select in the **Health Details** area.

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.

Name	Description
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.

Smart SSD

Beginning with release 3.1(3), Cisco UCS Manager supports monitoring SSD health. This feature is called Smart SSD. It provides statistical information about the properties like wear status in days, percentage life remaining, and so on. For every property, a minimum, a maximum and an average value is recorded and displayed. The feature also allows you to provide threshold limit for the properties.

Note The Smart SSD feature is supported only for a selected range of SSDs. It is not supported for any HDDs.

The SATA range of supported SSDs are:

- Intel
- Samsung
- Micron

The SAS range of supported SSDs are:

- Toshiba
- Sandisk
- Samsung
- Micron

Note

- Power Cycle Count is not available on SAS SSDs.
- Smart SSD feature is supported only on M4 servers and later.

Monitoring SSD Health

Procedure

Step 1 Navigate to **Equipment** > **Rack-Mounts** > **Servers** > **Server Number** > **Inventory** > **Storage**.

Step 2 Click the controller component for which you want to view the SSD health.

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

Step 4 Click the SSD for which you want to view the health properties.

You can view the values for

- **PercentageLifeLeft:** Displays the duration of life so action can be taken when required.
- **PowerCycleCount:** Displays the number of times the SSD is power cycled across the server reboot.
- **PowerOnHours:** Displays the duration for which the SSD is on. You can replace or turn the SSD off based on the requirement.

  Note If there is a change in any other property, updated **PowerOnHours** is displayed.

- **WearStatusInDays:** Provides guidance about the SSD wear based on the workload characteristics run at that time.

  Note These values are updated on an hourly basis.

  You can specify the threshold limit for the values and faults are raised when the value reaches or exceeds the threshold limit. Smart SSD feature tracks temperature and raises a fault as the temperature crosses the threshold limit (90°C) and moves the disk to the degraded state notifying the reason for degradation.

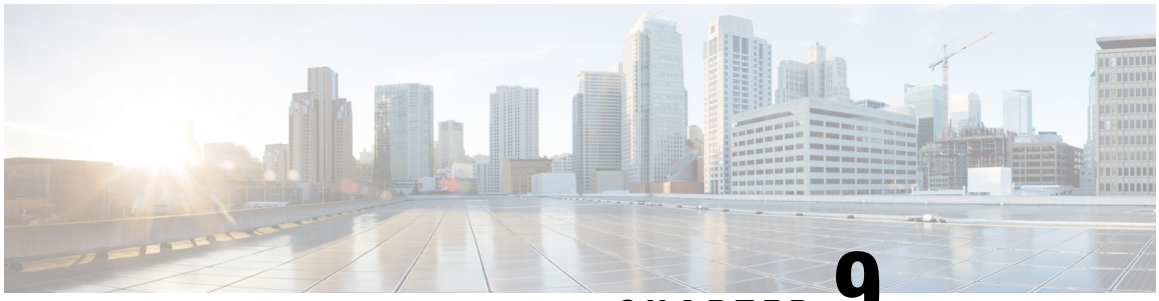

CHAPTER 9

Rack-Mount Server Hardware Management

- Rack-Mount Server Management, on page 75
- Guidelines for Removing and Decommissioning Rack-Mount Servers, on page 76
- Recommendations for Avoiding Unexpected Server Power Changes, on page 77
- Booting a Rack-Mount Server, on page 77
- Booting a Rack-Mount Server from the Service Profile , on page 78
- Determining the Boot Order of a Rack-Mount Server, on page 78
- Shutting Down a Rack-Mount Server, on page 79
- Shutting Down a Server from the Service Profile , on page 79
- Resetting a Rack-Mount Server, on page 80
- Resetting a Rack-Mount Server to Factory Default Settings, on page 81
- Persistent Memory Scrub, on page 81
- Reacknowledging a Rack-Mount Server, on page 82
- Deleting the Inband Configuration from a Rack-Mount Server, on page 83
- Decommissioning a Rack-Mount Server, on page 83
- Recommissioning a Rack-Mount Server, on page 84
- Renumbering a Rack-Mount Server, on page 84
- Removing a Non-Existent Rack-Mount Server from the Configuration Database, on page 85
- Turning the Locator LED for a Rack-Mount Server On and Off, on page 85
- Turning the Local Disk Locator LED on a Rack-Mount Server On and Off, on page 86
- Resetting the CMOS for a Rack-Mount Server, on page 86
- Resetting the CIMC for a Rack-Mount Server, on page 87
- Clearing TPM for a Rack-Mount Server, on page 87
- Issuing an NMI from a Rack-Mount Server, on page 88
- Viewing Health Events for a Rack-Mount Server, on page 88
- Viewing the POST Results for a Rack-Mount Server, on page 90
- Viewing the Power Transition Log, on page 90

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 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 Rack-Mount 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 multi tenancy, 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 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 multi tenancy, 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.

Resetting a Rack-Mount Server to Factory Default Settings

You can now reset a rack-mount server to its factory settings. By default, the factory reset operation does not affect storage, including 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** > **Rack Mounts** > **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, click **Reset to Factory Default**, then 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** checkbox 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.
- To delete all Persistent Memory storage, check the **Persistent Memory Scrub** checkbox.

Cisco UCS Manager resets the server to its factory default settings.

Persistent Memory Scrub

Persistent memory scrub allows you to remove the persistent memory configuration and data from the persistent memory modules on a server.

In Cisco IMC, you can scrub persistent memory by resetting the persistent memory modules to factory defaults.

In Cisco UCS Manager, you can scrub persistent memory by using one of the following methods:

- Disassociating the service profile and the scrub policy, which has the persistent memory scrub option set to yes
- Performing a **Reset to Factory Default** operation on the server with the persistent memory scrub option set to yes
- Deleting a goal

Excluding namespace and region configurations, persistent memory SKUs with encryption capabilities will have encryption keys changed internally by the device if a persistent memory policy is applied in Cisco UCS Manager. For SKUs without encryption capabilities or host-managed persistent memory modules, persistent memory data and configuration will be retained.

Persistent memory configuration is reset to factory default settings as described here:

- For B-Series and C-Series servers, 100% Memory Mode is applied. For S-Series servers, 0% Memory Mode and App Direct Non Interleaved type are applied.
- For UCS-Managed persistent memory modules, security will be disabled and frozen. For host-managed persistent memory modules, existing configuration will be retained.

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-Mount 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 **Servers**.

Step 2 Expand **Equipment** > **Rack Mounts** > **Servers** > ***Server Number***.

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 Under **Equipment**, 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**

Turning the Local Disk Locator LED on a Rack-Mount Server On and Off

Before you begin

- Ensure the server, on which the disk is located, is powered on. If the server is off, you are unable to turn on or off the local disk locator LED.

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 local disk locator LED on or off.

Step 4 In the **Work** pane, click the **Inventory** > **Storage** > **Disks** tabs.

The Storage Controller inventory appears.

Step 5 Click a disk.
The disk details appear.

Step 6 In the **Details** area, click **Toggle Locator LED**.
If the **Locator LED** state is **On**, it will turn **Off**. If the **Locator LED** state is **Off**, it will turn **On**.

Step 7 Click **Save Changes.**

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 the management controller of the blade 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 and higher 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**.

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.

Viewing Health Events for 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 for which you want to view health events.

Step 4 In the **Work** pane, click the **Health** tab

The health events triggered for this server appear. The fields in this tab are:

Name	Description
Health Summary area	
Health Qualifier field	Comma-separated names of all the heath events that are triggered for the component.

Name	Description
Health Severity field	Highest severity of all the health events that are triggered for the component. This can be one of the following: • **critical** • **major** • **minor** • **warning** • **info** • **cleared** **Note** The severity levels listed here are from highest to lowest severity.
Health Details area	
Severity column	Severity of the health event. This can be one of the following: • **critical** • **major** • **minor** • **warning** • **info** • **cleared** **Note** The severity levels listed here are from highest to lowest severity.
Name column	Name of the health event.
Description column	Detailed description of the health event.
Value column	Current value of the health event.
Details area	The **Details** area displays the **Name**, **Description**, **Severity**, and **Value** details of any health event that you select in the **Health Details** area.

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.

Viewing the Power Transition Log

You can view the **Power Transition Log**, which displays the last five server power transitions. The information provided includes the **Power Change Source** and the **Timestamp**.

Only unique power transition events are displayed. In case of a UCSM initiated power transition, the FSM causing the power transition is displayed.

Procedure

Step 1 Navigate to **Equipment** > **Rack-Mounts** > **Servers**

Step 2 Choose the server for which you want to view the power transition log.

The **Power Transition Log** is under the **General** tab.

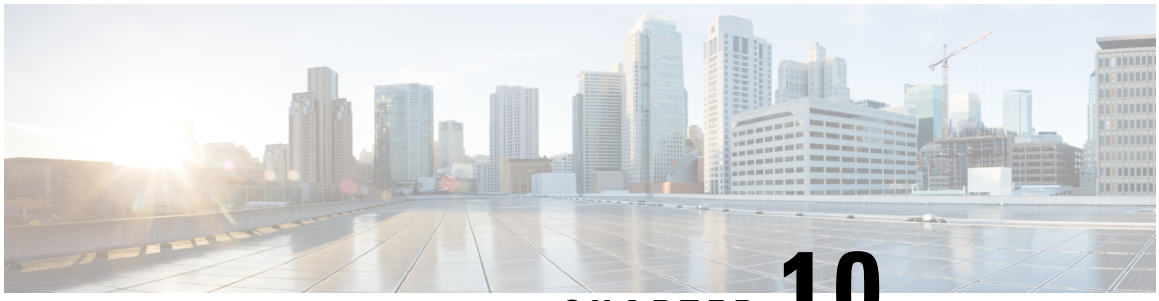

CHAPTER 10

S3X60 Server Node Hardware Management

- Cisco UCS S3260 Server Node Management, on page 91
- Booting a Cisco UCS S3260 Server Node, on page 92
- Booting a Cisco UCS S3260 Server Node from the Service Profile , on page 92
- Determining the Boot Order of a Cisco UCS S3260 Server Node, on page 92
- Shutting Down a Cisco UCS S3260 Server Node, on page 93
- Shutting Down a Cisco UCS S3260 Server Node from the Service Profile , on page 93
- Resetting a Cisco UCS S3260 Server Node, on page 94
- Resetting a Cisco UCS S3260 Server Node to Factory Default Settings, on page 95
- Reacknowledging a Cisco UCS S3260 Server Node, on page 96
- Removing a Cisco UCS S3260 Server Node from a Chassis, on page 96
- Deleting the Inband Configuration from a Cisco UCS S3260 Server Node, on page 97
- Decommissioning a Cisco UCS S3260 Server Node, on page 97
- Recommissioning a Cisco UCS S3260 Server Node, on page 98
- Reacknowledging a Server Slot in a S3260 Chassis, on page 98
- Removing a Non-Existent Cisco UCS S3260 Server Node from the Configuration Database, on page 99
- Turning the Locator LED for a Cisco UCS S3260 Server Node On and Off, on page 99
- Turning the Local Disk Locator LED on a Cisco UCS S3260 Server Node On and Off, on page 100
- Resetting the CIMC for a Cisco UCS S3260 Server Node, on page 100
- Resetting the CMOS for a Cisco UCS S3260 Server Node, on page 101
- Issuing an NMI from a Cisco UCS S3260 Server Node, on page 101
- Viewing the POST Results for a Cisco UCS S3260 Server Node, on page 102
- Viewing Health Events for a Cisco UCS S3260 Server Node, on page 102
- Health LED Alarms, on page 104

Cisco UCS S3260 Server Node Management

You can manage and monitor all Cisco UCS S3260 server nodes in a Cisco UCS domain through Cisco UCS Manager. You can perform some 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.

If a 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 rediscover the server in the slot.

Booting a Cisco UCS S3260 Server Node

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 Cisco UCS S3260 Server Node 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, or where the appropriate service profile already exists.

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 Cisco UCS S3260 Server Node

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.

Shutting Down a Cisco UCS S3260 Server Node

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 Cisco UCS S3260 Server Node 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 with the associated service profile.

Step 4 Choose the service profile associated with the server to be 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 Cisco UCS S3260 Server Node

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 Cisco UCS S3260 Server Node to Factory Default Settings

You can now reset a Cisco UCS S3260 Server Node to its factory settings. By default, the factory reset operation does not affect storage drives. This is to prevent any loss of data. However, you can choose to reset these devices to a known state as well.

The following guidelines apply to Cisco UCS S3260 Server Nodes when using scrub policies:

- For Cisco UCS S3260 Server Nodes, you cannot delete storage by using the scrub policy.
- Cisco UCS S3260 Server Nodes do not support FlexFlash drives.
- For Cisco UCS S3260 Server Nodes, you can only reset the BIOS by using the scrub policy.

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** check box.

  Note For Cisco UCS S3260 Server Nodes, you cannot delete storage using the scrub policy.

- To place all disks into their initial state after deleting all storage, check the **Create Initial Volumes** check box.

  You can check this check box only if you check the **Scrub Storage** check box. 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** box if you want to use storage profiles. Creating initial volumes when you are using storage profiles may result in configuration errors.

Cisco UCS Manager resets the server to its factory default settings.

Reacknowledging a Cisco UCS S3260 Server Node

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 Cisco UCS S3260 Server Node 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 reinstall the server, you must re-acknowledge the slot for Cisco UCS Manager to re-discover the server.

Deleting the Inband Configuration from a Cisco UCS S3260 Server Node

This procedure removes the inband management IP address configuration from a blade server. If this action is grayed out, no inband configuration was completed.

Procedure

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

Step 2 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 Cisco UCS S3260 Server Node

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 reinstall the server, you must re-acknowledge the slot for Cisco UCS Manager to rediscover the server.

Recommissioning a Cisco UCS S3260 Server Node

Procedure

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

Step 2 Click the **Chassis** node.

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

Step 4 On the row for each 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.

Reacknowledging a Server Slot in a S3260 Chassis

Perform the following procedure if you decommissioned a server without removing the physical hardware from the chassis, and you want Cisco UCS Manager to rediscover and reacknowledge 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.

Option	Description
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 Cisco UCS S3260 Server Node 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 Cisco UCS S3260 Server Node from a Chassis, on page 96.

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 Cisco UCS S3260 Server Node 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.

Turning the Local Disk Locator LED on a Cisco UCS S3260 Server Node On and Off

Before you begin

- Ensure that the disk is zoned. Turning the locator LED on and off cannot be done on disks that are not zoned.
- Ensure the server, on which the disk is located, is powered on. If the server is off, you are unable to turn on or off the local disk locator LED.

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 local disk locator LED on or off.

Step 4 In the **Work** pane, click the **Inventory** > **Storage** > **Disks** tabs.

The Storage Controller inventory appears.

Step 5 Click a disk.
The disk details appear.

Step 6 In the **Details** area, click **Toggle Locator LED**.
If the **Locator LED** state is **On**, it will turn **Off**. If the **Locator LED** state is **Off**, it will turn **On**.

Step 7 Click **Save Changes.**

Resetting the CIMC for a Cisco UCS S3260 Server Node

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 the management controller of the blade 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**.

Resetting the CMOS for a Cisco UCS S3260 Server Node

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**.

Issuing an NMI from a Cisco UCS S3260 Server Node

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.

Viewing the POST Results for a Cisco UCS S3260 Server Node

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.

Viewing Health Events for a Cisco UCS S3260 Server Node

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 health events.

Step 4 In the **Work** pane, click the **Health** tab

The health events triggered for this server appear. The fields in this tab are:

Name	Description
Health Summary area	
Health Qualifier field	Comma-separated names of all the heath events that are triggered for the component.

Name	Description
Health Severity field	Highest severity of all the health events that are triggered for the component. This can be one of the following: • **critical** • **major** • **minor** • **warning** • **info** • **cleared** **Note** The severity levels listed here are from highest to lowest severity.
Health Details area	
Severity column	Severity of the health event. This can be one of the following: • **critical** • **major** • **minor** • **warning** • **info** • **cleared** **Note** The severity levels listed here are from highest to lowest severity.
Name column	Name of the health event.
Description column	Detailed description of the health event.
Value column	Current value of the health event.
Details area	The **Details** area displays the **Name**, **Description**, **Severity**, and **Value** details of any health event that you select in the **Health Details** area.

Health LED Alarms

The server health LED is located on the front of each 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 server health LED blinks amber. This is indicated with a red dot. • Minor - The server 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 that 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.

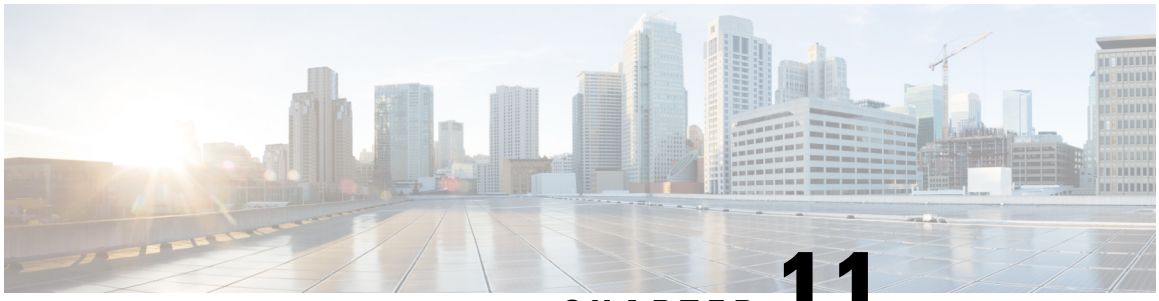

CHAPTER 11

Virtual Interface Management

- Virtual Circuits, on page 105
- Virtual Interfaces, on page 105
- Virtual Interface Subscription Management and Error Handling, on page 106
- Virtualization in Cisco UCS , on page 106

Virtual Circuits

A virtual circuit or virtual path refers to the path that a frame takes from its source vNIC to its destination virtual switch port (vEth) or from a source virtual switch port to its destination vNIC. There are many possible virtual circuits that traverse through a physical cable. Cisco UCS Manager uses virtual network tags (VN-TAG) to identify these virtual circuits and differentiate between them. The OS decides the virtual circuit that a frame must traverse on a basis of a series of decisions.

In the server, the OS decides the Ethernet interface from which to send the frame.

Note During service profile configuration, you can select the fabric interconnect to be associated with a vNIC. You can also choose whether fabric failover is enabled for the vNIC. If fabric failover is enabled, the vNIC can access the second fabric interconnect when the default fabric interconnect is unavailable. *Cisco UCS Manager Server Management Guide* provides more details about vNIC configuration during service profile creation.

After the host vNIC is selected, the frame exits the selected vNIC and, through the host interface port (HIF), enters the IOM to which the vNIC is pinned. The frame is then forwarded to the corresponding network Interface port (NIF) and then to the Fabric Interconnect to which the IOM is pinned.

The NIF is selected based on the number of physical connections between the IOM and the Fabric Interconnect, and on the server ID from which the frame originated.

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 maximum 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.

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 116 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.

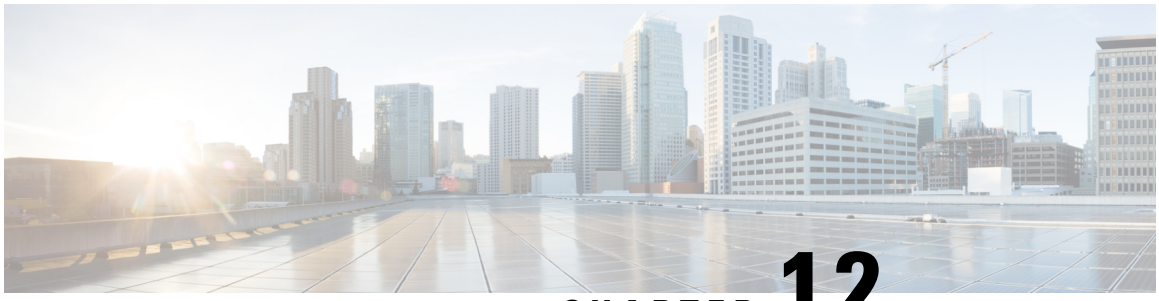

CHAPTER 12

Troubleshoot Infrastructure

- Recovering the Corrupt BIOS on a Blade Server, on page 109
- Recovering the Corrupt BIOS on a Rack-Mount Server, on page 110

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**.

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**.